Mechanism Design with Allocative, Informational and Learning Constraints

by

Abhinav Sinha

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2017

Doctoral Committee:

Associate Professor Achilleas Anastasopoulos, Chair
Associate Professor Jacob Abernethy
Professor Tilman M. Börgers
Professor Mingyan Liu

Abhinav Sinha

absi@umich.edu

ORCID iD: 0000-0003-2081-0580

Acknowledgments

I would like to express my sincere gratitude to my advisor Prof. Achilleas Anastasopoulos for the continuous support for the entire period of my PhD study and related research, for his patience, motivation and spectrum of knowledge. I learnt a lot through close association with him and his guidance was instrumental in shaping my research. In him I found the best advisor and mentor that I could have asked for.

I would also like to thank my thesis committee: Prof. Jacob Abernethy, Prof. Tilman Börgers and Prof. Mingyan Liu, for their insightful comments and encouragement, which were invaluable in the timely completion of this thesis and also pushed me to broaden my research horizon.

I would like to thank the many anonymous reviewers who refereed my journal and conference papers, their feedback was essential in the progress of my research as well as my thesis. I would also like to thank the many professors, students and staff at the University of Michigan who made life interesting and also provided many opportunities to learn new things here in Ann Arbor.

Finally, I would like to thank my family: my parents, Shri A.K. Sinha and Smt. Archana Sinha and to my sister, Anubha Kumar, for keeping my spirits up throughout my PhD.

Contents

Acknowledgments ii

List of Figures vii

List of Tables ix

Abstract x

Introduction 1

Part I. Mechanism Design 11

Chapter 1. A Practical Mechanism for Network Utility Maximization for Unicast Flows on the Internet 12

1. Introduction 12
2. Centralised Problem 15
   2.1. Assumptions 16
   2.2. Necessary and Sufficient Optimality conditions 17
   2.3. Taxation and Budget Balance 17
3. Mechanism with Weak Budget Balance 19
   3.1. Mechanism 19
   3.2. Results 21
   3.3. Comments 23
4. Mechanism with Strong Budget Balance 23
5. Summary and Comments 25
6. Appendix 26

Chapter 2. Mechanism Design for Resource Allocation in Networks with Intergroup Competition and Intragroup Sharing 36

1. Introduction 37

2. Centralized Problem 42
2.1. Assumptions 44
2.2. Necessary and Sufficient Optimality conditions 45
3. A mechanism with weak budget balance 47
3.1. Mechanism 47
3.2. Results 52
4. A mechanism with Strong Budget Balance 55
4.1. Mechanism 56
4.2. Results 57
5. Summary and Comments 59
6. Appendix 61

Chapter 3. A General Mechanism Design Methodology for Social Utility Maximization with Linear Constraints 71
1. Introduction 72
2. Motivation and Centralized Problem 75
2.1. Interesting Resource allocation problems in Communications . . 75
2.2. General Centralized Problem 78
2.3. Assumptions 78
2.4. KKT conditions 79
3. Mechanism 80
3.1. Allocation 81
3.2. Taxes 85
4. Equilibrium Results 85
4.1. Strong Budget Balance (at equilibrium) 88
5. Off-equilibrium Results 90
5.1. SBB off-equilibrium 91
6. Generalizations 93
7. Appendix 96

Chapter 4. Distributed Mechanism Design with Learning Guarantees 102
1. Introduction 103
2. Model 105
2.1. Private goods allocation problem 106
2.2. Public goods allocation problem 107
3. A mechanism for the private goods problem 109
3.1. Definitions 109
3.2. A distributed mechanism with scalable messages 110

3.3. A mechanism for the private goods problem 113
3.4. Results 114
4. A mechanism for the public goods problem 117
4.1. Results 118
5. Learning Guarantees 121
5.1. Results 123
5.2. Numerical Analysis of convergence 124
6. Summary and Comments 128
7. Appendix 131

Chapter 5. Mechanisms Design for Fairness 146
1. Introduction 146
2. Motivating Application and General Model 152
2.1. Demand-side management of smart grids 152
2.2. General Centralized Problem 155
3. Dominant Strategy Implementation 155
4. Bayesian Implementation 158
5. Full Bayesian Implementation 161
6. An Analytical Example for Dominant Implementation 164
7. Bayesian Implementation for smart grid demand-side management . . 166
8. Summary and Comments 169
9. Appendix 170

Part II. Dynamic Games 180

Chapter 6. Structured Perfect Bayesian Equilibrium in Infinite Horizon Dynamic Games with Asymmetric Information 181
1. Introduction 182
2. Model for Dynamic Game 184
3. Technical Lemmas in the finite horizon model 187
3.1. Finite Horizon problem and Backwards Recursion 187
3.2. Finite Horizon Result 189
4. SPBE in Infinite Horizon 190
4.1. Perfect Bayesian Equilibrium 190
4.2. Fixed Point Equation for Infinite Horizon 190
4.3. Relation between Infinite and Finite Horizon problems 192
4.4. Equilibrium Result 192

5. A Concrete Example 192
6. An existence result for the finite horizon fixed-point equation 197
    6.1. Points of Discontinuities and the Closed graph result 199
7. Conclusion and Future Work 201
8. Appendix 203

Bibliography 211

List of Figures

1. A network that imposes "unicast" constraints 4
2. A network that imposes "multicast/multirate" constraints 4
3. A public goods network with local interference 4
4. How does restricted communication affect design? 6
5. Range of modeling and their trade-offs. 6
6. An instance of the Unicast network 16
7. Network with "unicast" architecture 37
8. Network with "multicast/multirate" architecture 37
9. Radial projection allocation for a single-link multicast/multirate network with 2 multicast groups $\mathcal{K} = \{1, 2\}$ with a total of 3 agents i.e. $\mathcal{N} = \{11, 12, 21\}$. Capacity constraint is $\max\{x_{11}, x_{12}\} + x_{21} \leq 1$. For demand y allocation x is just the projection of y onto the feasible set boundary (first sub-case of (2.9c) applies here). 48
10. An illustration of the allocation for $N = 2$ 82
11. $n(i, j)$ and $d(i, j)$ for the strongly connected directed graph $\mathcal{G}$. 110
12. $||m_n - \widetilde{m}||_2$ vs. n for the private goods mechanism. 126
13. $||\widehat{x}(m_n) - x^*||_2$ vs. n for the private goods mechanism. 127
14. $||\left(\widehat{p}_i(m_n)\right)_{i\in\mathcal{N}} - (\lambda_1^*, \dots, \lambda_1^*)||_2$ vs. n for the private goods mechanism. . . . 127
15. $||m_n - \widetilde{m}||_2$ vs. n for the public goods mechanism. 128
16. Gini Coefficient of $\left(v(\hat{x}_\epsilon(\theta); \theta_i)\right)_{i\in\mathcal{N}}$ vs ϵ, $N = 50$. 165
17. Plot of quantities that determine ϵ_{max}. The difference between upper and lower bounds imposed by BSIC and $H(\epsilon)$ from eq. (5.12a). 166
18. Gini Coefficient (%) and Standard Deviation of $\left(v(\hat{x}_\epsilon(\theta); \theta_i)\right)_{i\in\mathcal{N}}$ vs ϵ. . . . 167
19. Standard Deviation of Taxes vs. ϵ, for 10% agents with type θ_H as well as mean with Bin(100,0.1) prior. 168

20. $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1, π_2) at $\delta = 0$. 195
21. $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1, π_2) at $\delta = 0.5$. 195
22. $\tilde{\gamma}^1(a^1 = 1 \mid x^H)$ vs. (π_1, π_2) at $\delta = 0.5$. 196
23. $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1, π_2) at $\delta = 0.95$. 196
24. $\tilde{\gamma}^1(a^1 = 1 \mid x^H)$ vs. (π_1, π_2) at $\delta = 0.95$. 196

List of Tables

1. (ξ, δ) parameter values for different graphs and problems. 125

2. $\Theta_i = \{5, 15\}$, $N = 100$, 10% agents with high type, $a = 0.3$, $m_i = 0.1$ 153

Abstract

Efficient allocation of network resources is a highly desirable goal, with applications of interest ranging from bandwidth allocation in unicast/multicast services on the internet, to power allocation in a wireless interference networks and spectrum allocation for cellular networks. In this thesis, we consider such network problems in the presence of strategic decision-makers and solve the various mechanism design and game theoretical problems that arise from it.

In the first four chapters, we present single-shot incentive mechanisms for strategic agents, geared towards the above applications, such that the outcome produced at any Nash equilibrium is socially efficient. The focus in the first three chapters is on developing a systematic approach to designing such mechanisms such that various applications that differ qualitatively can be combined under the same design umbrella. Our approach models each application with its own network allocation constraints. Additionally, the design ensures that the mechanism contract only promises feasible allocation regardless of what messages agents choose i.e., the allocation through the mechanism cannot violate the network constraints. To this end we introduce a new radial allocation scheme. The benefit from this design criterion is that it allows the mechanism to be played off-equilibrium e.g., when agents are still learning the Nash equilibrium.

Chapter 4 focuses on resolving two separate issues that most mechanisms in the Nash implementation literature overlook. The first is the distributed-ness of communication between agents – the message chosen by an agent is assumed to be observed only by that agent's neighbors on a communication graph (typical mechanism design model assumes that the messages are broadcasted). We demonstrate analytically that the distributed-ness of communication increases the dimensionality of the message space to the order of the average degree of the communication graph. The second issue is that of providing theoretical guarantees on how the one-shot Nash equilibrium can be learnt when agents only possess information about their

own utilities and not that of others. We provide positive convergence guarantees for the entire class of adaptive best-response learning dynamics. These are a set of strategies where actions are myopic and depend on the recent history of observed actions. The guarantee provided is such that agents do not need to coordinate with their strategies in order to converge. Resolving both the issues simultaneously further increases the dimensionality of the message space to the order of number of agents.

In Chapter 5, we consider a Bayesian model with uncertainty in agents' utilities. The objective of the mechanism designer is to achieve fairer allocation than the socially efficient one. We provide sufficient conditions on the fair allocation objective such that a truthful mechanism exists. The results are provided separately for Bayesian and dominant strategy incentive compatibility. Through a concrete example of demand-side management in smart grids we demonstrate numerically, the gains in fairness of the allocation as well as taxes collected through our method.

Chapter 6 considers infinite-horizon dynamic games of asymmetric information and develops a dynamic programming style characterization for calculating a subset of all perfect Bayesian equilibria (PBE). This subset of PBEs are rich enough to exhibit signaling behavior, which is also demonstrated through numerical analysis. The motivation for this work is the possible deployment of this equilibria characterization technique in designing dynamic mechanisms for systems where direct mechanisms are not desired, for instance due to privacy considerations.

Introduction

> Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves.
>
> *(Herbert Simon, The Sciences of the Artificial)*

As machines behave in the manner dictated, predicting outcomes for a well-designed system where machines are the decision makers is straightforward. This is not the case when the decision-makers in the system are human beings. Predicting how human beings behave is a complex, and as yet not fully understood, phenomenon which is typically investigated in the field of Decision Theory. One well-accepted model for human behavior is *rationality*.

Predicting outcomes in systems with multiple rational decision-makers is an even more complex task as each rational decision-maker acts in anticipation of how the other decision makers will act. For this reason, analysis of such systems is done via the study of *Games* with appropriately defined notion(s) of *Equilibrium*. With a set notion for analysis, a natural next step in development is design. Designing a system with multiple human decision-makers is studied in the field of Mechanism Design.

The framework of Mechanism design aims to bridge the informational gap between a designer, who wishes to achieve "efficient" allocation and rational agents, who posses private information relevant to determining the efficient allocation. This thesis focuses on designing mechanisms for a networked system of agents where the designer is interested in accommodating, apart from the main objective of efficient allocation, certain allocative, informational and learning constraints imposed by the network.

Motivating Design Principles

The motivation for designing mechanisms with inbuilt network constraints are listed below. In the subsequent section on contributions, the various motivations below are matched to problems - each presented in a different chapter later.

Unless mentioned otherwise, the solution concept used is Nash equilibrium and of interest is the design of full implementation mechanisms i.e., no inefficient Nash equilibria.

In the seminal paper [Ma99], necessary and sufficient conditions for Nash implementation (with three or more agents) are provided. Thus a complete characterization exists of the social choice functions (SCF) that are Nash implementable. Furthermore, for the SCFs that are Nash implementable, the constructive proof even suggests a "canonical" mechanism. However, this "canonical" mechanism involves a huge communication cost - each agent is required to quote every agent's preference, which is impractical in real world settings. Another impractical aspect of the proposed mechanism is that it leads agents to play a *modulo* game[1].

Our motivation stems from designing Nash implementation mechanisms, for specific applications, that are more structured and that can surpass the above criticisms of the "canonical" mechanism. Specific considerations are listed below.

(a) Reduction in the Communication overhead

A significant part of the literature on network protocols is interested in the Network Utility Maximization (NUM) problem. These problems aim to solve

$$\max_{x} \sum_{i=1}^{N} v_i(x) \quad s.t. \;\; x \in \mathcal{X}, \tag{0.1}$$

which is similar to the Utilitarian Social Welfare maximization problems in Economics, except for the *allocative* constraint $x \in \mathcal{X}$, imposed by the network. In the mechanism design model for this objective, the private information possessed by agents is the utility function $v_i(\cdot)$, which usually is a function from a subset of a Euclidean space to $\mathbb{R}$.

In such models, the use of *direct* mechanisms (such as the VCG mechanism) is prohibitive as it would require agents to communicate an entire function to the designer. Thus it is

[1] A game where each agent quotes an integer and the outcome is decided based on the remainder of the sum of the quotes, modulo N (the number of agents).

preferable to design mechanisms aimed at achieving the efficient allocation defined above which require agents to communicate "less" with the designer, e.g. it is preferable to quote a vector of real numbers than a function whose domain is a Euclidean space.

(b) Systematic Design

It is not surprising that design with the above mentioned reduction in communication is typically tuned to the allocative constraints imposed by the network (feasible set $\mathcal{X}$ in (0.1) above). The presented mechanisms in the literature reveal a case-by-case approach. One important motivation of this thesis is the unification of the design principles concerned with various instances of $\mathcal{X}$.

Depicted in Fig. 1, 2 and 3 are three such examples. The first two arise from the problem of bandwidth allocation between pairs of transmitters and receivers, where the data flow is restricted by capacity on the various links of the network. Fig. 1 and 2 consist of 4 pairs of receivers and transmitters: T1-R1, T1-R2, T2-R3, T2-R4. A unicast network is one where at each link the constraint on data flow is through the summation of data over each pair that uses that link. This means that the constraint on link AB is $x_1 + x_2 + x_3 + x_4 \leq c_{AB}$, where c_{AB} is the capacity of link AB. A multirate/multicast network is one where the constraint is through the summation over transmitters of only the maximum data flowing from each transmitter. This means that the constraint on link AB in the multirate/multicast network is $\max\{x_1, x_2\} + \max\{x_3, x_4\} \leq c_{AB}$. Such a protocol requires that on each link, for each transmitter, only one connection be setup - corresponding to the receiver who receives the highest quality data. Lower quality data can then be generated, at fixed nodes, by sampling higher quality data, so that the limited bandwidth on the links can be saved. Finally, the model depicted in Fig. 3 is that of power allocation on a wireless network with local interference. For instance, the power allocated to agent 1 causes interference to agent 2 (and vice-versa) since they are both part of the same neighborhood $\mathcal{N}_1$.

These problems represent different aspects of allocation models - public and private goods. Through a systematic design one hopes to combine both these aspects into the same design framework.

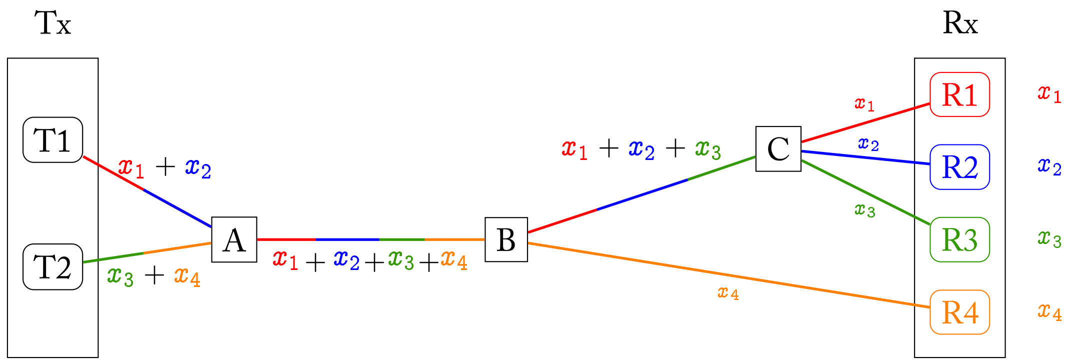

Figure 1.: A network that imposes "unicast" constraints

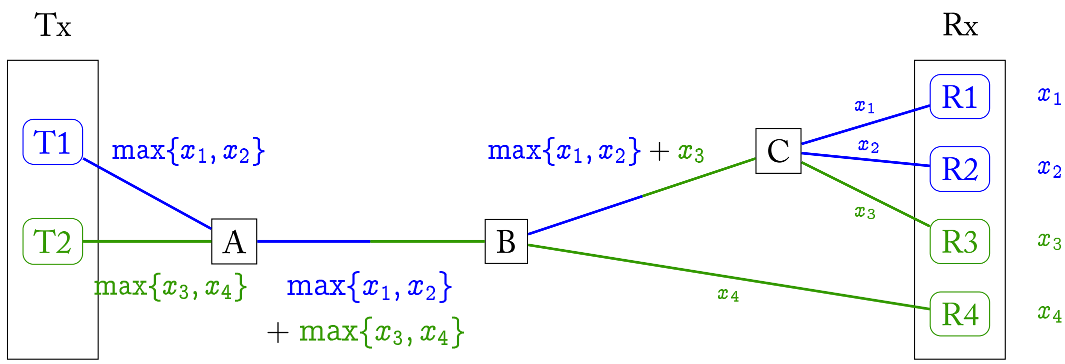

Figure 2.: A network that imposes "multicast/multirate" constraints

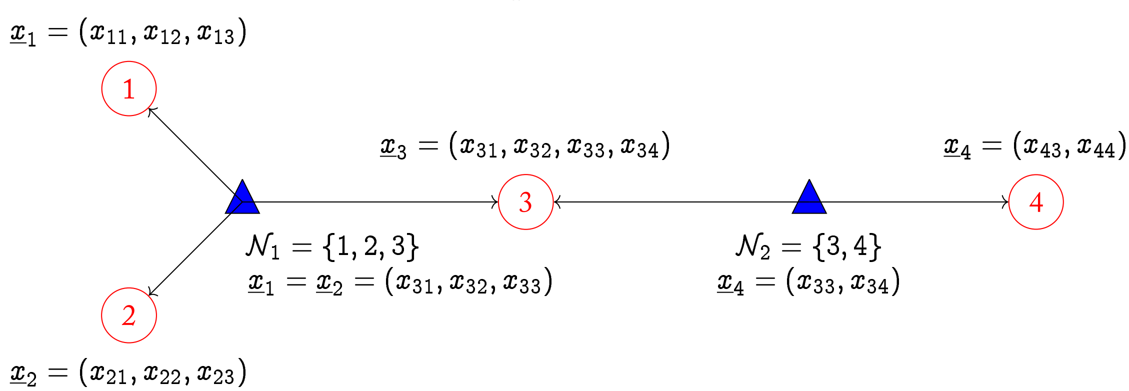

Figure 3.: A public goods network with local interference

(c) Informational constraints on messaging

The principle framework of Mechanism Design is such that it assumes the existence of central entity capable of communicating with each agent. Thus most mechanisms in literature define the contract (e.g., allocation and tax) such that messages from all agents need to be collected centrally in order to determine any agents' outcome. This is akin to assuming a broadcast structure of communication between agents, where an alternate interpretation is that agents broadcast their message to all other agents and then each agent can calculate their outcome by evaluating the prespecified contract at the overall message profile. What happens if certain channels of communication between agents are broken? For instance, Fig. 4 depicts a system with four agents where only four out of the six possible communication channels are still active. Agent 1 cannot send his/her message to agent 3 and vice-versa. Similarly, agent 1 cannot send his/her message to agent 4 and vice-versa. If a mechanism can be designed that obeys such a non-centralized communication structure, then it will be more robust.

The motivation for looking into a distributed communication structure whilst designing mechanisms comes from the literature on distributed optimization that aims to address similar informational constraints between **non-strategic** agents with private information and a designer interested in allocating efficiently, where the efficient allocation depends on the said private information.

(d) Learning of Nash equilibrium

Different models in mechanism design literature assume different extent of rationality for the agents. A more sophisticated model may assume fully rational agents who are prone to strategizing over long time horizons. This would mean that the outcome of repeated interaction of such agents is analyzed through the concept of dynamic games. On the other hand, a simpler model may assume rational agents who strategize only over a single time period, whereas for repeated interactions an appropriate model for their behavior is through a class of learning strategies, such as Bayesian learning. A simpler model would be one where agents are not rational and simply obey the command of the designer. As one moves from a sophisticated model to a compliant one, lesser practical

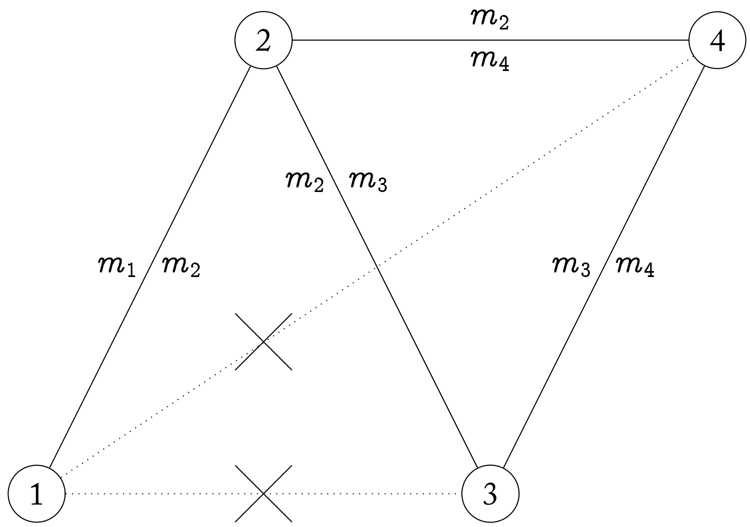

Figure 4.: How does restricted communication affect design?

scenarios are covered under the design paradigm and thus the guarantees that can be provided for practical applications reduce. On the other hand, the design space at the designer's disposal increases. Owing to this trade-off (Fig. 5), one may be inclined to investigate an appropriate model that leaves enough room for design but is capable of providing robust enough guarantees for applications in the real world. We describe the motivation behind one such framework below.

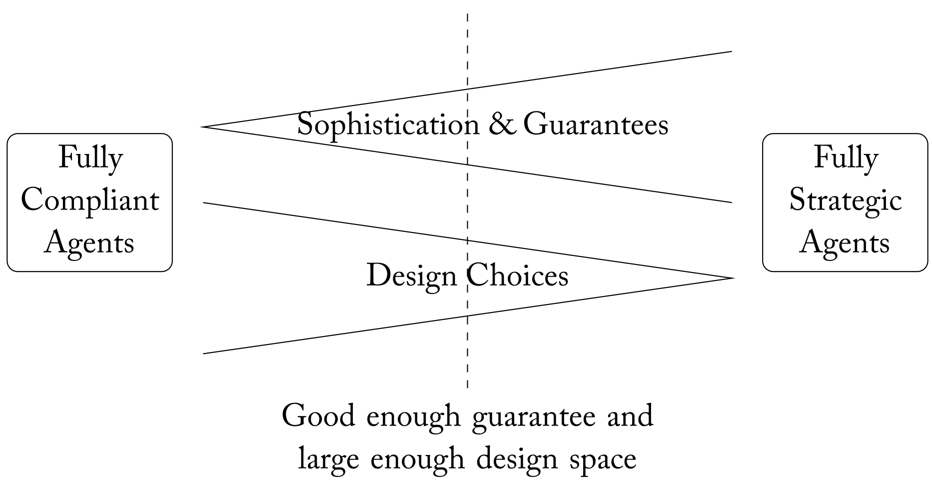

Figure 5.: Range of modeling and their trade-offs.

The perfect information solution concept of Nash equilibrium is used in mechanism design, even in systems where agents are practically not expected to be fully informed about other agents' private utilities. The common justification for use of Nash equilibrium in such cases is that it can be learned by repeated play and appropriate adjustment of strategies over time.

Learning guarantees i.e., theoretically proven convergence of a specific learning dynamics or a class of learning dynamics, is provided largely on a case-by-case basis. In most cases, learning guarantees are provided for a specific and model-tailored learning dynamic.

Another important motivation in this thesis is to design mechanisms with the design constraint of guaranteed convergence for a (large enough) class of learning dynamics that is not necessarily tailored to the model.

A slightly different motivation in this regard is also to design mechanisms that are "learning ready". For network problems with complicated allocative constraints, and where learning guarantees are in general difficult to provide, can there exist mechanisms that are at least conducive to repeated play that aims to learn the Nash equilibrium. A mechanism can be said to be conducive to learning if the allocation function designed by the mechanism designer produces feasible allocation, i.e., $\hat{x}(m) \in \mathcal{X}, \forall\, m$. This is a crucial consideration for network problems with allocative constraints since these constraints (e.g., imposed by the capacity of a link) are *hard* and absolutely cannot be violated in the short run.

(e) Designing beyond maximization of utilitarian social welfare

While many resource allocation problems are designed with NUM as their objective, there are problems of practical interest where the designer, for a variety of social reasons, may be interested in accommodating fairness in allocation even at the cost of utilitarian efficiency. Here fairness refers to a more equitable distribution of resources i.e., avoiding a situation where most of the resources end up in the hands of a few. In this case, a basic question to ask is whether there exist mechanisms that implement the fairness-aware efficient allocation? And if so, to what extent can fairness be accommodated?

There is a scarcity of results in mechanism design literature, with the objective of fair allocation. The prevalence of NUM type objective over fairness has strong mathematical foundations. The main reason is that for NUM, in models with additively separable taxes, an agent's individual utility can be aligned[2] directly with the overall social objective, so that when an agent maximizes his/her net utility they are simultaneously maximizing the overall social objective. The VCG mechanism is the most prominent example of this. The NUM problem gives the allocation function as $\hat{x}(\theta) = \mathtt{argmax}_{x \in \mathcal{X}} \sum_{j=1}^{N} v(x, \theta_j)$. For any agent i, his/her utility $v(x; \theta_i)$ is determined by the allocation x and their private type θ_i. For quoted message profile $\phi = (\phi_j)_{j \in \mathcal{N}}$, the utility of any agent i is $v(\hat{x}(\phi); \theta_i) - t_i(\phi)$, where VCG taxes $t_i(\phi) = -\sum_{j \neq i} v(\hat{x}(\phi); \phi_j) + f_i(\phi_{-i})$ precisely create a situation such that the user performs social objective maximization whilst maximizing self utility. As soon as one moves away from the NUM objective, due to the social and individual objectives not aligning with each other, basic design techniques like VCG mechanism are not useful. This is one of the main reasons why there are significantly fewer results on fairness in the mechanism design literature.

A second, more subtle, reason for investigating fairness in mechanism design is that with fairer allocation one expects there to be lesser variation in the utility of agents (derived from the allocation). It is thus natural to investigate whether this leads to lesser variation in the tax paid by agents. Such an outcome is desirable for many reasons, the prime among which is that due to implicit budgetary restrictions, a mechanism that realizes high tax variation may be considered impractical.

Overview of this thesis

Mechanism Design

Chapters 1, 2 and 3 consider the NUM objective with allocative constraints and present efficient mechanisms that accommodate restrictions consistent with motivations (a), (b), and a part of (d). Chapter 1 considers the general unicast network, which is a version of the private goods allocation problem. Chapter 2 considers the multirate/multicast network which possesses the qualitative properties of both types - private goods and public goods - of allocation

[2] through appropriate taxation

problems. This dual nature comes about due the unique nature of the network constraints in multirate/multicast. Finally, in Chapter 3, the unified and systematic approach developed in Chapters 1 and 2 is used to design mechanisms for any network with linear allocative constraints - even those that fall outside of unicast and multicast. The general approach applies to local interference power transmission problem depicted in Fig. 3. The mechanisms in these three chapters are all "learning-ready" in the sense that the allocation function is always feasible w.r.t. the allocative constraints - even off-equilibrium.

Chapter 4 considers the classical Walrasian (with money) and Lindahl allocation problems and presents efficient mechanisms that are consistent with motivations (a), (c) and (d). Communication between agents is restricted to a connected graph and the learning guarantees are provided for a subclass of adaptive strategies called adaptive best-response dynamics.

Chapter 5 considers mechanism design for fair allocation and is consistent with motivation (e). For two different cases, existence of dominant strategy incentive compatible and Bayesian strategy incentive compatible mechanisms are shown. The net effect on the outcome (fairness in allocation and taxes) is numerically analyzed through a concrete example of demand-side power management in smart grids.

Dynamic Games

Finally, Chapter 6 moves away from the theme of Mechanism design. It considers the problem of systematic evaluation of perfect Bayesian equilibria (PBE) in infinite horizon dynamic games with asymmetric information. While existence of PBE is guaranteed, till date there is no simplified characterization/algorithm of how to find one (or all). Technically, going by the definition, one has to solve a hugely complicated inter-temporal fixed point equation in strategies and beliefs to find the PBEs of a game. We propose a single-shot fixed-point equation in (one round) strategy along with a forward recursive construction to evaluate a subset of PBEs with "less complex" strategies i.e., where agents are not required to maintain ever-increasing history of observations. Furthermore, these equilibria are appealing because they exhibit *signaling* behavior.

The motivation for this work is the possible deployment of the systematic evaluation technique for designing dynamic mechanisms in systems where direct mechanisms are not desirable.

Part I.

Mechanism Design

Chapter 1.

A Practical Mechanism for Network Utility Maximization for Unicast Flows on the Internet

In this chapter we consider the scenario of unicast service on the Internet where a network operator wishes to allocate rates among strategic users in a way that maximizes overall user satisfaction while respecting capacity constraints on every link in the network. In particular, we construct two mechanisms that fully implement social welfare maximizing allocation in Nash equilibria (NE) for the above scenario when agents' utilities are their private information. We assume that the agents are strategic and thus the mechanism contract induces a game among the agents. The emphasis of this work is on full implementation which means that for the game induced by the mechanism, the allocation at all NE is the social welfare maximizing allocation and thus no extraneous/unwanted equilibria are created, as is the case in general mechanism design. The constructed mechanisms are amenable to learning, an essential requirement when using NE as a solution concept. This is achieved by ensuring that they result in feasible allocations on and off equilibrium and are budget balanced.

1. Introduction

As Internet applications become ubiquitous, there is a strong need for efficient usage of network resources currently at our disposal. Due to limited information capacity of network connections, each user can only be afforded lesser and lesser data-rate as the Internet becomes more congested. A typical way to ensure efficient allocation, especially in case of users with heterogeneous preferences (such as on the Internet), is via markets. Charging for services will distinguish between essential and non-essential usage for various types of agents.

When people are required to pay for services, it is natural to assume that they will be utility-maximizing, i.e., strategic. The designer should take this behavior into account and design protocols with the appropriate incentives in place, or else the design may lead to highly inefficient allocation of network resources. The appropriate theoretical framework for solving such problems is Mechanism Design or Implementation Theory. At the core of this framework is the solution to the information elicitation problem, where a designer – despite lacking full information required to implement a goal (such as network utility maximization) – is able to produce the right allocation by asking agents to make choices under appropriate incentives. Mechanism design has been studied extensively for market problems like implementation of Walrasian and Lindahl allocations [Hu79; GrLe77], and in auctions and bilateral trade contracts. In general mechanism design, it is required that at least one Nash equilibrium (NE) of the induced game is efficient [GaNaGu08a; GaNaGu08b; BöKrSt15; Kr09; HuRe06]. In this chapter our focus is on **full Nash implementation**, which puts more stringent constraints by requiring that all NE of the induced game must result in efficient allocation [Ja01].

Network utility maximization (NUM) or social welfare maximization are interchangeable terms in our context, and refer to allocation of data-rates subject to system constraints such that sum of utilities of all the agents in the system is maximized. We consider the unicast flow model on the Internet. In this set up, pairs of source and destination nodes are connected via links of a network where each link can be servicing multiple routes at the same time. This creates simultaneous (coupled) congestion on all links.

The goal of NUM for unicast networks has been central to resource allocation research for a long time. A number of works focus on aspects of complexity or distributed algorithms (see [SaTa02; KaTa06]) in the absence of strategic-ness. For the cases where strategic-ness is considered, the design specifics hinge on the solution concept employed for the model; examples of which are NE, Bayesian NE, dominant strategy (or ex-post), competitive equilibrium etc. In the seminal work [KeMaTa98], a mechanism for general unicast was designed with competitive equilibrium as the solution concept. This was modified to get a Nash implementation mechanism for the case of single-link unicast in [YaHa05; MaBa04]. A mechanism for unicast was developed in [JaWa10] with the property that at least one NE is efficient, but full implementation was not guaranteed. In [KaTe14], full implementation in NE was achieved, but the allocation was not feasible off equilibrium, thus making the mechanism inappropriate for practical implementation. Finally, VCG mechanisms [Kr09; BöKrSt15; GaNaGu08a] applied to NUM problems can yield dominant strategy implementation. However VCG mechanism require that agents broadcast messages that belong to their type space. This implies that when agents' private information is their entire valuation function, the corresponding messages have to be functions themselves, thus rendering this type of mechanism impractical due to excessive communication requirements on the Internet. A modification arising out of this requirement has been dealt with in [JoTs09; YaHa07], where one-dimensional messages are required per user. However, the problem of extraneous equilibria still persists and having message space too small can make learning NE impossible (see [Jo86]).

In this chapter, design was undertaken not only with the primary objective of full Nash implementation in mind but also practicality of using such a mechanism. This was dealt-with in two ways. Firstly, the dimension of message space was kept linear in number of agents so that the mechanism is scalable on a large network. Secondly, the concept of proportional allocation (first introduced in [KeMaTa98; YaHa05; MaBa04] for a single link) was generalized to an allocation function that *always* produces feasible allocation, thereby making the mechanism amenable to pre-game learning in a real-world setting. To see why feasibility off equilibrium is so important, recall that using NE implicitly requires agents in the system to be able to calculate it first. This however requires complete information for every agent in the system, which isn't justified in a physically and informationally decentralized system like the Internet. Use of NE in such cases can be justified since it has been empirically observed that agents dynamically learn the NE by playing the game repeatedly and coordinating as time goes on. During each time step in this dynamic process, agents broadcast their messages, and the mechanism allocates resources and taxes them accordingly. Thus, it is imperative that allocation of resources can be done throughout this process before equilibrium is reached, and this requires a mechanism that is feasible off equilibrium!

2. Centralised Problem

Consider a set $\mathcal{N} = \{1, 2, \ldots, N\}$, of N of Internet agents (an agent is considered a pair of source and destination users) that communicate over pre-specified routes on the Internet. Each agent $i \in \mathcal{N}$, communicates at an information rate $x_i \in \mathbb{R}_+$ (non-negative real numbers). Agents' valuation for an overall rate allocation $x = (x_i)_{i \in \mathcal{N}} \in \mathbb{R}_+^N$, is

$$\tilde{v}_i(x) = v_i(x_i) \qquad \forall\ i \in \mathcal{N} \tag{1.1}$$

where $v_i : \mathbb{R}_+ \to \mathbb{R}$, for all $i \in \mathcal{N}$. This indicates that agent i's satisfaction only depends on its own information rate allocation x_i. Due to capacity constraints on the utilized links, allocation to agents is constrained by a number of inequality constraints. It is assumed that although some agents may share information content (e.g., watch the same video stream), a separate data stream is transmitted for each agent. This transmission technique is referred to as unicast service. Each agent has a pre-determined route, $\mathcal{L}_i$ for agent i, which is the set of links that agent i uses for his communication, and $\mathcal{L} = \cup_{i \in \mathcal{N}} \mathcal{L}_i$ is the set of all available links. We also define the sets of agents utilizing link $l \in \mathcal{L}$ as $\mathcal{N}^l = \{i \in \mathcal{N} \mid l \in \mathcal{L}_i\}$. Finally, for any agent i and link l, we denote $L_i = |\mathcal{L}_i|$ and $N^l = |\mathcal{N}^l|$.

The network administrator is interested in maximizing the social welfare (or network utility) under the link capacity constraints. This centralized problem is

$$\max_x \sum_{i \in \mathcal{N}} v_i(x_i) \tag{1.2a}$$

$$\text{s.t.} \quad x_i \geq 0 \quad \forall\ i \in \mathcal{N} \tag{1.2b}$$

$$\text{and} \quad \sum_{j \in \mathcal{N}^l} \alpha_j^l x_j \leq c^l \quad \forall\ l \in \mathcal{L} \tag{1.2c}$$

Constraints in (1.2c) are the inequality constraints on allocation, which can be interpreted as capacity constraint for every link $l \in \mathcal{L}$, in the network. In this interpretation α_j^l would be representative of the QoS requirement of agent j combined with the specific architecture on link l. As an example, $\alpha_j^l = \dfrac{1}{R_j(1 - \epsilon_j^l)}$ for all links $l \in \mathcal{L}_j$, where ϵ_j^l represents the packet error probability for link l for a packet encoded with channel coding rate R_j.

An instance of the Unicast network is depicted in Fig. 6, with four agents (pairs of receivers

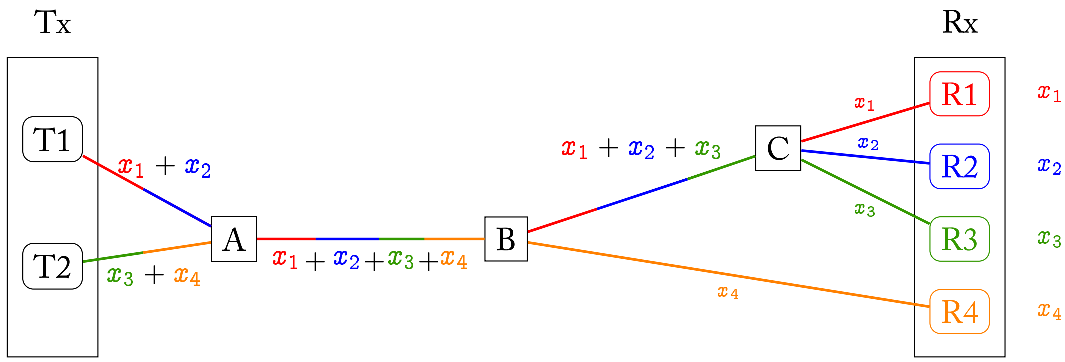

Figure 6.: An instance of the Unicast network

and transmitters): T1-R1, T1-R2, T2-R3, T2-R4. For simplicity assume that $\alpha_j^l = 1$ for all j, l. For all four agents, their data passes through the link AB and thus the constraint imposed by link AB is $x_1 + x_2 + x_3 + x_4 \leq c_{AB}$, where c_{AB} is the capacity of the link. Similarly, since link BC is used only by agents T1-R1, T1-R2 and T2-R3, the constraint imposed by it is $x_1 + x_2 + x_3 \leq c_{BC}$.

2.1. Assumptions

Following are assumptions on problem (1.2).

(A1) For all agents, $v_i(\cdot) \in \mathcal{V}_i$, where the sets $\mathcal{V}_i$ are arbitrary subsets of $\mathcal{V}_0$, the set of strictly increasing, strictly concave, twice differentiable functions $\mathbb{R}_+ \rightarrow \mathbb{R}$ with continuous second derivative.

(A2) $v_i'(0)$ is finite $\forall\ i \in \mathcal{N}$. This implies that $v_i'(x)$ is finite $\forall\ i$ and $\forall\ x$, since $v_i(\cdot)$'s are concave.

(A3) There are at least two agents on each link i.e. $N^l \geq 2, \forall\ l \in \mathcal{L}$.

(A4) The optimal solution of (1.2) has at least two non-zero components at each link i.e. if $S(x) := \{i \in \mathcal{N} \mid x_i > 0\}$ and $S^l(x) := S(x) \cap \mathcal{N}^l$ then we assume $|S^l(x^*)| \geq 2\ \ \forall\ l \in \mathcal{L}$. (where x^* solves (1.2))

In addition, the coefficients in (1.2c) are all strictly positive, i.e. $\alpha_j^l > 0\ \forall\ j \in \mathcal{N}^l, \forall\ l \in \mathcal{L}$. Also, for well-posedness of the problem we take $c^l > 0\ \ \forall\ l \in \mathcal{L}$.

Assumption (A1) is made in order for the centralized problem to have a unique solution and for this solution to be characterized by the KKT conditions. (A2) is a mild technical assumption that is required in the proof of Lemma 3.7. Assumption (A3) is made in order to avoid situations where there is a link constraint involving only one agent. Such case requires special handling in the design of the mechanism (since in such a case there is no contention at the link), and deviates from the basic idea that we want to communicate. Finally (A4) is related to (A3) and is made to simplify the exposition of the proposed mechanism, without having to define corner cases that are of minor importance.

2.2. Necessary and Sufficient Optimality conditions

We can write the following four KKT conditions, which are generally necessary, but in our case (due to all constraints being affine and strict concavity of v_i) they will also be sufficient. For optimal primal-dual solution (x^*, λ^*) we have

(a) Primal Feasibility: x^* satisfies (1.2b) and (1.2c).

(b) Dual Feasibility: $\lambda_l^* \geq 0 \;\; \forall\, l \in \mathcal{L}$.

(c) Complimentary Slackness:

$$\lambda_l^* \Big(\sum_{j \in \mathcal{N}^l} \alpha_j^l x_j^* - c^l \Big) = 0 \quad \forall \;\, l \in \mathcal{L} \tag{1.3}$$

(d) Stationarity:

$$v_i'(x_i^*) = \sum_{l \in \mathcal{L}_i} \lambda_l^* \alpha_i^l \quad \forall \;\, i \in \mathcal{N} \quad \text{s.t.} \quad x_i^* > 0 \tag{1.4}$$

$$v_i'(x_i^*) \leq \sum_{l \in \mathcal{L}_i} \lambda_l^* \alpha_i^l \quad \forall \;\, i \in \mathcal{N} \quad \text{s.t.} \quad x_i^* = 0 \tag{1.5}$$

2.3. Taxation and Budget Balance

The designer's task is to ensure the above optimum allocation is made. This clearly requires the knowledge of v_i even when constraints (1.2b) and (1.2c) are completely known. The premise of the problem, however, is that we are dealing with agents who are *strategic* and for whom, their valuation functions $\{v_i(\cdot)\}_{\mathcal{N}}$ is their private information. One way forward for the designer could be to simply ask each agent to report their private information and announce the solution of (1.2), with reported functions in place of v_i, for allocation. Apart from the fact that asking to report a function creates a practical communication problem, the main problem with this is that the agents could report untruthfully and end up getting a strictly better allocation (e.g., by reporting a v_i which has higher derivative than original at every point). In mechanism design terminology, the allocation function arising out of (1.2) isn't even partially implementable[1]. Restricting ourselves to a certain class of utility functions (quasi-linear utilities), provides the additional flexibility of penalising agents for reporting untruthfully by imposing taxes/subsidies. In this way, another related problem is created which is implementable, and which is equivalent to (1.2) as far as allocation is concerned. This leads us to the following additional assumption about agents' utilities:

(A5) All agents have quasi-linear utilities, i.e. the overall utility functions can be expressed as

$$u_i(x,t) = v_i(x_i) - t_i \qquad \forall \;\; i \in \mathcal{N} \tag{1.6}$$

where we have introduced taxes (or subsidies) $t = (t_i)_{i=1}^N \in \mathbb{R}^N$. Taxes affect utilities linearly and overall utility itself is valuation after adjustment for taxes (total monetary representation of one's state of happiness).

With taxation in the model, there are two reformulations of the NUM problem. One where taxes paid by agents go to the owner of the network (separate entity) and one where taxes are redistributed amongst agents who collectively own the network. It is easy to show that due to quasi-linear utilities the rate allocation problem is still same as (1.2). In addition, the first case would require that sum of taxes (which is the revenue for the network owner) be non-negative so that the owner of the network willingly participates. The second case would require that sum of taxes be zero. Henceforth we will refer to these reformulations as weak and strong budget balance, respectively.

In Section 3, we present a mechanism that fully implements (1.2) in Nash Equilibria (NE) with weak budget balance, while in Section 4 we will modify our mechanism to obtain strong budget balance.

[1]This can be deduced from the revelation principle. Indeed if there was a mechanism that even partially implements the allocation function arising out of (1.2), then there would exist also a truthful implementation. However, as shown with the above example, such an implementation will always fail.

3. Mechanism with Weak Budget Balance

We define a mechanism, in a way that doesn't require knowledge of $\{v_i(\cdot)\}$; whose game-form has NE in pure strategies where corresponding allocation for any NE is x^*, the unique optimizer of (1.2). In addition, agents are weakly better-off at NE than not participating at all.

Information assumptions Assume that $v_i(\cdot)$ is private information for agent i and nobody else knows it[2]. Let $\mathcal{I}_c$ be the set of common information between all agents, containing the information about full rationality of each agent. Finally, let $\mathcal{I}_d$ be the knowledge of the designer, containing the information about constraints (1.2b), (1.2c), the fact that $\mathcal{V}_i \subset \mathcal{V}_0,\ \forall i \in \mathcal{N}$.

3.1. Mechanism

Formally, we have an environment set $\mathcal{V} = \times_{i\in\mathcal{N}}\mathcal{V}_i$. We have seen from KKT, how each element of $\mathcal{V}$ can be mapped to an allocation x^* which maximises social welfare for that utility profile. In addition we also want tax to satisfy $\sum_{i\in\mathcal{N}} t_i \geq 0$. The designer defines a message space $\mathcal{S}_i$ for each agent $i \in \mathcal{N}$. Denote $\mathcal{S} = \times_{i\in\mathcal{N}}\mathcal{S}_i$ as the set of message profiles. The designer also announces the *contract* $h : \mathcal{S} \to \mathbb{R}_+^N \times \mathbb{R}^N$ that maps every message profile into an allocation vector and a tax vector. The designer then asks every agent $i \in \mathcal{N}$ to choose a message from the set $\mathcal{S}_i$, based on which allocations (and taxes) are made.

Specifically, agent i reports message $s_i = (y_i, p_i)$, $p_i = (p_i^l)_{l\in\mathcal{L}_i}$. This includes their demand for the good and the "price", for each constraint they are involved in, that they believe everyone should pay. This means $\mathcal{S}_i = \mathbb{R}_+ \times \mathbb{R}_+^{L_i}$. The mapping $s \mapsto (h_i(s)) = (h_{x,i}(s), h_{t,i}(s))$ $\forall\, i \in \mathcal{N}$ is as follows.

If the received demand vector is $y = (y_1, \ldots, y_N) = 0$ the allocation is $x = (x_1, \ldots, x_N) = 0$. Otherwise it is evaluated by first generating a scaling factor r

$$r = \min_{l\in\mathcal{L}} r^l, \tag{1.7}$$

[2]This is a realistic assumption, however, formally the use of Nash equilibrium as the solution concept requires complete information assumption amongst agents. Here we appeal to pre-game learning as a real-life technique for going from private to complete information.

$$r^l = \begin{cases} \frac{c^l}{\sum_{j\in\mathcal{N}^l} \alpha_j^l y_j}, & \text{if } |S^l(y)| \geq 2 \\ \frac{c^l}{\sum_{j\in\mathcal{N}^l} \alpha_j^l y_j} - f^l(y_i), & \text{if } S^l(y) = \{i\} \\ +\infty & \text{if } |S^l(y)| = 0 \end{cases} \tag{1.8}$$

with $f^l(y_i) = \frac{c^l}{\alpha_i^l y_i (y_i+1)}$. Using above quantities, the allocation and taxes are

$$h_{x,i}(s) = x_i = r y_i, \quad h_{t,i}(s) = t_i = \sum_{l\in\mathcal{L}_i} t_i^l \tag{1.9}$$

$$t_i^l = x_i \alpha_i^l \bar{p}_{-i}^l + (p_i^l - \bar{p}_{-i}^l)^2 + \eta\, \bar{p}_{-i}^l \left(p_i^l - \bar{p}_{-i}^l\right)\left(c^l - \sum_{j\in\mathcal{N}^l} \alpha_j^l x_j\right), \tag{1.10}$$

where η is a sufficiently small positive constant (formally chosen in the proof of Lemma 3.7). For any link $l \in \mathcal{L}_i$ we define

$$\bar{p}_{-i}^l \coloneqq \frac{1}{|\mathcal{N}^l \backslash \{i\}|} \sum_{j\in\mathcal{N}^l\backslash\{i\}} p_j^l = \frac{1}{N^l - 1} \sum_{j\in\mathcal{N}^l\backslash\{i\}} p_j^l. \tag{1.11}$$

The quantity $\bar{p}_{-i}^l$ is calculated by averaging the quoted prices for link l over all agents other than i who use that link (note that due to assumption (A3) every link has at least 2 agents). The interpretation of prices p_j^l in this mechanism is closely related to agent j's willingness to pay for consuming resource on link l. Since we have the problem of information elicitation for each agent's type (v_i), quoting of prices and demand is used as a way of eliciting $v_i'(x_i)$ by comparing it appropriately with prices.

The quantity $h_{x,i}(s)$ creates allocation by dilating/shrinking a given demand vector y on to one of the hyperplanes defined by the constraints in (1.2c), specifically, that hyperplane for which the corresponding allocation is the closest to origin (this can also happen to be at the intersection of multiple hyperplanes). The demand y is dilated/shrunk to the boundary of the feasible region defined by constraints in (1.2b), (1.2c). Additionally, the separate definition for r^l when $|S^l(y)| \leq 1$ is to ensure that there are no equilibria with $|S^l(y)| \leq 1$. This is required since we are only dealing with achieving solutions[3] to (1.2) which satisfy (A4).

The mechanism induces a one-shot game $\mathcal{G}$, played by agents in $\mathcal{N}$, with action set $(\mathcal{S}_i)_{i\in\mathcal{N}}$

[3]Note that for the given allocation function, $|S^l(y)| = 0, 1$ is equivalent to $|S^l(x)| = 0, 1$.

and utility

$$\hat{u}_i(s) = v_i(x_i) - t_i = v_i(h_{x,i}(s)) - h_{t,i}(s) \;\; \forall\, i \in \mathcal{N}. \tag{1.12}$$

3.2. Results

Theorem 3.1 (Full Implementation). *For game $\mathcal{G}$, there is a unique allocation, x, corresponding to all NE. Moreover, $x = x^*$, the maximiser of* (1.2). *In addition, individual rationality is satisfied for all agents.*

The theorem will be proved by a sequence of results, in which we will characterise all candidate NE of $\mathcal{G}$ by necessary conditions until we are left with only one family of NE candidates, each of whom gives allocation $x = x^*$. We will then show that there indeed exist NE in pure strategies for $\mathcal{G}$. Finally, individual rationality will be checked.

Lemma 3.2 (Primal Feasibility). *For any action profile $s = (y, P)$ of game $\mathcal{G}$, constraints* (1.2b) *and* (1.2c) *are satisfied for corresponding allocation $h_x(s)$.*

Proof. Please see Appendix 6.1 at the end of this chapter. □

Feasibility of allocation for action profiles is a direct consequence of using projections of demand y on to the feasible region. In the next result we will see that all agents using a link quote the same price for it at NE. This is brought about by the 2nd tax term $(p_i^l - \bar{p}_{-i}^l)^2$. Agents are penalized (with higher taxes) just for quoting a different price than average, at each link.

Lemma 3.3. *At any NE $s = (y, P)$ of $\mathcal{G}$, price for any link, is same for all agents, i.e., $p_i^l = p^l \;\; \forall\, i \in \mathcal{N}^l, \;\forall\, l \in \mathcal{L}$ (denote the common price at any link l by p^l).*

Proof. Please see Appendix 6.2 at the end of this chapter. □

With Lemma 3.3, we need only refer (in terms of prices) to the common price vector $p = (p^l)_{l \in \mathcal{L}}$ at NE. In fact we can identify each NE candidate profile $s = (y, P)$ with $s = (y, p)$. We will later see how p will take the place of dual variables λ when we compare equilibrium conditions with KKT conditions, hence we identify the following condition as dual feasibility.

Lemma 3.4 (Dual Feasibility). $p^l \geq 0 \quad \forall\, l \in \mathcal{L}$.

Proof. Please see Appendix 6.3 at the end of this chapter. □

Following is the property that solidifies the notion of prices as dual variables, since here we claim that inactive constraints do not contribute to payment at NE. This notion is very similar to the centralised problem, where if we know certain constraints to be inactive at the optimum then the same problem without these constraints would be equivalent to the original. The 3rd term in the tax function facilitates this property by charging extra taxes for inactive constraints when the agent is quoting higher prices than the average of remaining ones, thereby driving prices down.

Lemma 3.5 (Complimentary Slackness). *At any NE $s = (y, p)$ of game $\mathcal{G}$, allocation x satisfies*

$$p^l \Big(\sum_{k \in \mathcal{N}^l} \alpha_k^l x_k - c^l \Big) = 0 \quad \forall\, l \in \mathcal{L} \tag{1.13}$$

Proof. Please see Appendix 6.4 at the end of this chapter. □

Lemma 3.6 (Stationarity). *At any NE $s = (y, p)$ of game $\mathcal{G}$, and corresponding allocation, x, we have*

$$v_i'(x_i) = \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l \quad \forall\ i \in \mathcal{N} \quad \textit{if} \quad x_i > 0 \tag{1.14}$$

$$v_i'(x_i) \leq \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l \quad \forall\ i \in \mathcal{N} \quad \textit{if} \quad x_i = 0 \tag{1.15}$$

Proof. Please see Appendix 6.5 at the end of this chapter. □

Collecting above results, we see that every NE satisfies the KKT conditions. Therefore we have necessary conditions on the NE up to the point of having unique allocation. In the next Lemma we will verify the existence of the equilibria that we have claimed.

Lemma 3.7 (Existence). *For the game $\mathcal{G}$, there exists equilibria $s = (y, P)$, where price vectors are same for each agent i.e. $s = (y, p)$ and corresponding allocation-price pair (x, p) satisfy KKT with $\lambda = p$.*

Proof. Please see Appendix 6.6 at the end of this chapter. □

Lemma 3.8 (Individual Rationality). *At any NE $s = (y, p)$ of $\mathcal{G}$, with corresponding allocation x, we have*

$$u_i(x, t) \geq u_i(0, 0) \ \forall\, i \in \mathcal{N} \ \ \textit{and} \ \sum_{i \in \mathcal{N}} t_i \geq 0. \tag{1.16}$$

Proof. Please see Appendix 6.7 at the end of this chapter. □

All the Lemmas together complete the proof to Theorem 3.1, see Appendix 6.8.

3.3. Comments

Several comments are in order regarding the selection of the proportional allocation mechanism and in particular (1.8). If we use "pure" proportional allocation i.e. same expression for r^l for $|S^l(y)| \geq 2$ and ≤ 1, then irrespective of optimal solution of (1.2), for game $\mathcal{G}$ the "stationarity" property will not be satisfied for equilibria with $|S^l(y)| \leq 1$. Thus the mechanism will result in additional extraneous equilibria. For this reason we tweak the expression for r^l when $|S^l(y)| \leq 1$, so that we can eliminate these extraneous equilibria - irrespective of the solution of (1.2). With this tweak in the expression for r^l, all KKT conditions become necessary for all equilibria regardless of the value of $|S^l(y)|$. This however creates a problem in the proof of existence of equilibria. In particular, if x^* was such that it had links where $|S^l(x^*)| = 1$ then in our allocation this would require y at NE such that $|S^l(y)| = 1$. In this case the r^l used would be lower than what the proportional allocation requires (see second sub-case in (1.8)) and we actually would have the problem of possibly not having any y that creates x^* as allocation. Hence we have used (A4) to eliminate this case.

4. Mechanism with Strong Budget Balance

In this case we have the agents in $\mathcal{N}$, who are the owners as well as users that wish to allocate network resources amongst themselves in a way that maximises social welfare. Taxes facilitate efficient redistribution of an already available good. Here we will say that the mechanism fully implements maximising social welfare allocation if in addition to the previous conditions, we also have strong budget balance. To design a mechanism in this formulation, we have to find a way of redistributing the total tax paid by all the agents. From the last section, at NE total tax paid was

$$B = \sum_{i \in \mathcal{N}} \left(x_i \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l \right) = r \sum_{i \in \mathcal{N}} \left(y_i \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l \right) \tag{1.17}$$

since all other tax terms were zero at NE. We will redistribute taxes by modifying tax function for each agent in such a way that additional terms only use messages from other agents. This way, our NE calculations will be in line with Section 3, since deviations by an agent wouldn't affect his utility through this additional term. In view of this, we can express B as follows

$$B = r \sum_{i \in \mathcal{N}} \Big(\sum_{l \in \mathcal{L}_i} \frac{p^l}{N^l - 1} \sum_{j \in \mathcal{N}^l \setminus \{i\}} \alpha_j^l y_j \Big), \tag{1.18}$$

where each term of the outer summation depends only on demands of agents other than i. This means that each term in the parenthesis (scaled by the factor r) can now be used as the desired additional (negative) tax for user i. Observe, however, that r depends on each y_i. In lieu of this, the mechanism works by asking for an additional signal ρ_i from every agent and imposing an additional tax of $(\rho_i - r)^2$, thereby essentially ensuring that all agents agree on the value of r (via ρ_i's) at NE. Finally, we use $\bar{\rho}_{-i}$ (defined similarly to (1.11)) as a proxy for r in (1.18) - just like we did with $\bar{p}_{-i}^l$'s.

The message space $\mathcal{S}_i$ for agent i is $\mathbb{R}_+ \times \mathbb{R}_+^{L_i} \times \mathbb{R}_+$ with $s_i = (y_i, p_i, \rho_i)$. The designer announces the contract $h : \mathcal{S} \to \mathbb{R}_+^N \times \mathbb{R}^N$ where the only difference from before is in taxes.

$$\begin{aligned} h_{t,i}(s) = t_i &= \sum_{l \in \mathcal{L}_i} t_i^l + \zeta(\rho_i - r)^2 \\ t_i^l &= t_i^{l,\text{old}} - \frac{\bar{\rho}_{-i}\bar{p}_{-i}^l}{N^l - 1} \sum_{j \in \mathcal{N}^l \setminus \{i\}} \alpha_j^l y_j . \end{aligned} \tag{1.19}$$

where ζ, η are small enough positive constants. Here $\bar{p}_{-i}^l$ is as defined in (1.11) and $\bar{\rho}_{-i}$ similarly.

The only term in $\hat{u}_i$ that is affected by ρ_i is $-\zeta(\rho_i - r)^2$, so all the Lemmas from Section 3 will be valid with minor modifications. In addition ρ_i will all be equal to r at NE[4] thus proving the main result of full implementation (arguing similarly as before) as well as strong budget

[4]Please see the proof in Appendix 6.9 at the end of this chapter.

balance at NE[5]. Since we are only redistributing money, individual rationality continues to hold in this mechanism.

Among the lemmas from the previous section, only the proof of existence requires a non-trivial modification[6].

5. Summary and Comments

An important future research direction is to explicitly design a mechanism that is amenable to learning. This can be done by undertaking design in a way that makes a large class of learning (dynamic) algorithms stable. In such a case, each agent may chose his own learning strategy from this class independently of others. This way bounded rationality within learning can be explicitly modelled by creating an appropriate class whilst respecting strategic-ness of agents by not dictating how they should perform learning (i.e. without fixing the learning algorithm). Mechanism design along these lines can be seen in [HeMa12] for the case of Walrasian and Lindahl allocations and [Ch02] for Lindahl allocation.

[5] Please see the proof in Appendix 6.11 at the end of this chapter.
[6] Please see the proof of existence in Appendix 6.10 at the end of this chapter.

6. Appendix

6.1. Proof of Lemma 3.2

Proof. Constraint (1.2b) is clearly always satisfied. For $y = 0$, constraint (1.2c) is also clearly satisfied. We will now show (1.2c) is satisfied for any $y \neq 0$. In that case $r < +\infty$ (since there exists at least one link q with $|S^q(y)| \geq 1$ and thus $r^q < +\infty$). Now, for any link l, we have the following two cases. If $|S^l(y)| = 0$ then the allocation to agents on that link is clearly zero (since $x_i = ry_i$ and $y_i = 0$), so (1.2c) for those links is satisfied. If $|S^l(y)| \geq 1$ we have

$$\sum_{j \in \mathcal{N}^l} \alpha_j^l x_j = r \sum_{j \in \mathcal{N}^l} \alpha_j^l y_j \leq r^l \sum_{j \in \mathcal{N}^l} \alpha_j^l y_j \tag{1.20}$$

$$\leq \frac{c^l}{\sum_{j \in \mathcal{N}^l} \alpha_j^l y_j} \sum_{j \in \mathcal{N}^l} \alpha_j^l y_j = c^l \tag{1.21}$$

where the first inequality holds because r is the minimum of all r^l's. The second inequality will be equality if $|S^l(y)| \geq 2$ and will be strict only if $|S^l(y)| = 1$ (see second sub-case in eq. (1.8)). □

6.2. Proof of Lemma 3.3

Proof. Suppose there exist a link q for which prices $(p_i^q)_{i \in \mathcal{N}^q}$ at equilibrium are not all equal. Clearly then there is an agent $j \in \mathcal{N}^q$ for whom $p_j^q > \bar{p}_{-j}^q$ (this can be seen from eq. (1.11)). We will show that this agent can deviate by only reducing price for link q and be strictly better off, thereby contradicting the equilibrium condition.

Take the deviation by agent j, where $p'^q_j = \bar{p}_{-j}^q < p_j^q$ and his demand remains the same. Then the difference between utilities for agent j after and before deviation would arise only because of change in taxes (since allocations haven't changed for any agent) and moreover the difference would only arise from tax terms corresponding to link q (refer to (1.9)). We'll have

$$\Delta \hat{u}_j = -(p'^q_j - \bar{p}_{-j}^q)^2 - \eta\, \bar{p}_{-j}^q (p'^q_j - \bar{p}_{-j}^q)(c^q - \sum_{k \in \mathcal{N}^q} \alpha_k^q x_k)$$

$$+\,\eta\,\bar{p}^q_{-j}(p^q_j-\bar{p}^q_{-j})(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)+(p^q_j-\bar{p}^q_{-j})^2 \tag{1.22a}$$

$$=\underbrace{(p^q_j-\bar{p}^q_{-j})}_{>0}\Big[\eta\,\bar{p}^q_{-j}\underbrace{(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)}_{\geq 0\ by\ Lemma\ 3.2}+\underbrace{(p^q_j-\bar{p}^q_{-j})}_{>0}\Big] \tag{1.22b}$$

$$>0 \tag{1.22c}$$

which shows that the above deviation is a profitable one.

Hence at equilibrium, for any link, the price quoted for that link by any user using that link is the same, we denote the common price vector by $p=(p^l)_{l\in\mathcal{L}}$. □

6.3. Proof of Lemma 3.4

Proof. This is also by design, since any agent i is asked to select a price vector in $\mathbb{R}^{L_i}_+$. □

6.4. Proof of Lemma 3.5

Proof. Suppose there is a link q for which at NE $\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k<c^q$ and $p^q>0$. Again, we will show that deviation makes any agent $j\in\mathcal{N}^q$ better off. Take the deviation ${p'}^q_j=p^q-\epsilon>0$ and no deviation in the demand. We can write the difference after and before deviation as

$$\begin{aligned}\Delta\hat{u}_j&=-({p'}^q_j-p^q)^2-\eta\,p^q({p'}^q_j-p^q)(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)\\&\quad+\eta\,p^q(p^q-p^q)(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)+(p^q-p^q)^2\end{aligned} \tag{1.23}$$

$$\Rightarrow\quad \Delta\hat{u}_j=-(-\epsilon)^2-\eta\,p^q(-\epsilon)(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)+0=\epsilon\Big(-\epsilon+\eta\,p^q(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)\Big) \tag{1.24}$$

So if we take ϵ such that

$$\min\Big\{\underbrace{\eta\,p^q}_{>0}\underbrace{(c^q-\sum_{k\in\mathcal{N}^q}\alpha^q_k x_k)}_{>0},\quad\underbrace{p^q}_{>0}\Big\}>\epsilon>0 \tag{1.25}$$

then $\Delta\hat{u}_j>0$. Taking such an ϵ is possible because LHS above is positive. □

6.5. Proof of Lemma 3.6

Proof. At any NE s, agent i's utility $\hat{u}_i(s_i', s_{-i}) = v_i(h_{x,i}(s_i', s_{-i})) - h_{t,i}(s_i', s_{-i})$ as a function of his message $s_i' = (y_i', p_i')$, with s_{-i} fixed, should have a global maximum at $s_i = (y_i, p)$. This would mean that if this function was differentiable w.r.t. y_i' at s_i, the partial derivatives w.r.t. y_i' at s_i should be 0. However, since our allocation dilates/shrinks demand vector y' on to the feasible region, it could be the case that increasing and decreasing y_i' gives allocations lying on different hyperplanes, meaning that the transformation from y' to x' is different on both sides of y_i and therefore $\hat{u}_i$ may not be differentiable w.r.t y_i' at y_i. The important thing here however is to notice that right and left derivatives exist, it's just that they may not be equal. Hence we can take derivatives on both sides of y_i as (noting that derivative of the other two terms in utility will be zero due Lemma 3.3)

$$\left.\frac{\partial \hat{u}_i}{\partial y_i'}\right|_{y_i' \downarrow y_i} = \left(v_i'(x_i) - \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l \right) \left.\frac{\partial x_i'}{\partial y_i'}\right|_{y_i' \downarrow y_i}, \tag{1.26}$$

$$\left.\frac{\partial \hat{u}_i}{\partial y_i'}\right|_{y_i' \uparrow y_i} = \left(v_i'(x_i) - \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l \right) \left.\frac{\partial x_i'}{\partial y_i'}\right|_{y_i' \uparrow y_i}. \tag{1.27}$$

We will first show that $\partial x_i / \partial y_i$ term above (for either equation) is always positive. If $y = 0$ then clearly this is true, because if agent i demands $y_i = \epsilon > 0$ while $y_{-i} = 0$ then $x_i > 0$ (in fact the allocation is indeed differentiable at $y = 0$). If $y \neq 0$ then from (1.9), we can write

$$\beta := \frac{\partial x_i}{\partial y_i} = \frac{\partial (r y_i)}{\partial y_i} = r + y_i \frac{\partial r}{\partial y_i} = r^q + y_i \frac{\partial r^q}{\partial y_i} \tag{1.28}$$

where $r = r^q$.

From here on there are 3 cases: (A) $i \notin \mathcal{N}^q$; (B1) $i \in \mathcal{N}^q$ & $|S^q(y)| \geq 2$ and (B2) $i \in \mathcal{N}^q$ & $|S^q(y)| = 1$. (Note that $|S^q(y)| = 0$ isn't possible if $r = r^q$ and $y \neq 0$)

(A) Here clearly $\partial r^q / \partial y_i = 0$, and this makes $\beta = r^q > 0$.

(B1) Here we can calculate β as

$$\beta = \frac{(r^q)^2}{c^q} \sum_{j \in \mathcal{N}^q \setminus \{i\}} \alpha_j^q y_j \tag{1.29}$$

which is positive because $|S^l(y)| \geq 2$, so there is at least one positive term in the summation.

(B2) In this case we could have $S^q(y) = \{j\} \neq \{i\}$ or $S^q(y) = \{i\}$. In the first case, argument is same as (A). So take $S^q(y) = \{i\}$ and $r^q = c^q/(\alpha_i^q y_i) - f^q(y_i)$.

$$\beta = -f^q(y_i) - y_i \frac{\mathrm{d}f^q(y_i)}{\mathrm{d}y_i} = -\frac{\mathrm{d}\left(y_i f^q(y_i)\right)}{\mathrm{d}y_i} = \frac{c^q}{\alpha_i^q (y_i+1)^2}. \tag{1.30}$$

So we have $\beta > 0$ in all cases.

Now there are two further cases, the first term on RHS in both equations in (1.26) is positive or negative. If it's positive, then we can see from the first equation in (1.26) that by increasing y_i' from y_i (and therefore x_i' from x_i) agent i can increase his pay-off, which contradicts equilibrium. Now similarly consider the first term in (1.26) to be negative, then from the second equation in (1.26), agent i can reduce y_i' from y_i to get a better pay-off. But the downward deviation in y_i' is only possible if $y_i > 0$ ($\Leftrightarrow \quad x_i > 0$). So we conclude that

$$v_i'(x_i) = \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l, \quad \forall \ i \in \mathcal{N} \quad \text{if} \quad x_i > 0 \tag{1.31}$$

$$v_i'(x_i) \leq \sum_{l \in \mathcal{L}_i} p^l \alpha_i^l, \quad \forall \ i \in \mathcal{N} \quad \text{if} \quad x_i = 0. \tag{1.32}$$

□

6.6. Proof of Lemma 3.7

Proof. The proof is completed in two parts. Firstly we will check that for every x that can be a possible solution to (1.2) (while satisfying the assumptions, specifically (A4)) there is indeed at least one $y \in \mathbb{R}_+^N$ such that the allocation corresponding to y is x. (However we do not need to check the same for prices and Lagrange multipliers (λ) since there it is straightforward). Secondly we will check that for the claimed NE $s = (y, P)$, there are no unilateral deviation that are profitable.

In lieu of (A4), the optimal x^* is such that $|S^l(x^*)| \geq 2$ for all links; also it is clear that x^* is on the boundary of the feasible region defined by (1.2b) and (1.2c). For this, any vector y which is a scalar multiple of x^* will give allocation x^*. This completes the first part of the proof.

Now we will check for profitable deviations. For this we want to show that at any action profile where corresponding allocation is x^* and prices are equal and equal to λ^*, is a NE. Note that by construction, $\hat{u}_i(s)$ is continuous w.r.t. s_i. Our approach would be to check first and second order conditions for local maximum at the claimed NE points. In this process we will ensure that at any point where the gradient is 0, the Hessian is negative definite. With this we know that all extremum points are local maxima, which can only happen if there is one maximum.

First we will show that without the gradient being zero, there cannot be a local extremum (this needs to be shown here separately since $\hat{u}_i(s)$ is only piecewise differentiable w.r.t. s_i). Gradient of the utility function is

$$\nabla \hat{u}_i = \left(\frac{\partial \hat{u}_i}{\partial y_i}, \left(\frac{\partial \hat{u}_i}{\partial p_i^l} \right)_{l \in \mathcal{L}_i} \right) \tag{1.33}$$

Components in the gradient are

$$\frac{\partial \hat{u}_i}{\partial y_i} = \left(v_i'(x_i) - \sum_{l \in \mathcal{L}_i} \alpha_i^l \bar{p}_{-i}^l \right) \left(\frac{\partial x_i}{\partial y_i} \right) - \eta \sum_{l \in \mathcal{L}_i} \bar{p}_{-i}^l (p_i^l - \bar{p}_{-i}^l) \left(- \sum_{j \in \mathcal{N}^l} \alpha_j^l \frac{\partial x_j}{\partial y_i} \right) \tag{1.34}$$

$$\frac{\partial \hat{u}_i}{\partial p_i^l} = -2(p_i^l - \bar{p}_{-i}^l) - \eta \bar{p}_{-i}^l (c^l - \sum_{j \in \mathcal{N}^l} \alpha_j^l x_j), \quad \forall\, l \in \mathcal{L}_i \tag{1.35}$$

Since $\hat{u}_i$ is differentiable w.r.t. p_i^l (at all points), that component of the gradient indeed has to be 0 at a local extremum. Which implies equal prices and complimentary slackness properties and we can write

$$\frac{\partial \hat{u}_i}{\partial y_i} = \left(v_i'(x_i) - \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l \right) \left(\frac{\partial x_i}{\partial y_i} \right) \tag{1.36}$$

Note that in the proof of Lemma 3.6, we have shown that $\beta := \frac{\partial x_i}{\partial y_i} > 0$ always. So at the points of non-differentiability, β will have a jump discontinuity, however it will be positive on either side. It is then clear that without $v_i'(x_i) = \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l$, there cannot be a local extremum.

Going towards checking first and second order conditions at claimed NE points, we clearly have price derivatives (see (1.35)) equal to 0 at equilibrium, due to Lemmas 3.3 and 3.5. Now the second term in (1.34) is always zero at equilibrium (due to common prices). So if $x_i > 0$ at equilibrium then the 1st term is also zero as well (due to Stationarity) and if $x_i = 0$ ($\Leftrightarrow$

$y_i = 0$) the 1st term is either negative, making the whole derivative negative (this is fine since downward deviation isn't possible from $y_i = 0$), or equal to zero. For the remaining of the proof we will consider the case where the first order derivatives are equal to zero and consider second order derivatives.

Second order partial derivatives are

$$u_{pp} := \frac{\partial^2 \hat{u}_i}{\partial p_i^l \partial p_i^l} = -2 \qquad u_{lk} := \frac{\partial^2 \hat{u}_i}{\partial p_i^l \partial p_i^k} = 0 \qquad u_{py} := \frac{\partial^2 \hat{u}_i}{\partial p_i^l \partial y_i} = \eta\, \bar{p}_{-i}^l \left(\sum_{j \in \mathcal{N}^l} \alpha_j^l \frac{\partial x_j}{\partial y_i} \right) \tag{1.37a}$$

$$u_{yy} := \frac{\partial^2 \hat{u}_i}{\partial y_i^2} = \left(v_i'(x_i) - \sum_{l \in \mathcal{L}_i} \alpha_i^l \bar{p}_{-i}^l \right) \left(\frac{\partial^2 x_i}{\partial y_i^2} \right) + v_i''(x_i) \left(\frac{\partial x_i}{\partial y_i} \right)^2 - \eta \sum_{l \in \mathcal{L}_i} \bar{p}_{-i}^l (p_i^l - \bar{p}_{-i}^l) \left(- \sum_{j \in \mathcal{N}^l} \alpha_j^l \frac{\partial^2 x_j}{\partial y_i^2} \right). \tag{1.37b}$$

These derivatives will give us a Hessian H of size $(L_i+1) \times (L_i+1)$, where 1st row and column represent y_i and subsequent rows and columns represent p_i^l's for different l's. We want H to be negative definite at equilibrium. Now, 1st and 3rd terms in u_{yy} are zero at equilibrium, and the 2nd term is strictly negative due to strict concavity of v_i. This along with $u_{pp} = -2$ tells us that all diagonal entries in H are negative. Also notice that because of $u_{lk} = 0$, all off-diagonal entries other than the ones in first row and column are zero. Finally, note that due to assumption (A2), all prices are finite at equilibrium and so u_{py} will be finite. We will show that roots of the characteristic polynomial of H (i.e. its eigenvalues) all become negative if η is chosen sufficiently small. Here again, we use (A4). Since we only want to deal with NE which satisfies $|S^l(y)| \geq 2 \;\; \forall\, l$, in the following it is inherently assumed that we will be working with r that has the form defined in the first sub-case of (1.8).

For this, we take a generic matrix $A = \{a_{ij}\}$, which is similar in structure to H and has the same dependence on $|y|$ as H. So entries in A are

$$a_{11} = -\frac{a}{|y|^2} \quad a_{ij} = a_{ji} = 0 \;\; \forall \;\; i,j > 1, \;\; i \neq j \tag{1.38}$$

$$a_{ii} = -2 \quad a_{1i} = a_{i1} = \eta \, \frac{b_{i-1}}{|y|} \quad \forall \;\; 2 \leq i \leq L_i + 1 \tag{1.39}$$

where $a > 0$ (and we don't care about the sign of b_i's). The parameters a, b_i may not be completely independent of y but since the absolute value of y has been taken out of the scaling, their values are bounded. Magnitude of b_i's are bounded from above and a is bounded away from zero.

We can explicitly calculate $|A - \lambda I|$ and write the characteristic equation as

$$\begin{aligned} Q(\lambda) = & \left(-\frac{a}{|y|^2} - \lambda\right)(-2-\lambda)^{L_i} \\ & + \eta^2 \, \frac{\sum_{i=1}^{L_i}(-1)^i b_i^2}{|y|^2}(-2-\lambda)^{L_i - 1} = 0 \end{aligned} \tag{1.40}$$

So -2 is a repeated eigenvalue, $L_i - 1$ times. The equation for the remaining two roots can be written as

$$\left(-\frac{a}{|y|^2} - \lambda\right)(-2-\lambda) + \eta^2 \frac{C}{|y|^2} = 0 \tag{1.41}$$

Necessary and sufficient conditions for both roots of this quadratic to be negative are

$$\left(2 + \frac{a}{|y|^2}\right) > 0 \qquad \frac{2a}{|y|^2} + \eta^2 \frac{C}{|y|^2} > 0, \tag{1.42}$$

first of which is always true, since $a > 0$. The second one can be ensured by making η small enough, since a is bounded away from zero and magnitude of C is bounded from above.

Hence we have shown the Hessian H to be negative definite for η chosen to be small enough. $\square$

6.7. Proof of Lemma 3.8

Proof. Because of Lemma 3.3, the only non-zero term in t_i at equilibrium is $x_i \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l$, which is clearly non-negative. Hence $\sum_{i \in \mathcal{N}} t_i \geq 0$ at equilibrium. This is the seller's individual rationality condition.

Now if $x_i = 0$ then we know from Lemma 3.3 and (1.9) that $t_i = 0$ and so (1.16) is evident.

Now take $x_i > 0$ and define the function

$$f(z) = v_i(z) - z \sum_{l \in \mathcal{L}_i} \alpha_i^l p^l. \tag{1.43}$$

Note that $f(0) = u_i(0,0)$ and $f(x_i) = u_i(x,t)$, the utility at equilibrium. Since $f'(x_i) = 0$ (Lemma 3.6), we see that $\forall\ 0 < y < x_i$, $f'(y) > 0$ since f strictly concave (because of v_i). This clearly tells us $f(x_i) \geq f(0)$. $\square$

6.8. Proof of Theorem 3.1

Proof. We know that the four KKT conditions produce a unique solution x^* (and corresponding λ^*). For the game $\mathcal{G}$, from Lemmas 3.2–3.6 we can see that at any NE, allocation x and prices p satisfy the same conditions as the four KKT conditions and hence they give a unique $x = x^*$, as long as (A4) is satisfied. So we have that the allocation is x^* across all NE. This combined with individual rationality Lemma 3.8, proves Theorem 3.1. $\square$

6.9. Proof that $\rho_i = r, \forall\, i \in \mathcal{N}$ at NE

Proof. Suppose not, i.e. assume $\exists\, j \in \mathcal{N}$ such that $\rho_j \neq r$. In this case agent j can deviate with only changing $\rho_j' = r$ (which also means r is the same as before deviation, since demand y doesn't change). It's easy to see that this is a profitable deviation, since change in utility of agent j will be only through the term involving ρ_j.

$$\Delta \hat{u}_j = -\zeta(\rho_j' - r)^2 + \zeta(\rho_j - r)^2 = \zeta(\rho_j - r)^2 > 0. \tag{1.44}$$

$\square$

6.10. Outline of proof of NE existence in Section 4

Proof. Now we verify the existence of equilibria. The arguments here will be similar to the ones in the proof of Lemma 3.7. First order conditions can again be shown to be satisfied, the only difference is that here we will also use $\rho_i = r$ at equilibrium. The Hessian H here, for agent i, will be of order $(L_i + 2) \times (L_i + 2)$ where 1st, 2nd row and column represent

y_i, ρ_i respectively whereas the remaining rows and columns represent p_i^l's. The generic matrix $A = \{a_{ij}\}$ for H will then be

$$a_{11} = -\frac{a}{|y|^2} - \zeta\frac{d}{|y|^4} \qquad a_{12} = a_{21} = -\zeta\frac{e}{|y|^2} \tag{1.45}$$

$$a_{ij} = a_{ji} = 0 \quad \forall \quad i, j > 1, \;\; i \neq j \tag{1.46}$$

$$a_{22} = -2, \quad a_{ii} = -2, \tag{1.47}$$

$$a_{1i} = a_{i1} = \eta\frac{b_{i-1}}{|y|} \;\; \forall\, 3 \leq i \leq L_i + 2 \tag{1.48}$$

where $a, d, e > 0$. Writing the characteristic equation we will again get that -2 is a repeated eigenvalue, L_i times. And the equation for remaining two roots is

$$\lambda^2 + \lambda\left(2 + \frac{a}{|y|^2} + \zeta\frac{d}{|y|^4}\right) + \left(\frac{2a}{|y|^2} + \zeta\frac{2d}{|y|^4} - \zeta^2\frac{e^2}{|y|^4} + \eta^2\frac{C}{|y|^2}\right) = 0 \tag{1.49}$$

Necessary and sufficient conditions for the roots of above quadratic to be negative are again that coefficient of λ and the constant term are both positive. Coefficient of λ is clearly positive, and the constant term can also be made positive by choosing ζ, η small enough, irrespective of sign of C. Hence here also we get NE for all y (along a fixed direction) for ζ, η chosen to be small enough. □

6.11. Proof of Strong Budget Balance at NE

Proof. We now know that price vectors are equal at equilibrium for all agents and so we can write

$$\sum_{i\in\mathcal{N}} t_i = \sum_{i\in\mathcal{N}} x_i\left(\sum_{l\in\mathcal{L}_i} \alpha_i^l p^l\right) - r\sum_{l\in\mathcal{L}_i} \frac{p^l}{N^l - 1} \sum_{j\in\mathcal{N}^l\setminus\{i\}} \alpha_j^l y_j \tag{1.50}$$

$$\Rightarrow \quad \sum_{i\in\mathcal{N}} t_i = \sum_{i\in\mathcal{N}}\sum_{l\in\mathcal{L}_i}\left(x_i\alpha_i^l p^l - \frac{p^l}{N^l - 1}\sum_{j\in\mathcal{N}^l\setminus\{i\}} \alpha_j^l x_j\right) \tag{1.51}$$

Consider the coefficient of x_k for any agent k in the above expression

$$\sum_{l \in \mathcal{L}_k} \left(\alpha_k^l p^l - \sum_{i \in \mathcal{N}^l \setminus \{k\}} \frac{p^l}{N^l - 1} \alpha_k^l \right) = \sum_{l \in \mathcal{L}_k} \left(\alpha_k^l p^l - \frac{p^l \alpha_k^l}{N^l - 1} \left(N^l - 1 \right) \right) = 0, \tag{1.52}$$

which proves the claim. □

Chapter 2.

Mechanism Design for Resource Allocation in Networks with Intergroup Competition and Intragroup Sharing

We consider a network where strategic agents, who are contesting for allocation of resources, are divided into fixed groups. The network control protocol is such that within each group agents get to share the resource and across groups they contest for it. A prototypical example is the allocation of data rate on a network with multicast/multirate architecture. Compared to the unicast architecture (which is a special case), the multicast/multirate architecture can result in substantial bandwidth savings. However, design of a market mechanism in such a scenario requires dealing with both private and public good problems as opposed to just private goods in unicast.

The mechanism proposed in this work ensures that social welfare maximizing allocation on such a network is realized at all Nash equilibria (NE) i.e., full implementation in NE. In addition it is individually rational, i.e., agents have an incentive to participate in the mechanism. The mechanism, which is constructed in a quasi-systematic way starting from the dual of the centralized problem, has a number of useful properties. Specifically, due to a novel allocation scheme, namely "radial projection", the proposed mechanism results in feasible allocation even off equilibrium. This is a practical necessity for any realistic mechanism since agents have to "learn" the NE through a dynamic process. Finally, it is shown how strong budget balance at equilibrium can be achieved with a minimal increase in message space as an add-on to a

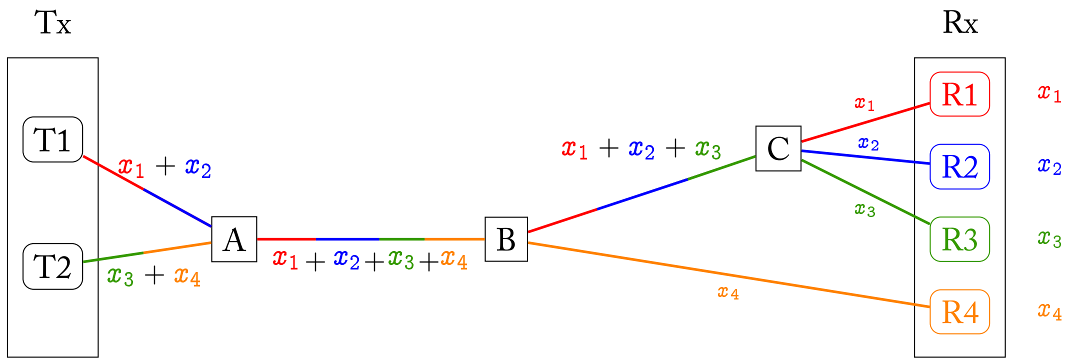

Figure 7.: Network with "unicast" architecture

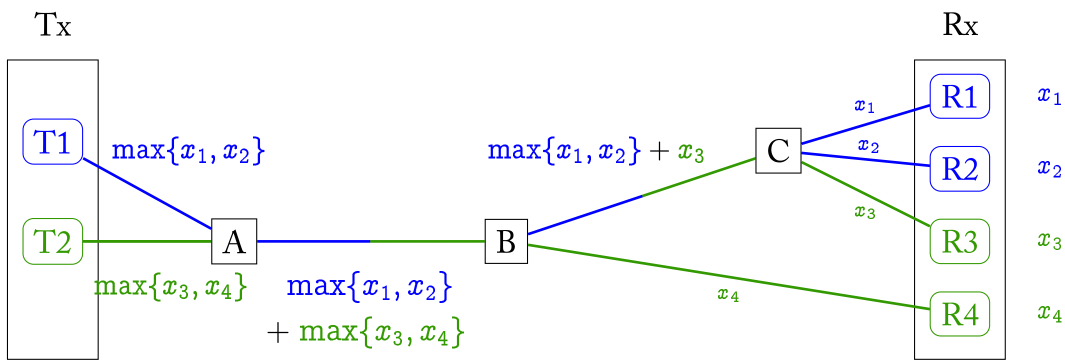

Figure 8.: Network with "multicast/multirate" architecture

weakly budget balanced mechanism.

1. Introduction

Design of mechanisms that fully implement Walrasian and Lindahl allocations in Nash equilibrium (NE) have been extensively studied in the literature, e.g. [GrLe77; Hu79; Ti89; PoWe89; Hu94; MaSj02]. These two ubiquitous examples present two different aspects in market design - private and public goods, respectively. Recently, generalizing the original contributions described above, a number of works addressed the problem of full implementation in NE of social utility maximization under linear inequality constraints [MaBa04; StLe06; JaWa10; KaTe14; BhKaChGu14; SiAn15]. In the area of computer/communication networks, the prototypical example of this problem is the allocation of rate (bandwidth) in a "unicast" network architecture on the Internet, as depicted in Fig. 7. In this context, an agent is a receiver (e.g., $\{\mathrm{Ri}\}_{\mathrm{i}=1,2,3,4}$ in Fig. 7), who communicates with his respective transmitter via a fixed route consisting of links (e.g., R2 communicates with T1 via the route consisting of the T1-A, A-B, B-C, C-R2 links, whereas R3 communicates with T2 via the T2-A, A-B, B-C, C-R3 links). The scenario depicted in Fig. 7 is such that receivers R1, R2 request the same data (e.g., the same movie), which is transmitted by transmitter T1. Similarly R3, R4 request the same data which is transmitted by transmitter T2. The term "unicast" refers to the fact that the network establishes separate connections for each agent in each link (even when agents communicate the same content), and thus each agent loads each link with the amount of bandwidth it is allocated. The inequality constraints quantify the capacity constraints at each link (e.g., in link A-B, $x_1 + x_2 + x_3 + x_4 \leq c_{AB}$, where c_{AB} is the capacity of link A-B and x_i's are the allocated rates for the four agents sharing this link). Thus the unicast problem is a pure private goods problem with one good (i.e., rate) to be allocated and multiple constraints on the same good (i.e., one capacity constraint per link).

There are however a number of interesting problems in economics and engineering that do not fit the above general model, as they involve agents forming groups and entail intergroup competition and intragroup sharing. We present this class of problems in the context of communication networks and in particular for the problem of rate allocation in a "multicast/multirate" architecture on the Internet, as depicted in Fig. 8. The model considered in this chapter is such that for each receiver there is exactly one transmitter who communicates with it, whereas for each transmitter there can be multiple receivers who communicate with it. Similarly to the unicast scenario, an agent is a receiver. Agents form multicast groups based on the content they communicate (e.g., {R1, R2} form a group and so do agents {R3, R4}). At the same time, agents within a group may request different bandwidth for the same content, where this differentiation is due to different quality of service requested by the users within a group, such as in high- vs standard-definition video. In the multicast protocol, in each link only a single connection is established for each multicast group carrying the corresponding content at the highest requested rate. This is motivated by the fact that any lower-quality content can always be derived from the higher-quality content and thus there is no need to further load the link other than with the highest rate (e.g., agents R1, R2, R3, R4 all communicate via link A-B, but only two distinct contents are transmitted on it, resulting in a capacity constraint of the form $\max\{x_1, x_2\} + \max\{x_3, x_4\} \leq c_{AB}$). Such a unique architecture makes the multicast/multirate allocation problem qualitatively different from the unicast since there is sharing of bandwidth as well as competition for it.

A second motivating scenario comes from the provision of data-security on server farms. Consider multiple server farms, each hosting data for several companies. Data security at each server farm is provided by simultaneous use of different security products out of a set $\mathcal{L}$ of possible products. Due to different companies having different security needs, for any company i in server farm k achieving security level x_{ki} requires a profile $\{\alpha_{ki}^{l} x_{ki}\}_{l \in \mathcal{L}}$ of quantities of different security products[1], where α_{ki}^{l} are positive constants denoting the effectiveness of product l for company i on server farm k. However, since at each server farm the effectiveness of the security product is not additive, the quantity of any security product required at a server farm is dictated only by the maximum quantity/quality of that security product demanded amongst companies in that farm. Finally, allocation of security products among server farms is constrained by the limited quantity of each security product being available.

Looking at the structure of both these problems, one sees both a private and a public goods aspect of market design in them. Referring to the multicast problem, due to the capacity constraints on the links of the network, allocation of rate for one content implies that such rate cannot be allocated to agents sharing this link, who have requested a different content - this is the private good aspect resulting in intergroup competition. On the other hand for agents requesting the same content, since the allocation via the capacity constraint is dictated only by the highest rate from that group, others in the group can be allocated additional rate (up to the maximum requested in that group on that link) without having to affect anybody else's allocation - this is the public goods aspect resulting in intragroup sharing and the inevitable "free-rider" problem [MaWhGr95, sec. 11.C], albeit in a problem where consumption is still private.

The goal of this chapter is to design appropriate incentives, through allocation and taxes, such that when acting strategically i.e. at NE, the corresponding allocation to agents is efficient. This efficient allocation is the solution to a convex optimization problem: maximization of sum of agents' utilities subject to network multicast constraints. The taxes imposed by the mechanism refer to actual money paid by agents and not to virtual signals, as may be the

[1]The final security level can be thought of as being achieved by use of different security products via a Leontief-like production function [MaWhGr95, p. 49].

case in works on distributed optimization. Therefore our model assumes the presence of an authority capable of collecting taxes from the agents or disbursing subsidies to the agents (as dictated by the mechanism).

For the problem of designing a mechanism that leads to a single-shot game, the well-accepted measure of complexity is the dimensionality of the message space of the proposed mechanism. Similar to some of the works mentioned earlier, the mechanism presented here requires agents to communicate via announcing relatively "small" signals, which only consist of demands and prices. However, contrary to some of the earlier works in communications, such as [JaWa10; YaHa07; JoTs09], the work here ensures full implementation of social welfare maximizing allocation, so the designer can guarantee that only the most efficient outcome will be reached and no other. This is the first contribution of this work. Note that we avoid using direct mechanisms (see [NiRoTaVa07; BöKrSt15] for definition) since for this problem agents' private information is their utility function and it would be impractical to ask agents to quote entire functions. This also means that we do not use such terms as "incentive compatibility" and "truth-telling" as they apply only to direct mechanisms.

In addition to the stated goal of getting optimal allocations at all Nash equilibria (NE) and ensuring individual rationality (IR), the mechanism presented here has two auxiliary properties. The first auxiliary property, and the second contribution of this work, is to achieve *feasibility on and off equilibrium* via a "radial projection" allocation function. This property implies that the contract promises to allocate rates to agents in such a way that the link capacity constraints are never violated, contrary to the works in [StLe06; KaTe14; KaTe13b; ShTe12]. While, clearly, feasibility is satisfied by definition at NE, the proposed mechanism achieves it even off equilibrium. Off-equilibrium feasibility has much deeper implications in the context of networks as described below. The use of NE as a solution concept in the implementation literature can lead to problems in practice since calculation of NE requires full information on part of the agents (which may not happen on an informationally and physically decentralized system like the Internet). Thus the justification for using NE (even for single shot games) is that for certain classes of learning dynamics, repeated play with the mechanism will ensure that the NE is "learned" eventually. Since information capacity constraints on a network are hard constraints and cannot be violated at any cost (in the short run), the above "learning" justification is applicable only when allocation is feasible throughout the learning process. Hence off-equilibrium feasibility becomes a necessity and the mechanism presented in this work is "learning ready" in that respect.

Regarding the specific technique used to achieve off-equilibrium feasibility, radial projection refers to agents' demands being converted into allocations by scaling of the overall demand vector to the boundary of the feasible region defined by the capacity constraints. This allocation method subsumes as a special case the allocation function introduced in [MaBa04; YaHa05] for communication on a single-link unicast network, and for stochastic control of networks in [KeMaTa98] (readers may be refer to [SiAn15] for a full implementation mechanism that uses a similar allocation concept in the unicast problem).

The second auxiliary property of the proposed mechanism and the final contribution of this work is to demonstrate how strong budget balance (SBB) at NE can be added to the non-budget balanced mechanism with the exchange of an extra signal (see Section 4). This is in contrast to [KaTe13b], where significant effort has been made to ensure budget balance off equilibrium for feasible allocations, while feasibility itself isn't ensured off equilibrium. With the proposed modification, achieving SBB at NE becomes a relatively straightforward task. This is to be contrasted with the mechanisms proposed in [YaHa07; JoTs09] where SBB becomes a very difficult property to guarantee.

The relevant literature on mechanism design is discussed below.

The earliest unicast mechanisms were inspired by the seminal work [KeMaTa98], which presents a mechanism with competitive equilibrium as the solution concept. Mechanism design work for unicast network include [MaBa03; MaBa04; JoTs04; YaHa05; JaWa10], and specifically, Nash implementation for general unicast has been studied in [StLe06; KaTe14; SiAn15]. In all the above works, the network structure is assumed to be fixed; alternately [Ne09] studies the relay problem where agents own links and can price the data going through them.

Mechanism design has also been studied for power allocation problems in wireless networks. These are public goods problems, since the quality of service achieved by any agent depends not only on his signal strength but also that of neighboring agents due to interference from surrounding transmissions (see [HoGa12; ShTe12; AlBoHoPo13b]).

Readers may refer to [Ch02; HeMa12] for mechanisms designed to fully implement Walrasian and Lindahl allocations that are guaranteed to converge to the NE with a large class of learning dynamics. Both [Ch02; HeMa12] rely on the off-equilibrium properties - specifically that of the best response correspondence - to ensure a "learning" property in their mechanism. Authors in [BhKaChGu14] ensure that their mechanism induces a best response correspondence that has a partial truth-telling property: the allocation and price (x_i, p_i) quoted by agent i, in her best response, are related to her utility $v_i(\cdot)$ as $p_i = v_i'(x_i)$ (even if x_i isn't equal to the optimal allocation x_i^*).

In multicast/multirate models without strategic agents, researchers have argued for different optimality criteria. One example is the *max-min* fairness for multicast used in [RuKuTo99; GrSrTo01; SaTa99; SaTa02]. On the other hand, the works in [KaSaTa01; KaTa06] have used integer and convex programming to get decentralized algorithms that maximize the sum of utilities. Stoenescu et al [StLiTe07] propose a realization algorithm which converges to optimal allocation.

The remainder of this chapter is structured as follows: in Section 2, the centralized problem that we wish to implement is stated and its solution is characterized. In Section 3, the mechanism is described and its properties are derived for the weak budget balance (WBB) case. This mechanism is modified in Section 4 to include strong budget balance at NE. Section 5 discusses the relevant literature and some salient features of the mechanism.

2. Centralized Problem

For concreteness, the exposition follows the prototypical example of rate allocation in a multicast/multirate network architecture on the Internet. The system consists of a set $\mathcal{N}$ of Internet agents who communicate over a fixed multicast network. Each agent here is considered as a pair of source and destination users and the agents are divided into disjoint multicast groups based on the content that they communicate. The set of agents is described as $\mathcal{N} = \{(k,i) \mid k \in \mathcal{K}, i \in \mathcal{G}_k\}$, where the set of multicast groups is denoted by $\mathcal{K} = \{1, 2, \ldots, K\}$ and within a group $k \in \mathcal{K}$, the set of agents is denoted by $\mathcal{G}_k = \{1, \ldots, G_k\}$. We use the notation ki instead of (k,i) to denote a generic agent.

We denote by $x_{ki} \geq 0$ the rate allocated to an agent ki and it refers to the data-rate that agent ki communicates with (with the data content being the same as that of any other agent from group k). Thus the overall allocation for the system is a vector $x = (x_{ki})_{ki \in \mathcal{N}}$ of rates allocated to all the agents.

Agents' have private valuation and thus for any agent ki, the valuation for allocation x_{ki} is $v_{ki}(x_{ki})$ where $v_{ki}(\cdot): \mathbb{R}_+ \to \mathbb{R}$.

The multicast network consists of links through which agents' data is transmitted from source to destination. The route $\mathcal{L}_{ki}$ of agent ki is the set of links that agent ki uses for his communication, and $\mathcal{L} = \cup_{ki\in\mathcal{N}}\mathcal{L}_{ki}$ is the set of all available links. We denote by $\mathcal{N}^l$ the set of agents utilizing a link $l \in \mathcal{L}$ and it is defined as $\mathcal{N}^l = \{ki \in \mathcal{N} \mid l \in \mathcal{L}_{ki}\}$. Also we define $\mathcal{G}_k^l \subseteq \mathcal{G}_k$ to be the set of agents from group k who use link l and $\mathcal{K}^l$ to be the set of groups that have at least one agent that uses link l i.e. $\mathcal{K}^l = \{k \in \mathcal{K} \mid \mathcal{G}_k^l \neq \emptyset\}$.

Finally, the cardinality of all the sets defined above are denoted as follows $N = |\mathcal{N}|$, $K = |\mathcal{K}|$, $G_k = |\mathcal{G}_k|$, $L = |\mathcal{L}|$, $L_{ki} = |\mathcal{L}_{ki}|$, $G_k^l = |\mathcal{G}_k^l|$ and $N^l = |\mathcal{N}^l|$.

In addition to group-wise ordering of agents, we also introduce a combined group and link-wise ordering of agents[2]. For group k and link l, any agent ki, can be identified by the index $g_k^l(i)$, where the mapping is defined as

$$ki \mapsto g_k^l(i) \quad \text{where} \quad 1 \leq g_k^l(i) \leq G_k^l, \quad \forall\, l \in \mathcal{L}_{ki}. \tag{2.1}$$

The above mapping is defined so that it preserves the original ordering i.e. $\forall\, i, j \in \mathcal{G}_k^l$ and $i > j$ we have $g_k^l(i) > g_k^l(j)$. Note that with this requirement the mapping is uniquely defined. For instance, if group k has four agents $\{k1, k2, k3, k4\}$ and only agents $k1, k3, k4$ use a particular link l then

$$k1 \mapsto g_k^l(1) = 1, \;\; k3 \mapsto g_k^l(3) = 2, \;\; k4 \mapsto g_k^l(4) = 3. \tag{2.2}$$

The network administrator is interested in maximizing the social welfare under the link capacity constraints. This centralized problem is

$$\max_x \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{G}_k} v_{ki}(x_{ki}) \tag{CP$_0$}$$

$$\text{s.t.} \quad x_{ki} \geq 0 \;\; \forall\, ki \in \mathcal{N} \tag{I$_1$}$$

[2]Given previous definitions, this reindexing is somewhat redundant; however, it will be used in the following to make the exposition clearer.

$$\text{and} \quad \sum_{k \in \mathcal{K}^l} \max_{j \in \mathcal{G}_k^l} \{\alpha_{kj}^l x_{kj}\} \leq c^l \ \ \forall\, l \in \mathcal{L}. \tag{I$_2$}$$

Constraints I_2 are the inequality constraints on allocation and represent the capacity constraint for every link $l \in \mathcal{L}$, in the network. Here α_{kj}^l are constants that represent the quality of service requirement of agent kj combined with the specific architecture on link l. So in order for agent kj to receive an actual data-rate of x_{kj} the bandwidth spent on link l is $\alpha_{kj}^l x_{kj}$. As an example, $\alpha_{kj}^l = \frac{1}{R_{kj}(1-\epsilon_{kj}^l)}$ for all links $l \in \mathcal{L}_{kj}$, where ϵ_{kj}^l represents the packet error probability for link l for a packet encoded with channel coding rate R_{kj}. Observe that due to the multicast/multirate architecture, only the maximum rate of each group k at each link l enters the capacity constraints.

2.1. Assumptions

The analysis will be done under the following assumptions.

(A1) For all agents, $v_{ki}(\cdot) \in \mathcal{V}_{ki}$, where the sets $\mathcal{V}_{ki}$ are arbitrary subsets of $\mathcal{V}_0$, the set of all strictly increasing, strictly concave, twice differentiable functions $\mathbb{R}_+ \to \mathbb{R}$ with continuous second derivative.

(A2) $v'_{ki}(0)$ is finite $\forall\, ki \in \mathcal{N}$. This also implies that $v'_{ki}(x)$ is finite and bounded $\forall\, ki$ and $\forall\, x \in \mathbb{R}_+$ since v_{ki}'s are concave.

(A3) Every link has at least two groups that use it i.e. $K^l \geq 2 \ \ \forall\, l \in \mathcal{L}$.

(A4) The optimal solution of the centralized problem is such that for every link there are at least 2 groups such that each has at least one non-zero component, i.e. if $S^l(x) := \{k \in \mathcal{K}^l \ \mid \ \exists\, i \in \mathcal{G}_k^l \ \text{ s.t. } x_{ki} > 0\}$ then the assumption says $|S^l(x^*)| \geq 2 \ \ \forall\, l \in \mathcal{L}$ (where x^* is the optimal solution of (CP$_0$)).

In addition, the coefficients are all strictly positive, i.e. $\alpha_{ki}^l > 0 \ \ \forall\, l \in \mathcal{L}_{ki}, \forall\, ki \in \mathcal{N}$. Also, for well-posedness of the problem we take $c^l > 0 \ \ \forall\, l \in \mathcal{L}$.

Assumption (A1) is made in order for the centralized problem to have a unique solution at the boundary of the feasible region and for this solution to be precisely characterized by the KKT conditions. (A2) is a mild technical assumption that is required in the proof of Lemma 3.7. Assumption (A3) is made in order to avoid situations where there is a link/constraint involving only one multicast group. Such a case requires special handling in the design of the mechanism (since in such a case there is no contention at that link), and destructs from the basic idea that we want to communicate. Finally (A4) is related to (A3) and is made in order to simplify the exposition of the proposed mechanism, without having to define corner cases that are of minor importance. One can state a number of mild sufficient conditions on the original optimization problem (CP_0) such that (A4) is satisfied. For instance if the optimal solution $x^* = \left(x^*_{ki}\right)_{ki\in\mathcal{N}}$ is such that $x^*_{ki} > 0$ for every agent ki then assumption (A4) is satisfied. Furthermore, this can be ensured by considering utility functions such that $v'_{ki}(0)$ is sufficiently large.

2.2. Necessary and Sufficient Optimality conditions

Under the stated assumptions, (CP_0) is a convex optimization problem. Following are the necessary KKT conditions for optimality, which under the stated assumptions are also sufficient (due to the strict concavity of $v_{ki}(\cdot)$). For this the centralized problem is first rewritten by restating the capacity constraints in a different form

$$\max_{x,m} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{G}_k} v_{ki}(x_{ki}) \tag{CP}$$

$$\text{s.t.} \quad x_{ki} \geq 0 \;\; \forall\, ki \in \mathcal{N} \tag{C$_1$}$$

$$\text{and} \quad \sum_{k\in\mathcal{K}^l} m^l_k \leq c^l \;\; \forall\, l \in \mathcal{L} \tag{C$_2$}$$

$$\text{and} \quad \alpha^l_{ki} x_{ki} \leq m^l_k \;\; \forall\, i \in \mathcal{G}^l_k,\; k \in \mathcal{K}^l,\; l \in \mathcal{L}. \tag{C$_3$}$$

The capacity constraints have been rewritten with the introduction of new variables. The virtual variables $m^l_k \in \mathbb{R}_+$ represent the weighted maximum requirement of group k on link l. It's easy to see that the solution of (CP) is the same as the solution of the original (CP_0) as far as optimal x is concerned. We introduce λ, μ, ν as the dual variables corresponding to constraints C_2, C_3 and C_1, respectively. The KKT conditions are stated without explicitly referring to ν_{ki}'s and just using the fact that $\nu^*_{ki} \geq 0$ and $\nu^*_{ki} x^*_{ki} = 0 \;\; \forall\, ki \in \mathcal{N}$. Note that with the assumptions above, the KKT conditions below will give rise to a unique x^* as the optimizer for (CP).

KKT conditions:

a) Primal Feasibility:

$$x_{ki}^* \geq 0 \;\; \forall \; ki \in \mathcal{N} \tag{2.3a}$$

$$\sum_{k \in \mathcal{K}^l} {m_k^l}^* \leq c^l \;\; \forall \; l \in \mathcal{L}; \tag{2.3b}$$

$$\alpha_{ki}^l x_{ki}^* \leq {m_k^l}^* \;\; \forall \; i \in \mathcal{G}_k^l, \; k \in \mathcal{K}^l, \; l \in \mathcal{L}. \tag{2.3c}$$

b) Dual Feasibility: $\lambda_l^* \geq 0 \;\; \forall \; l \in \mathcal{L}; \;\; \mu_{ki}^l \geq 0 \;\; \forall \; ki \in \mathcal{N}^l, l \in \mathcal{L}$.

c) Complimentary Slackness:

$$\lambda_l^* \left(\sum_{k \in \mathcal{K}^l} {m_k^l}^* - c^l \right) = 0 \;\; \forall \; l \in \mathcal{L}, \tag{2.4a}$$

$${\mu_{ki}^l}^* \left(\alpha_{ki}^l x_{ki}^* - {m_k^l}^* \right) = 0 \;\; \forall \; i \in \mathcal{G}_k^l, \; k \in \mathcal{K}^l, \; l \in \mathcal{L}. \tag{2.4b}$$

d) Stationarity:

$$v_{ki}'(x_{ki}^*) = \sum_{l \in \mathcal{L}_{ki}} {\mu_{ki}^l}^* \alpha_{ki}^l \; \forall \; ki \in \mathcal{N} \quad \text{if} \quad x_{ki}^* > 0 \tag{2.5a}$$

$$v_{ki}'(x_{ki}^*) \leq \sum_{l \in \mathcal{L}_{ki}} {\mu_{ki}^l}^* \alpha_{ki}^l \; \forall \; ki \in \mathcal{N} \quad \text{if} \quad x_{ki}^* = 0 \tag{2.5b}$$

and

$$\lambda_l^* = \sum_{i \in \mathcal{G}_k^l} {\mu_{ki}^l}^* \quad \forall \;\; k \in \mathcal{K}^l, \;\; l \in \mathcal{L}. \tag{2.6}$$

Looking at (2.4b), ${\mu_{ki}^l}^*$ will be non-zero only if $\alpha_{ki}^l x_{ki}^* = {m_k^l}^*$, so each ${\mu_{ki}^l}^*$ can be interpreted as the "price" paid only by those agents who receive maximum weighted allocation in group k at a given link l. Consequently, from (2.6), λ_l^* is the sum of ${\mu_{ki}^l}^*$ over those agents in group k that get maximum allocation within the group and it is the same for all groups. λ_l^* can then be interpreted as the common total price per unit of rate that each group is subject to at link l.

3. A mechanism with weak budget balance

In this section, we consider the case of weak budget balance, i.e., the taxes t_{ki} of all agents at equilibrium satisfy $\sum_{ki\in\mathcal{N}} t_{ki} \geq 0$. Note that the taxes used in the model refer to actual money paid by agents and not virtual signals, as may be the case in other works on distributed optimization.

3.1. Mechanism

The designer designs and announces the message space $\mathcal{S}_{ki}$ for each agent $ki \in \mathcal{N}$. The agents pick their message simultaneously and broadcast it. Based on the message profile received $s = (s_{ki})_{ki\in\mathcal{N}}$, the designer allocates rate $x = (x_{ki})_{ki\in\mathcal{N}}$ as $x_{ki} = h^x_{ki}(s)$. Similarly, the designer levies a tax (or subsidy) on each agent $t = (t_{ki})_{ki\in\mathcal{N}}$ as $t_{ki} = h^t_{ki}(s)$. The proposed mechanism is described below by defining the sets $\mathcal{S}_{ki}$ and functions $(h^x_{ki}(\cdot), h^t_{ki}(\cdot))$, for each $ki \in \mathcal{N}$.

This gives rise to a one-shot game

$$\mathfrak{G} = \Big(\mathcal{N}, \big(\mathcal{S}_{ki}\big)_{ki\in\mathcal{N}}, \big(\widehat{u}_{ki}(\cdot)\big)_{ki\in\mathcal{N}}\Big), \tag{2.7}$$

played by all the agents in $\mathcal{N}$, where action sets are $(\mathcal{S}_{ki})_{ki\in\mathcal{N}}$ and utilities are given by

$$\widehat{u}_{ki}(s) = v_{ki}(x_{ki}) - t_{ki} = v_{ki}(h^x_{ki}(s)) - h^t_{ki}(s), \quad \forall\, ki \in \mathcal{N}. \tag{2.8}$$

The resource allocation problem (CP) would be *fully implemented in NE*, if the outcomes (all possible NE) of this game produce allocation x^* and all agents in $\mathcal{N}$ are better-off participating in the mechanism than opting out (getting 0 allocation and taxes).

Message Space. The designer asks each agent to report a message $s_{ki} = (y_{ki}, Q_{ki})$ where $Q_{ki} = ({}^1Q_{ki}, {}^2Q_{ki})$ with ${}^1Q_{ki} = \left({}^1q^l_{ki}\right)_{l\in\mathcal{L}_{ki}}$ and ${}^2Q_{ki} = \left({}^2q^l_{ki}\right)_{l\in\mathcal{L}_{ki}}$; also for convenience denote $q^l_{ki} = \left({}^1q^l_{ki}, {}^2q^l_{ki}\right)$. The message s_{ki} includes a proxy for agent ki's demand for the rate, $y_{ki} \in \mathbb{R}_+$, and two "prices" ${}^1q^l_{ki} \in \mathbb{R}_+$, ${}^2q^l_{ki} \in \mathbb{R}_+$ for each link $l \in \mathcal{L}_{ki}$ that agent ki is involved in. For each agent ki and each link l, the first price ${}^1q^l_{ki}$ relates to the constraint $\alpha^l_{ki}x_{ki} \leq m^l_k$, while the second price ${}^2q^l_{ki}$ relates to the constraint $\alpha^l_{kj}x_{kj} \leq m^l_k$, where kj is the agent from group k on link l which can be identified by the index $g^l_k(i) + 1$ (see the group and link-wise indexing notation defined in (2.1)). As it turns out, our methodology makes quoting two prices necessary, as will be shown later. An intuitive explanation for the need of two prices is provided in the Discussion paragraphs after the tax definition in (2.13) and (2.14). From the above definitions, the message space for agent ki is $\mathcal{S}_{ki} = \mathbb{R}_+ \times \mathbb{R}_+^{2L_{ki}}$. For received messages $s = (s_{ki})_{ki \in \mathcal{N}} = (y, Q) = ((y_{ki})_{ki\in\mathcal{N}}, (Q_{ki})_{ki\in\mathcal{N}})$ the contract $h_{ki}(s) = (h_{ki}^x(s), h_{ki}^t(s))$ has an allocation component $h_{ki}^x(\cdot)$ and a tax component[3] $h_{ki}^t(\cdot)$ and is defined for each $ki \in \mathcal{N}$ as follows.

Allocation function. If the received demand vector is $y = (y_{11}, \ldots, y_{KG_K}) = 0$ then the allocation is $x = (x_{11}, \ldots, x_{KG_K}) = 0$ (also set $m = 0$). Otherwise it is evaluated by radially projecting the demand vector y on the boundary of the feasibility region as depicted in Fig. 9 for a two-group single-link case. More formally, the allocation function $h_{ki}^x(s)$ creates

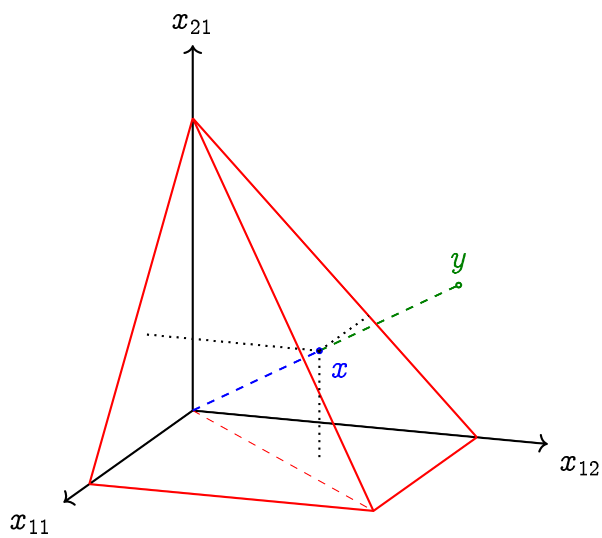

Figure 9.: Radial projection allocation for a single-link multicast/multirate network with 2 multicast groups $\mathcal{K} = \{1, 2\}$ with a total of 3 agents i.e. $\mathcal{N} = \{11, 12, 21\}$. Capacity constraint is $\max\{x_{11}, x_{12}\} + x_{21} \leq 1$. For demand y allocation x is just the projection of y onto the feasible set boundary (first sub-case of (2.9c) applies here).

allocation by first creating proxies $n_k^l \in \mathbb{R}_+$ for the weighted maximum at each link for each

[3] All utilities in this work are assumed to be quasi-linear.

group as follows

$$n_k^l \triangleq \max_{i \in \mathcal{G}_k^l} \{\alpha_{ki}^l y_{ki}\}, \quad \forall\, k \in \mathcal{K}^l,\ l \in \mathcal{L}. \tag{2.9a}$$

Then intermediate variables m_k^l's are created by dilating/shrinking n_k^l's on to one of the hyperplanes defined by the second set of constraints in C_2, specifically, that hyperplane for which the corresponding m_k^l's are the closest to origin (this could also be at the intersection of multiple hyperplanes). This is done through the introduction of a scaling factor r as follows

$$r = \min_{l \in \mathcal{L}} r^l \tag{2.9b}$$

$$r^l = \begin{cases} \dfrac{c^l}{\sum_{k \in \mathcal{K}^l} n_k^l}, & \text{if } |S^l(y)| \geq 2 \\ \dfrac{c^l}{\sum_{k \in \mathcal{K}^l} n_k^l} - f^l(n_k^l), & \text{if } S^l(y) = \{k\} \\ +\infty & \text{if } |S^l(y)| = 0 \end{cases} \tag{2.9c}$$

$$f^l(n_k^l) = \frac{c^l}{n_k^l(n_k^l + 1)}. \tag{2.9d}$$

Finally allocation x_{ki} is calculated by dilating/shrinking y_{ki} by the same factor.

$$h_{ki}^x(s) = x_{ki} = r\, y_{ki} \;\; \forall\, ki \in \mathcal{N}, \tag{2.10a}$$

$$m_k^l \triangleq r\, n_k^l, \;\; \forall\, k \in \mathcal{K}^l,\ l \in \mathcal{L}. \tag{2.10b}$$

In other words, the contract dilates/shrinks n_k^l to the boundary of the feasible region defined by the capacity constraints and then allocations within a group are made proportionally. Since all the α_{kj}^l's are positive, this means that all constraints in C_2 are satisfied for the allocation automatically, as shown later. In the above description of the allocation function, the separate definition for r^l when $|S^l(y)| < 2$ is to ensure (as it will be shown later) that there are no equilibria with allocation x^*, where $|S^l(x^*)| < 2$.

Tax function. For the taxes, we first define total prices, $w_k^l, \bar{w}_{-k}^l$ for any link l and group

$k \in \mathcal{K}^l$ as

$$w_k^l \triangleq \sum_{i \in \mathcal{G}_k^l} {}^1q_{ki}^l, \qquad \bar{w}_{-k}^l \triangleq \frac{1}{K^l - 1} \sum_{k' \in \mathcal{K}^l \setminus \{k\}} w_{k'}^l \tag{2.11}$$

where $\bar{w}_{-k}^l$ is well-defined due to assumption (A3). The tax is defined as the sum of taxes for each constraint in an agent's route

$$h_{ki}^t(s) = t_{ki} = \sum_{l \in \mathcal{L}_{ki}} t_{ki}^l \;\; \forall \, ki \in \mathcal{N}, \tag{2.12}$$

and each component t_{ki}^l is defined as follows. If $G_k^l \geq 2$ then consider agents kj and ke from group k and on link l who can be identified by the index $g_k^l(i) - 1$ and $g_k^l(i) + 1 \pmod{G_k^l}$, respectively. Then,

$$\begin{aligned} t_{ki}^l = x_{ki}\alpha_{ki}^l \, {}^2q_{kj}^l + ({}^2q_{ki}^l - {}^1q_{ke}^l)^2 + (w_k^l - \bar{w}_{-k}^l)^2 + \eta \, {}^2q_{kj}^l({}^1q_{ki}^l - {}^2q_{kj}^l)(m_k^l - \alpha_{ki}^l x_{ki}) \\ + \xi \, \bar{w}_{-k}^l(w_k^l - \bar{w}_{-k}^l)(c^l - \sum_{k' \in \mathcal{K}^l} m_{k'}^l), \end{aligned} \tag{2.13}$$

where η, ξ are sufficiently small positive constants (whose selection is outlined in the proof of Lemma 3.7). If $G_k^l = 1$ then

$$\begin{aligned} t_{ki}^l = x_{ki}\alpha_{ki}^l \bar{w}_{-k}^l + (w_k^l - \bar{w}_{-k}^l)^2 + \eta \, \bar{w}_{-k}^l({}^1q_{ki}^l - \bar{w}_{-k}^l)(m_k^l - \alpha_{ki}^l x_{ki}) \\ + \xi \, \bar{w}_{-k}^l(w_k^l - \bar{w}_{-k}^l)(c^l - \sum_{k' \in \mathcal{K}^l} m_{k'}^l). \end{aligned} \tag{2.14}$$

Discussion. The need for two prices is mainly a technical necessity (an unavoidable corner case whenever an inequality constraint involves only one user) and can be explained as follows. When there is a constraint involving at least two agents, then it suffices to ask each agent to quote a single price for it. The price quoted by agent kj would be used is used in two places in (2.13): it would determine the price paid by agent ki in his first tax term (this way agent ki does not control both his price and (indirectly) his allocation) and it will also be used in the fourth term of agent's kj taxes to make sure the complementary slackness condition is satisfied. Similarly for the price quoted by user ki. However, for the problem at hand, constraints (C_3) ($\alpha_{ki}^l x_{ki} \leq m_k^l$) involve only one agent, ki, per constraint. This changes things because this quoted price cannot be used both in his first and fourth tax terms! This necessitates that users quote two prices as follows. We require agent ki, to quote a price ${}^1q^l_{ki}$ that will be used in his fourth tax term to make sure the complementary slackness condition is satisfied. Clearly we cannot use the same quoted price for his first tax term (this would violate the condition that the same agent cannot control both his price and quantity). So we ask agent kj to quote a price ${}^2q^l_{kj}$ that is paid by agent ki (first term in t^l_{ki}). These two prices, ${}^2q^l_{kj}, {}^1q^l_{ki}$ (the second price quoted by kj and the first price quoted by ki) are thus used as a proxy for the same quantity: the optimal Lagrange multiplier ${\mu^l_{ki}}^*$ corresponding to the constraint $\alpha^l_{ki}x_{ki} \leq m^l_k$.

It should now be clear what the role of the other tax terms is. The second and third terms $({}^2q^l_{ki} - {}^1q^l_{ke})^2, (w^l_k - \bar{w}^l_{-k})^2$ are introduced to incentivize agents to quote same prices. For instance, the term $({}^2q^l_{ki} - {}^1q^l_{ke})^2$ drives agent ki to quote ${}^2q^l_{ki}$ such that it matches ${}^1q^l_{ke}$ - both are a proxy for the optimal Lagrange multiplier ${\mu^l_{ke}}^*$, as explained above (for the kj, ki pair of agents). Similarly, the purpose of the term $(w^l_k - \bar{w}^l_{-k})^2$ is to tax every agent from group k additionally (in a smooth manner) if the total group price $w^l_k = \sum_{i \in \mathcal{G}^l_k} {}^1q^l_{ki}$ doesn't match with the average of group price $\bar{w}^l_{-k}$ for other groups. This will drive agents to quote ${}^1q^l_{ki}$'s such that the two match and if the group price for all groups match with the average of others, then indeed all group prices are the same (as required by (2.6) in KKT). To satisfy the complimentary slackness conditions at NE we introduce the fourth and fifth terms which charge agent ki higher taxes if they quote non-zero prices for inactive constraints. This necessarily requires using an agent's own quoted price (i.e. price ${}^1q^l_{ki}$ for agent ki).

Finally, below we show that all tax terms except the first tax term are zero at NE and ${}^1q^l_{ki} = {}^2q^l_{kj} = {\mu^l_{ki}}^*$, where ${\mu^l_{ki}}^*$ is the optimal Lagrange multiplier (see (2.4b), (2.5a)). Consequently, the total tax paid by agent ki at NE is $t_{ki} = x^*_{ki} \sum_{l \in \mathcal{L}_{ki}} \alpha^l_{ki} {\mu^l_{ki}}^* = x^*_{ki} v'_{ki}(x^*_{ki})$ (see (2.5a)). The price paid by agent ki at NE is the true marginal valuation $v'_{ki}(x^*_{ki})$ of agent ki and thus the mechanism ensures fairness of taxes paid by agents in the sense that it they are exactly what expected in a free market.

Readers may refer to [Hu79; Ch02] for mechanisms with small message spaces which use similar technique for taxation in Nash implementing mechanisms for Lindahl correspondence (public goods).

There are two levels of interactions that this mechanism is dealing with, one among groups for allocation of maximum rate on each link and a second within each group. Agents contest for allocation that makes full use of the fact that from within a group only maximum at each link will give rise to a positive price on that link. At any link l and group k, total price w_k^l is the sum of prices quoted by all the agents in the group at link l. The quantity $\bar{w}_{-k}^l$ is calculated by averaging the total prices for link l over all other groups than k. Quoting of prices and demand is used as a way of eliciting $v'_{ki}(x_{ki})$ by comparing it appropriately with prices. In this vein, we do not wish to influence $\bar{w}_{-k}^l$ with prices quoted by groups whose agents aren't using the link at all, since the price then essentially doesn't contain any information.

3.2. Results

The basic result of this section is summarized in the following theorem.

Theorem 3.1 (Full Implementation). *For game $\mathfrak{G}$, there is a unique allocation, x, corresponding to all NE. Moreover, $x = x^*$, the maximizer of* (CP). *In addition, individual rationality is satisfied for all agents and so is weak budget balance.*

The theorem is proved by a sequence of results, in which all candidate NE of $\mathfrak{G}$ are characterized by necessary conditions until only one family of NE candidates is left. Subsequently, the existence of NE in pure strategies for $\mathfrak{G}$ is shown, and that all NE result in allocation $x = x^*$. Finally, individual rationality and WBB will be checked.

Lemma 3.2 (Primal Feasibility). *For any action profile $s = (y, Q)$ of game $\mathfrak{G}$, constraints C_1, C_2 and C_3 of* (CP) *are satisfied at the corresponding allocation.*

Proof. Please see Appendix 6.1 at the end of this chapter. □

Feasibility of allocation for action profiles is a direct consequence of the radial projection allocation function. Next, it is shown that all groups, using a link, quote the same total price w_k^l for that link at any equilibrium. This is brought about by the 3rd tax term $\sum_l (w_k^l - \bar{w}_{-k}^l)^2$. This is a way of threatening agents with higher taxes just for quoting a different price than average, at each link.

Lemma 3.3. *At any NE $s = (y, Q)$ of $\mathfrak{G}$, for any link $l \in \mathcal{L}$ we have*

$$w_k^l = w^l \quad \forall\, k \in \mathcal{K}^l. \tag{2.15}$$

Also, for any group k and link l such that $G_k^l \geq 2$ if we take any agents $i, e \in \mathcal{G}_k^l$ where agent ke can be identified by the index $g_k^l(i)+1$ then at equilibrium we will have ${}^2q_{ki}^l = {}^1q_{ke}^l$ (which will be denoted as p_{ke}^l).

Proof. Please see Appendix 6.2 at the end of this chapter. □

With Lemma 3.3, equilibria can now be referred to in terms of the common total price vector P rather than the two different total price vectors ${}^1Q$, and ${}^2Q$. In particular, any NE $s = (y, Q)$ with $Q = ({}^1Q, {}^2Q)$ can be characterized as $s = (y, P)$ with $P = (P_{ki})_{ki \in \mathcal{N}}$ and $P_{ki} = \left(p_{ki}^l\right)_{l \in \mathcal{L}_{ki}}$.

Later it will become clear how p_{ki}^l and w^l take the place of dual variables μ_{ki}^l and λ_l, respectively, when we compare equilibrium conditions with KKT conditions, hence we identify the following condition as dual feasibility.

Lemma 3.4 (Dual Feasibility). *$p_{ki}^l \geq 0$, $w^l \geq 0 \quad \forall\, i \in \mathcal{G}_k^l$, $\forall\, k \in \mathcal{K}^l$ and $\forall\, l \in \mathcal{L}$.*

Proof. This is by design, since agents are only allowed to quote non-negative prices and that w^l is the sum of such prices. □

Following is the property that solidifies the notion of prices as dual variables, since the claim here is that inactive constraints do not contribute to payment at equilibrium. This notion is very similar to the centralized problem, where if we know certain constraints to be inactive at the optimum then the same problem without these constraints would be equivalent to the original. The 4$^{\text{th}}$ and 5$^{\text{th}}$ terms in the tax function facilitate this by charging extra taxes for inactive constraints where the agent is quoting higher prices than the average of remaining ones, thereby driving prices down.

Lemma 3.5 (Complimentary Slackness). *At any NE $s = (y, P)$ of game $\mathfrak{G}$ with corresponding allocation x, for any agent $i \in \mathcal{G}_k^l$, group $k \in \mathcal{K}^l$ and link $l \in \mathcal{L}$ we have*

$$w^l \left(\sum_{k \in \mathcal{K}^l} m_k^l - c^l \right) = 0, \qquad p_{ki}^l \left(\alpha_{ki}^l x_{ki} - m_k^l \right) = 0. \tag{2.16}$$

Proof. Please see Appendix 6.3 at the end of this chapter. □

Lemma 3.6 (Stationarity). *At any NE $s = (y, P)$ of game $\mathfrak{G}$, and corresponding allocation x, we have*

$$v'_{ki}(x_{ki}) = \sum_{l \in \mathcal{L}_{ki}} p^l_{ki} \alpha^l_{ki} \;\; \forall\, ki \in \mathcal{N} \quad \textit{if} \quad x_{ki} > 0 \tag{2.17a}$$

$$v'_{ki}(x_{ki}) \leq \sum_{l \in \mathcal{L}_{ki}} p^l_{ki} \alpha^l_{ki} \;\; \forall\, ki \in \mathcal{N} \quad \textit{if} \quad x_{ki} = 0 \tag{2.17b}$$

and

$$w^l = \sum_{i \in \mathcal{G}^l_k} p^l_{ki} \;\; \forall\, k \in \mathcal{K}^l,\, l \in \mathcal{L}. \tag{2.18}$$

Proof. Please see Appendix 6.4 at the end of this chapter. □

Collecting the results of the above lemmas, we can conclude that every NE profile satisfies the KKT conditions of the (CP). This means there are now necessary conditions on the NE up to the point of having unique allocation (since KKT for (CP) is satisfied by a unique optimal x). In the next Lemma we verify the existence of the equilibria that we have claimed.

Lemma 3.7 (Existence). *For the game $\mathfrak{G}$, there exists an equilibrium.*

Proof. Please see Appendix 6.5 at the end of this chapter. □

Several comments are in order regarding the selection of the radial projection allocation function and in particular (2.9c). If one uses "pure" radial projection i.e. same expression for r^l for $|S^l(y)| \geq 2$ and ≤ 1, then irrespective of the optimal solution of (CP), for game $\mathfrak{G}$ the "stationarity" property will not be satisfied at equilibria with $|S^l(y)| \leq 1$. Thus the mechanism will result in additional extraneous equilibria. This is the reason why we tweak the expression for r^l when $|S^l(y)| \leq 1$, so as to eliminate these extraneous equilibria - irrespective of the solution of (CP). With this tweak in the expression for r^l, all KKT conditions become necessary for all equilibria regardless of the value of $|S^l(y)|$. This however creates a problem in the proof of existence of equilibria. In particular, if x^* was such that it had links where $|S^l(x^*)| = 1$ then with modified radial projection allocation this would require y at NE such that $|S^l(y)| = 1$. In this case the r^l used would be lower than what the radial projection requires (see second sub-case in (2.9c)) and we actually would have the problem of possibly not having any y that creates x^* as allocation. Hence we have used (A4) to eliminate this case.

Lemma 3.8 (Individual Rationality and WBB). *At any NE $s = (y, P)$ of $\mathfrak{G}$, with corresponding allocation x and taxes t, we have*

$$u_{ki}(x,t) \geq u_{ki}(0,0) \ \ \forall\, ki \in \mathcal{N} \tag{2.19}$$

$$\textit{and} \qquad \sum_{ki \in \mathcal{N}} t_{ki} \geq 0. \tag{2.20}$$

Proof. Please see Appendix 6.6 at the end of this chapter. □

With all the Lemmas characterizing NE in the same way as KKT conditions (and individual rationality), we can compare them to prove Theorem 3.1.

Proof of Theorem 3.1. We know that the four KKT conditions produce a unique solution x^* (and corresponding λ^*). For the game $\mathfrak{G}$, if NE exist, then from Lemmas 3.2–3.6 we can see that at any NE, allocation x and prices p satisfy the same conditions as the four KKT conditions and hence they give a unique $x = x^*$, as long as (A5) is satisfied. We conclude that the allocation is x^* across all NE. Existence of the claimed NE is established in Lemma 3.7. This combined with individual rationality Lemma 3.8, proves Theorem 3.1. □

4. A mechanism with Strong Budget Balance

We now present a modification of the previous mechanism that in addition to previous results also ensures *strong budget balance* (SBB) at NE i.e. $\sum_{ki \in \mathcal{N}} t_{ki} = 0$.

For creating a mechanism in this formulation, the main difference with the previous section, is that the designer has to find a way of redistributing the total tax paid by all the agents. In the last section it was shown that the total payment made at the equilibrium is

$$B = \sum_{ki \in \mathcal{N}} \left(x_{ki} \sum_{l \in \mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki} \right) = r \sum_{ki \in \mathcal{N}} \left(y_{ki} \sum_{l \in \mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki} \right), \tag{2.21}$$

since all other tax terms are zero at equilibrium. The method here, following in the spirit of [Hu79], is to redistribute taxes by modifying the tax function for each agent only using messages from other agents. This has the advantage of keeping equilibrium calculations in line with the previous mechanism, since deviations by an agent wouldn't affect his utility through this additional term. In view of this, an alternate expression for B is

$$B = r \sum_{ki\in\mathcal{N}} \left(\sum_{l\in\mathcal{L}_{ki}} \frac{1}{N^l - 1} \sum_{k'j\in\mathcal{N}^l\setminus\{ki\}} \alpha^l_{k'j} p^l_{k'j} y_{k'j} \right), \tag{2.22}$$

where each term of the outermost summation depends only on demands of agents other than the ki^{th} one. This means that each term in the parenthesis (scaled by the factor r) can now be used as the desired additional tax for user ki. Observe however that in our mechanism, each agent's demand affects the factor r as well. So, if all agents can agree on value of r then that signal can be used to create the term that facilitates budget balance.

In lieu of this, the mechanism here works by asking for an additional signal ρ_{ki} from every agent and imposing an additional tax of $(\rho_{ki} - r)^2$, thereby essentially ensuring that all agents agree on the value of r (via ρ_{ki}'s) at equilibrium. Finally, $\bar{\rho}_{-ki}$ (see (2.26)) is used as a proxy for r in (2.22) - somewhat similar to how $\bar{w}^l_{-k}$'s were used in the third term of (2.13).

4.1. Mechanism

The actions set for agent ki is now $\mathcal{S}_{ki} = \mathbb{R}_+ \times \mathbb{R}_+^{2L_{ki}} \times \mathbb{R}_+$ where an action is of the form $s_{ki} = (y_{ki}, Q_{ki}, \rho_{ki})$. The allocation and tax function announced by the designer are exactly as in the previous case, with the only exception that the tax for any agent ki is now defined as

$$h^t_{ki}(s) = t_{ki} = \zeta(\rho_{ki} - r)^2 + \sum_{l\in\mathcal{L}_{ki}} t^l_{ki}. \tag{2.23}$$

If $G^l_k \geq 2$, then again using agents kj and ke as described after (2.12), we have

$$\begin{aligned} t^l_{ki} = {} & x_{ki}\alpha^l_{ki}\,{}^2q^l_{kj} + ({}^2q^l_{ki} - {}^1q^l_{ke})^2 + (w^l_k - \bar{w}^l_{-k})^2 + \eta\,{}^2q^l_{kj}({}^1q^l_{ki} - {}^2q^l_{kj})(m^l_k - \alpha^l_{ki}x_{ki}) \\ & + \xi\,\bar{w}^l_{-k}(w^l_k - \bar{w}^l_{-k})(c^l - \sum_{k'\in\mathcal{K}^l} m^l_{k'}) - \frac{\bar{\rho}_{-ki}}{N^l - 1} \sum_{k'j\in\mathcal{N}^l\setminus\{ki\}} \alpha^l_{k'j}\,{}^2q^l_{k'j'} y_{k'j} \end{aligned} \tag{2.24}$$

where agent $k'j'$ is an agent from group k' on link l who can be identified by the index $g^l_{k'}(j) - 1$, if $G^l_{k'} \geq 2$ and ζ, η, ξ are small enough positive constants. However if $G^l_{k'} = 1$ we would use $\bar{w}^l_{-k'}$ instead of $q^l_{k'j'}$.

Similarly for $G_k^l = 1$, we have

$$\begin{aligned} t_{ki}^l &= x_{ki}\alpha_{ki}^l \bar{w}_{-k}^l + (w_k^l - \bar{w}_{-k}^l)^2 + \eta\, \bar{w}_{-k}^l ({}^1q_{ki}^l - \bar{w}_{-k}^l)(m_k^l - \alpha_{ki}^l x_{ki}) \\ &\quad + \xi\, \bar{w}_{-k}^l (w_k^l - \bar{w}_{-k}^l)(c^l - \sum_{k' \in \mathcal{K}^l} m_{k'}^l) - \frac{\bar{\rho}_{-ki}}{N^l - 1} \sum_{k'j \in \mathcal{N}^l \setminus \{ki\}} {\alpha_{k'j}^l}^2 q_{k'j'}^l y_{k'j}. \end{aligned} \tag{2.25}$$

In addition to previous definitions,

$$\bar{\rho}_{-ki} := \frac{1}{N-1} \sum_{k'j \in \mathcal{N} \setminus \{ki\}} \rho_{k'j}. \tag{2.26}$$

Denote the corresponding game by $\mathfrak{G}_0$. The implications of these modifications are discussed in the results section below.

4.2. Results

This new mechanism will also fully implement (CP). The only term in $\widehat{u}_{ki}$ that is affected by ρ_{ki} is $-(\rho_{ki} - r)^2$, so all the Lemmas from Section 3 are valid with minor modifications and the main result will follow using the same line of argumentation as for Theorem 3.1. Note here that, terms in $\widehat{u}_{ki}$ affected by ${}^1q_{ki}^l$, ${}^2q_{ki}^l$'s are the same as before but for y_{ki} there is a new term $-(\rho_{ki} - r)^2$ which is affected by it.

Theorem 4.1 (Full Implementation with SBB). *For game $\mathfrak{G}_0$, there is a unique allocation, x, corresponding to all NE. Moreover, $x = x^*$, the maximizer of* (CP), *where individual rationality is satisfied for all agents in $\mathcal{N}$. Also, SBB is satisfied at all NE.*

For economy of exposition, in the following, we only provide detailed proofs for the new properties of this mechanism, while we outline the proofs for the properties that are similar to those in the WBB mechanism.

- Primal Feasibility - Since allocation function is the same as before, this result holds here as well.

- Equal Prices at equilibrium - This was proved by taking price deviations only and keeping other parameters of the signal constant, so the same argument works here as well (noting that no new price related terms have been added in the new mechanism).

Before moving on to other results, we show that a common ρ_{ki}'s emerges at equilibrium.

Lemma 4.2. *At any NE $s = (y, P, \rho)$ of game $\mathfrak{G}_0$, we have $\rho_{ki} = r \ \forall \ ki \in \mathcal{N}$.*

Proof. Please see Appendix 6.7 at the end of this chapter. □

Note however that although ρ_{ki} are same for all ki at any equilibrium, that common value, r, will be different across equilibria. This is obvious since magnitude of vector y changes across equilibria.

Now we continue with other properties from Section 3.

- Dual Feasibility - This is obvious here as well.
- Complimentary Slackness - This was proved by taking only price deviations and hence the same argument works here as well.
- Stationarity - Compared to the WBB case, the additional term in the derivative is

$$\frac{\partial \widehat{u}_{ki}}{\partial y'_{ki}}\bigg|_{new} = \underbrace{\frac{\partial \widehat{u}_{ki}}{\partial y'_{ki}}\bigg|_{old}}_{T_1} \underbrace{- 2\zeta(\rho'_{ki} - r')\left(-\frac{\partial r'}{\partial y'_{ki}}\right)}_{T_2} \tag{2.27}$$

  The claim is, as before, that if T_1 is positive, agent ki can increase y'_{ki} from y_{ki} to be better-off. However agent ki has to ensure that he deviates with ρ'_{ki} simultaneously to make it equal to r', so that the contribution of the T_2 term to the derivative above continues to be zero. The only thing left to notice here is that the change in ρ'_{ki} is such that not only the term T_2 is zero but also that the contribution of term $-\zeta(\rho'_{ki} - r')^2$ to the utility is zero before and after deviation - so this deviation doesn't change other partial derivatives. Similar argument also works when T_1 is negative and we get the stationarity property here as well.

With the above properties, unique allocation x^* at all equilibria is guaranteed, and, as before, the prices are equal to λ^*.

Existence of equilibria is verified in Appendix 6.8 at the end of this chapter. The arguments used are similar to the ones in the proof of Lemma 3.7.

- Individual Rationality - This is obvious in here because money from the previous case is only being redistributed here, so if the mechanism there was individually rational it will be here too.

Lemma 4.3 (Strong Budget Balance). *At any NE* $s = (y, P, \rho)$ *of game* $\mathfrak{G}_0$*, with corresponding taxes* $\{t_{ki}\}_{ki\in\mathcal{N}}$*, we have* $\sum_{ki\in\mathcal{N}} t_{ki} = 0$.

Proof. Please see Appendix 6.9 at the end of this chapter. □

Proof of Theorem 4.1. By the preceding properties, at all equilibria the allocation is x^* and prices λ^*. Then SBB and individual rationality provide the desired full implementation. □

5. Summary and Comments

In this chapter, we present a mechanism that fully implements the sum of utilities maximizing allocation for agents who share data on a multicast/multirate network. The scope of application of this model goes beyond just data provision on networks. Another example (as mentioned in the Introduction) is the provision of security products for server farms. The design also encompasses two important auxiliary properties: off-equilibrium feasibility of allocation (using the radial projection function) and SBB at NE. In addition the overall size of the message space is linear in the number of agents, N.

The work in [KaTe13b] is similar to the work in this chapter except for two things. Firstly, the mechanism doesn't posses the off-equilibrium feasibility property and secondly there are instances of the centralized problem where the claimed NE profiles do indeed have profitable unilateral deviations, thereby contradicting the full implementation claim.

In the spirit of some recent works [Ch02; HeMa12] for Walrasian and Lindahl allocations, the main continuation of this work would be to design the allocation/tax functions where in addition to full implementation, convergence of certain classes of learning algorithms is also taken as a design objective. This way we can convert our "learning ready" mechanism into a guaranteed learning mechanism. Designing so that the ensuing game is a potential game or a super-modular game is too stringent for this problem. But the design technique used in [HeMa12], namely contractive best response correspondence, is slightly relaxed. Basic design parameters like the structure of the message space (consisting of demands and prices) can be kept the same in this case.

Another future direction, from an engineering perspective, can be to impose additional constraints on the mechanism design so that the messages in the resulting mechanism can be exchanged in a distributed manner and can possibly eliminate the need for a centralized coordinator to collect all quoted message and impose allocation and taxes. Naturally, such a development would require the assumption that even without the presence of a centralized coordinator agents will pay the taxes that the mechanism contract requires and that allocation (as defined by the mechanism) can be effected.

For the mechanism defined in this chapter the communication structure is as follows. The prices ${}^1q^l_{ki}, {}^2q^l_{ki}$ can be communicated only locally within group k and link l. The quantity $\bar{w}^l_{-k}$ requires communication between groups and for this a leader can be designated from each group, who will communicate with the other leaders. Calculating the allocation poses other issues. The leader of group k on link l can also calculate $n^l_k = \max_{i \in \mathcal{G}^l_k}\{\alpha^l_{ki} y_{ki}\}$ and communicate it with other groups on link l to calculate r^l and then communication is needed across links to get $r = \min_{l \in \mathcal{L}} r^l$.

6. Appendix

6.1. Proof of Lemma 3.2

Proof. Constraint C_1 is clearly always satisfied. For $y = 0$ we have $x = 0$ and $m = 0$, so constraints C_2 and C_3 are also clearly satisfied. Next we show C_2 and C_3 are satisfied for any $y \neq 0$ as demand. Firstly $r < +\infty$, since there exists at least one link d with $|S^d(y)| \geq 1$ and thus $r^d < +\infty$. Now, for any link l, there are the following two cases. If $|S^l(y)| = 0$, then the allocation for all agents in $\mathcal{N}^l$ is zero (along with the corresponding m_k^l's), so C_2 and C_3 for those links is satisfied. If $|S^l(y)| \geq 1$ then

$$\sum_{k \in \mathcal{K}^l} m_k^l = r \sum_{k \in \mathcal{K}^l} n_k^l \leq r^l \sum_{k \in \mathcal{K}^l} n_k^l \leq \frac{c^l}{\sum_{k \in \mathcal{K}^l} n_k^l} \sum_{k \in \mathcal{K}^l} n_k^l = c^l \tag{2.28}$$

where the first inequality holds because r is the minimum of all r^l's. The second inequality is equality if $|S^l(y)| \geq 2$ and is strict only if $|S^l(y)| = 1$ (see second sub-case in (2.9c)). For C_2, take any agent ki and link $l \in \mathcal{L}_{ki}$

$$\alpha_{ki}^l x_{ki} = r \alpha_{ki}^l y_{ki} \leq r n_k^l = m_k^l \tag{2.29}$$

where the inequality holds because n_k^l is the maximum over $\alpha_{ki}^l y_{ki}$'s for all $i \in \mathcal{G}_k^l$. □

6.2. Proof of Lemma 3.3

Proof. First we show the second part of the lemma, so suppose there are agents $i, e \in \mathcal{G}_k^l$ as above, for whom ${}^2q_{ki}^l \neq {}^1q_{ke}^l$. If agent ki deviates with ${}^2q_{ki}^{l\,\prime} = {}^1q_{ke}^l$ then we can write the difference in agent ki's utility after and before deviation by just comparing tax for link l (since allocation and tax for other links don't change)

$$\Delta \widehat{u}_{ki} = -({}^2q_{ki}^{l\,\prime} - {}^1q_{ke}^l)^2 + ({}^2q_{ki}^l - {}^1q_{ke}^l)^2 = ({}^2q_{ki}^l - {}^1q_{ke}^l)^2 > 0 \tag{2.30}$$

which means that the deviation was profitable. This gives us the second part of the lemma. (In addition to defining ${}^2q_{ki}^l = {}^1q_{ke}^l = p_{ke}^l$ when $G_k^l \geq 2$, we will also denote ${}^1q_{ki}^l = w_k^l = p_{ki}^l$ when $\mathcal{G}_k^l = \{i\}$).

For the first part, suppose there is a link l for which $(w^l_k)_{k\in\mathcal{K}^l}$ are not all equal, at equilibrium. Clearly then there is a group $k \in \mathcal{K}^l$ for which $w^l_k > \bar{w}^l_{-k}$ (this can be seen from (2.11)). We will show that some agent $i \in \mathcal{G}^l_k$ can deviate by reducing price ${}^1q^l_{ki}$ and be strictly better off, thereby contradicting the equilibrium condition. First we take the case when the group k is such that $G^l_k \geq 2$ and then $G^l_k = 1$.

Since $w^l_k > \bar{w}^l_{-k}$ we must have $w^l_k > 0$ and since $w^l_k = \sum_{i\in\mathcal{G}^l_k} {}^1q^l_{ki}$ there must be an agent $i \in \mathcal{G}^l_k$ for whom ${}^1q^l_{ki} > 0$. Take deviation by this agent ki as ${}^1q^l_{ki}{}' = {}^1q^l_{ki} - \epsilon > 0$, for which we can write the difference in utility, just as before, as

$$\begin{aligned}\Delta\hat{u}_{ki} &= -\epsilon^2 + 2\epsilon(w^l_k - \bar{w}^l_{-k}) + \eta\,\epsilon p^l_{ki}(m^l_k - \alpha^l_{ki}x_{ki}) + \xi\,\epsilon\bar{w}^l_{-k}(c^l - \sum_{k\in\mathcal{K}^l} m^l_k) \\ &= \epsilon\Big(-\epsilon + 2(w^l_k - \bar{w}^l_{-k}) + \eta\, p^l_{ki}(m^l_k - \alpha^l_{ki}x_{ki}) + \xi\,\bar{w}^l_{-k}(c^l - \sum_{k\in\mathcal{K}^l} m^l_k)\Big) = \epsilon(-\epsilon + a)\end{aligned} \tag{2.31}$$

where $a > 0$ because of Lemma 3.2 and the fact that $w^l_k > \bar{w}^l_{-k}$. So by taking ϵ such that $\min\{a, {}^1q^l_{ki}\} > \epsilon > 0$, the above deviation will be a profitable one for agent ki. This gives the result for $G^l_k \geq 2$.

For $G^l_k = 1$, say $\mathcal{G}^l_k = \{i\}$, we have that ${}^1q^l_{ki} = w^l_k > \bar{w}^l_{-k}$. This again means that ${}^1q^l_{ki} > 0$ and we take the deviation ${}^1q^l_{ki}{}' = {}^1q^l_{ki} - \epsilon > 0$ and get

$$\Delta\hat{u}_{ki} = \epsilon\Big(-\epsilon + 2(w^l_k - \bar{w}^l_{-k}) + \eta\,\bar{w}^l_{-k}(m^l_k - \alpha^l_{ki}x_{ki}) + \xi\,\bar{w}^l_{-k}(c^l - \sum_{k\in\mathcal{K}^l} m^l_k)\Big). \tag{2.32}$$

Following the same argument as above we will get our result here as well. □

6.3. Proof of Lemma 3.5

Proof. Suppose there is a link l for which $w^l > 0$ and $\sum_{k\in\mathcal{K}^l} m^l_k < c^l$. Take any group $k \in \mathcal{K}^l$ and an agent $i \in \mathcal{G}^l_k$ such that ${}^1q^l_{ki} = p^l_{ki} > 0$ (there is such an agent because $w^l = \sum_{i\in\mathcal{G}^l_k} {}^1q^l_{ki} > 0$). Take the deviation ${}^1q^l_{ki}{}' = {}^1q^l_{k,i} - \epsilon > 0$ and we get (using same arguments as in (2.31) and noting that $w^l_k = \bar{w}^l_{-k} = w^l$)

$$\Delta\widehat{u}_{ki} = \epsilon\Big(-\epsilon + \eta \underbrace{p^l_{ki}(m^l_k - \alpha^l_{ki}x_{ki})}_{\geq 0 \text{ by Lemma 3.2}} + \xi\, w^l(c^l - \sum_{k\in\mathcal{K}^l} m^l_k)\Big) = \epsilon(-\epsilon + a). \tag{2.33}$$

where $a > 0$ due to Lemma 3.2 and the assumption that $w^l(c^l - \sum_{k\in\mathcal{K}^l} m^l_k) > 0$. This gives us that $w^l(c^l - \sum_{k\in\mathcal{K}^l} m^l_k) = 0$ for all $l \in \mathcal{L}$ at equilibrium.

Now suppose there is an agent ki for whom ${}^1q^l_{ki} = p^l_{ki} > 0$ and $\alpha^l_{ki}x_{ki} < m^l_k$. Same as before, we take the deviation ${}^1q^l_{ki}{}' = {}^1q^l_{ki} - \epsilon > 0$,

$$\Delta\widehat{u}_{ki} = \epsilon\left(-\epsilon + \eta\, p^l_{ki}(m^l_k - \alpha^l_{ki}x_{ki})\right) = \epsilon(-\epsilon + a) \tag{2.34}$$

where $a > 0$ by assumption. This gives us that $p^l_{ki}(m^l_k - \alpha^l_{ki}x_{ki}) = 0$ for all $ki \in \mathcal{N}^l$ and $l \in \mathcal{L}$, at equilibrium. □

6.4. Proof of Lemma 3.6

Proof. (2.18) is true by construction since we defined $w^l_k = \sum_{i\in\mathcal{G}^l_k} p^l_{ki}$ and by Lemma 3.3 we have $p^l_{ki} = {}^1q^l_{ki}$ and $w^l_k = w^l$.

At any NE, agent ki's utility in the game $\widehat{u}_{ki}(s') = v_{ki}(h^x_{ki}(s')) - h^t_{ki}(s')$ as a function of his message $s'_{ki} = (y'_{ki}, Q'_{ki})$, with s_{-ki} fixed, should have a global maximum at $s_{ki} = (y_{ki}, P_{ki})$. This would mean that if this function was differentiable w.r.t. y'_{ki} at s, the partial derivatives w.r.t. y'_{ki} at s should be 0. However, since our allocation dilates/shrinks demand vector y' on to the feasible region, it could be the case that increasing and decreasing y'_{ki} gives allocations lying on different hyperplanes, meaning that the transformation from y' to x' is different on both sides of y_{ki} and therefore we conclude that $\widehat{u}_{ki}$ could be non-differentiable w.r.t y'_{ki} at s. Important thing here however is to notice that right and left derivatives exist, it's just that they may not be equal. Hence we can take derivatives on both sides of y_{ki} as (noting that derivative of the other terms in utility involving x_{ki} or involving m^l_k will be zero due Lemma 3.3)

$$\left.\frac{\partial \widehat{u}_{ki}{}'}{\partial y_{ki}}\right|_{y'_{ki}\downarrow y_{ki}} = \left(v'_{ki}(x_{ki}) - \sum_{l\in\mathcal{L}_{ki}} p^l_{ki}\alpha^l_{ki}\right) \left.\frac{\partial x'_{ki}}{\partial y'_{ki}}\right|_{y'_{ki}\downarrow y_{ki}} \tag{2.35a}$$

$$\left.\frac{\partial \widehat{u}_{ki}{}'}{\partial y_{ki}}\right|_{y'_{ki}\uparrow y_{ki}} = \left(v'_{ki}(x_{ki}) - \sum_{l\in\mathcal{L}_{ki}} p^l_{ki}\alpha^l_{ki}\right) \left.\frac{\partial x'_{ki}}{\partial y'_{ki}}\right|_{y'_{ki}\uparrow y_{ki}} \tag{2.35b}$$

We will first show that the $\partial x_{ki}/\partial y_{ki}$ term above (for either equation) is always positive. If $y = 0$ then clearly this is true, because if any agent ki demands $y_{ki} > \epsilon$ while $y_{-ki} = 0$ then clearly $x_{ki} > 0$ (in fact the allocation is differentiable at $y = 0$). If $y \neq 0$ from (2.10a), we can write

$$\beta := \frac{\partial x_{ki}}{\partial y_{ki}} = \frac{\partial (r y_{ki})}{\partial y_{ki}} = r + y_{ki}\frac{\partial r}{\partial y_{ki}} = r^q + y_{ki}\frac{\partial r^q}{\partial y_{ki}} \tag{2.36}$$

where $r = r^q$.

From here we divide our arguments into following cases: (A) $ki \notin \mathcal{N}^q$ ($\Leftrightarrow i \notin \mathcal{G}^q_k$); (B) $ki \in \mathcal{N}^q$, $i \notin \arg\max_{j\in\mathcal{G}^q_k}\{\alpha^q_{kj} y_{kj}\}$ and (C) $ki \in \mathcal{N}^q$, $i \in \arg\max_{j\in\mathcal{G}^q_k}\{\alpha^q_{kj} y_{kj}\}$.

(A) For this clearly $\partial r^q/\partial y_{ki} = 0$ and this makes $\beta = r^q > 0$.

(B) Since value of r^q depends only on the value of $(n^q_k)_{k\in\mathcal{K}^q}$, and in this case changes in y_{ki} don't affect n^q_k we can see that $\partial r^q/\partial y_{ki} = 0$ and so $\beta = r^q > 0$.

(C) We divide this case into two cases: $|S^q(y)| \geq 2$ or $|S^q(y)| = 1$. If $|S^q(y)| \geq 2$ then

$$\beta = r^q + y_{ki}\frac{\partial r^q}{\partial y_{ki}} = r^q + y_{ki}\left(-\frac{c^q}{(\sum_{k_0\in\mathcal{K}^q} n^q_{k_0})^2}\right)\frac{\partial n^q_k}{\partial y_{ki}}. \tag{2.37}$$

Now $\frac{\partial n^q_k}{\partial y_{ki}}$ is either α^q_{ki} or 0. If it is 0 then $\beta = r^q > 0$. Otherwise we have

$$\beta = \frac{(r^q)^2}{c^q}\sum_{k_0\in\mathcal{K}^q\setminus\{k\}} n^q_{k_0} \tag{2.38}$$

which is positive because $|S^q(y)| \geq 2$, since then there is at least one positive term in the summation. For $|S^q(y)| = 1$, we will consider $S^q(y) = \{k\}$; else in case $S^q(y) = \{k_0\} \neq \{k\}$, taking the derivative would give the same expression as above. For $S^q(y) = \{k\}$ we will get

$$r^q = \frac{c^q}{n^q_k} - \frac{c^q}{n^q_k(n^q_k+1)} = \frac{c^q}{n^q_k}\left(1 - \frac{1}{n^q_k+1}\right) = \frac{c^q}{n^q_k+1} \tag{2.39a}$$

$$\beta = r^q + y_{ki}\,\frac{\partial r^q}{\partial n^q_k}\,\frac{\partial n^q_k}{\partial y_{ki}} \tag{2.39b}$$

where $\frac{\partial n_k^q}{\partial y_{ki}}$ is either 0 or α_{ki}^q. If it is 0 then $\beta = r^q > 0$ and if it is α_{ki}^q then we have $\beta = \frac{(r^q)^2}{c^q}$. Hence we have $\beta > 0$ in all cases.

Referring to (2.35), there are two possibilities, the first term on RHS in both equations in (2.35) is positive or negative. If it's positive, then we can see from (2.35a) that by increasing y'_{ki} from y_{ki} (and therefore x'_{ki} from x_{ki}) agent ki can increase his pay-off, which contradicts equilibrium. Now similarly consider the first term in RHS of (2.35) to be negative, then from (2.35b), agent ki can reduce y'_{ki} from y_{ki} to get a better pay-off. But the downward deviation in y'_{ki} is only possible if $y_{ki} > 0$ ($\Leftrightarrow \; x_{ki} > 0$). Hence we conclude (2.17a), (2.17b) from the statement of this lemma. □

6.5. Proof of Lemma 3.7

Proof. We will show that $s = (y, Q)$ is a NE of the game $\mathfrak{G}$, where for each agent $ki \in \mathcal{N}$, $y_{ki} = \rho x_{ki}^*$ (for any $\rho > 0$) and

$$^1q_{ki}^l = {\mu_{ki}^l}^* \quad \text{and} \quad ^2q_{ki}^l = {\mu_{ke}^l}^*, \tag{2.40}$$

where agent ke is the one identified by index $g_k^l(i) + 1$. Here x^*, μ^* are the primal-dual variables that satisfy the KKT conditions of (CP).

Define the function $f : \mathcal{S}_{ki} \to \mathbb{R}$ as follows

$$\begin{aligned} f(s'_{ki}) &\triangleq \widehat{u}_{ki}(s'_{ki}, s_{-ki}) && (2.41a) \\ &= v_{ki}\Big(h_{ki}^x(s'_{ki}, s_{-ki})\Big) - h_{ki}^t(s'_{ki}, s_{-ki}). && (2.41b) \end{aligned}$$

The above claim will be shown in two steps: In Step 1 we show that message s results in allocation x^* and additionally we have $\nabla f(s_{ki}) = 0$. In Step 2 we show that for any agent ki, unilateral deviations from s_{ki} are not strictly profitable.

Step 1. Assumption (A4) implies that with $y = \rho x^*$, while calculating allocation $\Big(h_{ki}^x(s)\Big)_{ki \in \mathcal{N}}$, sub-case 1 in eq. (2.9c) will be valid. This is turn implies that $h^x(s) = x^*$. Thus, message $y = \rho x^*$ produces allocation x^*. Now, using the complimentary slackness property of KKT and the fact that in the claimed NE above, we have for any $ki \in \mathcal{N}$, $^2q_{ki}^l = {}^1q_{ke}^l = {\mu_{ke}^l}^*$, it can be easily shown that the partial derivative of f w.r.t. $^1q_{ki}^l$ evaluated at the point s_{ki} is zero and the same is true for the partial derivative w.r.t ${}^2q^l_{ki}$. Finally using the stationarity condition in KKT we deduce that the partial derivative of f w.r.t. y_{ki} evaluated at the point s_{ki} is zero. Thus, $\nabla f(s_{ki}) = 0$.

Step 2. We now check for profitable deviations. This task is accomplished in two steps. In Step 2a, we show that any interior local extrema $\tilde{s}_{ki} = \left(\tilde{y}_{ki}, {}^1\tilde{q}^l_{ki}, {}^2\tilde{q}^l_{ki}\right)$ of $f(\cdot)$ satisfies $\nabla f(\tilde{s}_{ki}) = 0$. Furthermore, we show that any interior local extrema $\tilde{s}_{ki}$ produces allocation $h^x(\tilde{s}_{ki}, s_{-ki})$ and prices that satisfy the KKT conditions as x^*, μ^*. In Step 2b, using Step 2a, we show that all interior local extrema are local maxima. As a result of this and using the fact that by construction, $f(s'_{ki})$ is continuous w.r.t. s'_{ki}, we conclude that $f(\cdot)$ can have only a single local maximum and this is also the global maximum. This is because all local extrema are necessarily characterized by $\nabla f(\tilde{s}_{ki}) = 0$ (Step 2a) and there cannot be multiple local maxima in the interior without a local minima between them. Furthermore, using this argument we know that this global maximizer has to be in the interior. Finally, using the fact that $\nabla f(s_{ki}) = 0$ (Step 1) we conclude that this maximizer is s_{ki}. In other words, agent ki cannot be strictly better off by unilaterally deviating from s_{ki}. As this is checked for any agent $ki \in \mathcal{N}$, we conclude that above mentioned message $s = (y, Q)$ is a NE.

Step 2a. First we will show that without the gradient being zero, there cannot be a local extremum. Note that by construction, $f(s'_{ki})$ is continuous w.r.t. s'_{ki}. Since $f(\cdot)$ is differentiable w.r.t. ${}^1q^l_{ki}, {}^2q^l_{ki}$ variables (at all points), derivatives w.r.t. these indeed have to be 0 at any interior local extremum. This implies equal prices and complimentary slackness properties (using arguments from respective proofs) and for the remaining term we can write

$$\frac{\partial f}{\partial y_{ki}} = \left(v'_{ki}(x_{ki}) - \sum_{l \in \mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki}\right) \left(\frac{\partial x_{ki}}{\partial y_{ki}}\right). \tag{2.42}$$

Note that in the proof of Lemma 3.6, we have shown that $\beta := \frac{\partial x_{ki}}{\partial y_{ki}} > 0$ always. So at the points of non-differentiability, β will have a jump discontinuity, however it will be positive on either side. It is then clear that without satisfying $v'_{ki}(x_{ki}) = \sum_{l \in \mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki}$, there cannot be a local extremum. This implies that $\nabla f(\tilde{s}_{ki}) = 0$ at any interior local extremum. Furthermore, KKT conditions must be satisfied at the corresponding allocation $h^x(\tilde{s}_{ki}, s_{-ki})$ and prices μ^*, λ^*. For this note that ${}^1\tilde{q}^l_{ki} = {}^2q^l_{kj}$ (where kj is the agent with index $g^l_k(i) - 1$) and from (2.40) we know that ${}^2q^l_{kj} = {\mu^l_{ki}}^*$. Thus ${}^1\tilde{q}^l_{ki} = {\mu^l_{ki}}^*$ and similarly ${}^2\tilde{q}^l_{ki} = {\mu^l_{ke}}^*$ (where ke is the agent with index $g^l_k(i)+1$). This combined with $v'_{ki}(x_{ki}) = \sum_{l\in\mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki}$ from above, implies that KKT is satisfied at the local extrema.

Step 2b. The Hessian H of $f(\cdot)$ w.r.t. s'_{ki} is of size $(2L_{ki}+1)\times(2L_{ki}+1)$, where the 1st row and column represent y'_{ki} and the two subsequent sets of L_{ki} rows and columns represent $\{{}^1{q^l_{ki}}'\}_{l\in\mathcal{L}_{ki}}$ and $\{{}^2{q^l_{ki}}'\}_{l\in\mathcal{L}_{ki}}$. We want H (evaluated at any local extremum) to be negative definite. Looking at the diagonal entries, and using what we have derived in the previous paragraph, it can be calculated that the only non-zero entries at any local extrema are negative $f_{{}^1q{}^1q} := \frac{\partial^2 f}{\partial^1 q^l_{ki}\partial^1 q^l_{ki}} = -2$, $f_{{}^2q{}^2q} := \frac{\partial^2 f}{\partial^2 q^l_{ki}\partial^2 q^l_{ki}} = -2$, $f_{yy} := \frac{\partial^2 f}{\partial y^2_{ki}} = v''_{ki}(x_{ki})\left(\frac{\partial x_{ki}}{\partial y_{ki}}\right)^2$ (due to strict concavity of v_{ki}). Also notice that all off-diagonal entries, except the first $L_{ki}+1$ in the first row and column, are zero. Finally, note that due to assumption (A2), all prices are finite at local extremum and so the aforementioned non-zero entries

$$f_{{}^1qy} := \frac{\partial^2 f}{\partial^1 q^l_{ki}\partial y_{ki}} = \eta\, {}^2q^l_{kj}\left(\alpha^l_{ki}\frac{\partial x_{ki}}{\partial y_{ki}} - \frac{\partial m^l_k}{\partial y_{ki}}\right) + \xi\, \bar{w}^l_{-k}\left(\sum_{k'\in\mathcal{K}^l}\frac{\partial m^l_{k'}}{\partial y_{ki}}\right) \tag{2.43}$$

are finite. We now show that the roots of the characteristic polynomial of H (i.e. its eigenvalues) all become negative for η, ξ chosen to be sufficiently small.

For this, we take a generic matrix A_0, which is similar in structure to H and whose entries have the same dependence on $|y|$ as H. So A_0 is of the form $A_0 = \begin{bmatrix} A & 0 \\ 0 & D \end{bmatrix}$ where $D = (-2)I_{L_{ki}}$. In this case we know that eigenvalues of A and D together will give us all the eigenvalues of A_0. Clearly eigenvalues of D are -2 repeated L_{ki} times, so all that we now need to do is check whether all eigenvalues of A are negative. Entries in A are $a_{11} = -\frac{a}{|y|^2}$, $a_{ij} = a_{ji} = 0\ \forall\ i,j > 1,\ i \neq j$ and $a_{ii} = -2$, $a_{1i} = a_{i1} = \eta\frac{b_{i-1}}{|y|}\ \forall\ 2 \leq i \leq L_{ki}+1$. where $a > 0$ (and we don't care about the sign of b_i's). The factor of η in front of a_{i1}, a_{1i} is to be taken as $\max(\eta,\xi)$, but since we can set either of them we take it simply as η here. The parameters a, b_i may not be completely independent of y but since the absolute value of y has been taken out of the scaling, their values are bounded. Magnitude of b_i's are bounded from above and a is bounded away from zero.

We can explicitly calculate $|A - \lambda I|$ and write the characteristic equation as

$$Q(\lambda) = \left(-\frac{a}{|y|^2} - \lambda\right)(-2-\lambda)^{L_{ki}} + \eta^2 \frac{\sum_{i=1}^{L_{ki}}(-1)^i b_i^2}{|y|^2}(-2-\lambda)^{L_{ki}-1} = 0. \tag{2.44}$$

So -2 is a repeated eigenvalue, $L_{ki} - 1$ times. The equation for the remaining two roots can be written as

$$\left(-\frac{a}{|y|^2} - \lambda\right)(-2-\lambda) + \eta^2 \frac{C}{|y|^2} = 0. \tag{2.45}$$

Necessary and sufficient conditions for both roots of this quadratic to be negative are

$$\left(2 + \frac{a}{|y|^2}\right) > 0, \qquad \frac{2a}{|y|^2} + \eta^2 \frac{C}{|y|^2} > 0, \tag{2.46}$$

the first of which is always true, since $a > 0$. The second one can be ensured by making η small enough, since a is bounded away from zero and the magnitude of C is bounded from above.

Hence the Hessian H is shown to be negative definite for η chosen to be small enough. □

6.6. Proof of Lemma 3.8

Proof. Because of Lemma 3.3, the only non-zero term in t_{ki} (see (2.12)) at equilibrium is $x_{ki} \sum_{l \in \mathcal{L}_{ki}} \alpha_{ki}^l p_{ki}^l$, which is clearly non-negative. Hence $\sum_{ki \in \mathcal{N}} t_{ki} \geq 0$ at equilibrium.

Now if $x_{ki} = 0$ then we know from Lemma 3.3 and (2.12) that $t_{ki} = 0$ and so (2.19) is evident. Now take $x_{ki} > 0$ and define the function

$$f(z) = v_{ki}(z) - z \sum_{l \in \mathcal{L}_{ki}} \alpha_{ki}^l p_{ki}^l. \tag{2.47}$$

Note that $f(0) = u_{ki}(0,0)$ and $f(x_{ki}) = u_{ki}(x,t)$, the utility at equilibrium. Since $f'(x_{ki}) = 0$ (Lemma 3.6), we see that $\forall\ 0 < y < x_{ki}$, $f'(y) > 0$ since f strictly concave (because of v_{ki}). This clearly implies $f(x_{ki}) \geq f(0)$. □

6.7. Proof of Lemma 4.2

Proof. Suppose not, i.e. assume $\exists\ kj \in \mathcal{N}$ such that $\rho_{kj} \neq r$. In this case agent kj can deviate with only changing $\rho'_{kj} = r$ (which also means r is the same as before deviation, since demand y doesn't change). It's easy to see that this is a profitable deviation, since change in utility of agent kj will be only through the term involving ρ_{kj}.

$$\Delta \hat{u}_{kj} = -\zeta(\rho'_{kj} - r)^2 + \zeta(\rho_{kj} - r)^2 = \zeta(\rho_{kj} - r)^2 > 0. \tag{2.48}$$

□

6.8. Proof of Existence in Section 4

Proof. First order conditions can again be shown to be satisfied, the only difference is that here we will also use $\rho_{ki} = r$ at local extremum. The Hessian H for any agent ki here will be of order $(2L_{ki}+2) \times (2L_{ki}+2)$ where 1st, 2nd row and column represent y_{ki}, ρ_{ki} respectively whereas the remaining rows and columns represent ${}^1q^l_{ki}$'s and ${}^2q^l_{ki}$'s. The generic matrix A_0 for H will then be of the form $A_0 = \begin{bmatrix} A & 0 \\ 0 & D \end{bmatrix}$ where $D = (-2)I_{L_{ki}}$ and matrix A, of order $(L_{ki}+2) \times (L_{ki}+2)$, has elements $a_{11} = -\frac{a}{|y|^2} - \zeta\frac{d}{|y|^4}$, $a_{12} = a_{21} = -\zeta\frac{e}{|y|^2}$, $a_{ij} = a_{ji} = 0$ $\forall\ i,j > 1,\ i \neq j$, $a_{22} = -2$, $a_{ii} = -2$, and $a_{1i} = a_{i1} = \eta\frac{b_{i-1}}{|y|}$ $\forall\ 3 \leq i \leq L_{ki}+2$. where $a, d, e > 0$. As before, it will be shown that all eigenvalues of A are negative (since that is clearly true for D). Writing the characteristic equation, it is the case again that -2 is a repeated eigenvalue, L_{ki} times. And the equation for remaining two roots is

$$\lambda^2 + \lambda\left(2 + \frac{a}{|y|^2} + \zeta\frac{d}{|y|^4}\right) + \left(\frac{2a}{|y|^2} + \zeta\frac{2d}{|y|^4} - \zeta^2\frac{E}{|y|^4} + \eta^2\frac{C}{|y|^2}\right) = 0 \tag{2.49}$$

Necessary and sufficient conditions for the roots of above quadratic to be negative are again that coefficient of λ and the constant term are both positive. Coefficient of λ is clearly positive, and the constant term can also be made positive by choosing ζ, η (& ξ) small enough, irrespective of sign of C. □

6.9. Proof of Lemma 4.3

Proof. Terms 2, 3, 4 and 5 in (2.24) and (2.25) are zero at equilibrium and so we can write (at equilibrium)

$$\sum_{ki\in\mathcal{N}} t_{ki} = \sum_{ki\in\mathcal{N}} x_{ki}\left(\sum_{l\in\mathcal{L}_{ki}} \alpha^l_{ki} p^l_{ki}\right) - r\sum_{l\in\mathcal{L}_{ki}} \frac{1}{N^l-1} \sum_{k'j'\in\mathcal{N}^l\setminus\{ki\}} \alpha^l_{k'j'} p^l_{k'j'} y_{k'j'} \tag{2.50a}$$

$$= \sum_{ki\in\mathcal{N}} \sum_{l\in\mathcal{L}_{ki}} \left(x_{ki}\alpha^l_{ki}p^l_{ki} - \frac{1}{N^l-1}\sum_{k'j'\in\mathcal{N}^l\setminus\{ki\}} \alpha^l_{k'j'}p^l_{k'j'}x_{k'j'}\right). \tag{2.50b}$$

The coefficient of x_{ki} for any agent ki in the above is

$$\sum_{l\in\mathcal{L}_{ki}} \left(\alpha^l_{ki}p^l_{ki} - \frac{1}{N^l-1}\sum_{k'j'\in\mathcal{N}^l\setminus\{ki\}} \alpha^l_{ki}p^l_{ki}\right) = 0 \tag{2.51}$$

which proves the claim. □

Chapter 3.

A General Mechanism Design Methodology for Social Utility Maximization with Linear Constraints

Social utility maximization refers to the process of allocating resources in such a way that the sum of agents' utilities is maximized under the system constraints. Such allocation arises in several problems in the general area of communications, including unicast (and multicast multi-rate) service on the Internet, as well as in applications with (local) public goods, such as power allocation in wireless networks, spectrum allocation, etc. Mechanisms that implement such allocations in Nash equilibrium have also been studied but either they do not possess full implementation property, or are given in a case-by-case fashion, thus obscuring fundamental understanding of these problems.

In this chapter we propose a unified methodology for creating mechanisms that fully implement, in Nash equilibria, social utility maximizing functions arising in various contexts where the constraints are convex. The construction of the mechanism is done in a systematic way by considering the dual optimization problem. In addition to the required properties of efficiency and individual rationality that such mechanisms ought to satisfy, three additional design goals are the focus of this chapter: a) the size of the message space scaling linearly with the number of agents (even if agents' types are entire valuation functions), b) allocation being feasible on and off equilibrium, and c) strong budget balance at equilibrium and also off equilibrium whenever demand is feasible.

1. Introduction

In the general area of communications, a number of decentralized resource allocation problems have been studied in the context of mechanism design. Such problems include unicast service on the Internet [YaHa05], [MaBa04], [KaTe13a], multi-rate multicast service on the Internet [KaTe13b], [SiAn13], power allocation in wireless networks [ShTe12], [Ha12a], [ZhGu11], spectrum allocation [HuBeHo06a] and pricing in a peer-to-peer network [Ne09]. The mechanism design framework is both appropriate and useful in the above problems since these problems are motivated by the designer's desire to allocate resources efficiently in the presence of strategic agents who possess private information about their level of satisfaction from the allocation.

Usually mechanism design solutions define a contract such that the induced game has at least one equilibrium (Nash, Bayesian Nash, dominant strategy, etc) that corresponds to the desired allocation. This is usually obtained with *direct* mechanisms by appealing to the revelation principle [BöKrSt15] [GaNaGu08a]. The drawbacks of these direct approaches is that they require agents to quote their types (which may be entire valuation functions) and that the induced game may have other extraneous equilibria that are not efficient. In this chapter, the focus is on **full Nash implementation**. Without going into a formal definition, full Nash implementation refers to the design of contracts such that only the designer's most preferred outcome is realized as a result of interaction between strategic agents (i.e. at Nash equilibria (NE)), as opposed to general mechanism design, where other less preferred outcomes are also possible. Thus implementation is more stringent and is capable of producing better allocations. Readers may refer to [Ja01] for a survey on implementation theory.

Concentrating on those proposed solutions in the literature that guarantee full Nash implementation, one observes a fragmented and case-by-case approach. One may ask how fundamentally different are for example the problems of unicast service on the Internet, multicast multi-rate service on the Internet, power allocation in wireless networks, etc, to justify a separately designed mechanism for each case. Alternatively, one may ask what are the common features in all these problems that can lead to a unified mechanism design approach. These questions provide the motivation for the work presented in this chapter.

In particular, our starting point is to state a class of problems as social utility maximization under linear inequality/equality constraints. We then proceed by characterizing their solution through Karush-Kuhn-Tucker (KKT) conditions. Analyzing the dual optimization problem is essential in our approach, because it hints at how the taxation part of the mechanism should be designed. Subsequently we define a general mechanism and show that it results in full Nash implementation when the agents' valuation functions are their private information[1]. Since the application domain of interest is in the area of communications, we give special emphasis on the size of the message space (as a consequence, VCG-type mechanisms, see [Kr09, Section 5.3], [BöKrSt15, Chapter 5], [ShLe09, Chapter 10], [GaNaGu08a], [GaNaGu08b], are inappropriate since they require quoting of types, which are entire valuation function in this set-up). A notable exception is the work [JoTs09, Section 5] and [YaHa07] that adapts VCG for a small message space (one dimensional message per user) and guarantees off-equilibrium feasibility. However in their work, full implementation can't be guaranteed and there is possibility of extraneous equilibria. In this chapter, the message space scales linearly with number of agents so that the proposed mechanism is scalable.

In addition to the stated goal of getting optimal allocations at all Nash equilibria, there are other auxiliary properties that are sought in a mechanism. An important such property is *individual rationality*; this requires that agents are weakly better off at NE than not participating in the mechanism at all. The purpose of this is to ensure that agents are willing to sign the contract in the first place. Another important property is *strong budget balance* (SBB) at NE. This requires that at NE, the total monetary transfer between agents (assuming quasi-linear utility functions) is zero. The unified mechanism proposed in this work possesses both these properties.

Finally, a mechanism may be endowed with auxiliary properties off equilibrium. These are meant to improve the applicability in practical settings. For instance, the NE is interpreted as the convergent outcome when agents in the system "learn" the game by playing it repeatedly, which implies that during this process the messages (and thus allocations and taxes) are off equilibrium. In our opinion, the most important auxiliary property is **feasibility of allo-**

[1]Use of Nash equilibrium as a solution concept itself requires some justification, since it applies only to games with complete information. This issue is discussed further in Section 5. However, in this work we accept the justification given by Reichelstein and Reiter in [ReRe88] and Groves and Ledyard in [GrLe77] based on the second interpretation of Nash equilibrium as the fixed point of an (possibly myopic) adjustment process, [Na50].

cation on and off equilibrium. This property is essential whenever the system constraints are hard constraints on resources and cannot be violated by any means. For instance, the contract should not promise rate allocations to agents that ever violate the capacity constraints of the network links because such a contract would be practically invalid since it promises something that can never be achieved. Note that, at equilibrium the allocation has to satisfy the constraints by definition, so feasibility always holds at NE. Another auxiliary property is SBB off equilibrium. Similarly this property guarantees that in a practical situation where a dynamic learning process converges to NE, each step of this process leaves a zero balance. One of the main features of the work in this chapter is that the allocation scheme is designed to guarantee both feasibility off equilibrium and SBB off equilibrium whenever demand is feasible. Specifically, feasibility is achieved by utilizing *proportional allocation*. Proportional allocation refers to the idea of using agents' demands to create their allocation by projecting their overall demand vector back down to the boundary of the feasible region whenever it is outside of it. This way everyone receives allocation that is proportional to their original demand. This allocation method generalizes the idea of proportional allocation that was introduced in mechanism design framework by [MaBa04], [YaHa05] for unicast network with one capacity constraint and for stochastic control of networks in [KeMaTa98], as well as by the authors in [SiAn14b], [SiAn13].

The main contributions of this work are summarized as follows.

- A unified framework is proposed for full implementation (in NE) mechanism design in social utility maximization problems with convex constraints.
- The proposed mechanism has small message space (although agents' types are entire valuation functions) and is scalable.
- The proposed mechanism is individually rational and strongly budget balanced at NE.
- The proposed mechanism lends itself to practical implementation (through learning algorithms) because allocation is feasible even off equilibrium (and is strongly budget balanced off equilibrium when demand is feasible).

The rest of the chapter is structured as follows. In Section 2 we describe three examples relevant in communications and formulate the general centralized allocation problem that we wish to implement. In Section 3 we describe the proposed mechanism. Section 4 contains all the proofs of full implementation. Section 5 contains off-equilibrium results and discusses their relevance. Finally in Section 6 we conclude with a discussion of important generalizations of this set-up.

2. Motivation and Centralized Problem

In this section we start by describing various important resource allocation problems that arise in communications and to which our generalized methodology applies (in Section 6 we discuss generalizations that will help solve an even larger class of problems than the ones described below). After the examples, we define a general centralized optimization problem, (CP), that covers all the examples. Following that we state general assumptions on (CP) (some additional technical assumptions will be made in Section 4).

2.1. Interesting Resource allocation problems in Communications

Unicast Transmission on the Internet

Consider agents on the Internet from set $\mathcal{N} = \{1, \ldots, N\}$, where each agent $i \in \mathcal{N}$ is a pair of source and destination that communicate via a pre-decided route consisting of links in $\mathcal{L}_i$. All the agents together communicate on the network consisting of links in $\mathcal{L} = \cup_{i \in \mathcal{N}} \mathcal{L}_i$. Since each link in the network has a limited capacity, this results in constraints on the information rate allocated to each agent. Considering a scenario where agents have (concave) utility functions/profiles $\{v_i(\cdot)\}_{i \in \mathcal{N}}$ that measure the satisfaction received by agents for various allocated rates, we can write the social utility optimization problem as

$$\max_{x \in \mathbb{R}_+^N} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{CP$_\mathrm{u}$}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{N}^l} \alpha_i^l x_i \leq c^l \quad \forall\, l \in \mathcal{L}. \tag{3.1}$$

In the above $\mathcal{N}^l \triangleq \{i \in \mathcal{N} \mid l \in \mathcal{L}_i\}$ is the set of agents present on link l, c^l is the capacity of link l and α_i^l are positive weights meant to differentiate between true information rate x_i and its imposition on capacity of the links of the network through coding rate, packet error etc.

The above feasible set is a polytope in the first quadrant of $\mathbb{R}_+^N$ and is created by faces that have outward normal vectors pointing away from the origin. For the details of a full implementation mechanism specifically for the unicast problem readers may refer to [JaWa10], [KaTe13a], [SiAn14b].

Public Goods

In contrast to the private consumption problem above, there are public goods problems where the resources are shared directly between agents instead of sharing via constraints. The unicast problem was a scenario where allocation of rate to one agent on a link would imply less available rate for other agents on the link (because of the capacity constraint) - the rate allocation thus is private in this case. In contrast to this, if there is a resource allocation problem where two or more agents can simultaneously use the same resource then it will be classified as a public goods problem. A well-studied example of this kind is the wireless transmission with interference, in which the power level of each agent affects any other agent through the signal-to-interference-and-noise ratio (SINR).

Here we consider a simplified scalar version of the public goods problem for a system of agents from the set $\mathcal{N}$ in the following form

$$\max_{x \in \mathbb{R}_+} \sum_{i \in \mathcal{N}} v_i(x) \tag{$\mathrm{CP_{pb}}$}$$

$$\text{s.t.} \quad x \in \mathcal{X}, \tag{3.2}$$

where $\mathcal{X}$ is a convex subset of $\mathbb{R}_+$. Note that the argument for all utility functions is the same; since there is one public good being simultaneously used by all the agents (for which x marks the usage level). We rewrite such a problem in the constrained form

$$\max_{\underline{x} \in \mathbb{R}_+^N} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{$\mathrm{CP_{pb}}$}$$

$$\text{s.t.} \quad x_1 = x_2 = x_3 = \ldots = x_N \tag{3.3a}$$

$$\text{s.t.} \quad x_1 \in \mathcal{X}. \tag{3.3b}$$

The treatment when x is a vector is not very different and we discuss this in the generalizations at the end. Implementation for the public goods problem is the most studied of all the examples in this chapter, see [GrLe77], [Hu79], [Ch02]. The distinction from private goods is that generally public goods problems require handling the "free-rider" problem [MaWhGr95, Section 11.C].

Local Public Goods

Another important resource allocation problem in communications is the local public goods problem. The basic idea of direct sharing of resources is same as above but in here the sharing is only among agents locally. So there are local groups of agents for whom the allocation has to be the same, but this common allocation can be different from one group to the next. If we divide the set of agents into disjoint local groups: $\mathcal{N} = \sqcup_{k \in \mathcal{K}} \mathcal{N}_k$, then the centralized problem can be written as

$$\max_{x \in \mathbb{R}_+^K} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}_k} v_i(x_k) \tag{CP$_{\text{lpb}}$}$$

$$\text{s.t.} \quad x_k \in \mathcal{X}_k \ \ \forall\, k \in \mathcal{K}, \tag{3.4}$$

with $(\mathcal{X}_k)_{k \in \mathcal{K}}$ being convex subsets of $\mathbb{R}_+$. As before, we would restate it as

$$\max_{x \in \mathbb{R}_+^K} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}_k} v_i(x_i) \tag{CP$_{\text{lpb}}$}$$

$$\text{s.t.} \quad x_i = x_j \ \ \forall\, i, j \in \mathcal{N}_k, \ \forall\, k \in \mathcal{K} \tag{3.5a}$$

$$\text{s.t.} \quad x_{j_k} \in \mathcal{X}_k \text{ for some } j_k \in \mathcal{N}_k, \ \forall\, k \in \mathcal{K}. \tag{3.5b}$$

Local public goods problems are relevant in those network settings where there is direct interaction between local agents. Wireless transmission is an example of this, each agent affects and is affected by other agents' transmission through interference (SINR) and it is reasonable to assume that this effect is only local and that agents situated far enough (either spatially or in the frequency domain) will not affect each other. Readers may refer to [ShTe12] for a specific mechanism for the local public goods problem.

2.2. General Centralized Problem

Here we state the generic form of the centralized optimization problem that we wish to fully implement. The resource allocation problem is defined for a system with agents indexed in the set $\mathcal{N} = \{1, 2, \ldots, N\}$, who have utility functions $\{v_i(\cdot)\}_{i\in\mathcal{N}}$. The objective is to find the optimum allocation of a single infinitely divisible good that maximizes the sum of utilities subject to constraints on the system. The allocation made to agents will be denoted by the vector $x \in \mathbb{R}_+^N$, with x_i being the allocation to agent i. The centralized optimization problem that we consider is

$$\max_x \sum_{i\in\mathcal{N}} v_i(x_i) \tag{CP}$$

$$\text{s.t.} \quad x \in \mathbb{R}_+^N \tag{C$_1$}$$

$$\text{s.t.} \quad A_l^\top x \le c_l \;\; \forall\, l \in \mathcal{L} \;\text{ where }\; A_l \in \mathbb{R}^N, c_l \ge 0. \tag{C$_2$}$$

The set $\mathcal{L} = \{1, 2, \ldots, L\}$ indexes all the constraints and A_l, c_l are all parameters of the optimization problem. It is easy to see that the above set-up covers equality constraints (such as from the public goods example), since we can always write $x_1 = \ldots = x_N$ as $x_1 \ge x_2 \ge \ldots \ge x_N \ge x_1$.

Denote by $\mathcal{C} \subset \mathbb{R}_+^N$ the above feasible set. Note that $\mathcal{C}$ is a polytope in the first quadrant of $\mathbb{R}^N$, possibly of a lower dimension than N (due to equality constraints). For convenience, we denote by $\mathcal{L}_i$ the set of constraints that involve agent i, i.e. $\mathcal{L}_i = \{l \in \mathcal{L} \mid A_{li} \neq 0\}$ with $L_i \triangleq |\mathcal{L}_i|$. Conversely, define $\mathcal{N}^l$ as the set of agents involved at link l i.e. $\mathcal{N}^l = \{i \in \mathcal{N} \mid l \in \mathcal{L}_i\}$ with $N^l \triangleq |\mathcal{N}^l|$.

2.3. Assumptions

Stated below are the assumptions on (CP) some of which restrict the environment $\{v_i\}_{i\in\mathcal{N}}$ and some the constraint set $\mathcal{C}$. Some additional technical assumptions will be introduced later to handle the degenerate cases.

(A1) For any $i \in \mathcal{N}$, $v_i(\cdot)$ is a strictly concave and continuously double differentiable function $\mathbb{R}_+ \to \mathbb{R}$.

The purpose of strict concavity is to have a (CP) whose solution can be described sufficiently by the KKT conditions (note that monotonicity is not assumed).

(A2) The optimal solution x^* is bounded such that $x^* \in \times_{i=1}^{N} (d_i, D)$ for some $0 < d_i < D$, with $d = (d_i)_{i=1}^{N} \in \mathcal{C}$ being arbitrarily close to $\underline{0}$ and D being large enough.

This assumption is used to eliminate corner cases of (CP), since they usually require special treatment and make the exposition unnecessarily convoluted. Note that we can always select a point $d \in \mathcal{C}$ that is arbitrarily close to $\underline{0}$ because of assumption (A3) below.

The next two assumptions restrict the constraint set $\mathcal{C}$.

(A3) The vector $\underline{0} \in \mathcal{C}$, i.e. $x = \underline{0}$ is feasible.

Since we are considering problems where all the variables x_i have a physical interpretation, it is natural to consider a constraint set whereby every agent getting 0 allocation is feasible. Note that this also explains the choice $c_l \geq 0$.

(A4) For any constraint $l \in \mathcal{L}$ in (C_2) there are two distinct $i, j \in \mathcal{N}$ such that $A_{li}, A_{lj} \neq 0$. This can also be stated as $N^l \geq 2 \; \forall \; l \in \mathcal{L}$ where N^l is as defined after (CP).

This ensures that there is indeed competition for all the constraints that could possibly be active at optimum. Again this is used to avoid special treatment of corner cases.

Denote by $\mathcal{V}_0$ the set of all possible functions $\{v_i(\cdot)\}_{i \in \mathcal{N}}$ that satisfy the above assumptions. Then $\mathcal{V}_0$ will be the environment for our mechanism design problem.

We will also make an assumption on the overall utility of agents in the system.

(A5) Apart from the valuation part $v_i(x_i)$ there is linear taxation component as well that affects agents' utilities. So overall utility of agent i is

$$u_i(x, t) = v_i(x_i) - t_i. \tag{3.6}$$

2.4. KKT conditions

The Lagrangian for the optimization problem (CP) is

$$L(x, \lambda, \mu) = \sum_{i \in \mathcal{N}} v_i(x_i) - \sum_{l \in \mathcal{L}} \lambda_l \left(A_l^\top x - c_l \right) + \sum_{i \in \mathcal{N}} \mu_i x_i \tag{3.7}$$

Due to assumption (A2), we can state the KKT conditions only in terms of λ^* and not involve μ^*.

1. Primal Feasibility: $x^* \in \mathcal{C}$.

2. Dual Feasibility: For all $l \in \mathcal{L}$, $\lambda_l^* \geq 0$.

3. Complimentary Slackness: For all $l \in \mathcal{L}$,

$$\lambda_l^* \left(A_l^\top x^* - c_l \right) = 0. \tag{3.8}$$

4. Stationarity: For all $i \in \mathcal{N}$,

$$v_i'(x_i^*) = \sum_{l \in \mathcal{L}} A_{li} \lambda_l^*. \tag{3.9}$$

We will see later that taxation will help us in achieving these KKT conditions when agents play the induced game from the mechanism. For that it will be important to think of the Lagrange multipliers above as "prices" where there will be one price per constraint.

3. Mechanism

In this section we describe the proposed mechanism. Description of the mechanism is divided into two parts - allocation and taxes. A mechanism in the Hurwicz-Reiter framework consists of an environment, an outcome space, a (centralized) correspondence between the environment and the outcome space, a message space and a contract from the message space to the outcome space. In our case the environment is set $\mathcal{V}_0$ of all possible valuation functions $\{v_i(\cdot)\}_{i \in \mathcal{N}}$. The outcome space is the Cartesian product of the set of all possible allocations and taxes, which is the set $\mathcal{C} \times \mathbb{R}^N \subset \mathbb{R}_+^N \times \mathbb{R}^N$. The correspondence between $\mathcal{V}_0$ and $\mathcal{C}$ is provided implicitly by the centralized problem (CP), where for each $\{v_i(\cdot)\}_{i \in \mathcal{N}}$ we get an optimal allocation x^* by solving (CP) (explicitly one can solve KKT to define x^*, together with corresponding Lagrange multipliers). This leaves the designer with the task of designing the message space and the contract.

The *message space* for our mechanism is $\mathcal{S} = \times_{i\in\mathcal{N}}\mathcal{S}_i$ with $\mathcal{S}_i = (d_i, +\infty) \times \mathbb{R}_+^{L_i}$ where messages from agents are of the form $s_i = (y_i, p_i)$ with $p_i = (p_i^l)_{l\in\mathcal{L}_i}$ and the total message is denoted by $s = (s_i)_{i\in\mathcal{N}} = (y, P)$ with $y = (y_i)_{i\in\mathcal{N}}$ and $P = (p_i)_{i\in\mathcal{N}}$. The message $s_i = (y_i, p_i)$ is to be interpreted as follows: y_i is the level of demand from agent i and p_i is the vector of prices that he believes everyone else should pay for the respective constraints. The *contract* $h : \mathcal{S} \to \mathbb{R}_+^N \times \mathbb{R}^N$ will specify allocation and taxes for all agents based on the message s, i.e., $h(s) = (h_{x,i}(s), h_{t,i}(s))_{i\in\mathcal{N}}$.

This contract along with agents' utilities will give rise to a one-shot game

$$\mathfrak{G} = (\mathcal{N}, \times_{i\in\mathcal{N}}\mathcal{S}_i, \{\widehat{u}_i\}_{i\in\mathcal{N}}) \tag{3.10}$$

between agents in $\mathcal{N}$ where action sets are $\mathcal{S}_i$ and utilities are

$$\widehat{u}_i(s) = u_i(x, t) = v_i(x_i) - t_i = v_i(h_{x,i}(s)) - h_{t,i}(s). \tag{3.11}$$

Information assumptions We assume that for any agent i, $v_i(\cdot)$ is his private information. The mechanism designer doesn't know $\{v_i(\cdot)\}_{i\in\mathcal{N}}$ but knows the set $\mathcal{V}_0$ to which they belong. Also the constraints (C_1) and (C_2) in (CP) (along with the assumptions) are common knowledge i.e. known to agents and the mechanism designer.

We say that the mechanism fully implements the centralized problem (CP) if **all** Nash equilibria of the induced game $\mathfrak{G}$ correspond to the unique allocation x^* - solution of (CP), and additionally individual rationality is satisfied i.e. agents are weakly better-off participating in the contract at equilibrium than not participating at all.

3.1. Allocation

We first describe the allocation in the case where the constraints in (C_2) do not have any effective degeneracy i.e. equality constraints. This distinction is based on whether the feasible set $\mathcal{C}$ has a proper interior or not. Clearly when there are no equality constraints, the constraint set will have an interior.

For allocation in this case, we first choose a point θ in the interior of the feasible set such

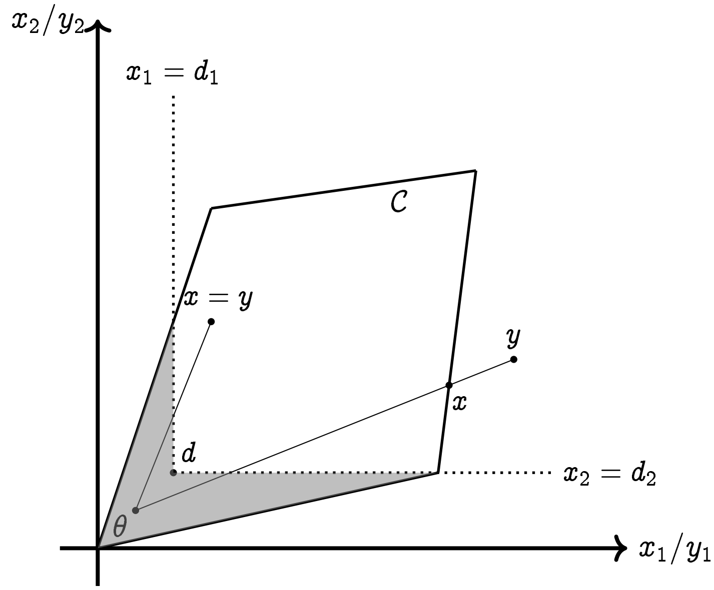

Figure 10.: An illustration of the allocation for $N = 2$

that

$$\theta \in \text{int}(\mathcal{C}) \cap \times_{i=1}^{N} (0, d_i) \tag{3.12}$$

Note that we can guarantee the existence of θ since by assumption (A3), $\underline{0} \in \mathcal{C}$ and clearly $\mathcal{C}$ being intersection of half-planes, is a convex set. Since d can be made arbitrarily close to $\underline{0}$, the same holds for θ as well.

Before formally defining the allocation, we define it informally with the help of Fig. 10. For any demand $y \in \mathcal{S}_y \triangleq \times_{i=1}^{N} (d_i, +\infty)$, the allocation x will be equal to y if y is inside the feasibility region $\mathcal{C}$. Otherwise the allocation will be the intersection point between the boundary of the feasibility region $\mathcal{C}$ and the line joining θ with y (that intersection point which lies between θ and y). The figure shows two different possible y's and their corresponding allocation x. The shaded region represents that part of $\mathcal{C}$ that can never be allocated. Note that since $||d||, ||\theta|| \approx 0$ this is a very small region and thus it doesn't significantly affect the generality of results presented in this chapter. Also this can be seen as a partial justification for why assumption (A2) was needed.

Formally, in case of feasibility set $\mathcal{C}$ not having any degeneracy, for a demand $y \in \mathcal{S}_y$, the allocation x generated by the contract is

$$x = \begin{cases} y & \text{if } y \in \mathcal{C} \\ y_0 & \text{if } y \notin \mathcal{C}, \end{cases} \tag{3.13}$$

where y_0 is the point on the boundary of $\mathcal{C}$ which is also on the line joining θ with y. Explicitly, if the above intersection happens on the hyperplane $\mathcal{F}_l = \{A_l^\top x = c_l\}$ then we can express y_0 as

$$y_0 - \theta = \alpha_0 (y - \theta) \quad \text{with} \quad \alpha_0 = \frac{c_l - A_l^\top \theta}{A_l^\top (y - \theta)} \tag{3.14}$$

The above allocation mapping is an extension of generalized proportional allocation idea (see [SiAn13] and [SiAn14b]), but modified in accordance with the generality of the problem (CP) and also so that points in the interior of $\mathcal{C}$ are covered as well. It is easy to verify in above that for $y \notin \mathcal{C}$, because of the way proportional allocation is defined, the expression for α_0 above is well-defined and positive. Also, if one extends the definition of y_0 to the boundary i.e $y \in \partial\mathcal{C}$ then it is easy to see that $y_0 = y$ and $\alpha_0 = 1$.

Another useful explicit way of defining α_0 is as follows.

$$\alpha_0 = \min_{\substack{l \in \mathcal{L} \\ \alpha_0^l > 0}} \alpha_0^l \quad \text{with} \quad \alpha_0^l = \frac{c_l - A_l^\top \theta}{A_l^\top (y - \theta)} \ \forall\, l \in \mathcal{L} \tag{3.15}$$

In the above, $c_l - A_l^\top \theta > 0$ for all $l \in \mathcal{L}$ (because $\theta \in \text{int}(\mathcal{C})$), but $A_l^\top (y - \theta)$ might be positive, negative or zero (since we are considering all $l \in \mathcal{L}$ and not restricting ourselves to appropriate regions as before). The purpose of taking minimum of α_0^l over positive ones is to find the "innermost" (or closest) constraint to θ in the (positive) direction of $y - \theta$.

Now we turn our attention to the case of degenerate feasible set. Since we are mainly interested in dealing with problems of (local) public goods nature, we define an alternate characterization of the constraint set only for those cases[2]. We can rewrite constraints from (C_2) to explicitly account for equalities. For this consider disjoint sets of agents $(\mathcal{N}_k)_{k \in \mathcal{K}}$, where $\mathcal{K} = \{1, \ldots, K\}$, such that $\mathcal{N} = \sqcup_{k \in \mathcal{K}} \mathcal{N}_k$, and the partition $(\mathcal{N}_k)_{k \in \mathcal{K}}$ of $\mathcal{N}$ serves the

[2] This generalizes in a straightforward way.

purpose of grouping agents by locality. Thus the rewritten constraints are

$$x_i = x_j \ \ \forall\, i, j \in \mathcal{N}_k, \ \forall\, k \in \mathcal{K} \tag{C$_{2a}$}$$

$$A_l^\top x \leq c_l \ \ \forall\, l \in \tilde{\mathcal{L}} \subseteq \mathcal{L}, \tag{C$_{2b}$}$$

where $\tilde{\mathcal{L}}$ is the subset of original constraints - the ones remaining after equality constraints have been separated. Without stating formally, we assume that the constraints in (C$_{2b}$) do not introduce any further degeneracy. Now one can consider a further refinement of above to state the constraints only in terms of free variables and thus containing only inequality constraints. For this define $\tilde{x} = (x_{j_k})_{k \in \mathcal{K}}$ where j_k represents the agent with lowest index from $\mathcal{N}_k$ (to make the representation unique). Clearly this is exactly the set of free variables from above and we can rewrite (C$_2$) as

$$\tilde{A}_l^\top \tilde{x} \leq c_l \ \ \forall\, l \in \tilde{\mathcal{L}}, \tag{C$_{2c}$}$$

with $\tilde{A}_l \in \mathbb{R}^K$ being derived from A_l by using (C$_{2a}$) explicitly. Recall that from assumption (A4) we had at least 2 agents on each constraint, in cases such as above we extend that to have at least 2 agents per constraint in the above definition. We denote by $\tilde{\mathcal{C}} \subset \mathbb{R}_+^K$ the effective polytope in the $\tilde{x}-$space, created by the constraints in (C$_{2c}$).

Coming back to allocation, for demand $y \in \mathcal{S}_y$ we first create modified demand $\tilde{y} \in \times_{i=1}^K (\tilde{d}_i, +\infty) \subset \mathbb{R}_+^K$ as follows:

$$\tilde{y}_k = \frac{1}{N_k} \sum_{i \in \mathcal{N}_k} y_i \ \ \forall\, k \in \mathcal{K}, \tag{3.17}$$

i.e. averaging quoted y's from each group to get one representative demand. Then proportional allocation, as above, is performed on $\tilde{y}$ and the feasibility set $\tilde{\mathcal{C}}$ to get $\tilde{x} \in \tilde{\mathcal{C}}$ (note that $\tilde{\theta}$ - the restriction of θ to appropriate coordinates, will be used in place of θ). This proxy allocation $\tilde{x}$ is finally converted back, using (C$_{2a}$), to get $x \in \mathcal{C}$.

3.2. Taxes

For any agent i, we define his tax t_i as

$$t_i = \sum_{l \in \mathcal{L}_i} t_i^l \tag{3.18a}$$

$$t_i^l = A_{li} x_i \bar{p}_{-i}^l + \left(p_i^l - \bar{p}_{-i}^l\right)^2 + \eta\, \bar{p}_{-i}^l p_i^l \left(c_l - A_l^\top x\right)^2 \tag{3.18b}$$

$$\bar{p}_{-i}^l \triangleq \frac{1}{N^l - 1} \sum_{\substack{j \in \mathcal{N}^l \\ j \neq i}} p_j^l \tag{3.18c}$$

With $\eta > 0$ being a positive constant.

In the results section we will see how the various terms in the tax above represent different KKT conditions. Note for example the price-taking nature that $A_{li} x_i \bar{p}_{-i}^l$ induces; here the agent is supposed to pay for his allocation but the price at which he has to pay is decided by others.

4. Equilibrium Results

In this section we state and prove lemmas that will give us the desired full implementation property for the mechanism defined in Section 3. Later in this section, the presented mechanism will be modified slightly to get an additional property of SBB at NE.

The main result of full implementation would require us to prove that all pure strategy NE of game $\mathfrak{G}$ result in allocation x^* - the unique solution of (CP), and to show individual rationality. The method for proving this result is as follows: firstly we show (Lemmas 4.1-4.5) that for any pure strategy NE the corresponding allocation and quoted prices must satisfy the KKT conditions from Section 2.4. Since by assumptions KKT conditions are both necessary and sufficient, this would mean that if pure strategy NE exist (unique or multiple), the corresponding allocation would have to be the solution of (CP) with quoted prices as the optimal Lagrange multipliers. Then we show existence (in Lemma 4.6) by explicitly writing out the messages that achieve the solution of (CP) and showing that there are no profitable unilateral deviations from those. Finally individual rationality will be shown (in Lemma 4.7).

For Lemma 4.5, we would need to distinguish between various cases arising from the general form of (CP). We will make specific assumptions when necessary.

Lemma 4.1 (Primal Feasibility). *For any message $s = (y, P) \in \mathcal{S}$ with corresponding allocation x, we have $x \in \mathcal{C}$ i.e. allocation is always feasible.*

Proof. Consider the case where $\mathcal{C}$ doesn't consist of equality constraints i.e. is not degenerate. For $y \in \mathcal{C}$, feasibility of allocation x is obvious. For $y \notin \mathcal{C}$, the allocation is chosen on the boundary of $\mathcal{C}$, hence it is feasible as well (note that $\mathcal{C}$ is closed).

In case when $\mathcal{C}$ is degenerate, the dummy allocation $\tilde{x}$ is feasible w.r.t. $\tilde{\mathcal{C}}$ using the above argument. But since the allocation x is produced from $\tilde{x}$ whilst satisfying the equality constraints, this means that x satisfies both the inequality and equality constraints and hence is feasible. □

Lemma 4.2 (Equal Prices). *At any NE $s = (y, P)$ of the game $\mathfrak{G}$, $\forall\, l \in \mathcal{L}$ and $\forall\, i \in \mathcal{N}^l$, we have $p_i^l = p^l$ i.e. all agents on a constraint quote the same price for that constraint.*

Proof. Please see Appendix 7.1 at the end of this chapter. □

Since we have shown that there is a common price p^l for any constraint $l \in \mathcal{L}$, here onwards we can refer only to p^l's instead of p_i^l's. Here onwards, whenever we refer to quoted price P, it is to be interpreted as $(p^l)_{l\in\mathcal{L}}$.

Lemma 4.3 (Dual Feasibility). *For any $l \in \mathcal{L}$, $p^l \geq 0$.*

Proof. This is obvious since quoted prices are all non-negative. □

Lemma 4.4 (Complimentary Slackness). *For any NE $s = (y, P)$ of game $\mathfrak{G}$ with corresponding allocation x, we have for any $l \in \mathcal{L}$*

$$p^l \left(c_l - A_l^\top x\right) = 0 \tag{3.19}$$

Proof. Please see Appendix 7.2 at the end of this chapter. □

Next we have the lemma for the stationarity property from KKT. For this we have to make assumptions on the feasible set $\mathcal{C}$.

(A6)

$$\widetilde{A}_l \in \mathbb{R}_+^K, \;\; \forall\, l \in \widetilde{\mathcal{L}} \tag{3.20}$$

where $\widetilde{A}_l$ was defined in (C_{2c}) (Section 3.1).

The above clearly translates into $A_l \in \mathbb{R}_+^N \;\; \forall\, l \in \mathcal{L}$ in case there is no degeneracy. But when there is degeneracy, the above assumption states that the effective polytope $\widetilde{\mathcal{C}}$ in the $\widetilde{x}-$space has faces whose outward normals are pointing away from the origin.

Lemma 4.5 (**Stationarity**). *At any NE $s = (y, P)$ of game $\mathfrak{G}$ with corresponding allocation x, we have for any agent $i \in \mathcal{N}$,*

$$v_i'(x_i) = \sum_{l \in \mathcal{L}_i} A_{li} p^l \tag{3.21}$$

Proof. Please see Appendix 7.3 at the end of this chapter. □

This completes the necessary part of our proof. The argument in the previous Lemmas applied to pure strategy NE if there exited any. Now we will prove the existence of pure strategy NE which will give the optimal allocation x^*.

Lemma 4.6 (**Existence**). *For game $\mathfrak{G}$, there exists NE $s = (y, P) \in \mathcal{S}$ such that the corresponding allocation x and prices $(p^l)_{l \in \mathcal{L}}$ satisfy the KKT conditions as x^* and $(\lambda_l^*)_{l \in \mathcal{L}}$, respectively.*

Proof. Please see Appendix 7.4 at the end of this chapter. □

Lemma 4.7 (**Individual Rationality**). *For any NE $s = (y, P)$ of game $\mathfrak{G}$ with corresponding allocation x and taxes t, we have for any $i \in \mathcal{N}$*

$$u_i(x, t) \geq u_i(0, 0) \tag{3.22}$$

Proof. Recall that $u_i(x, t) = v_i(x_i) - t_i$. For any agent $i \in \mathcal{N}$, define

$$f(z) \triangleq v_i(z) - z \sum_{l \in \mathcal{L}_i} A_{li} p^l \tag{3.23}$$

Note that $u_i(0, 0) = v_i(0) = f(0)$ and at NE $u_i(x, t) = v_i(x_i) - x_i \sum_{l \in \mathcal{L}_i} A_{li} p^l = f(x_i)$ (also recall that except the first term, all other tax terms go to zero at NE, refer (3.18b)).

By stationarity property we know $f'(x_i) = v_i'(x_i) - \sum_{l \in \mathcal{L}_i} A_{li} p^l = 0$ and by assumption (A1) it is clear that $f(\cdot)$ is strictly concave. Therefore $f'(y) > 0$ for $0 < y < x_i$ and we can claim by mean value theorem that $f(x_i) > f(0)$. □

Theorem 4.8 (Full Implementation). *All Nash equilibria $s = (y, P)$ of the game $\mathfrak{G}$ have the same corresponding allocation x. Furthermore, this allocation is identical to the unique solution of* (CP) *i.e. $x = x^*$. Also, all agents are weakly better-off at equilibrium than by not participating at all.*

Proof. Thus from Lemma 4.6 there exists at least one pure strategy NE. From Lemma 4.1-4.5, allocation x at any NE has to satisfy KKT conditions and knowing that KKT conditions are necessary and sufficient for optimality, it is clear that $x = x^*$. Finally, individual rationality was shown in Lemma 4.7 and with this we have our full implementation result. □

4.1. Strong Budget Balance (at equilibrium)

In this section we present a modification in the above mechanism so that in addition to above properties we also have SBB, i.e., $\sum_{i \in \mathcal{N}} t_i = 0$ at those equilibria where $y \in \mathcal{C}$.

The message space and allocation x is exactly same as before, the only modification will be to the taxes. Whenever the quoted demand is feasible, we have $x = y$. So from (3.18a), (3.18b) and previous lemmas we can write total tax at NE as

$$\sum_{i \in \mathcal{N}} t_i = \sum_{i \in \mathcal{N}} \sum_{l \in \mathcal{L}_i} t_i^l = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{N}^l} t_i^l = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{N}^l} A_{li} y_i p^l \tag{3.24}$$

To achieve SBB, our approach is to redistribute the above total tax amongst all the agents so that the amount of money received by an agent will not be controlled by that agent i.e. it will be only a function of y, p quoted by other agents. This will ensure that strategic decisions are still same as before and thus we would continue to have all the previous properties (Lemmas 4.1-4.6). Furthermore, we will redistribute taxes per constraint. So for any constraint $l \in \mathcal{L}$, $\sum_{i \in \mathcal{N}^l} t_i^l$ will be redistributed only amongst agents in $\mathcal{N}^l$. After appropriately redistributing the taxes, we will check individual rationality at equilibrium once again since the total tax paid by an agent might have a different value now. The reason we present SBB as a side result is that one can always use such a redistribution technique and thus SBB generally reduces to appropriate algebraic manipulation and can usually be added later on in the mechanism (for problems of these types).

Let T_i^l be the new tax where analogously the old tax (from (3.18b)) was t_i^l. Then the total tax for any agent i is $T_i = \sum_{l \in \mathcal{L}_i} T_i^l$. From above discussion, what we want is

$$T_i^l = t_i^l - f_i^l \left((y_j, p_j^l)_{\substack{j \in \mathcal{N}^l \\ j \neq i}} \right) \tag{3.25}$$

for any $i \in \mathcal{N}$ and $l \in \mathcal{L}_i$. The function f_i^l will be defined below.

For $l \in \mathcal{L}$ such that $N^l \geq 3$, in case where $A_{li} \geq 0 \; \forall \; i \in \mathcal{N}^l$, we have

$$f_i^l \equiv f_{i,1}^l = \frac{1}{N^l - 2} \sum_{\substack{j \in \mathcal{N}^l \\ j \neq i}} A_{lj} y_j \left(\bar{p}_{-i}^l - \frac{p_j^l}{N^l - 1} \right) \tag{3.26}$$

when there are negative coefficients in A_{li}, we can define f_i^l in the same manner as above, except that coefficients for those j that are involved in an equality constraint will be (scaled)[3] $\widetilde{A}_{lj}$ instead of A_{lj}.

For $l \in \mathcal{L}$ such that $N^l = 2$ (assuming $\mathcal{N}^l = \{i, j\}$), we divide the redistribution into 2 cases: first we consider the case $A_{li}, A_{lj} \geq 0$, for this we have

$$f_i^l \equiv f_{i,1}^l = A_{lj} y_j p_j^l. \tag{3.27}$$

If the coefficients A_{li} are not positive then we are essentially into the degenerate case, where $|A_{li}| = |A_{lj}| = 1$ with opposite signs. For this we have $f_i^l \equiv f_{i,1}^l = 0$ and same for agent j.

With the modified tax we will denote the new game as $\mathfrak{G}'$.

Lemma 4.9 (Strong Budget Balance at NE). *For the game $\mathfrak{G}'$, at all NE $s = (y, P)$ where $y \in \mathcal{C}$ we have*

$$\sum_{i \in \mathcal{N}} T_i = \sum_{i \in \mathcal{N}} \sum_{l \in \mathcal{L}_i} T_i^l = 0 \tag{SBB}$$

Proof. Please see Appendix 7.5 at the end of this chapter. □

After this, we have to check individual rationality. In cases when there is no degeneracy in

[3] the scaling factor being the inverse of the number of agents on l that are involved in constraints that make their allocation equal.

$\mathcal{C}$, due to assumption (A6), $A_{li} \geq 0 \; \forall \; l, i$. This means that with f_i^l we have redistributed taxes by paying agents back, thus the redistribution is indeed non-negative (can be seen explicitly from (3.26) and (3.27)). Thus if the NE was individually rational without this redistribution, it will continue to be so now.

For the cases involving degeneracy, for links associated with equality constraints the redistribution is zero. So individual rationality isn't affected. For other links, individual rationality follows from the argument in the previous paragraph after noting that effectively $\widetilde{A}_{li}$ were used in place of A_{li} for defining f_i^l (and by (A6), $\widetilde{A}_{li} \geq 0$).

5. Off-equilibrium Results

In this section we discuss (and prove) additional off-equilibrium properties of the mechanism. The original requirements of full implementation are only restricted to equilibrium properties. However we believe that in a realistic scenario, off-equilibrium properties are essential to justify working within a Nash-implementation framework. Nash equilibrium is an appropriate solution concept for complete information games i.e. games where it is assumed that all agents have complete knowledge about each others utility functions and also each agent knows that others know about his utility functions and so on (infinite hierarchy of beliefs). We are proposing this mechanism to be used in an informationally, and possibly physically, decentralized network setting like the Internet. The root of delving into mechanism design is that the designer doesn't have information about the utility functions, so to assume that agents themselves have all the information is impractical, especially for communications settings. In absence of complete information with agents, Nash implementation has still been used in literature - the justification being online or offline learning. In this formulation, before playing the actual game, agents participate in a multi-round learning process. After the agents have learnt about each others utilities (or more specifically the equilibrium action) they play the actual one-shot game.

In order to make the above learning model practical, there have to be real incentives involved for learning correctly. It is expected that while the learning process is still going on, in every round agents will quote demands and prices and receive allocation and taxes through the contract. We are interested in dealing with constrained resource allocation problems where the constraints could possibly be of a hard nature i.e. impossible to violate. An instance of this is the capacity constraint in unicast and multi-rate multicast examples. Hence if one assumes the learning model, then a necessary property for the mechanism would be that allocation is always (even off-equilibrium) feasible, since during the learning process agents will be playing off-equilibrium. To a lesser extent, if one is interested in SBB then same would have to be true of SBB as well. It is for these reasons that we introduced the proportional allocation, as a distinct feature of our mechanism, as opposed to the simple allocation used in [HeMa12], [KaTe13a], where allocation equals demand, i.e. $x = y$ everywhere.

Since the mid-90s, various learning models in Mechanism design and Game theory have been extensively studied. Readers may refer to [HeMa12] for studying how mechanism design and learning are handled together. Also one may refer to [Yo04], [Fu98], [Te12], [ShLe09], [MeOz11], [LaTe11] for a compendium of existing results regarding learning in a strategic set-up.

Now we extend the modification from Section 4.1 to achieve SBB off-equilibrium.

5.1. SBB off-equilibrium

Here we will define f_i^l from (3.25) differently so that we achieve the property of $\sum_{i \in \mathcal{N}} t_i = 0$ at all points where $y \in \mathcal{C}$, not just at NE. From the expressions in (3.24) as well as the description of f_i^l below, it will be clear that this method of redistribution only works when there are sufficient number of agents on every link. So here we will modify assumption (A4) so that

(A4′) For any $l \in \mathcal{L}$, $N^l \geq 5$ i.e. there are at least 5 agents on any constraint.

Also, we will only deal with the non-degenerate case where $A_{li} \geq 0 \; \forall \; l, i$. The corresponding results regarding degenerate cases would involve tedious case-by-case analysis and unnecessarily complicate the analysis of what should be a straightforward addition to the mechanism.

Now f_i^l is defined in three parts: $f_i^l\left((y_j, p_j^l)_{\substack{j \in \mathcal{N}^l \\ j \neq i}}\right) = f_{i,1}^l + f_{i,2}^l + f_{i,3}^l$. The three terms here are individually redistributing the three tax terms from (3.24).

$$f_i^l \equiv f_{i,1}^l = \frac{1}{N^l - 2} \sum_{\substack{j \in \mathcal{N}^l \\ j \neq i}} A_{lj} y_j \left(\bar{p}_{-i}^l - \frac{p_j^l}{N^l - 1} \right) \tag{3.28a}$$

$$f_{i,2}^l = \frac{N^l}{(N^l-1)^2(N^l-2)} \sum_{\substack{k>m \\ k,m\neq i}} \left(p_k^l - p_m^l\right)^2 \tag{3.28b}$$

$$f_{i,3}^l = \eta\left(f_{i,3a}^l + f_{i,3b}^l + f_{i,3c}^l\right) \tag{3.28c}$$

$$f_{i,3a}^l = \frac{2c_l^2}{(N^l-1)(N^l-2)} \sum_{\substack{k>m \\ k,m\neq i}} p_k^l p_m^l \tag{3.28d}$$

$$f_{i,3b}^l = \frac{2}{N^l-1} \sum_{\substack{j>q \\ j,q\neq i}} \left[\frac{1}{N^l-3} \sum_{\substack{k\in\mathcal{N}^l \\ k\neq i,j,q}} p_j^l p_q^l \phi(y_k) + \frac{1}{N^l-2} \sum_{\substack{k\in\mathcal{N}^l \\ k=j,q}} p_j^l p_q^l \phi(y_k)\right] \tag{3.28e}$$

where $\phi(y_k) = A_{kl}^2 y_k^2 - 2c_l A_{kl} y_k$. For $f_{i,3c}^l$ define the set $B \triangleq \{j,q\} \cap \{k,s\}$,

$$f_{i,3c}^l = \frac{4}{N^l-1} \sum_{\substack{j>q \\ j,q\neq i}} \left[\frac{1}{N^l-4} \sum_{\substack{k>s \\ k,s\neq i \\ |B|=0}} \psi + \frac{1}{N^l-3} \sum_{\substack{k>s \\ k,s\neq i \\ |B|=1}} \psi + \frac{1}{N^l-2} \sum_{\substack{k>s \\ k,s\neq i \\ k=j,s=q}} \psi\right] \tag{3.29}$$

where $\psi = \psi(p_j^l, p_q^l, y_k, y_s) = p_j^l p_q^l \left(A_{kl} y_k\right)\left(A_{sl} y_s\right)$. It is this last expression that necessitates the assumption (A4′). We will denote the game with the above modified taxes with $\mathfrak{G}''$.

Lemma 5.1 (Strong Budget Balance on and off-NE). *For the game $\mathfrak{G}''$, for all points in the message space $\mathcal{S}$ where $y \in \mathcal{C}$*

$$\sum_{i\in\mathcal{N}} T_i = \sum_{i\in\mathcal{N}} \sum_{l\in\mathcal{L}_i} T_i^l = 0 \tag{SBB}$$

Proof. Please see Appendix 7.6 at the end of this chapter. □

Now just as before, we check individual rationality (for the modified game $\mathfrak{G}''$). Overall utility for agent i after redistribution, at NE, is

$$\widehat{u}_i = v_i(x_i) - x_i \sum_{l\in\mathcal{L}_i} A_{li} p^l + \sum_{l\in\mathcal{L}_i} f_{i,1}^l + f_{i,2}^l + f_{i,3}^l \tag{3.30}$$

where the last summation is calculated at NE. At NE, we have equal prices, hence from (3.28b) it is clear that $\sum_{l\in\mathcal{L}_i} f_{i,2}^l$ term is zero. Since we are only dealing with the non-degenerate case in this section, it is clear from the definition of $f_{i,1}^l$ (see (3.28a)) that $\sum_{l\in\mathcal{L}_i} f_{i,1}^l \geq 0$. We know already, from Lemma 4.7, that $v_i(x_i) - x_i \sum_{l\in\mathcal{L}_i} A_{li} p^l - v_i(0) > 0$, furthermore since the functions v_i are strictly concave and $x_i > d_i > 0$ (by assumption (A2)) we can bound the difference in above inequality by a value bigger than zero.

Now note that $f_{i,3}^l$ is multiplied by a positive factor η which is still to be chosen. Since prices p^l, demand/allocation y_i and coefficients A_{li} are all absolutely bounded, we can always make the contribution of $\sum_{l\in\mathcal{L}_i} f_{i,3}^l$ small enough (without having the knowledge of equilibrium values in advance) so that the overall utility in (3.30) is bigger than $v_i(0)$. Hence we have individual rationality here as well.

6. Generalizations

Three quite interesting generalizations arise immediately from the set-up and analysis in this chapter. The first is a case where agents have utilities based on a vector allocation rather than a scalar allocation i.e., the multiple goods scenario. The objective in this case will be written as

$$\sum_{i\in\mathcal{N}} v_i(\underline{x}_i) \quad \text{with} \quad \underline{x}_i \in \mathbb{R}_+^{D_i}. \tag{3.31}$$

Note that the assumption of strict concavity can still be made. In a communications scenario, such an example can arise if the Internet agents have utility based on throughput as well as delay or packet error rate. In the case of the local public goods problem and in particular its instance that models the interaction between wireless agents, vector allocation for agent i models the power level of all users that can interfere with i. In this case the problem ($\mathrm{CP_{lpb}}$) can be restated as

$$\max_{x} \sum_{i\in\mathcal{N}} v_i(\underline{x}_i) \tag{C$_{\mathrm{vlpb}}$}$$

$$\text{s.t.} \quad \underline{x}_i \in \mathbb{R}_+^{D_i} \;\; \forall\, i\in\mathcal{N} \tag{3.32a}$$

$$\text{s.t.} \quad {E_i^k}^\top \underline{x}_i = {E_j^k}^\top \underline{x}_j \;\; \forall\, i,j\in\mathcal{N}_k,\; \forall\, k\in\mathcal{K} \tag{3.32b}$$

where $\mathcal{N}_k$ as before denote localities, however they needn't be disjoint now. Note that in the vector equations (3.32b), multiplication by matrices E_i^k accomplishes the task of selecting some coordinates from $\underline{x}_i$.

The second generalization is with problems which can be equivalently formulated in the form of (CP), but perhaps with the help of auxiliary variables. Consider the multi-rate multicast system, in which agents are divided into multicast groups, where agents within a multicast group communicate with exactly the same data but possibly at different quality (information rate) for each agent. This problem has both - private and public goods characteristics. Just like unicast, the limited capacity on links creates constraints on allocation here. However for saving bandwidth, only the highest demanded rate from each multicast group is transmitted on every link. This means that resources are directly shared within every group. If we index agents by a double-index $ki \in \mathcal{N}$ where k represents their multicast group and i is the sub-index within the multicast group k, we can state the optimization problem as

$$\max_{x \in \mathbb{R}_+^N} \sum_{ki \in \mathcal{N}} v_{ki}(x_{ki}) \tag{CP$_\mathrm{m}$}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}^l} \max_{i \in \mathcal{G}_k^l} \{\alpha_{ki}^l x_{ki}\} \leq c^l \quad \forall\, l \in \mathcal{L} \tag{3.33}$$

here $\mathcal{K}^l$ represents the set of multicast groups present on link l and $\mathcal{G}_k^l$ represents the set of sub-indices from multicast group k, that are present on link l. Mathematically, the important property to note above is that the feasibility region is indeed a polytope (since the constraints are piecewise linear). We can easily linearize the above by using auxiliary variables m_k^l as proxies for $\max\limits_{i \in \mathcal{G}_k^l} \{\alpha_{ki}^l x_{ki}\}$ indexed by group and link as follows

$$\max_{x, m} \sum_{ki \in \mathcal{N}} v_{ki}(x_{ki}) \tag{CP$'_\mathrm{m}$}$$

$$\text{s.t.} \quad \alpha_{ki}^l x_{ki} \leq m_k^l \quad \forall\, i \in \mathcal{G}_k^l,\ k \in \mathcal{K}^l,\ l \in \mathcal{L} \tag{3.34a}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}^l} m_k^l \leq c^l \quad \forall\, l \in \mathcal{L}. \tag{3.34b}$$

In the presence of auxiliary variables, the construction of the mechanism (esp. taxes) are slightly different in nature but follow the same philosophy as in this chapter. Readers may refer to [KaTe13b], [SiAn13] for a full implementing mechanism specifically for the multi-rate multicast problem. Other works for the multicast problem include finding decentralized algorithms to achieve social utility maximizing allocation [KaSaTa01], [StLiTe07] as well as determining optimal allocation via max-min fairness [SaTa99], [SaTa02]. The incorporation of this model into the unified design methodology is a research topic the authors are currently working on.

The third possible generalization would be where the constraint set, instead of being a polytope, is taken as general convex set (satisfying assumptions from Section 2.3). If one represents the centralized problem in the form a general convex optimization problem

$$\max_{x,m} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{3.35a}$$

$$\text{s.t.} \quad g_l(x) \leq 0 \quad \forall\, l \in \mathcal{L}, \tag{3.35b}$$

then, looking at the Lagrangian

$$L(x, \lambda) = \sum_{i \in \mathcal{N}} v_i(x_i) - \sum_{l \in \mathcal{L}} \lambda_l g_l(x), \tag{3.36}$$

one can construct the first term (payment) in the tax (see (3.18b)) in a similar way as before. After this, the rest of tax terms can be analogously constructed from the KKT conditions.

The aforementioned generalizations together can lead to fully implementing mechanisms with minimal message space that can solve an even larger class of problems of interest.

We conclude with a note on the robustness properties of these mechanisms that we would like to investigate. Recently there has been a lot of work related to robustness in mechanism design[4]. In addition to investigating the learning properties of our mechanism, we are also interested in investigating robustness of our mechanism (w.r.t. information i.e. beliefs of agents) and how the two might be related.

[4]see [AgFuHoKuTe12] for treatment of the same and [BeMo13] for a survey of the major robustness results in Mechanism design.

7. Appendix

7.1. Proof of Lemma 4.2

Proof. Suppose there exists a constraint $l \in \mathcal{L}$ such that $(p_j^l)_{j \in \mathcal{N}^l}$ are not all equal. Then there exists an agent $i \in \mathcal{N}^l$ such that $p_i^l > \bar{p}_{-i}^l$ (this is clear from definition in (3.18c)).

Consider the downwards price deviation $p_i^{l\prime} = \bar{p}_{-i}^l \geq 0$ for this agent (whilst keeping all other values in s_i the same). The difference in utility will be

$$\Delta \widehat{u}_i = \widehat{u}_i(s_i', s_{-i}) - \widehat{u}_i(s_i, s_{-i}) \tag{3.37a}$$

$$= -(p_i^{l\prime} - \bar{p}_{-i}^l)^2 + (p_i^l - \bar{p}_{-i}^l)^2 - \eta\, \bar{p}_{-i}^l p_i^{l\prime} \left(c_l - A_l^\top x\right)^2 + \eta\, \bar{p}_{-i}^l p_i^l \left(c_l - A_l^\top x\right)^2 \tag{3.37b}$$

$$= \underbrace{(p_i^l - \bar{p}_{-i}^l)^2}_{>0} + \eta\, \bar{p}_{-i}^l \underbrace{(p_i^l - \bar{p}_i^l)}_{>0} \left(c_l - A_l^\top x\right)^2 > 0 \tag{3.37c}$$

Hence such a unilateral deviation is profitable. Therefore at NE, all quoted prices for any constraint have to be the same. □

7.2. Proof of Lemma 4.4

Proof. Suppose there exists a constraint $l \in \mathcal{L}$ such that, at NE, $p^l > 0$ and $c^l > A_l^\top x$. We will take a downwards price deviation $p_i^{l\prime} = p^l - \varepsilon > 0$ for any agent $i \in \mathcal{N}^l$. Again we can write the difference in utility as

$$\Delta \widehat{u}_i = -(p_i^{l\prime} - \bar{p}_{-i}^l)^2 + 0 - \eta\, \bar{p}_{-i}^l p_i^{l\prime} \left(c_l - A_l^\top x\right)^2 + \eta\, \bar{p}_{-i}^l p_i^l \left(c_l - A_l^\top x\right)^2 \tag{3.38a}$$

$$= -(-\varepsilon)^2 + \eta\, p^l\,(\varepsilon) \left(c_l - A_l^\top x\right)^2 = -\varepsilon\,(\varepsilon - a) \tag{3.38b}$$

with $a = \eta\, p^l \left(c_l - A_l^\top x\right)^2 > 0$. Now choosing $0 < \varepsilon < \min(a, p^l)$ this can be made into a profitable deviation. □

7.3. Proof of Lemma 4.5

Proof. For any agent $i \in \mathcal{N}$ we can write

$$\frac{\partial \widehat{u}_i(s)}{\partial y_i} = \left(v_i'(x_i) - \frac{\partial t_i}{\partial x_i} \right) \frac{\partial x_i}{\partial y_i} \tag{3.39}$$

Note that $\beta \equiv \partial x_i / \partial y_i$ isn't always defined, since the allocation is continuous but only piecewise differentiable for $y \notin \mathcal{C}$. But right and left derivatives are always defined. Noting that $v_i'(x_i) - \frac{\partial t_i}{\partial x_i} = 0$ is equivalent to (3.21) (using previous lemmas characterizing NE), it will be sufficient for us if we show that $\beta > 0$. Since then without making $v_i'(x_i) - \frac{\partial t_i}{\partial x_i} = 0$ there is always an upwards or downwards deviation in y_i to make agent i strictly better-off.

For $y \in \mathcal{C}$ we have $x_i = y_i$, therefore $\beta = 1 \neq 0$. Otherwise, first consider the case where $\mathcal{C}$ is non-degenerate. Then

$$x_i = \theta_i + \alpha_0 \left(y_i - \theta_i \right) \quad \text{with} \quad \alpha_0 = \frac{c_l - A_l^\top \theta}{A_l^\top \left(y - \theta \right)} \tag{3.40a}$$

$$\frac{\partial x_i}{\partial y_i} = \alpha_0 + \left(y_i - \theta_i \right) \frac{\partial \alpha_0}{\partial y_i} \quad \text{with} \quad \frac{\partial \alpha_0}{\partial y_i} = \frac{-\alpha_0 A_{li}}{A_l^\top \left(y - \theta \right)} \tag{3.40b}$$

$$\Rightarrow \quad \frac{\partial x_i}{\partial y_i} = \alpha_0 \left(1 - \frac{A_{li} \left(y_i - \theta_i \right)}{A_l^\top \left(y - \theta \right)} \right) = \alpha_0 \Big(\sum_{j \in \mathcal{N}^l \setminus \{i\}} A_{lj} \left(y_j - \theta_j \right) \Big) \tag{3.40c}$$

Noting that $\alpha_0 > 0$ and that due to assumption (A4) there are always at least two agents on any constraint, clearly the above is positive (since $\theta_j < d_j < y_j \ \forall \ j$).

In the case when $\mathcal{C}$ does have equality constraints, note that allocations are created by composing the maps $y \overset{A}{\mapsto} \tilde{y} \overset{B}{\mapsto} \tilde{x} \overset{C}{\mapsto} x$ where B is proportional allocation on the set $\widetilde{\mathcal{C}}$ and A, C are linear with positive coefficients[5]. Thus using the fact that for proportional allocation we have $\beta > 0$ (from above) and that A, C have positive coefficients, we clearly have $\beta > 0$ here as well. □

7.4. Proof of Lemma 4.6

Proof. This proof will be done in two separate parts: first we will show that there exists $s \in \mathcal{S}$ such that the allocation through the contract is x^* and prices are Lagrange multipliers $(\lambda_l^*)_{l \in \mathcal{L}}$.

[5] Since they consist of averaging, (3.17), and assigning same value to multiple positions.

Secondly, we will show that for all claimed NE points, there doesn't exist a unilateral deviation that is profitable.

Existence of prices that match Lagrange multipliers λ_l^* is obvious, since agents can quote any price in $\mathbb{R}_+$ and dual feasibility says that $\lambda_l^* \geq 0$. For allocation to be the same as x^*, notice that due to assumption (A2) all possible solutions x^* lie in the set $\mathcal{Y} \triangleq \mathcal{C} \cap \mathcal{S}_y$ with $\mathcal{S}_y = \times_{i=1}^N (d_i, +\infty)$. Quoted demand y can be anywhere in the set $\mathcal{S}_y$ and for $y \in \mathcal{C}$ (i.e. $y \in \mathcal{Y}$) the allocation is $x = y$ and therefore each point $x \in \mathcal{Y}$ is achievable as allocation by quoting the same point $y = x \in \mathcal{Y}$. (However also note that points on the boundary of $\mathcal{Y}$ are achievable as allocation by many y's outside $\mathcal{Y}$ as well.)

Now we will check for unilateral profitable deviations. When quoted demand creates allocation as x^* and the quoted prices are equal and equal to λ^*, the utility for any agent i is

$$\hat{u}_i(s) = v_i(x_i^*) - x_i^* \sum_{l \in \mathcal{L}_i} A_{li} \lambda_l^* \tag{3.41}$$

Due to the fact that $f(x) = v_i(x) - x \sum_{l \in \mathcal{L}_i} A_{li} \lambda_l^*$ is strictly concave and $f'(x_i^*) = 0$ (Stationarity) we can conclude that x_i^* is a global maximizer of f. With this we have

$$\hat{u}_i(s) = v_i(x_i^*) - x_i^* \sum_{l \in \mathcal{L}_i} A_{li} \lambda_l^* \geq v_i(x_i) - x_i \sum_{l \in \mathcal{L}_i} A_{li} \lambda_l^* \tag{3.42a}$$

$$\geq v_i(x_i) - x_i \sum_{l \in \mathcal{L}_i} A_{li} \lambda_l^* - \sum_{l \in \mathcal{L}_i} \left(p_i^l - \lambda_l^*\right)^2 - \sum_{l \in \mathcal{L}_i} \eta\, p_i^l \lambda_l^* \left(c_l - A_l^\top x\right)^2 \tag{3.42b}$$

The above is true for any $(x_j)_{j \in \mathcal{N}}$ and $\left(p_i^l\right)_{l \in \mathcal{L}_i}$ non-negative and final inequality holds because the additional terms are non-positive. Now any unilateral deviation (s_i', s_{-i}) from agent i will result in utility for agent i which has the form as in (3.42b). Hence we have proved that with unilateral deviations from messages that have allocation x^* and prices λ^*, the corresponding agent can never be strictly better off. □

7.5. Proof of Lemma 4.9

Proof. Recall that due to Lemmas 4.2 and 4.4, at NE we have

$$\sum_{i \in \mathcal{N}} t_i = \sum_{i \in \mathcal{N}} \sum_{l \in \mathcal{L}_i} t_i^l = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{N}^l} t_i^l = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{N}^l} A_{li} y_i p^l \tag{3.43}$$

For completing the proof, it will be sufficient to show $\sum_{i \in \mathcal{N}^l} t_i^l = \sum_{i \in \mathcal{N}^l} f_i^l$ for any $l \in \mathcal{L}$, at NE. Let's begin with links with $N^l = 2$ (where $\mathcal{N}^l = \{i, j\}$). In case $A_{li}, A_{lj} \geq 0$ the total payment is

$$f_i^l + f_j^l = A_{lj} y_j p_j^l + A_{li} y_i p_i^l = \sum_{k \in \mathcal{N}^l} A_{lk} y_k p^l. \tag{3.44}$$

In the degenerate case, the tax paid by agents i, j at NE is $y_i p^l, y_j p^l$ with opposite signs. But note that $y_i = y_j$ at NE, since we are in the degenerate case. Hence the total tax is zero, which is the same as $f_i^l + f_j^l$. Now for other links,

$$\sum_{i \in \mathcal{N}^l} \frac{1}{N^l - 2} \sum_{\substack{j \in \mathcal{N}^l \\ j \neq i}} A_{jl} y_j \left(\bar{p}_{-i}^l - \frac{p_j^l}{N^l - 1} \right) = \frac{1}{N^l - 2} \sum_{i \in \mathcal{N}^l} \sum_{\substack{j \in \mathcal{N}^l \\ j \neq i}} A_{jl} y_j \left(\bar{p}_{-j}^l - \frac{p_i^l}{N^l - 1} \right) \tag{3.45a}$$

$$= \frac{N^l - 1}{N^l - 2} \sum_{j \in \mathcal{N}^l} A_{jl} y_j \bar{p}_{-j}^l - \frac{1}{N^l - 2} \sum_{j \in \mathcal{N}^l} \sum_{\substack{i \in \mathcal{N}^l \\ i \neq j}} A_{jl} y_j \frac{p_i^l}{N^l - 1} \tag{3.45b}$$

$$= \frac{N^l - 1}{N^l - 2} \sum_{j \in \mathcal{N}^l} A_{jl} y_j \bar{p}_{-j}^l - \frac{1}{N^l - 2} \sum_{j \in \mathcal{N}^l} A_{jl} y_j \bar{p}_{-j}^l = \sum_{j \in \mathcal{N}^l} A_{jl} y_j \bar{p}_{-j}^l = \sum_{j \in \mathcal{N}^l} A_{jl} y_j p^l \tag{3.45c}$$

(the proof when A_{lj} are possibly negative also follows from above after explicitly using that equality constraints would make some allocations equal to each other at NE). $\square$

7.6. Proof of Lemma 5.1

Proof. By rearranging the above sum: $\sum_{i \in \mathcal{N}} \sum_{l \in \mathcal{L}_i} T_i^l = \sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{N}^l} T_i^l$, we can see that it will be sufficient if we show that for any constraint $l \in \mathcal{L}$,

$$\sum_{i \in \mathcal{N}^l} T_i^l = 0 \quad \Leftrightarrow \quad \sum_{i \in \mathcal{N}^l} t_i^l = \sum_{i \in \mathcal{N}^l} f_i^l \left((y_j, p_j^l)_{\substack{j \in \mathcal{N}^l \\ j \neq i}} \right) \tag{3.46}$$

We will again show this term-by-term; we have already shown in the proof of Lemma 4.9 that the sum of $f_{i,1}^l$ is equal to the sum of payment (first) term from t_i^l, similarly the sum of $f_{i,2}^l$ and $f_{i,3}^l$ will be shown to be equal to the sum of the second and third terms from t_i^l, respectively. Starting from the second term

$$\sum_{i\in\mathcal{N}^l}\left(p_i^l-\bar{p}_{-i}^l\right)^2=\sum_{i\in\mathcal{N}^l}(p_i^l)^2+\left(\bar{p}_{-i}^l\right)^2-2p_i^l\bar{p}_{-i}^l \tag{3.47a}$$

$$=\frac{N^l}{N^l-1}\sum_{i\in\mathcal{N}^l}(p_i^l)^2-\frac{2N^l}{(N^l-1)^2}\sum_{\substack{k,m\in\mathcal{N}^l\\k>m}}p_k^lp_m^l \tag{3.47b}$$

$$=\frac{N^l}{(N^l-1)^2}\left[(N^l-1)\sum_{i\in\mathcal{N}^l}(p_i^l)^2-2\sum_{\substack{k,m\in\mathcal{N}^l\\k>m}}p_k^lp_m^l\right] \tag{3.47c}$$

$$=\frac{N^l}{(N^l-1)^2}\sum_{\substack{k,m\in\mathcal{N}^l\\k>m}}\left(p_k^l-p_m^l\right)^2 \tag{3.47d}$$

Now observing

$$\sum_{\substack{k,m\in\mathcal{N}^l\\k>m}}\left(p_k^l-p_m^l\right)^2=\frac{1}{N^l-2}\sum_{i\in\mathcal{N}^l}\sum_{\substack{k,m\in\mathcal{N}^l\\k>m\\k,m\neq i}}\left(p_k^l-p_m^l\right)^2 \tag{3.48}$$

we get the equality of second terms. For the third term

$$\sum_{i\in\mathcal{N}^l}\eta\, p_j^l\bar{p}_{-j}^l\left(c_l-\sum_{k\in\mathcal{N}^l}\alpha_k^ly_k\right)^2=\eta\, c_l^2\sum_{j\in\mathcal{N}^l}p_j\bar{p}_{-j}^l+\sum_{j\in\mathcal{N}^l}\eta\, p_j\bar{p}_{-j}^l\left(\sum_{k\in\mathcal{N}^l}\phi(y_k)\right) \tag{3.49a}$$

$$+2\eta\sum_{j\in\mathcal{N}^l}p_j\bar{p}_{-j}^l\left(\sum_{k>m}A_{kl}A_{ml}y_ky_m\right)\equiv\eta\,(T_1+T_2+T_3) \tag{3.49b}$$

where as before $\phi(y_k)=A_{kl}^2y_k^2-2c_lA_{kl}y_k$. Here we will equate T_1,T_2,T_3 with the sum of $f_{i,3a}^l, f_{i,3b}^l, f_{i,3c}^l$, respectively.

$$c_l^2\sum_{j\in\mathcal{N}^l}p_j\bar{p}_{-j}^l=\frac{2c_l^2}{N^l-1}\sum_{k>m}p_k^lp_m^l=\frac{2c_l^2}{(N^l-1)(N^l-2)}\sum_{i\in\mathcal{N}^l}\sum_{\substack{k>m\\k,m\neq i}}p_k^lp_m^l \tag{3.50}$$

This completes $f_{i,3a}$, now for $f_{i,3b}$

$$\sum_{j\in\mathcal{N}^l}p_j\bar{p}_{-j}^l\left(\sum_{k\in\mathcal{N}^l}\phi(y_k)\right)=\frac{2}{N^l-1}\sum_{j>m}\sum_{k\in\mathcal{N}^l}p_j^lp_m^l\phi(y_k) \tag{3.51a}$$

$$= \frac{2}{N^l - 1} \sum_{j>m} \Big(\sum_{k \neq j,m} p_j^l p_m^l \phi(y_k) + \sum_{k=j,m} p_j^l p_m^l \phi(y_k) \Big) \tag{3.51b}$$

$$= \frac{2}{N^l - 1} \sum_{i \in \mathcal{N}^l} \sum_{\substack{j>m \\ j,m \neq i}} \left(\frac{1}{N^l - 3} \sum_{k \neq i,j,m} p_j^l p_m^l \phi(y_k) + \frac{1}{N^l - 2} \sum_{k=j,m} p_j^l p_m^l \phi(y_k) \right) \tag{3.51c}$$

This completes $f_{i,3b}$ and finally the sum for $f_{i,3c}$ can be shown in the same way as above.

Hence after comparing the sum of all three terms from t_i^l we have indeed proved (3.46). $\square$

Chapter 4.

Distributed Mechanism Design with Learning Guarantees

Mechanism design for fully strategic agents commonly assumes broadcast nature of communication between agents of the system. Moreover, for mechanism design, the stability of Nash equilibrium (NE) is demonstrated by showing convergence of specific pre-designed learning dynamics, rather than for a class of learning dynamics. In this chapter we consider two common resource allocation problems: sharing K infinitely divisible resources among strategic agents for their private consumption (private goods), and determining the level for an infinitely divisible public good with P features, that is shared between strategic agents. For both cases, we present a distributed mechanism for a set of agents who communicate through a given network. In a distributed mechanism, agents' messages are not broadcast to all other agents as in the standard mechanism design framework, but are exchanged only in the local neighborhood of each agent. The presented mechanisms produce a unique NE and fully implement the social welfare maximizing allocation. In addition, the mechanisms are budget-balanced at NE. It is also shown that the mechanisms induce a game with contractive best-response, leading to guaranteed convergence for all learning dynamics within the Adaptive Best-Response dynamics class, including dynamics such as Cournot best-response, $k-$period best-response and Fictitious Play. We also present a distributed mechanism that does not possess any learning guarantees where the average dimension of the message space scales sub-linearly w.r.t. the number of agents. Finally, we present a numerically study of convergence under repeated play, for various communication graphs and learning dynamics.

1. Introduction

The framework of Mechanism design aims to bridge the informational gap between a designer, who wishes to achieve "efficient" allocation and agents, who are strategic and posses private information relevant to determining the efficient allocation. The basic premise of mechanism design has been applied to a variety of applications and as a result, recent works have focused on designing mechanisms that are also practically implementable.

In this vein, there are two very important features of practical mechanisms that haven't been addressed in the literature. (a) The first is the informational constraint between agents - most mechanisms define the contract (e.g., allocation and tax) such that messages from all agents need to be collected centrally in order to determine each agents' outcome. This is akin to assuming a broadcast structure of communication between agents. (b) The second is related to the fact that, when Nash equilibrium (NE) is used as the solution concept, there is little to no guarantee that agents can learn about each others' private information in order to calculate the NE. More specifically, dynamic learning guarantees on convergence to NE induced by the mechanism are typically provided for specific learning dynamics, which can vary from model to model. Addressing the two features together can make a mechanism ready for application to scenarios where agents are distributed - both physically and informationally.

The motivation for feature (a), i.e., designing a distributed mechanism comes from the literature on "distributed optimization", [RaNo04; NeOz09; BoPaChPeEc11; DuAgWa12], and the motivation for feature (b) comes from the literature on "learning in games", [MiRo90; MoSh96; Fu98; HoSa02; Yo04].

Regarding problem (a), the literature on distributed optimization aims to address the informational constraints between **non-strategic** agents who possess local private information relevant to the centralized optimization. This model, in-part, is analogous to mechanism design, where in many cases efficient allocation is indeed described by a centralized optimization. Thus a natural question to ask is whether efficient mechanisms can be designed such that they obey the informational constraints of distributed agents. The fact that mechanisms have to deal with strategic agents, means that this is a new and non-trivial question.

It is important to note here that the centralized allocation problem (see (4.4) and (4.6)) itself is completely oblivious to the informational constraints that the mechanism designer faces. For instance, resolving the issue (a) for the public goods problem (Section 4) is not to be confused with the problem of provisioning **local** public goods [BeCo03; ShTe12]. In such models, each agents' (intrinsic) utility is assumed to depend only on his/her neighbors allocation thereby naturally aligning the informational constraints with the utility structure. Utilities for the public goods problem do not require any such assumption. Instead, we require that the exchanged messages implied by the mechanism satisfy the network communication constraints.

Regarding problem (b), learning in games is motivated by the fact that NE is, theoretically, a complete information solution concept. Use of NE in models that don't necessarily assume perfect information among agents is typically justified[1] by showing that under certain natural learning strategies, such as evolutionary dynamics [BöSa97], or under specifically designed learning strategies, such as regret-minimizing algorithms[2], agents are guaranteed to learn each others' private information and in turn the NE. The larger the class of learning dynamics for which convergence can be guaranteed, the stronger the implied justification.

In this paper, our objective as the designer is to implement the social welfare maximizing allocation for the private goods (Walrasian allocation with money) and the public goods (Lindahl allocation) problems in the presence of strategic agents. To achieve this we design a mechanism which induces a unique and efficient NE that incorporates both the aforementioned features. Thus, the mechanism determines outcomes for each agent using only his/her neighbors' messages and doesn't require collecting all agents' messages centrally. Also, it is shown that for a sufficiently large class of learning dynamics - *adaptive best-response* dynamics (ABR) - there is guaranteed convergence to NE. The ABR class contains several well-known learning dynamics such as "Cournot best-response" and "fictitious play", [Br51a], among others. Also, the induced game is supermodular with non-compact action space, this can possibly lead to convergence of the *adaptive dynamics* (AD) class, [MiRo90], of learning strategies as well (see Section 5 for discussion).

Related works in the literature that focus on learning in games are as follows. In their seminal work [MiRo90], Milgrom and Roberts show that for a *supermodular* game with compact

[1] Learning of NE can also be justified with the *Evolutive* interpretation of NE, [Na50; OsRu94], which implies that NE in a single-shot game can be thought of as the stationary point of a dynamic adjustment process.

[2] For instance, for the case of zero-sum games, Daskalakis, Deckelbaum and Kim in [DaDeKi15] propose a new learning algorithm based on regret minimization and show that it converges at a linear rate.

action space, any learning dynamic within the AD class is eventually bounded between the two most extreme Nash equilibria. Chen in [Ch02] then presents a mechanism with non-compact action space (without any informational constraints between agents) for the Lindahl allocation problem such that the induced game is supermodular and has a unique NE. Following this development, Healy and Mathevet in [HeMa12] show that for a *contractive* game, all learning dynamics within the ABR dynamics class converge to the unique NE. The authors also present a mechanism for the Lindahl allocation problem (again without any informational constraints between agents) such that the induced game is contractive. As contraction is a more stringent property than supermodularity, ABR class is broader than AD class.

In general, "adaptive" learning strategies are broadly a class of learning strategies where, at each time, agents respond optimally to some combination of empirical distribution arising out of the past observed actions. For example, fictitious play is an adaptive learning strategy. Hofbauer and Sandholm in [HoSa02] provide convergence results specifically for fictitious play in the case of zero-sum games, supermodular games and potential games. Monderer and Shapley in [MoSh96] provide convergence results for fictitious play in identical interests games, i.e., where best-response is equivalent to that of a game where agents have the same utility functions.

The structure of this chapter is as follows: Section 2 describes the two centralized allocation problems and their optimality conditions. Section 3 defines some mechanism design basics and then presents the mechanism for the private goods problem. Before presenting the mechanism that resolves both issues, (a) and (b), we present a mechanism that resolves only issue (a) i.e., distributed communication. Section 4 presents the mechanism for the public goods problem. Section 5 introduces learning-related properties and contains the result of guaranteed convergence of any learning dynamic within the ABR class. Finally, Section 5 also contains a numerical study of the convergence pattern of various learning dynamics for different underlying communication graphs.

2. Model

There are N strategic agents, denoted by the set $\mathcal{N} = \{1, \ldots, N\}$. A directed communication graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ is given, where the vertexes correspond to the agents and an edge from vertex i to j indicates that agent i can "listen" to agent j. It is assumed that the given graph $\mathcal{G}$ is strongly connected. In this chapter, we are interested in two different types of allocation problems: private goods and public goods, we describe each model below. In the language of Economics literature, these are also known as Walrasian and Lindahl allocation problems [Hu79; MaWhGr95].

2.1. Private goods allocation problem

There are K infinitely divisible goods, denoted by set $\mathcal{K} = \{1, \dots, K\}$, that are to be distributed among the agents. Each agent receives a utility $v_i(x_i)$ based on the profile $x_i = (x_i^1, \dots, x_i^K)$ of quantity of each good that he/she receives. Since for each agent, its utility depends only on privately consumed allocation x_i and not on other agents' allocation, this is the private goods model.

It is assumed that $v_i : \mathbb{R}^K \to \mathbb{R}$ is a continuously twice differentiable, strictly concave function that for some $\eta > 1$ satisfies, $\forall\ k \in \mathcal{K}$,

$$-\eta < H_{kk}^{-1} + \sum_{l \in \mathcal{K},\, l \neq k} \left| H_{kl}^{-1} \right| < 0, \tag{4.1a}$$

$$H_{kk}^{-1} < -\frac{1}{\eta}, \tag{4.1b}$$

where H^{-1} is the inverse of the Hessian $H = \left[(\partial^2 v_i(\cdot)) \big/ (\partial x_i^k \partial x_i^l) \right]_{k,l}$. To understand the significance of this assumption consider the case of $K = 1$, then this condition is the same as

$$v_i''(\cdot) \in \left(-\eta, -\frac{1}{\eta} \right). \tag{4.2}$$

It is already assumed that $v_i(\cdot)$ is strictly concave, the only additional imposition made by this assumption is that the second derivative of $v_i(\cdot)$ is bounded away from 0 and $-\infty$. More generally if the utility is separable, $v_i(x_i) = \sum_{k \in \mathcal{K}} v_{i,k}(x_i^k)$, then the condition in (4.1) is the same as

$$v_{i,k}''(\cdot) \in \left(-\eta, -\frac{1}{\eta} \right), \quad \forall\ k \in \mathcal{K}. \tag{4.3}$$

The above mentioned properties of the utility function are assumed to be common knowledge between agents and the designer. However, the utility function $v_i(\cdot)$ itself is known only to agent i and is not known to other agents or the designer. The designer wishes to allocate available goods such that the sum of utilities is maximized subject to availability constraints, i.e., to solve the following centralized allocation problem,

$$x^* = \underset{x \in \mathbb{R}^{KN}}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{4.4a}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{N}} x_i^k = c_k, \quad \forall\, k \in \mathcal{K}, \tag{4.4b}$$

where $c_k \in \mathbb{R}$ is the total available amount of good $k \in \mathcal{K}$. The allocation x^* is also called the efficient allocation and we assume that it is finite i.e., the optimization is well-defined. The efficient allocation is unique since the utilities are strictly concave. Further, the necessary and sufficient condition for optimality are

$$\frac{\partial v_i(x_i^*)}{\partial x_i^k} = \lambda_k^*, \quad \forall\, k \in \mathcal{K},\, \forall\, i \in \mathcal{N}, \tag{4.5a}$$

$$\sum_{i \in \mathcal{N}} x_i^{k*} = c_k, \quad \forall\, k \in \mathcal{K}, \tag{4.5b}$$

where $\left(\lambda_k^*\right)_{k \in \mathcal{K}} \in \mathbb{R}^K$ are the (unique) optimal dual variables for each constraint in (4.4b).

2.2. Public goods allocation problem

There is a single infinitely divisible public good with P features, with the set of features denoted by $\mathcal{P} = \{1, \ldots, P\}$. Each agent receives a utility $v_i(x)$ based on the quantity of the public good $x = (x^p)_{p \in \mathcal{P}} \in \mathbb{R}^P$. Since for each agent, its utility depends on the common allocation x, this is the public goods model. It is assumed that $v_i : \mathbb{R}^P \to \mathbb{R}$ is a continuously double-differentiable, strictly concave function that satisfies (4.1) with the Hessian $H = \left[(\partial^2 v_i(\cdot))/(\partial x^p \partial x^q)\right]_{p,q}$. If $P = 1$, then this condition is the same as in (4.2). As in the private goods model, the properties of the utility function are assumed to be common knowledge between agents and the designer. However, the utility function $v_i(\cdot)$ itself is known only to agent i and is not known to other agents or the designer. The designer wishes to allocate the public good such that the sum of utilities is maximized, i.e., to solve the following centralized allocation problem,

$$x^* = \underset{x \in \mathbb{R}^P}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} v_i(x). \tag{4.6}$$

The allocation x^* is also called the efficient allocation and we assume it is finite i.e., the optimization is well-defined. The efficient allocation is unique due to strictly concave utilities. The necessary and sufficient optimality conditions can be written as

$$\frac{\partial v_i(x^*)}{\partial x^p} = \mu_i^{p*}, \quad \forall\, p \in \mathcal{P},\ \forall\, i \in \mathcal{N}, \tag{4.7a}$$

$$\sum_{i \in \mathcal{N}} \mu_i^{p*} = 0, \qquad \forall\, p \in \mathcal{P}, \tag{4.7b}$$

where $\left(\mu_i^{p*}\right)_{p \in \mathcal{P}, i \in \mathcal{N}}$ are the (unique) optimal dual variables. The dual variables arise in this unconstrained optimization problem because of the standard technique rewriting the public goods problem (4.6) as a private goods problem with equality constraints,

$$\left(x_1^*, x_2^*, \ldots, x_N^*\right) = \underset{x_1, \ldots, x_N \in \mathbb{R}^P}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{4.8a}$$

$$\text{s.t.} \quad x_1 = x_2,\ x_2 = x_3, \cdots, x_N = x_1. \tag{4.8b}$$

The model above, leading up to (4.6), captures the possibility that one of the agents in $\mathcal{N}$ is a seller and thus his/her utility $v_i(x) = -c_i(x)$ is the negative of the cost of producing quantity x of the public good. In this case, a convex cost of production leads to a concave utility function. In general, for the public goods problem one can also assume a seller in the system who produces the quantity x and for whom the cost of production is a known function. In this case, the social welfare maximizing allocation contains the utility of the seller as well. For keeping exposition clear and focus on the similarities between the public and private goods problems, we do not consider the seller in our model. If needed, this can accommodated in a straightforward manner.

In the next two sections, we present a mechanism for each of the above problems.

3. A mechanism for the private goods problem

3.1. Definitions

A one-shot mechanism is defined by the triplet,

$$\Big(\mathcal{M} = \mathcal{M}_1 \times \cdots \times \mathcal{M}_N, \big(\widehat{x}_1(\cdot), \ldots, \widehat{x}_N(\cdot)\big), \big(\widehat{t}_1(\cdot), \ldots, \widehat{t}_N(\cdot)\big)\Big) \tag{4.9}$$

which consists of, for each agent $i \in \mathcal{N}$, the message space $\mathcal{M}_i$, the allocation function $\widehat{x}_i : \mathcal{M} \to \mathbb{R}^K$ and the tax function $\widehat{t}_i : \mathcal{M} \to \mathbb{R}$. Given a mechanism, a game $\mathfrak{G}_{pvt}$ is setup between the agents in $\mathcal{N}$, with action space $\mathcal{M}$ and utilities

$$u_i(m) = v_i(\widehat{x}_i(m)) - \widehat{t}_i(m). \tag{4.10}$$

For this game, $\widetilde{m} \in \mathcal{M}$ is a Nash equilibrium if

$$u_i(\widetilde{m}) \geq u_i(m_i, \widetilde{m}_{-i}), \quad \forall\, m_i \in \mathcal{M}_i,\ \forall\, i \in \mathcal{N}. \tag{4.11}$$

The mechanism is said to *fully implement* the efficient allocation if

$$\Big(\widehat{x}_1(\widetilde{m}), \ldots, \widehat{x}_N(\widetilde{m})\Big) = x^*, \quad \forall\, \widetilde{m} \in \mathcal{M}_{NE}, \tag{4.12}$$

where $\mathcal{M}_{NE} \subseteq \mathcal{M}$ is the set of all Nash equilibria of game $\mathfrak{G}_{pvt}$ and x^* is the efficient allocation from (4.4). Furthermore, the mechanism is said to be *budget balanced* at Nash equilibrium if

$$\sum_{i \in \mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0, \quad \forall\, \widetilde{m} \in \mathcal{M}_{NE}. \tag{4.13}$$

Finally, we call the mechanism *distributed* if for any agent $i \in \mathcal{N}$, the allocation function $\widehat{x}_i(\cdot)$ and tax function $\widehat{t}_i(\cdot)$ instead of depending on the entire message $m = (m_j)_{j\in\mathcal{N}}$, depend only on m_i and $\big(m_j\big)_{j\in\mathcal{N}(i)}$, i.e., agent i and his/her immediate neighbors. Here $\mathcal{N}(i)$ are all the "out"-neighbors of i i.e., there exists an edge from i to j in graph $\mathcal{G}$ iff $j \in \mathcal{N}(i)$.

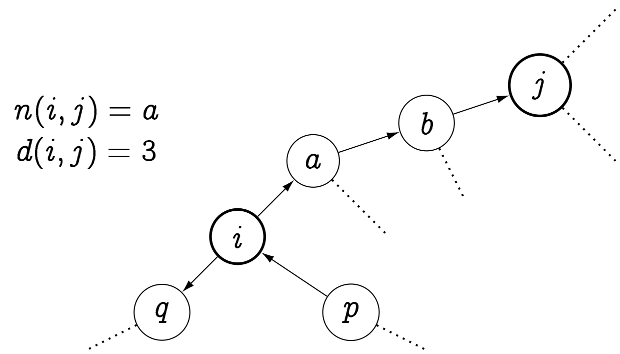

Figure 11.: $n(i, j)$ and $d(i, j)$ for the strongly connected directed graph $\mathcal{G}$.

3.2. A distributed mechanism with scalable messages

Before presenting the main mechanism below in Section 3.3, we present here a mechanism without any learning guarantees and where each agent's allocation and tax is determined only by his/her message and those of his/her neighbors. Thus this mechanism resolves issue (a) from the Introduction but not issue (b). The purpose is to set a theoretical benchmark regarding the dimensionality of the message space for designing distributed mechanisms, and compare how the additional criterion of learning imposes a burden on the dimension of the message space.

For the purpose of maintaining clarity in the exposition, we present here our mechanism for the special case of a single private good[3] i.e., $K = 1$ and with $c_1 = 0$. For the mechanism we present in this section, the average dimension of the message space per agent scales as the average degree of the graph whereas for the mechanism in the next section the average message space scales linearly with N i.e., the additional criterion of learning increases the dimension of the message space.

We consider an arbitrary communication graph that is undirected and connected. We define a spanning tree, $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, on the communication graph and the mechanism presented below works on this spanning tree. This implies that the mechanism prohibits communication along edges that were in the original communication graph but not in the spanning tree. The assumption that the communication graph is undirected is made for simplicity of exposition.

Since the spanning tree $\mathcal{G}$ is strongly connected, for any pair of vertexes $i, j \in \mathcal{N}$, the

[3]An analogous construction for the public goods problem works as well, but we do not present it.

following two quantities are well-defined. $d(i,j)$ is the length of the shortest path from i to j and $n(i,j) \in \mathcal{N}(i)$ is the neighbor of i such that the shortest path from i to j goes through $n(i,j)$. The two quantities are depicted in Fig. 11 for the more general case of a directed graph.

Additionally, for any agent $i \in \mathcal{N}$ denote by $\phi(i)$ an arbitrarily chosen neighbor and denote by I_i as the set of agents for which i was the chosen neighbor i.e., $I_i = \{k \in \mathcal{N} \mid \phi(k) = i\}$.

The message space for any agent i is $\mathcal{M}_i = \mathbb{R}^{1+|\mathcal{N}(i)|+|I_i|}$. A generic message is structured as follows: $m_i = (y_i, n_i, q_i)$, where $y_i \in \mathbb{R}$ is the demand for agent i. Here $n_i = \left(n_i^j\right)_{j\in\mathcal{N}(i)} \in \mathbb{R}^{|\mathcal{N}(i)|}$ is supposed to represent the profile of demand aggregates, where n_i^j represents the aggregate of demands over all agents k such that the shortest path from i to k goes through j, including j itself. Finally, $q_i = \left(q_i^j\right)_{j\in I_i} \in \mathbb{R}^{|I_i|}$ is supposed to represent the profile of neighbor demand proxies, where q_i^j represents the demand of neighbor j.

The average size (per agent) of the $n-$component of the message is the average degree of the graph $\overline{deg}$ and the average size of the $q-$component is 1 (since for each agent, the mapping ϕ choses exactly one neighbor). This implies that the average dimension of the message space is $2 + \overline{deg}$.

The allocation and tax function are

$$\widehat{x}_i(m) = y_i - \frac{1}{N-1} \sum_{j\in\mathcal{N}(i)} \left(y_j + \sum_{\substack{k\in\mathcal{N}(j)\\ k\neq i}} n_j^k \right) \tag{4.14a}$$

$$\widehat{t}_i(m) = \widehat{p}_i(m_{-i})\widehat{x}_i(m) + \sum_{j\in\mathcal{N}(i)} \left(n_i^j - y_j - \sum_{\substack{k\in\mathcal{N}(i)\\ k\neq i}} n_j^k \right)^2 + \sum_{j\in I_i} \left(q_i^j - y_j\right)^2 \tag{4.14b}$$

$$\widehat{p}_i(m_{-i}) = q^i_{\phi(i)} + \sum_{j\in\mathcal{N}(i)} \left(y_j + \sum_{\substack{k\in\mathcal{N}(j)\\ k\neq i}} n_j^k \right). \tag{4.14c}$$

Since agents are connected through a given graph, they can only communicate with a restricted set of agents i.e., their neighbors. Yet, as indicated in the optimality conditions, (4.5), there needs to be two kinds of global consensus at the efficient allocation. Firstly, the total allocation to all the agents must equal the total available amount, $\sum_{i\in\mathcal{N}} x_i^* = 0$. Secondly, agents also need to agree on a common "price", λ_1^*. The essential idea behind this design is to create an allocation function that at NE is equal to $\widehat{x}_i = y_i - \frac{1}{N-1}\sum_{j\neq i} y_j$ and a tax function where at NE the price is equal to $\widehat{p}_i = \sum_j y_j$. The intuition behind this specific expression for $\widehat{x}_i(m)$ and $\widehat{p}_i(m)$ at NE has also been used in prior works [GrLe77; Hu79; HeMa12].

In both expressions, $\widehat{x}_i(\cdot), \widehat{p}_i(\cdot)$, the aim as the designer is to find an appropriate substitutes for $\sum_{j\neq i} y_j$ such that it depends only on agent i's neighbors' messages. This is where the aggregate demand proxies n_i^j's come in handy. The second term in the tax (4.14b) incentivizes agents to aggregate demands onto the variable n_i^j such that at NE it is equal to the sum total of demands y_k of all agents k for whom the shortest path in $\mathcal{G}$ from i goes through i's neighbor j.

One standard design principle for mechanisms is that if an agent partially controls their own allocation (such as here, since $\widehat{x}_i(m)$ depends on y_i) then they shouldn't be able to control the price. The surrogate variables q_i are needed so that the price $\widehat{p}_i$ can be defined without using y_i. The third term in the tax, (4.14b), is designed for incentivizing agents to achieve the duplication of y_j onto q_i^j for this purpose.

The mechanism induces the game $\mathfrak{G}_s = \left(\mathcal{N}, \mathcal{M}_1 \times \cdots \times \mathcal{M}_N, \left(\widehat{u}_1, \ldots, \widehat{u}_N\right)\right)$ where the utilities are $\widehat{u}_i(m) = v_i(\widehat{x}_i(m)) - \widehat{t}_i(m)$.

Theorem 3.1 (Full Implementation). *The game $\mathfrak{G}_s$ has a unique Nash equilibrium $\widetilde{m} \in \mathcal{M}$ and the corresponding allocation is efficient i.e., $\widehat{x}_i(\widetilde{m}) = x_i^*, \forall\, i \in \mathcal{N}$. Further, the total tax paid at Nash equilibrium is zero i.e., $\sum_{i\in\mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0$.*

Proof. Please see Appendix 7.1 at the end of this chapter. □

The proof technique used in the learning result, Theorem 5.4 in Section 5, shows the best-response as a contraction mapping. Although we don't formally present a negative result, it is fairly straightforward to check that the best-response in the game induced by the above mechanism is far from contractive and thus this design does not provide any learning guarantees. One important difference in the best-response comes about because of the use of aggregate demand proxies n_i^j's. Each summand in the second term on the RHS in (4.14b) contains three terms whereas the corresponding terms in (4.16a) for the mechanism in Section 3.3 contain only two.

3.3. A mechanism for the private goods problem

For the purpose of maintaining clarity in the exposition, we present here our mechanism for the special case of a single private good i.e., $K = 1$. At the end of this section we comment on the natural extension of the presented mechanism to the general case of K goods. Also, kindly note that in this section, the communication graph $\mathcal{G}$ is only assumed to be directed and strongly connected.

Since $\mathcal{G}$ is strongly connected, for any pair of vertexes $i, j \in \mathcal{N}$, the following two quantities are well-defined. $d(i, j)$ is the length of the shortest path from i to j and $n(i, j) \in \mathcal{N}(i)$ is the out-neighbor of i such that the shortest path from i to j goes through $n(i, j)$. The two quantities are depicted in Fig. 11.

For any agent $i \in \mathcal{N}$, the message space is $\mathcal{M}_i = \mathbb{R}^{N+1}$. The message $m_i = (y_i, q_i)$ consists of agent i's demand $y_i \in \mathbb{R}$ for the allocation of the single good and a surrogate/proxy $q_i = \left(q_i^1, \ldots, q_i^N\right) \in \mathbb{R}^N$ for the demand of all the agents (including himself/herself).

The allocation function is defined as

$$\widehat{x}_i(m) = y_i - \frac{1}{N-1} \sum_{r \in \mathcal{N}(i)} \frac{q_r^r}{\xi} - \frac{1}{N-1} \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}} + \frac{c_1}{N}, \quad \forall\, i \in \mathcal{N}, \tag{4.15}$$

where $\xi \in (0, 1)$ is an appropriately chosen *contraction* parameter and its selection is discussed in Section 5 on Learning Guarantees, proof of Theorem 5.4. The tax function is defined as

$$\begin{aligned} \widehat{t}_i(m) = \widehat{p}_i(m_{-i}) \left(\widehat{x}_i(m) - \frac{c_1}{N}\right) + \left(q_i^i - \xi y_i\right)^2 + \sum_{r \in \mathcal{N}(i)} \left(q_i^r - \xi y_r\right)^2 \qquad & (4.16\text{a}) \\ + \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \left(q_i^r - \xi q_{n(i,r)}^r\right)^2, \qquad & \end{aligned}$$

$$\widehat{p}_i(m_{-i}) = \frac{1}{\delta} \left(\frac{q_{n(i,i)}^i}{\xi} + \sum_{r \in \mathcal{N}(i)} \frac{q_r^r}{\xi} + \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}} \right), \quad \forall\, i \in \mathcal{N}, \tag{4.16b}$$

where $n(i, i) \in \mathcal{N}(i)$ is an arbitrarily chosen neighbor of i and $\delta > 0$ is an appropriately chosen parameter. Both ξ, δ are tuned simultaneously in the proof of Theorem 5.4 from Section 5.

The quantities, $n(\cdot,\cdot)$ and $d(\cdot,\cdot)$, are based on the graph $\mathcal{G}$. The only property of relevance here is that the two are related recursively i.e., $d(n(i,r),r) = d(i,r)-1$. Thus if a mechanism designer wishes to avoid calculating the shortest path (possibly due to the high complexity) then n, d can be replaced by any valid neighbor and distance mapping, respectively, as long as they are related recursively as above.

3.4. Results

Fact 3.2 (Distributed)**.** The mechanism defined in (4.15) and (4.16) is distributed.

The distributed-ness of the mechanism follows from the fact that the expressions in (4.15) and (4.16) depend only on m_i and $\left(m_r\right)_{r\in\mathcal{N}(i)}$.

Since agents are connected through a given graph, they can only communicate with a restricted set of agents i.e., their neighbors. Yet, as indicated in the optimality conditions, (4.5), there needs to be two kinds of global consensus at the efficient allocation. Firstly, the total allocation to all the agents must equal the total available amount, $\sum_{i\in\mathcal{N}} x_i^* = c_1$. Secondly, agents also need to agree on a common "price", λ_1^*. To facilitate this, the message space consists of the surrogate variables $q_i = (q_i^1,\dots,q_i^N)$ which are known locally to agent i and are expected at equilibrium to be representative of the global demand $y = (y_1,\dots,y_N)$. Specifically, the second, third and fourth terms in the tax, (4.16a), are designed for incentivizing agents to achieve the aforementioned duplication of global demand y to the local surrogate q_i.

To motivate the choice of the allocation function and the remaining part of the tax function consider the case of $\xi = 1$ and take into account the duplication, i.e., $q_r = y, \forall\, r \in \mathcal{N}$. Since $\xi = 1$, all the factors involving ξ become 1 for the expressions in (4.15) and (4.16). We design the allocation $\widehat{x}_i(m)$ as a function of y_i and $(q_r)_{r\in\mathcal{N}(i)}$ such that after taking into account the duplication it becomes $y_i - \frac{1}{N-1}\sum_{j\neq i} y_j + \frac{c_1}{N}$. This facilitates the first global consensus, $\sum_{i\in\mathcal{N}} x_i = c_1$. One standard design principle for mechanisms is that if an agent partially controls their own allocation (such as here, since $\widehat{x}_i(m)$ depends on y_i) then they shouldn't be able to control the price. This is the reason that $\widehat{p}_i(\cdot)$ doesn't depend on m_i. It is function of $(q_r)_{r\in\mathcal{N}(i)}$ and is designed such that after taking into account the duplication, the price for any agent is proportional to $\sum_j y_j$. This facilitates the second consensus - common price for all agents. The intuition behind the specific expression for $\widehat{x}_i(m)$ and $\widehat{p}_i(m)$ at NE (after taking into account the duplication) has also been used in prior works [GrLe77; Hu79; HeMa12].

Finally, we set $\xi < 1$ and adjust the allocation and tax function accordingly so that the game $\mathfrak{G}_{pvt}$ can be contractive (see Section 5).

Before proceeding to the main results, we define the best-response of any agent i,

$$\beta_i(m_{-i}) = \left(\tilde{y}_i(m_{-i}), \tilde{q}_i(m_{-i})\right) \triangleq \underset{m_i \in \mathcal{M}_i}{\operatorname{argmax}}\ u_i(m). \tag{4.17}$$

Note that $\beta_i(\cdot)$ is a set-valued function in general.

Lemma 3.3 (Concavity). *For any agent $i \in \mathcal{N}$ and $m_{-i} \in \mathcal{M}_{-i}$, the utility $u_i(m)$, defined in* (4.10), *for the game $\mathfrak{G}_{pvt}$ is strictly concave in $m_i = (y_i, q_i)$. Thus, the best-response of agent i is unique and is defined by the first order conditions.*

Proof. Please see Appendix 7.2 at the end of this chapter. □

Concavity of the induced utility in the game $\mathfrak{G}_{pvt}$ largely follows from the tax terms being quadratic. The second tax term in (4.16a) is the only source of cross derivatives across components of message m_i. We prove concavity by verifying that the Hessian matrix is negative definite.

Theorem 3.4 (Full Implementation and Budget Balance). *For the game $\mathfrak{G}_{pvt}$, there exists a unique Nash equilibrium, $\widetilde{m} \in \mathcal{M}$, and the allocation at Nash equilibrium is efficient, i.e., $\widehat{x}_i(\widetilde{m}) = x_i^*$, $\forall\ i \in \mathcal{N}$. Further, the total tax paid at Nash equilibrium $\widetilde{m}$ is zero, i.e.,*

$$\sum_{i \in \mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0. \tag{4.18}$$

Proof. Please see Appendix 7.3 at the end of this chapter. □

The proof of Theorem 3.4 can be intuitively explained as follows. Since the optimality conditions in (4.5) are sufficient, we start by showing that at any Nash equilibrium the allocation and price necessarily satisfy the optimality conditions. Thereby ensuring that if Nash equilibrium exists (unique or multiple) the corresponding allocation is efficient. Then we show existence and uniqueness by showing a one-to-one map between message at Nash equilibrium and (x^*, λ^*) arising out of the optimization in (4.4).

Generalizing to multiple goods ($K > 1$)

For the general problem we use notation on a per good basis. The message space is,

$$\mathcal{M}_i = \bigtimes_{k \in \mathcal{K}} \mathcal{M}_i^k = \bigtimes_{k \in \mathcal{K}} \mathbb{R}^{(N+1)} = \mathbb{R}^{(N+1)K}. \tag{4.19}$$

Any message $m_i = \left(m_i^k\right)_{k \in \mathcal{K}} = (y_i, q_i)$ contains separate demands and proxies for each good $k \in \mathcal{K}$. We denote it as follows

$$y_i = \left(y_i^k\right)_{k \in \mathcal{K}} \in \mathbb{R}^K, \tag{4.20a}$$

$$q_i = \left(q_i^{r,k}\right)_{r \in \mathcal{N}, k \in \mathcal{K}} \in \mathbb{R}^{NK}, \tag{4.20b}$$

where for any good $k \in \mathcal{K}$, demand and proxies $\left(y_i^k, (q_i^{r,k})_{r \in \mathcal{N}}\right) \in \mathbb{R}^{N+1}$ have the same interpretation as in the presented mechanism with only one good. The allocation function $\widehat{x}_i^k(\cdot)$, for any $k \in \mathcal{K}$, depends only on the messages pertaining to good k i.e.,

$$\widehat{x}_i^k : \bigtimes_{j \in \mathcal{N}} \mathcal{M}_j^k \to \mathbb{R}, \tag{4.21}$$

where the expression for $\widehat{x}_i^k(\cdot)$ is the same as in the presented mechanism, (4.15), replacing $y_i, q_{n(i,r)}^r, q_r^r$ with $y_i^k, q_{n(i,r)}^{r,k}, q_r^{r,k}$, respectively. The tax function too can be divided on a per-good basis as follows,

$$\widehat{t}_i(m) = \sum_{k \in \mathcal{K}} \widehat{t}_i^k(m), \quad \forall\, m \in \mathcal{M},\ \forall\, i \in \mathcal{N}. \tag{4.22}$$

Here $\widehat{t}_i^k$ has the same domain and co-domain as $\widehat{x}_i^k$ in (4.21) and the expression for it is the same as in the presented mechanism, (4.16), replacing all variables with their corresponding good k variables. Also, $\widehat{x}_i$ is replaced by $\widehat{x}_i^k$ and instead of denoting price by $\widehat{p}_i$ we denote it by $\widehat{p}_i^k$.

The generalized mechanism is distributed as well, since all allocation and tax functions still depend only on message of the agent and that of its neighbors. Owing to the design with a good-wise separation, we have the following form for utility

$$u_i(m) = v_i\left(\left(\widehat{x}_i^k(m^k)\right)_{k\in\mathcal{K}}\right) - \sum_{k\in\mathcal{K}} \widehat{t}_i^k(m^k). \tag{4.23}$$

Concavity of u_i follows from arguments similar to Lemma 3.3, since the only cross-derivatives across goods are of the form $(\partial^2 v_i)\big/(\partial x_i^k \partial x_i^l)$ and for multiple goods such terms are part of the Hessian of $v_i(\cdot)$ (which is assumed to be strictly concave). Since the optimality conditions in (4.5) are sufficient, the properties of efficiency, existence and uniqueness of Nash equilibrium follow from the first order conditions for optimality in the best-response. Finally, the Budget Balance result holds true on a per-good basis, so it also holds for the total tax which is the sum of per-good taxes.

4. A mechanism for the public goods problem

For the purpose of maintaining clarity in the exposition, we present here our mechanism for the special case of a single feature in the public good i.e., $P = 1$. At the end of this section we comment on the natural extension of the presented mechanism to the general case of P features. Also the definition of full implementation here is as before, except (4.12) which is replaced by

$$\widehat{x}_1(\widetilde{m}) = \widehat{x}_2(\widetilde{m}) = \cdots = \widehat{x}_N(\widetilde{m}) = x^*, \qquad \forall\, \widetilde{m} \in \mathcal{M}_{NE}. \tag{4.24}$$

For any agent $i \in \mathcal{N}$, the message space is $\mathcal{M}_i = \mathbb{R}^{N+1}$. The message $m_i = (y_i, q_i)$ consists of agent i's contribution $y_i \in \mathbb{R}$ to the common public good and a surrogate/proxy $q_i = \left(q_i^1, \ldots, q_i^N\right) \in \mathbb{R}^N$ for the contributions of all the agents (including himself/herself).

The allocation function is defined as

$$\widehat{x}_i(m) = \frac{1}{N}\left(y_i + \sum_{r\in\mathcal{N}(i)} \frac{q_r^r}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}}\right), \quad \forall\, i \in \mathcal{N}. \tag{4.25}$$

The tax function is

$$\widehat{t}_i(m) = \widehat{p}_i(m_{-i})\widehat{x}_i(m) + \left(q_i^i - \xi y_i\right)^2 + \sum_{r \in \mathcal{N}(i)} \left(q_i^r - \xi y_r\right)^2 + \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \left(q_i^r - \xi q_{n(i,r)}^r\right)^2, \\ + \frac{\delta}{2}\left(q_{n(i,i)}^i - \xi y_i\right)^2 \tag{4.26a}$$

$$\widehat{p}_i(m_{-i}) = \delta(N-1)\left(\frac{q_{n(i,i)}^i}{\xi} - \frac{1}{N-1}\sum_{r \in \mathcal{N}(i)} \frac{q_r^r}{\xi} - \frac{1}{N-1}\sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}}\right), \quad \forall\, i \in \mathcal{N}, \tag{4.26b}$$

where $n(i,i) \in \mathcal{N}(i)$ is an arbitrarily chosen neighbor of i and $\xi \in (0,1)$, $\delta > 0$ are appropriately chosen parameters and their selection is discussed in Section 5 on Learning Guarantees, proof of Theorem 5.5.

In the following, we denote the induced game by $\mathfrak{G}_{pub}$ and define, as in the previous section, the utility in the game by (4.10) and the best-response by (4.17), .

4.1. Results

Fact 4.1 (Distributed)**.** The mechanism defined in (4.25) and (4.26) is distributed.

Generally, for a public goods problem one expects the allocation function of the form $\hat{x} : \mathcal{M} \to \mathbb{R}$ instead of $\left(\hat{x}_i : \mathcal{M} \to \mathbb{R}\right)_{i \in \mathcal{N}}$ i.e, a single common allocation instead of N different allocation functions, one for each agent. However, owing to the informational constraints of the model, there does not exist[4] a single function $\hat{x} : \mathcal{M} \to \mathbb{R}$ such that it depends only on m_i and $(m_j)_{j \in \mathcal{N}(i)}$, for every $i \in \mathcal{N}$. Since in general there are N different neighborhoods (one for each agent), the informational constraints necessitate the use of N allocation functions $\left(\hat{x}_i\right)_{i \in \mathcal{N}}$. The multiple allocation functions are consistent with the model in (4.8). As each $\hat{x}_i$ represents the level of the same public good, at NE we must have $\hat{x}_i(\widetilde{m}) = \hat{x}_j(\widetilde{m})$ for all $i, j \in \mathcal{N}$. The fact that $\hat{x}_i(m) \neq \hat{x}_j(m)$ for all $m \in \mathcal{M}$ means that the allocation is not feasible off-equilibrium.

One interpretation for N allocation functions is as follows: at the time of signing the mech-

[4] apart from the trivial constant function, which cannot give a full implementation result.

anism contract, agents are aware of the informational constraints that the mechanism and its implied messaging must satisfy. Thus, each agent i can only use messages $(m_j)_{j\in\mathcal{N}(i)}$ to calculate their own best-response. The mechanism designer thus designs allocation function $\hat{x}_i$ for agent i such that appropriate proxies[5] can be used for messages $(m_j)_{j\notin\mathcal{N}(i),j\neq i}$ of agents not in the neighborhood of i. Effectively thus, the mechanism designer accounts for the informational constraints through the appropriate design of proxies. A second interpretation is through the simulated learning process of NE. In this each agent "acts virtually" during the learning phase i.e., each agent takes only virtual actions during the learning phase, updating his/her action from one round to the next by observing only his/her neighbors message. Once the stationary point of the learning process is reached (NE), the "real action" is taken only once. The virtual interpretation is also used in most of the distributed optimization literature [RaNo04; NeOz09; BoPaChPeEc11; DuAgWa12].

The specific design of allocation and taxes in (4.25) and (4.26) is discussed below. The optimality conditions in (4.7) require that agents have global consensus on two aspects: allocation must be the same for all and sum of prices should be equal to zero. As in the private goods mechanism, we design the second, third and fourth tax terms in (4.26a) such that agents are incentivized to duplicate global demand y onto locally available variables q_i. To motivate the allocation function and the remaining part of the tax function consider the case of $\xi = 1$ and take into account the duplication $q_r = y$, $\forall\ r \in \mathcal{N}$. In this case all the factors involving ξ in (4.25) and (4.26) are 1. The allocation function $\widehat{x}_i(m)$ depends on $y_i, (q_r)_{r\in\mathcal{N}(i)}$ and is designed such that after taking into account the duplication it is proportional to $\sum_j y_j$. This facilitates the first consensus - all agents' allocation must be the same. The price $\widehat{p}_i(m_{-i})$ is designed such that it depends only on $(q_r)_{r\in\mathcal{N}(i)}$ and after taking into account the duplication it is proportional to $y_i - \frac{1}{N-1}\sum_{j\neq i} y_j$. This facilitates the second consensus - sum of prices over all agents is zero.

With the above design principles, all the results of this section follow. We then introduce an additional fifth term in the tax, (4.26a), just for the purpose of achieving contraction in Section 5 (see proof of Theorem 5.5). Incentives provided by this term are in line with those already provided to the neighboring agent $n(i, i)$ through his/her third tax term, hence it

[5] described in the paragraph below.

doesn't interfere with the equilibrium results in this section. Only the proof of concavity changes slightly.

Finally, we set $\xi < 1$ and adjust everything in the allocation and tax function correspondingly so that the game $\mathfrak{G}_{pub}$ can be contractive (see Section 5).

For the next two results, the basic idea behind the proofs are similar to the corresponding results from the previous section.

Lemma 4.2 (Concavity). *For any agent $i \in \mathcal{N}$ and $m_{-i} \in \mathcal{M}_{-i}$, the utility $u_i(m)$ for the game $\mathfrak{G}_{pub}$ is strictly concave in $m_i = (y_i, q_i)$. Thus, the best-response of agent i is unique and is defined by the first order conditions.*

Proof. Please see Appendix 7.4 at the end of this chapter. □

Theorem 4.3 (Full Implementation and Budget Balance). *For the game $\mathfrak{G}_{pub}$, there exists a unique Nash equilibrium, $\widetilde{m} \in \mathcal{M}$, and the allocation at Nash equilibrium is efficient, i.e., $\widehat{x}(\widetilde{m}) = x^*$. Further, the total tax paid at Nash equilibrium $\widetilde{m}$ is zero, i.e.,*

$$\sum_{i \in \mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0. \tag{4.27}$$

Proof. Please see Appendix 7.5 at the end of this chapter. □

Generalizing to multiple features ($P > 1$)

With an idea similar to the extension for the private goods mechanism, we extend the presented mechanism by first increasing the message space such that for each agent we have $m_i = \left(m_i^p\right)_{p \in \mathcal{P}} = (y_i, q_i)$. The allocation in this case is $P-$dimensional and the expression for $\widehat{x}_i^p(m)$ is the same as in the presented mechanism with y_i, $\left(q_r\right)_{r \in \mathcal{N}(i)}$ replaced by y_i^p, $\left(q_r^p\right)_{r \in \mathcal{N}(i)}$. The tax function is

$$\widehat{t}_i(m) = \widehat{t}_i^p(m), \tag{4.28}$$

where the expression for $\widehat{t}_i^p$ is the same as in the presented mechanism, replacing m by m^k.

The results above follow using analogous arguments to those mentioned in the previous section regarding generalization to multiple goods, $K > 1$.

5. Learning Guarantees

In this section we provide results for guaranteed convergence for a class of learning algorithms when the mechanisms defined in Sections 3 and 4 are played repeatedly. As discussed in the Introduction, such results act as a measure of robustness of a mechanism, w.r.t. information available to agents, and thus makes the mechanism ready for practical applications.

A learning dynamic is represented by

$$(\mu_n)_{n\geq 1} \subseteq \bigtimes_{i\in\mathcal{N}} \Delta(\mathcal{M}_i), \tag{4.29}$$

where μ_n is a mixed strategy profile with product structure to be used at time n. Denote by $S(\mu_n) \subseteq \mathcal{M}$ the support of the mixed strategy profile μ_n and denote by $m_n \in S(\mu_n)$ the realized action. Authors in [HeMa12] define the Adaptive Best-Response (ABR) dynamics class by restricting the support $S(\mu_n)$ in terms of past observed actions. Define the history $H_{n',n} = (m_{n'}, m_{n'+1}, \ldots, m_{n-1})$ as the set of observed actions between n' and $n-1$. Denote by $\widetilde{m}$ the unique Nash equilibrium of the game and define $B\left(\mathcal{M}'\right)$ as the smallest closed ball centered at $\widetilde{m}$ that contains the set $\mathcal{M}' \subset \mathcal{M}$. The closed ball is defined with any valid metric d on the message space $\mathcal{M}$.

Informally, a learning dynamic is in the ABR class if any point in the support of the action (distribution) at any time n is at most as far away from the Nash equilibrium than the best-response to any action that has been observed in some finite past $\{n', \ldots, n-1\}$. Furthermore, for the above comparison, the only relevant measure is the distance to Nash equilibrium and thus the actions observed in some finite past $\{n', \ldots, n-1\}$ can be replaced by other actions that are no further from the Nash equilibrium. Formally, we have the following.

Definition 5.1 (Adaptive Best-Response Learning Class [HeMa12]). A learning dynamic is an adaptive best-response dynamic if $\forall\, n'$, $\exists\, \hat{n} > n'$, s.t. $\forall\, n \geq \hat{n}$,

$$S(\mu_n) \subseteq B\left(\beta\left(B\left(H_{n',n}\right)\right)\right), \tag{4.30}$$

where $\beta : \mathcal{M} \to \mathcal{M}$ is the best-response of the game.

The above is satisfied for instance if every agent puts belief zero over actions further from

Nash equilibrium than the ones that he/she has observed in the past. Some well-known learning dynamics are in the ABR class, following are a few examples.

Cournot best-response is defined as $S(\mu_n) = m_n = \beta(m_{n-1})$, i.e., at every time agents best-respond to the last round's action. This gives rise to a deterministic strategy at each time. More generally, $k-$period best-response is defined as the learning dynamic where at any time n, an agent i's strategy is a best-response to the mixed strategy of agents $j \neq i$ which are created using the observed empirical distribution from the actions of the previous $k-$rounds i.e., $\{m_{j,n-k}, \ldots, m_{j,n-1}\}$. In fact, the generalization of this is also in the ABR class, where at each time n, an agent i's strategy is the best-response to the mixed strategy of agents $j \neq i$ that is formed by taking any convex combination of the empirical distributions of actions observed in the previous $k-$rounds. Finally, Fictitious Play [Br51b; Fu98], which maintains empirical distribution of all the past actions (instead of k most recent ones) is also in ABR. The additional requirement for this is that the utility in the game should be strictly concave, which is true for the presented mechanisms (see Lemmas 3.3 and 4.2).

Definition 5.2 (Contractive Mechanism). Let d be any metric defined on the message space $\mathcal{M}$ such that $(\mathcal{M}, d)$ is a complete metric space. A mechanism is called contractive if for any profile of utility function $\left(v_i(\cdot)\right)_{i\in\mathcal{N}}$ that satisfy the assumptions of the model, the induced game $\mathfrak{G}_{pvt}$ or $\mathfrak{G}_{pub}$ (depending on the allocation problem) has a single-valued best-response function $\beta : \mathcal{M} \to \mathcal{M}$ that is a $d-$contraction mapping.

A function $h : \mathcal{M} \to \mathcal{M}$ is a $d-$contraction mapping if there exists $c < 1$ such that $d\left(h(x), h(y)\right) \leq cd(x, y)$ for all distinct $x, y \in \mathcal{M}$. For the results below, we use the following well-known check for contraction mapping: h is a contraction mapping (for some metric d) if the Jacobian has norm less than one, i.e., $||\nabla h|| < 1$, where any matrix norm can be considered. Specifically, we consider the *row-sum* norm.

For the game induced by a contractive mechanism, by definition, there is a unique Nash equilibrium. This is due to the Banach fixed-point theorem, which gives that the best-response iteration for the induced game converges to a unique point. As shown in the previous sections, the game $\mathfrak{G}_{pvt}$ and $\mathfrak{G}_{pub}$ already have a unqiue Nash equilibrium.

Fact 5.3 ([HeMa12, Theorem 1]). If a game is contractive, then all ABR dynamics converge to the unique Nash equilibrium.

The idea behind the above proof is to show that after a finite time, under any ABR dynamic the distance to equilibrium gets smaller (exponentially) between successively rounds due to the best-response being a contraction mapping. It is shown below that the presented mechanism for both, the private and public goods problems, is contractive and thus owing to the above result there is guaranteed convergence for all learning dynamics in the ABR class.

Contraction ensures convergence for the ABR class; this result is in the same vein as the one in the seminal work [MiRo90]. Milgrom and Roberts show that Supermodularity ensures convergence for the Adaptive Dynamics class of learning algorithms (also defined in [MiRo90]). Supermodularity requires that the best-response of any agent i is non-decreasing in the message m_j of any other agent $j \neq i$. The aim in this chapter is to get guarantees for the ABR class, however it is shown below that the game $\mathfrak{G}_{pvt}$ induced by the private goods mechanism is also supermodular and thus has guaranteed convergence for the Adaptive Dynamics class as well.

Before proceeding, kindly note that contraction is somewhat more stringent a condition than supermodularity and consequently the ABR class is broader than the adaptive dynamics class. Thus in order to have better learning guarantees, our principle aim to ensure the property of contraction.

5.1. Results

Theorem 5.4 (Contraction). *The game $\mathfrak{G}_{pvt}$ defined in Section 3.3 is contractive but not supermodular. Thus, all learning dynamics within the ABR dynamics class converge to the unique efficient Nash equilibrium.*

Proof. Please see Appendix 7.6 at the end of this chapter. □

The intuition behind the proof of Theorem 5.4 can be motivated as follows. By selecting parameter $\xi < 1$, contraction of best-response for the variables $(\tilde{q}_i^r)_{r \neq i}$ is already ensured. However, for best-response in $\tilde{y}_i, \tilde{q}_i^i$ we need tuning of ξ, δ. Due to the specific nature of the mechanism it turns out that in order to accommodate any given value of $\eta > 1$, ξ needs to be selected close to 1 and correspondingly δ needs to be selected as a function of the chosen ξ. Since the best-response $\tilde{y}_i$ and $\tilde{q}_i^i$ are decreasing in $q_{n(i,i)}^i$ the game is not supermodular.

Generalizing to $K > 1$ The proof of Theorem 5.4 relies on inverting $v_i' : \mathbb{R} \to \mathbb{R}$ and bounding it appropriately. For the general case this is the same as inverting $\nabla v_i : \mathbb{R}^K \to \mathbb{R}^K$. Strict concavity of $v_i(\cdot)$ ensures that the determinant of Hessian of v_i is never zero, hence an inverse for ∇v_i exists in the general case. Of course, to bound the derivatives of $(\nabla v_i)^{-1}$, instead of using the condition in (4.2) we use the more general condition from (4.1). Finally, the proof is completed by tuning the parameters ξ, δ in exactly the same manner as in the proof of Theorem 5.4.

Theorem 5.5 (Contraction). *The game $\mathfrak{G}_{pub}$ defined in Section 4 is both contractive and supermodular. Thus, all learning dynamics within the ABR dynamics class converge to the unique Nash equilibrium.*

Proof. Please see Appendix 7.7 at the end of this chapter. □

Initially the parameters have the following bound: $\xi \in (0, 1)$ and $\delta > 0$. In order to get contraction, a further restriction $\xi \in \left(\sqrt{(N-1)/N}, 1\right)$ needs to be imposed. This also gives that each best-response is non-decreasing in the message of every other agent and thus the game is supermodular. Finally, here too in order to accommodate any value of $\eta > 1$ the final tuning of ξ requires it to be chosen very close to 1 and consequently δ is selected as a function of the chosen ξ.

Generalizing to $P > 1$, the essential idea of inverting ∇v_i from above works here too and from there onwards the steps follow analogously to the ones in the proof above.

5.2. Numerical Analysis of convergence

For numerical analysis we consider the private goods problem with one good ($K = 1$) and the public goods problem with one feature ($P = 1$). We consider $N = 31$, $\eta = 25$ and the agents' utility function as quadratic, i.e.,

$$v_i(x_i) = \theta_i x_i^2 + \sigma_i x_i, \qquad \text{(Private)}$$

$$v_i(x) = \theta_i x^2 + \sigma_i x. \qquad \text{(Public)}$$

An example of quadratic utility function can be found in [SaMoScWo12], for the model of demand side management in smart-grids. In each case the second derivative of $v_i(\cdot)$ is $2\theta_i$

	Private goods	Public goods
Binary Tree	$\xi = 1 - (1.831 \times 10^{-4})$, $\delta = 1005.6$.	$\xi = 1 - (2.515 \times 10^{-4})$, $\delta = 0.9505$.
Erdős-Reńyi	$\xi = 1 - (7.324 \times 10^{-4})$, $\delta = 932.3$.	$\xi = 1 - (10^{-3})$, $\delta = 0.8744$.

Table 1.: (ξ, δ) parameter values for different graphs and problems.

and thus for any agent i the value for θ_i is chosen uniformly randomly in the range $(-\frac{\eta}{2}, -\frac{1}{2\eta})$. As the model doesn't impose any restriction on the first derivative, the value for σ_i is chosen uniformly randomly in the range $(10, 20)$. From the proof of Theorems 5.4 and 5.5, one can numerically calculate the value of parameters ξ, δ. For the particular instance of the random θ, σ generated to be used for the plots below (the same values for θ, σ are used for both public and private goods examples), the parameter values are listed in Table 1.

For numerical analysis, quadratic utilities provide the following simplification. Since the utility functions $v_i(\cdot)$ are quadratic, this gives that $(v_i')^{-1}(\cdot)$ is affine. Thus for any agent $i \in \mathcal{N}$, the best-response $\tilde{y}_i$ and $\tilde{q}_i$ are affine in other agents' message. This gives that for learning strategies such as $k-$period best-response or Fictitious Play, instead of maintaining the empirical distribution over other agents' past actions, every agent can simply maintain the empirical average.

Since it has been shown that the best-response is a contraction mapping, one expects that any learning strategy that best-responds to some convex combination of past actions from finitely many rounds, converges at an exponential rate. Indeed, this is exactly observed from Fig. 12 and 15, where the absolute distance $||m_n - \widetilde{m}||_2$ of the action in round n to the Nash equilibrium $\widetilde{m}$ of the game $\mathfrak{G}_{pvt}$ and $\mathfrak{G}_{pub}$ is plotted versus n, respectively. Here a straight line in the semi-log plot indicates that the logarithm of the distance is linear in n. The four cases considered in each figure are through two possible graphs and two possible learning dynamics. For the graph $\mathcal{G}$, the first case is a full binary tree (lower average degree) and second a sample a Erdős-Reńyi random graph with only one connected component, where any two edges are connected with probability $p = 0.3$ (higher average degree). The same instance of Erdős-Reńyi random graph is used for both private and public goods examples. For the learning dynamic, one case considered is when the action taken by any agent i at time

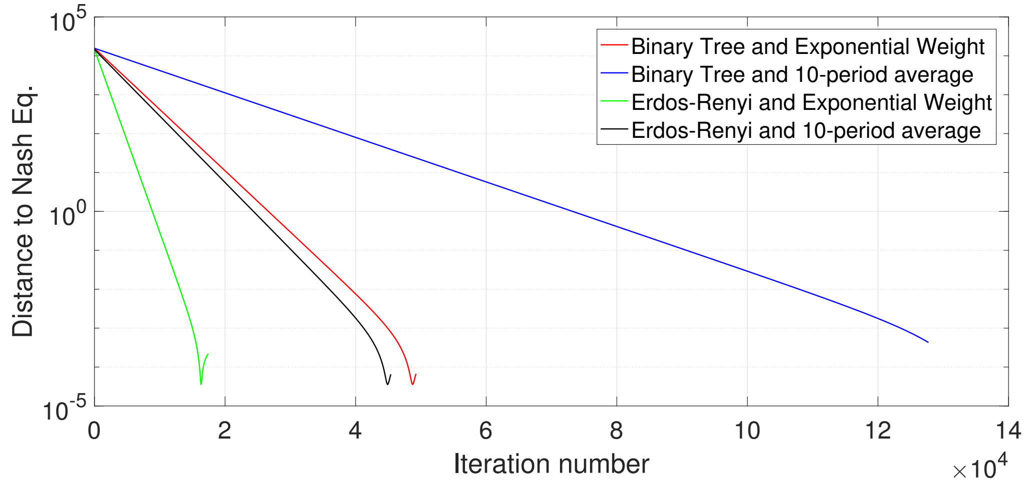

Figure 12.: $||m_n - \widetilde{m}||_2$ vs. n for the private goods mechanism.

n is the best-response to an exponentially weighed average of past actions, i.e.,

$$m_n = \beta \left(\frac{m_{n-1}}{2} + \frac{r_{n-1}}{2} \right), \tag{4.31a}$$

$$r_n = \frac{m_n}{2} + \frac{r_{n-1}}{2}. \tag{4.31b}$$

The second learning dynamic is to best-respond to the average of past 10 rounds.

Since the learning iterations are essentially conducting information exchange, one expects that the convergence to be faster for a graph that is more connected. From both the figures this is evident, as for each learning dynamic higher average degree Erdős-Reńyi random graph shows faster convergence than lower average degree Binary Tree. In fact, for both learning dynamics the convergence for the Erdős-Reńyi random graph is faster than either learning dynamic for Binary Tree. Comparing the two learning dynamics among themselves, we observe that the more aggressive exponential weighing leads to faster convergence compared to the learning dynamic that puts equal weight on each of the previous 10 actions.

Finally, for Fig. 12, in each case the relative distance to Nash equilibrium, defined as $||m_n - \widetilde{m}||_2 / ||\widetilde{m}||_1$, is in the order of 10^{-9} when the absolute distance to Nash equilibrium is 10^{-3}. For Fig. 15, the relative distance in each case is of the order of 10^{-8} when the absolute distance to Nash equilibrium is 10^{-5}.

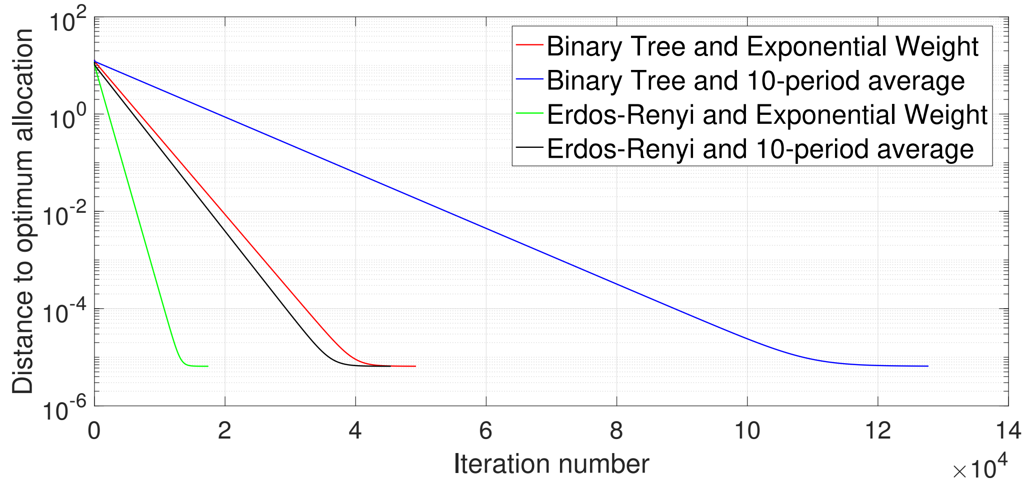

Figure 13.: $\|\widehat{x}(m_n) - x^*\|_2$ vs. n for the private goods mechanism.

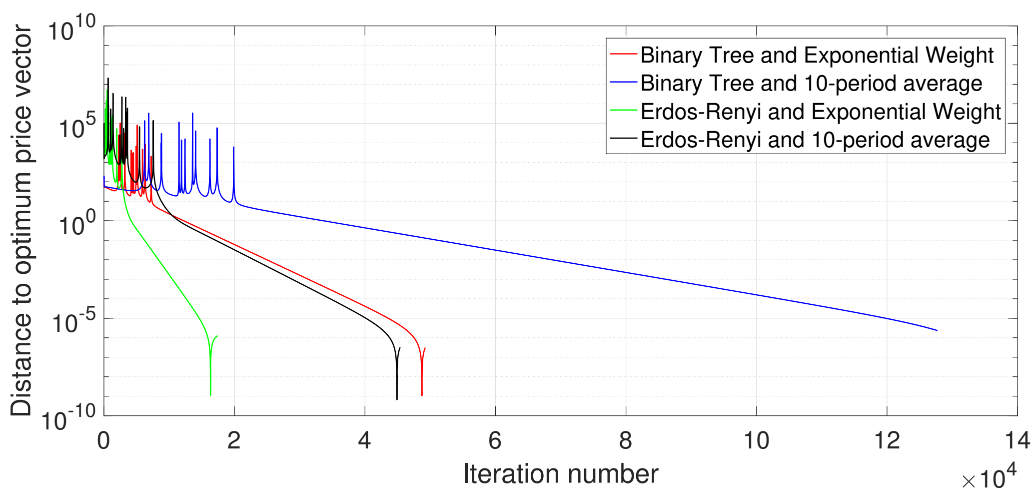

Figure 14.: $\|\left(\widehat{p}_i(m_n)\right)_{i \in \mathcal{N}} - (\lambda_1^*, \ldots, \lambda_1^*)\|_2$ vs. n for the private goods mechanism.

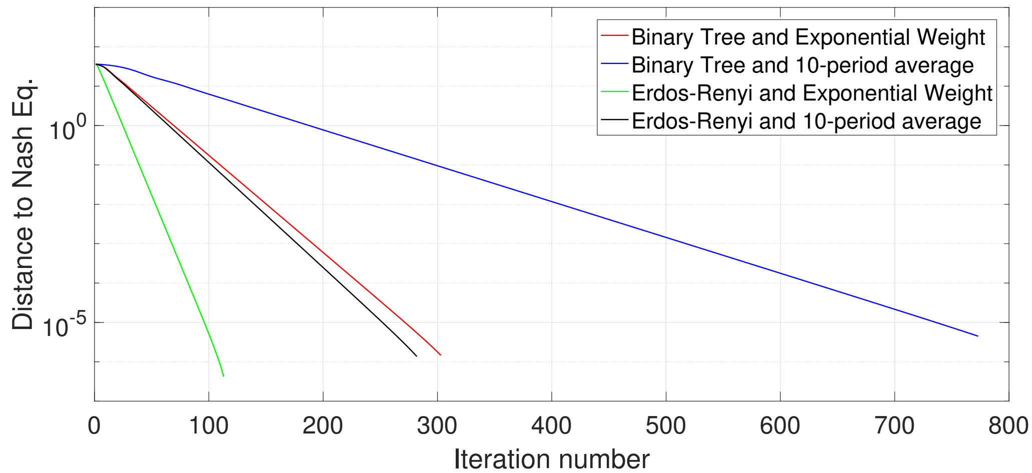

Figure 15.: $||m_n - \widetilde{m}||_2$ vs. n for the public goods mechanism.

6. Summary and Comments

In this chapter, we present a distributed mechanism where agents only need to exchange messages locally with their neighbors. While for models with non-strategic agents, extensive research has been done in the field of distributed learning and optimization, this is not the case with mechanism design where agents are fully strategic. For every profile $(v_i(\cdot))_{i\in\mathcal{N}}$ of utility functions, the induced game is shown to have a unique NE. The allocation at equilibrium is efficient and taxes are budget balanced. Then we establish informational robustness of the mechanism by showing that the best-response in the induced game is a contraction mapping. This establishes that every learning dynamic within the ABR dynamics class converges to the unique and efficient NE when the game is played repeatedly. The ABR class contains learning dynamics such as Cournot best-response, $k-$period best-response and Fictitious Play.

Future Research Directions

A significant scope for improvement in the presented mechanism is the reduction of the size of the message space. A more scalable mechanism, as in Section 3.2, would be one where on average the dimension of each agent's message space scales sub-linearly w.r.t. N. However, such an attempt might possibly require restrictions on either the underlying graph $\mathcal{G}$ or the upper bound on η that is admissible under the model. For the presented mechanism, the only restriction on the graph is that it is connected and there is no upper bound on η. Another improvement can be that of considering more complicated constraint sets for the optimization (4.4). However, as can be seen from previous attempts at mechanism design for general constraint sets, [JoTs09; JaWa10; KaTe12; SiAn17; SiAn14a], such an extension is not straightforward. Finally, for preventing inter-temporal exchange of money during the learning phase, one can adjust taxes such that there is budget balance for all messages, rather than just at NE.

Learning in summary style mechanism Our preliminary work strongly suggests that with a summary-style distributed mechanism from Section 3.2, contractive best-response cannot be guaranteed for arbitrary connected graphs. In such a case, one way to ensure learning guarantees is to restrict the set of communication graphs allowed under the model. We are currently investigating the possibility of getting similar learning guarantees as in Section 5 when the communication is restricted to a *regular* graph.

We are currently also investigating into the impossibility results for designing distributed efficient mechanisms with guaranteed convergence of the ABR class of dynamics, where the message space scales sub-linearly with N (for each agent) and where allocation and tax functions are restricted to be polynomial function of the messages.

Better learning guarantees in non-summary distributed mechanism

We are currently also investigating the possibility of improving the ambit of learning guarantees presented. The result presented in Theorems 5.4 and 5.5, guarantee convergence for a subclass of adaptive learning strategies - adaptive best-response (ABR) dynamics class. The result was proved by showing that the best-response in the induced game is a contraction mapping and then using the result from [HeMa12], where authors show that for a game with contractive best-response any ABR dynamic converges.

The class of learning strategies for which we are currently investigating convergence results are *regret-minimizing* learning algorithms[6] where at each time t, an agent plays an action

[6]such as the Multiplicative Weights algorithm.

according to a rule that ensures that the regret

$$\mathcal{R}_t = \sum_{n=1}^{t} \hat{u}_i(m_n) - \max_{m} \sum_{n=1}^{t} \hat{u}_i(m_n) \tag{4.32a}$$

$$= \sum_{n=1}^{t} \Big(v_i(\hat{x}(m_n)) - \hat{t}_i(m_n) \Big) - \max_{m} \sum_{n=1}^{t} \Big(v_i(\hat{x}(m)) - \hat{t}_i(m) \Big) \tag{4.32b}$$

grows sub-linearly with time.

It is in order to mention here that since we are interested in multiple agents, the precise definition of regret can be adapted by appropriate modification (see [Ro16, Ch. 17]). Since agents are interested in maximization of an objective that depends on their overall utility in the game $\hat{u}_i(m) = v_i(\hat{x}(m)) - \hat{t}_i(m)$, we are investigating convergence of regret-minimizing algorithms for the mechanism presented in Sections 3 and 4, which already have a contractive best-response. It is guaranteed that regret-minimizing algorithms converge to the set of *Correlated equilibria* of a game[7] and that this set contains all the Nash equilibria (see [Ro16, Ch. 17]).

[7] A subtle distinction that needs to be emphasized here is that minimizing internal regret leads to convergence to the set of correlated equilibria (CE) and minimizing the external regret leads to convergence to the set of coarse-correlated equilibria (CCE). The set of CCE contains the set of CE which in turn contains the set of NE.

7. Appendix

7.1. Proof of Theorem 3.1

Proof. The utility in the game, $u_i(m_i, m_{-i})$, is continuously twice differentiable in m_i for any m_{-i}. By evaluating the matrix of second derivatives of w.r.t. component of m_i it can be shown that the utility is strictly concave[8] in m_i for any m_{-i}. This implies that the best-response is unique and that necessarily at any NE, for every agent $i \in \mathcal{N}$ the first order derivative of $u_i(m_i, m_{-i})$ w.r.t. m_i must be zero. This gives that at any NE $m = (y_i, n_i, q_i)_{i \in \mathcal{N}}$ we have, for any $i \in \mathcal{N}$,

$$v_i'(\widehat{x}_i(m)) = \widehat{p}_i(m_{-i}), \tag{4.33a}$$

$$n_i^j = y_j + \sum_{\substack{k \in \mathcal{N}(i) \\ k \neq i}} n_j^k, \quad \forall\, j \in \mathcal{N}(i), \tag{4.33b}$$

$$q_i^j = y_j, \quad \forall\, j \in I_i. \tag{4.33c}$$

The condition in (4.33b) implies for any $i \in \mathcal{N}$,

$$n_i^j = \sum_{\substack{k \in \mathcal{N} \\ n(i,k) = j}} y_k, \quad \forall\, j \in \mathcal{N}(i). \tag{4.34}$$

For verifying this we substitute (4.34) in (4.33b)[9]. This means we have to verify

$$\sum_{\substack{k \in \mathcal{N} \\ n(i,k) = j}} y_k = y_j + \sum_{\substack{k \in \mathcal{N}(i) \\ k \neq i}} \sum_{\substack{r \in \mathcal{N} \\ n(j,r) = k}} y_r. \tag{4.35}$$

The condition $n(j, r) = k$ creates a partition of nodes and thus for no $r \in \mathcal{N}$ is there a repetition of y_r in the RHS of above. So to verify the above we can simply show that the sets A, B defined below are the same (note that $i \in \mathcal{N}$ and $j \in \mathcal{N}(i)$ are kept fixed in the

[8]The Hessian is a diagonal matrix with all negative entries.

[9]Uniqueness of solution can be verified easily in this case. Since we work with the spanning tree $\mathcal{G}$, we can start at one level above the leaf nodes. Consider j as a leaf node and i as its parent, in this case (4.33b) implies that $n_i^j = y_j$. Next consider node k that is one level above i, in this case (4.33b) implies n_k^i can be written explicitly in terms of demand y's. Uniqueness is shown by following similar construction all the way back to the root node and establishing that the system of equations (4.33b) gives n_i^j's explicitly in terms of y_j's.

argument below).

$$A = \{l \in \mathcal{N} \mid n(i,l) \neq j\}, \tag{4.36a}$$

$$B = \{k \in \mathcal{N} \mid n(j,k) = i\}. \tag{4.36b}$$

Consider any $k \in B$, this gives $n(j,k) = i$ and thus the shortest path from i to k is a sub-path of the shortest path from j to k. This means that the shortest path from i to k cannot contain j i.e., $n(i,k) \neq j$. This gives $B \subseteq A$. Consider any $l \in A$ i.e., shortest path from i to l does not contain j. Now suppose that the shortest path from j to l does not contain i. Since i and j are connected, this means that there exists a cycle containing i, j, l, which is a contradiction. Thus the shortest path from j to l does contain i i.e., $A \subseteq B$. This completes the verification of (4.34).

Substituting (4.33c) and (4.34) in (4.33a) and the expression for $\widehat{x}_i(m)$ from (4.14a) gives that at any NE we have

$$v_i'(\widehat{x}_i(m)) = \widehat{p}_i(m_{-i}) = \sum_{j \in \mathcal{N}} y_j, \quad \forall\, i \in \mathcal{N}, \tag{4.37}$$

$$\sum_{i \in \mathcal{N}} \widehat{x}_i(m) = \sum_{i \in \mathcal{N}} \left(y_i - \frac{1}{N-1} \sum_{j \neq i} y_j \right) = 0. \tag{4.38}$$

These are precisely the KKT optimality conditions of private goods optimization problem (4.4). Since the utilities $v_i(\cdot)$ are strictly concave, the KKT conditions are both necessary and sufficient for optimality. This implies that at any NE m the allocation is efficient i.e., $\widehat{x}_i(m) = x_i^*$, $\forall\, i \in \mathcal{N}$.

Existence and uniqueness of NE is verified constructively by specifying the message m at NE in terms of the unique optimal solution (including the optimal dual variable) (x^*, λ^*) to the optimization problem (4.4). The steps are exactly as in the proof of Theorem 3.4 in Appendix 7.3 later. Proof of budget balance at NE also follows the same steps as in Appendix 7.3.

$\square$

7.2. Proof of Lemma 3.3 (Concavity - Private goods)

Proof. Since the allocation and tax functions are smooth and $v_i(\cdot)$ is continuously double-differentiable, to establish concavity we show that the Hessian of $u_i(m)$ w.r.t. m_i is negative definite i.e., $H \prec 0$. Once this is established, the optimization in (4.17) has a strictly concave objective and an unbounded constraint set. Thus it has a unique maximizer, defined by the first order derivative conditions.

The Hessian is of size $(N+1) \times (N+1)$ and we have

$$H_{11} = \frac{\partial^2 u_i(m)}{\partial y_i^2} = v_i''(\widehat{x}_i(m)) - 2\xi^2, \tag{4.39a}$$

$$H_{(j+1)1} = H_{1(j+1)} = \frac{\partial^2 u_i(m)}{\partial y_i \partial q_i^j} = \begin{cases} 0 & \text{for } j \in \mathcal{N},\ j \neq i, \\ 2\xi & \text{for } j = i, \end{cases} \tag{4.39b}$$

$$H_{(j+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial (q_i^j)^2} = -2, \quad \forall\, j \in \mathcal{N}, \tag{4.39c}$$

$$H_{(r+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial q_i^r \partial q_i^j} = 0, \quad \forall\, j, r \in \mathcal{N},\ j \neq r. \tag{4.39d}$$

The characteristic equation, $\mathsf{Det}\,(H - xI) = 0$, becomes

$$(x+2)^{N-1} \left((x+2)(x - H_{11}) - 4\xi^2 \right) = 0. \tag{4.40}$$

This implies that $N-1$ eigenvalues of H are -2 and the remaining two eigenvalues satisfy $x^2 + (2 - H_{11})x - 2v_i''(\widehat{x}_i(m)) = 0$. Since H is a symmetric matrix, all its eigenvalues are real. Due to $v_i''(\cdot) < 0$, the product of roots in the above quadratic equation is positive and the sum of roots is negative. This gives that the remaining two eigenvalues of H are also negative. $\square$

7.3. Proof of Theorem 3.4 (Full Implementation - Private goods)

Proof. For the private goods problem in (4.4), the optimality conditions in (4.5) are sufficient. Thus in order to prove that the corresponding allocation at Nash equilibrium is efficient, we show that at any Nash equilibrium $\overline{m} = (\overline{y}, \overline{q}) \in \mathcal{M}$, the allocation $\left(\widehat{x}_i(\overline{m})\right)_{i \in \mathcal{N}}$ and prices $\left(\widehat{p}_i(\overline{m})\right)_{i \in \mathcal{N}}$ satisfy the optimality conditions as x^* and λ_1^*, respectively. Then using an invertibility argument we show existence and uniqueness of Nash equilibrium.

Using Lemma 3.3, at any Nash equilibrium $\overline{m}$ we have: $\nabla_{m_i} u_i(\overline{m}) = 0, \forall\, i \in \mathcal{N}$. This gives

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial y_i} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial y_i} = 0, \quad \forall\, i \in \mathcal{N}, \tag{4.41a}$$

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial q_i^r} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial q_i^r} = 0, \quad \forall\, r \in \mathcal{N},\ i \in \mathcal{N}. \tag{4.41b}$$

Using the definitions in (4.15) and (4.16), this becomes

$$v_i'(\widehat{x}_i(\overline{m})) - \widehat{p}_i(\overline{m}_{-i}) + 2\xi(\overline{q}_i^i - \xi\overline{y}_i) = 0, \quad \forall\, i \in \mathcal{N}, \tag{4.42a}$$

$$\overline{q}_i^r = \begin{cases} \xi\overline{y}_i & \text{for } r = i, \\ \xi\overline{y}_r & \text{for } r \in \mathcal{N}(i), \\ \xi\overline{q}^r_{n(i,r)} & \text{for } r \notin \mathcal{N}(i) \text{ and } r \neq i, \end{cases} \qquad \forall\, i \in \mathcal{N}. \tag{4.42b}$$

For any distinct pair of vertexes i, r, denote by $\{i, i_1, i_2, \ldots, i_{d(i,r)} = r\}$ the ordered vertexes in the shortest path between i and r, where $i_1 = n(i,r) \in \mathcal{N}(i)$. Since the shortest path between i and r contains the shortest path between i_k and r, for any $k < d(i,r)$, we have $n(i_k, r) = i_{k+1}$. Using the third sub-equation in (4.42b) repeatedly, replacing i by i_k gives,

$$\overline{q}_i^r = \xi\, \overline{q}_{i_1}^r = \xi^2\, \overline{q}_{i_2}^r = \cdots = \xi^{d(i,r)-1}\, \overline{q}^r_{i_{d(i,r)-1}}. \tag{4.43}$$

Now using the second sub-equation of (4.42b), replacing i by $i_{d(i,r)-1}$ and noting that $r \in \mathcal{N}(i_{d(i,r)-1})$, gives $\overline{q}^r_{i_{d(i,r)-1}} = \xi\, \overline{y}_r$. This combined with the above equation gives that (4.42b) implies

$$\overline{q}_i^r = \begin{cases} \xi\overline{y}_i & \text{for } r = i, \\ \xi^{d(i,r)}\overline{y}_r & \text{for } r \neq i, \end{cases} \qquad \forall\, i \in \mathcal{N}. \tag{4.44}$$

Using the above and then combining (4.42a) with (4.15) and (4.16b) gives, $\forall\, i \in \mathcal{N}$,

$$v_i'(\widehat{x}_i(\overline{m})) = \widehat{p}_i(\overline{m}_{-i}), \tag{4.45a}$$

$$\widehat{x}_i(\overline{m}) = \overline{y}_i - \frac{1}{N-1}\sum_{j \neq i} \overline{y}_j, \tag{4.45b}$$

$$\widehat{p}_i(\overline{m}_{-i}) = \frac{1}{\delta}\sum_{j \in \mathcal{N}} \overline{y}_j. \tag{4.45c}$$

(4.45b) implies $\sum_{i\in\mathcal{N}} \widehat{x}_i(\overline{m}) = 0$ and combining (4.45a) and (4.45c) gives $v_i'(\widehat{x}_i(\overline{m})) = \frac{1}{\delta}\sum_{j\in\mathcal{N}} \overline{y}_j$, $\forall\, i \in \mathcal{N}$. Thus, the allocation-price pair

$$\left(\Big(\overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j\Big)_{i\in\mathcal{N}},\ \frac{1}{\delta}\sum_{j\in\mathcal{N}}\overline{y}_j \right) \tag{4.46}$$

satisfy the optimality conditions, (4.5), as (x^*, λ_1^*). Since the optimality conditions are sufficient, the allocation at any Nash equilibrium $\overline{m}$ is the efficient allocation x^*.

For existence and uniqueness, consider the following set of linear equations that must be satisfied at any Nash equilibrium $\overline{m}$,

$$x_i^* = \overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j, \quad \forall\, i \in \mathcal{N}, \tag{4.47a}$$

$$\lambda_1^* = \frac{1}{\delta}\sum_{j\in\mathcal{N}}\overline{y}_j. \tag{4.47b}$$

Here $\left(\overline{y}_j\right)_{j\in\mathcal{N}}$ are the variables and (x^*, λ_1^*) are fixed - since they are uniquely defined by the optimization (4.4). The above equations can be inverted to give the unique solution as,

$$\overline{y}_i = \frac{N-1}{N}x_i^* + \frac{\delta\lambda_1^*}{N}, \quad \forall\, i \in \mathcal{N}. \tag{4.48}$$

Furthermore, using above and (4.44), the values for $\left(\overline{q}_i^r\right)_{i,r\in\mathcal{N}}$ can also be calculated uniquely. Since a solution for $\overline{m} = (\overline{y}, \overline{q})$ in terms of x^*, λ_1^* exists, existence of Nash equilibrium is guaranteed. Also, since this solution is unique, there is a unique Nash equilibrium.

For Budget Balance, we have the following. From the above characterization, at Nash Equilibrium $\widetilde{m}$ all tax terms from (4.16a), other than $\widehat{p}_i(\widetilde{m}_{-i})\left(\widehat{x}_i(\widetilde{m}) - \dfrac{c_1}{N}\right)$, are zero. Furthermore, the prices are all equal to λ_1^* and allocations are equal to x_i^*. Thus,

$$\sum_{i\in\mathcal{N}}\widehat{t}_i(\widetilde{m}) = \sum_{i\in\mathcal{N}}\lambda_1^*\left(\widehat{x}_i(\widetilde{m}) - \frac{c_1}{N}\right) = \lambda_1^*\left(\sum_{i\in\mathcal{N}}x_i^* - c_1\right) = \lambda_1^*\cdot 0 = 0, \tag{4.49}$$

since the efficient allocation x^* satisfies the constraint $\sum_{i\in\mathcal{N}} x_i^* = c_1$. $\square$

7.4. Proof of Lemma 4.2 (Concavity - Public Goods)

Proof. Since the allocation and tax functions are smooth and $v_i(\cdot)$ is continuously double-differentiable, to establish concavity we show that the Hessian of $u_i(m)$ w.r.t. m_i is negative definite i.e., $H \prec 0$. Once this is established, the optimization in (4.17) has a strictly concave objective and an unbounded constraint set. Thus it has a unique maximizer, defined by the first order derivative conditions.

The Hessian is of size $(N+1) \times (N+1)$ and we have

$$H_{11} = \frac{\partial^2 u_i(m)}{\partial y_i^2} = \frac{v_i''(\widehat{x}_i(m))}{N^2} - (2+\delta)\xi^2, \tag{4.50a}$$

$$H_{(j+1)1} = H_{1(j+1)} = \frac{\partial^2 u_i(m)}{\partial y_i \partial q_i^j} = \begin{cases} 0 & \text{for } j \in \mathcal{N},\ j \neq i, \\ 2\xi & \text{for } j = i, \end{cases} \tag{4.50b}$$

$$H_{(j+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial (q_i^j)^2} = -2, \quad \forall\, j \in \mathcal{N}, \tag{4.50c}$$

$$H_{(r+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial q_i^r \partial q_i^j} = 0, \quad \forall\, j, r \in \mathcal{N},\ j \neq r. \tag{4.50d}$$

The characteristic equation, $\mathsf{Det}\,(H - xI) = 0$, becomes

$$(x+2)^{N-1}\left((x+2)(x - H_{11}) - 4\xi^2\right) = 0. \tag{4.51}$$

This implies that $N-1$ eigenvalues of H are -2 and the remaining two eigenvalues satisfy $x^2 + (2 - H_{11})x + 2\delta\xi^2 - \frac{2}{N^2} v_i''(\widehat{x}_i(m)) = 0$. Since H is a symmetric matrix, all its eigenvalues are real. Due to $v_i''(\cdot) < 0$, the product of roots in the above quadratic equation is positive and the sum of roots is negative. This gives that the remaining two eigenvalues of H are also negative. □

7.5. Proof of Theorem 4.3 (Full Implementation - Public goods)

Proof. For the public goods problem in (4.6), the optimality conditions in (4.7) are sufficient. Thus in order to prove that the corresponding allocation at Nash equilibrium is efficient, we show that at any Nash equilibrium $\overline{m} = (\overline{y}, \overline{q}) \in \mathcal{M}$, the allocation $\left(\widehat{x}_i(\overline{m})\right)_{i \in \mathcal{N}}$ and prices $\left(\widehat{p}_i(\overline{m})\right)_{i \in \mathcal{N}}$ satisfy the optimality conditions as x^* and μ^{1*}, respectively. Then using an invertibility argument we show existence and uniqueness of Nash equilibrium.

Using Lemma 4.2, at any Nash equilibrium $\overline{m}$ we have: $\nabla_{m_i} u_i(\overline{m}) = 0, \forall\, i \in \mathcal{N}$. This gives

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial y_i} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial y_i} = 0, \quad \forall\, i \in \mathcal{N}, \tag{4.52a}$$

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial q_i^r} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial q_i^r} = 0, \quad \forall\, r \in \mathcal{N},\ i \in \mathcal{N}. \tag{4.52b}$$

Using the definitions in (4.25) and (4.26), this becomes, $\forall\, i \in \mathcal{N}$,

$$\frac{1}{N}\Big(v_i'(\widehat{x}_i(\overline{m})) - \widehat{p}_i(\overline{m}_{-i})\Big) + 2\xi(\overline{q}_i^i - \xi\overline{y}_i) + \delta\xi(\overline{q}_{n(i,i)}^i - \xi\overline{y}_i) = 0, \tag{4.53a}$$

$$\overline{q}_i^r = \begin{cases} \xi\overline{y}_i & \text{for } r = i, \\ \xi\overline{y}_r & \text{for } r \in \mathcal{N}(i), \\ \xi\overline{q}_{n(i,r)}^r & \text{for } r \notin \mathcal{N}(i) \text{ and } r \neq i, \end{cases} \quad . \tag{4.53b}$$

For any distinct pair of vertexes i, r, denote by $\{i, i_1, i_2, \ldots, i_{d(i,r)} = r\}$ the ordered vertexes in the shortest path between i and r, where $i_1 = n(i,r) \in \mathcal{N}(i)$. Since the shortest path between i and r contains the shortest path between i_k and r, for any $k < d(i,r)$, we have $n(i_k, r) = i_{k+1}$. Using the third sub-equation in (4.53b) repeatedly, replacing i by i_k gives,

$$\overline{q}_i^r = \xi\,\overline{q}_{i_1}^r = \xi^2\,\overline{q}_{i_2}^r = \cdots = \xi^{d(i,r)-1}\,\overline{q}_{i_{d(i,r)-1}}^r. \tag{4.54}$$

Now using the second sub-equation of (4.53b), replacing i by $i_{d(i,r)-1}$ and noting that $r \in \mathcal{N}(i_{d(i,r)-1})$, gives $\overline{q}_{i_{d(i,r)-1}}^r = \xi\,\overline{y}_r$. This combined with the above equation gives that (4.53b) implies

$$\overline{q}_i^r = \begin{cases} \xi\overline{y}_i & \text{for } r = i, \\ \xi^{d(i,r)}\overline{y}_r & \text{for } r \neq i, \end{cases} \quad \forall\, i \in \mathcal{N}. \tag{4.55}$$

Using the above and then combining (4.53a) with (4.25) and (4.26b) gives, $\forall\, i \in \mathcal{N}$,

$$v_i'(\widehat{x}_i(\overline{m})) = \widehat{p}_i(\overline{m}_{-i}), \tag{4.56a}$$

$$\widehat{x}_i(\overline{m}) = \frac{1}{N}\sum_{j \in \mathcal{N}} \overline{y}_j, \tag{4.56b}$$

$$\widehat{p}_i(\overline{m}_{-i}) = \delta(N-1)\left(\overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j\right). \tag{4.56c}$$

(4.56b) implies $\widehat{x}_i(\overline{m}) = \widehat{x}_r(\overline{m})$ for any $i, r \in \mathcal{N}$ and (4.56c) gives $\sum_{r\in\mathcal{N}} \widehat{p}_i(\overline{m}_{-i}) = 0$. Thus, the allocation-price pair

$$\left(\frac{1}{N}\sum_{j\in\mathcal{N}}\overline{y}_j\,,\ \left(\delta(N-1)\left(\overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j\right)\right)_{i\in\mathcal{N}}\right) \tag{4.57}$$

satisfy the optimality conditions, (4.7), as (x^*, μ^{1^*}). Since the optimality conditions are sufficient, the allocation at any Nash equilibrium $\overline{m}$ is the efficient allocation x^*.

For existence and uniqueness, consider the following set of linear equations that must be satisfied at any Nash equilibrium $\overline{m}$,

$$x^* = \frac{1}{N}\sum_{j\in\mathcal{N}}\overline{y}_j, \tag{4.58a}$$

$$\mu_i^{1^*} = \delta(N-1)\left(\overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j\right), \quad \forall\, i \in \mathcal{N}. \tag{4.58b}$$

Here $\left(\overline{y}_j\right)_{j\in\mathcal{N}}$ are the variables and (x^*, μ^{1^*}) are fixed - since they are uniquely defined by the optimization, (4.6). The above equations can be inverted to give the unique solution as,

$$\overline{y}_i = x^* + \frac{\mu_i^{1^*}}{\delta N}, \quad \forall\, i \in \mathcal{N}. \tag{4.59}$$

Furthermore, using above and (4.55), the values for $\left(\overline{q}_i^r\right)_{i,r\in\mathcal{N}}$ can also be calculated uniquely. Since a solution for $\overline{m} = (\overline{y}, \overline{q})$ in terms of x^*, μ^{1^*} exists, existence of Nash equilibrium is guaranteed. Also, since this solution is unique, there is a unique Nash equilibrium.

For Budget Balance, we have the following. By the characterization from above we know that at Nash Equilibrium $\widetilde{m}$, all tax terms from (4.26a), other than $\widehat{p}_i(\widetilde{m}_{-i})\widehat{x}_i(\widetilde{m})$, are zero. Furthermore, the prices are equal to $\mu_i^{1^*}$ and each allocation is equal to x^*. Thus,

$$\sum_{i\in\mathcal{N}}\widehat{t}_i(\widetilde{m}) = \sum_{i\in\mathcal{N}}\mu_i^{1^*}x^* = x^*\sum_{i\in\mathcal{N}}\mu_i^{1^*} = x^*\cdot 0 = 0, \tag{4.60}$$

since the optimal dual variables $\left(\mu_i^{1*}\right)_{i\in\mathcal{N}}$ satisfy, (4.7b), $\sum_{i\in\mathcal{N}} \mu_i^{1*} = 0$. $\square$

7.6. Proof of Theorem 5.4 (Contraction - Private goods)

Proof. The game is contractive if the matrix norm of the Jacobian of best-response $\beta = \left(\beta_i\right)_{i\in\mathcal{N}} = \left(\tilde{y}_i, \tilde{q}_i\right)_{i\in\mathcal{N}}$ is smaller than unity, i.e., $||\nabla\beta|| < 1$. We use the row-sum norm for this, and in this proof verify specifically the following set of conditions,

$$\sum_{r\in\mathcal{N},\, r\neq i} \left(\left|\frac{\partial \tilde{y}_i}{\partial y_r}\right| + \sum_{j\in\mathcal{N}} \left|\frac{\partial \tilde{y}_i}{\partial q_r^j}\right| \right) < 1, \quad \forall\, i\in\mathcal{N}, \tag{4.61a}$$

$$\sum_{r\in\mathcal{N},\, r\neq i} \left(\left|\frac{\partial \tilde{q}_i^w}{\partial y_r}\right| + \sum_{j\in\mathcal{N}} \left|\frac{\partial \tilde{q}_i^w}{\partial q_r^j}\right| \right) < 1, \quad \forall\, w\in\mathcal{N},\ \forall\, i\in\mathcal{N}. \tag{4.61b}$$

The summation can be performed simply over the indexes $r\in\mathcal{N}(i)$ instead of $r\neq i$ because our defined mechanism is distributed and hence the best-response of agent i depends only on $\left(m_j\right)_{j\in\mathcal{N}(i)}$.

Consider any agent $i\in\mathcal{N}$, for the best-response $\tilde{q}_i$ we have

$$\tilde{q}_i^w = \begin{cases} \xi\tilde{y}_i & \text{for } w = i, \\ \xi y_w & \text{for } w\in\mathcal{N}(i), \\ \xi q_{n(i,w)}^w & \text{for } w\notin\mathcal{N}(i) \text{ and } w\neq i. \end{cases} \tag{4.62}$$

Thus, by choosing $\xi\in(0,1)$, all conditions within (4.61b) are satisfied where $w\neq i$. Next, we verify conditions in (4.61a). Once this is done, then in conjunction with $\xi\in(0,1)$, the conditions from (4.61b) with $w = i$ are also automatically verified.

For the best-response $\tilde{y}_i$, we have

$$\tilde{y}_i = \frac{1}{N-1}\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}} + \frac{1}{N-1}\sum_{r\in\mathcal{N}(i)} \frac{q_r^r}{\xi} + \left(v_i'\right)^{-1}\left(\widehat{p}_i(m_{-i})\right). \tag{4.63}$$

where $\widehat{p}_i(m_{-i})$ is defined in (4.16b). Thus, we have,

$$\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} = \begin{cases} \dfrac{1}{\delta\xi}\dfrac{1}{v_i''(\cdot)} & \text{for } r = i, \\ \dfrac{1}{N-1}\dfrac{1}{\xi} + \dfrac{1}{\delta\xi}\dfrac{1}{v_i''(\cdot)} & \text{for } r \in \mathcal{N}(i), \\ \dfrac{1}{N-1}\dfrac{1}{\xi^{d(i,r)-1}} + \dfrac{1}{\delta\xi^{d(i,r)-1}}\dfrac{1}{v_i''(\cdot)} & \text{for } r \notin \mathcal{N}(i),\ r \neq i. \end{cases} \tag{4.64}$$

where in each expression above $v_i''(\cdot)$ is evaluated at $\widehat{p}_i(m_{-i})$. Also, in the notation used above, for any $r \in \mathcal{N}(i)$, $n(i,r) = r$. All other partial derivative of $\tilde{y}_i$ are zero. With all this condition in (4.61a) becomes,

$$\left|\frac{1}{\delta\xi}\frac{1}{v_i''(\cdot)}\right| + \left|1 + \frac{N-1}{\delta v_i''(\cdot)}\right| \left(\sum_{r\in\mathcal{N}(i)} \frac{1}{(N-1)\xi}\right) + \left|1 + \frac{N-1}{\delta v_i''(\cdot)}\right| \left(\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{(N-1)\xi^{d(i,r)-1}}\right) < 1. \tag{4.65}$$

To simplify the above, we utilize the upper bound from (4.2), $v_i''(\cdot) \in (-\eta, -\frac{1}{\eta})$. Set

$$\eta < \frac{\delta}{N-1}, \tag{4.66}$$

so that the expressions inside absolute value operator for the second and third terms on the LHS in (4.65) are guaranteed to be positive. With this, (4.65) becomes

$$\begin{aligned} &\left[\frac{-1}{v_i''(\cdot)}\right] \left(\frac{1}{\xi} - \sum_{r\in\mathcal{N}(i)} \frac{1}{\xi} - \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{\xi^{d(i,r)-1}}\right) \\ &\qquad < \frac{\delta}{N-1}\left(N - 1 - \left[\sum_{r\in\mathcal{N}(i)} \frac{1}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{\xi^{d(i,r)-1}}\right]\right) \end{aligned} \tag{4.67}$$

Since any agent has at least one neighbor i.e., $|\mathcal{N}(i)| \geq 1$, the LHS above is negative. For the RHS, note that the expression inside the square brackets has exactly $N-1$ terms and each term is of the form ξ^{-k}, for some $k \geq 1$. Since $\xi < 1$, this gives that even the RHS is negative. Utilizing the lower bound from (4.2), a sufficient condition to verify (4.67) is

$\eta < \frac{N-1}{\delta}\left|\frac{C_i}{D_i}\right|$, where C_i, D_i are the expression inside the curved bracket on the LHS and RHS of (4.67), respectively. Combining this with the condition in (4.66), a sufficient condition for verifying (4.61a) is

$$\eta < \min\left(\frac{\delta}{N-1}, \frac{N-1}{\delta}\left|\frac{C_i}{D_i}\right|\right), \quad \forall\, i \in \mathcal{N}. \tag{4.68}$$

Without any further tuning of parameters ξ, δ, the proof is complete as long as η satisfies above. However, in our model we would like to accommodate any value of $\eta > 1$ and this requires tuning of parameters ξ, δ. Set

$$\delta = (N-1)\sqrt{\min_{i\in\mathcal{N}}\left|\frac{C_i}{D_i}\right|} > 0, \tag{4.69}$$

in this to get the sufficient condition as $\eta^2 < \min_{i\in\mathcal{N}}\left|\frac{C_i}{D_i}\right|$, i.e.,

$$\eta^2 < \min_{i\in\mathcal{N}}\left\{\left|\frac{1}{\xi} - \sum_{r\in\mathcal{N}(i)}\frac{1}{\xi} - \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right| \Bigg/ \left|N-1-\left[\sum_{r\in\mathcal{N}(i)}\frac{1}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right]\right|\right\}. \tag{4.70}$$

We want to tune $\xi \in (0,1)$ such that the RHS above can be made arbitrarily large. For this, first note that, for any $i \in \mathcal{N}$ the numerator of the RHS is bounded away from zero for ξ in the neighborhood of 1. Second, the denominator can be made arbitrarily close to 0 by choosing ξ close enough to 1,

$$N-1-\left[\sum_{r\in\mathcal{N}(i)}\frac{1}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right] = \sum_{r\in\mathcal{N}(i)}\left(1-\frac{1}{\xi}\right) + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\left(1-\frac{1}{\xi^{d(i,r)-1}}\right), \tag{4.71}$$

where for any given $k \geq 1, \epsilon > 0$, choose $\xi \in \left(\left(\frac{1}{1+\epsilon}\right)^{1/k}, 1\right)$ to have $|1-\xi^{-k}| < \epsilon$. Finally, it is clear from above that the denominator can be made arbitrarily small concurrently for all $i \in \mathcal{N}$.

That the game is not supermodular follows from the first sub-equation of (4.64), where the derivate of best-response $\tilde{y}_i$ w.r.t. $q^i_{n(i,i)}$ is clearly negative. Also, convergence of every learning dynamic within the ABR class is guaranteed by Fact 5.3. □

7.7. Proof of Theorem 5.5 (Contraction - Public goods)

Proof. The game is contractive if the matrix norm of the Jacobian of best-response $\beta = \big(\beta_i\big)_{i\in\mathcal{N}} = \big(\tilde{y}_i, \tilde{q}_i\big)_{i\in\mathcal{N}}$ is smaller than unity, i.e., $||\nabla\beta|| < 1$. We use the row-sum norm for this, and in this proof verify specifically the following set of conditions,

$$\sum_{r\in\mathcal{N},\, r\neq i}\left(\left|\frac{\partial \tilde{y}_i}{\partial y_r}\right| + \sum_{j\in\mathcal{N}}\left|\frac{\partial \tilde{y}_i}{\partial q_r^j}\right|\right) < 1, \quad \forall\, i \in \mathcal{N}, \tag{4.72a}$$

$$\sum_{r\in\mathcal{N},\, r\neq i}\left(\left|\frac{\partial \tilde{q}_i^w}{\partial y_r}\right| + \sum_{j\in\mathcal{N}}\left|\frac{\partial \tilde{q}_i^w}{\partial q_r^j}\right|\right) < 1, \quad \forall\, w \in \mathcal{N},\, \forall\, i \in \mathcal{N}. \tag{4.72b}$$

The summation can be performed simply over the indexes $r \in \mathcal{N}(i)$ instead of $r \neq i$ because our defined mechanism is distributed and hence the best-response of agent i depends only on $\big(m_j\big)_{j\in\mathcal{N}(i)}$.

Consider any agent $i \in \mathcal{N}$, for the best-response $\tilde{q}_i$ we have

$$\tilde{q}_i^w = \begin{cases} \xi\tilde{y}_i & \text{for } w = i, \\ \xi y_w & \text{for } w \in \mathcal{N}(i), \\ \xi q^w_{n(i,w)} & \text{for } w \notin \mathcal{N}(i) \text{ and } w \neq i. \end{cases} \tag{4.73}$$

Thus, by choosing $\xi \in (0, 1)$, all conditions within (4.72b) are satisfied where $w \neq i$. Next, we verify conditions in (4.72a). Once this is done, then in conjunction with $\xi \in (0, 1)$, the conditions from (4.72b) with $w = i$ are also automatically verified.

For the best-response $\tilde{y}_i$, we have the following relation

$$\frac{1}{N}\Big(v_i'(\widehat{x}_i(m)) - \widehat{p}_i(m_{-i})\Big) + 2\xi(\tilde{q}_i^i - \xi\tilde{y}_i) + \delta\xi(q^i_{n(i,i)} - \xi\tilde{y}_i) = 0, \tag{4.74a}$$

$$\Rightarrow \quad \frac{1}{N}\Big(v_i'(\widehat{x}_i(m)) - \widehat{p}_i(m_{-i})\Big) + \delta\xi(q^i_{n(i,i)} - \xi\tilde{y}_i) = 0. \tag{4.74b}$$

In the above relation, $\widehat{x}_i(m)$ is evaluated at $\tilde{y}_i$ instead of y_i. Also, this relation implicitly defines $\tilde{y}_i$. Differentiating this equation w.r.t. $\left(q^r_{n(i,r)}\right)_{r\in\mathcal{N}}$ gives

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,i)}} - \frac{\delta(N-1)}{N\xi} + \delta\xi\left(1-\xi\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,i)}}\right) = 0, \qquad r = i, \tag{4.75a}$$

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\left(\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} + \frac{1}{\xi}\right) + \frac{\delta}{N\xi} - \delta\xi^2\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} = 0, \quad \forall\, r \in \mathcal{N}(i), \tag{4.75b}$$

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\left(\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} + \frac{1}{\xi^{d(i,r)-1}}\right) + \frac{\delta}{N\xi^{d(i,r)-1}} - \delta\xi^2\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} = 0, \quad \forall\, r \notin \mathcal{N}(i),\ r \neq i, \tag{4.75c}$$

which implies

$$\frac{\partial \tilde{y}_i}{\partial q^r_{n(i,r)}} = \frac{1}{\frac{v_i''(\widehat{x}_i(m))}{N^2} - \delta\xi^2} \times \begin{cases} \frac{\delta(N-1)}{N\xi} - \delta\xi & \text{for } r = i, \\ -\frac{\delta}{N\xi} - \frac{v_i''(\widehat{x}_i(m))}{N^2\xi} & \text{for } r \in \mathcal{N}(i), \\ -\frac{\delta}{N\xi^{d(i,r)-1}} - \frac{v_i''(\widehat{x}_i(m))}{N^2\xi^{d(i,r)-1}} & \text{for } r \notin \mathcal{N}(i),\ r \neq i. \end{cases} \tag{4.76}$$

In the notation used above, for any $r \in \mathcal{N}(i)$, $n(i,r) = r$. All other partial derivative of $\tilde{y}_i$ are zero. With all this condition in (4.72a) becomes,

$$\begin{aligned} &\left|\frac{\delta(N-1)}{N\xi} - \delta\xi\right| + \left|-\frac{\delta}{N} - \frac{v_i''(\widehat{x}_i(m))}{N^2}\right|\left(\sum_{r\in\mathcal{N}(i)}\frac{1}{\xi}\right) \\ &\qquad + \left|-\frac{\delta}{N} - \frac{v_i''(\widehat{x}_i(m))}{N^2}\right|\left(\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right) < \delta\xi^2 - \frac{v_i''(\widehat{x}_i(m))}{N^2}. \end{aligned} \tag{4.77}$$

We impose the condition

$$\xi \in \left(\sqrt{\frac{N-1}{N}}, 1\right) \tag{4.78}$$

so that the expression inside the first absolute value term in above is negative. To simplify the other expressions containing absolute value, we utilize the lower bound from (4.2), $v_i''(\cdot) \in$

$(-\eta, -\frac{1}{\eta})$. Set

$$\eta < N\delta, \tag{4.79}$$

so that the remaining expressions inside absolute value in (4.77) are guaranteed to be negative. With this, (4.77) becomes

$$\begin{aligned}\Big[-v_i''(\widehat{x}_i(m))\Big]&\left(1+\sum_{r\in\mathcal{N}(i)}\frac{1}{\xi}+\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right)\\ &> N\delta\left(\frac{1}{\xi}\left[\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-2}}-(N-|\mathcal{N}(i)|-1)\right]+N\xi(1-\xi)\right),\end{aligned} \tag{4.80}$$

where $N-|\mathcal{N}(i)|-1$ is the number of agents in the system except agent i and all his/her neighbors in $\mathcal{N}(i)$. Clearly the LHS above is positive. For any $r\in\mathcal{N}(i)$ and $r\neq i$, we have $d(i,r)\geq 2$. On the RHS, inside the square brackets there are exactly $N-|\mathcal{N}(i)|-1$ terms in the summation and each term is of the form ξ^{-k}, for some $k\geq 0$. Since $\xi<1$, this gives that even the RHS is positive. Utilizing the upper bound from (4.2), a sufficient condition to verify (4.80) is $\eta<\frac{1}{N\delta}\frac{C_i}{D_i}$, where C_i, D_i are the expression inside the curved bracket on the LHS and RHS of (4.80), respectively. Combining this with the condition in (4.79), a sufficient condition for verifying (4.72a) is

$$\eta<\min\left(N\delta\,,\,\frac{1}{N\delta}\frac{C_i}{D_i}\right),\quad \forall\, i\in\mathcal{N}. \tag{4.81}$$

Without any further tuning of parameters ξ,δ, the proof is complete as long as η satisfies above. However, in our model we would like to accommodate any value of $\eta>1$ and this requires tuning of parameters ξ,δ. Set

$$\delta=\frac{1}{N}\sqrt{\min_{i\in\mathcal{N}}\left(\frac{C_i}{D_i}\right)}>0, \tag{4.82}$$

in this to get the sufficient condition as $\eta^2<\min_{i\in\mathcal{N}}\left(\frac{C_i}{D_i}\right)$, i.e.,

$$\eta^2 < \min_{i\in\mathcal{N}}\left\{ \left(1+\sum_{r\in\mathcal{N}(i)}\frac{1}{\xi}+\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right) \Big/ \left(\frac{1}{\xi}\left[\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-2}}-(N-|\mathcal{N}(i)|-1)\right]+N\xi(1-\xi)\right)\right\}. \quad (4.83)$$

We want to tune ξ such that the RHS above can be made arbitrarily large, whilst satisfying (4.78). For this, firstly note that, for any $i \in \mathcal{N}$ the numerator of the RHS is greater than 1, hence it is bounded away from zero. Secondly, the denominator can be made arbitrarily close to 0 by choosing ξ close enough to 1,

$$\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-2}}-(N-|\mathcal{N}(i)|-1)=\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\left(\frac{1}{\xi^{d(i,r)-2}}-1\right)=\sum_{\substack{r\in\mathcal{N}\\ d(i,r)\geq 3}}\left(\frac{1}{\xi^{d(i,r)-2}}-1\right), \quad (4.84)$$

where for any given $k \geq 1, \epsilon > 0$, choose $\xi \in \left(\left(\frac{1}{1+\epsilon}\right)^{1/k}, 1\right)$ to have $\left(\xi^{-k}-1\right) < \epsilon$. Note that this is consistent with (4.78). The remaining term $N\xi(1-\xi)$ can also made made arbitrarily small by choosing ξ close enough to 1. Finally, it is clear from above that the denominator can made arbitrarily small concurrently for all $i \in \mathcal{N}$.

The fact that the game is supermodular follows from the preceding analysis, where the parameters are chosen such that each expression in (4.76) is positive. □

Chapter 5.

Mechanisms Design for Fairness

Mechanism design for incentivizing strategic agents to maximize their sum of utilities (SoU) is a well-studied problem in the context of resource allocation in networks. There are, however, a number of network resource allocation problems of interest where a designer may have a different objective than maximization of the SoU. The obvious reason for seeking a different objective is that this notion of efficiency does not account for fairness of allocation. A second, more subtle, reason for demanding fairer allocation is that it indirectly implies less variation in taxes paid by agents. This is desirable in a situation where implicit individual agent budgetary constraints make payment of large taxes unrealistic.

In this chapter we study a family of social utilities that provide fair allocation (with SoU being subsumed as an extreme case) and derive conditions under which Bayesian and dominant strategy implementation is possible. Furthermore, it is shown how a modification of the above mechanism by adding just one message per agent can guarantee full Bayesian implementation i.e., no extraneous equilibria. We consider the problem of demand-side management in smart grids as a specific motivating application and through numerical analysis it is demonstrated that in this application the proposed method can result in significant gains in fairness of allocation and a reduction in tax variation among agents.

1. Introduction

Mechanism design or incentive design is a well-established framework for dealing with decentralized resource allocation problems in the presence of strategic agents. The corresponding literature is vast, especially in the domain of dominant, Nash and Bayesian implementation, [HuRe06; MaSj02; Ja01; Kr09; AlBoHoPo13a].

There are several engineering applications where mechanism design has been studied for resource allocation in the presence of strategic agents. Provision of bandwidth for networks has been studied for different networks, such as unicast and multicast [MaBa06; YaHa07; JaWa10; SiAn14a; SiAn16]. Power allocation in wireless networks has been considered in [FuSc07], where the VCG mechanism was adapted for competition between stations, and in [ShTe12] where a mechanism that achieves Lindahl equilibrium allocation was presented. A similar problem of throughput maximization via mechanism design was solved in [SoZhFa08]. Truthful mechanisms for efficient spectrum allocation to selfish secondary users in a dynamic setting has been studied in [WaWuJiLiCl08; Wa10], while the authors in [HuBeHo06b] proposed a mechanism that achieved $\max - \min$ fair allocation. Demand-side management in smart grids using incentives has been studied in [FaAl00; CaKe10; SaMoScWo12]. Finally, an online auction for pricing of electric vehicle charging has been studied in [Ge11], [BhKaChGu14].

In the area of Networks, a majority of the works start with the assumption that the social objective is the sum of individual utilities (SoU) of all the system's agents. There are, however, a number of problems of theoretical and practical interest where a designer may have a different objective than maximization of the SoU. Consider for example, a network where optimizing SoU results in one (or a few) agents receiving almost all the available resources (e.g. bandwidth) and everyone else receiving an appreciably lower portion. In such a case, the social planner may want to introduce fairness in the allocation process even if it means a slight reduction in revenue. In such a market where efficiency is defined by allocating through maximization of SoU, an agent can receive disproportionate allocation as long as he/she is able to pay for it. Even if agents have this capability, a designer may desire to move away from this notion of efficiency.

A second–more subtle–reason for a departure from the SoU social objective is the fact that the standard mechanism design framework does not provide any formal way of limiting the range of the taxes/subsidies required at equilibrium. This implies that the magnitude of the monetary transfers (taxes/subsidies) can vary greatly among agents (with those who benefit more from the allocation having to contribute more as well). This can be a significant practical problem since it does not take into account the budget constraints of individual agents.

The prevalence of SoU over fairness considerations in most of the mechanism design literature has strong mathematical reasons. The main reason is that for SoU, in conjunction with quasi-linear utilities, agents' individual goals can be aligned directly with the overall social objective, so that when agents maximize their net utility they are simultaneously maximizing the overall social objective. The VCG mechanism [Gr73] is the most prominent example of this. For any agent i, his/her utility $v(x;\theta_i)$ is determined by the allocation x and their private type θ_i. The SoU social objective is $\sum_{i=1}^{N} v(x;\theta_i)$, which upon maximization determines the optimal allocation $\hat{x}(\theta)$. For quoted message profile $\phi = (\phi_j)_{j\in\mathcal{N}}$, the utility of any agent i is $v(\hat{x}(\phi);\theta_i) - t_i(\phi)$, where VCG taxes $t_i(\phi) = -\sum_{j\neq i} v(\hat{x}(\phi);\phi_j) + f_i(\phi_{-i})$ precisely create a situation such that the user performs social objective maximization whilst maximizing self utility.

As soon as one moves away from the SoU objective, due to the social and individual objectives not aligning with each other, basic design techniques like VCG mechanism are not useful. This is one of the main reasons why there are significantly fewer results on fairness in the mechanism design literature.

In [CoGkGo13], authors use the concept of "proportional fairness" (same as maximizing the sum of log of utilities) to reduce disparity. The focus is on a tax-less mechanism where the contract proposes to throw away existing resources ("resource-burning") in order to tax untruthful agents. This mechanism achieves at least $\frac{1}{e}$ fraction of proportionally fair allocation. In [PaVo12], optimal auctions in the Bayesian set up are derived, such that instead of efficiency maximization or revenue maximization, a linearly combined metric is maximized which favors exchange of goods at low prices (thereby ensuring fairer trade).

"Envy-freeness" is another well-known criterion for equitable allocations (see [Va74]). This notion was originally proposed for exchange economies where an allocation is called envy-free if no agent is strictly better off by taking someone else's allocation instead of their own[1]. Such a notion was argued in terms of the stability it provides, since each agent may be content with what they have comparing to the possible option of acquiring someone else's allocation. This, however, is an ex-post notion and in general imposes quite stringent constraints on the

[1]Note that this exchange refers to both allocation of good and taxes.

design. For a large enough environment it may indeed be impossible to achieve envy-freeness in optimal allocation.

Max-min fairness is the one notion of fairness that has received the most attention in the literature. In [PoShTe04] authors consider the mechanism design problem under perfect information with weak budget balance constraint and focus on achieving $k-$fairness[2] which is a generalization of $\max - \min$. Possibility results are shown for $k = 1, 3$ whereas $k = 2$ leads to an impossibility result. In [AtYe08] authors provide tight bounds on inefficiency arising from $k-$fairness in terms of budget deficit bound. In [WaXuAhCo15] authors consider a form of VCG mechanism with the AGV (C. d'Aspremont and L.-A. Gérard-Varet) [dAGé79] taxform (an idea arising out of [dAGé79; dACrGé90]). The problem is of Bayesian mechanism design with strong budget balance and efficient allocations. The authors formulate $\max - \min$ in this scenario as a parameter optimization problem and present an algorithm to solve it.

The work of [WeMaOz14] is motivated by the second reason mentioned above for departing from the SoU social objective, in the context of dynamic pricing for electricity markets. The dual version of the resource allocation optimization problem is considered and additional structure is introduced on the dual variables to ensure less variation between them. This would typically ensure less fluctuation in prices (and thus in the monetary transfers) since it is well-known that the dual variables for resource allocation optimization problems act as prices in the corresponding markets.

In [RaBo07] authors provide a centralized (non-strategic, non-decentralized) algorithm for $\max - \min$ fair allocation, with applications to wireless network problems. The proposed algorithm generalizes the well-known water-filling solutions. The authors demonstrate that the $\max - \min$ fair allocation is derived from the geometric properties of the feasible allocation set. Recently, in the area of green communications, analysis has focused on strategic behavior in large-scale, multi-agent systems with decentralization of information. In [BeAlHoAmAl15] authors introduce a fairness-based metric for energy efficiency in wireless communication (non-strategic), which results in a $\max - \min$ optimization problem. Readers may refer to [LaKaChSa10] for a comparative study on different fairness metrics for resource allocation in networking.

---

[2] $k-$fairness is said to have been achieved if at equilibrium every agents achieves utility of at least $-\frac{v_{[k]}}{N}$, where $v_{[k]}$ is the k^{th} lowest cost and N is the number of agents.

We would like to clarify that "fairness" when used in this chapter refers to allocating resources such that all agents receive utilities as close to each other as possible. This is different from the notion of "proportional fairness" [KeMaTa98; MaBa06] in resource allocation problems which is a technique to translate demands into allocation in a proportional manner. It is also different from the idea of fairness used in scheduling algorithms for resource allocation [KuWh04; RaGuGoPs06; HuSuBeAg09; SaAlEr12], which refers to ensuring average processor time assigned to each agent is proportional to their cost. The challenge for the mechanism design problem such as ours is to elicit information about private utility functions in order to make utilities closer at equilibrium. Also the model considered here is of non-cooperation and thus fairness based on cooperative models [BaLaSaSa16] that provide solutions like bargaining are not considered.

In this chapter, we ask if and how we can design mechanisms that implement social objectives that are especially designed for fairness and go beyond the standard objective of SoU. We seek a methodology that is flexible enough to create space for the designer when other notions of fairness might be too stringent. We consider the problem of demand-side management in smart grids as a specific motivating application.

The contributions of this work are summarized as follows.

We consider the form of the social objective given by the additive function $\sum_{i=1}^{N} g_\epsilon(v(x;\theta_i))$ where $v(x;\theta_i)$ is the utility of the i-th user with allocation x and type θ_i, and $g_\epsilon(z)$ is a family of concave functions parametrized by $\epsilon > 0$. This is taken such that at $\epsilon = 0$, the objective is simply the SoU i.e., $g_0(z) = z$; while as the parameter ϵ increases, $g_\epsilon(\cdot)$ becomes more concave. Therefore, as ϵ increases, more fairness is built into the allocation and comparison from the baseline case of SoU ($\epsilon = 0$) can be made. Note that unlike agents' utilities, the choice of the specific function g_ϵ is in the designer's hand, as long as it serves the design objective of fairer allocation.

Within this framework we ask for which values of the parameter ϵ and under what conditions for $g_\epsilon(\cdot)$ is dominant strategy implementation possible and when is Bayesian Nash Equilbrium (BNE) implementation possible. We show that indeed mechanism design is possible provided ϵ is not too large, by providing an upper bound on the range of ϵ. This is done by formulating the incentive compatibility constraints as a set of linear inequalities on the design variables and checking whether this system (together with strong budget balance and/or individual rationality) is feasible. Our proving techniques follow closely the work of [dACrGé90; dACrGé04]. Not surprisingly, the results in the Bayesian set-up are derived under certain assumptions on the prior beliefs, $p_i(\theta_{-i}|\theta_i)$, of agents, which generalize the notion of independent beliefs.

For the case of Bayesian mechanism design, we go one step further and propose a modification to our mechanism which ensures not only that truth-telling is a BNE but also that it is the only BNE. This provides robustness to the mechanism as prediction/focusing towards a specific desired equilibrium is no longer needed. For general Bayesian mechanism design, authors in [MoRe90] have derived a set of sufficient conditions under which augmenting the message space results in full implementation. Our approach is different from that of [MoRe90] in that the sufficient condition is easier to verify and a precise modification to the mechanism is presented. This modification doesn't make significant changes to the mechanism and only requires exchange of one additional message.

We finally demonstrate that the range of ϵ induced by our methodology (from Section 3 and 4) is sufficient to provide quite significant gains in fairness - as measured by the decrease in several measures of disparity. This is done for dominant implementation through analytically solving for relevant quantities for the scenario of two-type agents. Similarly, for Bayesian implementation, we carry out a numerical study for the smart grid application introduced in Section 2.1.

The remaining of this chapter is organized as follows: Section 2.1 describes a simplified model for demand-side management of smart grids motivating the need for fairness in mechanism design. Section 2.2 defines the general centralized problem which has been modified for fairness. The next two sections prove the existence of mechanisms that implement the centralized problem for type sets of size two in dominant strategy (Section 3) and for general finite type sets in BNE (Section 4). Section 5 presents the modification for full Bayesian implementation. The analytical derivation for dominant strategy implementation is worked out in Section 6 and the numerical study for the smart grid application is presented in Section 7. Finally, Section 8 discusses future work and immediate extensions of the main results in this chapter.

2. Motivating Application and General Model

2.1. Demand-side management of smart grids

Below we discuss a practical application, as a motivation for our work, from demand-side management of smart grids using a simplified version of the model described in [SaMoScWo12].

We assume a power system with N users and a single supplier. The supplier does not have any decision making power and is considered a passive agent. Each user has a decision to make namely, the power level he/she wants to consume. It is assumed that the user loads are *price-elastic* [St02] i.e., users can decide the power level they want for themselves after taking price into consideration. The system runs for a single time slot[3]. Let $x = (x_1, \ldots, x_N)$ denote the power consumption profile of all the users. Owing to the power demand of *must-run* loads, for each user, the power consumption has a minimum m_i limit, which is considered public information[4]. This gives the feasible set of power allocation as $\mathcal{X} = \{x \in \mathbb{R}_+^N \mid 1 \leq i \leq N, m_i \leq x_i\}$. Users are assumed to be heterogeneous and are modeled by private utility functions [FaAl00], $v(x_i, \theta_i)$, where x_i is the allocation and parameter $\theta_i \in \Theta_i$ represents each user's value for power. Following [SaMoScWo12] we assume quadratic utilities of the form

$$v(x_i, \theta_i) = \begin{cases} \theta_i x_i - \frac{x_i^2}{4} & \text{if } 0 \leq x_i < 2\theta_i \\ \theta_i^2 & \text{if } x_i \geq 2\theta_i. \end{cases} \tag{5.1}$$

The power supplier, although is a passive agent, incurs a cost for supplying the total power demanded. This cost is assumed to be an increasing, quadratic (convex) function, $C(y) = ay^2$, where a is a technology dependent parameter and $y = \sum_{i=1}^N x_i$ is the total demanded power.

Most of the literature on mechanism design for efficient allocation of power defines efficiency from a market point of view i.e., power is allocated such that the sum of utilities is maximized i.e.,

$$\max_{x \in \mathcal{X}} \sum_{i=1}^N v(x_i, \theta_i) - C\left(\sum_{j=1}^N x_j\right). \tag{5.2}$$

[3] This is a simplification over [SaMoScWo12], where $K \geq 1$ time slots are considered. This simplification retains all the salient features of the model, since in our work the focus is on allocation across agents and not across time.

[4] Again, this is a mild simplification compared to [SaMoScWo12], where m_i is part of agent i's private information.

Table 2.: $\Theta_i = \{5, 15\}$, $N = 100$, 10% agents with high type, $a = 0.3$, $m_i = 0.1$

$\hat{x}_H$	$\hat{x}_L$	σ σ/μ (%) GC (%)	$v(\hat{x}_H, \theta_H)$	$v(\hat{x}_L, \theta_L)$	σ σ/μ (%) GC (%)
1.477	0.100	0.4152 175% 53%	21.608	0.498	6.36 244% 74%

$\hat{t}_H$ (VCG)	$\hat{t}_L$ (VCG)	σ σ/μ (%) GC (%)
21.01	1.43	5.90 174% 53%

Setting aside for the time being the issue of how this problem can be solved in a decentralized way in the presence of strategic agents, let us study the solution of the centralized problem (5.2). For this define $\hat{x}(\theta)$ as the maximizer in (5.2), where $\theta = (\theta_1, \ldots, \theta_N)$.

For illustration, we consider a specific instance of the problem in (5.2) with the following parameters. Users can be of two types, either high type (i.e., users who have a high valuation for power) with parameter $\theta_H = 15$ or low type (i.e., users who have a low valuation for power) with parameter $\theta_L = 5$. There are $N = 100$ users and 10% of them are of high type. Due to the symmetry of the system the solution $\hat{x}(\theta)$ to (5.2) can be expressed by just two numbers: allocation to agents with high type, say $\hat{x}_H$, and allocation to agents with low type, say $\hat{x}_L$. These results are provided in Table 2. We also tabulate the corresponding utilities achieved at this allocation, together with different measures of disparity: (a) the standard deviation, σ; (b) the normalized standard deviation by the mean, σ/μ; and (c) the Gini coefficient (GC)[5]. It is evident that there is a great disparity in the solution of the SoU problem (5.2) regardless of the measure of disparity one considers.

Now consider the problem of mechanism design when allocation is through (5.2). We

[5]Defined as the ratio of mean of the difference between every possible pair of data points with mean size. The lower the value of GC, the more equitable the allocation. GC of 0% is perfectly fair and GC of 100% is absolutely unfair - everyone except an individual receiving 0. GC is independent of scale, hence can also be used in comparing fairness across different settings.

choose as our baseline case the widely used and accepted Vickerey-Clarke-Groves (VCG) mechanism [Gr73]. Indeed, VCG is proven to be a truthful mechanism when allocation is based on SoU and moreover, it provides implementation in dominant strategies and thus is robust with respect to information available to agents. Each agent is taxed based on the quoted type profile by all agents. If agents quote $\phi = (\phi_1, \ldots, \phi_N)$, then any agent i's net utility is $v(x_i; \theta_i) - t_i(\phi)$. Due to the symmetry of the system,the VCG taxes can be expressed by two quantities, $\hat{t}_H$ (taxes paid by users that quote type as x_H) and $\hat{t}_L$ (taxes paid by users that quote type as x_L). In the last three columns in Table 2 we present results on disparity in taxes paid by different agents by the same three measures mentioned above. As above, there is disparity in the tax paid between agents, an expected result since in a market, agents who pay more can get more.

Now suppose that a designer wants to impose a fairer or less disparate allocation of power and consequently taxes. The main question is how can this be achieved? We outline our approach below in the context of the aforementioned application and develop this approach in the general case in the rest of the chapter.

We propose that the designer solves the following optimization problem instead of (5.2)

$$\max_{x \in \mathcal{X}} \sum_{i=1}^{N} g(v(x_i, \theta_i)) - C\left(\sum_{i=1}^{N} x_i\right) \tag{5.3}$$

where $g(\cdot)$ is a concave, increasing function. Due to the concavity of $g(\cdot)$ the law of diminishing returns ensures that the resulting optimal allocation is less disparate than the one from (5.2). Consequently, the allocation resulting from (5.3) is preferable for a designer interested in fairer allocation. However, this allocation cannot be implemented through the VCG mechanism. The reason is that although the designer allocates through (5.3), each user optimizes his/her true net utility, and as a result the objectives of the designer and the user are not aligned through the VCG taxes. Hence, for a designer who wants to implement a fairer allocation from (5.3), the question is whether there is a truthful taxation mechanism for implementing this allocation in the presence of strategic agents. In a nutshell, this chapter provides sufficient conditions under which the answer to the above question is positive.

2.2. General Centralized Problem

For a system of agents $\mathcal{N} = \{1, \ldots, N\}$, allocation is defined via the following optimization problem:

$$\hat{x}_\epsilon(\theta) = \arg\max_{x \in \mathcal{X}} \sum_{i \in \mathcal{N}} g_\epsilon(v(x; \theta_i)), \tag{5.4}$$

where $\mathcal{X} \subset \mathbb{R}_+^N$ is the constraint set, $\theta = (\theta_i)_{i \in \mathcal{N}} \in \Theta \triangleq \times_{i \in \mathcal{N}} \Theta_i$ is the *type profile* of agents and Θ_i is the discrete type set for agent i with $|\Theta_i| = L_i$. The utility function $v(x; \theta_i)$ measures agent i's satisfaction at allocation x with private type being θ_i. A concave transformation $g_\epsilon(\cdot)$ is applied for making the allocation *fairer* compared to the SoU setup. It is further assumed that the functions $g_\epsilon : \mathbb{R} \to \mathbb{R}$ and $v(\cdot; \theta_i) : \mathbb{R}_+^N \to \mathbb{R}$ are such that the optimization has a unique solution (e.g., if $v(\cdot; \theta_i)$, $g_\epsilon(\cdot)$, $g_\epsilon(v(\cdot; \theta_i))$ are concave and increasing and $\mathcal{X}$ is a convex set).

Specifically the form $g_\epsilon(z) = z - \epsilon f(z)$ for $\epsilon \geq 0$ is considered, where $f(\cdot)$ is assumed to be a convex function. Note that at $\epsilon = 0$ optimization (5.4) becomes the SoU problem, while as ϵ increases from 0 the function g_ϵ has a stronger concave component. The form $g_\epsilon(z) = z - \epsilon f(z)$ is considered without significant loss generality, since it closely emulates Taylor's series (w.r.t. ϵ) of many interesting families of concave functions. Consider for example the family $h_\epsilon(z) = z^{1-\epsilon}$ whose Taylor's series is $h_\epsilon(z) = z - \epsilon(z \log(z)) + o(\epsilon)$, where $z \log(z)$ is indeed convex. In any case, the validity of a particular choice of the function $f(\cdot)$ has to be justified in a specific application by the reduction of an appropriate disparity measure it achieves.

3. Dominant Strategy Implementation

The Mechanism design problem in this section is to find a message space $\mathcal{M} = \times_{i \in \mathcal{N}} \mathcal{M}_i$ and allocation, tax functions $(\tilde{x}, t) : \mathcal{M} \to \mathcal{X} \times \mathbb{R}^N$ such that the induced game for agents in $\mathcal{N}$ with action space $\mathcal{M}$ and quasi-linear utilities

$$\tilde{u}_i(m; \theta_i) = v(\tilde{x}(m); \theta_i) - t_i(m) \quad \forall\, m \in \mathcal{M},\ i \in \mathcal{N} \tag{5.5}$$

has a dominant strategy equilibrium[6] m^* for which $\tilde{x}(m^*) = \hat{x}_\epsilon(\theta)$, where $\theta = (\theta_i)_{i \in \mathcal{N}}$ is the true type profile. This is known as Dominant Strategy Incentive Compatibility (DSIC).

In general, dominant strategy implementation is very restrictive (note that the well studied VCG mechanisms are no longer applicable since this is not the maximization of SoU). Following [dAGé79], the special case of $L_i = 2 \; \forall \; i \in \mathcal{N}$ is considered in this section. In particular, each agent is assumed to have a high and a low type, i.e., $\Theta_i = \{\theta_i^H, \theta_i^L\} \; \forall \; i \in \mathcal{N}$.

The proposed mechanism is a direct mechanism, thus $\mathcal{M}_i = \Theta_i \; \forall \; i \in \mathcal{N}$. Agents report their types (possibly untruthfully) $\phi = (\phi_i)_{i \in \mathcal{N}}$ and the allocation they receive on the basis of this is the optimal allocation $\hat{x}_\epsilon(\phi)$ for the quoted type profile ϕ, where $\hat{x}_\epsilon(\cdot)$ is defined in (5.4).

Assuming that for not participating in the mechanism, an agent receives 0 utility value (including 0 tax), the voluntary participation condition for dominant strategy implementation is $\forall \; (\theta_i, \theta_{-i}) \in \Theta, \; i \in \mathcal{N}$,

$$v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_i) - t_i(\theta_i, \theta_{-i}) \geq 0. \tag{5.6}$$

This is the *ex-post* version of the individual rationality (IR).

We now explain the intuition behind the main result in this section and subsequently formalize it Theorem 3.1. It has been shown in the literature [dAGé79] that for the SoU optimization problem (under certain conditions) there exists a mechanism which implements the optimal allocation in dominant strategies. However, when the optimization problem changes to the one in (5.4), implementation may not be possible since the designer's objective (objective in (5.4)) is not aligned with the individual users' objectives which are always of the form $v(\hat{x}_\epsilon(\theta); \theta_i) - t_i(\theta)$. In fact, as ϵ increases from zero, the two objectives become more and more different. In Theorem 3.1 we are trying to quantify this effect by showing that for sufficiently small ϵ, the first order differences $-\epsilon f(v(x; \theta_i))$ do not affect implementability. We do that by summarizing this effect in the quantity $K(\cdot)$ defined in Condition (A_1). Intuitively, as ϵ increases the feasibility set defining taxes that satisfy the DSIC and IR conditions shrinks and at some point vanishes.

[6]For any θ, message $m^* \in \mathcal{M}$ is a dominant strategy equilibrium if it satisfies $\tilde{u}_i(m_i^*, m_{-i}; \theta_i) \geq \tilde{u}_i(m_i, m_{-i}; \theta_i) \quad \forall \, m_i \in \mathcal{M}_i, \; \forall \, m_{-i} \in \mathcal{M}_{-i}, \; \forall \, i \in \mathcal{N}$.

Next we state the assumption on (5.4) under which the main result in this section is derived. For the optimization (5.4), first define $\forall\ i \in \mathcal{N}$, $\mathcal{F}_i := \{(\psi_i, \xi_i) \in \Theta_i^2 \mid \psi_i \neq \xi_i\}$. With this define the difference, $\forall\ i \in \mathcal{N}$, $(\theta_i, \phi_i) \in \mathcal{F}_i$, $\theta_{-i} \in \Theta_{-i}$, $\epsilon > 0$,

$$K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) := \sum_{j \in \mathcal{N}} g_\epsilon(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j)) - \sum_{j \in \mathcal{N}} g_\epsilon(v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j)). \tag{5.7}$$

Optimality conditions from (5.4) imply that the above difference is always non-negative, since $\hat{x}_\epsilon(\theta_i, \theta_{-i}), \hat{x}_\epsilon(\phi_i, \theta_{-i}) \in \mathcal{X}$ and $\hat{x}_\epsilon(\theta_i, \theta_{-i})$ is the maximizer of $\sum_{j \in \mathcal{N}} g_\epsilon(v(x; \theta_j))$ subject to $x \in \mathcal{X}$. However, due to the finite type spaces, the difference above is expected to be strictly positive. We formalize this expectation into the following condition.

Condition (A$_1$) Assume that $\exists\ \epsilon_{max} > 0$ such that for all $0 \leq \epsilon < \epsilon_{max}$, $\forall\ i \in \mathcal{N}$, $(\theta_i, \phi_i) \in \mathcal{F}_i$, $\theta_{-i} \in \Theta_{-i}$,

$$K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) + K_i(\phi_i, \theta_i, \theta_{-i}, \epsilon) > 0. \tag{5.8}$$

Condition (A$_1$) need only be checked at $\epsilon = 0$. By continuity of the optimization on parameter ϵ, this will imply that these conditions continue to hold for all $0 \leq \epsilon < \epsilon_{max}$ for some $\epsilon_{max} > 0$.

The first contribution of this chapter is summarized in the following theorem.

Theorem 3.1. *If Condition (A$_1$) is satisfied then $\exists\ \tilde{\epsilon}_{max} > 0$ such that for all $0 \leq \epsilon < \tilde{\epsilon}_{max}$ there exist taxes $\big(t_i(\phi)\big)_{\phi \in \Theta, i \in \mathcal{N}}$ that satisfy DSIC (which implies implementation in dominant strategies) and IR.*

Proof. Please see Appendix 9.1 at the end of this chapter. □

The relation between ϵ_{max} and $\tilde{\epsilon}_{max}$ is defined in the proof, specifically via the relation in (5.25).

The theorem is proved in two parts, firstly the DSIC condition is stated as a system of linear inequality constraints on the taxes (which act as design variables). In the spirit of the technique used in [dACrGé90], Farkas Lemma is used to get an alternate condition which is equivalent to satisfying DSIC. This alternate condition can be rewritten so that dual variables arising from the Farkas Lemma are eliminated and thus the condition contains no unknowns and only the parameters of the problem. Finally these conditions are compared with Condition (A_1) to get the result and the precise value of $\tilde{\epsilon}_{max}$.

Note that the stricter, but easier to verify condition

$$\exists\, \epsilon_{max} > 0 \text{ s.t. } \forall\, 0 \leq \epsilon < \epsilon_{max}, \qquad \min_{i \in \mathcal{N}} \min_{(\theta_i, \phi_i) \in \mathcal{F}_i} \min_{\theta_{-i} \in \Theta_{-i}} K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) > 0, \tag{5.9}$$

is sufficient to guarantee that Condition (A_1) is satisfied. Note that the above condition and similarly condition (A_1) do not impose a restriction on the agents' types but only on the maximum value of ϵ for which dominant implementation can be guaranteed. This implies that the functions $g_\epsilon(z) = z - \epsilon f(z)$, cannot be made to be as concave as the designer might have wanted in order to have a fairer allocation, while still achieving implementation in dominant strategies.

4. Bayesian Implementation

In this section we drop the assumption from the previous section that the type sets are binary and we consider type sets of arbitrary size. Dominant strategy implementation is too restrictive for the general scenario, hence the next best reasonable solution concept - Bayesian implementation, is considered.

In a Bayesian set up, agents have a prior distribution on the type profile. For agent i, prior is $p_i \in \Delta(\Theta)$. For basic regularity assume that the prior gives non-zero probability on all points of Θ (this is only a technical condition and the ensuing results can be proved without it as well). These priors are assumed to be *common knowledge* between agents and designer - hence there is no need to introduce second order beliefs over the priors and so on.

The mechanism used is a direct mechanism with allocation function $\hat{x}_\epsilon(\cdot)$, as defined in (5.4). Given the allocation and tax functions $(\hat{x}_\epsilon, t) : \Theta \to \mathcal{X} \times \mathbb{R}^N$, the utility function in the Bayesian set up for strategy profile $\sigma = \left(\sigma_j : \Theta_j \to \Theta_j\right)_{j \in \mathcal{N}}$ is given by, $\forall\, i \in \mathcal{N}$,

$$\tilde{u}_i(\sigma \mid \theta_i) = \mathbb{E}_{p_i}\Big[v\Big(\hat{x}_\epsilon(\sigma(\theta)); \theta_i\Big) - t_i\Big(\sigma(\theta)\Big) \mid \theta_i\Big] \tag{5.10}$$

where θ_i is the true type of agent i and the expectation is conditioned on it. The Bayesian implementation condition–also known as Bayesian Strategy Incentive Compatibility (BSIC)–

for the direct mechanism is that the truthful strategy $\sigma_i^*(\theta_i) = \theta_i, \ \forall\, \theta_i \in \Theta_i, \ \forall\, i \in \mathcal{N}$ must be a Bayesian Nash equilibrium (BNE)[7] for the induced Bayesian game.

In addition, it is required the tax function to have the *Strong Budget Balance* (SBB) property, i.e.,

$$\sum_{i \in \mathcal{N}} t_i(\psi) = 0, \quad \forall\, \psi \in \Theta. \tag{5.11}$$

We restrict attention to optimization (5.4) and priors $\big(p_i(\cdot)\big)_{i \in \mathcal{N}}$ that satisfy the following conditions.

Condition (A$_2$) Assume that $\exists\, \epsilon_{max} > 0$ such that $H(\epsilon) > 0$ for all $0 \leq \epsilon < \epsilon_{max}$, where

$$H(\epsilon) := \min_{i \in \mathcal{N}} \min_{(\theta_i, \phi_i) \in \mathcal{F}_i} \mathbb{E}_{p_i}\Big[h_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) \mid \theta_i\Big], \tag{5.12a}$$

$$h_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) := \sum_{j \in \mathcal{N}} \Big(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j) - v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j) \Big). \tag{5.12b}$$

This is the most general form of the assumption needed about the optimization. More relaxed assumptions are discussed later in this section.

Condition (B) Assume that for any non-zero vector $R := (R(\psi))_{\psi \in \Theta}$ with $R(\psi) \in \mathbb{R}$, there does not exist any $\lambda := (\lambda_k(\theta_k, \phi_k))_{k \in \mathcal{N}, (\theta_k, \phi_k) \in \mathcal{F}_k}$ with $\lambda_k(\theta_k, \phi_k) \in \mathbb{R}_+$ such that $\forall\, i \in \mathcal{N}, \forall\, \psi \in \Theta$,

$$p_i(\psi_{-i} \mid \psi_i) \sum_{\substack{\phi_i \in \Theta_i \\ \phi_i \neq \psi_i}} \lambda_i(\phi_i, \psi_i) - \sum_{\substack{\theta_i \in \Theta_i \\ \theta_i \neq \psi_i}} p_i(\psi_{-i} \mid \theta_i) \lambda_i(\theta_i, \psi_i) = R(\psi). \tag{5.13}$$

Note that Condition (B) involves only the priors. This condition was first introduced in [dAGé79] and is a generalization of independent priors, i.e., $p_j(\psi_{-j} \mid \psi_j) = p_j(\psi_{-j})$ (refer to [dAGé79] for an example of priors which are not conditionally independent but still satisfy the condition above).

The second contribution of this chapter is summarized in the following theorem.

[7] Strategy $\sigma^* = \big(\sigma_i^* : \Theta_i \to \Theta_i\big)_{i \in \mathcal{N}}$ is a BNE if $\forall i \in \mathcal{N}, \ \forall\, \theta_i \in \Theta_i, \ \forall\, \sigma_i' : \Theta_i \to \Theta_i$; $\tilde{u}_i(\sigma_i^*, \sigma_{-i}^* \mid \theta_i) \geq \tilde{u}_i(\sigma_i', \sigma_{-i}^* \mid \theta_i)$.

Theorem 4.1. *If Conditions (A_2) and (B) are satisfied then for all* $0 \leq \epsilon < \epsilon_{max}$*, there exist taxes* $\big(t_i(\phi)\big)_{\phi \in \Theta, i \in \mathcal{N}}$ *that satisfy BSIC (which implies implementation in BNE) and SBB.*

Proof. Please see Appendix 9.3 at the end of this chapter. □

The proof of this theorem begins by modifying tax variables t to new variables z where the relation between the two is such that SBB is always satisfied (without losing generality); this is the famous AGV form of taxes [dACrGé90]. As in Section 3, the BSIC condition on the design variables z is stated as a system of linear inequalities which is further converted into an alternate expression by the use of Farkas Lemma. Unlike Section 3 though the dual variables can't be eliminated in this case. With the help of Condition (B) the alternate condition is further simplified to a convenient form which is then rewritten (by exchange of summation order) in a way that can be compared to Condition (A_2). This leads to the result and also the precise value of $\tilde{\epsilon}_{max}$

For ease of verification, instead of checking (5.12) one can alternatively check the stricter condition

$$\min_{i \in \mathcal{N}} \min_{(\theta_i, \phi_i) \in \mathcal{F}_i} \min_{\theta_{-i} \in \Theta_{-i}} h_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) > 0. \tag{5.14}$$

or even

$$K(\epsilon) + \epsilon \, s_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) > 0, \tag{5.15a}$$

$$\text{where} \quad K(\epsilon) := \min_{i \in \mathcal{N}} \min_{(\theta_i, \phi_i) \in \mathcal{F}_i} \min_{\theta_{-i} \in \Theta_{-i}} K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon), \tag{5.15b}$$

$$s_i(\theta_i, \phi_i, \theta_{-i}, \epsilon) := \sum_{j \in \mathcal{N}} f(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j)) - f(v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j)). \tag{5.15c}$$

These conditions are easier to verify since they are prior-independent and depend only on the optimization problem (5.4). Also, it is straightforward to check that (5.15a) $\Rightarrow$ (5.14) $\Rightarrow$ Condition (A_2). We finally mention that as in the case discussed in Section 3, the above condition, due to continuity, need only be checked at $\epsilon = 0$.

Note that the SBB condition in (5.11) is an additional property that we impose on the design of the mechanism and not an additional assumption. In other words, if one removes this additional requirement, all our results will still hold, and furthermore, the range of allowable values of ϵ will further increase.

5. Full Bayesian Implementation

In accordance with the majority of the literature on Bayesian implementation (see for instance [BöKrSt15; Ja91; Ja01; Ha12b]), the main result in the previous section aimed at BSIC condition for implementation. This ensures that truth-telling is a BNE but gives no information about other possible BNE. In general this may be a problem, since when the Bayesian game is played, the designer cannot predict in advance which BNE will be achieved. The justification used in such a situation is that of "focusing"[8]. Along with the mechanism, the selected BNE is also announced by the designer. All the agents are then focused towards this particular BNE and in anticipation that others will be playing according to it, they too choose to play it.

In this section we modify our mechanism by augmenting each agents' message space to include one additional message - which takes values in a continuous space - with the specific aim of achieving full implementation, i.e., truth-telling as the only BNE. This will add to the robustness of the mechanism and will put to rest any equilibrium selection issues. Such a modification only requires that the BSIC constraints (see (5.37)) can be satisfied with strict inequality. This gives the designer room to alter taxes.

For general Bayesian mechanism design, authors in [MoRe90] have derived sufficient conditions under which augmenting the message space results in full implementation. We choose however not to follow this technique, since it only provides existence results and the sufficient condition itself is difficult to check since it requires searching over functions. Instead, a concrete augmentation of message space and a concrete modification of the mechanism is proposed below.

In our analysis we restrict attention to private goods i.e., any agent i's utility is affected only by level of his/her consumption level x_i and the scenario where the utility functions $v_i(\cdot)$ are differentiable. Extending to general public goods is possible, but will make the proofs more technical and obfuscate the basic idea behind them. Similarly, mean value theorem can be used (to replace derivatives) but it will make the analysis more technical.

[8] This is the same justification as the one used for working without loss of generality, with the "revelation principle" and direct mechanisms.

The new message space, allocation and taxes for any agent i are

$$\mathcal{M}_i = \Theta_i \times [-\delta, +\delta], \quad m_i = (\phi_i, y_i). \tag{5.16a}$$

$$\tilde{x}_{\epsilon,i}(m) = \hat{x}_{\epsilon,i}(\phi) + y_i, \tag{5.16b}$$

$$\tilde{t}_i(m) = \begin{cases} \hat{t}_i(\phi) + y_i v'(\hat{x}_{\epsilon,i}(\phi); \phi_i) & \text{if } y \in (-\delta, +\delta), \\ B & \text{if } y = \pm\delta. \end{cases} \tag{5.16c}$$

where $\hat{x}_\epsilon(\phi)$ is the allocation function defined in (5.4) and $\hat{t}_i(\phi)$ is a tax designed for BSIC and SBB (from the previous section). Also, $\delta, B > 0$ are constants chosen by the designer. In particular, δ is a small constant, the choice of which will be clear in the proof of Theorem 5.1. The constant B is chosen to be large enough so that no rational agent will ever choose message $y_i = \pm\delta$. This is possible because due to the type sets being discrete, the utilities are bounded. The modification above allows agents to change their allocation by a small amount $y_i \in [-\delta, +\delta]$. For this increase/decrease in allocation they are charged/subsidized at the "market" price $v'(\hat{x}_{\epsilon,i}(\phi); \phi_i)$ corresponding to the truth-telling strategy. This modification serves the purpose of giving agents more delicate control of their allocation and utility than in a discrete set-up. Technically, this allows the designer to disrupt any BNE at which the price an agent is capable of paying (i.e., his/her derivative without tax terms) doesn't match the market price designed for truthful strategies.

Theorem 5.1. *Assuming that the conditions of Theorem 4.1 are satisfied and the space of taxes has a proper interior (i.e., all incentive inequalities are satisfied strictly), then for the mechanism defined in* (5.16) *at any true type profile* $\theta = (\theta_i)_{i \in \mathcal{N}}$*,*

(1) message $m^* = \left(m_i^*\right)_{i \in \mathcal{N}} = \left((\theta_i, 0)\right)_{i \in \mathcal{N}}$ *is a BNE.*

Furthermore, assuming $\forall\, i \in \mathcal{N},\ (\theta_i, \phi_i) \in \mathcal{F}_i$*, the sign of the quantity* $v'(\hat{x}_{\epsilon,i}(\phi_i, \phi_{-i}); \theta_i) - v'(\hat{x}_{\epsilon,i}(\phi_i, \phi_{-i}); \phi_i)$ *remains the same* $\forall\, \phi_{-i} \in \Theta_{-i}$*,*

(2) any message other than m^* *is not a BNE.*

Hence the mechanism in (5.16) *achieves full Bayesian implementation.*

Proof. Please see Appendices 9.4 and 9.5 at the end of this chapter. □

Several comments are in order regarding the statement and the proof of theorem. The condition just prior to Part (2) of the theorem can be interpreted as follows. Start with a stricter condition: sign of $v'(x;\theta_i) - v'(x;\phi_i)$ is the same for all $x \in \mathcal{X}$. Based on this the types in Θ_i can be ordered solely based on the derivative they induce. This ordering, which is otherwise not possible (since the expression can in general have different signs as x varies), has further significance. In SoU problems (like (5.4) at $\epsilon = 0$), optimal allocation is dependent only on the rate of growth of utilities[9]. Thus for ordered types, from (5.4) at $\epsilon = 0$, optimal allocation for an agent depends monotonically on their type irrespective of others' types. The actual condition prior to Part (2) of the theorem is a simpler version of the stricter condition and requires checking for only finite number of arguments.

The proof of the first part of the theorem begins by establishing that truth-telling is better for any agent given that others are playing m^*_{-i}. Then, further optimization on the y_i variable, with the use of calculus (since the utility depends on this continuously), proves that truth-telling and $y_i = 0$ is indeed a BNE.

The proof of the second part of the theorem relies on establishing a profitable deviation with the y_i variable (for some agent i) whenever the message is different from m^*. The gap in the incentive inequalities (assumed at the beginning of the theorem) determines the range for y_i i.e. value of the δ.

In general, the proof of Part (2) of the theorem requires for all θ_i, i and σ_{-i}, the quantities

$$\sum_{\theta_{-i}\in\Theta_{-i}} p_i(\theta_{-i} \mid \theta_i)\Big[v'(\hat{x}_i(\sigma(\theta));\theta_i) - v'(\hat{x}_i(\sigma(\theta));\sigma_i(\theta_i))\Big] \tag{5.17}$$

to remain non-zero whenever $\sigma_i(\theta_i) \neq \theta_i$. Any condition that ensures that for any non-truthful strategy σ there exist an agent i and type θ_i for which the above expression is non-zero, would be sufficient for the theorem to hold. For example, given type sets, utilities and the optimization, a condition on the priors may then be sufficient. Furthermore, such condition wouldn't significantly reduce the set of possible priors since the space of priors is essentially equivalent to $\mathbb{R}^{\prod_{i=1}^{N} L_i - 1}$ while only a finite set of hyperplanes are eliminated from it.

Finally, the mechanism presented has the SBB property only at BNE. However, as in the

[9]This is because at optimum, trade-offs between giving infinitesimal allocation to one agent vs. another will be determined by the slope of their utilities at those points.

proof of Theorem 4.1, using the AGV form for budget balanced taxes i.e., $\tilde{t}_i(m) = \tilde{z}_i(m) - \frac{1}{N-1}\sum_{j\neq i}\tilde{z}_j(m)$ a similar modification can make the mechanism in (5.16) budget balanced off-equilibrium as well. In order to effectively convey the salient features of the augmentation for full implementation we presented the mechanism without this modification.

6. An Analytical Example for Dominant Implementation

This section contains analysis of an example where we are interested in evaluating the gain in overall fairness w.r.t. allocation and taxes, attained with the proposed method. The example considered provides closed form expressions for the relevant quantities from the optimization (5.4), which enables a closed form expression for the gains in fairness.

Consider $\Theta_i = S \triangleq \{\theta_H, \theta_L\}\ \forall\, i \in \mathcal{N}$, where $\theta_H > \theta_L > 0$ and linear utilities with *private consumption* i.e., $v(x;\theta_i) \coloneqq \theta_i x_i$ (this is a simplification over piecewise linear utilities, for e.g. [Mo16] uses it in an energy storage application, [CuNa05] uses it in resource allocation on wireless networks). The constraint set is $\mathcal{X} = \{x \in \mathbb{R}_+^N \mid \sum_{i=1}^N x_i \leq 1\}$ and $g_\epsilon(z) = z - \epsilon z^2$. Thus the centralized optimization problem is

$$\hat{x}_\epsilon(\theta) = \arg\max_{x\in\mathcal{X}} \sum_{i\in\mathcal{N}} \left(\theta_i x_i - \epsilon\theta_i^2 x_i^2\right). \tag{5.18}$$

At $\epsilon = 0$ the objective is exactly the SoU. As ϵ starts increasing from 0 onwards, the optimization problem is transformed such that higher utilities will be weighed less than lower ones - thereby giving closer to equal distribution of allocation.

Due to the symmetric nature of the problem, the solution can be described by the pair (x_H, x_L), which denotes the allocation to any agent with type θ_H, θ_L, respectively. Generically we denote by $m \in \{0,\ldots,N\}$ the number of agents with type θ_H.

At $\epsilon = 0$ a solution (there are multiple) to the above is $x_H = \frac{1}{m}$ and $x_L = 0$. For $0 < \epsilon < \frac{m}{2}\frac{\theta_H - \theta_L}{\theta_H^2}$, the unique solution is the same. For $\epsilon > \frac{m}{2}\frac{\theta_H-\theta_L}{\theta_H^2}$, the unique solution is

$$x_H = \frac{\frac{(N-m)}{2\epsilon}(\theta_H - \theta_L) + \theta_L^2}{m\theta_L^2 + (N-m)\theta_H^2}, \qquad x_L = \frac{\frac{m}{2\epsilon}(\theta_L - \theta_H) + \theta_H^2}{m\theta_L^2 + (N-m)\theta_H^2}. \tag{5.19}$$

Dominant implementation condition from the proof of Theorem 3.1 is satisfied if $\tilde{\epsilon}_{max}$ is

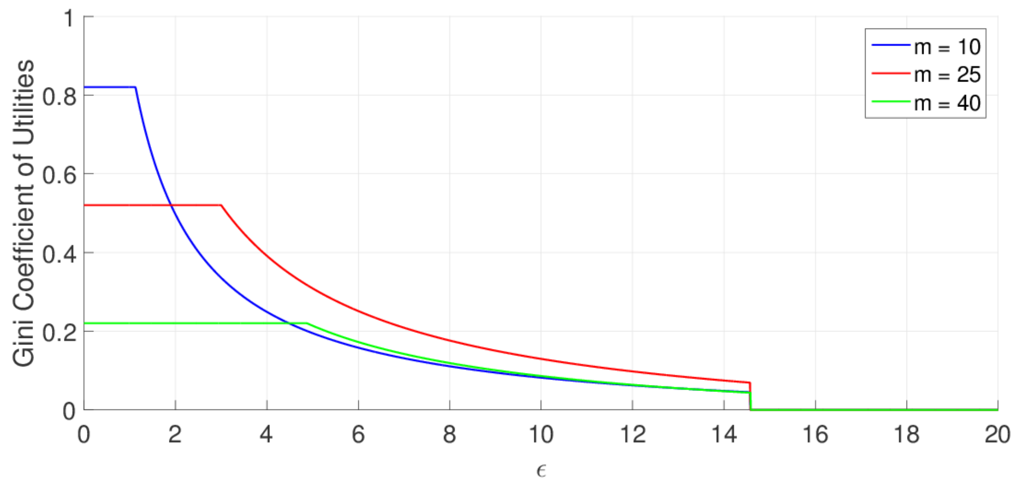

Figure 16.: Gini Coefficient of $\left(v(\hat{x}_\epsilon(\theta);\theta_i)\right)_{i\in\mathcal{N}}$ vs ϵ, $N = 50$.

such that

$$2\tilde{\epsilon}_{max} = \frac{N+1}{\theta_H+\theta_L} + \frac{\theta_L^2}{(\theta_H+\theta_L)\Big(m\theta_L^2 + (N-m-1)\theta_H^2\Big)} \tag{5.20a}$$

$$\geq \frac{N+1}{\theta_H+\theta_L} + \frac{\theta_L^2}{(\theta_H+\theta_L)\Big((N-1)\theta_H^2\Big)}. \tag{5.20b}$$

The inequality is calculated so that the expression is independent of m.

The main quantity of interest is the variation in the utility profile $\left(v(\hat{x}_\epsilon(\theta);\theta_j)\right)_{j=1}^N$. We consider the well known *Gini coefficient* (GC) as a measure of disparity in allocation.

$$GC = \frac{m(N-m)(\theta_H x_H - \theta_L x_L)}{N(m\theta_H x_H + (N-m)\theta_L x_L)} \tag{5.21}$$

where GC depends on ϵ via x_H, x_L.

Fig. 16 plots GC vs. ϵ for $N = 50, (\theta_H,\theta_L) = (1, 0.75)$ and three cases $m \in \{9, 24, 39\}$. In this example $\tilde{\epsilon}_{max} = 14$. As is evident from Fig. 16, there is considerable reduction in the GC in the range of allowable ϵ. The gains are particularly significant for smaller values of m i.e., $m = 9, 24$ since in those cases the inequality in allocation is significant to begin with (at $\epsilon = 0$).

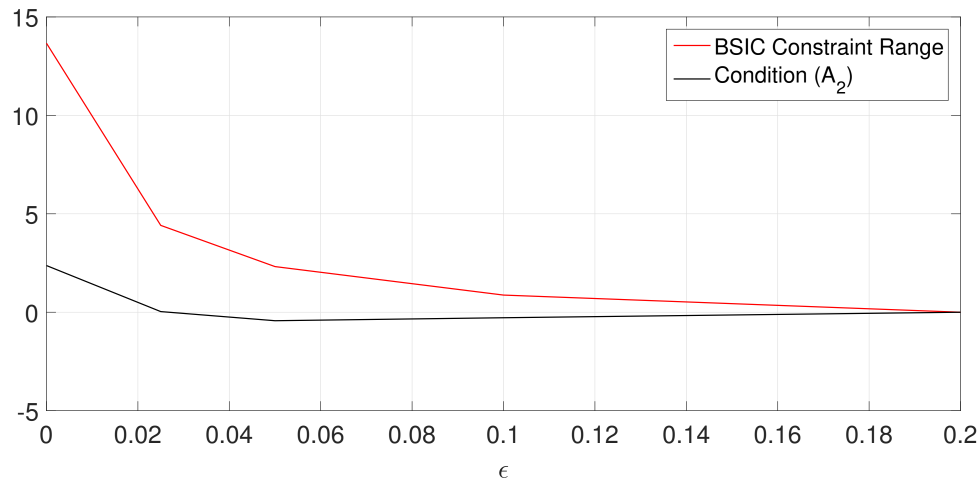

Figure 17.: Plot of quantities that determine ϵ_{max}. The difference between upper and lower bounds imposed by BSIC and $H(\epsilon)$ from eq. (5.12a).

7. Bayesian Implementation for smart grid demand-side management

This section considers the case of Bayesian implementation and further elaborates on the smart grid application presented in Section 2.1.

For the numerical analysis we consider parameters (as in Section 2.1), $N = 100$, $\theta_H = 15$, $\theta_L = 5$, $a = 0.3$, $m_i = 0.1$. In this application we have symmetric agents and hence any profile of types $\theta = (\theta_1, \ldots, \theta_{100})$ can be represented by a single number $m \in \{0, \ldots, 100\}$ which represents the number of agents with type θ_H. A common prior Binomial(100, 0.1) is assumed on m, so the typical number of high-type users is 10. Also note that both allocation and taxes, which are a function of type profile θ can now be represented as functions of the variable m i.e., $\hat{x}_i(\theta) = x_H(m)$ if $\theta_i = \theta_H$, else $\hat{x}_i(\theta) = x_L(m)$. For appropriate comparison with the results in the literature we do not consider the SBB constraints for this example.

Figure 17 depicts two quantities as a function of ϵ. First, $H(\epsilon)$ defined in (5.12) in Condition (A_2) in Section 4. We find that $H(\epsilon) > 0$ for all $\epsilon \leq 0.025$. For this specific example, since the types are binary and agents are symmetric, the BSIC constraints imposed on the taxes (as shown in (5.37)) simplify to a single inequality of the form $A \geq 0$. The second quantity plotted in Figure 17 is A. From this plot, we ascertain that taxes satisfying BSIC exist

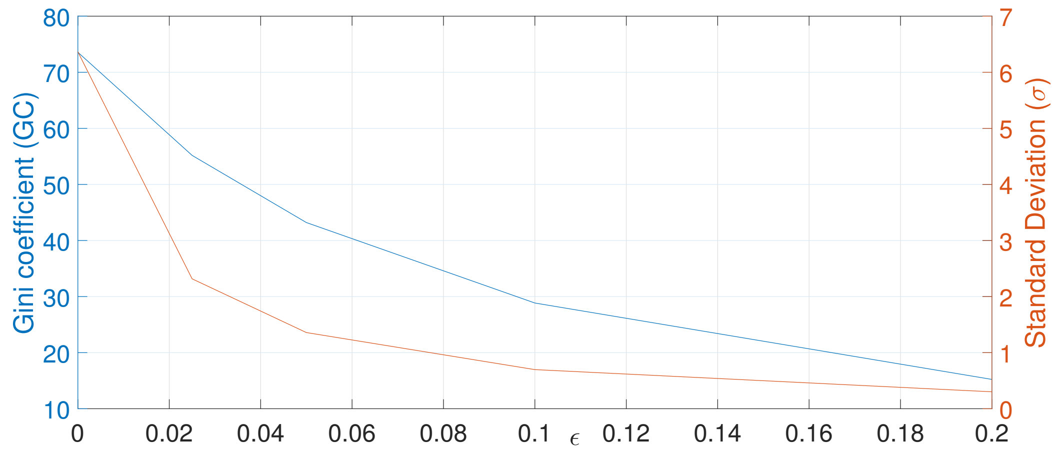

Figure 18.: Gini Coefficient (%) and Standard Deviation of $\left(v(\hat{x}_\epsilon(\theta);\theta_i)\right)_{i\in\mathcal{N}}$ vs ϵ.

for up to $\epsilon \leq 0.2$. This implies that the sufficient condition (A_2) imposes almost an order of magnitude less ϵ_{max} compared to the one imposed by the looser condition (5.37) required in the proof. We make two comments regarding this discrepancy. First, in cases where agents aren't symmetric or the number of types is more than two, one cannot represent feasibility in the simplified form $A \geq 0$ and so one has to rely on sufficient condition (A_2). Second, as will be shown in the subsequent analysis, even this reduced range of ϵ provides significant gains in fairness.

Figure 18 depicts GC and the standard deviation, σ (dual y-axes) of the utility at optimal allocation, $\left(v(\hat{x}_{\epsilon,i}(\theta),\theta_i)\right)_{i\in\mathcal{N}}$ vs. ϵ, for the typical case of $m = 10$. The GC value is 73% at $\epsilon = 0$, 55% at $\epsilon = 0.025$ and 15% at $\epsilon = 0.2$. Recall that GC is a normalized quantity with 100% being extreme case of unfairness and 0% being fairest. This demonstrates that even under the stringent sufficient condition (A_2) we can have reduction of GC by 18 percentage points (and a further reduction by another 40 percentage points if we choose the extreme value of $\epsilon_{max} = 0.2$ dictated by (5.37)). A similar decreasing trend is also evident for the standard deviation, as ϵ increases.

The results in Section 4 guarantee existence of taxes such that BSIC are satisfied. This means that there is a non-empty feasible set of taxes (shown to be a polytope in the proof, Appendix 9.3). In order to fully specify taxes a designer can choose any point within this non-empty set. In the following, we consider two such choices detailed below.

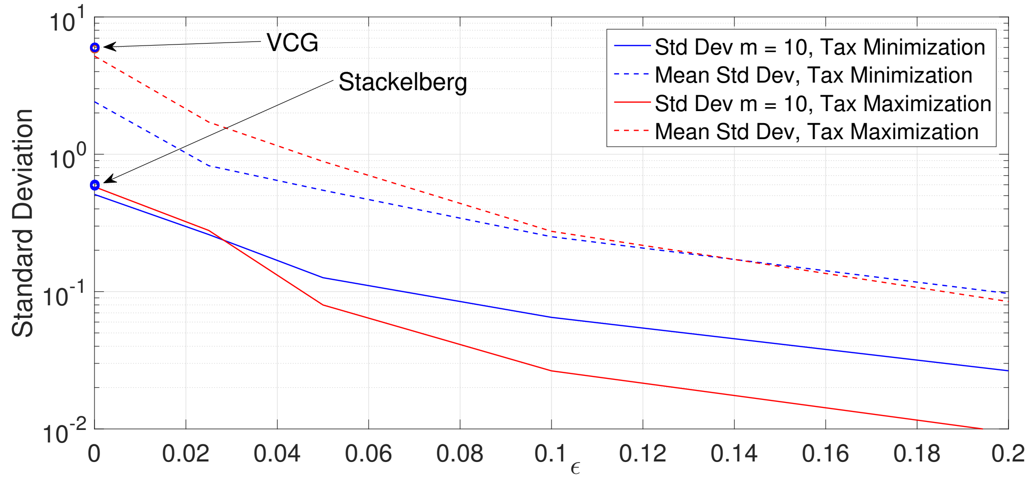

Figure 19.: Standard Deviation of Taxes vs. ϵ, for 10% agents with type θ_H as well as mean with Bin(100,0.1) prior.

- The first one minimizes the expected total tax paid by agents over the non-empty feasible set. In this case, we impose two further practical constraints. First, all taxes are non-negative and second, the expected payoff of the supplier (revenue minus cost) must be positive. If this were not true then the supplier would be better off not participating.

  Figure 19 depicts standard deviation σ in the tax vector $\left(t_1(\theta), \ldots, t_N(\theta)\right)$ with $m = 10$ vs. ϵ, as well as the mean w.r.t. m. For the case of $m = 10$, the σ value is 0.51 at $\epsilon = 0$, 0.26 at $\epsilon = 0.025$ and 0.026 at $\epsilon = 0.2$, and similar trends are observed for the mean σ. This justifies one of the stated goals of this work, which is to lower tax fluctuation between agents.

- The second one maximizes the expected total tax paid by agents (i.e., revenue for the producer). In this case, we impose the non-negative constraint on taxes as above, and a complementary constraint of individual rationality of each user. This ensures that at equilibrium users are better off than not participating.

  Figure 19 also depicts the standard deviation in the tax vector $\left(t_1(\theta), \ldots, t_N(\theta)\right)$ with $m = 10$ vs. ϵ, as well as the mean w.r.t. m, with similar trends observed as above.

The above results can be compared to two well-accepted baseline designs. The first one is the VCG which results in standard deviation of the tax vector equal to 5.9 (typical i.e., $m = 10$)

and equal to 6 (average). As the second baseline design we consider a Stackelberg Game played between the supplier and users [MaWhGr95, Chapter 12, Bertrand Competition]. Supplier plays first by setting the price, agents play next by setting their demand equal to the level that maximizes their own utility (for this to be well-posed we assume minimum demand $m_i = 0$). Anticipating this the supplier sets that price which maximizes his/her expected payoff, revenue minus cost. This is the same as Peak Load Pricing (PLP) scheme mentioned in [SaMoScWo12], applied to a single-shot model. This scheme results in standard deviation of the tax vector equal to 0.59 (typical i.e., $m = 10$) and equal to 0.6 (average).

8. Summary and Comments

Fairness for resource allocation problems in the presence of strategic agents has received relatively less attention due to lack of mathematical convenience in analysis. This chapter proposes a new concept of fairness in resource allocation problems that goes beyond SoU maximization. This fairness aspect is adjustable (through the selection of the parameter ϵ and the function $f(\cdot)$) in a family of functions, thus giving a wide range that a designer may exploit at their own discretion. The main result proves the existence of incentive mechanisms that implement the fairer allocation proposed in dominant and Bayesian equilibria (in respective cases). Analytical and numerical results indicate that through the proposed techniques there are significant gains in fairness of allocation, within the permissible limits of the design method. The theoretical results are justified by discussing a specific model for demand-side management in smart grids. Finally a modification of our mechanism is presented, which guarantees truth-telling as the only BNE. This is done by adding one continuous message per agent, other than his/her type. This modification is especially useful in situations where selection of equilibria is too complex to predict.

Although the form considered in (5.4) is $g_\epsilon(z) = z - \epsilon f(z)$, it is easy to see that the results can be extended to cases where f depends on ϵ; as long as terms of the form $\epsilon f(z, \epsilon) \to 0$ as $\epsilon \to 0$. Ideally one would like to consider the class of $g_\epsilon(z) = z^{1-\epsilon}$, so as to reconcile with known fair social utilities such as the geometric mean and $\min$ utility. This however may not be a practical necessity since as indicated by the results in Sections 6, 7, even the form $g_\epsilon(z) = z - \epsilon z^2$ provides a significant reduction in disparity of allocation, utilities and taxes.

9. Appendix

9.1. Proof of Theorem 3.1

Proof. We begin with the following Lemma:

Lemma 9.1. *For any $\epsilon \geq 0$, there exist taxes $\left(t_i(\phi)\right)_{\phi\in\Theta, i\in\mathcal{N}}$ that satisfy the corresponding DSIC (which implies implementation in dominant strategies) and IR conditions if and only if, $\forall\ \theta_{-i} \in \Theta_{-i},\ i \in \mathcal{N}$,*

$$v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^H) + v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^L) - v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^L) - v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^H) \geq 0. \tag{5.22}$$

Proof. Please see Appendix 9.2 below. □

From the definition in (5.7), we have

$$\begin{aligned} &v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_i) - \epsilon f\left(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_i)\right) + \sum_{j\neq i} v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j) - \epsilon f\left(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j)\right) \\ &= v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_i) - \epsilon f\left(v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_i)\right) + \sum_{j\neq i} v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j) - \epsilon f\left(v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j)\right) \\ &\qquad + K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon). \end{aligned} \tag{5.23}$$

Using the above twice, first with $(\theta_i, \phi_i) = (\theta_i^H, \theta_i^L), \theta_{-i} = \theta_{-i}$ and then with $(\theta_i, \phi_i) = (\theta_i^L, \theta_i^H), \theta_{-i} = \theta_{-i}$, and adding the two results in (using the notation from proof of Theorem 3.1)

$$\begin{aligned} &A_{HH} - \epsilon f(A_{HH}) + A_{LL} - \epsilon f(A_{LL}) \\ &= K_i(\theta_i^H, \theta_i^L, \theta_{-i}, \epsilon) + K_i(\theta_i^L, \theta_i^H, \theta_{-i}, \epsilon) + A_{HL} - \epsilon f(A_{HL}) + A_{LH} - \epsilon f(A_{LH}). \end{aligned} \tag{5.24}$$

This can be rewritten as

$$A_{HH} + A_{LL} - A_{HL} - A_{LH}$$

$$= K_i(\theta_i^H, \theta_i^L, \theta_{-i}, \epsilon) + K_i(\theta_i^L, \theta_i^H, \theta_{-i}, \epsilon) + \epsilon\Big(f(A_{HH}) + f(A_{LL}) - f(A_{HL}) - f(A_{LH})\Big). \tag{5.25}$$

Thus it is sufficient to prove that the RHS above is non-negative. Owing to Condition (A_1), $\exists\, \epsilon_{max} > 0$ such that the sum of the first two terms in RHS is strictly positive for all $0 \leq \epsilon < \epsilon_{max}$ and clearly the second term can be made arbitrarily small in magnitude (by choosing a smaller ϵ_{max}). Hence the condition in (5.22) is satisfied[10] for all $0 \leq \epsilon < \tilde{\epsilon}_{max}$. The value $\tilde{\epsilon}_{max}$ is bigger or smaller than ϵ_{max} depending on whether $f(A_{HH}) + f(A_{LL}) - f(A_{HL}) - f(A_{LH})$ is positive or negative in the range $(0, \epsilon_{max})$. □

9.2. Proof of Lemma 9.1

Proof. Note that the taxes are a finite collection of variables, since $\Theta, \mathcal{N}$ are both finite sets. For the DSIC constraints to be satisfied, the following constraints must hold $\forall\; i \in \mathcal{N}$, $\forall\; \theta_{-i} \in \Theta_{-i}$,

$$v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^H) - t_i(\theta_i^H, \theta_{-i}) \geq v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^H) - t_i(\theta_i^L, \theta_{-i}), \tag{5.26a}$$

$$v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^L) - t_i(\theta_i^L, \theta_{-i}) \geq v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^L) - t_i(\theta_i^H, \theta_{-i}). \tag{5.26b}$$

This gives truth-telling as a dominant strategy for agent i regardless of types of others. For IR, the following constraints must be satisfied $\forall\; i \in \mathcal{N}$, $\forall\; \theta_{-i} \in \Theta_{-i}$,

$$t_i(\theta_i^H, \theta_{-i}) \leq v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^H), \tag{5.27a}$$

$$t_i(\theta_i^L, \theta_{-i}) \leq v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^L). \tag{5.27b}$$

From the above sets of constraints it is clear that one can design $\Big(t_i(\theta_i^H, \theta_{-i}), t_i(\theta_i^L, \theta_{-i})\Big)$ separately for each $i \in \mathcal{N}$, $\theta_{-i} \in \Theta_{-i}$. So for any i, θ_{-i}, the constraints can be rewritten in

[10] Overall the behaviour of $K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon)$ w.r.t. ϵ will dictate the value of $\epsilon_{max}, \tilde{\epsilon}_{max}$. This in turn will effect the usefulness of this method, since a designer might want to ensure certain minimum gains in fairness for which he/she might want to choose ϵ as large as possible.

the form

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} t_1 \\ t_2 \end{bmatrix} \leq \begin{bmatrix} A_{HH} - A_{LH} \\ A_{LL} - A_{HL} \\ A_{HH} \\ A_{LL} \end{bmatrix} \tag{5.28}$$

where $(t_1, t_2) = \Big(t_i(\theta_i^H, \theta_{-i}), t_i(\theta_i^L, \theta_{-i})\Big)$ and

$$A_{HH} = v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^H), \quad A_{LL} = v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^L), \tag{5.29}$$

$$A_{HL} = v(\hat{x}_\epsilon(\theta_i^H, \theta_{-i}); \theta_i^L), \quad A_{LH} = v(\hat{x}_\epsilon(\theta_i^L, \theta_{-i}); \theta_i^H). \tag{5.30}$$

Using the Farkas Lemma, the above system is feasible in t if and only if $\forall\ \lambda \in \mathbb{R}_+^4$

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}^\top \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \end{bmatrix} = 0 \ \Rightarrow \ \begin{bmatrix} A_{HH} - A_{LH} \\ A_{LL} - A_{HL} \\ A_{HH} \\ A_{LL} \end{bmatrix}^\top \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \end{bmatrix} \geq 0. \tag{5.31}$$

The equality constraints on λ give that $\lambda_1 + \lambda_3 = \lambda_2$ and $\lambda_2 + \lambda_4 = \lambda_1$. These give $\lambda_3 = \lambda_4 = 0$. Hence $\lambda \in \mathbb{R}_+^4$ can be parametrized as $\lambda = (\xi, \xi, 0, 0)^\top$ for $\xi \in \mathbb{R}_+$. Thus for feasibility, using Farkas Lemma, the condition that must be satisfied is

$$\Big(A_{HH} - A_{LH}, A_{LL} - A_{HL}, A_{HH}, A_{LL}\Big) \cdot (\xi, \xi, 0, 0) \geq 0 \tag{5.32}$$

$$\Leftrightarrow \quad \xi\Big(A_{HH} + A_{LL} - A_{HL} - A_{LH}\Big) \geq 0 \tag{5.33}$$

$$\Leftrightarrow \quad A_{HH} + A_{LL} - A_{HL} - A_{LH} \geq 0. \tag{5.34}$$

□

9.3. Proof of Theorem 4.1

Proof. The utility for any agent i when other agents are truth-telling is

$$\tilde{u}_i(\phi_i \mid \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \Big[v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_i) - t_i(\phi_i, \theta_{-i}) \Big] \tag{5.35}$$

where agent i's true and quoted types are $\theta_i, \phi_i \in \Theta_i$, respectively. We consider taxes in the AGV form [dACrGé90],

$$t_i(\psi) = z_i(\psi) - \frac{1}{N-1} \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} z_j(\psi) \quad \forall\, i \in \mathcal{N},\ \psi \in \Theta. \tag{5.36}$$

Taxes in this form always satisfy SBB and any tax function which satisfies SBB can be written in this form. Therefore WLOG, the design variables henceforth are $\big(z_j(\psi)\big)_{\substack{j \in \mathcal{N} \\ \psi \in \Theta}}$.

BSIC constraints can be written as: $\forall\, i \in \mathcal{N},\ (\theta_i, \phi_i) \in \mathcal{F}_i$,

$$\tilde{u}_i(\theta_i \mid \theta_i) \geq \tilde{u}_i(\phi_i \mid \theta_i). \tag{5.37a}$$

$$\begin{aligned} \Leftrightarrow \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \Big[& z_i(\theta_i, \theta_{-i}) - \frac{1}{N-1} \sum_{j \neq i} z_j(\theta_i, \theta_{-i}) - z_i(\phi_i, \theta_{-i}) \\ & + \frac{1}{N-1} \sum_{j \neq i} z_j(\phi_i, \theta_{-i}) \Big] \\ \leq \sum_{\theta_{-i} \in \Theta_{-i}} & p_i(\theta_{-i} \mid \theta_i) \Big[v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_i) - v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_i) \Big]. \end{aligned} \tag{5.37b}$$

So the condition for design of taxes is a linear system of inequalities and can be written in the form $Az \leq b$ where the indexing is as follows: $z = \Big(z_j(\psi)\Big) \in \mathbb{R}^{D_1}$, $b = \Big(b(i, \theta_i, \phi_i)\Big) \in \mathbb{R}^{D_2}$ and $A = \Big(A(i, \theta_i, \phi_i \mid j, \psi)\Big) \in \mathbb{R}^{D_2 \times D_1}$, for $D_1 = N \prod_{j=1}^N L_j$ and $D_2 = \prod_{i=1}^N (L_i^2 - L_i)$. The rest of the proof will be to show that this linear system is feasible in variable z, using the *Farkas Lemma*[11] [Ro70, pg.201].

[11] Relevant version of the Farkas alternative result: The system

$$x \in \mathbb{R}^N, \quad Ax \leq b; \quad A \in \mathbb{R}^{M \times N}, b \in \mathbb{R}^M \tag{5.38}$$

is feasible iff all solutions $\lambda \in \mathbb{R}_+^M$ of $A^\top \lambda = 0$ satisfy $b^\top \lambda \geq 0$.

Consider any $\lambda = \Big(\lambda_i(\theta_i, \phi_i)\Big) \in \Lambda \triangleq \{\lambda_k(\theta_k, \phi_k) \in \mathbb{R}_+ \mid k \in \mathcal{N}, (\theta_k, \phi_k) \in \mathcal{F}_k\}$ that satisfies $A^\top \lambda = 0$, i.e., $\forall\, j \in \mathcal{N},\ \psi \in \Theta$,

$$\sum_{i \in \mathcal{N}} \sum_{(\theta_i, \phi_i) \in \mathcal{F}_i} A(i, \theta_i, \phi_i \mid j, \psi)\, \lambda_i(\theta_i, \phi_i) = 0$$
$$\Leftrightarrow \quad p_j(\psi_{-j} \mid \psi_j) \sum_{\phi_j \neq \psi_j} \lambda_j(\psi_j, \phi_j) - \sum_{\theta_j \neq \psi_j} p_j(\psi_{-j} \mid \theta_j)\, \lambda_j(\theta_j, \psi_j)$$
$$- \frac{1}{N-1} \sum_{k \neq j} \Big[p_k(\psi_{-k} \mid \psi_k) \sum_{\phi_k \neq \psi_k} \lambda_k(\psi_k, \phi_k) - \sum_{\theta_k \neq \psi_k} p_k(\psi_{-k} \mid \theta_k)\, \lambda_k(\theta_k, \psi_k) \Big] = 0. \tag{5.39}$$

The above equation can be rearranged to give that $\forall\, j \in \mathcal{N},\ \psi \in \Theta$,

$$p_j(\psi_{-j} \mid \psi_j) \sum_{\phi_j \neq \psi_j} \lambda_j(\psi_j, \phi_j) - \sum_{\theta_j \neq \psi_j} p_j(\psi_{-j} \mid \theta_j)\, \lambda_j(\theta_j, \psi_j)$$
$$= \frac{1}{N} \sum_{k \in \mathcal{N}} \Big[p_k(\psi_{-k} \mid \psi_k) \sum_{\phi_k \neq \psi_k} \lambda_k(\psi_k, \phi_k) - \sum_{\theta_k \neq \psi_k} p_k(\psi_{-k} \mid \theta_k)\, \lambda_k(\theta_k, \psi_k) \Big]. \tag{5.40}$$

Denote the RHS above by $R(\psi)$ and note that it depends only on ψ and not j.

Lemma 9.2. *If* $\lambda \in \Lambda$ *satisfies* (5.40) *then for any* $j \in \mathcal{N},\ \psi_j \in \Theta_j$,

$$\sum_{\substack{\theta_j \in \Theta_j \\ \theta_j \neq \psi_j}} \lambda_j(\theta_j, \psi_j) = \sum_{\substack{\phi_j \in \Theta_j \\ \phi_j \neq \psi_j}} \lambda_j(\psi_j, \phi_j). \tag{5.41}$$

Proof.

$$\sum_{\substack{\theta_j \in \Theta_j \\ \theta_j \neq \psi_j}} \lambda_j(\theta_j, \psi_j) - \sum_{\substack{\phi_j \in \Theta_j \\ \phi_j \neq \psi_j}} \lambda_j(\psi_j, \phi_j) \tag{5.42a}$$
$$= \sum_{\psi_{-j}} \Big[p_j(\psi_{-j} \mid \psi_j) \sum_{\theta_j \neq \psi_j} \lambda_j(\theta_j, \psi_j) - \sum_{\phi_j \neq \psi_j} p_j(\psi_{-j} \mid \phi_j)\, \lambda_j(\psi_j, \phi_j) \Big] \tag{5.42b}$$
$$= \sum_{\psi_{-j}} R(\psi_j, \psi_{-j}) \tag{5.42c}$$
$$= \sum_{\psi_{-j}} \Big[p_k(\psi_{-k} \mid \psi_k) \sum_{\theta_k \neq \psi_k} \lambda_k(\theta_k, \psi_k) - \sum_{\phi_k \neq \psi_k} p_k(\psi_{-k} \mid \phi_k)\, \lambda_k(\psi_k, \phi_k) \Big], \quad k \neq j \tag{5.42d}$$
$$= 0. \tag{5.42e}$$

Here (5.42c), (5.42d) follows by application of (5.40) and other equations are just by rearranging summation terms. $\square$

With the application of above Lemma, one can rewrite the LHS of (5.40) to get that $\forall\, j, \psi$,

$$p_j(\psi_{-j} \mid \psi_j) \sum_{\phi_j \neq \psi_j} \lambda_j(\phi_j, \psi_j) - \sum_{\theta_j \neq \psi_j} p_j(\psi_{-j} \mid \theta_j)\, \lambda_j(\theta_j, \psi_j) = R(\psi). \tag{5.43}$$

Condition (B) on priors states that there exist no $\lambda \in \Lambda$ such that above holds for a non-zero R. Hence $A^\top \lambda = 0$ implies that $R \equiv 0$, therefore (by (5.40)) $\forall\, j \in \mathcal{N}, \forall\, \psi \in \Theta$,

$$p_j(\psi_{-j} \mid \psi_j) \sum_{\phi_j \neq \psi_j} \lambda_j(\psi_j, \phi_j) = \sum_{\theta_j \neq \psi_j} p_j(\psi_{-j} \mid \theta_j)\, \lambda_j(\theta_j, \psi_j). \tag{5.44}$$

Next we show that for all $\lambda \in \Lambda$ that satisfy (5.44) we have $b^\top \lambda \geq 0$ i.e.,

$$\sum_{i \in \mathcal{N}} \sum_{(\theta_i, \phi_i) \in \mathcal{F}_i} b(i, \theta_i, \psi_i)\, \lambda_i(\theta_i, \phi_i) \geq 0 \tag{5.45a}$$

$$\Leftrightarrow \sum_{\substack{i \in \mathcal{N}, \\ (\theta_i, \phi_i) \in \mathcal{F}_i}} \lambda_i(\theta_i, \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \cdot \Big[v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_i) - v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_i) \Big] \geq 0. \tag{5.45b}$$

Proving this will finish the proof by Farkas Lemma.

For any $i, \theta_i, \phi_i, \theta_{-i}$, denote $x = \hat{x}_\epsilon(\theta_i, \theta_{-i})$, $\tilde{x} = \hat{x}_\epsilon(\phi_i, \theta_{-i})$. Rearranging terms from (5.7), gives

$$\begin{aligned} & v(x; \theta_i) - v(\tilde{x}; \theta_i) \\ & \quad = \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} \Big(v(\tilde{x}; \theta_j) - v(x; \theta_j) \Big) + \epsilon \sum_{j \in \mathcal{N}} \Big(f(v(x; \theta_i)) - f(v(\tilde{x}; \theta_i)) \Big) + K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon). \end{aligned} \tag{5.46}$$

Denote the RHS expression in (5.46) as $\eta = \eta_1 + \eta_2$, where

$$\eta_1 = \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} \Big(v(\tilde{x}; \theta_j) - v(x; \theta_j) \Big), \tag{5.47a}$$

$$\eta_2 = \epsilon \sum_{j \in \mathcal{N}} \Big(f(v(x;\theta_j)) - f(v(\tilde{x};\theta_j)) \Big) + K_i(\theta_i, \phi_i, \theta_{-i}, \epsilon). \tag{5.47b}$$

Now continuing from the LHS of (5.45b),

$$\text{LHS of (5.45b)} = \sum_{\substack{i \in \mathcal{N} \\ (\theta_i,\phi_i) \in \mathcal{F}_i}} \lambda_i(\theta_i, \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i)(\eta_1 + \eta_2). \tag{5.48}$$

For any fixed i, consider the summation with only η_1 first

$$\begin{aligned}
\sum_{(\theta_i,\phi_i) \in \mathcal{F}_i} \lambda_i(\theta_i, \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} \Big(v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j) - v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j) \Big) \\
= \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} \sum_{\substack{\phi_i \in \Theta_i, \\ \theta_{-i} \in \Theta_{-i}}} v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j) \sum_{\theta_i \neq \phi_i} p_i(\theta_{-i} \mid \theta_i) \lambda_i(\theta_i, \phi_i) \\
- \sum_{\substack{j \in \mathcal{N} \\ j \neq i}} \sum_{\substack{(\theta_i,\phi_i) \in \mathcal{F}_i \\ \theta_{-i} \in \Theta_{-i}}} \lambda_i(\theta_i, \phi_i)\, p_i(\theta_{-i} \mid \theta_i)\, v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j).
\end{aligned} \tag{5.49}$$

By (5.44), the inside summation in the first term is equal to $p_i(\theta_{-i} \mid \phi_i) \sum_{\psi_i} \lambda(\phi_i, \psi_i)$. Incorporating this and changing variables of summation appropriately gives the overall summation from (5.49) equal to 0. Now consider the term in the RHS of (5.48) with η_2

$$\sum_{\substack{i \in \mathcal{N} \\ (\theta_i,\phi_i) \in \mathcal{F}_i}} \lambda_i(\theta_i, \phi_i) \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \cdot \eta_2. \tag{5.50}$$

Rearranging terms in η_2, we can write

$$\eta_2 = \sum_{j \in \mathcal{N}} \Big(v(\hat{x}_\epsilon(\theta_i, \theta_{-i}); \theta_j) - v(\hat{x}_\epsilon(\phi_i, \theta_{-i}); \theta_j) \Big). \tag{5.51}$$

Therefore by Condition (A_2), $\exists\, \epsilon_{max} > 0$ such that for all $0 \leq \epsilon < \epsilon_{max}$, the inside summation in (5.50) is non-negative $\forall\, i \in \mathcal{N}$, $(\theta_i, \phi_i) \in \mathcal{F}_i$. This finishes the proof by Farkas Lemma, since the expression in (5.45a) is now shown to be non-negative for all positive $0 \leq \epsilon < \epsilon_{max}$. □

9.4. Proof of Part (1), Theorem 5.1

Proof. For strategy

$$(\sigma,\rho)=\Big(\big(\sigma_i:\Theta_i\to\Theta_i\big),\ \big(\rho_i:\Theta_i\to[-\delta,+\delta]\big)\Big)_{i\in\mathcal{N}},\tag{5.52}$$

utility is

$$\begin{aligned}u_i(\sigma,\rho\mid\theta_i)=&\sum_{\theta_{-i}}p_i(\theta_{-i}\mid\theta_i)\Big[v\Big(\hat{x}_i(\sigma(\theta))+\rho_i(\theta_i);\theta_i\Big)-\hat{t}_i\Big(\sigma(\theta)\Big)\\&-\rho_i(\theta_i)v'\Big(\hat{x}_{\epsilon,i}(\sigma(\theta));\sigma_i(\theta_i)\Big)\Big]\end{aligned}\tag{5.53}$$

First in this proof we establish that for any agent i, true type θ_i, when other agents use strategy $m^*_{-i}=(\theta_{-i},\underline{0})$, strategy $\sigma_i(\theta_i)=\theta_i$ gives higher utility than $\sigma_i(\theta_i)\neq\theta_i$ at any value of $\rho_i(\theta_i)$. For this start by considering the utility in (5.53) with $\sigma_{-i}(\theta_{-i})=\theta_{-i}$ and $\rho_i(\theta_i)=0$,

$$\tilde{u}_i(\sigma_i(\theta_i)\mid\theta_i)=\sum_{\theta_{-i}}p_i(\theta_{-i}\mid\theta_i)\Big[v\Big(\hat{x}_{\epsilon,i}(\sigma_i(\theta_i),\theta_{-i});\theta_i\Big)-\hat{t}_i\Big(\sigma_i(\theta_i),\theta_{-i}\Big)\Big].\tag{5.54}$$

From the result in Section 4, allocation and taxes $(\hat{x}_\epsilon,\hat{t})$ satisfy BSIC condition. Hence the utility in (5.54) is maximized at $\sigma_i(\theta_i)=\theta_i$. The minimum difference between the value of such an expression at $\sigma_i(\theta_i)=\theta_i$ and $\sigma_i(\theta_i)\neq\theta_i$ is strictly positive since the tax function $\hat{t}$ satisfies BSIC with strict inequality. Now the utility u_i in (5.53) clearly depends continuously on $\rho_i(\theta_i)$, hence by choosing a small enough range for $\rho_i(\theta_i)$ (i.e., by choosing a small enough δ) it can be ensured that even in (5.53), for $\sigma_{-i}(\theta_{-i})=\theta_{-i}$, u_i is maximzed at $\sigma_i(\theta_i)=\theta_i$.

Now the only thing that remains to be proven is that when other agents quote message m^*_{-i} and $\sigma_i(\theta_i)=\theta_i$, then $\rho_i(\theta_i)=0$ is optimal. For agent i and true type θ_i, the utility in (5.53) only depends on ρ_i through $\rho_i(\theta_i)$. Optimizing w.r.t. $\rho_i(\theta_i)$ (for calculating equilibrium) gives

$$\frac{\partial u_i}{\partial\rho_i(\theta_i)}=\sum_{\theta_{-i}\in\Theta_{-i}}p_i(\theta_{-i}\mid\theta_i)\Big[v'(\hat{x}_{\epsilon,i}(\theta)+\rho_i(\theta_i);\theta_i)-v'(\hat{x}_{\epsilon,i}(\theta);\sigma_i(\theta_i))\Big]=0.\tag{5.55}$$

At equilibrium, the above derivative will have to be zero since the end points $\pm\delta$ cannot be maximizers - since they incur a huge tax B. Also note that since $v(x_i;\theta_i)$ is assumed to be concave in x_i (for any θ_i) the utility $u_i(\sigma, \rho|\theta_i)$ in (5.53) is concave in $\rho_i(\theta_i)$ and thus the condition in (5.55) is both necessary and sufficient for optimality w.r.t. $\rho_i(\theta_i)$.

The expression in (5.55) is clearly equal to 0 when $\sigma_i(\theta_i) = \theta_i$ and $\rho_i(\theta_i) = 0$. Hence applying the above result for all $\theta_i \in \Theta_i$ and $i \in \mathcal{N}$ gives that $m^* = (\theta, \underline{0})$ is a BNE. □

9.5. Proof of Part (2), Theorem 5.1

Proof. Consider the derivative similar to (5.55) but with general strategies

$$\frac{\partial u_i}{\partial \rho_i(\theta_i)} = \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \Big[v'(\hat{x}_{\epsilon,i}(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})) + \rho_i(\theta_i); \theta_i) \\ - v'(\hat{x}_{\epsilon,i}(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})); \sigma_i(\theta_i)) \Big]. \quad (5.56)$$

This being zero is a necessary at BNE. Now we begin by considering the same with $\rho_i(\theta_i) = 0$,

$$\Psi = \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \Big[v'(\hat{x}_{\epsilon,i}(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})); \theta_i) - v'(\hat{x}_{\epsilon,i}(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})); \sigma_i(\theta_i)) \Big]. \quad (5.57)$$

By the assumption in the theorem statement, for $\sigma_i(\theta_i) \neq \theta_i$, the term inside the brackets above is either positive for all θ_{-i} or negative for all θ_{-i}. Hence the expression Ψ cannot be zero for any i, θ_i when $\sigma_i(\theta_i) \neq \theta_i$. Since only discrete variables are involved in the expression Ψ, it means that all the values that Ψ can take (i.e., $\forall\ i, \theta_i, \sigma$ with $\sigma_i(\theta_i) \neq \theta_i$) are a discrete set (say, D) of points not containing zero.

The only difference between the expression in (5.56) and Ψ is the introduction of the variable $\rho_i(\theta_i) \in [-\delta, +\delta]$ - on which the expression in (5.56) depends continuously. This means that the set of all values (say, E) taken by $\frac{\partial u_i}{\partial \rho_i(\theta_i)}$ is the union of continuous sets centered at the points from set D. The range of the continuous sets around each point in D can be continuously controlled by changing δ with the sets becoming a discrete point as $\delta \to 0$. Hence it is clear that one can find δ small enough so that even this new set E doesn't contain zero. Once such a δ is chosen it is clear that any strategy σ for which $\exists\ i, \theta_i$ such that $\sigma_i(\theta_i) \neq \theta_i$ cannot be a BNE since the expression in (5.56) being zero was a necessary condition for a BNE.

Note that each part of the proof of Theorem 5.1 prescribes a positive upper limit for δ. The designer will eventually choose a δ that satisfies both.

Now consider a strategy of the form $m_i^* = (\theta_i, y_i)$ where there is truth-telling but $y_i \neq 0$. It is clear that this cannot be a BNE, since by strict concavity of $v(x_i; \theta_i)$, at $\sigma_i(\theta_i) = \theta_i$ the expression in (5.56) goes to zero only for $\rho_i(\theta_i) = y_i = 0$. $\square$

Part II.

Dynamic Games

Chapter 6.

Structured Perfect Bayesian Equilibrium in Infinite Horizon Dynamic Games with Asymmetric Information

In dynamic games with asymmetric information structure, the widely used concept of equilibrium is perfect Bayesian equilibrium (PBE). This is expressed as a strategy and belief pair that simultaneously satisfy sequential rationality and belief consistency. Unlike symmetric information dynamic games, where subgame perfect equilibrium (SPE) is the natural equilibrium concept, to date there does not exist a universal algorithm that decouples the interdependence of strategies and beliefs over time in calculating PBE. In this chapter we find a subset of PBE for an infinite horizon discounted reward asymmetric information dynamic game. We refer to it as Structured PBE or SPBE; in SPBE, any agents' strategy depends on the public history only through a common public belief and on private history only through the respective agents' latest private information (his private type). The public belief acts as a summary of all the relevant past information and it's dimension does not increase with time. The motivation for this comes the common information approach proposed in Nayyar et al. (2013) for solving decentralized team (non-strategic) resource allocation problems with asymmetric information and Vasal et al. (2016) where SPBE for finite horizon games are characterized. We calculate SPBE by solving a single-shot fixed-point equation and a corresponding forward recursive algorithm. The class of SPBE contain equilibria that show *signaling* behavior and we demonstrate this (and our methodology) by means of a public goods investment example.

Finally, analytical and numerical tools for checking the existence of SPBEs is discussed.

1. Introduction

Dynamic games with symmetric information among agents has been studied well in game theory literature [FuTi91; BaOl99; MaSa06; KrZa14]. The appropriate equilibrium concept is subgame perfect equilibrium (SPE) which is a refinement of Nash equilibrium. A strategy profile is SPE if its restriction to any subgame of the original game is also an SPE of the subgame. However there are models of interest where strategic agents interact repeatedly whilst observing private signals. For instance, relay scheduling with private queue length information at different relays [VaAn14]. Another example is that of Bayesian learning games [BiHiWe92; AcDaLoOz11; LeSuBe14; VaAn16b] such as when customers post reviews on websites such as Amazon; clearly agents posses their own experience of the object they bought and this constitutes private information. These examples lead to a model with asymmetric information, where the notion of SPE is ineffective.

Appropriate equilibrium concepts in such a case consist of strategy profiles and beliefs. These include weak perfect Bayesian equilibrium, perfect Bayesian equilibrium (PBE) and Sequential equilibrium. PBE is the most commonly used notion. The requirement of PBE is sequential rationality of strategy as well as belief consistency. Due to the fact that future beliefs depend on past strategy and strategies are sequentially rational based on specific beliefs, finding a pair of strategy and belief that satisfy PBE requirement is a difficult problem and usually reduces to a large fixed-point equation over the entire time horizon. To date, there is no universal algorithm that provides simplification by decomposing the aforementioned fixed-point equation for calculating PBEs.

Our motivation stems from the work of Nayyar et al. [NaMaTe13]. Authors in [NaMaTe13] consider a decentralized dynamic team problem (non-strategic agents) with asymmetric information. Each agents observes a private type which evolves with time as a controlled Markov process. They introduce a common information based approach, whereby each agent calculates a belief on every agents' current private type. This belief is defined in such a manner that despite having asymmetric information, agents can agree on it. They proceed to show optimality of policies that depend on private history only through this belief and respective agents' current private type.

The common information based approach has been studied for finite horizon dynamic games with asymmetric information in [GuNaLaBa14; NaGuLaBa14; VaAn16a; VaAn15; VaAn16b]. Of these, [GuNaLaBa14; NaGuLaBa14] consider models where the aforementioned belief from the common information approach is updated independent of strategy. This implies that the simultaneous requirements of sequential rationality and belief consistency can be decoupled, resulting in a simplification in calculating PBEs. However these models produce non-signaling PBEs. For the models studied in [VaAn16a; VaAn15; VaAn16b] belief updates depend on strategies. Consequently the resulting PBEs are signaling equilibria. The problem of finding them though, becomes more complicated.

Other than the common information based approach, Li et al. [LiSh14] consider a finite horizon zero-sum dynamic game, where at each time only one agent out of the two knows the state of the system. The value of the game is calculated by formulating an appropriate linear program.

All works listed above for dynamic games consider a finite horizon. In this chapter, we deal with an infinite horizon discounted reward dynamic game. Cole et al. [CoKo01] consider an infinite horizon discounted reward dynamic game where actions are only privately observable. They provide a fixed-point equation for calculating a subset of sequential equilibrium, which is referred to as Markov private equilibrium (MPrE). In MPrE strategies depend on history only through the latest private observation.

Contributions

In this work we follow the models of [VaAn16a; VaAn15; VaAn16b] where beliefs are updated dependent on strategy and thus signaling PBEs are considered.

This chapter provides a one-shot fixed-point equation and a forward recursive algorithm that together generate a subset of perfect Bayesian equilibria for the infinite horizon discounted reward dynamic game with asymmetric information. The common information based approach of [NaMaTe13] is used and the strategies are such that they depend on the public history only through the common belief (which summarizes the relevant information) and on private history only through respective agents' current private type. The PBE thus generated are referred to as Structured PBE or SPBE and posses the property that the mapping between belief and strategy is stationary (w.r.t. time).

The provided methodology (consisting of solving a one-shot fixed-point equation) provides a decomposition of the interdependence between belief and strategy in PBEs and enables a systematic evaluation of SPBEs.

The model allows for signaling amongst agents as beliefs depend on strategies. This is because the source of asymmetry of information in our model is that each agent has a privately observed type that affects the utility function of all the agents. These types evolve as a controlled Markov process and depend on actions of all agents (the actions are assumed to be common knowledge).

Our methodology for proving our results is as follows: first we extend the finite horizon model of [VaAn16a] to include belief-based terminal reward. This does not change the proofs significantly but helps us simplify the exposition and proofs of the infinite horizon case. In the next step, the infinite horizon results are obtained with the help of the finite horizon ones through continuity arguments as horizon $T \to \infty$.

The remainder of this chapter is organized as follows: Section 2 introduces the model for the dynamic game with private types and discounted reward. We consider two versions of the problem, finite and infinite horizon. Section 3 defines the backward recursive algorithm for calculating SPBEs in the finite horizon model. The results in this section are an adaptation of the finite horizon results from [VaAn16a]. This is done so that the finite horizon results can be appropriately used in Section 4. Section 4 contains the central result of this chapter, namely the fixed-point equation that defines the set of SPBE strategy for the infinite horizon game. Section 5 discusses a concrete example of a multi-stage public goods game with two agents. This example is an infinite horizon version of the finite horizon public goods example from [FuTi91, ch. 8, Example 8.3]. Finally, Section 6 discusses the question of existence of SPBEs and provides analytical as well as numerical techniques for checking the same.

2. Model for Dynamic Game

We consider a dynamical system with N strategic agents, denoted by the set $\mathcal{N}$. Associated with each agent i at any time $t \geq 1$, is a type $x_t^i \in \mathcal{X}^i$. The set of types is denoted by $\mathcal{X}^i$ and is assumed to be finite. We assume that each agent i can observe their own type x_t^i but not that of others. Denote the profile of types at time t by $x_t = (x_t^i)_{i \in \mathcal{N}}$.

Each agent i at any time $t \geq 1$, after observing their private type x_t^i takes action $a_t^i \in \mathcal{A}^i$. The set of available actions to agent i is denoted by $\mathcal{A}^i$ and is assumed to be finite. It is assumed that the action profile $a_t = (a_t^i)_{i \in \mathcal{N}}$ at time t is publicly observed and is common knowledge.

The type of each agent evolves over time as a controlled Markov process independent of other agents given the action profile a_t. We assume a time-homogeneous kernel Q^i for the evolution of agent i's type i.e.,

$$\mathbb{P}(x_{t+1} \mid x_{1:t}, a_{1:t}) = \mathbb{P}(x_{t+1} \mid x_t, a_t) = \prod_{i=1}^{N} Q^i(x_{t+1}^i \mid x_t^i, a_t). \tag{6.1}$$

The notation we use is $x_{1:t} = (x_k)_{k=1,\ldots,t}$ and $a_{1:t} = (a_k)_{k=1,\ldots,t}$. Initial belief, at time 1, is assumed as $\mathbb{P}(x_1) = \prod_{i=1}^{N} Q_0^i(x_1^i)$.

Associated with each agent i is an instantaneous reward function R^i that depends on the type and action profiles i.e., reward at time t for agent i is $R^i(x_t, a_t)$. Furthermore, rewards are accumulated over time in a discounted manner with a common discount factor $\delta \in (0, 1)$. This means that player i's objective is to maximize

$$J^{i,g} \triangleq \mathbb{E}^g \left\{ \sum_{t=1}^{\infty} \delta^{t-1} R^i(X_t, A_t) \right\}. \tag{6.2}$$

where $\mathbb{E}^g$ denotes the expectation given strategy profile g (described below) used by all agents. The transition kernels, initial belief and reward functions are assumed to be common knowledge.

We consider two versions of the problem - finite horizon (T) and infinite horizon, where the main result is for the infinite horizon problem. For the finite horizon problem we introduce a terminal reward and in this regard we first define public beliefs $\underline{\pi}_t$ since the terminal reward is defined as a function of these beliefs.

Define the public history h_t^c at time t as the set of publicly observed actions i.e., $h_t^c = a_{1:t-1} \in (\times_{i=1}^{N} \mathcal{A}^i)^{t-1}$. Denote the set of public histories by $\mathcal{H}_t^c$. Define private history h_t^i of any agent i at time t as the set of privately observed types $x_{1:t}^i \in (\mathcal{X}^i)^t$ of agent i and publicly observed action profiles $a_{1:t-1}$ i.e., $h_t^i = (x_{1:t}^i, a_{1:t-1})$. Denote the set of private histories by $\mathcal{H}_t^i$.

Note that private history h_t^i contains public history h_t^c. Also denote the overall history at time t by $h_t = (x_{1:t}, a_{1:t-1})$ and the set of all such histories by $\mathcal{H}_t$. To preserve the above information structure, we also assume that the instantaneous rewards $R_t^i(x_t, a_t)$ are not observed until the end of the game (or else it might reveal partial information about the state x_t).

Any strategy used by agent i can be represented as $g^i = (g_t^i)_{t \geq 1}$ with $g_t^i : \mathcal{H}_t^i \to \Delta(\mathcal{A}^i)$. Here $\Delta(\mathcal{A}^i)$ is the set of all probability distributions over the elements of the set $\mathcal{A}^i$.

The public belief $\underline{\pi}_t \in \times_{j=1}^N \Delta(\mathcal{X}^j)$, at any time t, depends on the strategy profile $g_{1:t-1} = (g_{1:t-1}^j)_{j \in \mathcal{N}}$. Given a strategy profile $g_{1:t-1}$, we define the belief as follows: $\underline{\pi}_t(x_t) = \left(\pi_t^i(x_t^i)\right)_{i \in \mathcal{N}}$, with

$$\pi_t^i(x_t^i) \triangleq \begin{cases} \mathbb{P}^g(x_t^i \mid h_t^c) & \text{if } \mathbb{P}^g(h_t^c) > 0, \\ \sum_{x_{t-1}^i \in \mathcal{X}^i} Q^i(x_t^i \mid x_{t-1}^i, a_{t-1}) \pi_{t-1}^i(x_{t-1}^i) & \text{if } \mathbb{P}^g(h_t^c) = 0. \end{cases} \tag{6.3}$$

Initial belief is $\pi_1^i(x_1^i) = Q_0^i(x_1^i)$. The second sub-case in the above definition implies that for histories with zero probability of occurrence, under a particular strategy profile g, the public belief at time t is an updated version of the public belief at time $t-1$, taking into account the forward transition kernel Q^i and assuming that the distribution over the latest action a_{t-1} does not depend on private type x_{t-1}^i. It ensures that the belief random variable Π_t, at any time t, is measurable, even if the set of histories resulting in any instance π_t has zero measure under strategy g.

The second sub-case in (6.3) is, however, only one of many possible updates than can be used here. Dynamics that govern the evolution of public beliefs at histories with zero probability of occurrence, affect equilibrium strategies. Thus, the construction proposed for calculating perfect Bayesian equilibria in this chapter will produce a different set of equilibria if one changes the second sub-case above. The most well-known example of another such update is the *intuitive criterion* proposed in [ChKr87] for Nash equilibria, later generalized for sequential equilibria in [Ch87]. Intuitive criterion assigns zero probability to states that can be excluded based on data available to all players (in our case action profile history $a_{1:t-1}$). Another example of belief update is *universal divinity*, proposed in [BaSo87].

In this chapter we stick with the update presented in (6.3), however the proofs of all technical lemmas and the main theorems will go through for other updates as well. An interesting direction for future work is to isolate and analyze the effect of using other belief updates on the equilibrium calculation technique proposed in this chapter and thus predict the qualitative properties of the equilibrium strategies calculated.

Finally, in the finite horizon game, for each agent i there is a terminal reward $G^i(\pi_{T+1}, x^i_{T+1})$ that depends on the terminal type of agent i and the terminal belief. It is assumed that $G^i(\cdot)$ is absolutely bounded.

3. Technical Lemmas in the finite horizon model

In this section, we state and prove properties of the finite horizon backwards recursive algorithm for calculating SPBE, adopted from [VaAn16a] with minor modifications to accommodate for belief-based terminal rewards. The results from this section are used in Section 4 to prove the infinite horizon equilibrium result.

To distinguish quantities defined in this section with infinite horizon, we add a superscript T to all the quantities.

3.1. Finite Horizon problem and Backwards Recursion

Consider a finite horizon, $T > 1$, problem in the dynamic game model described in Section 2.

Define the value functions $\left(V_t^{i,T} : \times_{j\in\mathcal{N}}\Delta(\mathcal{X}^j) \times \mathcal{X}^i \to \mathbb{R}\right)_{i\in\mathcal{N},t\in\{1,\ldots,T\}}$ and strategy $\left(\tilde{\gamma}_t^{i,T}\right)_{i\in\mathcal{N},t\in\{1,\ldots,T\}}$ backwards inductively as follows[1].

1. $\forall\, i \in \mathcal{N}$, set $V_{T+1}^{i,T} \equiv G^i$.

2. For any $t \in \{1,\ldots,T\}$ and $\underline{\pi}_t$, solve the following fixed-point equation in $(\tilde{\gamma}_t^{i,T})_{i\in\mathcal{N}}$. $\forall\, i \in \mathcal{N}, x^i \in \mathcal{X}^i$,

$$\tilde{\gamma}_t^{i,T}(\cdot \mid x_t^i) \in \underset{\gamma_t^i(\cdot|x_t^i)\in\Delta(\mathcal{A}^i)}{\operatorname{argmax}} \mathbb{E}^{\gamma_t^i(\cdot|x_t^i),\tilde{\gamma}_t^{-i,T},\underline{\pi}_t^{-i}}\Big[R^i(X_t, A_t) + \delta V_{t+1}^{i,T}\Big(\underline{F}(\underline{\pi}_t, \tilde{\gamma}_t^T, A_t), X_{t+1}^i\Big) \mid \underline{\pi}_t, x_t^i\Big] \quad (6.4)$$

(see below for the various quantities involved in the above expression).

[1]This construction differs from the one proposed in [VaAn16a] only in Step 1., since it has to accommodate the terminal reward.

3. Then define,

$$V_t^{i,T}(\underline{\pi}_t, x_t^i) = \mathbb{E}^{\tilde{\gamma}_t^{i,T}(\cdot|x_t^i),\tilde{\gamma}_t^{-i,T},\underline{\pi}_t^{-i}}\Big[R^i(X_t, A_t) + \delta V_{t+1}^{i,T}\Big(\underline{F}(\underline{\pi}_t, \tilde{\gamma}_t^T, A_t), X_{t+1}^i\Big) \mid \underline{\pi}_t, x_t^i\Big]. \quad (6.5)$$

Note that the strategy $\tilde{\gamma}_t^{i,T}$ generated by above depends on π_t. Denote by θ_t^i the mapping $\underline{\pi}_t \mapsto \tilde{\gamma}_t^{i,T}$ i.e., $\tilde{\gamma}_t^{i,T} = \theta_t^i[\underline{\pi}_t]$.

In (6.4), the expectation is with the following distribution

$$(X_t^{-i}, A_t^i, A_t^{-i}, X_{t+1}^i) \sim \pi_t^{-i}(x_t^{-i})\gamma_t^{i,T}(a_t^i \mid x_t^i)\tilde{\gamma}_t^{-i,T}(a_t^{-i} \mid x_t^{-i})Q^i(x_{t+1}^i \mid x_t^i, a_t) \quad (6.6)$$

and $\underline{F}(\underline{\pi}_t, \tilde{\gamma}_t^T, A_t) = \Big[F(\pi_t^j, \tilde{\gamma}_t^{j,T}, A_t)\Big]_{j=1}^N$, where

$$F(\pi^j, \gamma^j, a)(x'^j) \triangleq \begin{cases} \dfrac{\sum_{x^j \in \mathcal{X}^j} \pi^j(x^j)\gamma^j(a^j \mid x^j)Q^j(x'^j \mid x^j, a)}{\sum_{\tilde{x}^j} \pi^j(\tilde{x}^j)\gamma^j(a^j \mid \tilde{x}^i)} & \text{if Den.} > 0 \\ \sum_{x^j \in \mathcal{X}^j} \pi^j(x^j)Q^j(x'^j \mid x^j, a) & \text{if Den.} = 0 \end{cases} \quad (6.7)$$

Note here that the second sub-case above is consistent with (6.3). Below we define strategy-belief pair (β^*, μ^*),

$$\beta^* = \left(\beta_t^{i,*}\right)_{t\in\{1,\ldots,T\}, i\in\mathcal{N}} \qquad \beta_t^{i,*} : \mathcal{H}_t^i \to \Delta(\mathcal{A}^i) \quad (6.8a)$$

$$\mu^* = \left(\mu_t^{i,*}\right)_{t\in\{1,\ldots,T\}, i\in\mathcal{N}} \qquad \mu_t^{i,*} : \mathcal{H}_t^i \to \Delta(\mathcal{H}_t) \quad (6.8b)$$

based on the mapping θ produced by the above algorithm.

Define belief μ^* inductively as follows: set $\mu_1^*(x_1) = \prod_{i=1}^N Q_0^i(x_1^i)$. For $t \in \{1, \ldots, T\}$,

$$\mu_{t+1}^{i,*}\Big[h_{t+1}^c\Big] = F\Big(\mu_t^{i,*}[h_t^c], \theta_t^i\Big[\underline{\mu}_t^*[h_t^c]\Big], a_t\Big) \quad (6.9a)$$

$$\mu_{t+1}^*\Big[h_{t+1}^c\Big](x_{t+1}) = \prod_{i=1}^N \mu_{t+1}^{i,*}\Big[h_{t+1}^c\Big](x_{t+1}^i) \quad (6.9b)$$

Denote the profile of beliefs by $\underline{\mu}_t^*[h_t^c] = \left(\mu_t^{i,*}[h_t^c]\right)_{i\in\mathcal{N}}$. Denote the strategy arising out of

$\tilde{\gamma}_t^{i,T}$ by $\beta_t^{i,*}$ i.e.,

$$\beta_t^{i,*}(a_t^i \mid h_t^i) = \theta_t^i\Big[\mu_t^*[h_t^c]\Big](a_t^i \mid x_t^i) \tag{6.10}$$

3.2. Finite Horizon Result

In this section, we presented three lemmas. The first two are technical results needed in the proof of the third. The result of the third lemma is used in Section 4.

Define the reward-to-go $W_t^{i,\beta^i,T}$ for any agent i and strategy β^i as

$$W_t^{i,\beta^i,T}(h_t^i) = \mathbb{E}^{\beta^i,\beta^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{T} \delta^{n-t} R^i(X_n, A_n) + \delta^{T+1-t} G^i(\underline{\Pi}_{T+1}, X_{T+1}^i) \mid h_t^i\Big]. \tag{6.11}$$

Here agent i's strategy is β^i whereas all other agents use strategy $\beta^{-i,*}$ defined above. Since $\mathcal{X}^i, \mathcal{A}^i$ are assumed to be finite and G^i absolutely bounded, the reward-to-go is finite $\forall$ i, t, β^i, h_t^i.

Lemma 3.1. *For any $t \in \{1, \dots, T\}$, $i \in \mathcal{N}$, h_t^i and β^i,*

$$V_t^{i,T}(\mu_t^*[h_t^c], x_t^i) \geq \mathbb{E}^{\beta^i,\beta^{-i,*},\mu_t^*[h_t^c]}\Big[R^i(X_t, A_t) + \delta V_{t+1}^{i,T}\Big(\underline{F}(\underline{\mu}_t^*[h_t^c], \beta_t^*, A_t), X_{t+1}^i\Big) \mid h_t^i\Big] \tag{6.12}$$

Proof. Please see Appendix 8.1 at the end of this chapter. □

Lemma 3.2.

$$\begin{aligned}
&\mathbb{E}^{\beta_{t+1:T}^i,\beta_{t+1:T}^{-i,*},\mu_{t+1}^*[h_t^c,a_t]}\Big[\sum_{n=t+1}^{T} \delta^{n-(t+1)} R^i(X_n, A_n) + \delta^{T+1-t} G^i(\underline{\Pi}_{T+1}, X_{T+1}^i) \mid h_t^i, a_t, x_{t+1}^i\Big] \\
&= \mathbb{E}^{\beta_{t:T}^i,\beta_{t:T}^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t+1}^{T} \delta^{n-(t+1)} R^i(X_n, A_n) + \delta^{T+1-t} G^i(\underline{\Pi}_{T+1}, X_{T+1}^i) \mid h_t^i, a_t, x_{t+1}^i\Big]
\end{aligned} \tag{6.13}$$

Proof. Please see Appendix 8.2 at the end of this chapter. □

The result below shows that the value function from the backwards recursive algorithm is higher than any reward-to-go.

Lemma 3.3. *For any* $t \in \{1, \dots, T\}$, $i \in \mathcal{N}$, h_t^i *and* β^i,

$$V_t^{i,T}(\underline{\mu}_t^*[h_t^c], x_t^i) \geq W_t^{i,\beta^i,T}(h_t^i) \tag{6.14}$$

Proof. Please see Appendix 8.3 at the end of this chapter. □

4. SPBE in Infinite Horizon

In this section we consider the infinite horizon dynamic game, with naturally no terminal reward.

4.1. Perfect Bayesian Equilibrium

A perfect Bayesian equilibrium is the pair of strategy and belief (β^*, μ^*), where

$$\beta^* = \left(\beta_t^{i,*}\right)_{t\geq 1, i\in\mathcal{N}} \qquad \beta_t^{i,*} : \mathcal{H}_t^i \to \Delta(\mathcal{A}^i) \tag{6.15a}$$

$$\mu^* = \left(\mu_t^{i,*}\right)_{t\geq 1, i\in\mathcal{N}} \qquad \mu_t^{i,*} : \mathcal{H}_t^i \to \Delta(\mathcal{H}_t) \tag{6.15b}$$

such that sequential rationality is satisfied: $\forall\, i \in \mathcal{N}$, β^i, $t \geq 1$, $h_t^i \in \mathcal{H}_t^i$,

$$\mathbb{E}^{\beta^{i,*},\beta^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{\infty} \delta^{n-t} R^i(X_n, A_n) \mid h_t^i\Big] \;\geq\; \mathbb{E}^{\beta^{i},\beta^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{\infty} \delta^{n-t} R^i(X_n, A_n) \mid h_t^i\Big] \tag{6.16}$$

and beliefs satisfy certain consistency conditions (please refer [FuTi91, pp. 331] for the exact conditions).

Definition 4.1 (Structured Perfect Bayesian Equilibrium). A structured perfect Bayesian equilibrium has strategy β^* that depends on public history h_t^c only through public belief π_t and on private observation $x_{1:t}^i$ only through the latest observation x_t^i.

4.2. Fixed Point Equation for Infinite Horizon

Below we state the fixed-point equation that defines the value function and strategy mapping for the infinite horizon problem. This is analogous to the backwards recursion ((6.4) and (6.5)) that defined the value function and θ mapping for the finite horizon problem. The fixed-point equation construction below along with the forward recursion given later calculate perfect Bayesian equilibria of the infinite horizon dynamic game.

Define the set of functions $V^i : \times_{j=1}^{N} \Delta(\mathcal{X}^j) \times \mathcal{X}^i \to \mathbb{R}$ and strategies $\tilde{\gamma}^i : \mathcal{X}^i \to \Delta(\mathcal{A}^i)$ (which is generated formally as $\tilde{\gamma}^i = \theta^i[\underline{\pi}]$ for given $\underline{\pi}$) via the following fixed-point equation: $\forall\, i \in \mathcal{N}$, $x^i \in \mathcal{X}^i$,

$$\tilde{\gamma}^i(\cdot \mid x^i) \in \underset{\gamma^i(\cdot|x^i)\in\Delta(\mathcal{A}^i)}{\operatorname{argmax}} \mathbb{E}^{\gamma^i(\cdot|x^i),\tilde{\gamma}^{-i},\pi^{-i}}\Big[R^i(X,A) + \delta V^i\Big(\underline{F}(\underline{\pi},\tilde{\gamma},A), X^{i,\prime}\Big) \mid \underline{\pi}, x^i\Big], \tag{6.17a}$$

$$V^i(\underline{\pi}, x^i) = \mathbb{E}^{\tilde{\gamma}^i(\cdot|x^i),\tilde{\gamma}^{-i},\pi^{-i}}\Big[R^i(X,A) + \delta V^i\Big(\underline{F}(\underline{\pi},\tilde{\gamma},A), X^{i,\prime}\Big) \mid \underline{\pi}, x^i\Big]. \tag{6.17b}$$

Note that the above is a joint fixed-point equation in $(V, \tilde{\gamma})$, unlike the backwards recursive algorithm earlier which required solving a fixed-point equation only in $\tilde{\gamma}$. Here the unknown quantity is distributed as

$$(X^{-i}, A^i, A^{-i}, X^{i,\prime}) \sim \pi^{-i}(x^{-i})\gamma^i(a^i \mid x^i)\tilde{\gamma}^{-i}(a^{-i} \mid x^{-i})Q^i(x^{i,\prime} \mid x^i, a). \tag{6.18}$$

and $\underline{F}(\cdot)$ is defined through (6.7). The belief update for histories with zero probability is part of the definition in (6.7). The choice of this update affects the value function and strategy calculated by the above fixed-point equation and thus the signaling equilibria that are calculated in this section (also see the discussion below (6.3)). The particular form used in (6.7) is just one instance of many that can be used here, proofs in this chapter go through for other belief updates as well.

Define belief μ^* inductively as follows: set $\mu_1^*(x_1) = \prod_{i=1}^N Q_0^i(x_1^i)$. Then for $t \geq 1$,

$$\mu_{t+1}^{i,*}\Big[h_{t+1}^c\Big] = F\Big(\mu_t^{i,*}[h_t^c], \theta^i\Big[\mu_t^*[h_t^c]\Big], a_t\Big), \tag{6.19a}$$

$$\mu_{t+1}^*\Big[h_{t+1}^c\Big](x_{t+1}) = \prod_{i=1}^N \mu_{t+1}^{i,*}\Big[h_{t+1}^c\Big](x_{t+1}^i). \tag{6.19b}$$

By construction the belief defined above satisfies the consistency condition needed for a Perfect Bayesian Equilibrium. Denote the stationary strategy arising out of $\tilde{\gamma}$ by β^* i.e.,

$$\beta_t^{i,*}(a_t^i \mid h_t^i) = \theta^i\Big[\mu_t^*[h_t^c]\Big](a_t^i \mid x_t^i). \tag{6.20}$$

4.3. Relation between Infinite and Finite Horizon problems

The following result states the consequence of the similarity between the fixed-point equation in infinite horizon and the backwards recursion in the finite horizon.

Lemma 4.2. *Consider the finite horizon game with* $G^i \equiv V^i$. *Then* $V_t^{i,T} = V^i$, $\forall\ i \in \mathcal{N}$, $t \in \{1, \dots, T\}$ *satisfies the backwards recursive construction* (6.4) *and* (6.5).

Proof. Please see Appendix 8.4 at the end of this chapter. □

4.4. Equilibrium Result

The central result of this chapter is that the strategy-belief pair (β^*, μ^*) constructed from the solution of the fixed-point equation (6.17) and the forward recursion of (6.19) and (6.20) indeed constitutes a PBE.

Theorem 4.3. *Assuming that the fixed-point equation* (6.17) *admits an absolutely bounded solution* V^i *(for all* $i \in \mathcal{N}$*), the strategy-belief pair* (β^*, μ^*) *defined in* (6.19) *and* (6.20) *is a PBE of the infinite horizon discounted reward dynamic game i.e.,* $\forall\ i \in \mathcal{N}$, β^i, $t \geq 1$, $h_t^i \in \mathcal{H}_t^i$,

$$\mathbb{E}^{\beta^{i,*},\beta^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{\infty} \delta^{n-t} R^i(X_n, A_n) \mid h_t^i\Big] \geq \mathbb{E}^{\beta^{i},\beta^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{\infty} \delta^{n-t} R^i(X_n, A_n) \mid h_t^i\Big] \tag{6.21}$$

Remark Note that by definition in (6.9), μ^* already satisfies the consistency conditions required for perfect Bayesian equilibrium.

Proof. Please see Appendix 8.5 at the end of this chapter. □

5. A Concrete Example

In this section, we consider an infinite horizon version of the public goods game from [FuTi91, ch. 8, Example 8.3]. We solve the corresponding fixed point equation (arising out of (6.17)) numerically to calculate the mapping θ (which in turn generates the perfect Bayesian equilibrium (β^*, μ^*)).

The example consists of two symmetric agents. The type space and action sets are $\mathcal{X}^1 = \mathcal{X}^2 = \{x^H, x^L\}$ and $\mathcal{A}^1 = \mathcal{A}^2 = \{0, 1\}$. Each agents' type is static and does not vary with time.

The actions represents whether agents are willing to contribute for a common public good. If at least one agent contributes then both agents receive utility 1 and the agent(s) that contributed receive cost equal to their type. If no one contributes then both agents receive utility 0. Thus the reward function is

$$R^i(x, a) = \begin{cases} 1 - x^i & \text{if } a^i = 1 \\ a^{-i} & \text{if } a^i = 0 \end{cases} \tag{6.22}$$

where a^{-i} represents the action taken by the agent other than i.

We use the following values $x^H = 1.2$, $x^L = 0.2$ and consider three values $\delta = 0, 0.5, 0.95$. Since type sets have two elements we can represent the distribution $\pi_1(\cdot) \in \Delta(\mathcal{X}^1)$ with only $\pi_1(x^H) \in [0, 1]$, similarly for agent 2. For any $\underline{\pi} = (\pi_1, \pi_2) \in [0, 1]^2$, the mapping $\theta[\underline{\pi}]$ produces $\tilde{\gamma}^i(\cdot \mid x^i)$ for every $i \in \{1, 2\}$ and $x^i \in \mathcal{X}^i = \{x^H, x^L\}$. Since the action space contains two elements, we can represent the distribution $\tilde{\gamma}^i(\cdot \mid x^i)$ by $\tilde{\gamma}^i(a^i = 1 \mid x^i)$ i.e., the probability of taking action 1. We solve the fixed-point equation by discretizing the $\underline{\pi}-$space $[0, 1]^2$ and all solutions that we find are symmetric w.r.t. agents i.e., $\tilde{\gamma}^1(\cdot \mid x^L)$ for $\underline{\pi} = (\pi_1, \pi_2)$ is the same as $\tilde{\gamma}^2(\cdot \mid x^L)$ for $\underline{\pi}' = (\pi_2, \pi_1)$ and similarly for type x^H.

For $\delta = 0$, the game is instantaneous and for the values considered, we have $1 - x^H = -0.2 < 0$. This implies that whenever agent 1's type is x^H, it is instantaneously profitable not to contribute. This gives $\tilde{\gamma}^1(a^1 = 1 \mid x^H) = 0$, for all $\underline{\pi}$. Thus we only plot $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$; in Fig. 20. For $\delta = 0$ the fixed-point equation (6.17) is only for the variable $\tilde{\gamma}$ and not V, and can be solved analytically. Refer to [VaAn16a, eq. (20) and Fig. (1)], where this solution is stated. There are multiple solutions to the fixed-point equation and our result from Fig. 20 matches with the one of the results in [VaAn16a].

Intuitively, with type x^L the only value of $\underline{\pi}$ for which agent 1 would not wish to contribute is if he anticipates agent 2's type to be x^L with high probability and rely on agent 2 to contribute. This is why for lower values of π_2 (i.e., agent 2's type likely to be x^L) we see $\tilde{\gamma}^1(a^1 = 1 \mid x^L) = 0$ in Fig. 20.

Now consider $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ plotted in Fig. 20, 21 and 23. As δ increases, future rewards attain more priority and signaling comes into play. So while taking an action, agents not only look for their instantaneous reward but also how their action affects the future public belief π about their private type. It is evident in the figures that as δ increases, at high π_1, up to larger values of π_2 agent 1 chooses not to contribute when his type is x^L. This way he intends to send a "wrong" signal to agent 2 i.e., that his type is x^H and subsequently force agent 2 to invest. This way agent 1 can free-ride on agent 2's investment.

Now consider Fig. 22 and 24, where $\tilde{\gamma}^1(a^1 = 1 \mid x^H)$ is plotted. Coordination via signaling is evident here. Although it is instantaneously not profitable to contribute if agent 1's type is x^H, by contributing at higher values of π_2 (i.e., agent 2's type is likely x^H) and low π_1, agent 1 coordinates with agent 2 to achieve net profit greater than 0 (reward when no one contributes). This can be done since the loss of contributing is -0.2 whereas profit from free-riding on agent 2's contribution is 1.

Under the equilibrium strategy, beliefs $\underline{\Pi}_t$ form a Markov chain. One can trace this Markov chain to study the signaling effect at equilibrium. On numerically simulating this Markov chain for the above example (at $\delta = 0.95$) we observe that for almost all initial beliefs, within a few rounds agents completely learn each other's private type truthfully (or at least with very high probability). In other words, agents manage to reveal their private type via their actions at equilibrium and to such an extent that it negates any possibly incorrect initial belief about their type.

As a measure of cooperative coordination at equilibrium one can perform the following calculation. Compare the value function $V^1(\cdot, x)$ of agent 1 arising out of the fixed-point equation, for $\delta = 0.95$ and $x \in \{x^H, x^L\}$ (normalize it by multiplying with $1 - \delta$ so that it represents per-round value) with the best possible attainable single-round reward under a symmetric mixed strategy with a) full coordination and b) no coordination. Note that the two cases need not be equilibrium themselves, which is why this will result in a bound on the efficiency of the evaluated equilibria.

In case a), assuming both agents have the same type x, full coordination can lead to the best possible reward of $\frac{1+1-x}{2} = 1 - \frac{x}{2}$ i.e., agent 1 contributes with probability 0.5 and agent

2 contributes with probability 0.5 but in a coordinated manner so that it doesn't overlap with agent 1 contributing.

In case b) when agents do not coordinate and invest with probability p each, then the expected single-round reward is $p(1-x)+p(1-p)$. The maximum possible value of this expression is $(1-\frac{x}{2})^2$.

For $x = x^L = 0.2$, the range of values of $V^1(\pi_1,\pi_2,x^L)$ over $(\pi_1,\pi_2) \in [0,1]^2$ is $[0.865, 0.894]$. Whereas full coordination produces 0.9 and no coordination 0.81. It is thus evident that agents at equilibrium end up achieving reward close to the best possible and gain significantly compared to the strategy of no coordination.

Similarly for $x = x^H = 1.2$ the range is $[0.3, 0.395]$. Whereas full coordination produces 0.4 and no coordination 0.16. The gain via coordination is evident here too.

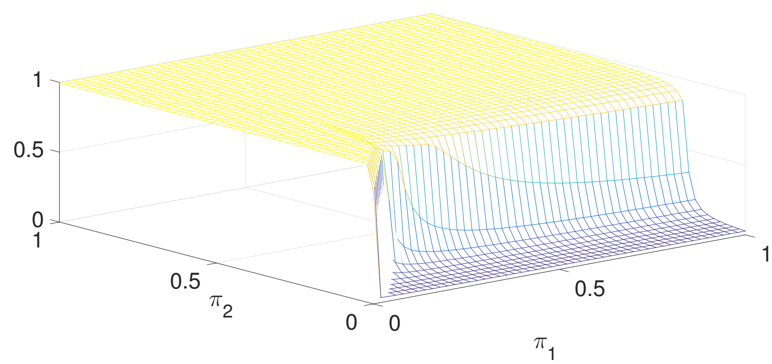

Figure 20.: $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1,π_2) at $\delta = 0$.

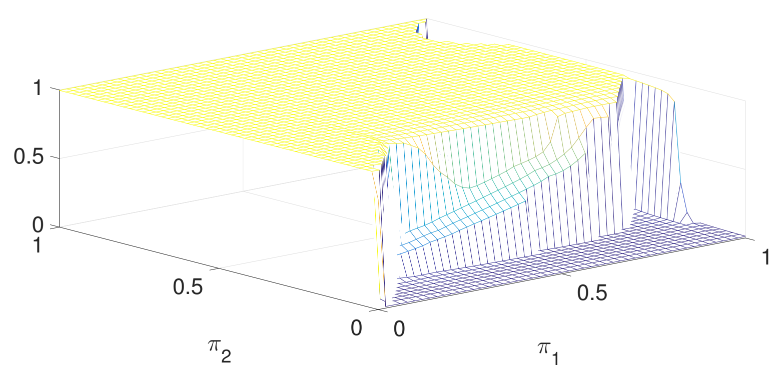

Figure 21.: $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1,π_2) at $\delta = 0.5$.

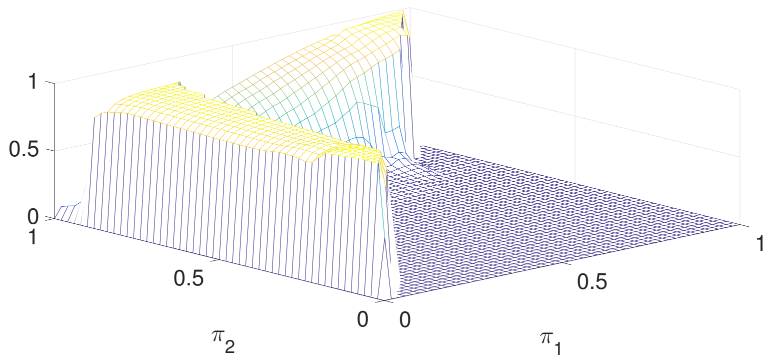

Figure 22.: $\tilde{\gamma}^1(a^1 = 1 \mid x^H)$ vs. (π_1, π_2) at $\delta = 0.5$.

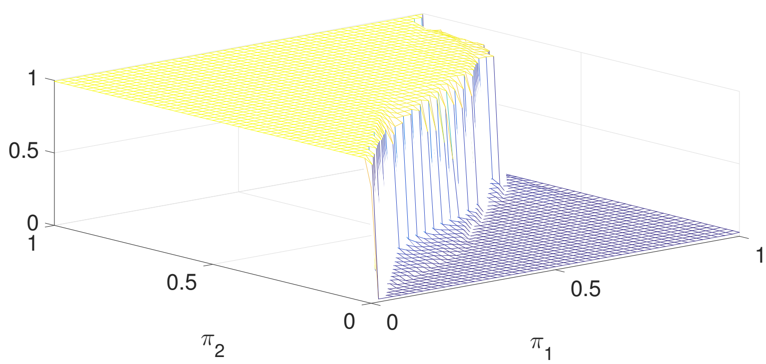

Figure 23.: $\tilde{\gamma}^1(a^1 = 1 \mid x^L)$ vs. (π_1, π_2) at $\delta = 0.95$.

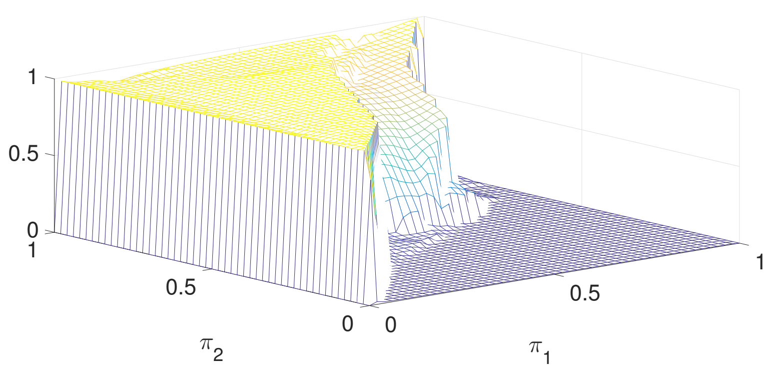

Figure 24.: $\tilde{\gamma}^1(a^1 = 1 \mid x^H)$ vs. (π_1, π_2) at $\delta = 0.95$.

6. An existence result for the finite horizon fixed-point equation

While it is known that for any finite dynamic game with asymmetric information and perfect recall, there always exists a PBE [OsRu94, Prop. 249.1], existence of SPBE is not guaranteed. It is clear from our algorithm that existence of SPBEs boils down to existence of a solution to the fixed-point equation (6.17). In this section we focus on the corresponding fixed-point equation in the finite horizon version (6.4) and discuss analytical and computational techniques to check existence of a solution. We focus on the finite horizon fixed-point equation to avoid existence issues related to the Value function V^i in the infinite horizon fixed-point equation. At each time t given the functions V^i_{t+1} for all $i \in \mathcal{N}$ from the previous round (in the backwards recursion), (6.4) must have a solution $\tilde{\gamma}^i_t$ for all $i \in \mathcal{N}$.

Generally, existence of equilibria is shown through Kakutani's fixed point theorem, similar to proving existence of a mixed strategy Nash equilibrium of a finite games [OsRu94; Na50]. Specifically, this is done by showing existence of a fixed point for the best-response correspondences of the game. Among other conditions, it requires the "closed graph" property of the correspondences, which is usually implied by the continuity property of the utility functions involved. For (6.4) establishing existence is not straightforward due to: (a) potential discontinuity of the π_t update function F when the denominator in the Bayesian update is 0 and (b) potential discontinuity of the value functions, V^i_{t+1}. In the following we provide sufficient conditions that can be checked at each time t to establish the existence of a solution.

It is in order to mention here the work in [Re99], generalized in [Re16], where the author identifies the property of *better-reply security* in games with discontinuous utilities to establish existence of pure strategy Nash equilibrium. Analogously, the author in [Re16] identifies the more general property of *point security*. There is a qualitative difference in the nature of the fixed-point equation encountered in this chapter and the one encountered in calculating pure strategy Nash equilibrium in normal form games (as in [Re99; Re16]), which is why we focus on developing a different set of existence results that are particularly geared towards the fixed-point equation encountered in this chapter.

In (6.4), agent i's optimal strategy $\tilde{\gamma}^i_t(\cdot|x^i_t)$ is the maximizer of the current reward plus reward-to-go from $t+1$. However this reward-to-go contains an updated belief Π_{t+1} at time $t+1$ and since we are interested in signaling equilibria, the belief update depends on the equilibrium strategy profile $\tilde{\gamma}_t$. Thus agent i's own strategy appears on the LHS as well as RHS and this makes the fixed-point equation (6.4) qualitative different from the best-response fixed-point equation encountered in calculating Nash equilibria in normal form games.

We consider a generic fixed-point equation similar to the one encountered in Section 3 and Section 4 and state conditions under which they are guaranteed to have a solution. To concentrate on the essential aspects of the problem we consider a simple case with $N = 2$, type sets $\mathcal{X}^i = \{x^H, x^L\}$ and action sets $\mathcal{A}^i = \{0, 1\}$. Furthermore, types are static and instantaneous rewards $R^i(x, a)$ do not depend on x^{-i}.

Given public belief $\underline{\pi} = (\pi^1, \pi^2) \in \times_{i=1}^2 \Delta(A^i)$, value functions V^1, V^2, one wishes to solve the following system of equations for $\left(\tilde{\gamma}^i(\cdot \mid x^i)\right)_{x^i \in \{x^H, x^L\}, i \in \{1,2\}}$.

$$\begin{aligned}\tilde{\gamma}^i(\cdot \mid x^i) \in \underset{\gamma^i(\cdot|x^i) \in \Delta(\mathcal{A}^i)}{\operatorname{argmax}}\; & \mathbb{E}^{\gamma^i(\cdot|x^i), \tilde{\gamma}^{-i}} \Big[R^i(x^i, A) \\ & + V^i\Big(\Big(F^1(\pi^1, \tilde{\gamma}^1, A^1), F^2(\pi^2, \tilde{\gamma}^2, A^2)\Big), x^i\Big) \mid x^i, \underline{\pi} \Big] \end{aligned} \tag{6.23}$$

where the expectation is evaluated using the probability distribution

$$(A^1, A^2) \sim \gamma^i(a^i \mid x^i)\Big[\pi^j(x^H)\tilde{\gamma}^j(a^j \mid x^H) + \pi^j(x^L)\tilde{\gamma}^j(a^j \mid x^L)\Big]. \tag{6.24}$$

The probabilistic policy $\tilde{\gamma}$ can be represented by the 4-tuple $\underline{p} = \left(\tilde{p}^{1L}, \tilde{p}^{2L}, \tilde{p}^{1H}, \tilde{p}^{2H}\right)$ where $\tilde{p}^{iH} = \gamma^i(a^i = 1 \mid x^H)$ and $\tilde{p}^{iL} = \gamma^i(a^i = 1 \mid x^L)$, $i = 1, 2$.

The fixed-point equation of interest reduces to

$$\begin{aligned}\tilde{p}^{1H} \in \underset{a \in [0,1]}{\operatorname{argmax}}\, a \Big[& \Big(\pi^2 \tilde{p}^{2H} + (1-\pi^2)\tilde{p}^{2L}\Big) \times \\ & \Big(V^1(F_1(\pi^1, \tilde{p}^1), F_1(\pi^2, \tilde{p}^2), x^H) - V^1(F_0(\pi^1, \tilde{p}^1), F_1(\pi^2, \tilde{p}^2), x^H)\Big) \\ + \Big(1 - \pi^2 \tilde{p}^{2H} - (1-\pi^2)\tilde{p}^{2L}\Big) & \Big(V^1(F_1(\pi^1, \tilde{p}^1), F_0(\pi^2, \tilde{p}^2), x^H) - V^1(F_0(\pi^1, \tilde{p}^1), F_0(\pi^2, \tilde{p}^2), x^H)\Big) \\ & + \Big(\pi^2 \tilde{p}^{2H} + (1-\pi^2)\tilde{p}^{2L}\Big)\Big(R^1(x^H, 1, 1) - R^1(x^H, 0, 1)\Big) \\ & + \Big(1 - \pi^2 \tilde{p}^{2H} - (1-\pi^2)\tilde{p}^{2L}\Big)\Big(R^1(x^H, 1, 0) - R^1(x^H, 0, 0)\Big)\Big] \end{aligned} \tag{6.25}$$

and three other similar equations for $\tilde{p}^{1L}, \tilde{p}^{2H}, \tilde{p}^{2L}$. Here

$$F_1(\pi, (p^H, p^L)) \triangleq \frac{\pi p^H}{\pi p^H + \bar{\pi} p^L} \tag{6.26a}$$

$$F_0(\pi, (p^H, p^L)) \triangleq \frac{\pi(1 - p^H)}{\pi(1 - p^H) + \bar{\pi}(1 - p^L)} \tag{6.26b}$$

and in both definitions, if the denominator is 0 then the RHS is taken as π.

6.1. Points of Discontinuities and the Closed graph result

Equation (6.25) and the other three similar equations are essentially of the form (for a given $\underline{\pi}$)

$$x \in \underset{a \in [0,1]}{\operatorname{argmax}} \; a f_1(x, y, w, z) \tag{6.27a}$$

$$y \in \underset{b \in [0,1]}{\operatorname{argmax}} \; b f_2(x, y, w, z) \tag{6.27b}$$

$$w \in \underset{c \in [0,1]}{\operatorname{argmax}} \; c f_3(x, y, w, z) \tag{6.27c}$$

$$z \in \underset{d \in [0,1]}{\operatorname{argmax}} \; d f_4(x, y, w, z) \tag{6.27d}$$

with x, y, z, w as $\tilde{p}^{1H}, \tilde{p}^{1L}, \tilde{p}^{2H}, \tilde{p}^{2L}$, respectively.

Define $\mathcal{D}_i \subseteq [0, 1]^4$ as the set of discontinuity points of f_i and $\mathcal{D} \triangleq \cup_{i=1}^4 \mathcal{D}_i$.

For any point $\underline{x}_0 \in \mathcal{D}$, define $S(\underline{x}_0)$ as the subset of indexes $i \in \{1, 2, 3, 4\}$ for which $f_i(\underline{x})$ is discontinuous at $\underline{x}_0$.

Assumption (E1) At any point $\underline{x}_0 \in \mathcal{D}$, $\forall\, i \in S(\underline{x}_0)$ one of the following is satisfied:

1. $f_i(\underline{x}_0) = 0$, or
2. $\exists\, \epsilon > 0$ such that $\forall\, \underline{x} \in B_\epsilon(\underline{x}_0)$ (inside an ϵ-ball of $\underline{x}_0$) the sign of $f_i(\underline{x})$ is same as the sign of $f_i(\underline{x}_0)$.

In the following we provide a sufficient condition for existence.

Theorem 6.1. *Under Assumption (E1), there exists a solution to the fixed-point equation* (6.27).

Proof. Please see Appendix 8.6 at the end of this chapter. □

The above set of results provide us with an analytical tool for establishing existence of a solution to the concerned fixed-point equation.

While the above analytical result is useful in understanding a theoretical basis for existence, it doesn't cover all instances. For instance, fixed-point equation [VaAn16a, eq. (21)] does not satisfy assumption (E1). In the following we provide a more computationally orientated approach to establishing existence and/or solving the generic fixed-point equation (6.27).

We motivate this case-by-case approach with the help of an example. Suppose we hypothesize that the solution to (6.27) is such that $x = 0, w = 0$ and $y, z \in (0, 1)$. Then (6.27) effectively reduces to checking if there exists $y^*, z^* \in (0, 1)$ such that

$$y^* \in \operatorname*{argmax}_b b f_2(0, y^*, 0, z^*) \tag{6.28a}$$

$$z^* \in \operatorname*{argmax}_d d f_4(0, y^*, 0, z^*) \tag{6.28b}$$

$$f_1(0, y^*, 0, z^*) \le 0 \tag{6.28c}$$

$$f_3(0, y^*, 0, z^*) \le 0. \tag{6.28d}$$

Thus the 4-variable system reduces to solving a 2-variable system and 2 conditions to verify. For instance, if $f_2(0, y, 0, z)$, $f_4(0, y, 0, z)$ as functions of y, z satisfy the conditions of Theorem 6.1 then the sub-system (6.28a), (6.28b) has a solution. If one of these solution is also consistent with (6.28c), (6.28d) then this sub-case indeed provides a solution to (6.27).

Generalizing the simplification provided in the above example, we divide solutions into $3^4 = 81$ cases[2] based on whether each of x, y, w, z are in $\{0\}, (0, 1), \{1\}$. There are

(1) 16 corner cases where none are in the strict interior $(0, 1)$;

(2) 32 cases where exactly one is in the strict interior $(0, 1)$;

(3) 24 cases where 2 variables are in the strict interior $(0, 1)$;

(4) 8 cases where 3 variables are in the strict interior $(0, 1)$; and

(5) 1 case where all 4 variables are in the strict interior $(0, 1)$.

Similar to the calculations above, for each of the 81 cases one can write a sub-system to

[2] In general there are $3^{\sum_{i=1}^{N} M_i}$ cases, where M_i is the number of types of agent i.

which the problem (6.27) effectively reduces to. Clearly, if any one of the 81 sub-systems has a solution then the problem (6.27) has a solution. Furthermore, searching for a solution reduces to an appropriate sub-problem depending on the case.

The approach then is to enumerate each of these 81 cases (as stated above) and check them in order. The computational simplification arises out of the fact that whenever a variable, say y, is not in the strict interior $(0, 1)$ then the corresponding equation (6.27b) need not be solved, since one only needs to verify the sign at a specific point. Hence, all sub-cases of (1) reduce to simply checking the value of functions f_i at corner points - no need for solving a fixed-point equation. All sub-cases of (2) reduce to solving a 1-variable fixed-point equation and three corresponding conditions to verify, etc.

7. Conclusion and Future Work

This chapter considers the infinite horizon discounted reward dynamic game with private types i.e., where each agent can only observe their own type. The types evolve as a controlled Markov process and are conditionally independent across agents given the action profile. Asymmetry of information between agents exists in this model, since each agent only knows their own private type.

To date, there exists no universal algorithm for calculating PBE in models with asymmetry of information that decouples, w.r.t. time, the calculation of strategy. Section 4 provides a single-shot fixed-point equation for calculating the equilibrium generating function θ, which in conjunction with the forward recursion in (6.19) and (6.20) gives a subset of PBEs (β^*, μ^*) of this game. The method proposed in this chapter finds PBE of a certain type i.e., where for any agent i, his strategy at equilibrium depends on the public history only through the common belief $\underline{\pi}_t$ and on private history only through agent i's current private type x_t^i.

Finally, we demonstrate our methodology by a concrete example of a two agent symmetric public goods game and observe the signaling effect in agents' strategies at equilibrium as discount factor is increased. The signaling effect implies that agents take into account how their actions affect future public beliefs $\underline{\pi}$ about their private type.

One important direction for future work is characterization of games where the proposed SPBE exists. This boils down to existence of a solution to the fixed-point equation in Section 4, as any SPBE must satisfy this equation.

Privacy-aware Dynamic Mechanisms

A possible direction for future research is investigation into mechanisms for dynamic environments encountered in the above model. The search for dynamic mechanisms can be aided by the fact that the above analysis provided a sequential decomposition algorithm for equilibrium analysis. The additional property of privacy is motivated by the following three works.

Authors in [BeVä10; AtSe13] present an (per round) ex-post incentive compatible direct dynamic mechanism, for the above model, which is efficient i.e., it achieves allocation that maximizes the infinite horizon discounted sum of utilities of the agents in the system.

$$\mathbb{E}\Big[\sum_{t=1}^{\infty}\delta^{t-1}\sum_{i=1}^{N}v_i(a_t,\theta_t^i)\Big]. \tag{6.29}$$

At equilibrium, in each round, each agent reveals their private observation/type θ_t^i truthfully. One of the main criticisms of truthful direct mechanisms, that applies here too, is that agents may not wish reveal their private observation completely, thereby making predictions derived from such mechanisms less robust. This motivation drives investigation into privacy-preserving mechanisms, where direct mechanisms are not considered as a design choice.

A special case of the model described above is when the uncertainty in state transition are Gaussian and the costs are quadratic (LQG). Authors in [OlRoYo15], define a framework of privacy for a system with Gaussian noise and quadratic costs that applies to dynamic single-price market clearing auctions and through appropriately revealing summary signals (such as only the price or only the total allocation) show that at the perfect Bayesian equilibrium (which summarizes entire history through appropriate beliefs) the allocation is efficient.

8. Appendix

8.1. Proof of Lemma 3.1

Proof. We use proof by contradiction. Suppose $\exists\, i, t, \hat{h}^i_t, \hat{\beta}^i$ such that (6.12) is violated. Construct strategy $\hat{\gamma}^i_t$ as

$$\hat{\gamma}^i_t(a^i_t \mid x^i_t) = \begin{cases} \hat{\beta}^i_t(a^i_t \mid \hat{h}^i_t) & \text{if } x^i_t = \hat{x}^i_t \\ \frac{1}{|\mathcal{A}^i|} & \text{if } x^i_t \neq \hat{x}^i_t \end{cases} \tag{6.30}$$

Then

$$\begin{aligned}
& V^{i,T}_t(\mu^*_t[\hat{h}^c_t], \hat{x}^i_t) && (6.31\text{a}) \\
& \geq \mathbb{E}^{\hat{\gamma}^i_t(\cdot|\hat{x}^i_t), \beta^{-i,*}_t, \mu^*_t[\hat{h}^c_t]}\Big[R^i(X_t, A_t) + \delta V^{i,T}_{t+1}\Big(\underline{F}(\underline{\mu}^*_t[\hat{h}^c_t], \beta^*_t(\cdot \mid \hat{h}^c_t, \cdot), A_t), X^i_{t+1}\Big) \mid \mu^*_t[\hat{h}^c_t], \hat{x}^i_t\Big] && (6.31\text{b}) \\
& = \mathbb{E}^{\hat{\beta}^i_t, \beta^{-i,*}_t, \mu^*_t[\hat{h}^c_t]}\Big[R^i(X_t, A_t) + \delta V^{i,T}_{t+1}\Big(\underline{F}(\underline{\mu}^*_t[\hat{h}^c_t], \beta^*_t(\cdot \mid \hat{h}^c_t, \cdot), A_t), X^i_{t+1}\Big) \mid \hat{h}^i_t\Big] && (6.31\text{c}) \\
& > V^{i,T}_t(\mu^*_t[\hat{h}^c_t], \hat{x}^i_t) && (6.31\text{d})
\end{aligned}$$

The first inequality above follows from the algorithm definition in (6.4) and (6.5), the second equality follows from the definition above in (6.30) and finally the last inequality follows from the assumption at the beginning of this proof.

Since the above is clearly a contradiction, the result follows. □

8.2. Proof of Lemma 3.2

Proof. This result relies on the structure of the update in (6.9), specifically that $\mu^{-i,*}_{t+1}[h^c_{t+1}]$ is a deterministic function of $\mu^{-i,*}_t[h^c_t]$, $\beta^{-i,*}_t$, a_t and does not depend on β^i.

Consider the joint pmf-pdf of random variables involved in the expectation

$$\mathbb{P}^{\beta^i_{t:T}, \beta^{-i,*}_{t:T}, \mu^*_t[h^c_t]}\Big(x^{-i}_{t+1}, a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t, a_t, x^i_{t+1}\Big). \tag{6.32}$$

This can be written as $\frac{A}{B}$ with

$$A = \sum_{x_t^{-i}} \mathbb{P}^{\beta^i_{t:T},\beta^{-i,*}_{t:T},\mu^*_t[h^c_t]}\Big(x_t^{-i}, a_t, x_{t+1}, a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t\Big) \tag{6.33a}$$

$$B = \sum_{\tilde{x}_t^{-i}} \mathbb{P}^{\beta^i_{t:T},\beta^{-i,*}_{t:T},\mu^*_t[h^c_t]}\Big(\tilde{x}_t^{-i}, a_t, x^i_{t+1} \mid h^i_t\Big) \tag{6.33b}$$

Using causal decomposition we can write

$$A = \sum_{x_t^{-i}} \mathbb{P}^{\beta^i_{t:T},\beta^{-i,*}_{t:T},\mu^*_t[h^c_t]}(x_t^{-i} \mid h^i_t)\beta^i_t(a^i_t \mid h^i_t)\beta^{-i,*}_t(a^{-i}_t \mid h^c_t, x^{-i}_t)Q(x_{t+1} \mid x_t, a_t)$$
$$\mathbb{P}^{\beta^i_{t:T},\beta^{-i,*}_{t:T},\mu^*_t[h^c_t]}(a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t, a_t, x^{-i}_t, x_{t+1}) \tag{6.34a}$$

$$= \sum_{x_t^{-i}} \mu^{-i,*}_t[h^c_t](x^{-i}_t)\beta^i_t(a^i_t \mid h^i_t)\beta^{-i,*}_t(a^{-i}_t \mid h^c_t, x^{-i}_t)Q^i(x^i_{t+1} \mid x^i_t, a_t)Q^{-i}(x^{-i}_{t+1} \mid x^{-i}_t, a_t)$$
$$\mathbb{P}^{\beta^i_{t+1:T},\beta^{-i,*}_{t+1:T},\mu^*_{t+1}[h^c_{t+1}]}(a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t, a_t, x^{-i}_t, x_{t+1}) \tag{6.34b}$$

where the second equality follows from the fact that given $h^i_t, a_t, x^{-i}_t, x_{t+1}$ and $\mu^*_t[h^c_t]$, the probability of $(a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1})$ depends on $h^i_t, a_t, x^{-i}_t, x_{t+1}, \mu^*_{t+1}[h^c_{t+1}]$ only through $\beta^i_{t+1:T}, \beta^{-i,*}_{t+1:T}$. Also the second equality above uses the fact that types evolve conditionally independent given action. Performing similar decomposition of the denominator B and substituting back in the expression from (6.32) allows us to cancel the terms $\beta^i(\cdot)$ and $Q^i(\cdot)$. Using the belief update from (6.9), this gives that the expression in (6.32) is

$$\mu^{-i,*}_{t+1}[h^c_{t+1}](x^{-i}_{t+1})\mathbb{P}^{\beta^i_{t+1:T},\beta^{-i,*}_{t+1:T},\mu^*_{t+1}[h^c_{t+1}]}(a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t, a_t, x_{t+1})$$
$$= \mathbb{P}^{\beta^i_{t+1:T},\beta^{-i,*}_{t+1:T},\mu^*_{t+1}[h^c_{t+1}]}(x^{-i}_{t+1}, a_{t+1:T}, x_{t+2:T}, x^i_{T+1}, \pi_{T+1} \mid h^i_t, a_t, x^i_{t+1}) \tag{6.35a}$$

The above equality follows directly from definition.

This completes the proof. □

8.3. Proof of Lemma 3.3

Proof. We use backward induction for this. At time T, using the maximization property from (6.4),

$$
\begin{aligned}
& V_T^{i,T}(\mu_T^*[h_T^c], x_T^i) && \text{(6.36a)}\\
& \triangleq \mathbb{E}^{\tilde{\gamma}_T^{i,T}(\cdot|x_T^i),\tilde{\gamma}_T^{-i,T},\mu_T^*[h_t^c]}\Big[R^i(X_T,A_T) + \delta G^i\Big(\underline{F}(\underline{\mu}_T^*[h_T^c],\tilde{\gamma}_T^T,A_T),X_{T+1}^i\Big) \mid \mu_T^*[h_T^c],x_T^i\Big] && \text{(6.36b)}\\
& \geq \mathbb{E}^{\gamma_T^{i,T}(\cdot|x_T^i),\tilde{\gamma}_T^{-i,T},\mu_T^*[h_t^c]}\Big[R^i(X_T,A_T) + \delta G^i\Big(\underline{F}(\underline{\mu}_T^*[h_T^c],\tilde{\gamma}_T^T,A_T),X_{T+1}^i\Big) \mid \mu_T^*[h_T^c],x_T^i\Big] && \text{(6.36c)}\\
& = W_T^{i,\beta^i,T}(h_T^i) && \text{(6.36d)}
\end{aligned}
$$

Here the second inequality follows from (6.4) and (6.5) and the final equality is by definition in (6.11).

Assume that the result holds for all $n \in \{t+1,\ldots,T\}$, then at time t we have

$$
\begin{aligned}
& V_t^{i,T}(\mu_t^*[h_t^c], x_t^i) && \text{(6.37a)}\\
& \geq \mathbb{E}^{\beta_t^i,\beta_t^{-i,*},\mu_t^*[h_t^c]}\Big[R^i(X_t,A_t) + \delta V_{t+1}^{i,T}\Big(\underline{F}(\underline{\mu}_t^*[h_t^c],\beta_t^*,A_t),X_{t+1}^i\Big) \mid h_t^i\Big] && \text{(6.37b)}\\
& \geq \mathbb{E}^{\beta_t^i,\beta_t^{-i,*},\mu_t^*[h_t^c]}\Big[R^i(X_t,A_t) + \delta \mathbb{E}^{\beta_{t+1:T}^i,\beta_{t+1:T}^{-i,*},\mu_{t+1}^*[h_t^c,A_t]}\Big[\sum_{n=t+1}^{T}\delta^{n-(t+1)}R^i(X_n,A_n) && \text{(6.37c)}\\
& \qquad + \delta^{T-t}G^i(\Pi_{T+1},X_{T+1}^i) \mid h_t^i,A_t,X_{t+1}^i\Big] \mid h_t^i\Big]\\
& = \mathbb{E}^{\beta_{t:T}^i,\beta_{t:T}^{-i,*},\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{T}\delta^{n-t}R^i(X_n,A_n) + \delta^{T+1-t}G^i(\Pi_{T+1},X_{T+1}^i) \mid h_t^i\Big] && \text{(6.37d)}\\
& = W_t^{i,\beta^i,T}(h_t^i) && \text{(6.37e)}
\end{aligned}
$$

Here the first inequality follows from Lemma 3.1, the second inequality from the induction hypothesis, the third equality follows from Lemma 3.2 and the final equality by definition (6.11). □

8.4. Proof of Lemma 4.2

Proof. Use backward induction for this. Consider the finite horizon algorithm at time $t = T$, noting that $V^{i,T}_{T+1} \equiv G^i \equiv V^i$,

$$\tilde{\gamma}^{i,T}_T(\cdot \mid x^i_T) \in \underset{\gamma^i_T(\cdot|x^i_T)\in\Delta(\mathcal{A}^i)}{\operatorname{argmax}} \mathbb{E}^{\gamma^i_T(\cdot|x^i_T),\tilde{\gamma}^{-i,T}_T,\underline{\pi}^{-i}_T}\Big[R^i(X_T, A_T) \\ + \delta V^i\Big(\underline{F}(\underline{\pi}_T, \tilde{\gamma}^T_T, A_T), X^i_{T+1}\Big) \mid \underline{\pi}_T, x^i_T\Big] \quad (6.38a)$$

$$V^{i,T}_T(\underline{\pi}_T, x^i_T) = \mathbb{E}^{\tilde{\gamma}^{i,T}_T(\cdot|x^i_T),\tilde{\gamma}^{-i,T}_T,\underline{\pi}^{-i}_T}\Big[R^i(X_T, A_T) \\ + \delta V^i\Big(\underline{F}(\underline{\pi}_T, \tilde{\gamma}^T_T, A_T), X^i_{T+1}\Big) \mid \underline{\pi}_T, x^i_T\Big] \quad (6.38b)$$

Comparing the above set of equations with (6.17), we can see that the pair $(V, \tilde{\gamma})$ arising out of (6.17) satisfies the above. Now assume that $V^{i,T}_n \equiv V^i$ for all $n \in \{t+1, \ldots, T\}$. At time t, in the finite horizon construction from (6.4), (6.5), substituting V^i in place of $V^{i,T}_{t+1}$ from the induction hypothesis, we get the same set of equations as (6.38). Thus $V^{i,T}_t \equiv V^i$ satisfies it. □

8.5. Proof of Theorem 4.3

Proof. We divide the proof into two parts: first we show that the value function V^i is at least as big as any reward-to-go function; secondly we show that under the strategy β^*_i, reward-to-go is V^i.

Part 1. For any $i \in \mathcal{N}$, β^i define the following reward-to-go functions

$$W^{i,\beta^i}_t(h^i_t) = \mathbb{E}^{\beta^i,\beta^{-i,*},\mu^*_t[h^c_t]}\Big[\sum_{n=t}^{\infty} \delta^{n-t} R^i(X_n, A_n) \mid h^i_t\Big] \quad (6.39a)$$

$$W^{i,\beta^i,T}_t(h^i_t) = \mathbb{E}^{\beta^i,\beta^{-i,*},\mu^*_t[h^c_t]}\Big[\sum_{n=t}^{T} \delta^{n-t} R^i(X_n, A_n) + \delta^{T+1-t} V^i(\underline{\Pi}_{T+1}, X^i_{T+1}) \mid h^i_t\Big] \quad (6.39b)$$

Since $\mathcal{X}^i$, $\mathcal{A}^i$ are finite sets the reward R^i is absolutely bounded, the reward-to-go $W_t^{i,\beta^i}(h_t^i)$ is finite $\forall\ i, t, \beta^i, h_t^i$.

For any $i \in \mathcal{N}$, $h_t^i \in \mathcal{H}_t^i$,

$$V^i\Big(\mu_t^*[h_t^c], x_t^i\Big) - W_t^{i,\beta^i}(h_t^i) = \\ \Big(V^i\Big(\mu_t^*[h_t^c], x_t^i\Big) - W_t^{i,\beta^i,T}(h_t^i)\Big) + \Big(W_t^{i,\beta^i,T}(h_t^i) - W_t^{i,\beta^i}(h_t^i)\Big) \quad (6.40)$$

Combining results from Lemma 3.3 and 4.2, the term in the first bracket in RHS of (6.40) is non-negative. Using (6.39), the term in the second bracket is

$$\Big(\delta^{T+1-t}\Big)\,\mathbb{E}^{\beta^i,\beta^{-i,*},\mu_t^*[h_t^c]}\Big[-\sum_{n=T+1}^{\infty}\delta^{n-(T+1)}R^i(X_n, A_n) + V^i(\underline{\Pi}_{T+1}, X_{T+1}^i) \mid h_t^i\Big]. \quad (6.41)$$

The summation in the expression above is bounded by a convergent Geometric series. Also, V^i is bounded. Hence the above quantity can be made arbitrarily small by choosing T appropriately large. Since the LHS of (6.40) does not depend on T, this gives that

$$V^i\Big(\mu_t^*[h_t^c], x_t^i\Big) \geq W_t^{i,\beta^i}(h_t^i) \quad (6.42)$$

Part 2. Since the strategy β^* generated in (6.20) is such that $\beta_t^{i,*}$ depends on h_t^i only through $\mu_t^*[h_t^c]$ and x_t^i, the reward-to-go $W_t^{i,\beta^{i,*}}$, at strategy β^*, can be written (with abuse of notation) as

$$W_t^{i,\beta^{i,*}}(h_t^i) = W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) = \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[\sum_{n=t}^{\infty}\delta^{n-t}R^i(X_n, A_n) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.43)$$

For any $h_t^i \in \mathcal{H}_t^i$,

$$W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) = \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[R^i(X_t, A_t) \\ + \delta W_{t+1}^{i,\beta^{i,*}}\Big(\underline{F}(\underline{\mu}_t^*[h_t^c], \theta[\mu_t^*[h_t^c]], A_{t+1}), X_{t+1}^i\Big) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.44a)$$

$$V^i(\mu_t^*[h_t^c], x_t^i) = \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[R^i(X_t, A_t) \\ + \delta V^i\Big(\underline{F}(\underline{\mu}_t^*[h_t^c], \theta[\mu_t^*[h_t^c]], A_{t+1}), X_{t+1}^i\Big) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.44b)$$

Repeated application of the above for the first n time periods gives

$$W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) = \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[\sum_{m=t}^{t+n-1} \delta^{m-t} R^i(X_t, A_t) + \delta^n W_{t+n}^{i,\beta^{i,*}}\Big(\Pi_{t+n}, X_{t+n}^i\Big) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.45a)$$

$$V^i(\mu_t^*[h_t^c], x_t^i) = \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[\sum_{m=t}^{t+n-1} \delta^{m-t} R^i(X_t, A_t) + \delta^n V^i\Big(\Pi_{t+n}, X_{t+n}^i\Big) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.45b)$$

Here Π_{t+n} is the $n-$step belief update under strategy and belief prescribed by β^*, μ^*.

Taking difference gives

$$W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) - V^i(\mu_t^*[h_t^c], x_t^i) = \delta^n \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[W_{t+n}^{i,\beta^{i,*}}\Big(\Pi_{t+n}, X_{t+n}^i\Big) - V^i\Big(\Pi_{t+n}, X_{t+n}^i\Big) \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.46a)$$

Taking absolute value of both sides then using Jensen's inequality for $f(x) = |x|$ and finally taking supremum over h_t^i gives us

$$\sup_{h_t^i} \Big|W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) - V^i(\mu_t^*[h_t^c], x_t^i)\Big| \leq \delta^n \sup_{h_t^i} \mathbb{E}^{\beta^*,\mu_t^*[h_t^c]}\Big[\Big|W_{t+n}^{i,\beta^{i,*}}(\Pi_{t+n}, X_{t+n}^i) - V^i(\mu_t^*[h_t^c], x_t^i)\Big| \mid \mu_t^*[h_t^c], x_t^i\Big] \quad (6.47)$$

Now using the fact that W_{t+n}, V^i are bounded and that we can choose n arbitrarily large, we get $\sup_{h_t^i} |W_t^{i,\beta^{i,*}}(\mu_t^*[h_t^c], x_t^i) - V^i(\mu_t^*[h_t^c], x_t^i)| = 0$. $\square$

8.6. Proof of Theorem 6.1

Proof. Denote the vector correspondence defined by the RHS of (6.27) by

$$\phi(\underline{x}) = \begin{pmatrix} \phi_1(\underline{x}) \\ \vdots \\ \phi_4(\underline{x}) \end{pmatrix} = \begin{pmatrix} \operatorname{argmax}_a a f_1(\underline{x}) \\ \vdots \\ \operatorname{argmax}_d d f_4(\underline{x}) \end{pmatrix} \tag{6.48}$$

where $\underline{x} = (x, y, w, z)$. For any $\underline{x} \in [0, 1]^4$, $\phi(\underline{x})$ is non-empty and closed, since the argmax solution always exists and is one of $\{0\}, \{1\}, [0, 1]$. If in addition ϕ also has a *closed graph* then by Kakutani Fixed Point Theorem there exists a solution to (6.27).

Consider any sequence $(\underline{x}_n, a_n, b_n, c_n, d_n) \to (\underline{x}_0, a_0, b_0, c_0, d_0)$ such that $\forall\ n \geq 1$,

$$a_n \in \operatorname*{argmax}_{a\in[0,1]} a f_1(\underline{x}_n) \tag{6.49a}$$

$$b_n \in \operatorname*{argmax}_{b\in[0,1]} b f_2(\underline{x}_n) \tag{6.49b}$$

$$c_n \in \operatorname*{argmax}_{c\in[0,1]} c f_3(\underline{x}_n) \tag{6.49c}$$

$$d_n \in \operatorname*{argmax}_{d\in[0,1]} d f_4(\underline{x}_n). \tag{6.49d}$$

We need to verify that (6.49) also holds for the limit $(\underline{x}_0, a_0, b_0, c_0, d_0)$.

If $\underline{x}_0 \notin \mathcal{D}$ then due to continuity, (6.49) indeed holds at the limit.

For $\underline{x}_0 \in \mathcal{D}$, for any $i \in S(\underline{x}_0)$ if $f_i(\underline{x}_0) = 0$ then in the relation to be verified, the requirement is either of $a_0, b_0, c_0, d_0 \in [0, 1]$, which is always true.

For $\underline{x}_0 \in \mathcal{D}_1 \cap \mathcal{D}_2^{\complement} \cap \mathcal{D}_3^{\complement} \cap \mathcal{D}_4^{\complement}$, if $f_1(\underline{x}_0) > 0$ then for any sequence $\underline{x}_n \to \underline{x}_0$, for large n the points in the sequence are within $B_\epsilon(\underline{x}_0)$ and thus $f_1(\underline{x}_n) > 0$ for large n. This means that the relation from (6.49) holds at the limit (noting that f_2, f_3, f_4 are continuous at $\underline{x}_0$ in this case).

Similarly if $f_1(\underline{x}_0) < 0$ and for any $\underline{x}_0 \in \mathcal{D}_1^{\complement} \cap \mathcal{D}_2 \cap \mathcal{D}_3^{\complement} \cap \mathcal{D}_4^{\complement}$.

For $\underline{x}_0 \in \mathcal{D}_1 \cap \mathcal{D}_2 \cap \mathcal{D}_3^{\complement} \cap \mathcal{D}_4^{\complement}$ if $f_1(\underline{x}_0) > 0$ and $f_2(\underline{x}_0) < 0$ then there exists an $\epsilon > 0$ such that $\forall\ \underline{x} \in B_\epsilon(\underline{x}_0)$ we have $f_1(\underline{x}) > 0$ and $f_2(\underline{x}) < 0$. From this it follows that the relation (6.49) holds at the limit. Similar argument works for any other sign combination of

f_1, f_2, f_3, f_4.

The two arguments above cover all cases. $\square$

Bibliography

[AcDaLoOz11] Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. "Bayesian learning in social networks". In: *The Review of Economic Studies* 78.4 (2011), pp. 1201–1236.

[AgFuHoKuTe12] Philippe Aghion, Drew Fudenberg, Richard Holden, Takashi Kunimoto, and Olivier Tercieux. "Subgame-Perfect Implementation Under Information Perturbations*". In: *The Quarterly Journal of Economics* 127.4 (2012), pp. 1843–1881.

[AlBoHoPo13a] Tansu Alpcan, Holger Boche, Michael L Honig, and H Vincen Poor. *Mechanisms and games for dynamic spectrum allocation*. Cambridge University Press, 2013.

[AlBoHoPo13b] Tansu Alpcan, Holger Boche, Michael L Honig, and H Vincent Poor. *Mechanisms and games for dynamic spectrum allocation*. Cambridge University Press, 2013.

[AtSe13] Susan Athey and Ilya Segal. "An efficient dynamic mechanism". In: *Econometrica* 81.6 (2013), pp. 2463–2485.

[AtYe08] Murat Atlamaz and Duygu Yengin. "Fair groves mechanisms". In: *Social Choice and Welfare* 31.4 (2008), pp. 573–587.

[BaLaSaSa16] G. Bacci, S. Lasaulce, W. Saad, and L. Sanguinetti. "Game Theory for Networks: A tutorial on game-theoretic tools for emerging signal processing applications". In: *IEEE Signal Processing Magazine* 33.1 (Jan. 2016), pp. 94–119. ISSN: 1053-5888. DOI: 10.1109/MSP.2015.2451994.

[BaOl99] Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. Vol. 23. Siam, 1999.

[BaSo87] Jeffrey S. Banks and Joel Sobel. "Equilibrium selection in signaling games". In: *Econometrica* 55.3 (1987), pp. 647–661.

[BeAlHoAmAl15] E. Bedeer, A. Alorainy, M.J. Hossain, O. Amin, and M.-S. Alouini. "Fairness-Aware Energy-Efficient Resource Allocation for AF Co-Operative OFDMA Networks". In: 33.12 (Dec. 2015), pp. 2478–2493.

[BeCo03] Timothy Besley and Stephen Coate. "Centralized versus decentralized provision of local public goods: a political economy approach". In: *Journal of public economics* 87.12 (2003), pp. 2611–2637.

[BeMo13] Dirk Bergemann and Stephen E Morris. *An Introduction to Robust Mechanism Design*. Now, 2013.

[BeVä10] Dirk Bergemann and Juuso Välimäki. "The dynamic pivot mechanism". In: *Econometrica* 78.2 (2010), pp. 771–789.

[BhKaChGu14] Saptarshi Bhattacharya, Koushik Kar, Joe H Chow, and Aparna Gupta. "Progressive second price auctions with elastic supply for PEV charging in the smart grid". In: *Proc. of NeTGCoop* (2014).

[BiHiWe92] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. "A theory of fads, fashion, custom, and cultural change as informational cascades". In: *Journal of political Economy* (1992), pp. 992–1026.

[BöKrSt15] Tilman Börgers, Daniel Krahmer, and Roland Strausz. *An introduction to the theory of mechanism design*. Oxford University Press, 2015.

[BoPaChPeEc11] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. "Distributed optimization and statistical learning via the alternating direction method of multipliers". In: *Foundations and Trends® in Machine Learning* 3.1 (2011), pp. 1–122.

[BöSa97] Tilman Börgers and Rajiv Sarin. "Learning through reinforcement and replicator dynamics". In: *Journal of Economic Theory* 77.1 (1997), pp. 1–14.

[Br51a] George W Brown. "Iterative solution of games by fictitious play". In: *Activity analysis of production and allocation* 13.1 (1951), pp. 374–376.

[Br51b] George W Brown. "Iterative solution of games by fictitious play". In: *Activity analysis of production and allocation* 13.1 (1951), pp. 374–376.

[CaKe10] S. Caron and G. Kesidis. "Incentive-Based Energy Consumption Scheduling Algorithms for the Smart Grid". In: *First IEEE International Conference on Smart Grid Communications (SmartGridComm), 2010*. Oct. 2010, pp. 391–396.

[Ch02] Yan Chen. "A family of supermodular Nash mechanisms implementing Lindahl allocations". In: *Economic Theory* 19.4 (2002), pp. 773–790.

[Ch87] In-Koo Cho. "A refinement of sequential equilibria". In: *Econometrica* 55.6 (1987), pp. 1367–1389.

[ChKr87] In-Koo Cho and David M. Kreps. "Signaling games and stable equilibria". In: *The Quarterly Journal of Economics* 102.2 (1987), pp. 179–221.

[CoGkGo13] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. "Mechanism Design for Fair Division: Allocating Divisible Items Without Payments". In: *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. New York, NY, USA: ACM, 2013, pp. 251–268. DOI: `10.1145/2482540.2482582`. URL: `http://doi.acm.org/10.1145/2482540.2482582`.

[CoKo01] Harold L Cole and Narayana Kocherlakota. "Dynamic games with hidden actions and hidden states". In: *Journal of Economic Theory* 98.1 (2001), pp. 114–126.

[CuNa05] Calin Curescu and Simin Nadjm-Tehrani. "Price/utility-based optimized resource allocation in wireless ad hoc networks". In: *SECON*. 2005, pp. 85–95.

[dACrGé04] Claude d'Aspremont, Jacques Crémer, and Louis-André Gérard-Varet. "Balanced bayesian mechanisms". In: *Journal of Economic Theory* 115.2 (2004), pp. 385–396.

[dACrGé90] Claude d'Aspremont, Jacques Cremer, and Louis-André Gérard-Varet. "Incentives and the existence of Pareto-optimal revelation mechanisms". In: *Journal of Economic Theory* 51.2 (1990), pp. 233–254.

[DaDeKi15] Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. "Near-optimal no-regret algorithms for zero-sum games". In: *Games and Economic Behavior* 92 (2015), pp. 327–348.

[dAGé79] Claude d'Aspremont and Louis-Andre Gérard-Varet. "Incentives and incomplete information". In: *Journal of Public economics* 11.1 (1979), pp. 25–45.

[DuAgWa12] John C Duchi, Alekh Agarwal, and Martin J Wainwright. "Dual averaging for distributed optimization: Convergence analysis and network scaling". In: *IEEE Transactions on Automatic control* 57.3 (2012), pp. 592–606.

[FaAl00] M. Fahrioglu and F. L. Alvarado. "Designing incentive compatible contracts for effective demand management". In: *IEEE Transactions on Power Systems* 15.4 (Nov. 2000), pp. 1255–1260.

[Fu98] Drew Fudenberg. *The theory of learning in games*. Vol. 2. MIT press, 1998.

[FuSc07] F. Fu and M. van der Schaar. "Noncollaborative Resource Management for Wireless Multimedia Applications Using Mechanism Design". In: *IEEE Transactions on Multimedia* 9.4 (June 2007), pp. 851–868.

[FuTi91] Drew Fudenberg and Jean Tirole. *Game theory*. MIT press, 1991.

[GaNaGu08a] Dinesh Garg, Y Narahari, and Sujit Gujar. "Foundations of mechanism design: A tutorial Part 1-Key concepts and classical results". In: *Sadhana (Academy Proceedings in Engineering Sciences)*. Vol. 33. 2. Indian Academy of Sciences. 2008, pp. 83–130.

[GaNaGu08b] Dinesh Garg, Y Narahari, and Sujit Gujar. "Foundations of mechanism design: A tutorial Part 2-Advanced concepts and results". In: *Sadhana* 33.2 (2008), pp. 131–174.

[Ge11] Enrico H Gerding, Valentin Robu, Sebastian Stein, David C Parkes, Alex Rogers, and Nicholas R Jennings. "Online mechanism design for electric vehicle charging". In: *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems. 2011, pp. 811–818.

[Gr73] Theodore Groves. "Incentives in teams". In: *Econometrica: Journal of the Econometric Society* (1973), pp. 617–631.

[GrLe77] Theodore Groves and John Ledyard. "Optimal allocation of public goods: A solution to the "free rider" problem". In: *Econometrica: Journal of the Econometric Society* (1977), pp. 783–809.

[GrSrTo01] Emily E Graves, R Srikant, and Don Towsley. *Decentralized computation of weighted max-min fair bandwidth allocation in networks with multicast flows*. Springer, 2001.

[GuNaLaBa14] Abhishek Gupta, Ashutosh Nayyar, Cedric Langbort, and Tamer Başar. "Common information based Markov perfect equilibria for linear-Gaussian games with asymmetric information". In: *SIAM Journal on Control and Optimization* 52.5 (2014), pp. 3228–3260.

[Ha12a] Zhu Han. *Game theory in wireless and communication networks: theory, models, and applications*. Cambridge University Press, 2012.

[Ha12b] Jason D Hartline. "Bayesian mechanism design". In: *Theoretical Computer Science* 8.3 (2012), pp. 143–263.

[HeMa12] Paul J Healy and Laurent Mathevet. "Designing stable mechanisms for economic environments". In: *Theoretical Economics* 7.3 (2012), pp. 609–661.

[HoGa12] Mingyi Hong and Alfredo Garcia. "Mechanism design for base station association and resource allocation in downlink OFDMA network". In: 30.11 (2012), pp. 2238–2250.

[HoSa02] Josef Hofbauer and William H Sandholm. "On the global convergence of stochastic fictitious play". In: *Econometrica* 70.6 (2002), pp. 2265–2294.

[Hu79] Leonid Hurwicz. "Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points". In: *The Review of Economic Studies* 46.2 (1979), pp. 217–225.

[Hu94] Leonid Hurwicz. "Economic design, adjustment processes, mechanisms, and institutions". In: *Review of Economic Design* 1.1 (1994), pp. 1–14.

[HuBeHo06a] Jianwei Huang, Randall A Berry, and Michael L Honig. "Auction-based spectrum sharing". In: *Mobile Networks and Applications* 11.3 (2006), pp. 405–418.

[HuBeHo06b] Jianwei Huang, Randall A Berry, and Michael L Honig. "Auction-based spectrum sharing". In: *Mobile Networks and Applications* 11.3 (2006), pp. 405–418.

[HuRe06] Leonid Hurwicz and Stanley Reiter. *Designing economic mechanisms*. Cambridge University Press, 2006.

[HuSuBeAg09] Jianwei Huang, Vijay Subramanian, Randall Berry, and Rajeev Agrawal. "Scheduling and resource allocation in OFDMA wireless systems". In: *Orthogonal Frequency Division Multiple Access Fundamentals and Applications* (2009).

[Ja01] Matthew O Jackson. "A crash course in implementation theory". In: *Social choice and welfare* 18.4 (2001), pp. 655–708.

[Ja91] Mathew O Jackson. "Bayesian implementation". In: *Econometrica: Journal of the Econometric Society* (1991), pp. 461–477.

[JaWa10] Rahul Jain and Jean Walrand. "An efficient Nash-implementation mechanism for network resource allocation". In: *Automatica* 46.8 (2010), pp. 1276–1283.

[Jo86] James S Jordan. "Instability in the implementation of Walrasian allocations". In: *Journal of Economic Theory* 39.2 (1986), pp. 301–328.

[JoTs04] Ramesh Johari and John N Tsitsiklis. "Efficiency loss in a network resource allocation game". In: *Mathematics of Operations Research* 29.3 (2004), pp. 407–435.

[JoTs09] Ramesh Johari and John N Tsitsiklis. "Efficiency of scalar-parameterized mechanisms". In: *Operations Research* 57.4 (2009), pp. 823–839.

[KaSaTa01] Koushik Kar, Saswati Sarkar, and Leandros Tassiulas. "Optimization based rate control for multirate multicast sessions". In: *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*. Vol. 1. IEEE. 2001, pp. 123–132.

[KaTa06] Koushik Kar and Leandros Tassiulas. "Layered multicast rate control based on Lagrangian relaxation and dynamic programming". In: 24.8 (2006), pp. 1464–1474.

[KaTe12] A. Kakhbod and D. Teneketzis. "An Efficient Game Form for Unicast Service Provisioning". In: *IEEE Transactions on Automatic Control* 57.2 (2012), pp. 392–404. ISSN: 0018-9286.

[KaTe13a] A. Kakhbod and D. Teneketzis. "Correction to "An Efficient Game Form for Unicast Service Provisioning"". In: (2013). arXiv: `0910.5502`. URL: `http://arxiv.org/abs/0910.5502`.

[KaTe13b] Ali Kakhbod and Demosthen Teneketzis. "Correction to "An Efficient Game Form for Multi-Rate Multicast Service Provisioning"". In: 31.7 (July 2013), pp. 1355–1356. ISSN: 0733-8716.

[KaTe14] A. Kakhbod and D. Teneketzis. "Correction to "An Efficient Game Form for Unicast Service Provisioning"". In: PP.99 (2014), pp. 1–1. ISSN: 0018-9286.

[KeMaTa98] Frank P Kelly, Aman K Maulloo, and David KH Tan. "Rate control for communication networks: shadow prices, proportional fairness and stability". In: *Journal of the Operational Research society* 49.3 (1998), pp. 237–252.

[Kr09] Vijay Krishna. *Auction theory*. Academic press, 2009.

[KrZa14] Vlastimil Krivan and Georges Zaccour. *Advances in dynamic games*. Springer, 2014.

[KuWh04] H. J. Kushner and P. A. Whiting. "Convergence of proportional-fair sharing algorithms under general conditions". In: *IEEE Transactions on Wireless Communications* 3.4 (July 2004).

[LaKaChSa10] T. Lan, D. Kao, M. Chiang, and A. Sabharwal. "An Axiomatic Theory of Fairness in Network Resource Allocation". In: *INFOCOM, 2010 Proceedings IEEE*. Mar. 2010.

[LaTe11] Samson Lasaulce and Hamidou Tembine. *Game theory and learning for wireless networks: fundamentals and applications*. Access Online via Elsevier, 2011.

[LeSuBe14] Tho Ngoc Le, Vijay G Subramanian, and Randall A Berry. "The impact of observation and action errors on informational cascades". In: *53rd IEEE Conference on Decision and Control*. IEEE. 2014, pp. 1917–1922.

[LiSh14] L. Li and J. Shamma. "LP formulation of asymmetric zero-sum stochastic games". In: *53rd IEEE Conference on Decision and Control*. Dec. 2014, pp. 1930–1935.

[Ma99] Eric Maskin. "Nash equilibrium and welfare optimality". In: *The Review of Economic Studies* 66.1 (1999), pp. 23–38.

[MaBa03] Rajiv T Maheswaran and Tamer Başar. "Nash equilibrium and decentralized negotiation in auctioning divisible resources". In: *Group Decision and Negotiation* 12.5 (2003), pp. 361–395.

[MaBa04] Rajiv T Maheswaran and Tamer Başar. “Social welfare of selfish agents: motivating efficiency for divisible resources”. In: *43rd IEEE Conference on Decision and Control, 2004. CDC.* Vol. 2. IEEE. 2004, pp. 1550–1555.

[MaBa06] Rajiv Maheswaran and Tamer Başar. “Efficient signal proportional allocation (ESPA) mechanisms: Decentralized social welfare maximization for divisible resources”. In: *IEEE Journal on Selected Areas in Communications* 24.5 (2006), pp. 1000–1009.

[MaSa06] George J Mailath and Larry Samuelson. *Repeated games and reputations: long-run relationships*. Oxford university press, 2006.

[MaSj02] Eric Maskin and Tomas Sjöström. “Implementation theory”. In: *Handbook of social Choice and Welfare* 1 (2002), pp. 237–288.

[MaWhGr95] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R Green. *Microeconomic theory*. Vol. 1. Oxford university press New York, 1995.

[MeOz11] Ishai Menache and Asuman Ozdaglar. “Network Games: Theory, Models, and Dynamics”. In: *Synthesis Lectures on Communication Networks* 4.1 (2011), pp. 1–159.

[MiRo90] Paul Milgrom and John Roberts. “Rationalizability, learning, and equilibrium in games with strategic complementarities”. In: *Econometrica: Journal of the Econometric Society* (1990), pp. 1255–1277.

[Mo16] H. Mohsenian-Rad. “Coordinated Price-Maker Operation of Large Energy Storage Units in Nodal Energy Markets”. In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016), pp. 786–797.

[MoRe90] Dilip Mookherjee and Stefan Reichelstein. “Implementation via augmented revelation mechanisms”. In: *The Review of Economic Studies* (1990), pp. 453–475.

[MoSh96] Dov Monderer and Lloyd S Shapley. “Fictitious play property for games with identical interests”. In: *Journal of economic theory* 68.1 (1996), pp. 258–265.

[Na50] Jonh F. Nash. “Non-cooperative Games”. PhD thesis. Princeton University, May 1950.

[NaGuLaBa14] A. Nayyar, A. Gupta, C. Langbort, and T. Başar. “Common Information Based Markov Perfect Equilibria for Stochastic Games With Asymmetric Information: Finite Games”. In: *IEEE Transactions on Automatic Control* 59.3 (Mar. 2014), pp. 555–570.

[NaMaTe13] A. Nayyar, A. Mahajan, and D. Teneketzis. “Decentralized Stochastic Control with Partial History Sharing: A Common Information Approach”. In: *IEEE Transactions on Automatic Control* 58.7 (July 2013), pp. 1644–1658.

[Ne09] Michael J Neely. “Optimal pricing in a free market wireless network”. In: *Wireless Networks* 15.7 (2009), pp. 901–915.

[NeOz09] Angelia Nedić and Asuman Ozdaglar. “Distributed subgradient methods for multi-agent optimization”. In: *IEEE Transactions on Automatic Control* 54.1 (2009), pp. 48–61.

[NiRoTaVa07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Vol. 1. Cambridge University Press Cambridge, 2007.

[OlRoYo15] Mariann Ollár, Marzena Rostek, and Ji Hee Yoon. *Privacy-Preserving Market Design*. Tech. rep. 2015. URL: `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.701.8490&rep=rep1&type=pdf`.

[OsRu94] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.

[PaVo12] Mallesh M Pai and Rakesh Vohra. *Auction design with fairness concerns: Subsidies vs. set-asides*. Tech. rep. Discussion Paper, Center for Mathematical Studies in Economics and Management Science, 2012.

[PoShTe04] Ryan Porter, Yoav Shoham, and Moshe Tennenholtz. “Fair imposition”. In: *Journal of Economic Theory* 118.2 (2004), pp. 209–228.

[PoWe89] Andrew Postlewaite and David Wettstein. "Feasible and continuous implementation". In: *The Review of Economic Studies* 56.4 (1989), pp. 603–611.

[RaBo07] Bozidar Radunović and Jean-Yves Le Boudec. "A unified framework for max-min and min-max fairness with applications". In: *IEEE/ACM Transactions on Networking (TON)* 15.5 (2007), pp. 1073–1083.

[RaGuGoPs06] Sumit Rangwala, Ramakrishna Gummadi, Ramesh Govindan, and Konstantinos Psounis. "Interference-aware fair rate control in wireless sensor networks". In: *ACM SIGCOMM Computer Communication Review* 36.4 (2006), pp. 63–74.

[RaNo04] Michael Rabbat and Robert Nowak. "Distributed optimization in sensor networks". In: *Proceedings of the 3rd international symposium on Information processing in sensor networks*. ACM. 2004, pp. 20–27.

[Re16] Philip J. Reny. "Nash equilibrium in discontinuous games". In: *Economic Theory* 61.3 (Mar. 2016), pp. 553–569. ISSN: 1432-0479. DOI: `10.1007/s00199-015-0934-3`. URL: `https://doi.org/10.1007/s00199-015-0934-3`.

[Re99] Philip J Reny. "On the existence of pure and mixed strategy Nash equilibria in discontinuous games". In: *Econometrica* 67.5 (1999), pp. 1029–1056.

[ReRe88] Stefan Reichelstein and Stanley Reiter. "Game forms with minimal message spaces". In: *Econometrica: Journal of the Econometric Society* (1988), pp. 661–692.

[Ro16] Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.

[Ro70] R Tyrrell Rockafellar. *Convex analysis*. 28. Princeton university press, 1970.

[RuKuTo99] Dan Rubenstein, Jim Kurose, and Don Towsley. "The impact of multicast layering on network fairness". In: *ACM SIGCOMM Computer Communication Review*. Vol. 29. 4. ACM. 1999, pp. 27–38.

[SaAlEr12] Yunus Sarikaya, Tansu Alpcan, and Ozgur Ercetin. "Dynamic pricing and queue stability in wireless random access games". In: *IEEE Journal of Selected Topics in Signal Processing* 6.2 (2012), pp. 140–150.

[SaMoScWo12] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W. S. Won. "Advanced Demand Side Management for the Future Smart Grid Using Mechanism Design". In: *IEEE Transactions on Smart Grid* 3.3 (Sept. 2012), pp. 1170–1180.

[SaTa02] S. Sarkar and L. Tassiulas. "Fair allocation of utilities in multirate multicast networks: a framework for unifying diverse fairness objectives". In: 47.6 (2002), pp. 931–944. ISSN: 0018-9286.

[SaTa99] Saswati Sarkar and Leandros Tassiulas. "Fair allocation of utilities in multirate multicast networks". In: vol. 37. 1999, pp. 718–727.

[ShLe09] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2009.

[ShTe12] Shrutivandana Sharma and Demosthenis Teneketzis. "Local public good provisioning in networks: A Nash implementation mechanism". In: *IEEE Journal on Selected Areas in Communication* 30.11 (2012), pp. 2105–2116.

[SiAn13] Abhinav Sinha and Achilleas Anastasopoulos. "Mechanism design for resource allocation in networks with intergroup competition and intragroup sharing". In: *CoRR* abs/1307.2569 (2013). URL: `http://arxiv.org/abs/1307.2569`.

[SiAn14a] Abhinav Sinha and Achilleas Anastasopoulos. "A general mechanism design methodology for social utility maximization with linear constraints". In: *ACM SIGMETRICS Performance Evaluation Review* 42.3 (2014), pp. 12–15. DOI: `10.1145/2695533.2695538`.

[SiAn14b] Abhinav Sinha and Achilleas Anastasopoulos. *Generalized Proportional Allocation Mechanism Design for Unicast Service on the Internet*. Tech. rep. 2014. URL: `http://arxiv.org/abs/1401.1760`.

[SiAn15] Abhinav Sinha and Achilleas Anastasopoulos. "A practical mechanism for network utility maximization for unicast flows on the Internet". In: June 2015, pp. 5679–5684.

[SiAn16] Abhinav Sinha and Achilleas Anastasopoulos. "Mechanism design for resource allocation in networks with intergroup competition and intragroup sharing". In: *IEEE Transactions on Control of Network Systems* (Under revision (2016)).

[SiAn17] A. Sinha and A. Anastasopoulos. "Mechanism design for resource allocation in networks with intergroup competition and intragroup sharing". In: *IEEE Transactions on Control of Network Systems* PP.99 (2017), pp. 1–1. ISSN: 2325-5870. DOI: `10.1109/TCNS.2017.2680319`.

[SoZhFa08] Y. Song, C. Zhang, and Y. Fang. "Joint Channel and Power Allocation in Wireless Mesh Networks: A Game Theoretical Perspective". In: *IEEE Journal on Selected Areas in Communications* 26.7 (2008), pp. 1149–1159.

[St02] Steven Stoft. "Power system economics". In: *Journal of Energy Literature* 8 (2002), pp. 94–99.

[StLe06] Tudor M Stoenescu and John O Ledyard. "A pricing mechanism which implements in Nash equilibria a rate allocation problem in networks". In: *45th IEEE Conference on Decision and Control*. IEEE. 2006, pp. 1270–1277.

[StLiTe07] Tudor Mihai Stoenescu, Mingyan Liu, and Demosthenis Teneketzis. "Multirate multicast service provisioning I: an algorithm for optimal price splitting along multicast trees". In: *Mathematical Methods of Operations Research* 65.2 (2007), pp. 199–228.

[Te12] Hamidou Tembine. *Distributed strategic learning for wireless engineers*. CRC Press, 2012.

[Ti89] Guoqiang Tian. "Implementation of the Lindahl correspondence by a single-valued, feasible, and continuous mechanism". In: *The Review of Economic Studies* 56.4 (1989), pp. 613–621.

[Va74] Hal R Varian. "Equity, envy, and efficiency". In: *Journal of economic theory* 9.1 (1974), pp. 63–91.

[VaAn14] D. Vasal and A. Anastasopoulos. "Stochastic Control of Relay Channels With Cooperative and Strategic Users". In: *IEEE Transactions on Communications* 62.10 (Oct. 2014), pp. 3434–3446.

[VaAn15] Deepanshu Vasal and Achilleas Anastasopoulos. *Signaling equilibria for dynamic LQG games with asymmetric information*. Submitted to IEEE CDC 2016. 2015. URL: `http://www-personal.umich.edu/~dvasal/papers/cdc16.pdf`.

[VaAn16a] D. Vasal and A. Anastasopoulos. "A systematic process for evaluating structured perfect Bayesian equilibria in dynamic games with asymmetric information". In: *American control conference*. July 2016. URL: `http://arxiv.org/abs/1508.06269`.

[VaAn16b] D. Vasal and A. Anastasopoulos. *Decentralized Bayesian learning in dynamic games*. Submitted to IEEE Allerton Conference 2016. 2016. URL: `http://web.eecs.umich.edu/~anastas/docs/NetEcon16.pdf`.

[Wa10] ShiGuang Wang, Ping Xu, XiaoHua Xu, ShaoJie Tang, XiangYang Li, and Xin Liu. "TODA: truthful online double auction for spectrum allocation in wireless networks". In: *IEEE Symposium on New Frontiers in Dynamic Spectrum*. 2010, pp. 1–10.

[WaWuJiLiCl08] B. Wang, Y. Wu, Z. Ji, K. J. R. Liu, and T. C. Clancy. "Game theoretical mechanism design methods". In: *IEEE Signal Processing Magazine* 25.6 (Nov. 2008), pp. 74–84.

[WaXuAhCo15] T. Wang, Y. Xu, S. D. Ahipasaoglu, and C. Courcoubetis. "Max-min fairness of generalized AGV mechanisms". In: *2015 54th IEEE Conference on Decision and Control (CDC)*. Dec. 2015, pp. 5170–5177.

[WeMaOz14] Ermin Wei, Azarakhsh Malekian, and Asuman Ozdaglar. "Competitive Equilibrium in Electricity Markets with Heterogeneous Users and Price Fluctuation Penalty". In: *LIDS report* 2918 (2014).

[YaHa05] Sichao Yang and Bruce Hajek. *Revenue and stability of a mechanism for efficient allocation of a divisible good*. Tech. rep. University of Illinois Urbana-Champaign, 2005. URL: `http://www.ifp.illinois.edu/~hajek/Papers/YangHajek06.pdf`.

[YaHa07] Sichao Yang and Bruce Hajek. "VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals". In: *IEEE Journal on Selected Areas in Communications* 25.6 (2007), pp. 1237–1243.

[Yo04] H Peyton Young. *Strategic learning and its limits*. Oxford University Press, 2004.

[ZhGu11] Yan Zhang and Mohsen Guizani. *Game theory for wireless communications and networking*. CRC Press, 2011.