Towards Improving Allocative Efficiency in Games and Markets

By

Jian Lou

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Computer Science

August 9, 2019

Nashville, Tennessee

Approved:

Yevgeniy Vorobeychik, Ph.D.

Xenofon Koutsoukos, Ph.D.

Abhishek Dubey, Ph.D.

Bradley Malin, Ph.D.

Myrna Wooders, Ph.D.

I dedicate this work to my parents.

ACKNOWLEDGMENTS

It is an incredible experience to move to the United States and pursue a Ph.D. with all the help from my colleagues, friends, and family. My thanks are due first and foremost to my advisor, Prof. Yevgeniy Vorobeychik, for his guidance and continuous support throughout my Ph.D. career. I cannot imagine that I could finish my Ph.D. without Prof. Vorobeychik's consistently patient guidance and good advice. His sharp sense for research and positive attitude will always be an inspiration to me. I also want to thank my committee members, Prof. Myrna Wooders, Prof. Xenofon Koutsoukos, Prof. Bradley Malin and Prof. Abhishek Dubey. They are always patient with my questions and taught me how to think about problems in different perspectives. Without the discussion with them, I will not be able to come up with all these wonderful ideas.

I would like to express my heartfelt thanks for the help I received during my internship at Alibaba Group. Special thanks go to Huan Xu who was my mentor during my internship for offering me numerous help on my research projects. I also would thank all the help from my coworkers at Alibaba, including Sen Yang, Hao Yu, Zhenliang Zhang, Zirun Zhang, Ming Lin, Yuanhong Xu, Qingsong Wen, Mingrui Liu, Li Wang, Yi Xu, Yuanhang Su, Pichao Wang, Yuhong, etc.

I am always enjoying discussing problems and chatting with my lab mates. It is a wonderful experience to work with these awesome guys, including Bo Li, Ayan Mukhopadhyay, Haifeng Zhang, Yi Li, Liang Tong, Sixie Yu, Swetasudha Panda, Rajgopal Venkatasaramani, Liyiming Ke, etc. I also appreciate the help from my coauthors, Yang Liu, Andrew Smith, Aron Laszka, Chen Hajaj, Greg Leo, Martin Van der Linden, Matthew Chambers.

I will always value the friendship during my Ph.D., and I want to thank my friends, including Qishen Zhang, Chao Yan, Rui Wang, Xuanli Deng, Wenxuan Li, Shunxing Bao, Ming Yan, Weizhuang Peng, Mengtang Li, Xiao Li, Tianjiao Wang, Xiaochen Yang, Dongqing Zhang, Dana Zhang, Fangzhou Sun, Yongtai Liu, Yalong Li, Junren Wang,

Sifang Zhao, Dun Liu, Junyuan Lin, Hui Su, Huijie Li, Can Zhang, Yizhou Wang, Pengfei Wang, Jianing Wang, Weihan Wang, Huahong Zhang, Feiyang Cai, etc.

Finally, I am grateful to my family for having supported me for my doctoral endeavor no matter what difficulties I face.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

Introduction

My focus in the thesis is the efficiency/inefficiency introduced by players' selfish behavior in games and markets, and the way to deal with the inefficiency. To model users' selfish behavior, I adopt *game-theoretic* analysis, which is widely used to model outcomes of strategic interactions among multiple self-interested players. To deal with the inefficiency due to players' selfish behavior, I employ the *mechanism/market design* paradigm, which studies how to design the rules of encounter which lead to socially desirable equilibrium outcomes, such as to achieve high allocative efficiency. In the thesis, my goal is to study some situations which lead to inefficiency in equilibrium outcomes, as well as approaches for mitigating such inefficiency.

Broadly speaking, game theory studies mathematical models of strategic interaction among rational decision-makers with perfect hindsight. In the thesis, I mainly consider *non-cooperative games*, in which players are self-interested. Due to the selfish behavior of players, the efficiency of a system may degrade. One well-known example of the inefficiency is the *prisoner's dilemma*, which shows why two completely rational individuals might not cooperate, even if it appears that it is in their best interests to do so. In the thesis, I *quantitatively* measure how the efficiency of a system degrades due to the selfish behavior of its agents in some security domains, in terms of the *price of anarchy*, which is also a general notion that can be extended to diverse systems and notions of efficiency. I also note that the equilibrium introduced by selfish behavior can be efficient in some security domain, which will be shown with full details in later chapters.

Having observed some examples where equilibria can be very inefficient, I explore the problem of designing the rules of encounter, such as mechanisms and markets, that lead to greater efficiency. Specifically, mechanism/market design paradigm takes an engineering approach to designing economic mechanisms or markets, toward desired objectives, in

strategic settings, where players act rationally. The objectives of mechanism/market design consist of both economic and computational requirement. The economic requirement includes efficiency, fairness, incentive compatibility, etc; and the computational objectives consist of computational complexity, time and space efficiency, etc. Note that it may be computationally hard for some mechanisms with good economic properties, and the trade-off between economic and computational requirement is often considered when we are designing mechanisms or markets in the real application. In the thesis, I consider both economic and computational efficiency in designing mechanisms to partition a collection of individuals into teams (coalitions). I also leverage some computational techniques to reduce the computational burden for some team formation mechanisms with good economic properties but dissatisfactory in computation. Furthermore, I also study how it could improve efficiency by introducing a secondary market among previously independent players. By employing game theoretical analysis, I show that introducing a second market can often improve efficiency.

In the following, I will mainly introduce the motivation and background of the problems I study, and my contribution to these problems.

I.1 Towards Efficiency in Security Games

In the thesis, I first focus on efficiency/inefficiency in security scenarios. To analyze a security domain, game theory has come to play an important role in it, with considerable modeling and algorithmic advances, as well as the actual deployment of security systems in practice that is based on such models and algorithms, including LAX Airport [3, 4], US Coast Guard [5], and the Federal Air Marshals Service [6, 7, 8], among others.

A popular game-theoretic model of security that has received much attention both in the research and in practice is a Stackelberg game between a single defender and a single attacker, in which the defender commits to a randomized strategy, while the attacker, upon learning this strategy, chooses an optimal target or a subset of targets to attack [9]. In

most of the associated literature, it is assumed that a single defender is responsible for all the targets that need protection, and that she has control over all of the security resources. However, there are many domains in which there are multiple defender agencies who are in charge of different subsets of all targets. While sometimes such agencies can be aligned to follow the same set of goals, in general, different defender entities exhibit at least some disparities in goals. In the thesis, I focus on the scenarios in which there are multiple defenders who are rational and selfish, and analyze the efficiency/inefficiency introduced by defenders' strategic behavior.

I.1.1 Multi-Defender Security Games

I consider the security scenarios with multiple defenders, where each defender protects multiple targets. In this setting, I theoretically characterize Nash and approximate Nash equilibria, as well as their efficiency.

In the literature on security games, a defender is typically responsible (financially, politically, or legally) for targets in their direct charge, rather than other targets that may have social importance. This is certainly the case for the private sector, where different corporations secure their own resources without necessarily much concern for those of others, but is also common for the public sector, with different government agencies held accountable for their own assets, and not for those of others. In such *non-cooperative* security scenarios, the typical single-defender Stackelberg game model is clearly inadequate. Instead, we must consider the consequences of *strategic interactions* among *multiple defenders*, each charged with protecting their assets from common adversaries. An important consideration in such games is the *negative externalities* that security decisions impose on others: specifically, when a defender chooses a high level of security investment, budget-constrained attackers are more likely to choose others to attack. The resulting dynamics are likely to lead to over-investment in security, a phenomenon observed in several related efforts [10]. In other words, the selfish behavior among defenders may incur severe degradation of effi-

ciency.

In the thesis, I characterize Nash and approximate equilibrium among defenders and theoretically analyze the efficiency degradation due to defenders' selfish behavior in multi-defender security games.

I.1.2 Multi-Defender Against Spear-Phishing Attacks

I also study the strategic interaction among multiple defenders in spear-phishing attacks. Spear-phishing attacks pose a serious threat to sensitive computer systems, since they sidestep technical security mechanisms by exploiting the carelessness of authorized users. A number of high-profile targets have fallen victim to spear-phishing attacks. In 2013, Target, the second largest general merchandise retailer in the US, suffered a massive data breach due to a spear-phishing attack [11]. As a consequence, Target had to pay Visa issuers $67 million as reimbursement, and it is reportedly working on a similar deal with MasterCard [12]. In 2014, the corporate network of a German steel mill was infiltrated by a spear-phishing attack [13]. The attackers manipulated and disrupted control systems, resulting in massive physical damage. Further examples include one of the White House internal networks [14], computers at the Nuclear Regulatory Commission [15], and Oak Ridge National Laboratory [16].

To mitigate spear-phishing attacks, an organization may set up an e-mail filter, which assigns a maliciousness score to each incoming e-mail and delivers only those that are below a given threshold [17]. Unfortunately, scoring is inevitably imperfect, and threshold choice must necessarily balance security (risk of delivering malicious e-mails) and usability (blocking of benign traffic).

Unlike non-targeted malicious e-mails, such as spam, spear-phishing e-mails must be customized to their targets, which means that an attacker must spend a substantial amount of effort on each target [18]. Consequently, attackers can only target a limited number of users in any spear-phishing campaign. This limitation implies that an attacker must select a

subset of targets to maximize expected yield from an attack. Moreover, resource limitation on the attacker links the decisions of otherwise independent defenders: filtering decisions by some may result in others being targeted. If a single organization were responsible for setting filtering thresholds for all users, it could optimally account for such interdependencies, as shown in prior work [1, 2]. Realistically, however, numerous organizations are typically targeted, and their goals are generally distinct. The externalities that users impose upon one another therefore become strategically significant, and no work to date analyzes the resulting strategic dynamics in the spear-phishing context, even though prior work has considered other, quite different, interdependent security problems [19, 20, 21, 22].

In the thesis, I study players' strategic behavior when there are multiple selfish defenders against spear-phishing attacks, analyze the existence and the efficiency of equilibrium, and propose a polynomial algorithm to find such an equilibrium if it exists.

I.1.3 Decentralization and Security in Traffic Light Control

I consider the decentralized control system as a real application of the multi-defender security game model. Effective design of large-scale complex traffic control systems, involving many controlled intersections, is fundamental in modern urban centers. As a result, this problem has been considered extensively in prior literature spanning fields such as transportation, operations research, economics, and computer science. One broad class of approaches involve the study of self-organized phenomena in many-particle systems, such as traffic flows on highways [23, 24]. In order to explain phenomena such as the emergence of traffic jams or stop-and-go waves, a variety of different traffic flow models have been proposed, including follow-the-leader models [25] and fluid-dynamic traffic models in both discrete and continuous space. Recently, research has focused on network traffic, extending one-dimensional traffic models in order to cope with situations, where traffic flows merge or intersect. These models can explain how jam fronts propagate backward over network nodes.

One grand challenge in this connection is the optimization of traffic lights in urban road networks [26], especially the coordination of vehicle flows and traffic lights. A common goal in this literature is to minimize the travel times in a traffic network. In previous work, it was shown that a further improvement of the traffic flow requires us to apply more flexible strategies than fixed-time controls [27]. Gershenson and Rosenblueth [28], for example, showed for a regular network with periodic boundary conditions that their control strategy synchronizes traffic lights even without explicit communication between them. Lammer et al. [29] proposed to represent the traffic lights by locally coupled phase oscillators, whose frequencies adapt to the minimum cycle of all nodes in the network. Other algorithms perform parameter adaptations by means of neural networks [30], genetic reinforcement learning [31], fuzzy logic[32], or swarm algorithms [33].

Although adaptive, state-aware strategies can offer tremendous gains in traffic control efficiency, they expose an attack surface that can be exploited to substantially increase congestion. For example, a common kind of adaptive control logic involves state captured by vehicle queue lengths in each direction, with light switching between red and green as a function of relative queue lengths. While such state-aware switching can significantly increase efficiency, they also expose the vulnerability of controllers to attacks on sensors from which queue length information is derived.

An additional consideration which is crucial in modern complex traffic networks is that traffic lights on the network are often designed by multiple actors (e.g., municipalities). Consequently, while in principle we may be able to design extremely efficient and resilient controllers for a particular traffic network, this is impractical due to misalignment of interests among the different parties that actually control such networks. In the thesis, I employ game theoretic analysis on the traffic control system and design scalable algorithm for approximating a Nash equilibrium in the system.

I.2 Toward Efficiency in Coalition Formation

The second part of the thesis is about designing *efficient* coalition formation mechanisms among multiple selfish players. Besides the economic efficiency, the mechanisms need also to meet some other economic and computational requirements, such as fairness, incentive compatibility, and computational efficiency.

Division of individuals into groups is a common task, important in a multitude of economic and social problems. Examples include dividing students into study groups or dorms, forming teams for a basketball game, or forming groups for carpooling. The issue of team, or coalition, formation in domains with hedonic preferences (players only care about the members of their own team) is commonly studied from the perspective of stability, where the focus is on characterizing or computing solutions of the game, such as the core [34, 35, 36]. In the thesis, I consider the coalition(team) formation problem mainly from the (centralized) mechanism design perspective. Then I specifically discuss the roommates problem, a special case of coalition formation, and adopt automated mechanism design approach to deal with it.

I.2.1 Coalition Formation Mechanisms

In the thesis, I consider coalition formation as a mechanism design problem in which central authority is in charge of forming coalitions based on players' reported hedonic preferences. A challenging aspect of this mechanism design problem is that players may seek to benefit by misreporting their true preferences. It is further complicated if the mechanism is required to satisfy additional desiderata, such as individual rationality (players have incentives to participate), matching of soulmates (any collection of players who mostly prefer to be with one another are always matched), efficiency, and fairness. Indeed, in general, this combination of desiderata is impossible to achieve, even in two-sided matching problems [37], which is a subclass of the coalition formation mechanism design problem. Alcalde and Barberà [38] point out that without restrictions on the sets of admissible pref-

erences, there is no matching mechanism that is Pareto efficient, individually rational, and incentive compatible. In general, it is also impossible to design coalition formation mechanisms that are both incentive compatible and match soulmates [39].

Two special cases of the coalition formation mechanism design problem have received considerable attention: two-sided matching markets [37], such as matching medical school residents with residency programs, and one-sided matching or assignment problems [40, 41], such as school choice and course allocation (the latter abstracted as combinatorial assignment). Well-known mechanisms for one- and two-sided matching, such as the deferred acceptance mechanism [37], possess many of the desired properties, but even generalizing to combinatorial assignment runs into numerous impossibility results [41]. Indeed, even in the roommate problem [42] where arbitrary teams of pairs can be constructed, few known positive results exist. In the general coalition formation problems, *random serial dictatorship (RSD)* is, to our knowledge, the only mechanism which is incentive compatible and ex post Pareto efficient [43, 44]. Although Wright and Vorobeychik [44] present several other mechanisms, these do not satisfy any of the mentioned desiderata, making RSD the only theoretically grounded mechanism known for general coalition formation.

In the thesis, I mainly consider the coalition formation problem from the perspective of (centralized) mechanism design. Our mechanisms are constructed using subgame perfect Nash equilibria (SPNE) of an *accept-reject game (ARG)* in which players propose coalitions in a predetermined order. Unlike RSD, prospective teammates may choose to either accept or reject the proposals in ARGs. Chamber et al. [45] show that the SPNE of ARG is individual rational and implements *iterated matching of soulmates* [39], where soulmate coalitions are matched in an iterative fashion. Chamber et al. [45] also demonstrate that the SPNE of Rotating Proposer Game (RPG), which is a special case of ARG, is Pareto efficient. Rotating Proposer Mechanism (RPM) is a mechanism that implements the SPNE of the corresponding RPG. As the SPNE is highly nontrivial to compute, I mainly focus on the computation of the equilibrium. Other than the computational issues, I evaluate the

RPM by employing empirical methods. And I show that RPM has very good efficiency and incentive properties.

I.2.2 Mechanism Design for the Roommates Problem

Matching problem is a special case of coalition formation problem, and has received considerable attention across multiple research areas in economics, operation research, and computer science, starting with the seminal paper by Gale and Shapley [46] which first introduced two-sided matching, along with the *deferred acceptance (DA)* algorithm for finding a stable matching. Among its most prominent applications is the National Resident Matching Program (NRMP). In addition to the matching problem, Gale and Shapley [46] also introduced a generalization, the *roommates* problem, in which any player can match with any other. While potential applications of the roommates problem abound, such as pairing police officers on patrols or pilots on flights [47], holiday home exchanges [48], kidney exchange [49], students to share double rooms in colleges, and course project teams, there is still no widely accepted mechanism design solution for it.

I consider a new perspective on the roommates problem based on *automated mechanism design (AMD)* [50]. In a prototypical AMD setup, one obtains preferences from the players, and then solves an optimization problem (for example, an integer linear program) in which constraints ensure incentive compatibility. However, applying AMD to matching problems in general, and the roommates problem in particular, faces a number of challenges. First, it is conventional to consider ordinal, rather than cardinal preference reports by the players. Second, incentive compatibility is often incompatible with other highly desirable properties, such as stability and, for cardinal preferences, optimal social welfare. Third, standard AMD methods explicitly compute the full mapping from preference reports to outcomes, which is intractable for even small roommates problems.

To address the challenges, I implement a rank-preserving transformation from ordinal to cardinal preferences, after ordinal preferences are received, and before the AMD approach

is applied. Then I consider relaxing incentive compatibility, which bounds the most one can gain from lying, and propose several approaches to construct integer linear programs for computing outcomes. I will show that these approaches have good theoretical or empirical properties.

I.3 Towards Efficiency by Introducing a Secondary Market

In the third part of the thesis I study how a free-trading secondary market could mitigate players' demand uncertainty in the primary market, and make the outcome more efficient.

Consider the following setting, increasingly common in international business. A computationally intense corporation is allocating computational resources (e.g. cloud services, CPUs, GPUs, or memories, etc.) to its business units (BUs). These BUs are affiliated with the corporation, but decentralized management practices ensure that these are operating *independently* in various decision-making tasks. The goal of the computational resource manager is to meet BUs' requirement instead of making decisions for them. At the beginning of each season, each BU reports the number of resources it needs to the computational resource manager of the corporation. Based on BUs' requests, the resource manager orders the corresponding amount of resources and allocates them to each BU based on the reported amount, and each BU is responsible for paying for the resources they get. Since these orders entail commitments at the beginning of a season, BUs are unable to perfectly predict their need as the season progresses: actual realized demand depends on a host of stochastic factors, such as growth or recession, expansion of the BU, changes in employee structure, etc. However, if BUs run short of resources during the season, the company will often impose stiff penalties, financial and otherwise (e.g., poor performance evaluation of the BU), as this has both significant reputational and financial repercussions for the corporation. As a consequence of the combination of uncertainty in true demand, and stiff penalties following a shortage, BUs have a strong incentive to *over-request* resources from the resources manager. This naturally leads to resource abundance among the BUs, but also

to waste and unnecessary costs for the company.

In the thesis, I model the BUs as players, and I leverage game-theoretic analysis to study the influence of a *secondary market* on players' decisions and social welfare in the resource allocation problem above.[1] Secondary markets have been successfully used in several other settings, such as spectrum market [51], emission trading market [52], and energy market [53]. Unlike the approaches above, however, I study the extent to which the secondary market can mitigate demand uncertainty when information is *complete* and all players are *strategic*. To this end, I use a well-known newsvendor model as a starting point to construct a two-stage game model:

1. In the first stage (i.e. primary market), players report their orders to the authority, and then pay and get the requested resources. In this stage, actual demand is uncertain.

2. In the second stage (i.e. secondary market), demand uncertainty is resolved, all demanded resources used, and remaining resources can be traded freely.

Crucially, if no secondary market exists, the game devolves into a collection of independent decisions by all players in the newsvendor framework. The secondary market, on the other hand, creates an implicit dependency among optimal decisions in the first stage. Our ultimate focus is on characterizing equilibrium decisions in this first stage of the two-stage game.

In the thesis, by employing game theoretic analysis, I show that introducing a secondary market among these players could indeed mitigate the demand uncertainty, sometimes significantly. In particular, the resource allocation can be also much more efficient with a secondary market than without.

[1]The focus on the free-trading market also distinguishes our work from those market design works that involve invoking mechanism design approaches.

CHAPTER II

Related Work

II.1 Stackelberg Security Games

Stackelberg game is a strategic game in which the leader moves first and then the follower moves sequentially, and it is named after the German economist Heinrich Freiherr von Stackelberg. In the early work, von Stackelberg [54] shows that in Cournot' duopoly model [55], if one firm is able to commit to a production quantity first, that firm will do much better than in the simultaneous-move (Nash) solution. In the computer science community, Conitzer and Sandholm [9] consider the *Stackelberg model* and study how to compute optimal strategies to commit to under both commitments to pure strategies and commitment to mixed strategies, in both normal-form and Bayesian games. They show that in two-player normal-form games, an optimal mixed strategy to commit to can be found in polynomial time using linear programming. However, finding an optimal pure strategy to commit to in two-player Bayesian games is NP-hard, even when the follower has only a single type.

In the security domain, it is reasonable to assume that a defender firstly commits to a (mixed) strategy and an attacker best responds it. Considering the uncertainty of players, these domains are commonly modeled as Bayesian games. Due to the NP-hardness of computing the optimal strategy for defenders, there are many works that try to develop efficient algorithms to deal with it. Paruchuri et al. [56] provide an efficient heuristic approach for security against multiple adversaries. In another work, Paruchuri et al. [57] present an efficient exact algorithm for finding the optimal strategy for the leader to commit to in Bayesian Stackelberg games. The algorithm, DOBSS, is also at the heart of the ARMOR system that was being deployed for security scheduling at the Los Angeles International Airport (LAX) [58]. Afterwards, Kiekintveld et al. [59] develop new models and algorithms that scale to much more complex instances of security games. The key idea is to

use a compact model of security games, which allows exponential improvements in both memory and runtime relative to previous algorithms. Based on the same idea of strategic randomization in [59], Tsai et al. [60] implement Intelligent Randomization In Scheduling (IRIS) system, a software scheduling assistant for the Federal Air Marshals (FAMs) that provide law enforcement aboard U.S. commercial flights.

The success of Stackelberg security games in ARMOR (in LAX) and IRIS (in FAMs) attract significant interest in game-theoretic approaches to security. And the Stackelberg model has been adopted to a lot of real-world situations, such as protection of fisheries [61], patrolling to protect ferries [62], protection of forest land [63] and wildlife [64], etc. Other than those applications, Stackelberg model has been also used to study some problems those look nothing like typical security games, such as adversarial machine learning [65], privacy-preserving data sharing [66], vaccine design [67], etc.

II.2 Security Games with Multiple Defenders

It is assumed that there are two players in the Stackelberg game, and a group of players are often modeled as a single agent when there are multiple players in the game. However, there are some works studying the scenarios that are not suitable to be modeled as two-player games. Jiang et al. [68] considered (mis)-coordination in cases where there are multiple defenders who are responsible for different sets of targets and share the common utility function over all targets. In this work, the defenders are fundamentally cooperative (sharing identical goals). Bachrach et al. [10] examined non-cooperative security games among many defenders, in a two-stage model, but imposed strong assumptions on the model structure, and only considered one-dimensional continuous "security investment" strategies for the defender (departing significantly from the typical structure of Stackelberg security games, in which defensive strategies are discrete protection choices).

Among the earliest multi-defender models is in the literature on *interdependent security games* [69], in which interactions among multiple defenders are modeled as an n-player,

2-action game, where a player decides whether to invest in security; however, no attacker is considered. More recently, time-dependent scenarios where coordination of defender resources amongst multiple defenders is assumed to have been studied using Markov decision processes [70]. Since total cooperation is assumed, this model effectively reduces to a single defender game in which the defender controls all resources. A natural extension of interdependent security games, *interdependent defense games* [71], does consider an attacker who acts *simultaneously* with the defenders, rather than after observing the joint defense configuration. Interdependent defense games have also been studied in the context of traffic infrastructure defense [72]. Two recent efforts studying mult-idefender games explicitly model interdependence among targets through a probabilistic contagion process [73, 74]. Like our work, they consider attackers who observe the joint defense prior to making a decision, but each defender is restricted to secure a single node, and the strategy space is assumed to be continuous. Vorobeychik et al. [75] attempt to study strategic settings related to security in which each player's decision space is combinatorial. However, this work does not consider a strategic attacker.

In security games with multiple defenders, players' decisions are often correlated with each other, and the interdependence among players can be often modeled as a network structure. They motivate research on methods for improving information security, such as network design. Individuals derive benefits from their connections, but these may expose them to external threats. Cerdeiro et al. [76, 77] propose a model to explore the tension between connectivity and exposure to an external threat when security choices are decentralized. They find that faced with an intelligent adversary who seeks to minimize network value, both over-investment and under-investment in security are possible. Social welfare may be maximized in sparsely connected networks when under-investment pressures are present, and fragmented networks when over-investment pressures prevail. Their result is very similar to that in [78], though they are analyzing the problem from different perspectives.

II.3 Security Games for Defending Against Cyber Attacks

Recently, there are more and more works that apply security game theory to the defense against cyber attacks. One representative example is applying the Stackelberg game to solve the spear-phishing attack. Jain et al. [1] model the decision problem faced by a single defender who has to protect multiple users against targeted and non-targeted malicious e-mail. They focuses on characterizing and computing optimal defense strategies, and they use numerical results to demonstrate that strategic threshold selection can substantially decrease losses compared to naïve thresholds. Zhao et al. [79] study a variant of the previous model: they assume that the targeting attacker can launch an unlimited number of costly spear-phishing attacks in order to learn a secret, which only a subset of the users know. This work also focuses on the computational aspects of finding an optimal defense strategy; however, this variant of the model does not consider non-targeted malicious e-mails.

Classifying e-mails in order to estimate their likelihood of being malicious has been extensively studied [80]. Note that these results are complementary to the strategic threshold-selection problem, since the latter builds on an exogenously given classifier. Potentially malicious e-mails can be classified based on many attributes. For example, Fette et al. [81] build a classifier for detecting phishing e-mails using a variety of features, such as the number of links in the e-mail and the age of the linked-to domain names. When evaluated on a real-world dataset, the false negative rate of the classifier was less than 4%, while its false positive rate was around 0.1%. As another example, Bergholz et al. [82] design an e-mail classifier for detecting spam and phishing e-mails, and they describe a number of novel features, such as design elements of known brands and intentional distortion of content not perceivable by the reader.

Other than spear-phishing attacks, researchers also adopt security game theoretical model to other cyber attack defense problems. Laszka et al. [83] model intrusion-detection systems as an attacker-defender security game and study the problem of finding optimal intrusion detection thresholds. Intrusion-detection systems can play a key role in protect-

ing sensitive computer systems. However, an over-sensitive intrusion-detection system, which produces a large number of false alarms, imposes prohibitively high operational costs. Laszka et al. [83] try to optimize the sensitivity of intrusion detection systems and try to balance between maximizing security and minimizing costs. Schlenker et al. [84] study the scenario in which the attacker could use deceptive techniques to attack a cyber network system, and introduce a game theoretical model of deceptive interactions between a defender and a cyber attacker, which is called the Cyber Deception Game. Furthermore, they consider the computational complexity issues for both defender and attacker. Li et al. [85] propose a game theoretical model in Man-in-the-Middle (MITM) attack, and model the strategic interaction between the Man-in-the-Middle (MITM) attacker and multiple defenders as a simultaneous-move game. They also provide the theoretical analysis of the uniqueness of Nash equilibrium, and propose practical learning algorithms for the defenders and the attacker.

II.4 Hedonic Coalition Formation

In recent years, hedonic coalition formation has been extensively investigated in Economics and AI literature. A hedonic coalition formation game [34, 86] is a game that models the formation of coalitions (teams) of players when players have preferences over which group they belong to.

Coalition formation problem has been studied since the early age of multi-agent systems. When agents communicate in the multi-agent systems, they may decide to cooperate on a given task or for a given amount of time. Zlotkin and Rosenschein [87] consider the coalition formation problem from the perspective of agents negotiation. They present a coalition formation mechanism that uses cryptographic techniques for subadditive task oriented domains. Since this work is conducted in a game theory framework, agents consider the utility of joining a coalition in which they are bound to try to advance the utility of other members in exchange for reciprocal consideration. Shehory and Kraus [88] present a dis-

tributed algorithm for task allocation when coalitions are either needed to perform tasks or more efficient than single agents. Sandholm and Lesser [89] use a vehicle routing domain to illustrate a method by which agents can form valuable coalitions when it is intractable to discover the optimal coalitions. There are also some more following works that consider agents' cooperation in coalition formation (such as [90, 91, 92]).

Besides the settings with cooperative agents, much research is also focused on self-interested players aiming to maximize their utility. Bogomolnaia and Jackson [34] study the existence of stable coalition structures, and provide, among other results, restrictions on preference profiles that ensure the non-emptiness of the core. Besides the notion of core stability – a traditional notion of stability – they also consider some other notion of stability, such as individual stability, Nash stability, and contractual individual stability. The problem of existence of (core, Nash, individually and contractually individually) stable coalitions is also considered in other work, such as [93]. A potentially infinitely long coalition formation process in the context of hedonic games was studied in [94]. There are some more works those study the stability in coalition formation, and more literature can be seen in [86].

The literature considering coalition formation from a mechanism design perspective has been relatively limited. Several efforts consider this problem within a restricted set of coalitions. Different from dealing with preference constraints in [34], Pápai [95] considers it as a generalization of more specific coalition formation model, such as the marriage [46] and roommate models [42]. Rodriguez-Alvarez [96] introduces *single-lapping* property, which is the sufficient and necessary condition of unique stability. Moreover, she also shows that single-lapping rules are the only rules that satisfy strategy-proofness, individual rationality, and Pareto efficiency when agents' preferences over coalitions are not restricted. Banerjee et al. [97] introduce the *top-coalition* property and they prove that top-coalition property is sufficient to guarantee the existence of a unique core coalition structure. Aziz et al. [98] identify a close structural connection between Pareto optimality and perfection that has various algorithmic consequences for coalition formation. Based on this insight, they

formulate an algorithm that computes an individually rational and Pareto optimal outcome in hedonic games. Wright and Vorobeychik [44] consider the general coalition formation problem with cardinality being the only constraints on coalitions. They also propose several mechanisms for this problem; however, none are incentive compatible or Pareto efficient with the exception of random serial dictatorship [99], which is both. The problem of incentive compatibility has been also considered in some other literature [100, 101, 96, 102], albeit either in very restricted domains (such as single-lapping coalitions), or when monetary transfers are allowed.

II.5 Matching and Roommates Problem

Since the seminal paper by Gale and Shapley [46], stable matching problems have been well studied in economics and recently in computer science and artificial intelligence community. Algorithms for finding solutions to the stable matching problem have applications in a variety of real-world situations. One of the best-known application of the matching problem is the National Resident Matching Program (NRMP), which assigns graduating medical students to their first hospital appointments.

In a marriage problem (two-sided matching), agents are divided into two disjoint groups where agents can only be matched to an agent in the other group. Gale and Shapley [46] prove that stable outcomes always exist in two-sided matching problem. And they presented *deferred acceptance (DA) algorithm* to find a stable outcome. Even though DA mechanism is not incentive compatible, the stability property is very desirable and it is stronger than individual rationality or Pareto efficiency. In fact, it has been shown that incentive compatibility and stability are not compatible with two-sided matching [103]. What's more, Alcalde and Salvador Barberà [38] prove that, without restrictions on the sets of admissible preferences, there is no matching mechanism that is Pareto efficient, individually rational, and incentive compatible. In reality, the original marriage model can be generalized to many-to-one [104] or many-to-many two-sided matching [105]. And in-

centive compatibility is not a big issue for the two-sided matching model in a real-world situation. For instance, Kojima and Pathak [106] analyze the scope for manipulation in many-to-one matching markets, and they show that the fraction of participants with incentives to misrepresent their preferences when others are truthful approaches zero as the market becomes large.

Gale and Shapley [46] also defined the roommates problem, in which all agents are from a single group, and any agent can be matched with any other. Comparing with extensive literature on the marriage problem, roommates problem has been much less studied. However, there are many applications that can be modeled as a roommates problem: pairing police officers on patrols or pilots on flights [47], holiday home exchanges [48], kidney exchange [49], students to share double rooms in colleges, class project teams assignment, etc. What's more, roommates problem also boils down to hedonic coalition formation problem [34, 86]. Gale and Shapley [46] also notes that stable outcomes do not necessarily exist in roommates problem. Irving [42] presents a polynomial algorithm to find a stable outcome if it exists in roommates problem. In some literature, it is common practice to restrict the analysis to those problems in which a stable matching exists (see for instance, [107, 108, 109, 110]). However, restricting attention to the roommate problem which has stable outcomes means ignoring a subclass of problems without stable matchings. In fact, in [111] they show that as the number of agents increases, the probability of not resulting in stable outcomes for a roommate problem increases fairly steeply, which makes predicting the outcome of the problem much more challenging than marriage problem. Some other works in roommates problem try to propose some other solution concept, which is guaranteed to exist, in roommates problem, such as Q-stable matching [112].

Given the incompatibility between incentive compatibility (i.e., agents report their preferences truthfully) and stability in marriage problems, it is even harder to design incentive compatible mechanism for the roommates problem. Thus, things can be more severe if we are designing some centralized roommates mechanism in which truthful preferences are

expected to be received. Some literature tries to study the constrained manipulation in the two-sided matching problem. In two-sided matching, the most commonly studied model of manipulation in this literature is *truncation* [113], whereby one removes some of the least preferred partners from a rank order. More pertinent to our work is *permutation manipulation* [114], in which a player can permute her true preference. Vaish and Garg [114] study such manipulations in the context of deferred acceptance, and they show that deferred acceptance can be manipulated by both permutations and promotions (where a player shifts someone up in their preference order).

II.6 Secondary Markets

The idea of introducing a secondary market is not new. In the literature, auctions [115] [116] have been used extensively as the mechanism for conducting trades in a secondary market [51]. One such example is the secondary market design for *spectrum auction* [117]. In practice, a significant amount of wireless spectrum is under-used by current owners. To enable better use of the spectrum, the auction approach was extensively adopted to dynamically allocate the spectrum in a secondary market. In this line of works, different designs of auctions have been proposed to fulfill the different requirement of spectrum auction, such as improving social welfare, guaranteeing strategyproofness, fairness, etc (see [118, 119, 120]). For instance, Kash et al. [51] propose an auction approach that leverages dynamic spectrum access techniques to allocate spectrum in a secondary market. These are markets where spectrum owners can either sell or lease spectrum to other parties.

Another salient application of secondary market is that of emission trading, which is a market-based approach to controlling pollution by providing economic incentives for achieving reductions in the emissions of pollutants [121]. Greenhouse gas emissions trading schemes (ETSs) are operational in several countries. The trade in carbon permits or credits within and between ETSs is growing [122]. In the ETSs, the authority (usually the government) allocates or sells permits to discharge specific quantities of a specific pollutant

per time period [123]. Polluters who want to increase their emissions could buy permits from others that are willing to sell in a secondary market [124]. Although there is some difference between the emission trading setting and mine, the results in my analysis could also give some managerial insight and implication for emission trading system.

My work is also closely related to game theoretical analysis in inventory management and supply chain. To my best knowledge, Parlar [125] was the first to analyze the game theoretic version of the newsvendor problem with two retailers competing on product availability, which is also one of the first articles modeling inventory management in a game theoretic framework. Following this line of research, Bernstein and Federgruen [126] investigate the equilibrium behavior of decentralized supply chains with competing retailers under demand uncertainty. Other than the noncooperative game, Hartman et al. [127] consider a cooperative inventory-"centralized" game among n stores with single-item and single-period demands. They examine the conditions under which such an inventory centralization game has a nonempty core. Muller et al. [128] prove that the core is always nonempty for all possible joint distributions of the random demands in a cooperative newsvendor game. Cachon and Netessine [129] survey the applications of game theory to supply chain analysis and outlines game-theoretic concepts that have a potential for real-world application. More recently, Fiestras-Janeiro et al. [130] provide a review of the applications of cooperative game theory in the management of centralized inventory systems.

Lee and Whang [131] discuss the impact of the secondary market on supply chain problem, which is very relevant to mine. Lee and Whang develop a two-period newsvendor model with a single manufacturer and many resellers. At the beginning of the first period, resellers order and receive products from the manufacturer; then in the second period, resellers can trade inventories among themselves in the secondary market. One main difference that distinguishes my work from [131] is that we aim to understand how secondary market mitigates the uncertainty in demands when allocating resources; in contrast, the main goal of [131] is to understand the influence of secondary market in a supply chain

and inventory management. Another main difference is that, in [131], the demand in the secondary period remains *stochastic*, and thus the second period can still be modeled as a newsvendor model - this to a certain degree helped simplify the analysis. To understand how secondary market mitigates uncertainty, we assume that players' demands will be realized in the secondary market and it cannot be modeled as a newsvendor problem anymore. Furthermore, [131] only discussed symmetric equilibrium (even in a large market), which is hardly the case in my setting - if we assume that players' demands have been realized in the secondary market, then symmetric equilibrium does not exist almost surely in my case (for large market). We have obtained and explained different results with my model: for instance, [131] shows that the price in the secondary market is strictly lower than the original purchase price - but this is not the case with my model. We'd like to emphasize that both [131] and my work studied the influence of secondary market instead of the mechanism/market design question, which is different from many of the secondary markets work mentioned earlier.

Part I

Toward Efficiency: Security Game

CHAPTER III

Multi-Defender Security Games

In this chapter, I will analyze the equilibrium in general multi-defender security games. I will characterize both Nash equilibrium and approximate equilibrium in the games, and analyze the inefficiency introduced by defenders' selfish behavior by the *price of anarchy* analysis.

I consider a problem with multiple defenders protecting a collection of homogeneous targets. Each defender chooses a probability distribution over protection levels for all targets in their charge. A single attacker then best responds to the defenders' action by attacking the target with the lowest probability to be protected, breaking ties uniformly at random.

My analysis is focused on three models of such multi-defender games, with defenders acting non-cooperatively in all of these. I show that a Nash equilibrium among defenders in this two-stage game model need not always exist, even when the defenders utilize randomized strategies (i.e., probability distributions over target protection levels); this is distinct from a model in which the attacker moves simultaneously with the defenders, where a mixed strategy equilibrium is guaranteed to exist. When an equilibrium does exist, I show that the defenders protect all of their targets with probability 1 in all three models, whereas the socially optimal protection levels are generally significantly lower. When no equilibrium exists, I characterize the best approximate Nash equilibrium (that is, one in which defenders have the least gain from deviation), showing that over-investment is substantial in this case as well. Our *price of anarchy (PoA)* analysis, which relies on the unique equilibrium when it exists, and the approximate equilibrium otherwise, demonstrates a surprising finding: whereas PoA is unbounded in the simpler models, increasing linearly with the number of defenders, the more general model shows this to be an atypical special case

achieved when several parameters are exactly zero. More generally, PoA tends to a constant as the number of defenders increases.

This work has been published in [22], [78].

III.1 Problem Setting

In the mult-idefender security game model, there are a collection of defenders $N = \{1, 2, 3, \ldots, n\}$, and a single attacker. A collection of targets T will be protected by these defenders. Each defender i is in charge of a set of targets T_i, such that $T_i \subseteq T$. I assume $T_i \cap T_{i'} = \emptyset$ when $i \neq i'$, and $\cup_{i \in N} T_i = T$.

Strategies Suppose that each defender i can choose from a finite set $O = \{o_1, o_2, \ldots, o_{|O|}\}$ of security configurations for each target $t \in T_i$. A *pure strategy of defender i* is $\mathbf{o}_i =< o_{i,t_{i1}}, o_{i,t_{i2}}, \ldots, o_{i,t_{ik}} \ldots, o_{i,t_{i|T_i|}} >$, in which t_{ik} is the k_{th} target of defender i, and $o_{i,j}$ (here $j = t_{i1}, t_{i2}, \ldots$, etc.) means defender i's security configuration on target j s.t. $j \in T_i$. I assume the attacker is resource constrained and can only attack one target in the game. That is, a pure strategy of the attacker is j, s.t. $j \in T$.

A *Mixed Strategy of a defender i* is a matrix

$$\mathbf{q}_i = \begin{pmatrix} q_{i,t_{i1}}^{o_1} & q_{i,t_{i2}}^{o_1} & \cdots & q_{i,t_{i|T_i|}}^{o_1} \\ q_{i,t_{i1}}^{o_2} & q_{i,t_{i2}}^{o_2} & \cdots & q_{i,t_{i|T_i|}}^{o_2} \\ \vdots & \vdots & \ddots & \vdots \\ q_{i,t_{i1}}^{o_{|O|}} & q_{i,t_{i2}}^{o_{|O|}} & \cdots & q_{i,t_{i|T_i|}}^{o_{|O|}} \end{pmatrix}$$

In which, $q_{i,j}^{o}$ (here $o = o_1, o_2, \ldots, o_{|O|}$ and $j = t_{i1}, t_{i2}, \ldots, t_{i|T_i|}$) is the probability that the defender i chooses o at target j, and $\sum_{o \in O} q_{i,j}^{o} = 1$.

In the model, I assume a single strategic attacker that observes the defenders' coverage probabilities and chooses a target that maximizes its utility. A mixed strategy of the attacker can be denoted by $\mathbf{p} =< p_{t_1}, p_{t_2}, \ldots, p_{t_k} \ldots, p_{t_{|T|}} >$, in which, t_k is the k_{th} target in target set T, and p_j (here $j = t_1, t_2, \ldots, t_{|T|}$) is the probability of attacking target $j \in T$.

Let $\mathbf{q} =< \mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_n >$ denote the strategy profile of the defenders, and $(\mathbf{q}, \mathbf{p})$ denote the strategy profile of the defenders and the attacker.

Payoffs A configuration $o \in O$ for target $j \in T_i$ incurs a cost c_j^o to the defender i. If the attacker attacks a target $j \in T$ while configuration o is in place, the expected value to a defender i is denoted by $U_{i,j}^o$, while the attacker's value is V_j^o. We assume in this model that each player's utility depends only on the target attacked and its security configuration [132, 133].

Solution Concepts Traditionally, in single defender Stackelberg security games, the solution concept used is Strong Stackelberg Equilibrium (SSE). A SSE is characterized by an assumption that the attacker breaks ties in defender's favor. However the notion of "breaking ties in defender's favor" is no longer well defined when there are multiple defenders, as we must specify *which* defender will receive the favor. In the thesis, I adopt a natural tie-breaking rule in which the attacker chooses a target uniformly at random from the set of all best responses. We call the corresponding solution concept (which is a refinement of the subgame perfect equilibrium of our game) the *Average-case Stackelberg Equilibrium (ASE)*.

Definition III.1.1. *(Average-case Stackelberg Equilibrium) A strategy profile* $(\mathbf{q}, \mathbf{p})$ *is ASE if each defender's strategy is a best response, taking other defenders' strategies as given and assuming that the attacker will always play a best-response strategy, breaking ties uniformly at random if there are multiple best-response strategies.*

As I demonstrate below, ASE is not guaranteed to exist, in which case I focus on ε-ASE (a refinement of ε-equilibrium), in which no defender gains more than ε by deviating; in particular, I will consider ε-ASE with the smallest attainable ε.

To measure how the efficiency of the game degrades due to selfish behavior of the defenders, I consider *Utilitarian Social Welfare* and (ε)-*Price of Anarchy* in the thesis. *Utilitarian Social Welfare* is the sum of all defenders' payoffs. For the smallest attainable

ε, I define ε-Price of Anarchy (ε-PoA) as follows:

$$\varepsilon\text{-}PoA = \frac{SW_O}{\varepsilon\text{-}SW_E}$$

where SW_O is the optimal (utilitarian) social welfare that can be obtained (i.e., if there was a single defender), and ε-SW_E is the worst-case (utilitarian) social welfare in ε-ASE. An underlying assumption of this definition is that the value of SW_O and ε-SW_E are both positive. If they are both negative, then ε-PoA will be the reciprocal of the above equation. Note that the ordinary *Price of Anarchy* is a special case of ε-Price of Anarchy with $\varepsilon = 0$.

III.2 Equilibrium Analysis

In this chapter, I consider scenarios in which the values of the targets are *independent* and *homogeneous* among the defenders. Our equilibrium and Price of Anarchy analysis will show that a Nash equilibrium among defenders in the Stackelberg game model (equivalently, ASE) [1] need not always exist, even when the defenders utilize randomized strategies (i.e., probability distributions over target protection levels). For cases when there is no Nash equilibrium, I make use of approximate Nash (ASE) equilibrium and the associated (ε)-Price of Anarchy.

III.2.1 Equilibrium Analysis on a Baseline Model

I start with a model which most reflects the related literature: in particular, this model involves n defenders and a single attacker, with each defender engaged in protecting a single target. Each target has the same value to the defender $v > 0$. Suppose that the defender has two discrete choices: to protect the target, or not. In addition, the defender can randomly choose among these; our focus is on these *coverage probabilities* (i.e., the probability of protecting, or covering, the target), which we denote by s_i for a given defender i. The attacker is strategic and could observe the defenders' strategies to choose a target so as to

[1] If we treat attacker as an externality, we could see an ASE as a Nash equilibrium among defenders. For ease of exposition, we will also use "Nash equilibrium among defenders" to denote ASE in the chapter.

maximize the damage. We assume that the attacker is indifferent among the targets, and attacks the target with the lowest coverage probability, breaking ties uniformly at random. In a given scenario, for all defenders, the attacker's strategy is a vector of probabilities $P =< p_1, p_2, ..., p_n >$, where p_i is the probability of attacking target i, with $\sum_{i=1}^{n} p_i = 1$.

We assume that if the attacker chooses to attack a target corresponding to defender i and defender i chooses to protect the target, then the utility of the defender i is 0, and if the attacker attacks the target but it is not protected, then the utility of the defender is $-v$. If a defender chooses to cover a target, it will incur a cost $c > 0$. Additionally, we assume that the defender gets a utility of zero whenever another defender's target is attacked. We can thus define the expected utility of a defender i as

$$u_i = p_i u_i^a + (1 - p_i) u_i^u,$$

where u_i^a is the utility of i if it is attacked, and u_i^u is the utility of i if it is not attacked. By the assumptions above,

$$u_i^a = -(1 - s_i)v - s_i c = -v + s_i(v - c)$$

$$u_i^u = -s_i c.$$

Our first result presents necessary and sufficient conditions for the existence of a Nash equilibrium in the baseline model, and characterizes it when it does exist.

Theorem III.2.1. *In the* Baseline model*, Nash equilibrium exists* if and only if $v \geq c$. *In this equilibrium, all targets are protected with probability 1.*

Proof. Firstly, we claim that Nash equilibrium among defenders can appear *only if* all targets have the same coverage probability s to be protected. Otherwise, some defender j who has the possibility of 0 to be attacked has the incentive to decrease her s_j. To find the Nash equilibria, we need only consider strategy profiles in which all targets have the same

coverage probabilities to be protected.

When all defenders have the same possibility s to cover their targets. For each defender, her expected utility is

$$u = \frac{(v - cn)s - v}{n}$$

If $s < 1$, some defender i could slightly increase s to $s + \delta$ (δ is a very small positive real number) to make sure herself not be attacked and get utility $u' = -(s + \delta)c$,

$$u' - u = \frac{v(1 - s) - nc\delta}{n}$$

As δ can be very small, $u' - u > 0$ when $s < 1$. We could know that the defender has incentive to improve s when $s < 1$. So the Nash equilibrium can appear *only if* $s_i = 1$ for all defender i.

When all defenders have the same possibility $s = 1$ to cover their targets. For each defender, her expected utility is

$$u = -c$$

If a defender i decreases her coverage probability to $s' < 1$, then her target will have the probability of 1 to be attacked, and she gets expected utility $u' = -v + s'(v - c)$,

$$u' - u = (v - c)(s' - 1)$$

If $v \geq c$, then $u' - u \leq 0$, all defenders do not have the incentive to deviate, so it is a Nash equilibrium. If $v < c$, then $u' - u > 0$, the defender has the incentive to deviate, so it is not a Nash equilibrium. To sum up, Nash equilibrium exists *if and only if* $v \geq c$, in which all defenders have the same probability 1 to protect their targets. □

Thus, if a Nash equilibrium does exist, it is unique, with all defenders always protecting their target. But what if the equilibrium does not exist? Next, we characterize the (unique) ε-equilibrium with the minimal ε that arises in such a case. We will use this approximate

equilibrium strategy profile as a *prediction* of the defenders' strategies.

Theorem III.2.2. *In the* Baseline model*, if $v < c$, the optimal ε-equilibrium is for all defenders to cover their target with probability $\frac{v}{c}$. The corresponding ε is $\frac{v(c-v)}{cn}$.*

Proof. We firstly consider strategy profiles in which all targets have the same possibility s to be protected. Then for each defender, her expected utility is

$$u = \frac{(v-cn)s - v}{n}$$

Assume $0 \leq s < 1$. If some defender i slightly increase s to $s+\delta_1$, then she could get utility $u' = -(s+\delta_1)c$,

$$u' - u = \frac{v(1-s) - nc\delta_1}{n} < \frac{v(1-s)}{n}$$

Assume $0 < s \leq 1$. If some defender i slightly decreases s to $s-\delta_2$, then she could get the utility $u'' = -v + (s-\delta_2)(v-c)$

$$u'' - u = \frac{v(1-s)(1-n) + \delta_2 n(c-v)}{n}$$

As $\delta_2 \leq s$, we could get

$$u'' - u \leq \frac{v(1-s)(1-n) + sn(c-v)}{n} = \frac{v(1-s)}{n} + (sc - v)$$

Let $d_1 = \frac{v(1-s)}{n}$, $d_2 = \frac{v(1-s)}{n} + (sc-v)$. For $s = 0$, a defender could deviate to get an increased value which is less than $\frac{v}{n}$, so it is $\frac{v}{n}$-equilibrium. For $s = 1$, a defender could deviate to get an increased value which is less or equal to $(c-v)$, then it is $(c-v)$-equilibrium.

When $0 < s \leq \frac{v}{c}$ and $d_2 \leq d_1$, it is d_1-equilibrium. When $\frac{v}{c} < s < 1$ and $d_2 > d_1$, it is d_2-equilibrium.

To sum up, for ε-equilibrium,

$$\varepsilon = \begin{cases} \frac{v(1-s)}{n}, & \text{if } 0 \leq s \leq \frac{v}{c}; \\ \frac{v(1-s)}{n} + (sc - v), & \text{if } \frac{v}{c} < s \leq 1. \end{cases}$$

When $s = \frac{v}{c}$, we could get the minimal $\varepsilon = \frac{v(c-v)}{cn}$. And it is the only $\frac{v(c-v)}{cn}$-equilibrium in strategy profiles of all defenders having the same coverage probabilities.

We claim that the $\frac{v(c-v)}{cn}$-equilibrium *could only* exist in a profile of all defenders having the same coverage probability s. Otherwise, assume defenders have different probabilities to cover their targets, then there are α defenders ($1 \leq \alpha < n$) who have the same minimal probability s' to protect their targets. The expected utility for each defender among these α defenders is:

$$u_e = \frac{(v - c\alpha)s' - v}{\alpha}$$

When $\frac{v}{c} < s' \leq 1$, some defender i among these α defenders could decrease her probability to 0 to get value $u_1 = -v$,

$$u_1 - u_e = \frac{v(1-s')}{\alpha} + (s'c - v) > \frac{v(1-s')}{n} + (s'c - v)$$

When $0 \leq s' \leq \frac{v}{c}$, some defender i among these α defenders could slightly increase her probability to $s' + \delta_3$ to get the utility $u_2 = -(s' + \delta_3)c$

$$u_2 - u_e = \frac{v(1-s') - \alpha c\delta_3}{\alpha} > \frac{v(1-s')}{n}$$

The above inequation holds because δ_3 can be very small. Then we could know that it cannot be a $\frac{v(c-v)}{cn}$-equilibrium.

So we could know that it is the only $\frac{v(c-v)}{cn}$-equilibrium when all defenders have the equal probability $\frac{v}{c}$ to cover their targets. And it is the optimal approximate equilibrium.

□

Armed with a complete characterization of predictions of strategic behavior among the defenders, we can now consider how this behavior related to socially optimal protection decisions. Since the solutions are unique, there is no distinction between the notions of *price of anarchy* and *price of stability*; we term the ratio of socially optimal welfare to welfare in equilibrium as the price of anarchy for convenience.

First, we characterize the socially optimal outcome.

Theorem III.2.3. *In the* Baseline model*, the optimal social welfare SW_O is*

$$SW_O = \begin{cases} -cn, & \text{if } v \geq cn; \\ -v, & \text{if } v < cn. \end{cases}$$

Proof. We firstly claim that we could get optimal social welfare *only if* all defenders have the same probability s to protect their targets. Otherwise, their coverage probabilities are different, and some defender j has the probability of 0 to be attacked. Then we could decrease s_j to get better social welfare. Therefore we just need to look for the identical coverage probability s which makes the optimal social welfare. The function of social welfare over s is as follows:

$$SW(s) = -v + s(v-c) + (n-1)(-sc) = (v-cn)s - v$$

Then we could get the optimal social welfare as the theorem shown. □

From this result, it is already clear that defenders systematically over-invest in security, except when values of the targets are quite high. This stems from the fact that the attacker creates a *negative externality* of protection: if a defender protects his target with higher probability than others, the attacker will have an incentive to attack another defender. In such a case, we can expect a "dynamic" adjustment process with defenders increasing their security investment well beyond what is socially optimal. To see just how much the

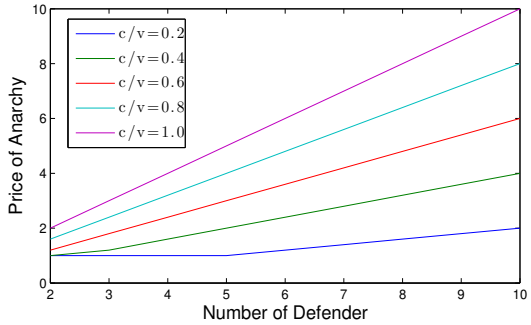

Figure III.1: Price of Anarchy when $v \geq c$

defenders lose in the process, we now characterize the price of anarchy of our game.

If $v \geq c$, it is one and only one Nash equilibrium when all defenders have the coverage probability 1 for their targets. And the corresponding social welfare is

$$SW_E = -cn$$

Because it is the only Nash equilibrium, we could get the *Price of Anarchy* as follows:

$$PoA = \begin{cases} 1, & \text{if } v \geq cn; \\ \frac{nc}{v}, & \text{if } c < v < cn. \end{cases}$$

Figure III.1 shows the relationship among Price of Anarchy, the number of defenders, and the ratio of cost c and value v. From the figure we could find that when the number of defenders and the ratio of c and v are small enough (e.g. $n \leq 5$ and $\frac{c}{v} = 0.2$), the price of anarchy is close to 1. Otherwise, the price of anarchy is unbounded, growing linearly with n.

If $v < c$, there is no Nash equilibrium. However, we could get the optimal ε-equilibrium when all defenders have the same coverage probability $\frac{v}{c}$ for their targets. The corresponding Social Welfare is

$$SW_E = (v - cn)\frac{v}{c} - v$$

Similarly, we could get the $\frac{v(c-v)}{cn}$-Price of Anarchy as follows,

$$\frac{v(c-v)}{cn}\text{-}PoA = \frac{cn+c-v}{c},$$

which is, again, linear in n.

III.2.2 Equilibrium Analysis of the General Model

In this model, we assume all targets in T are homogeneous, and each target has the same value to the defender. In the game model, each defender protects k targets, i.e. $|T_1| = |T_2| = \ldots = |T_n| = k$. The security configuration space is $O = \{0,1\}$, i.e., the defender's decision is binary. For example, 1 can correspond to the decision to protect an asset, while configuration 0 would leave the asset unprotected. The *pure strategy* of defender i is $\mathbf{o}_i =< o_{i,t_{i1}}, o_{i,t_{i2}}, \ldots, o_{i,t_{ik}} >$, in which $o_{i,j}$ (here $j = t_{i1}, t_{i2}, \ldots$) is a binary value. The mixed strategy of a defender i is $\mathbf{q}_i =< q_{i,t_{i1}}, q_{i,t_{i2}}, \ldots q_{i,t_{ik}} >$, in which $q_{i,j}$ is the probability of protecting target j for defender i (coverage probability). The cost to defend each target is denoted by c.

If the attacker chooses to attack a target controlled by defender i and the defender chooses to protect the target, we define the value of the target to defender i to be U^c, and if the attacker attacks the target but it is not protected, then the value of the target to the defender is U^u. It is reasonable to assume that $U^c \geq U^u$. If the target of defender i is not attacked, the value of the target for defender i is $\Omega \geq U^c$. In this setting, we assume that the attacker aims to maximize expected damage to the defender, so that the attacker's utility is $-U^u$, $-U^c$, and $-\Omega$ for the three outcomes above, respectively. Since these values are uniform across targets, equivalently the attacker attacks a target with lowest coverage probability (breaking ties uniformly at random).

Our first result presents necessary and sufficient conditions for the existence of a Nash equilibrium among defenders (ASE) in the setting, and characterizes it when it does exist.

Theorem III.2.4. *In the* Independent Multidefender setting*, Nash equilibrium among defenders (ASE) exists* if and only if $U^c - U^u \geq kc - \frac{(n-1)(\Omega - U^c)}{n}$. *In this equilibrium all targets are protected with probability 1.*

Proof. We firstly claim that Nash equilibrium can appear *only if* coverage probabilities of all of targets t_{ij} are identical. Otherwise, there will be a target t_{ik} which has the probability 0 of being attacked, and the defender i has an incentive to decrease q_{ik}. To determine a Nash equilibrium, we therefore need only consider scenarios in which all targets have the same coverage probability.

When all targets have the same coverage probability q to be protected, the utility of each defender is

$$u = \frac{(U^c - U^u - nkc)q + U^u + (nk-1)\Omega}{n}.$$

If $q < 1$, then some defender i could increase q to $q + \delta$ for all of her targets to ensure none of them are attacked, and obtain utility of $u' = k\Omega - k(q+\delta)c$, so that

$$u' - u = \frac{(U^c - U^u)(1-q) + (\Omega - U^c) - nkc\delta}{n}.$$

As $U^c \geq U^u$, $\Omega \geq U^c$, and δ can be arbitrarily small, $u' - u > 0$ when $q < 1$, which means that this cannot be a Nash equilibrium. Thus, the only possible equilibrium can be $q_{ij} = 1$ for all targets t_{ij}.

When all targets have the same coverage probability $q = 1$, each defender's utility is

$$u = \frac{U^c - nkc + (nk-1)\Omega}{n}.$$

We claim that if a defender i has an incentive to deviate, it is optimal for this defender to use the same coverage probability for all her targets. Otherwise, for some target t_{ik} which has probability 0 of being attacked, she could decrease q'_{ik} to obtain higher utility. If probabilities of targets protected by defender i are all q' $(0 \leq q' < 1)$, then her expected

utility is $u' = (U^c - U^u - c)q' + U^u + (k-1)(\Omega - q'c)$, and

$$u' - u = (U^c - U^u - kc)(q' - 1) + \frac{(n-1)(U^c - \Omega)}{n}.$$

We therefore have two cases:

1) If $U^c - U^u \geq kc$, then $u' - u \leq 0$, and $q = 1$ for all targets is a Nash equilibrium.

2) If $U^c - U^u < kc$, the maximal value of $u' - u$ corresponds to $q' = 0$:

$$\max_{0 \leq q' < 1} u' - u = -(U^c - U^u - kc) - \frac{(n-1)(\Omega - U^c)}{n}.$$

If $kc - \frac{(n-1)(\Omega - U^c)}{n} \leq U^c - U^u < kc$, $u' - u \leq 0$, it is a Nash equilibrium; otherwise, it is not.

To sum up, a Nash equilibrium exists *if and only if* $U^c - U^u \geq kc - \frac{(n-1)(\Omega - U^c)}{n}$, and the equilibrium corresponds to all targets having probability 1 of being protected. □

Thus, if a Nash equilibrium does exist, it is unique, with all defenders always protecting their targets. But what if the equilibrium does not exist? Next, we characterize the (unique) ε-equilibrium (ε-ASE) with the minimal ε that arises in such a case. We will use this approximate equilibrium strategy profile as a *prediction* of the defenders' strategies.

Theorem III.2.5. *In* Independent Multidefender setting*, in the optimal ε-equilibrium (ε-ASE) all targets are protected with probability $\frac{\Omega - U^u}{kc}$. The corresponding ε is $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$.*

Proof. When all targets have the same coverage probability q, the expected utility of each defender is

$$u = \frac{(U^c - U^u - nkc)q + U^u + (nk-1)\Omega}{n}.$$

Suppose $0 \leq q < 1$. If some defender i increases q to $q + \delta_{ij}$ for each of her target t_{ij}, then

she would obtain utility $u' = \sum_{j=1}^{k} \Omega - (q + \delta_{ij})c$, and

$$\begin{aligned} u' - u &= \frac{\Omega - (U^c - U^u)q - U^u}{n} - \sum_{j=1}^{k} \delta_{ij} c \\ &\leq \frac{\Omega - (U^c - U^u)q - U^u}{n}. \end{aligned} \tag{III.1}$$

Now we consider scenarios in which a defender i could obtain higher utility by decreasing protection probability. We claim that if a defender i has an incentive to deviate, it is optimal for this defender to use the same coverage probability for all her targets. Otherwise, for some target t_{ik} which has probability 0 of being attacked, she could decrease q'_{ik} to obtain higher utility. Thus, we need only consider cases in which a defender deviates by decreasing coverage probabilities for all her targets to $q - \delta$. Her utility will become $u'' = (U^c - U^u - kc)(q - \delta) + U^u + (k-1)\Omega$. Since $U^c - U^u < kc$, $\delta = q$ (the maximal value of δ) maximizes $u'' - u$:

$$\max_{0<\delta\leq q} u'' - u = \frac{\Omega - (U^c - U^u)q - U^u}{nk} + kcq + U^u - \Omega. \tag{III.2}$$

By comparing the value of equation (III.1) and equation (III.2), we get different values of ε for ε-equilibrium:

$$\varepsilon = \begin{cases} \frac{\Omega - (U^c - U^u)q - U^u}{n}, & \text{if } 0 \leq q \leq \frac{\Omega - U^u}{kc}; \\ \frac{\Omega - (U^c - U^u)q - U^u}{n} + kcq + U^u - \Omega, & \text{if } \frac{\Omega - U^u}{kc} < q \leq 1. \end{cases}$$

When $q = \frac{\Omega - U^u}{kc}$, we get the minimal $\varepsilon = \frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$.

We claim that the $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-equilibrium can appear *only if* all targets have the same coverage probability q. We prove this by contradiction. Suppose that targets have different coverage probabilities. This gives rise to two cases: 1) Each defender uses an identical coverage probability for each target she owns (these may differ between defenders); and 2) Some defender has different coverage probabilities for her targets. In case 1),

there exist β defenders ($1 \leq \beta < n$) who have the same minimal coverage probability q'. The expected utility for each defender among these β is

$$u = \frac{(U^c - U^u - k\beta c)q' + U^u + (k\beta - 1)\Omega}{\beta}.$$

When $\frac{\Omega - U^u}{kc} < q' \leq 1$, some defender i among these β could decrease the coverage probability of all her targets to 0 and obtain the utility of $u_1 = U^u + (k-1)\Omega$, so that

$$\begin{aligned} u_1 - u &= \frac{\Omega - (U^c - U^u)q' - U^u}{\beta} + kcq' + U^u - \Omega \\ &> \frac{\Omega - (U^c - U^u)q' - U^u}{n} + kcq' + U^u - \Omega. \end{aligned}$$

When $0 \leq q' \leq \frac{\Omega - U^u}{kc}$, some defender i among these β can increase coverage probabilities of all her targets to $q' + \delta_3$ to obtain utility of $u_2 = k\Omega - k(q' + \delta_3)c$, with

$$\begin{aligned} u_2 - u &= \frac{\Omega - (U^c - U^u)q' - U^u - k\beta c\delta_3}{\beta} \\ &> \frac{\Omega - (U^c - U^u)q' - U^u}{n}, \end{aligned}$$

where the inequality holds because δ_3 can be arbitrarily small. Thus, no profile in case 1) can be a $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-equilibrium. In case 2), any defender who has different coverage probabilities among her targets can always increase her payoff by decreasing the coverage probabilities of the targets with higher coverage to yield identical coverage for all targets. Consequently, no profile in case 2) can be a $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-equilibrium. □

Armed with a complete characterization of predictions of strategic behavior among the defenders, we can now consider how this behavior is related to socially optimal protection decisions. Since the solutions are unique, there is no distinction between the notions of *price of anarchy* and *price of stability*; we term the ratio of socially optimal welfare to welfare in equilibrium as the price of anarchy for convenience.

Theorem III.2.6. *In the* Independent Multidefender setting*, the optimal social welfare*

SW_O is

$$SW_O = \begin{cases} U^c - nkc + (nk-1)\Omega, & \text{if } U^c - U^u \geq nkc; \\ U^u + (n-1)\Omega, & \text{if } U^c - U^u < nkc. \end{cases}$$

Proof sketch. First, we claim that we could get optimal social welfare *only if* all targets have the same coverage probability q. Otherwise, some target j, which is influenced by defender i has probability 0 of being attacked, and we can decrease $q_{i,j}$ to improve social welfare. Consequently, we need only to consider an optimal symmetric coverage probability q to maximize social welfare, which can be done in a relatively straightforward way. □

If $U^c - U^u \geq kc - \frac{(n-1)(\Omega - U^c)}{n}$, the Nash equilibrium is unique, with all targets protected with probability 1. The corresponding social welfare is

$$SW_E = U^c - nkc + (nk-1)\Omega.$$

So far we have not yet added any constrains to value of Ω, U^c, and U^u (except that $\Omega \geq U^c \geq U^u$). In order to make *Price of Anarchy* well-defined, we need to add constraints that values of Ω, U^c, and U^u are all non-positive or all non-negative. We add constraints that U^c, U^u and Ω are all non-positive (little changes if all are non-negative).

In the case of a unique Nash equilibrium, the price of anarchy is

$$PoA = \begin{cases} 1, & \text{if } U^c - U^u \geq nkc; \\ \frac{U^c - U^u - nkc}{U^u + (nk-1)\Omega} + 1, & \text{if } kc - \frac{(n-1)(\Omega - U^c)}{n} \leq \\ & U^c - U^u < nkc. \end{cases}$$

If $U^c - U^u < kc - \frac{(n-1)(\Omega - U^c)}{n}$, there is no Nash equilibrium. The Social Welfare in the

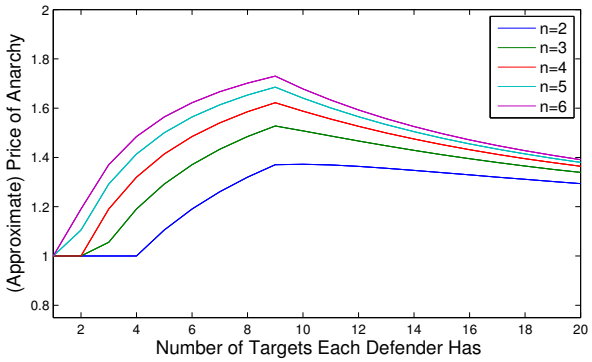

Figure III.2: (Approximate) Price of Anarchy when $c = 1, \Omega = -1, U^c = -2$ and $U^u = -10$

optimal approximate equilibrium is

$$\varepsilon\text{-}SW_E = (U^c - U^u - nkc)\frac{\Omega - U^u}{kc} + U^u + (nk-1)\Omega,$$

and the $\frac{(\Omega - U^u)(kc - U^c + U^u)}{cnk}$-Price of Anarchy is $\frac{(U^c - U^u - nkc)(\Omega - U^u)}{kcU^u + (nk-1)kc\Omega} + 1$.

From this result, it is already clear that defenders systematically over-invest in security. This stems from the fact that the attacker creates a *negative externality* of protection: if a defender protects his target with higher probability than others, the attacker will have an incentive to attack another defender. In such a case, we can expect a "dynamic" adjustment process with defenders increasing their security investment well beyond what is socially optimal.

We now analyze the relationship between (ε-)PoA and the values of n and k. First, we consider (ε-)PoA as the function of n. If $\Omega = 0$, (ε-)PoA linearly increases in n, and is therefore unbounded. However, if $\Omega \neq 0$, while PoA and ε-PoA are increasing in n, as $n \to \infty$, they approach $1 - \frac{c}{\Omega}$ and $1 + \frac{U^u - \Omega}{k\Omega}$, respectively. In other words, PoA (exact and approximate) is bounded by a constant, for a constant k.

III.3 Conclusion

I examined a non-cooperative multi-defender security game in which defenders may protect multiple targets, offering complete characterization Average-case Stackelberg Equilibrium (or equivalently, Nash equilibrium among defenders) and approximate equilibria, socially optimal solutions, and price of anarchy. The results show that defenders generally over-protect the targets in this model, but different modeling assumptions give rise to qualitatively different outcomes: a simpler model gives rise to an unbounded price of anarchy, whereas a more general model sees PoA converge to a constant when the number of defenders increases.

CHAPTER IV

Multi-Defender Strategic Filtering Against Spear-Phishing Attacks

In this chapter, I address the problem of strategic e-mail threshold selection by a collection of independent users, faced with a threat of both spear-phishing and non-targeted (e.g., spam) malicious e-mail campaigns. I consider strategic dynamics by appealing to a Stackelberg multi-defender equilibrium concept. I offer a characterization of the equilibria, and present a polynomial-time algorithm for computing the Stackelberg multi-defender equilibrium. Remarkably, I demonstrate that Stackelberg multi-defender equilibria need not exist, and it is socially optimal if it exists, which is very different from the outcome I get from general multi-defender security games.

This work is published in [134].

IV.1 Problem Settings

The model is based on the model introduced in [1], which we now extend for independent and self-interested defenders. I model the strategic interactions of spear-phishing as a game between multiple *users* and a targeting *attacker*. Note that I refer to the defending players as users; however, these players can naturally model groups of users having the same e-mail filtering policy, or even entire organizations.

Users may receive three types of e-mails: *non-malicious*, *malicious non-targeted*, and *malicious targeted*. If a non-malicious e-mail is filtered out, which we call a *false positive* (FP), then the user suffers usability loss. If a malicious e-mail is not filtered out, which we call a *false negative* (FN), then the user might open that e-mail and suffer loss from the attack. We assume that the attainable false-positive and false-negative probability pairs are given by a function $FP : [0,1] \mapsto [0,1]$, where $FP(f)$ is the probability of false positives when the probability of false negatives is f. In any practical e-mail classifier, $FP(f)$ is a non-increasing function of f (see Figure IV.1 for an illustration). For analytical tractability, we further assume that $FP(f)$ is a continuous, strictly decreasing, and strictly convex

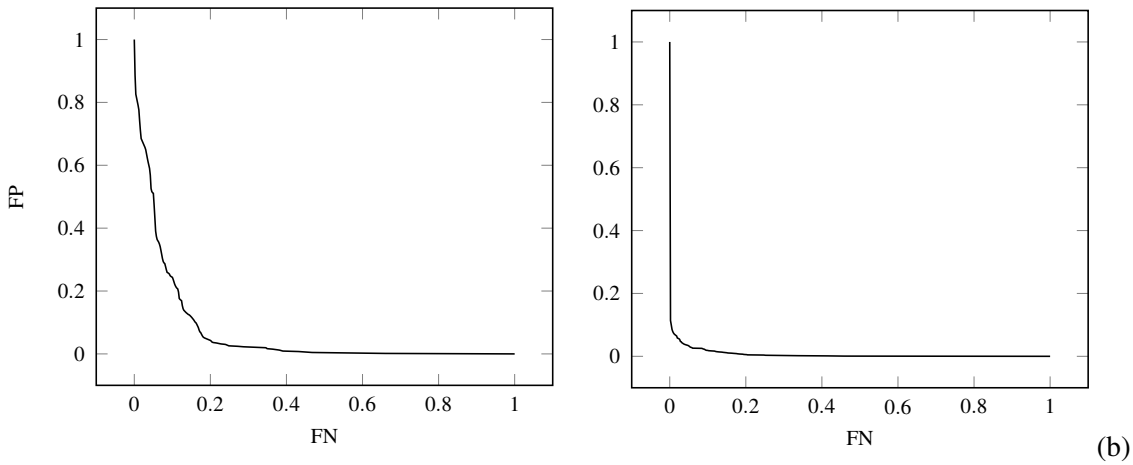

(a) UCI dataset (b) Enron dataset

Figure IV.1: False-negative to false-positive tradeoff curves for the two datasets used in [1] and [2].

function of f. Note that these assumptions hold approximately in practice.

Malicious e-mails are divided into two categories: targeted and non-targeted. The former includes spear-phishing and whaling e-mails sent by the targeting attacker, while the latter includes spam and non-targeted phishing e-mails. Since the senders of non-targeted e-mails do not choose their targets in a strategic way in practice, we model them as *non-strategic* actors instead of game-theoretic players (see constant N_u below).

Strategies A *pure strategy of user u* is a false-negative probability f_u, and we let $\boldsymbol{f}$ denote the strategy profile of the users. Note that we do not have to consider thresholds explicitly in our model, since there is a bijection between false-negative probabilities and thresholds values.

A *pure strategy of the attacker* is a set of users $\mathscr{A}$, who will be attacked. Since targeted e-mails have to be customized, which requires spending a considerable amount of effort on each target, the number of users that can be targeted is limited. Formally, the attacker's strategy is subject to $|\mathscr{A}| \leq A$. For the same reason, we also assume that the attacker is *lazy* in the sense that she does not target a user when she would receive zero payoff for targeting the user.

We will also consider mixed strategies, which are defined naturally: a *mixed strategy of the attacker* is a distribution over subsets of users, while a *mixed strategy of user u* is a distribution over false-negative values from $[0,1]$.

Payoffs For a given pure-strategy profile $(\boldsymbol{f}, \mathscr{A})$, the attacker's payoff is

$$\mathscr{U} = \sum_{u \in \mathscr{A}} f_u L_u, \tag{IV.1}$$

where $L_u > 0$ is the expected amount of damage when user u falls victim to a targeted attack.

If user u is targeted by the attacker, then her loss (i.e., inverse payoff) is

$$\mathscr{L}_u^1(f_u) = f_u(L_u + N_u) + FP(f_u)C_u, \tag{IV.2}$$

and if user u is not targeted, her loss is

$$\mathscr{L}_u^0(f_u) = f_u N_u + FP(f_u)C_u, \tag{IV.3}$$

where $N_u > 0$ is the loss of user u for delivering non-targeted malicious e-mails, and $C_u > 0$ is the loss for not delivering non-malicious e-mails. Payoffs for mixed-strategies are defined naturally as the expected payoff.

Solution Concepts In our analysis, we will study both short-term and long-term strategic dynamics of the e-mail filtering problem. As is typical in the literature, we study them using two different solution concepts, *Stackelberg multi-defender equilibrium* and *Nash equilibrium*.

In the short-term model, the game has two stages: in the first stage, the users make their strategic decision simultaneously; while in the second stage, the attacker makes its decision knowing which strategies the users have chosen. We solve this model using the concept of *Stackelberg multi-defender equilibrium* (SMDE), which is defined as follows.

Definition IV.1.1 (Stackelberg multi-defender equilibrium)**.** *A strategy profile is an SMDE if each user's strategy is a best response, taking the user's strategies as given and assuming that the attacker will always play a best-response strategy.*

IV.2 Equilibrium Analysis

First, in Section IV.2.1, we provide necessary conditions on the equilibria and introduce additional notation to facilitate our analysis. Then, we study and characterize the Stackelberg multi-defender of the game in Sections IV.2.2.

IV.2.1 Preliminaries

I begin the analysis by providing a necessary condition on the users' mixed-strategy best responses, which applies to SMDE.

Lemma IV.2.1. *The best-response strategy for a user is always a pure strategy.*

As a consequence, for the remainder of this chapter, I will consider only pure strategies (i.e., single false-negative values) for the users.

Proof. Suppose that we are given a mixed strategy that is not a pure strategy (i.e., its support consist of more than one false-negative value); then, we show that the expected false-negative value is a better strategy than the distribution. Firstly, it is easy to see that the other players' payoffs (and, hence, their best responses) remain the same if the user changes its strategy from a distribution to the expected value. Secondly, since the function FP is strictly convex, we have from Jensen's inequality that the user's loss is strictly less for the expected value than for the distribution. Therefore, for every mixed strategy that is not a pure strategy, there exists a strictly better pure strategy. □

Next, we introduce a simpler notation for the attacker's mixed strategies. Let a_u be the probability that user u is targeted by the attacker, that is, the probability that u is an element

of a subset chosen randomly according to the attacker's mixed strategy. Using this notation, we can express the attacker's expected payoff as

$$\mathscr{U} = \sum_u a_u f_u L_u \tag{IV.4}$$

and user u's expected loss as

$$\mathscr{L}_u^{a_u}(f_u) = f_u(a_u L_u + N_u) + FP(f_u)C_u. \tag{IV.5}$$

For every mixed strategy of the attacker, we can easily compute the corresponding vector of probabilities $\boldsymbol{a}$, which must satisfy $\sum_u a_u \leq A$. Furthermore, it is also easy to see that for every vector of probabilities $\boldsymbol{a}$ satisfying $\sum_u a_u \leq A$, there exists a mixed-strategy whose marginal is $\boldsymbol{a}$. For the remainder of this chapter, we will represent the attacker's mixed-strategies as vectors of probabilities.

Now, we introduce some additional notation to facilitate our analysis. Let $f_u^{a_u}$ denote user u's optimal false-negative probability given that the attacker targets it with probability a_u, that is, let $f_u^{a_u}$ be the f_u which minimizes $\mathscr{L}_u^{a_u}(f_u)$. It is easy to see that $f_u^{a_u}$ is well defined for any a_u, and it is a non-increasing and continuous function of a_u.

Finally, consider the value f_u^0, which is the optimal false-negative probability given that the attacker never targets user u (i.e., given $a_u = 0$). If $f_u^0 = 0$ for user u, then it is easy to see that the user will always play the strategy $f_u = 0$, regardless of the other players' strategies or the solution concept used. Furthermore, such users do not affect the other players' strategic choices either, since an attacker will never target user u if $f_u = 0$. Consequently, for the remainder of the chapter, we can disregard these users and assume that $f_u^0 > 0$ for every user u.

IV.2.2 Stackelberg Multi-Defender Equilibrium

In this subsection, I characterize the Stackelberg multi-defender equilibrium (SMDE) and design an algorithm to find it. First, I show that in an SMDE, the attacker plays a pure strategy and the users play f_u^1 or f_u^0. Then, I show that the SMDE is unique if it exists, and provide an efficient algorithm for computing it. However, I also find that the SMDE does not necessarily exist, but our algorithm can return "there is no SMDE" if it does not exist.

The following lemma shows that the attacker always plays a pure strategy in an SMDE.

Lemma IV.2.2. *A strategy profile is an SMDE* only if *for each user u, either* $a_u = 0$ *or* $a_u = 1$ *holds.*

Proof sketch. We prove the claim by contradiction. If $0 < a_u < 1$ for some user u, then her expected loss is

$$\mathscr{L}_u^{a_u}(f_u) = f_u(a_u L_u + N_u) + FP(f_u)C_u. \tag{IV.6}$$

Since the attacker's strategy is a best response, $0 < a_u < 1$ implies that there exists some user $v \neq u$ with $0 < a_v < 1$ and $f_u L_u = f_v L_v$. Hence, if user u changes her strategy to $f_u - \varepsilon$ (where ε is an arbitrarily small positive number), the attacker will target user v (or some other user) instead of user u. Then, the loss of user u will be

$$\mathscr{L}_u^0(f_u - \varepsilon) = (f_u - \varepsilon)N_u + FP(f_u - \varepsilon)C_u, \tag{IV.7}$$

since she is no longer targeted. We can compute the decrease in her loss due to deviating from her strategy as

$$\begin{aligned} &\mathscr{L}_u^{a_u}(f_u) - \mathscr{L}_u^0(f_u - \varepsilon) \\ =&a_u f_u L_u + \varepsilon N_u + [FP(f_u) - FP(f_u - \varepsilon)]\, C_u. \end{aligned} \tag{IV.8}$$

Clearly, $\mathscr{L}_u^{a_u}(f_u) - \mathscr{L}_u^0(f_u - \varepsilon)$ can be greater than 0, as ε can be arbitrarily small and $FP(f_u) - FP(f_u - \varepsilon)$ can be arbitrarily close to 0. Hence, user u can decrease her loss by

deviating from the strategy f_u, which leads to a contradiction with our initial assumption that the strategy profile is an SMDE. Therefore, the claim of the lemma has to hold. □

From Lemmas IV.2.1 and IV.2.2, we know that in an SMDE, both the users and the attacker play pure strategies. Now, we further constrain the users' equilibrium strategies by showing that user u plays either f_u^1 or f_u^0 in an SMDE.

Lemma IV.2.3. *A strategy profile is an SMDE* only if $f_u = f_u^1$ *for every user u who is targeted, and* $f_u = f_u^0$ *for every user u who is not targeted.*

Proof. We prove the claim by contradiction. Suppose that $\exists u$ such that $f_u \neq f_u^1$ and $f_u \neq f_u^0$. Based on Lemma IV.2.2, if the profile is an SMDE, then either $a_u = 0$ or $a_u = 1$, i.e., user u is targeted with probability 1 or 0.

1) First, assume that user u is targeted. Then, we show that strategy f_u^1 is better for user u than strategy f_u. First, if user u is still attacked after she switches to f_u^1, then we have by definition that f_u^1 is better since it minimizes $\mathscr{L}_u^1$. On the other hand, if the attacker no longer targets user u, then we have that the users' loss is even lower: $\mathscr{L}_u^0(f_u^1) \leq \mathscr{L}_u^1(f_u^1) < \mathscr{L}_u^1(f_u)$. Hence, f_u cannot be a best response.

2) If user u is not targeted, there are two cases: $f_u > f_u^0$ or $f_u < f_u^0$. If $f_u > f_u^0$, she can switch to f_u^0 to lower her loss without becoming a target of the attacker. If $f_u < f_u^0$, then we consider another user $v = \text{argmin}_{u \in \mathscr{A}} f_u L_u$, i.e., the user in the targeted set that makes attacker get lower payoff. Using an argument similar to the one used in the proof of Lemma IV.2.2, we can show that $f_u L_u < f_v L_v$; otherwise, user v could lower her loss by decreasing f_v with an arbitrarily small value. On one hand, when $f_u^0 L_u < f_v L_v$, user u can switch to f_u^0 to lower her loss without becoming a target of the attacker. On the other hand, when $f_u^0 L_u \geq f_v L_v$, user u can switch to some value f_u' such that $f_u L_u < f_u' L_u < f_v L_v$ and still not be targeted. Then, based on characteristics of $\mathscr{L}_u^0(f_u)$, we have that user u can lower her loss by switching to strategy f_u'.

Consequently, user u has incentives to deviate from her strategy in both cases, which implies that there is no SMDE in which $f_u \neq f_u^1$ and $f_u \neq f_u^0$ for some user u. □

Based on the above results, I first provide a necessary and sufficient condition for a strategy profile being an SMDE, and then present an algorithm to find an SMDE.

Theorem IV.2.1. *A strategy profile* $(\boldsymbol{f}, \mathscr{A})$ *is an SMDE* if and only if

1) $\forall u \in \mathscr{A}$*:* $f_u = f_u^1$,

2) $\forall u \notin \mathscr{A}$*:* $f_u = f_u^0$,

3) $F_1 > F_0$,

4) $\forall u \in \mathscr{A}$*:* $\mathscr{L}_u^1(f_u^1) \leq \mathscr{L}_u^0(\frac{F_0}{L_u})$,

where $F_1 = \min_{u \in \mathscr{A}} f_u^1 L_u$ *and* $F_0 = \max_{u \notin \mathscr{A}} f_u^0 L_u$.

Proof sketch. First, we prove that the conditions of the theorem are necessary. From Lemma IV.2.3, we readily have that $f_u = f_u^1$, $\forall u \in \mathscr{A}$, and $f_v = f_v^0$, $\forall v \notin \mathscr{A}$ hold in an SMDE. Next, since the attacker's best response must target the users with the highest $f_u L_u$ values, we also have that $\min_{u \in \mathscr{A}} f_u^1 L_u \geq \max_{u \notin \mathscr{A}} f_u^0 L_u$ has to hold in an SMDE. Furthermore, this inequality has to be strict, otherwise a user in $\mathscr{A}$ could decrease her loss by decreasing her strategy by an arbitrarily small amount. Finally, in an SMDE, users in $\mathscr{A}$ do not have the incentive to deviate from their strategy. If some user $u \in \mathscr{A}$ were to deviate, then she would pick a strategy that would divert attacks to another user, that is, she would consider a strategy $f_u \leq \frac{F_0}{L_u}$ (otherwise, following f_u^1 would obviously be better). Since $f_u^0 > f_u^1 > \frac{F_0}{L_u}$, her best choice would have to be $\frac{F_0}{L_u}$. Therefore, if user u has no incentive to deviate, then $\mathscr{L}_u^1(f_u^1) \leq \mathscr{L}_u^0(\frac{F_0}{L_u})$ has to hold.

Second, we prove that the conditions of the theorem are sufficient. For $\forall u \notin \mathscr{A}$, based on characteristics of the functions $\mathscr{L}_u^1(f_u)$ and $\mathscr{L}_u^0(f_u)$, we have that $\mathscr{L}_u^0(f_u^0)$ is the minimal loss user u can ever get, so she has no incentive to deviate. For $\forall u \in \mathscr{A}$, $\mathscr{L}_u^1(f_u^1)$ is

the minimal loss of user u given that she is targeted. Hence, the only way that she could decrease her loss is to avoid being targeted by the attacker. In order to avoid being targeted, she has to pick a strategy $f_u \leq F_0/L_u$. From the convexity of $\mathscr{L}_u^0$ and $f_u^0 > f_u^1 > F_0/L_u$, we have that $\mathscr{L}_u^0(f_u)$ is a decreasing function when $f_u < F_0/L_u$. Hence, her best strategy that avoids being targeted is F_0/L_u; however, it follows from $\forall u \in \mathscr{A}$: $\mathscr{L}_u^1(f_u^1) \leq \mathscr{L}_u^0(F_0/L_u)$ that this is inferior to f_u^1. Therefore, the users' strategies are best responses under the conditions of the theorem. Finally, it follows readily from $F_1 > F_0$ that the attacker's strategy is also the best response. □

In Theorem IV.2.1, we provided conditions for determining whether targeting a given set of users is an SMDE. In order to find an equilibrium, we could enumerate every subset $\mathscr{A}$ of users subject to $|\mathscr{A}| = A$, and check whether targeting $\mathscr{A}$ is an SMDE using Theorem IV.2.1. However, the running time of this approach grows exponentially as a function of A, and quickly becomes prohibitively large. We now provide a rather strong and surprising result which states that in an SMDE, the attacker will always target the set of A users with the highest value of $f_u^1 L_u$.

Lemma IV.2.4. *Let $\mathscr{A}$ be a subset of users such that $|\mathscr{A}| = A$ and $\min_{u\in\mathscr{A}} f_u^1 L_u > \max_{u\notin\mathscr{A}} f_u^1 L_u$. In an SMDE, all of the users in $\mathscr{A}$ will be targeted.*

Proof. We prove the claim by contradiction. Suppose that there is an SMDE such that some user $v \notin \mathscr{A}$ is targeted. Then, user v plays f_v^1, and there exists some user $w \in \mathscr{A}$ who is not targeted and plays f_w^0. Since the attacker's strategy is a best response, we have that $f_v^1 L_v \geq f_w^0 L_w$. From $\min_{u\in\mathscr{A}} f_u^1 L_u > \max_{u\notin\mathscr{A}} f_u^1 L_u$, we obtain $f_v^1 L_v < f_w^1 L_w$. However, since $\forall u : f_u^1 \leq f_u^0$, we also have $f_v^1 L_v < f_w^1 L_w \leq f_w^0 L_w$, which contradicts $f_v^1 L_v \geq f_w^0 L_w$. Hence, the original claim must hold. □

Then, based on Theorem IV.2.1 and Lemma IV.2.4, we propose Algorithm 4 for finding an SMDE. We also find that an SMDE may not necessarily exist, and we provide an example for this case below. Furthermore, we find that the SMDE is unique if it exists. To see

Algorithm 1 Find a Stackelberg Multi-Defender Equilibrium (SMDE)

input: a set of users $\mathbf{U}$, L_u, $\mathscr{L}_u^1(f_u)$ and $\mathscr{L}_u^0(f_u)$ for every user u, and A for attacker
return: a SMDE or "there is no SMDE"

for each user u **do**
    compute f_u^1 and f_u^0 based on $\mathscr{L}_u^1(f_u)$ and $\mathscr{L}_u^0(f_u)$
end for
if $|\mathbf{U}| \leq A$ **then**
    $\mathscr{A} \leftarrow \mathbf{U}$
    $F_0 \leftarrow 0$
else
    $\mathscr{A} \leftarrow$ the set of A users with highest $f_u^1 L_u$ value
    $F_0 \leftarrow \max_{u \notin \mathscr{A}} f_u^0 L_u$
end if
$F_1 \leftarrow \min_{u \in \mathscr{A}} f_u^1 L_u$
if $F_1 > F_0$ and $\forall u \in \mathscr{A}$, $\mathscr{L}_u^1(f_u^1) \leq \mathscr{L}_u^0(\frac{F_0}{L_u})$ **then**
    return profile $(\boldsymbol{f}, \mathscr{A})$ in which $\forall u \in \mathscr{A}$: $f_u = f_u^1$, o.w. $f_u = f_u^0$
else
    return "there is no SMDE"
end if

this, recall that in an SMDE, the attacker always plays a pure strategy targeting the set of users with the highest values of $f_u^1 L_u$, and this set is obviously unique. Algorithm 1 always finds the unique SMDE if it exists, and returns "no SMDE" if there is no SMDE. Finally, it is also easy to see that the running time of the algorithm is polynomial in the number of users.

Numerical Example Consider a game consisting of two users (user 1 and user 2) and an attacker, who can target only a single user (i.e., $A = 1$). Let $L_1 = L_2 = 1$, $N_1 = N_2 = \frac{1}{2}$, $C_1 = 1$, and $C_2 = 2$. Finally, let $FP(f) = (1-f)^2$, which obviously satisfies our assumptions about FP. Then, $f_1^1 = \frac{1}{4}$, $f_1^0 = \frac{3}{4}$, $f_2^1 = \frac{5}{8}$, and $f_2^0 = \frac{7}{8}$.

Now, we show that the game does not have an SMDE. From Lemma IV.2.2, we know that the attacker will target either user 1 or user 2. First, suppose that the attacker targets user 1 (i.e., $\mathscr{A} = \{1\}$). Then, from Theorem IV.2.1, we have that the users' strategies must be $f_1 = f_1^1 = \frac{1}{4}$ and $f_2 = f_2^0 = \frac{7}{8}$. However, this contradicts that $\mathscr{A}$ is a best response,

since $f_1 L_1 = \frac{1}{4} < \frac{7}{8} = f_2 L_2$. Second, suppose that the attacker targets user 2. Then, the user's strategies must be $f_1 = f_1^0 = \frac{3}{4}$ and $f_2 = f_2^1 = \frac{5}{8}$, which contradicts that $\mathscr{A}$ is a best response, since $f_1 L_1 = \frac{3}{4} > \frac{5}{8} = f_2 L_2$.

For now, we have characterized SMDE, which are formed by the users' selfish decisions. A crucial question regarding these equilibria is how close they are to the social optimum, i.e., to the strategies chosen by a social planner who is interested in minimizing the players? losses. The details of the discussion can be seen in [134].

IV.3 Conclusion

In order to mitigate the serious threat posed by spear-phishing attacks, defenders can deploy e-mail filters. However, the strategic nature of these attacks and the independent configuration of the e-mail filters may lead to a coordination problem. In this chapter, I studied this coordination problem by extending previous work on the strategic threshold-selection. I considered the dynamics from the perspective of Stackelberg model. I consider the multi-defender Stackelberg equilibrium in the model, and propose a polynomial algorithm to find such an equilibrium.

CHAPTER V

Decentralization and Security in Dynamic Traffic Light Control

In the chapter, I propose to systematically address the decentralized control problems described in the introduction by considering a multi-intersection scenario in which a) traffic light controllers take into account relative queue lengths to determine red-green state of the traffic lights at an intersection, b) controllers for all lights must be designed to work jointly so as to optimize overall traffic network performance, c) sensors feeding data into the controllers are vulnerable to denial-of-service attacks, and d) intersections are partitioned among a set of players, with own goals pertaining to congestion within their local municipal region, which are in general misaligned with global interests of the entire traffic network. In particular, I make the following contributions:

1. A scalable local search algorithm for multi-intersection controller design,
2. a game theoretic model of resilient control in the face of denial-of-service attacks,
3. a scalable algorithm for resilient control,
4. a game theoretic model of decentralized traffic light control involving multiple self-interested parties (e.g., municipalities), and
5. a scalable algorithm for approximating a Nash equilibrium for decentralized control games in both baseline and resilient (i.e, accounting for sensor attacks) settings.

I use the "Simulation of Urban MObility" (SUMO) [135] platform to implement, illustrate, and evaluate our approach.

V.1 Traffic Network Model

In the section, I introduce the control logic I use in the chapter, and define the metrics to measure the efficiency of a traffic system. The control logic is adapted and revised from

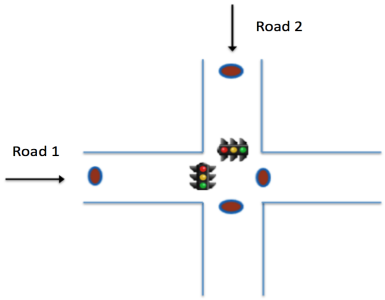

Figure V.1: Intersection

[136]. However, they only considered a single-intersection scenario. In the chapter, I will generalize the control logic into cases with multiple intersections and correspondingly multiple traffic lights.

Traffic Control System Consider a traffic network consisting of n intersections $I_1, I_2, \ldots, I_n$. We assume that each intersection is a cross of two "one-way" roads, and has no left turns (see Figure V.1).[1] In addition, we assume (as is common) that yellow light cycles are counted as a part of red lights cycles. Each direction j ($j = 1$, or 2) of intersection I_i $(1 \leq i \leq n)$ has an exogenously specified minimum *green* light cycle length $\Omega_{i,j,min}$ and maximum *green* light cycle length $\Omega_{i,j,max}$. We assume that each intersection I_i has two sensors in each direction j allowing us to count the number of vehicles, $m_{i,j}(t)$, queued at that intersection in direction j (specifically, an ingress sensor counts incoming vehicles, and an egress outgoing vehicles, with the difference giving us the queue length). For each direction j, we also define a clock variable $c_{i,j}(t)$, which measures the time since the last switch from *red* to the *green* of the traffic light for direction j.

Controllers For a given intersection I_i, we adopt a two-parameter control logic model from [136], which determines behavior based on a comparison of queue lengths $m_{i,j}(t)$ and corresponding thresholds s_{ij}. Intuitively, when queue length in a particular direction j exceeds the corresponding threshold, this direction is viewed as high-priority and congested.

[1] Allowing for "two-way" streets and left turns is a relatively straightforward generalization.

We thereby obtain four distinct cases:

- $M_{i,1} = \{(m_{i,1}, m_{i,2}) : m_{i,1}(t) < s_{i,1}, m_{i,2}(t) < s_{i,2}\}$;
- $M_{i,2} = \{(m_{i,1}, m_{i,2}) : m_{i,1}(t) < s_{i,1}, m_{i,2}(t) \geq s_{i,2}\}$;
- $M_{i,3} = \{(m_{i,1}, m_{i,2}) : m_{i,1}(t) \geq s_{i,1}, m_{i,2}(t) < s_{i,2}\}$;
- $M_{i,4} = \{(m_{i,1}, m_{i,2}) : m_{i,1}(t) \geq s_{i,1}, m_{i,2}(t) \geq s_{i,2}\}$.

Let $\Delta_i(t)$ denote the traffic light state for intersection I_i. $\Delta_i(t) = 1$ means the road 1 is *green* and road 2 is *red*; and $\Delta_i(t) = 2$ means road 2 is *green* and road 1 is *red*. Then we could get different control logic when it is in different state spaces.

For $M_{i,1}$ and $M_{i,4}$,

$$\Delta_i(t) = \begin{cases} 1 & \text{if } c_{i,1}(t) \in (0, \Omega_{i,1,max}) \text{ and } c_{i,2}(t) = 0, \\ 2 & \text{otherwise} \end{cases} \tag{V.1}$$

For $M_{i,2}$,

$$\Delta(t) = \begin{cases} 1 & \text{if } c_{i,1}(t) \in (0, \Omega_{i,1,min}) \text{ and } c_{i,2}(t) = 0, \\ 2 & \text{otherwise} \end{cases} \tag{V.2}$$

For $M_{i,3}$,

$$\Delta(t) = \begin{cases} 2 & \text{if } c_{i,2}(t) \in (0, \Omega_{i,2,min}) \text{ and } c_{i,1}(t) = 0, \\ 1 & \text{otherwise} \end{cases} \tag{V.3}$$

The equation V.1 shows that at some time t, if the number of vehicles in both directions is lower or higher than the corresponding thresholds, then they have the same priority for the two directions, and we will wait for the current green cycle to reach the maximum green cycle length. However, if one direction exceeds the threshold but not the other, then the former has a higher priority. Generally speaking, the above control logic is "interrupt" based, which could ensure that the road i always receives a minimum *green* light cycle

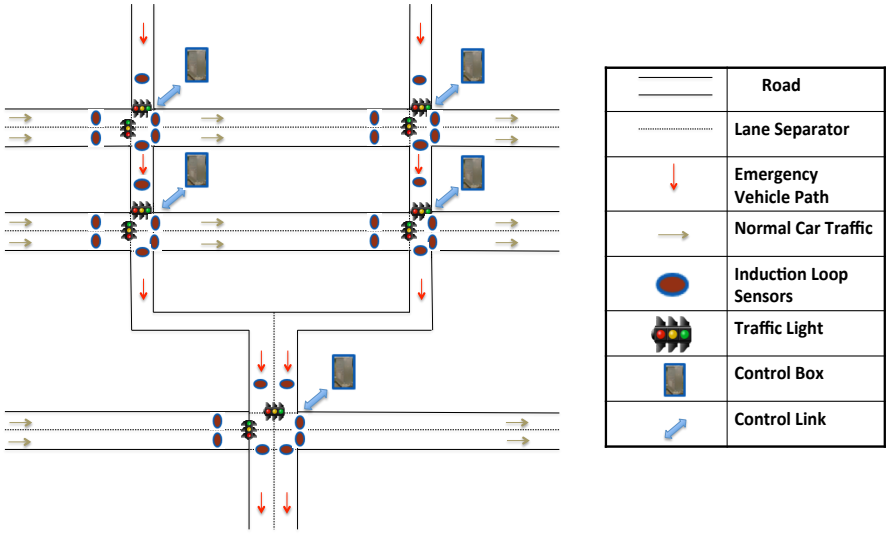

Figure V.2: Emergency Vehicle Scenario

$\Omega_{i,min}$, and the *green* light cycle may be dynamically interrupted anytime after $\Omega_{i,min}$ based on the vehicle information gathered from sensors. However, when the *green* light cycle reaches the $\Omega_{i,max}$, then the cycle has to be terminated and switched to red light cycle (correspondingly green light cycle for another direction).

Objective: Weighted Average Latency Assume there is a vehicle set V in the system, s.t. $|V| = d$. We assume $V = \{v_1, v_2, \ldots, v_d\}$, and every vehicle v_i $(1 \leq i \leq d)$ has a corresponding weight w_i which denotes the relative importance of the vehicle. For instance, an ambulance may have a higher weight than a common personal car. For each vehicle v_i traveling in the traffic system, *latency* l_i measures the time consumed for the car from entering the system to leaving the system. We can define *Weighted Average Latency (denote as $\mathscr{L}$)* as follows, and minimizing it is also the main goal for the manager of the system.

$$\mathscr{L} = \frac{\sum_{i=1}^{d} w_i l_i}{\sum_{i=1}^{d} w_i}$$

Optimization Problem

Consider a set of n intersections, $\{I_1, \ldots, I_n\}$, and associated threshold parameters

$$\mathbf{s} = \{\langle s_{1,1}, s_{1,2}\rangle, \langle s_{2,1}, s_{2,2}\rangle, \ldots, \langle s_{n,1}, s_{n,2}\rangle\}$$

in which $s_{i,j} \in \mathbb{R}^+$ $(1 \leq i \leq n, j = 1,2)$. Our goal is to choose the parameters of all intersections s so as to minimize overall weighted latency, for a given weight vector w:

$$\min_{\mathbf{s}} \mathscr{L}(\mathbf{s}; w). \tag{V.4}$$

V.2 Optimizing Traffic Network Configuration

The optimization problem in Equation V.4 is intractable because the objective function is a challenge to evaluate even for a fixed parameter vector **s**, let alone optimize (typically, as below, it is evaluated by running simulations). Rather than exhaustive search, we propose a *coordinate greedy* (or just *CGA*) local search method for efficiently computing an approximately optimal configuration **s**. The proposed algorithm, Algorithm 2, works by

Algorithm 2 Coordinate Greedy Algorithm (CGA)

input: Starting Parameter set $\hat{\mathbf{s}}$
return: Local Minimal Parameter set $\mathbf{s}^*$

1: Copy $\hat{\mathbf{s}}$ to $\mathbf{s}^*$
2: **while** There exists an intersection, such that we could change parameters of the intersection to make $\mathscr{L}$ smaller **do**
3: Make the change to $\mathbf{s}^*$
4: **end while**
5: **return** $\mathbf{s}^*$

first discretizing parameters for each traffic light i, and then iteratively choosing a particular traffic light, and finding the optimal configuration of parameters of this light, *keeping configuration of the rest fixed*.

V.3 Resilient Traffic Network Control

Game Theoretic Model

In order to formally investigate the consequence of DoS attacks on sensors, as well as the associated problem of resilient traffic network control (i.e., designing control parameters of all intersections so as to endow the network with a degree of resilience against DoS attacks), we consider a Stackelberg game model. In this model, the controller (defender) D first chooses the parameter vector $\mathbf{s}$, and the attacker $\mathscr{A}$ chooses a single sensor r_{ik} at a single intersection i to disable in order to maximally disrupt traffic, where $k = 0$ corresponds to an ingress and $k = 1$ an egress sensor. Formally, the defender's goal is to minimize weighted latency, $\mathscr{L}$, which the attacker aims to maximize.

Definition V.3.1. *A* Stackelberg equilibrium *of the* resilient network control *game* Γ *is* $(\mathbf{s}^*, r^*(\mathbf{s}^*))$*, such that* $\sum_{i,k} r_{ik} = 1$*,* r^* *maximizes latency* $\mathscr{L}(\mathbf{s}^*, r(\mathbf{s}))$*, and* $\mathbf{s}^*$ *minimizes the resulting maximal latency* $\mathscr{L}(\mathbf{s}, r^*(\mathbf{s}))$.

Resilient Control Algorithm

The goal of resilient network control is to choose $\mathbf{s}^*$ which is a part of a Stackelberg equilibrium accounting for the attacker's best response. Since, again, exhaustive search is clearly intractable, we propose an augmented version of the CGA algorithm, shown as Algorithm 3 (RCGA).

Algorithm 3 Resilient CGA (RCGA)

input: Local Optimal Parameter set $\hat{\mathbf{s}}$ given no Attacker
return: Resilient Parameter set $\mathbf{s}^*$

1: Given $\hat{\mathbf{s}}$ is applied in the system, enumerate sensors and find a sensor α, by attacking it we could get maximal $\mathscr{L}$.
2: Given α is attacked, we run Algorithm 2 starting from $\hat{\mathbf{s}}$
3: **return** resulting parameter set $\mathbf{s}^*$

V.4 Decentralized Control

Game Theoretic Model

I now present the most general model, in which multiple defenders determine configurations non-overlapping subsets of traffic lights. There may or may not be an attacker. If

the attacker is not considered, we view it as a baseline decentralized control game, whereas consideration of an attacker extends the model to a resilient decentralized control game.

Formally, assume there is a set of defenders D ($|D| \leq n$) who are in charge of different districts and corresponding traffic lights in a traffic network. Each defender d is only concerned about the Weighted Average Latency $\mathscr{L}_d$ for her own district d. Let $\mathbf{s}_d$ be the set of parameters controlled by defender $d \in D$, then $\mathbf{s}_d \cap \mathbf{s}_{\mathbf{d}'} = \emptyset$ for $d \neq d'$, and $\bigcup_d \mathbf{s}_d = \mathbf{s}$. Assume there is an attacker $\mathscr{A}$ who can attack a sensor in the system, and her goal is to increase the overall $\mathscr{L}$ of the system. We define *Multi-Defender Traffic Control Game* as follows:

Definition V.4.1. *A* Multi-Defender Traffic Control Game *is defined by a tuple* $\Gamma = [P, S(\cdot), v(\cdot)]$, *where*

- *There is a set of players* $P = \{D, \mathscr{A}\}$, *in which D is set of defenders in the traffic system,* $\mathscr{A}$ *is an attacker;*
- *For each* $d \in D$, *the strategy space of defender d is the space of parameters under her control, i.e.* $S(d) = \{\mathbf{s}_d\}$; *for the attacker* $\mathscr{A}$, *her strategy space is the set of sensors in the system.*
- *For each* $d \in D$, *she only cares about the Weighted Average Latency in her own district, i.e. the payoff of d is* $v(d) = -\mathscr{L}_d$; *for the attacker* $\mathscr{A}$, *her payoff is the Weighted Average Latency in the overall system, i.e. the payoff of attacker is* $v(\mathscr{A}) = \mathscr{L}$.

Approximating Equilibrium in Decentralized Control

As previously mentioned, a traffic system may consist of different districts, which may be owned and controlled by different agents. Sometimes, these agents may have varying or sometimes conflicting interests. When there are multiple defenders, I consider the game as a Normal-Form game among multiple agents, and my goal is to compute or approximate a *Nash equilibria* among them.

Definition V.4.2. *A profile of traffic light parameters* $\mathbf{s}$ *is a* Nash equilibrium *if no defender* $d \in D$ *can reduce its latency* $\mathscr{L}_d$ *by unilaterally reconfiguring its lights through changing* $\mathbf{s}_d$.

Since computing a Nash equilibrium in our setting is intractable, I propose a simple iterative best response algorithm (Algorithm 4), in which each traffic light is chosen in a given iteration, and the associated defender d optimizes parameters of this traffic light only to minimize $\mathscr{L}_d$, fixing all other parameters. We refer to this algorithm as *BRA* (best response algorithm).

Algorithm 4 Best Response Algorithm (BRA)

input: Starting Parameter set $\hat{\mathbf{s}}$
return: Equilibrium Parameter set $\mathbf{s}^*$

1: Copy $\hat{\mathbf{s}}$ to $\mathbf{s}^*$
2: **while** There exists an defender d, such that we could change parameters of the intersection controlled by defender d to make $\mathscr{L}_d$ smaller **do**
3: Make the change to $\mathbf{s}^*$
4: **end while**
5: **return** $\mathbf{s}^*$

Approximating Equilibrium in Resilient Decentralized Control

Finally, I consider the decentralized setting, but now allowing for an attacker who will optimally respond to the joint configuration of all traffic lights by all defenders, $\mathbf{s}$. Formally, each Defender d wants to minimize $\mathscr{L}_d$ to improve resilience, and the attacker wants to maximize overall $\mathscr{L}$ by attacking a sensor. I propose a resilient extension of *BRA*, shown

Algorithm 5 Finding Resilient Parameter Set When There are Multiple Defenders

input: Equilibrium Parameter set $\hat{\mathbf{s}}$ given no Attacker
return: Resilient Parameter set $\mathbf{s}^*$

1: Given $\hat{\mathbf{s}}$ is applied in the system, enumerate sensors and find a sensor α, by attacking it we could get maximal $\mathscr{L}$.
2: Given α is attacked, we run Algorithm 4 starting from $\hat{\mathbf{s}}$
3: **return** resulting parameter set $\mathbf{s}^*$

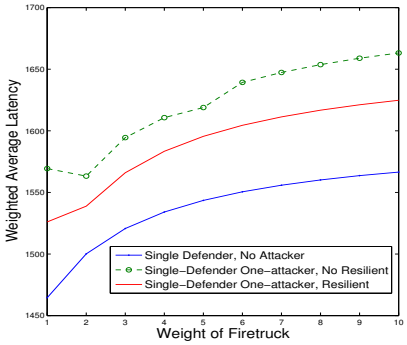

Figure V.3: Comparison with single-defender

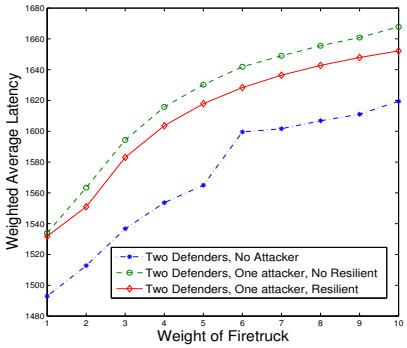

Figure V.4: Comparison with no attacker (baseline) configuration

in Algorithm 6, in which each best response iteration now accounts for the attacker's sensor DoS attack strategy.

Finally, I consider the decentralized setting, but now allowing for an attacker who will optimally respond to the joint configuration of all traffic lights by all defenders, $\mathbf{s}$. Formally, each Defender d wants to minimize $\mathscr{L}_d$ to improve resilience, and the attacker wants to maximize overall $\mathscr{L}$ by attacking a sensor. I propose a resilient extension of *BRA*, shown in Algorithm 6, in which each best response iteration now accounts for the attacker's sensor DoS attack strategy.

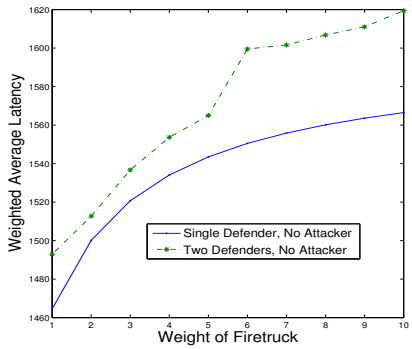

Figure V.5: Comparison with resilient single-defender configuration

Algorithm 6 Finding Resilient Parameter Set When There are Multiple Defenders

input: Equilibrium Parameter set $\hat{\mathbf{s}}$ given no Attacker
return: Resilient Parameter set $\mathbf{s}^*$

1: Given $\hat{\mathbf{s}}$ is applied in the system, enumerate sensors and find a sensor α, by attacking it we could get maximal $\mathscr{L}$.
2: Given α is attacked, we run Algorithm 4 starting from $\hat{\mathbf{s}}$
3: **return** resulting parameter set $\mathbf{s}^*$

V.5 Evaluation and Results

To implement the traffic control algorithm and perform simulation, I employ a simulation suit called SUMO (short for "Simulation of Urban MObility"). SUMO [135] is an open source, highly portable, microscopic road traffic simulation package designed to handle large road networks. SUMO also provides a Traffic Control Interface (TraCI) to let external controllers control the traffic. In our work, we use a Python script to control the simulation through TraCI and implement our control algorithm.

In the chapter, I consider an Emergency Vehicle Scenario (see Figure V.2). In the scenario, there are some common cars traveling from west to east, and some emergency vehicles (firetrucks) traveling from north to south. Assume that common cars have weight 1, and emergency vehicles have higher weights. There are some traffic lights that can be controlled in the intersections of the scenario. Before each direction in an intersection, there are two sensors that count the number of vehicles.

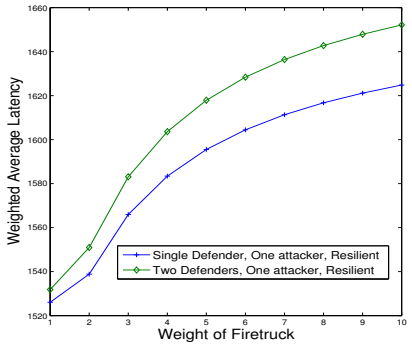

Figure V.6: Comparison with decentralized solutions

The experiment results can be seen in Figure V.3, V.4, V.5 and V.6, which shows the overall Weighted Average Latency $\mathscr{L}$ as a function of firetruck weights. When there is a single defender (Figure V.3) and no attacker, we obtain a relatively low $\mathscr{L}$ (applying Algorithm 2). However, the figure shows that an attack on the non-resilient configuration can substantially elevate $\mathscr{L}$: ignoring the possibility of a DoS attack can be disastrous for traffic in this scenario. On the other hand, resilient configuration (applying Algorithm 3) performs substantially better under attack.

Next, I split the scenario into two parts, in the upper part, one defender is in charge of the upper two intersections and another defender is in charge of the lower three intersections. And each defender only cares about Weighted Average Latency of her own district. The result is shown in Figure V.4. As we can observe, considering resilience is beneficial for the defenders; however, the benefit is smaller than if there were only one defender.

From Figures V.5 and V.6, we can observe how decentralization impacts the efficiency of a system. By comparing single-defender cases and two-defender cases, we find that the overall $\mathscr{L}$ in two-defender cases is higher than that in single defender cases, with and without attacker. It comes from the *negative externalities* introduced to the system when there are multiple selfish defenders, which make the overall system behavior inefficient.

V.6 Conclusion

I considered decentralization and security issues in dynamic traffic light control as a multi-defender security game. I proposed a game theoretic model and simulation-based optimization and equilibrium approximation algorithms to address the problem. I then implemented and evaluated our algorithms on the SUMO platform.

There are a number of future research directions that can be considered. One such direction is to investigate the scalability of our approach to significantly larger and more complex scenarios. Additionally, I only consider DoS attack on sensors. In future work, it will be important to evaluate resilience in the context of integrity attacks as well.

Part II

Toward Efficiency: Coalition Formation Mechanism

CHAPTER VI

Mechanism Design in Coalition Formation

In this chapter, I discuss the mechanism design in coalition formation (team formation). In [45], we consider an accept-reject game (ARG) in which players propose coalitions in a predetermined order, and prospective teammates may choose to either accept or reject the proposals in ARGs. We have shown some theoretical properties of the subgame perfect Nash equilibrium (SPNE), such as individual rationality and iterated matching of soulmates [39], where soulmate coalitions are matched in an iterative fashion. In [45], we also demonstrate that the SPNE of Rotating Proposer Game (RPG), which is a special case of ARG, is Pareto efficient. Rotating Proposer Mechanism (RPM) is a mechanism that implements the SPNE of the corresponding RPG. As the SPNE is highly nontrivial to compute, I focus primary on the computation of the equilibrium and evaluate the RPM by empirical methods in the thesis.

One significant challenge in implementing RPM in practice is the combinatorial complexity of backward induction. I address this issue in two ways. First, I use the IMS (iterated matching of soulmates) as an efficient preprocessing and pruning procedure. I show experimentally that this significantly reduces the computational burden of performing backward induction. Second, I develop a method to approximate RPM on the roommate problem which allows us to trade off computation time and quality of approximation of the subgame perfect equilibrium in RPM. My experiments demonstrate that there is a natural tradeoff point which allows us to retain most of the positive properties of RPM at a significantly reduced computational overhead. To enable scalability on problems with larger coalitions, I propose *Heuristic Rotating Proposer Mechanism (HRPM)* which uses heuristics both to determine which coalitions are proposed and whether they are accepted.

In extensive experiments, I evaluate the economic properties of RPM in both exact and approximate versions, compared to the RSD and a recent alternative, *one-player-one-pick*

(OPOP) shown previously to be highly effective. I observe that in all instances, exact and approximate RPM is significantly more efficient (in terms of social welfare) and more equitable than RSD in the roommate problem. I also observe that HRPM significantly outperforms both RSD and OPOP in settings where coalitions' size can be at most 3 on both of these metrics in nearly all cases. Moreover, using an algorithm for finding an upper bound on untruthful players, I show that RPM and its approximate versions introduce few incentives to lie.

VI.1 Problem Setting

I consider the standard model described in [97] of an environment populated a set of players $N = \{1, \ldots, n\}$ who need to be partitioned into coalitions. A team (coalition) $T \in 2^N$ is a set of players, and a partition (coalition structure) π is a collection of coalitions such that: 1) for any distinct $T, T' \in \pi$, $T \cap T' = \emptyset$, and 2) $\cup_{T \in \pi} T = N$. For a player i, let π_i be the coalition in the partition π containing i. In many tasks, coalitions have some feasibility constraints; for example, one could constrain coalitions to consist of at most k individuals. Generically, let $\mathscr{T}$ denote the set of *feasible* coalitions, which we assume to always include singleton coalitions, $\{i\}$. For a player i, we denote a subset $\mathscr{T}_i \subset \mathscr{T}$ of coalitions that include i by $\mathscr{T}_i$. Each player $i \in N$ has a strict hedonic preference ordering $\succeq_i$ over $\mathscr{T}_i$. A profile of preferences $\succeq$ (or *profile* for short) is a list of preferences for every $i \in N$. Given a profile $\succeq$, the list of preferences for all players except i is denoted by $\succeq_{-i}$. A *coalition formation mechanism M* maps every preference profile $\succeq$ to a partition π, i.e. $\pi = M(\succeq)$.

VI.2 Team Formation Games and Rotating Proposer Mechanism

We begin by considering a natural sequential non-cooperative coalition formation game with complete information about hedonic preferences of all players, which will serve as the core component of the coalition formation mechanism below. We term such games *accept-reject games (ARGs)*, because they proceed through an exogenously specified order of players, with each player making a proposal of a coalition, and all prospective members

having a chance to accept or reject this proposal.

Formally, an ARG is defined by a set of players N, a preference profile $\succeq$, a set of feasible coalitions $\mathscr{T}$, and an ordered list of players $O = (o_1, o_2, \ldots, o_m)$, in which each player $i \in N$ is included at least once. The game proceeds through a series of rounds. In each round the next player i in the order list O proposes to a coalition $T \in \mathscr{T}_i$, with the constraint that i cannot have proposed to T in any prior round. Given a proposal T made by i, all $j \in T \setminus i$ sequentially decide whether to accept or reject the proposal.[1] If any player rejects, the entire proposal is rejected, and we proceed with the next round. If all $j \in T \setminus i$ accept, the coalition T is added to the partition π, all players are removed from the game and from O, and the game proceeds to the next round, unless no players remain (in which case the game ends with a partition π). If after m rounds there are players remaining, they each become singleton coalitions, completing the partition. Algorithm 7 describes the game procedure more precisely. And the Example VI.2.1 shows an example of the game.

Example VI.2.1. *Consider an ARG with four players $N = \{1,2,3,4\}$, and the order of proposers $O = (1, 2, 3, 4)$ in which the size of each coalition is at most two. Suppose that the profile is as follows:*

$$
\begin{aligned}
1: \quad & \{1,4\} \succeq_1 \{1,2\} \succeq_1 \{1,3\} \succeq_1 \{1\} \\
2: \quad & \{2,1\} \succeq_2 \{2,4\} \succeq_2 \{2,3\} \succeq_2 \{2\} \\
3: \quad & \{3,2\} \succeq_3 \{3,1\} \succeq_3 \{3,4\} \succeq_3 \{3\} \\
4: \quad & \{4,3\} \succeq_4 \{4,2\} \succeq_4 \{4,1\} \succeq_4 \{4\}
\end{aligned}
$$

The following is an example scenario:

1. *Player* 1 *proposes to* $\{1,4\}$*, and* 4 *rejects the proposal.*

2. 2 *proposes to* $\{2,1\}$ *and* 1 *accepts the proposal.* 1 *and* 2 *are removed from the game.*

[1] The order of this sequence can be arbitrary; for example, it can follow the order of the players' first appearance in O.

Algorithm 7 Construction of ARG

input: $(N, \succeq, \mathscr{T}, O)$

return: Coalition formation outcome π

$\pi = \emptyset$
while O is non-empty **do**
    $i \leftarrow$ the first player in O
    Player i proposes to a coalition $T \in \mathscr{T}_i(N)$
    All the players in T sequentially decide whether to accept i's proposal (if $T = \{i\}$, the proposal is automatically accepted)
    if All players in T accept player i's proposal **then**
        $\pi \leftarrow \pi \cup \{T\}$
        $N \leftarrow N \backslash T$
        for each $j \in T$ **do**
            $O \leftarrow O \backslash \{j\}$
        end for
    else
        Get new O by removing the first instance i in O
    end if
    for each player $i \in N$ **do**
        for each feasible coalition $T \in \mathscr{T}_i$ **do**
            if $T \not\subset N$ **then**
                $\mathscr{T}_i \leftarrow \mathscr{T}_i \backslash \{T\}$
            end if
        end for
    end for
end while
while N is non-empty **do** ▷ add singletons into the outcome.
    pick an arbitrary instance i from N
    $\pi \leftarrow \pi \cup \{i\}$
    $N \leftarrow N \backslash \{i\}$
end while
return π

3. 3 *propose to* $\{3,4\}$ *and* 4 *accepts the proposal.* $3,4$ *are removed.*

The partition that results from this sequence is $\pi = \{\{1,2\},\{3,4\}\}$.

In our working paper [45], we demonstrate that there are several important properties that hold in any subgame perfect Nash equilibrium of an arbitrary accept-reject game:

- individual rationality (players are not a part of any coalition if they would prefer to be by themselves),
- matching of soulmates (players who all prefer to be together are matched) (see [39] for more details), and
- when the game is "IMS-complete" [39], the outcomes are in the core of the derived cooperative game.

However, we cannot guarantee the Pareto efficiency of the SPNE of any ARG. In [45], we show that if we add some constraints to the proposing order O, then we can make sure the Pareto efficiency of the ARM. We propose a class of ARGs which we term rotating proposer games (RPG). In the RPG, the order O over players is such that each player i can make $|\mathscr{T}_i| + 1$ proposals before we move on to another player.

We also introduced a centralized coalition formation mechanism, termed Rotating Proposer Mechanism (RPM), which implements the subgame perfect Nash equilibrium of the RPG in which all proposals are accepted. In this equilibrium, whenever it's a player i's turn to propose, i makes a proposal to her most preferred coalition among those that would be accepted. For any profile, if all players report their preferences truthfully, equilibrium outcomes of the game have a number of good properties which are thereby inherited by RPM. Of particular note is that RPM is individually rational, Pareto optimal, and implements IMS. However, it is also immediate from known results that the RPM mechanism is not in general strategyproof [39].

VI.3 Implementing RPM

In [45], we showed that RPM has important theoretical advantages. However, it is computationally challenging to implement. In particular, the size of the backward induction search tree is $O(2^{\Sigma_{i=1}^{n}|\mathscr{T}_i|})$. Even in roommate problem, in which the size of coalitions is at most two, computing SPE is $O(2^{n^2})$. We address this challenge in three ways: (1) preprocessing and pruning to reduce the search space, (2) approximation for the roommate problem, and (3) a general heuristic implementation.

VI.3.1 Preprocessing and Pruning

One of the central properties of RPM is that it implements iterative matching of soulmates. In fact, it does so in every subgame in the backward induction process. Now, observe that computing the subset of coalitions produced through IMS is $O(n^3)$ in general, and $O(n^2)$ for the roommate problem, and is typically much faster in practice. We therefore use it as a preprocessing step both initially (reducing the number of players we need to consider in backward induction), and in each subgame of the backward induction search tree (thereby pruning irrelevant subtrees).

VI.3.2 Approximate RPM for the Roommate Problem

Using IMS for preprocessing and pruning does not sufficiently speed up RPM computation in large-scale problem instances. Next, we developed a parametric approximation of RPM which allows us to explicitly trade-off computational time and approximation quality. We leverage the observation that the primary computational challenge of applying RPM to the roommate problem is determining whether a proposal is to be accepted or rejected. If we are to make this decision without exploring the full game subtree associated with it, considerable time can be saved. Our approach is to use a heuristic to evaluate the "likely" opportunity of getting a better teammate in later stages: if this heuristic value is very low, the offer is accepted; if it is very high, the offer is rejected; and we explore the full subgame in the balance of instances.

More precisely, consider an arbitrary offer from i to another player j. Given the subgame of the corresponding RPM, let $\mathscr{U}_j(i)$ denote the set of feasible teammates that j prefers to i, and let $\mathscr{U}_j(j)$ be the set of feasible teammates who j prefers to be alone. We can use these to heuristically compute the likelihood $R_j(i)$ that j can find a better teammate than the proposer i:

$$R_j(i) = \frac{|\mathscr{U}_j(i)|}{|\mathscr{U}_j(j)|} \cdot \frac{1}{|\mathscr{U}_j(i)|} \sum_{k \in \mathscr{U}_j(i)} \left(1 - \frac{|\mathscr{U}_k(j)|}{|\mathscr{U}_k(k)|}\right) = \frac{1}{|\mathscr{U}_j(j)|} \sum_{k \in \mathscr{U}_j(i)} \left(1 - \frac{|\mathscr{U}_k(j)|}{|\mathscr{U}_k(k)|}\right) \tag{VI.1}$$

Intuitively, we first compute the proportion of feasible teammates that j prefers to i. Then, for each such teammate k, we find at the proportion of feasible teammates who are not more preferred by k than the receiver j. Our heuristic then uses an exogenously specified threshold, α, ($0 \leq \alpha \leq 0.5$) as follows. If $R_j(i) \leq \alpha$, player j accepts the proposal, while if $R_j(i) \geq 1 - \alpha$, the proposal is rejected. In the remaining cases, our heuristic proceeds with evaluating the subgame at the associated decision node. Consequently, when $\alpha = 0$, it is equivalent to the full backward induction procedure, and computes the exact RPM. Note that for any α, this approximate RPM preserves IR, and we also maintain IMS by running it as a preprocessing step.

VI.3.3 Heuristic Rotating Proposer Mechanism (HRPM)

Unlike the roommate problem, general coalition formation problems have another source of computational complexity: the need to iterate through the combinatorial set of potential coalitions to propose to. Moreover, evaluating acceptance and rejection becomes considerably more challenging. We therefore develop a more general heuristic which scales far better than the approaches above, but no longer has the exact RPM as a special case. We term the resulting approximate mechanism *Heuristic Rotating Proposer Mechanism (HRPM)*, and it assumes that the sole constraint on coalitions is their cardinality and that preferences can be represented by an additively separable utility function [97]. With the latter assumptions, we allow preferences over coalitions to be represented simply as pref-

erence orders over potential teammates, avoiding the combinatorial explosion in the size of the preference representation.

In HRPM, each proposer i attempts to add a single member to their team at a time in the order of preferences over players. If the potential teammate j accepts i's proposal, j is added to i's coalition, and i proposes to the next prospective teammate until either the coalition size constraint is reached, or no one else who i prefers to be alone is willing to join the coalition. Player j's decision to accept or reject i's proposal is based on calculating $R_j(l)$ for each member l of i's current coalition T using Equation VI.1, and then computing the average for the entire coalition, $R_j(T) = \frac{1}{|T|}\sum_{l \in T} R_j(l)$ (see Algorithm 8 for the fully precise description of HRPM). We then use an exogenously specified threshold $\beta \in [0,1]$, where j accepts if $R_j(T) \leq \beta$ and rejects otherwise. The advantage of HRPM is that the coalition partition can be found in $O(\omega n^2)$, where ω is the maximum coalition size. The disadvantage, of course, is that it only heuristically implements RPM. Crucially, it does preserve IR, and IMS is implemented as a preprocessing step.

VI.4 Experiment

My evaluation considers two coalition formation settings: (1) the *roommate problem*, where coalitions are capped at 2, and (2) the *trio-roommate* problem, with coalitions of at most 3. We note that both of these problems are essentially open from a mechanism design perspective: in either case, RSD is the only known mechanism which is either Pareto efficient or incentive compatible even in a well-understood restricted setting. No mechanism is known for these problems which are both IR and Pareto efficient, or IR and implements IMS. We benchmark RPM and its approximate variants to RSD in the roommate problem, and additionally to the One-Player-One-Pick (OPOP) mechanism [44] in the trio-roommate setting (OPOP and RSD are equivalent in the roommate problem). OPOP first chooses a set of captains, and then captains choose a single teammate at a time, while non-captains choose a coalition to join following a heuristic evaluation function. In either case, such "proposals"

Algorithm 8 Heuristic Rotating Proposer Mechanism (HRPM)

input: $(N, \succeq, O)$, ω, β
return: Team formation outcome π

$\pi = \emptyset$
while O is non-empty **do**
    $i \leftarrow$ the first player in O
    $\pi_i \leftarrow \{i\}$
    while $|\pi_i| < \omega$ **do**
        if $\succeq_i$ is empty or the first player in $\succeq_i$ is i **then**
            $O \leftarrow O \backslash \{i\}$
            break
        end if
        Player i proposes to the first player j in $\succeq_i$
        for each $l \in \pi_i$ **do**
            Calculate $R_j(l)$ based on equation VI.1
        end for
        Calculate $R_j(\pi_i) = \frac{1}{|\pi_i|} \sum_{l \in \pi_i} R_j(l)$
        if $R_j(\pi_i) \leq \beta$ **then** ▷ player j accepts the proposal
            $\pi_i \leftarrow \pi_i \cup \{j\}$
            $O \leftarrow O \backslash \{j\}$
            $N \leftarrow N \backslash \{j\}$
        end if
        Delete j from $\succeq_l$ for each player $k \in N$
    end while
    $O \leftarrow O \backslash \{i\}$
    $N \leftarrow N \backslash \{i\}$
    Delete i from $\succeq_l$ for each player $k \in N$
end while
while N is non-empty **do** ▷ add singletons into the outcome.
    pick an arbitrary instance i from N
    $\pi \leftarrow \pi \cup \{i\}$
    $N \leftarrow N \backslash \{i\}$
end while
return π

are always accepted. OPOP was previously shown to outperform several others, including RSD, in terms of social welfare and fairness [44], but does not satisfy any of our desiderata. In all mechanisms, players are ordered randomly.

VI.4.1 Data Sets

For evaluating our proposed mechanisms, we use both synthetic and real hedonic preference data. In both cases, preferences were generated based on a social network structure in which a player i is represented as a node and the total order over neighbors is then generated randomly. And non-neighbors represent undesirable teammates (i would prefer being alone to being teamed up with them). The networks used for our experiments were generated using the following models:

- **Scale-free network:** We adopt the Barabási-Albert model ([137]) to generate scale-free networks. For each (n,m), where n is the number of players, m denotes the density of the network, we generate $1,000$ instances of networks and profiles.

- **Karate-Club Network [138]:** This network represents an actual social network of friendships between 34 members of a karate club at a US university, where links correspond to neighbors. We generate 100 preference profiles based on the network.

Finally, we used a *Newfrat* dataset [139] which contains 15 matrices recording weekly sociometric preference rankings from 17 men attending the University of Michigan. In order to quantitatively evaluate both the exact and approximate variants of RPM, the ordinal preferences $\succeq_i$ have to be converted to cardinal ones $u_i(\cdot)$, upon which both mechanisms operate. For this purpose, we introduce a *scoring function* suggested by Bouveret and Lang [140] to measure a player's utility. The scoring function is a non-increasing function $g : [1..k] \to \mathbb{R}$ for some k and $[1..k]$ is the list of integers from 1 to k. To compute a player i's utility of player j we adopt *normalized Borda scoring function*, defined as $u_i(j) = g(r) = 2(k-r+1)/k-1$, where k is the number of i's neighbors, and $r \in [1..k]$ is the rank of j in

i's preference list. Without loss of generality, for every player i we set the utility of being a singleton $u_i(i) = 0$. As mentioned above, in trio-roommate problem we assume that the preferences of players are additively separable [97], which means that a player i's utility of a coalition T is $u_i(T) = \sum_{j \in T} u_i(j)$.

VI.4.2 Computing and Approximating RPM

We begin by investigating the relationship between running time and the approximation quality of our approaches to the roommate problem. Our simulations were performed on Mac OS 10.11 with a 2.6 GHz Intel Core i5 processor.

IMS Preprocessing

First, we show the computational value of IMS in preprocessing and pruning using synthetic preference profiles based on the generative scale-free model.

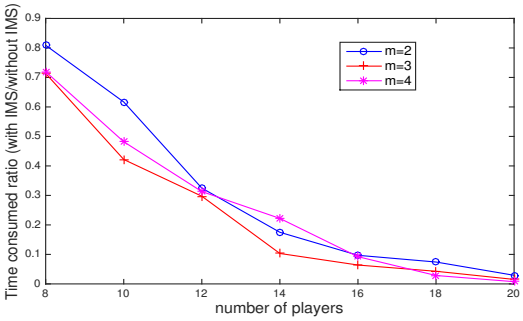

Figure VI.1: Time consumed ratio (with IMS/without IMS) for RPM on scale-free networks

Figure VI.1 shows the ratio of time consumed by RPM with IMS to that without IMS. In all cases, we see a clear trend that using IMS in preprocessing and pruning has increasing importance with increased problem size. The key takeaway is that *implementing IMS has both important economic and computational consequences*.

Approximating RPM

The parameter α of our approximation method for RPM in the roommate problem allows us to directly evaluate the tradeoff between running time and quality of approxi-

mation: small α will lead to less aggressive use of the acceptance/rejection heuristic, with most evaluations involving actual subgame search, while large α yields an increasingly heuristic approach for computing RPM, with few subgames fully explored.

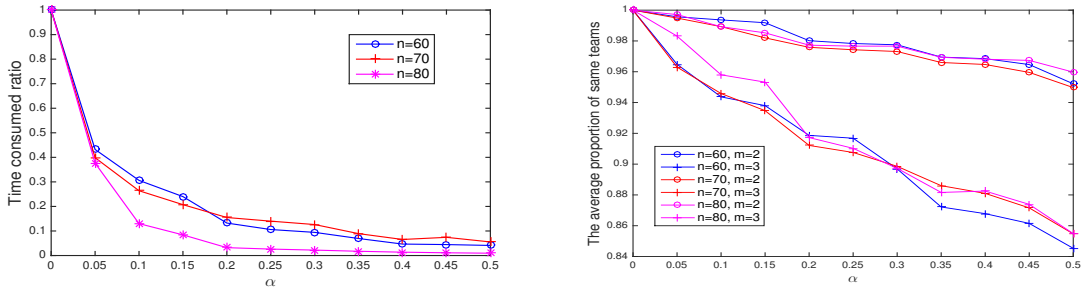

(a) Time consumed ratio (b)

Figure VI.2: Time consumed and average proportion of same coalitions

Figure VI.2a depicts the fraction of time consumed by RPM with different values of α compared to exact RPM (when $\alpha = 0$) on scale-free networks ($m = 3$). Based on this figure, even a comparatively small value of α dramatically decreases computation time. Figure VI.2b compares the similarity of the final coalition partition when using the heuristic compared to the exact RPM. Notice that even for high values of α, there is a significant overlap between the outcomes selected by RPM with and without the heuristic. $\alpha = 0.1$ appears to trade off approximation quality and running time particularly well: for comparatively sparse networks (i.e., $m = 2$) it yields over 99% overlap with exact RPM (this proportion is only slightly worse for denser networks), at a small fraction of the running time. Henceforth, we use $\alpha = 0.1$ when referring to the approximate RPM.

VI.4.3 Utilitarian Social Welfare

Ex post Pareto efficiency, satisfied by both RSD and RPM, is a very weak criterion. Conversion of ordinal to cardinal preferences allows us to consider empirically *utilitarian social welfare*, a much stronger criterion commonly used in mechanism design with cardinal preferences. We define social welfare as $\frac{1}{|N|}\sum_{i \in N} u_i(\pi_i)$, where π_i is the coalition that i was assigned to by the mechanism.

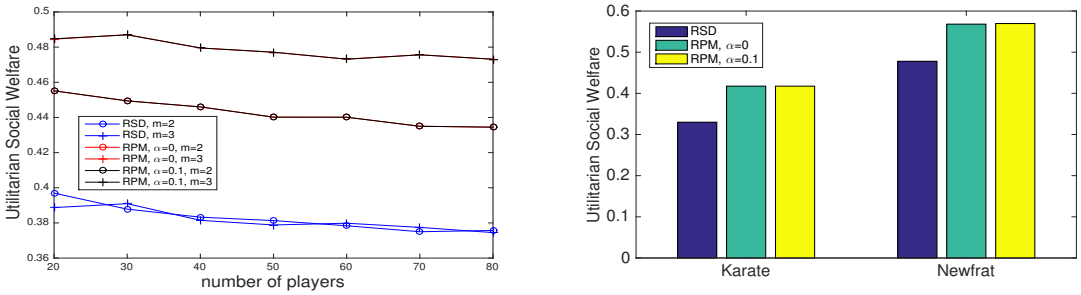

(a) scale-free networks (b) Karate and Newfrat

Figure VI.3: Utilitarian social welfare for roommate problem

Figures VI.3a and VI.3b depict the average utilitarian social welfare for RSD and RPM in the roommate problem on scale-free networks, Karate club networks, and the Newfrat data. In all cases, RPM yields significantly higher social welfare than RSD, with $15\% - 20\%$ improvement in most cases. These results are statistically significant ($p < 0.01$). Furthermore, there is virtually no difference between the exact and approximate RPM.

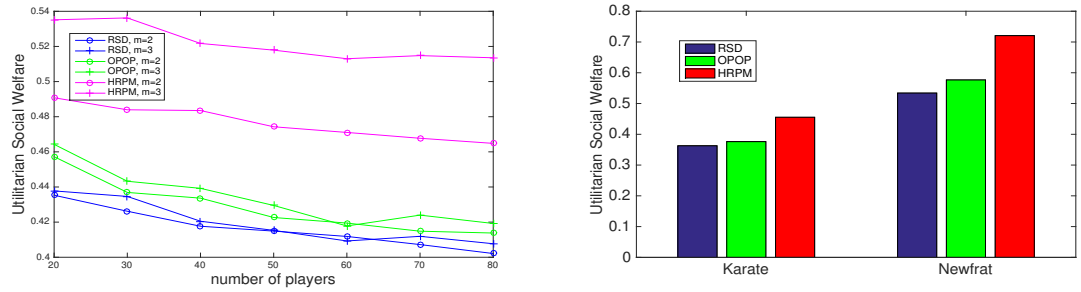

(a) scale-free networks (b) scale-free networks

Figure VI.4: Utilitarian social welfare for trio-roommate problem

For the trio-roommate problem, we compare HRPM ($\beta = 0.6$) with RSD and OPOP on the same data sets. Figures VI.4a and VI.4b show that HRPM yields significantly higher social welfare than both RSD and OPOP in all instances, and HPRM performs even better when the network is comparatively dense ($m = 3$ in the scale-free network). All results are statistically significant ($p < 0.01$).

VI.4.4 Fairness

A number of measures of fairness exist in prior literature. One common measure, envy-freeness, is too weak to use, especially for the roommates problem: every player who is not matched with his most preferred other will envy someone else. Indeed, because RPM matches soulmates—in contrast to RSD, which does not—it already guarantees the fewest number of players with envy. We consider two alternative measures, which aim to capture different and complementary aspects of fairness: maximum coalition utility difference, and the correlation between utility and rank in the random proposer order. Maximum coalition utility difference measures the difference in utility between teammates in each coalition T in a partition π, and takes the largest such difference over all coalitions. Formally, it computes $\max_{T\in\pi}(\max_{i\in T} u_i(T) - \min_{i\in T} u_i(T))$. Correlation between utility and rank considers each random ranking of players in O used for both RSD and RPM, along with corresponding utilities $u_i(\pi)$ of players for the partition π generated by the mechanism, and computes the correlation between these. It thereby captures the relative advantage that someone has by being earlier (or later) in the order to propose than others, and is a key cause of ex-post inequity in RSD.

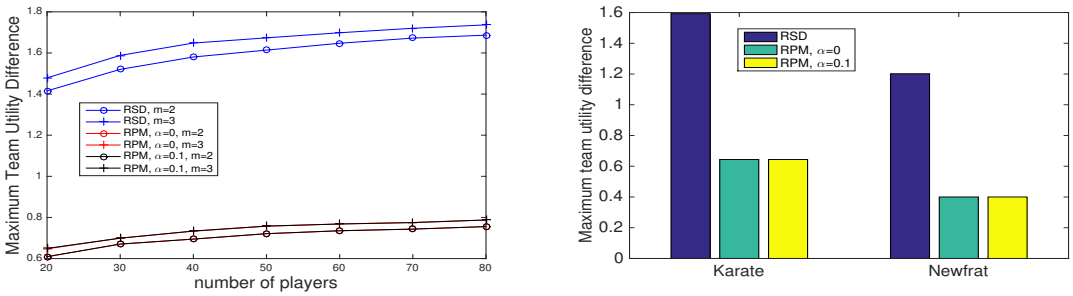

(a) scale-free network

(b) Karate and Newfrat

Figure VI.5: Maximum team utility difference for the roommate problem

Our experiments on the roommate problem show that RPM is significantly more equitable than RSD on scale-free networks (Figures VI.5a and VI.6a), as well as on the Karate club network and Newfrat dataset (Figures VI.5b and VI.6b). The differences between

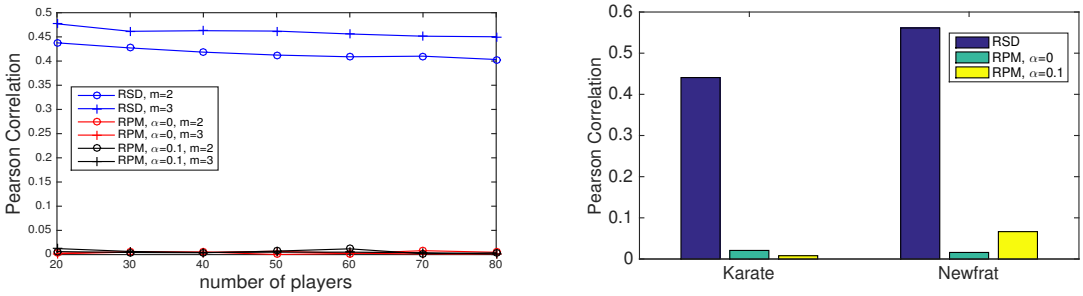

(a) scale-free network (b) Karate and Newfrat

Figure VI.6: Pearson Correlation for the roommate problem

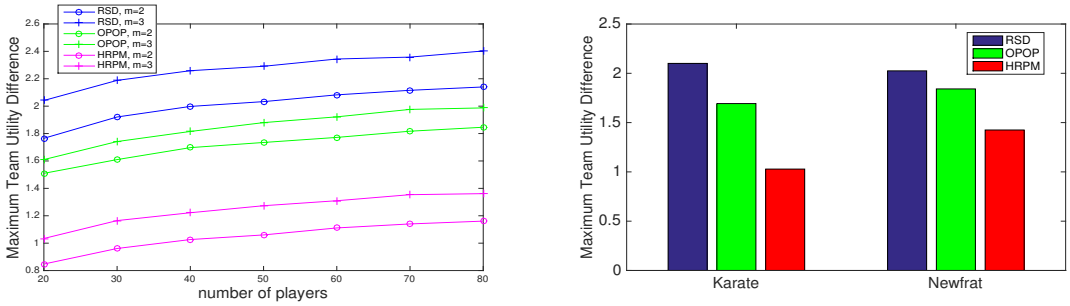

(a) scale-free network (b) Karate and Newfrat

Figure VI.7: Maximum Coalition Utility Difference for the trio-roommate problem

exact and approximate RPM are negligible in most instances.

In the trio-roommate problem, HRPM ($\beta = 0.6$) is much more equitable than both RSD and OPOP, except for correlation on Newfrat data, in which OPOP is better, as shown in Figures VI.7 and VI.8. These results are statistically significant ($p < 0.01$).

VI.4.5 Incentive Compatibility

Incentive compatibility is where RSD has a clear advantage over RPM. Although RPM is strongly incentive compatible on IMS-complete domains, we now explore its incentive properties empirically in more general settings. We focus on the roommate problem, because here we can compute an upper bound on the number of players with an incentive to lie using Algorithm 9, where we use $\mathcal{T}_i$ to denote the set of feasible teammates (since coalitions are of size at most 2). At the high level, this algorithm considers all the play-

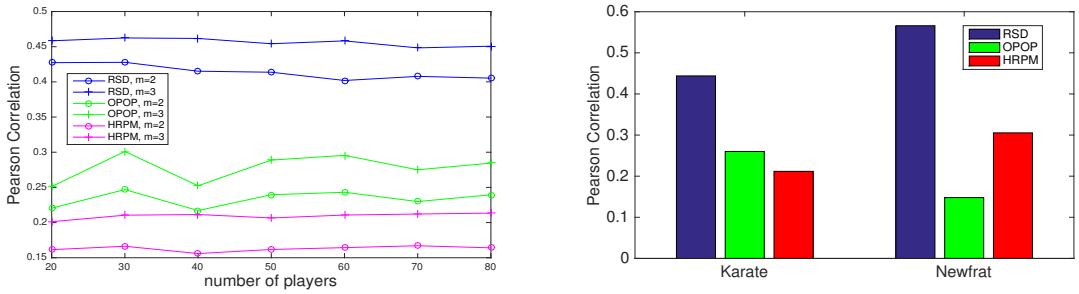

(a) scale-free network

(b) Karate and Newfrat

Figure VI.8: Pearson Correlation for the trio-roommate problem

ers who have accepted or rejected a proposal, and checks whether reversing this decision improves their outcomes. The following theorem shows that this method indeed finds the upper bound of untruthful players.

Theorem VI.4.1. *Algorithm 9 returns an upper bound on the number of players who can gain by misreporting their preferences.*

Proof. We divide the players into *proposers* and *receivers*. Proposers are those who proposed in RPM and were thus teamed up (including singleton coalitions). Receivers accepted someone's offer.

There are 4 cases:

1. *A proposer i untruthfully reveals its preference and remains a proposer.* In RPM, a proposer proposes to other players in order of preference. When i proposes to j, all others more preferred by i must have already rejected. Consequently, i cannot improve the utility by lying.

2. *A receiver j untruthfully reveals its preference and is still a receiver.* In this case, if j has an incentive to lie, there has to be a proposer i' who prefers j to its teammate under RPM, while j must prefer i' to its teammate. Steps $4-7$ in Algorithm 9 count all such instances.

Algorithm 9 Computing Upper Bound of Untruthful Players

input: $(N, \succeq, \mathscr{T}, O)$, teammate vector $teammate[]$ which results from RPM
return: number of potential untruthful players *Sum*

1: $Sum \leftarrow 0$
2: **while** $|O| \geq 2$ **do**
3: $proposer \leftarrow$ the first player in O
4: $receiver \leftarrow teammate[proposer]$
5: **for** player $i \in \mathscr{T}_{proposer}$ **do**
6: **if** $i \succeq_{proposer} receiver$ and $proposer \succeq_i teammate[i]$ **then**
7: $Sum \leftarrow Sum + 1$ ▷ i is potentially untruthful
8: **end if**
9: **end for**
10: **for** player $j \in \mathscr{T}_{receiver}$ **do**
11: **if** $j \succeq_{receiver} proposer$ and $receiver \succeq_j teammate[j]$ **then**
12: $Sum \leftarrow Sum + 1$ ▷ *receiver* is potentially untruthful
13: **end if**
14: **end for**
15: remove *proposer* and *receiver* from N, O and $\mathscr{T}$
16: **end while**
17: return *Sum*

3. *A proposer i untruthfully reveals her preference and becomes a receiver.* In this case, if i has an incentive to untruthfully reveal her preference, there has to be a proposer i' who prefer i to their teammate under RPM, and who i also prefers to its teammate. Steps $4-7$ in Algorithm 9 count all such instances.

4. *A receiver j untruthfully reveals its preference and becomes a proposer.* In this case, if j has an incentive to misreport its preference, there must be a receiver j' who prefers j to its teammate, while j must prefer j' to its teammate. Steps $8-10$ in Algorithm 9 count all such instances.

□

Table VI.1 presents the upper bound on the number of players with an incentive to lie, as a proportion of all players, on scale-free networks. We can observe that the upper bound is always below 0.2%, and is even lower when the networks are sparse ($m = 2$). On the

Table VI.1: Average Upper Bound of Untruthful Players for (Approximate) RPM

n	20	30	40	50	60	70	80
$m = 2$, $\alpha = 0$	0.015%	0.013%	0.013%	0.002%	0.008%	0.011%	0.010%
$m = 2$, $\alpha = 0.1$	0.015%	0.010%	0.015%	0.004%	0.022%	0.029%	0.036%
$m = 3$, $\alpha = 0$	0.105%	0.107%	0.072%	0.038%	0.037%	0.024%	0.023%
$m = 3$, $\alpha = 0.1$	0.115%	0.103%	0.085%	0.076%	0.065%	0.074%	0.093%

Table VI.2: Lower Bound of Profiles Where Every Player Is Truthful for (Approximate) RPM

n	20	30	40	50	60	70	80
$m = 2$, $\alpha = 0$	99.7%	99.6%	99.5%	99.9%	99.6%	99.2%	99.2%
$m = 2$, $\alpha = 0.1$	99.7%	99.7%	99.4%	99.8%	98.8%	98.1%	97.2%
$m = 3$, $\alpha = 0$	97.9%	96.8%	97.1%	98.1%	97.8%	98.4%	98.3%
$m = 3$, $\alpha = 0.1$	97.8%	96.9%	96.8%	96.2%	96.3%	95.1%	92.9%

Karate club data, we did not find any player with an incentive to lie in test cases when we apply (Approximate) RPM. On the Newfrat data, the upper bounds are less than 0.4% and 7% when we apply RPM without and with heuristics, respectively. In addition, we also computed the lower bound on the fraction of preference profiles where truth-telling is a Nash equilibrium (Table VI.2). We find that without the heuristic, when $m = 2$ (sparse networks), RPM is incentive compatible in more than 99% of the profiles; and when $m = 3$ (the networks are comparatively dense), RPM is truthful at least 96% of the time.

Table VI.3: Average Upper Bound of Untruthful Players for HRPM

n	20	30	40	50	60	70	80
$m = 2$, $\beta = 0.5$	1.44%	1.77%	1.71%	2.00%	2.09%	2.16%	2.06%
$m = 2$, $\beta = 0.6$	1.62%	1.83%	1.96%	2.09%	2.25%	2.11%	2.11%
$m = 3$, $\beta = 0.5$	2.99%	3.36%	3.76%	3.90%	4.18%	4.02%	4.33%
$m = 3$, $\beta = 0.6$	3.44%	3.69%	3.97%	3.98%	4.40%	4.24%	4.52%

Table VI.3 presents the upper bound on the number of untruthful players for HRPM

(still for the roommate problem). Even with this heuristic, we can see that less than 5% of the players have any incentive to misreport preferences.

VI.5 Conclusion

I mainly address the computational challenges in implementing *rotating proposer mechanism*, which implements a subgame perfect Nash equilibrium in the corresponding rotating proposer game, and evaluate the mechanism by empirical methods. To address the challenges, I introduce preprocessing and pruning, as well as approximate versions of RPM, one tailored to the roommate problem (with coalitions of at most two), and another for coalitions of arbitrary size. The experiments show that even the approximate versions of RPM significantly outperforms several alternative mechanisms for coalition formation in terms of social welfare and fairness, are do not introduce significant incentives to misreport preferences.

CHAPTER VII

Automated Mechanism Design for Roommates Problem

In the thesis, I consider a new perspective on the roommates problem based on *automated mechanism design (AMD)* [50]. In a prototypical AMD setup, one obtains preferences from the players, and then solves an optimization problem (for example, an integer linear program) in which constraints ensure incentive compatibility. However, applying AMD to matching problems in general, and the roommates problem in particular, faces a number of challenges. First, it is conventional to consider ordinal, rather than cardinal preference reports by the players. Second, incentive compatibility is often incompatible with other highly desirable properties, such as stability and, for cardinal preferences, optimal social welfare. Third, standard AMD methods explicitly compute the full mapping from preference reports to outcomes, which is intractable for even small roommates problems.

I address the first challenge by implementing a rank-preserving transformation from ordinal to cardinal preferences, after ordinal preferences are received, and before the AMD approach is applied. To address the second challenge, I propose relaxing incentive compatibility along two dimensions: a) restrictions on the set of *salient* deviations, and b) approximation, which bounds the most one can gain from lying. Specifically, I consider three restricted forms of incentive compatibility: permutation IC (where manipulations are permutations of the preference ranking, and exclude truncations), promotion IC (which allows promotions of prospective roommates), and promotion-one IC (which only considers promotions to the top position in the preference order). DA is known not to even be promotion IC, but I show, surprisingly, that it is promotion-one IC.

To address the final challenge, I propose several approaches to construct integer linear programs for computing outcomes for a *specific* set of preference reports, which extends social welfare maximization by introducing constraints which aim to achieve approximate restricted IC. For one of these approaches, I am able to show that the solution, in fact,

guarantees a bound on incentives to lie in the restricted space of manipulations. Two others, however, lack such guarantees, but show superior performance in the experimental evaluation.

VII.1 The Roommates Model

The roommates problem, a generalization of two-sided matching [46], involves a set of players $N = \{1, 2, \ldots, n\}$ who are to be grouped into a collection of teams of at most two each. Each $i \in N$ has a set of *feasible* partners $R_i \subseteq N$, and a preference ranking $\succ_i$ (i.e., a complete, anti-symmetric, and transitive relation) over R_i. We assume player i would rather be a singleton than be with a player in $N \backslash (R_i \cup \{i\})$. We say that a player j has a higher *rank* than player k in player i's preference *if* $j, k \in R_i$ and $j \succ_i k$. Preferences are assumed to be strict, and $x \succeq_i y$ means that either $x \succ_i y$ or $x = y$.

A *profile* of preferences $\succ$ (or simply *profile*) is a list of preferences for every $i \in N$. Given a profile, the list of preferences for all players except i is denoted by $\succ_{-i}$. A roommates matching π is a function $\pi : N \to N$, such that $\pi(\pi(i)) = i$ for any player i, and $\pi(i) \neq \pi(j)$ if $i \neq j$. We also assume that $\pi(i) = i$ if player i is a singleton in the matching. The well-known two-sided matching (marriage) problem is a special case of roommates model in which players are separated into two disjoint sets M and W, s.t. $R_i \subset W \ \forall i \in M$ and $R_j \subset M \ \forall j \in W$. A *roommates mechanism* $\mathscr{M}$ maps every preference profile $\succ$ to a roommates matching π, i.e. $\pi = \mathscr{M}(\succ)$.[1] We denote the roommate of player i generated by the mechanism $\mathscr{M}$ by $\mathscr{M}(\succ, i)$. Finally, we let $\succ_i^j$ denote a preference ranking of i modified by promoting another player j to be most preferred by i.

Throughout, we make use of both ordinal and cardinal preference notions. We leverage ordinal preferences for two reasons: a) such preferences are easier to express in the context of matching problems, and b) it restricts inputs to the mechanism itself. The mechanism transforms ordinal preferences into cardinal in order to (a) consider social welfare as an

[1] To simplify notation, we also let $\succ_i$ denote the preference over several matching outcomes, i.e. $\pi \succ_i \pi'$ denotes $\pi(i) \succ_i \pi'(i)$.

objective, which is a much stronger notion of efficiency than, say, Pareto optimality, and (b) consider quantitative relaxations of incentive compatibility to allow us to consider the problem in the *automated mechanism design* framework; we discuss this in greater detail below. We define $u_i(\pi(i))$ as the cardinal utility of the partner assigned to player i in the matching π. If the partner of i is j, we use u_{ij} to denote $u_i(j)$.

VII.1.1 Incentive Compatibility

In mechanism design, a crucial criterion is *incentive compatibility*, which aims to eliminate any incentives for players to misreport their true preferences (in our case, over partners). Formally, we say that a mechanism $\mathscr{M}$ is *ex-post incentive compatible (IC)* if $\forall \succ, i, \succ_i'$, $\mathscr{M}((\succ_i, \succ_{-i}), i) \succeq_i \mathscr{M}((\succ_i', \succ_{-i}), i)$. When preferences are cardinal transformations of $\succ_i$, we can define an *approximate* IC in cardinal form, for a given additive approximation ε, as

$$\forall \succ, i, \succ_i' \colon\ u_i(\mathscr{M}((\succ_i, \succ_{-i}), i)) \geq u_i(\mathscr{M}((\succ_i', \succ_{-i}), i)) - \varepsilon.$$

VII.1.2 Individual Rationality

Another important criterion in mechanism design is individual rationality. In the roommates problem, this can be represented by constraints $\forall \succ, i, \mathscr{M}(\succ, i) \succeq_i \{i\}$, which, for cardinal preferences, becomes $u_i(\mathscr{M}(\succ, i)) \geq u_i(\{i\})$

VII.1.3 Social Welfare

Social welfare is a notion of efficiency defined for cardinal preferences. Formally, the social welfare of a matching π is $sw(\pi) = \sum_i u_i(\pi)$. Our goal below will be to maximize social welfare for a cardinal transformation of ordinal preferences, subject to approximate incentive compatibility (in a restricted form, discussed below). Formally, we aim to solve

the following problem:

$$\max_{\mathscr{M}} \quad \sum_{\succ \in \mathscr{D}} \sum_{i \in N} [(1-\alpha) u_i(\mathscr{M}(\succ)) - \alpha |N| \varepsilon] \tag{VII.1a}$$

s.t.

$\forall \succ, i, \succ_i'$:

$$u_i(\mathscr{M}((\succ_i, \succ_{-i}), i)) \geq u_i(\mathscr{M}((\succ_i', \succ_{-i}), i)) - \varepsilon. \tag{VII.1b}$$

It trades off relaxations of incentive compatibility and optimizing social welfare, with α the associated tradeoff parameter. The key technical challenge is that this problem requires us to optimize over all possible mechanisms $\mathscr{M}$, and even *representing a given mechanism* is intractable (since it maps all possible preference profiles to matchings). Instead, we would like to leverage the structure of the problem in computing a matching for a *given* preference profile. As we show presently, we can indeed accomplish this. We do so by appealing to relaxed notions of IC, which we turn to next.

VII.2 Restricted Incentive Compatibility

It is well-known that when payments are not allowed and/or utility is not transferable, incentive compatibility is, in general, in conflict with social welfare optimality [141]. One classical approach to address this tension is to additively relax incentive compatibility, requiring instead that no player can gain from lying about preferences more than a small amount, ε [142]. This relaxation is justified by suggesting that agents typically face implicit costs (actual or cognitive) from lying, or gaming the system.

However, traditional concepts which focus on small gains from lying do not account for another barrier to manipulation: complexity. A common approach in this vein is to consider the computational hardness of manipulation [143]. However, when manipulators are *human*, computational complexity may not be appropriate. We propose several alternative notions of cognitive salience in considering the space of possible manipulations. The

central idea is to *restrict the space of feasible manipulations* to those which are cognitively natural. We present three classes of such restrictions: *promotion-one IC*, where salient manipulations involve promoting a player to the top position; *promotion IC*, in which agents consider *promoting* a prospective roommate in their reported ranking to an arbitrary position; and the far more general *permutation IC*, in which agents may report a permutation of their true preferences.

VII.2.1 Promotion-One Incentive Compatibility

We now present the first restriction of manipulating to allow players to promote *anyone* to the first position in their preference ranking. We term this *promotion-one incentive compatibility (POIC)*.

Definition VII.2.1. *A mechanism $\mathscr{M}$ is* promotion-one incentive compatible (POIC) *for a profile domain $\mathscr{D}$ if for any profile $\succ \in \mathscr{D}$, $i, j \in N$, $\mathscr{M}(\succ) \succeq_i \mathscr{M}(\succ_i^j, \succ_{-i})$.*

Translation to cardinal preferences is direct.

I now present one of our main results, which demonstrates the value of POIC as a restriction of incentive compatibility. In particular, I show that the most common mechanism for two-sided matching, deferred acceptance [46], is POIC. [2]

Theorem VII.2.1. *Deferred acceptance mechanism is promotion-one incentive compatible.*

Proof. Without loss of generality, we consider the women-proposing deferred acceptance mechanism (DA for short). We have known that DA is incentive compatible for the proposing side (i.e. women), so we will show that it is also promotion-one incentive compatible for the men.

The proof is by contradiction. Suppose that a man m matches with a woman w under the true profile $\succ$, and matches with a more preferred woman w' (i.e. $w' \succ_m w$) when reporting

[2] [37] shows a similar result.

some promotion-one manipulating preference $\succ_m^{w^*}$, where the promoted woman is w^*. As $DA(\succ) \neq DA(\succ_m^{w^*}, \succ_{-m})$, there must be a round of DA that differs when DA is applied to $\succ$ and $(\succ_m^{w^*}, \succ_{-m})$. Let r^* be the first such round.

Let S_m denote the set of women held by m, which includes m's current mate and unmatched women that propose to m in the round r^*. As r^* is the first round in which the outcome of DA is different in the two profiles, we know that the set S_m is the same at the beginning of round r^* under both profiles. Because m is the only player whose preferences are different in the two profiles, m must be the man whose mate is different in the two profiles at the end of round r^*. Assume $\hat{w}$ is the mate of m under the true profile $\succ$ in the round r^*, based on the property of DA , $\hat{w} \succ_m \tilde{w}$, for all $\tilde{w} \in S_m$.

Also, because w^* is the only woman whose ranking changes in $\succ_m$ and $\succ_m^*$, set S_m must contain woman w^*. (Otherwise, m would match with the same players at the end of round r^* under both profiles, contradicting the fact that DA is different in the round r^*.) Furthermore, w^* must be the woman that m matches at the end of round r^*, and $w^* \neq \hat{w}$. (Otherwise, because S_m is the same under both profiles and w^* is the only woman the ranking of which changes in the preference of m, m would again match the same players at the end of round r^* under both profiles.) Then we could know that $\hat{w} \succ_m w^*$.

As w^* has been promoted to the first place in m's preference $\succ_m^*$, m cannot expect to get a better mate than w^*. At the same time, the receiving side players can only improve (or remain the same) with respect to their reported preferences as the rounds of DA progress. So $DA_m(\succ_m^{w^*}, \succ_{-m}) = w^*$, and $DA_m(\succ) \succeq_m \hat{w}$. It implies $DA_m(\succ) \succeq_m \hat{w} \succ_m w^* = DA_m(\succ_m^{w^*}, \succ_{-m})$, which contradicts with the assumption that m matches with a better player by reporting preference $\succ_m^{w^*}$. □

This result may help explain the success this mechanism has had in practice, with rather little concern about preference manipulation: it appears that it is incentive compatible under a *highly salient* set of preference manipulations.

VII.2.2 Promotion and Permutation Incentive Compatibility

POIC is still a rather restrictive set of manipulations, and it is natural to consider further generalizations. The first generalization is *promotion incentive compatibility*, where promotion can be to an arbitrary position in the preference ranking. To formally define it, let $\succ_i^{j\to l}$ denote a manipulation of an original preference ranking of i, $\succ_i$, in which j is promoted to a position l. Let $p(\succ_i, j)$ be the position of j in i's original preference ranking. Suppose that position $l < k$ in a preference order $\succ_i$ means that a player in position l is more preferred than one in position k.

Definition VII.2.2. *A mechanism $\mathscr{M}$ is* promotion incentive compatible (PIC) *for a profile domain $\mathscr{D}$ if for any profile $\succ \in \mathscr{D}$, for every player i, prospective partner j, and position $l < p(\succ_i, j)$, $\mathscr{M}(\succ) \succeq_i \mathscr{M}(\succ_i^{j\to l}, \succ_{-i})$.*

The final relaxation of IC we consider is *permutation incentive compatibility*.

Definition VII.2.3. *A mechanism $\mathscr{M}$ is* permutation incentive compatible (Permutation IC) *for a profile domain $\mathscr{D}$ if for any profile $\succ \in \mathscr{D}$, for every player i, $\mathscr{M}(\succ) \succeq_i \mathscr{M}(\succ_i', \succ_{-i})$, in which $\succ_i'$ is any permutation of $\succ_i$.*

Again, we can translate these two definitions to cardinal preferences directly. Observe that if players are required to submit full preference orders among feasible roommates, permutation IC is equivalent to general IC; consequently, this is a very general notion of incentive compatibility.

Perhaps surprisingly, Example 2 in [114] shows that deferred acceptance mechanism is not even *promotion IC*. Thus, our positive result above is tight.

One of the major successes of the deferred acceptance has been the belief that it doesn't incentivize manipulation *in practice* [144]. Showing that it is IC for the restricted subset of promotion-one manipulations, coupled with this anecdotal claim, offers evidence that promotion-one manipulations may be the most *salient* manipulations in matching settings.

Consequently, our experimental evaluation below focuses on promotion-one incentives to measure benefits from misreporting preferences.

VII.3 Automated Mechanism Design for Roommates Problem

I now proceed to leverage the (restricted) notions of incentive compatibility in devising an automated mechanism design (AMD) approach [50] for the roommates problem. First, I always treat individual rationality as a hard constraint. Second, I aim to maximize social welfare subject to constraints involving approximate notions of IC.

In this section, I describe several automated mechanism design approaches for the roommates problem. I present integer linear programming methods for this problem that trade off social welfare and an upper bound of ε in a manner similar to Problem (VII.1), but *for a specific instance of reported preferences*; consequently, the mechanism is implicitly specified through solutions of such programs for all possible preference profiles.

Throughout, I consider the transformation from ordinal preference to the cardinal utility. A transformation is a *mapping* from an ordinal profile $\succ$ to a set of utility functions $\{u_i(\cdot)\}$, s.t. $i \in N$. Specifically, $u_i(j)$ means the utility of player i for matching with player j. For any $j, j' \in R_i$, if $j \succ_i j'$, then $u_i(j) \geq u_i(j')$, or, in our notation introduced above, $u_{ij} \geq u_{ij'}$.

VII.3.1 AMD That Maximizes Social Welfare

As a benchmark, I first present the program that could maximize the social welfare of players with only individual rationality constraints. This also provides the main building block for the AMD approaches that follow.

First, I introduce integer variables $x_{ij} \in \{0,1\}$ which represent matching a player i with j. Since the "roommates" matching is symmetric (if i is j's roommate, then j is i's roommate), $x_{ij} = x_{ji}$ for all i, j. Moreover, since a player can have at most one roommate, and can be a roommate of at most one player I have in addition the constraints $\sum_i x_{ij} \leq 1$ and $\sum_j x_{ij} \leq 1$ for all i, j. Finally, the objective of social welfare becomes $\sum_{i,j} x_{ij} u_{ij}$. The

full mechanism then becomes Program (VII.2).

$$P_{MSW} = \max_{x} \quad \sum_{i \in N} \sum_{j \in N} x_{ij} u_{ij} \tag{VII.2a}$$

$$\text{s.t.}: \quad \sum_{j \in N} x_{ij} \leq 1, \quad \forall i \in N \tag{VII.2b}$$

$$\sum_{i \in N} x_{ij} \leq 1, \quad \forall j \in N \tag{VII.2c}$$

$$x_{ij} = x_{ji}, \quad \forall i, j \in N \tag{VII.2d}$$

$$IR: \quad \sum_{j \in N} x_{ij} u_{ij} \geq u_{ii}, \quad \forall i \in N \tag{VII.2e}$$

$$x_{ij} \in \{0, 1\}, \quad \forall i, j \in N \tag{VII.2f}$$

VII.3.2 AMD with Approximate Permutation Incentive Compatibility

Next, I consider the problem of maximizing social welfare with restricted forms of incentive compatibility captured by a collection of constraints. As I had remarked earlier, the central challenge in doing so is that, in general, incentive compatibility must consider *mechanism outcomes* for alternative manipulation, something that can be difficult to capture without explicitly defining a general roommates mechanism (a clearly intractable proposition). In this subsection, I demonstrate that the special structure of the roommates problem allows us to overcome this challenge.

As our principle approach, I present an integer linear program that trades off social welfare and an upper bound of ε. And the resulting mechanism is ε-permutation incentive compatible. The key idea for developing this approach is to consider each possible manipulation by a player i. The worst case, from the mechanism designer's standpoint, is that the manipulation succeeds, and i matches with a better roommate than her current mate. I can very conservatively guard against *all* such possible deviations by simply introducing a collection of constraints that each player i obtains at least u_{ij} for any possible partner j they may have. This yields our first integer program for automated mechanism design,

presented in Program (VII.3). In the program, α is set to trade off between social welfare and incentive compatibility and $0 \leq \alpha \leq 1$.

$$P_{IC} = \max_{x,\varepsilon} \quad (1-\alpha)\sum_{i,j} x_{ij}u_{ij} - \alpha|N|\varepsilon, \tag{VII.3a}$$

$$\text{s.t. : constraints (VII.2b)} - \text{(VII.2e)}$$

$$\varepsilon\text{-IC} : \quad \sum_{k\in N} x_{ik}u_{ik} \geq u_{ij} - \varepsilon, \forall i \in N, j \in R_i \tag{VII.3b}$$

$$\varepsilon \geq 0 \tag{VII.3c}$$

$$x_{ij} \in \{0,1\}, \quad \forall i,j \in N \tag{VII.3d}$$

The above optimization problem is clearly conservative, as it introduces constraints about prospective roommates *whether or not they can be realized through a unilateral deviation*. I deal with this presently, but for the moment, this provides our first principled approach.

Let us denote by $\mathscr{M}_{IC}(\succ)$ the mechanism (implicitly) implemented by the integer program (VII.3). As we now observe, the optimal solution to this program yields an upper bound on the most any player can gain from arbitrary permutations of their preferences—that is, the result is approximately permutation IC (and, consequently, promotion and promotion-one IC).

Theorem VII.3.1. *Assume the optimal solution of Program VII.3 is (x^*, ε^*), then the mechanism $\mathscr{M}_{IC}$ is ε^*-Permutation IC, i.e. no player can gain more than ε^* via permutation of her true preference.*

Proof. Due to the IR constraints (VII.2e), player i never team up with a player that is not in R_i. When player i untruthfully report her preference by permutation, she will still match with one player in R_i (or be a singleton). The constraint (VII.3b) can make sure that player i cannot gain more than ε^* by matching with any player in R_i. □

While this program appears simplistic, it is a useful step as enables us to parametrically

trade off social welfare and incentives to lie by tuning the parameter α. As our experiments demonstrate, this yields a non-trivial tradeoff with respect to the highly salient promotion-one-IC deviations.

We note that many of the constraints in the above integer program are unnecessary, and, indeed, significantly over-constrain the problem. In the program, all players in the feasible set R_i are treated as "potential" teammates when i is trying to manipulate the mechanism, and some players in R_i are far less preferred than the final teammates. Consequently, these constraints will never be relevant.

I use this intuition to develop an iterative approach to generating the approximate permutation IC constraints, allowing us to focus only on those which actually matter. We start with a program that maximizes social welfare, and use it to obtain an initial roommates assignment π. Then we add constraints to make sure that player i cannot gain more than ε by matching with any player j s.t. $j \succ_i \pi(i)$. We then solve the program with the newly added constraints (replacing the objective with (VII.3a) and adding the constraint that $\varepsilon \geq 0$), obtain a new assignment, and repeat the process until convergence (which is guaranteed in the quadratic time since the set of possible constraints is quadratic). This approach is shown in Algorithm 10, and the resulting program is, again, ε-permutation IC (details omitted due to space constraints).

Algorithm 10 ε-Permutation IC Program

input: initial program (VII.2)
return: ε-Permutation IC Program

program $P_{IC} \leftarrow$ program (VII.2)
repeat
    solve the program and get the matching assignment π
    for i, and $j \in R_i$ **do**
        if $j \succ_i \pi(i)$ and corresponding constraint has not been added **then**
            add $\sum_{k \in N} x_{ik} u_{ik} \geq u_{ij} - \varepsilon$ into P_{IC}
        end if
    end for
until the value ε converges
return program P_{IC}

VII.3.3 Heuristic Approaches with Promotion-One Manipulations

While the approaches described above are principled in the sense that they yield provable guarantees, even the iterative approach is likely to introduce too many constraints (thereby compromising social welfare which could have been achieved). One major reason for this is that it still accounts for the full space of permutation manipulations, rather than the more salient restricted space of manipulations in which a player only promotes another to the top position in her order. We now introduce two iterative heuristic approaches which directly consider this smaller set of manipulations, albeit losing theoretical guarantees.

The intuition behind our first heuristic approach is to iteratively allow each player to promote another to the top position, check if the result yields *strictly higher* social welfare than current allocation, if only in this case add the corresponding constraint. The full approach is shown in Algorithm 11. Here $P_{MSW}(\cdot)$ denotes the assignment which maximizes social welfare (i.e., solves program (VII.2)), and $sw(\cdot)$ denotes the social welfare of an assignment.

Algorithm 11 Heuristic 1

input: initial program (VII.2)
return: Heuristic ε-POIC Program

1: program $P_{POIC} \leftarrow$ program (VII.2)
2: Replacing the objective with (VII.3a)
3: **repeat**
4: **for** i, and $j \in R_i$ **do**
5: **if** $sw(P_{MSW}(\succ_i^j, \succ_{-i})) > sw(P_{POIC}(\succ))$ and corresponding constraint has not been added **then**
6: add $\sum_{k \in N} x_{ik} u_{ik} \geq u_{ij} - \varepsilon$ into P_{POIC}
7: **end if**
8: **end for**
9: **until** the value ε converges
10: return program P_{POIC}

Our second heuristic approach even further relaxes the IC constraints. In the process of constraint generation, we still compute the social welfare of the manipulated profile. However, unlike from Algorithm (11), we let player i match with j, compute the social

welfare of the remaining players in the current program, and check if the utility of i and j, along with social welfare of the match among remaining players, is thereby increased; if it is, we add the constraint, since this is likely the salient manipulation. The full approach for the second heuristic is described in Algorithm 12. Here U_i is the utility i receives from ranking a player (j in this case) in the first position, and $\succ_{-\{i,j\}}$ is the profile of all players other than i and j.

Algorithm 12 Heuristic 2

input: initial program (VII.2)
return: Heuristic ε-POIC Program

1: program $P_{POIC} \leftarrow$ program (VII.2)
2: **for** i, and $j \in R_i$ **do**
3: **if** $U_i + u_{ji} + sw(P_{POIC}(\succ_{-\{i,j\}})) > sw(P_{POIC}(\succ))$ **then**
4: add $\sum_{k \in N} x_{ik} u_{ik} \geq u_{ij} - \varepsilon$ into P_{POIC}
5: **end if**
6: **end for**
7: return program P_{POIC}

VII.4 Experiments

I evaluate the proposed approaches (Algorithm 1–3) on preference profiles generated by social networks with respect to the social welfare and (approximate) promotion-one incentive compatibility (POIC). In the three approaches introduced, $\alpha = 0$ is equivalent to Program (VII.2), which maximizes social welfare; this will be the baseline for our approaches. I consider two measures to capture incentives to lie: 1) maximum benefit from deviation to any player (a relatively standard metric), and 2) proportion of players who benefit from deviation (used for matching settings [144]). I adopt Erdös-Rényi (ER) [145] and Barabási-Albert (BA) [137] models to generate *random networks*. For ER, the probability of an edge is $p = 0.2$. For BA, each new node is connected to 2 nodes when generating the network (see [137] for more details). For each network, I generate preference profiles among players as follows. We assume that a player i is only willing to match with neighbors, i.e. $j \in R_i$ if and only if i and j are neighbors, and rank i's neighbors in random

order. To map ordinal preference to the cardinal, we use *Borda scoring functions* suggested by [140]. Specifically, I adopt the *normalized* Borda scoring function, defined as $u_i(j) = g(r) = (k - r + 1)/k$ where k is the number of feasible roommates, and $r \in [1 \dots k]$ is the rank of j in i's preference. Thus, utilities are in the $[0, 1]$ interval.

For all the results, we take the average over 100 samples, and consider games with 20 players. I solve mixed integer linear programs using CPLEX 12.6.1.

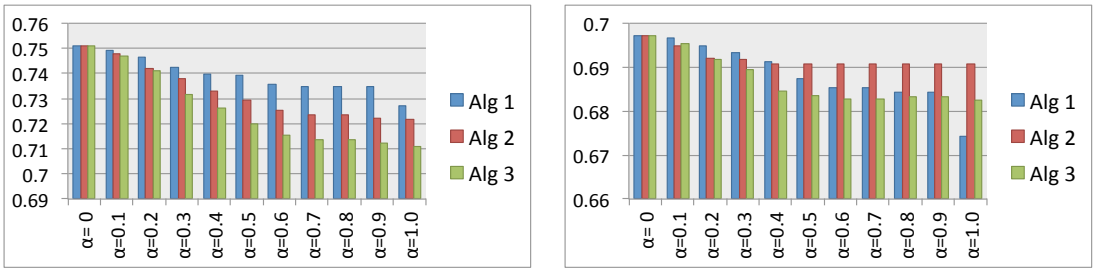

Figure VII.1: Social welfare on ER (left) and BA (right) networks.

Figure VII.1 shows the (normalized) social welfare on ER and BA models, respectively. As expected, social welfare decreases with α (the weight of the incentive term). Our key observation here, however, is that our three algorithmic approaches (Algorithm 1-3) yield similar social welfare when $\alpha \leq 0.3$ (within 0.01).

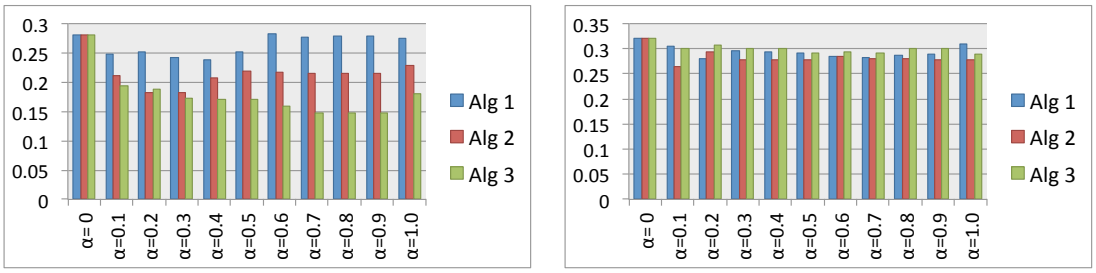

Figure VII.2: Maximum benefit from deviation on ER (left) and BA (right) networks.

Figures VII.2 and VII.3 present the results regarding incentives to lie. It is perhaps surprising that these are not monotonically decreasing with α, but note that our approaches minimize an upper bound on the gain from lying, which can lead to over-constrained programs and be actually counterproductive. The sweet spot appears to be $\alpha = 0.2$. For ER

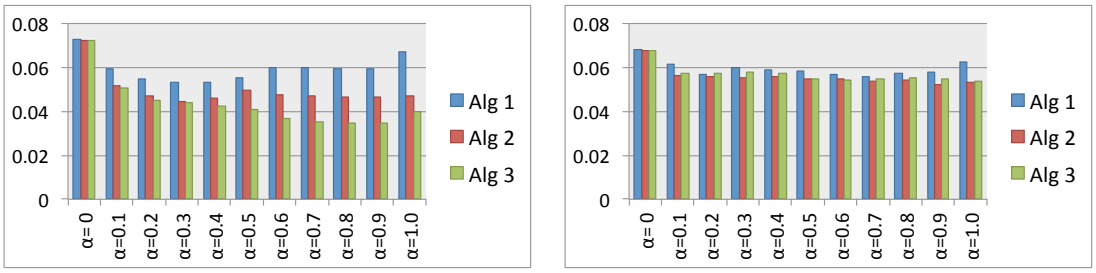

Figure VII.3: Proportion of players who can benefit from deviation on ER (left) and BA (right) networks.

networks, Algorithm 2 and 3 are both considerably better than Algorithm 1 in terms of incentives to lie. However, the difference is not that significant for BA networks, even though we can also see the superiority of Algorithm 2 and Algorithm 3. Moreover, there doesn't appear to be an appreciable difference between Algorithms 2 and 3.

VII.5 Conclusion

I present the first treatment of the roommates problem from an automated mechanism design (AMD). The proposed approaches and analysis discussed through the chapter, are mainly based on two restricted incentive compatibility: promotion-one incentive compatibility, in which the manipulation consists of promoting a single player to the top position in the manipulator's preference list; and permutation incentive compatibility, in which the manipulation is a permutation of the preference. Specifically, I prove that the well-known *deferred acceptance* mechanism in two-sided matching is promotion-one IC, which in some sense explains the success of this mechanism and also shows the usefulness of the restriction. Finally, I propose several automated mechanism design approaches to the roommates problem. I prove that the first approach is ε-permutation IC. While the other two don't have theoretical guarantees, empirical results show that they have even better incentive properties.

Part III

Toward Efficiency: Secondary Market

CHAPTER VIII

Secondary Market Mitigates Demand Uncertainty

In the chapter, I leverage game-theoretic analysis to study the influence of a *secondary market* on players' decisions and social welfare in the resource allocation problem introduced in the introduction. I use a well-known newsvendor model as a starting point to construct a two-stage game model:

1. In the first stage (i.e. primary market), players report their orders to the authority, and then pay and get the requested resources. In this stage, actual demand is uncertain.

2. In the second stage (i.e. secondary market), demand uncertainty is resolved, all demanded resources used, and remaining resources can be traded freely.

I focus on the following two representative settings, one for a large market and another for a small one:

- **Large markets:** In the large market model, I analyze the influence of a secondary market asymptotically as the number of players n approaches infinity. In this setting, our solution concept for the secondary market is a competitive equilibrium, with backward induction allowing us to characterize the (asymptotic) Nash equilibrium strategies in the first stage. I find that an individual's influence on the price in the secondary market is bounded by $O\left(\frac{1}{\min(\sqrt{n},\pi(n))}\right)$, where $\pi(n) \leq n$ and $\lim_{n\to+\infty} \pi(n) = +\infty$, which validates the price-taking assumption of the competitive equilibrium solution concept. Moreover, I show that the secondary market is able to significantly mitigate the uncertainty of demand by aligning with the price in the primary market. Furthermore, for both social welfare and total orders in the primary market, the difference between the optimal and equilibrium outcomes is bounded by $O(\max(\sqrt{n}, \frac{n}{\pi(n)}))$, a significant improvement over the $\Theta(n)$ difference between the optimum and the outcome with only the newsvendor model (i.e., without the secondary market).

- **Small markets (2 players):** In the small market model, I assume players are *bargaining* on the price in the secondary market, and adopt the *Nash bargaining* solution concept [146]. In this model, I provide a sufficient condition that guarantees the existence of pure strategy Nash equilibrium and subsequently focus on characterizing symmetric equilibrium outcomes when the game is symmetric. Under this equilibrium, I prove that social welfare is always no worse than in the newsvendor model (i.e., without the secondary market).

In summary, I conclude that secondary market is indeed able to mitigate demand uncertainty and improve social welfare.

The rest of this chapter organizes as follows. I formalize our model in Section VIII.1. The results for a large market are presented in Section VIII.2, followed by the results for a small market in Section VIII.3. This chapter is concluded with discussion and future work in Section VIII.4.

VIII.1 Model

In this section, I go through our market model with and without a secondary market, and formalize the setup of the induced game.

VIII.1.1 Background: The Newsvendor Model

My point of departure is the standard *newsvendor model* [147]. The focus of this model is an agent who decides on the quantity of homogeneous and divisible resources to order to satisfy uncertain future demand. To formalize, let c be the unit cost of the resource, $x \geq 0$ the number of units of the resource ordered, and Q the random variable representing the uncertain demand, distributed according to a continuous cumulative distribution function $F_Q(\cdot)$. Finally, let c' be the marginal penalty of resource shortage. The expected utility of the agent is then modeled by

$$\mathbf{U}(x) = \mathbb{E}_Q\left[-cx - c'(Q-x)^+\right], \tag{VIII.1}$$

where $(Q-x)^+ = \max(Q-x, 0)$. Thus, the agent trades off the immediate incurred cost of ordering x units, and the uncertain future deleterious consequences of failing to fulfill demand.

Based on Littlewood's rule [148], the optimal solution to the newsvendor problem is characterized by

$$x^* = F_Q^{-1}\left(1 - \frac{c}{c'}\right) \tag{VIII.2}$$

where $F_Q^{-1}(\cdot)$ is the inverse cumulative distribution function of Q. An equivalent characterization is

$$\mathbb{P}(Q \geq x^*) = \frac{c}{c'}, \tag{VIII.3}$$

where $\frac{c}{c'}$ is known as the *critical fractile*.

VIII.1.2 The Newsvendor Model with a Secondary Market

I consider an extension of the standard newsvendor model in which there exists a secondary market at which players can trade excess supply of the resource *after demand uncertainty has been resolved*.

To formalize, let $N = \{1, 2, \ldots, n\}$ be the set of players. For each player $i \in N$, the marginal penalty of resource shortage is c_i', whereas the unit cost of procuring the resource is c for all players. In addition, each player $i \in N$ has a stochastic demand Q_i distributed according to $F_{Q_i}(\cdot)$, and we assume that $0 \leq Q_i \leq q_{max}$, where q_{max} is the maximum possible demand. To simplify notation, let $\mathbf{Q} = \{Q_1, Q_2, \ldots, Q_n\}$. Let x_i denote the number of units of the resource i orders, and assume that this is bounded by x_{max}.

So far, the decisions of all players are entirely independent. We now suppose that there is a secondary market in which the players can trade resources after their individual demands Q_i have been realized and the ordered resources have been consumed. This ability to trade in the future then connects the initial decisions by the players about how many resources to purchase. The consequence is the following two-stage game:

- In the first stage (*primary market*), all players' demands are *unknown* (but their distributions are assumed to be common knowledge) and each player $i \in N$ orders x_i units of the resource at unit price c.

- In the second stage (*secondary market*), the demand of each player $i \in N$ is realized and order resources are consumed. We denote these realizations of demand by q_i, and let $\mathbf{q} = \{q_1, q_2, \ldots, q_n\}$. We assume that these realizations become common knowledge. At this point, the resource can be exchanged on the secondary market, where sellers are players i with $x_i - q_i > 0$ (excess supply) and buyers are players i with $x_i - q_i < 0$ (excess demand). The secondary market then determines the price p for exchanging the resource among players; I discuss the specific solution concepts used to characterize the market price below.

The price in the secondary market is influenced by the excess demand and supply in the market, which are jointly determined by $\mathbf{x}$ and $\mathbf{Q}$. Let $P(\mathbf{Q},\mathbf{x})$ denote the price in the secondary market. For each player i, the amount of the resource purchased or sold in the secondary market is then jointly determined by $\mathbf{x}$, $\mathbf{Q}$ and $P(\mathbf{Q},\mathbf{x})$. Let $I_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x}))$ and $O_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x}))$ denote the amount of the resource purchased and sold by player i in the secondary market, respectively (note that at most one of these is non-zero). Then the (expected) utility of player i ($i \in N$) *in the full game*, as a function of the joint resource ordering decisions $x = (x_i, x_{-i})$ in the first stage is defined as

$$\begin{aligned}\mathbf{U}_i(x_i,x_{-i}) =& \mathbb{E}_{\mathbf{Q}}\Big\{-cx_i + P(\mathbf{Q},\mathbf{x})O_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x})) - P(\mathbf{Q},\mathbf{x})I_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x})) \\ & - c_i'\Big[Q_i - x_i - I_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x})) + O_i(\mathbf{Q},\mathbf{x},P(\mathbf{Q},\mathbf{x}))\Big]^+\Big\}\end{aligned} \tag{VIII.4}$$

In the thesis I study the pure strategy Nash equilibrium of the game in the first stage, with behavior in the secondary market characterized using an associated solution concept that I deal with shortly (observe that the game can be treated as complete information, as is common in related literature [129], since there is no information asymmetry among the

players).

Definition VIII.1.1. *$(x_1, x_2, \ldots, x_n)$ is a pure strategy Nash equilibrium if for any player i, given x_{-i}, we have $\mathbf{U}_i(x_i, x_{-i}) \geq \mathbf{U}_i(x_i', x_{-i})$ for any $0 \leq x_i' \leq x_{max}$.*

VIII.2 Large Markets: Asymptotic Analysis

I begin our study with a limiting case where the number of players n approaches to infinity. Let $\pi(\cdot)$ be a function of n such that $\pi(n) \leq n$ for any n and $\pi(n) \to +\infty$ as $n \to +\infty$. The assumptions can be formalize as follows,

Assumption VIII.2.1. *Assume that the number of players $n \to +\infty$. Furthermore, we assume that the domain of marginal penalty of resource shortage c' is $[c'_{min}, c'_{max}]$, in which c'_{min} and c'_{max} are the minimum and maximum marginal penalty of resource shortage, respectively, with $c'_{min} \geq c$. Finally, for any given range $[\underline{b}, \overline{b}]$ with $c'_{min} \leq \underline{b} < \overline{b} \leq c'_{max}$, we have $\lim_{n \to +\infty} \sum_{i \in N} \mathbb{1}(c'_i \in [\underline{b}, \overline{b}]) = \Omega(\pi(n))$, where $\pi(n) \leq n$ and $\lim_{n \to +\infty} \pi(n) = +\infty$.*

Intuitively, we consider the case where players are distributed based on their marginal penalty of resource shortage, and the distribution does not exhibit significant skew in the valid range. This assumption will be important when we study the beneficial impact that the secondary market has on social welfare.

Under the asymptotic assumption, we focus on the ε-Nash equilibrium (where $\varepsilon \to 0$ as $n \to +\infty$) of the game in the first stage. We term it as "asymptotic Nash equilibrium".

Definition VIII.2.1. *$(x_1, x_2, \ldots, x_n)$ is an asymptotic Nash equilibrium if for any player i, given x_{-i}, we have $\mathbf{U}_i(x_i, x_{-i}) \geq \mathbf{U}_i(x_i', x_{-i}) - \varepsilon$ for any $0 \leq x_i' \leq x_{max}$, where $\varepsilon \to 0$ as $n \to +\infty$.*

VIII.2.1 Market Clearing Price in Secondary Market

In the limiting case when the number of players approaches infinity, it is natural to consider the competitive equilibrium as the solution concept for the secondary market, in which the price $P(\mathbf{Q}, \mathbf{x})$ is the *market clearing* price.

Let $N_+ = \{i \in N \,|\, x_i - q_i > 0\}$ be the set of players with excess supply, $N_- = \{i \in N \,|\, x_i - q_i < 0\}$ the set of players with excess demand. We define competitive equilibrium as follows,

Definition VIII.2.2. *Given the orders* $\mathbf{x}$ *(in the primary market) and the realized demands* $\mathbf{q}$ *(in the secondary market), the price* $P(\mathbf{q},\mathbf{x})$ *and* $\{O_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})),\, I_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x}))\}_{i\in N}$ *form a competitive equilibrium (CE) of the secondary market if the following conditions hold:*

- *For all players* $i \in N_-$ *with* $c_i' > P(\mathbf{q},\mathbf{x})$, $0 \le I_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) \le q_i - x_i$, *and* $O_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) = 0$. *That is to say, these players are willing to buy resources at price* $P(\mathbf{q},\mathbf{x})$.
- *For all players* $i \in N_+$, $I_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) = 0$, *and* $0 \le O_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) \le x_i - q_i$. *That is to say, these players are willing to sell redundant resources at price* $P(\mathbf{q},\mathbf{x})$.
- *For all players* $i \in N_-$ *with* $c_i' < P(\mathbf{q},\mathbf{x})$, $I_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) = O_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) = 0$. *That is to say, these players are not willing to sell or buy resources.*
- **Market Clears:** $\sum_{i\in N} I_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x})) = \sum_{i\in N} O_i(\mathbf{q},\mathbf{x},P(\mathbf{q},\mathbf{x}))$.

We observe that the competitive equilibrium may be not unique, and we define the *maximum* market clearing price under a competitive equilibrium in the following.

Let $z_+ = \sum_{i\in N_+}(x_i - q_i)$ be the aggregate supply, and $z_- = \sum_{i\in N_-}(q_i - x_i)$ be the aggregate demand (in the secondary market). Recall that $P(\mathbf{q},\mathbf{x}) \ge 0$ is the market clearing price given the strategy profile $\mathbf{x}$ and realized demands $\mathbf{q}$. We divide the relation between supply and demand in the secondary market into two cases:

1. $z_+ \ge z_-$, i.e. supply exceeds demand. Sellers always have an incentive to decrease prices when $P(\mathbf{q},\mathbf{x}) > 0$, which means that in this case, $P(\mathbf{q},\mathbf{x}) = 0$.

2. $z_+ < z_-$, i.e. demand exceeds supply. Given a price $p = P(\mathbf{q}, \mathbf{x})$, players in $\{i \in N_- \,|c_i' \geq p\}$ will purchase the resource on the market, while those with $c_i' < p$ will not.

We now characterize the market clearing price in case 2 above. First, let us sort players $i \in N_-$ in increasing order of c_i'; let $c'_{(j)}$ denote the jth marginal penalty of resource shortage in this order. Let j^* be the smallest j such that

$$\sum_{i:\{i \in N_- \,|c_i' \geq c'_{(j)}\}} (q_i - x_i) = \bar{z} \leq z_+. \tag{VIII.5}$$

In other words, the total demand by players, including (j^*)th, who purchase the resource with the market price $\bar{z}$ is no larger than the supply z_+. Next, observe that if $\bar{z} = z_+$, then $c'_{(j^*)}$ is the market clearing price, by definition. Otherwise (i.e., if $\bar{z} < z_+$), if we reduce the price to $c'_{(j^*-1)}$ (noting that $j > 1$ since otherwise we are in case 1), we can clear the market by having $(j^* - 1)$th player purchase the rest of the available resource (since this player is indifferent between purchasing and not). Consequently, in case 2, a market clearing price is

$$P(\mathbf{q}, \mathbf{x}) = \begin{cases} c'_{(j^*)} & \text{if } \bar{z} = z_+ \\ c'_{(j^*-1)} & \text{o.w.} \end{cases} \tag{VIII.6}$$

The choice of the competitive equilibrium solution concept and market clearing price are subsequently justified when we show that an individual player' influence on the market price indeed vanishes, and price-taking is indeed a Nash equilibrium in the limit that $n \to +\infty$.

We now could get players' (expected) utility function as follows,

- When $P(\mathbf{Q}, \mathbf{x}) = 0$, $\mathbf{U}_i(x_i, x_{-i}) = -cx_i$ for all player i.
- When $P(\mathbf{Q}, \mathbf{x}) > 0$,
  - player $i \in N_+$(i.e. sellers) could sell all her redundant resources in the secondary

market at price $P(\mathbf{Q},\mathbf{x})$, i.e. $\mathbf{U}_i(x_i,x_{-i}) = \mathbb{E}_{\mathbf{Q}}\left[-cx_i + (x_i - Q_i)P(\mathbf{Q},\mathbf{x})\right]$;

- player $i \in N_-$ s.t. $c_i' > P(\mathbf{Q},\mathbf{x})$ would buy as many of the resources as she needs, i.e. $\mathbf{U}_i(x_i,x_{-i}) = \mathbb{E}_{\mathbf{Q}}\left[-cx_i - (Q_i - x_i)P(\mathbf{Q},\mathbf{x})\right]$;
- player $i \in N_-$ s.t. $c_i' = P(\mathbf{Q},\mathbf{x})$ (including the (j^*-1)th player mentioned above) may buy some resources in the secondary market while bear the cost of resource shortage for the rest, as $c_i' = P(\mathbf{Q},\mathbf{x})$, we could get $\mathbf{U}_i(x_i,x_{-i}) = \mathbb{E}_{\mathbf{Q}}\left[-cx_i - (Q_i - x_i)c_i'\right]$;
- player $i \in N_-$ s.t. $c_i' < P(\mathbf{Q},\mathbf{x})$ does not buy resources in the secondary market, so $\mathbf{U}_i(x_i,x_{-i}) = \mathbb{E}_{\mathbf{Q}}\left[-cx_i - (Q_i - x_i)c_i'\right]$.

To sum them up,

$$\mathbf{U}_i(x_i,x_{-i}) = \mathbb{E}_{\mathbf{Q}}\left[-cx_i + (x_i - Q_i)^+ P(\mathbf{Q},\mathbf{x}) - (Q_i - x_i)^+ \min\left(P(\mathbf{Q},\mathbf{x}), c_i'\right)\right] \qquad \text{(VIII.7)}$$

VIII.2.2 Players' Influence on the Price

It is commonly assumed that players are *price-takers* in a market with perfect competition. Under the price-taking assumption, sellers and buyers could sell or buy resources at the market price without affecting that price. This is in general clearly not the case in our setting. However, in this subsection, we will show that in the asymptotic model, an individual's influence on the expected price in the secondary market is $O\left(\frac{1}{\min(\sqrt{n},\pi(n))}\right)$ as $n \to +\infty$, where $\pi(n)$ measures the density of players' distribution in any finite range of marginal penalty of resource shortage. In the model, we assume that players' demand are independent and identically distributed with finite mean and variance.

Note that $P(\mathbf{Q},\mathbf{x})$ is jointly determined by $\mathbf{Q}$ and $\mathbf{x}$. We will firstly show that the influence of an individual's order in the primary market on the market clearing price in the secondary market is upper bounded by $O\left(\frac{x_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)$ as $n \to +\infty$, where x_{max} is maximum possible amount of resources an individual could order in the primary market. Then we show that the influence of different realization of an individual's demand on the ex-

pected price is $O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)$ as $n \to +\infty$, where q_{max} is the maximum possible realized demand of an individual. The two results together validate the price-taking assumption adopted in the following subsection.

Before introducing the two results, we firstly present a lemma that will be useful in the proof of the two results.

Lemma VIII.2.1. *Let* $\{Y_1, Y_2, \ldots, Y_n\}$ *be a sequence of* independent and identically distributed (i.i.d.) *random variables drawn from a distribution of expected value given by* μ *and finite variance given by* σ^2*, and* $S_n = \sum_{i \in n} Y_i$*. For any sequence of numbers* $\{z_1, z_2, \ldots, z_n\}$ *and a nonnegative number b,* $\mathbb{P}(z_n \leq S_n \leq z_n + b)$ *is bounded by* $O\left(\frac{b}{\sigma\sqrt{n}}\right)$ *as* $n \to \infty$*.*

Proof. Based on central limit theorem, as $n \to \infty$,

$$\frac{1}{n} S_n \sim N\left(\mu, \frac{\sigma^2}{n}\right) \tag{VIII.8}$$

After normalizing,

$$\frac{\frac{1}{n} S_n - \mu}{\frac{\sigma}{\sqrt{n}}} \sim N(0,1) \tag{VIII.9}$$

So

$$\begin{aligned}
&\mathbb{P}(z_n \leq S_n \leq z_n + b) \\
=&\mathbb{P}\left(\frac{\frac{1}{n} z_n - \mu}{\frac{\sigma}{\sqrt{n}}} \leq \frac{\frac{1}{n} S_n - \mu}{\frac{\sigma}{\sqrt{n}}} \leq \frac{\frac{1}{n}(z_n + b) - \mu}{\frac{\sigma}{\sqrt{n}}}\right) \\
=&\mathbb{P}\left(\frac{z_n}{\sigma\sqrt{n}} - \frac{\sqrt{n}\mu}{\sigma} \leq \frac{\frac{1}{n} S_n - \mu}{\frac{\sigma}{\sqrt{n}}} \leq \frac{z_n}{\sigma\sqrt{n}} - \frac{\sqrt{n}\mu}{\sigma} + \frac{b}{\sigma\sqrt{n}}\right) \\
=&O\left(\frac{b}{\sigma\sqrt{n}}\right)
\end{aligned} \tag{VIII.10}$$

□

Theorem VIII.2.1. *Assume that* Q_i $(i \in N)$ *are independent and identically distributed with finite mean* μ *and variance* σ^2*. For a player* $i \in N$*, for any* x' *and* x''*, s.t.* $0 \leq x', x'' \leq x_{max}$*, we could get* $\left|\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}) | (x_i = x')] - \mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}) | (x_i = x'')]\right| = O\left(\frac{x_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)$*, as* $n \to +\infty$*, given* $\{x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n\}$*.*

Proof. Let $z_n = \sum_{j=1}^{n} x_j - x_i$ and we know that z_n has been determined based on the assumption. Firstly, we note that given any realization of players' demands $\mathbf{q}$, the market-clearing price $P(\mathbf{q},\mathbf{x})$ in the secondary market is *monotonically non-increasing* with the increasing of x_i. Take the expectation of the players' demands, the expected price is also monotonically non-increasing with the increasing the x_i. So we obtain the following:

$$\begin{aligned} &\left|\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|_{(x_i=x')}\right] - \mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|_{(x_i=x'')}\right]\right| \\ \leq &\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|_{(x_i=0)}\right] - \mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|_{(x_i=x_{max})}\right] \end{aligned} \tag{VIII.11}$$

Let $p_1 = P(\mathbf{Q},\mathbf{x})|_{(x_i=x_{max})}$, $p_2 = P(\mathbf{Q},\mathbf{x})|_{(x_i=0)}$, and $\Delta p = p_2 - p_1$, and we consider different realizations of $\mathbf{Q}$ as follows,

- When $\sum_i q_i \leq z_n$, no matter how much player i purchased in the primary market, the supply is always over (or equal to) demand, the market-clearing price is always 0, so

$$\Delta p = 0 \tag{VIII.12}$$

- When $\sum_i q_i > z_n + x_{max}$, the demand is always over supplies. When x_i changes from x_{max} to 0, the supply in the secondary market will be decreased by x_{max}, which will push the price increase somehow. Specifically[1],

$$x_{max} - 0 \approx \sum_{i \in N} \mathbb{1}(c_i' \in [p_1, p_2])(q_i - x_i)^+ \tag{VIII.13}$$

  Based on Assumption VIII.2.1, $\sum_{i \in N} \mathbb{1}(c_i' \in [p_1, p_2]) = \Omega(\pi(n))$. As $(q_i - x_i)^+ = \Theta(1)$, so

$$\Delta p = \frac{x_{max}}{\Omega(\pi(n))\Theta(1)} = O\left(\frac{x_{max}}{\pi(n)}\right) \tag{VIII.14}$$

- When $z_n < \sum_i q_i \leq z_n + x_{max}$, the demand can be either over or under the supply. If

[1] The consideration of corner cases is similar as that in Equation VIII.6, and it does not influence the order in asymptotic analysis

$x_i = x_{max}$, the supply is over demand, and the market-clearing price is 0; if $x_i = 0$, the demand is over supply and the price is some value between c'_{min} and c'_{max}. So

$$\Delta p = \Theta(1) \tag{VIII.15}$$

Based on Lemma VIII.2.1, $Pr(z_n < \sum_i q_i \leq z_n + x_{max}) = O\left(\frac{x_{max}}{\sigma\sqrt{n}}\right)$.

Combining the three cases of the realization of $\mathbf{Q}$, the following can be derived:

$$\begin{aligned}
&\left|\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i = x')\right] - \mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i = x'')\right]\right| \\
\leq &\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i = 0)\right] - \mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i = x_{max})\right] \\
= &\mathbb{P}\left(\sum_{i\in N} Q_i \leq z_n\right)\cdot 0 + \mathbb{P}\left(\sum_{i\in N} Q_i > z_n + x_{max}\right)\cdot O\left(\frac{x_{max}}{\pi(n)}\right) \\
&+ \mathbb{P}\left(z_n < \sum_{i\in N} Q_i \leq z_n + x_{max}\right)\cdot \Theta(1) \\
= &\Theta(1)\cdot 0 + \Theta(1)\cdot O\left(\frac{x_{max}}{\pi(n)}\right) + O\left(\frac{x_{max}}{\sigma\sqrt{n}}\right)\cdot\Theta(1) \\
= &O\left(\frac{x_{max}}{\sigma\sqrt{n}} + \frac{x_{max}}{\pi(n)}\right) \\
= &O\left(\frac{x_{max}}{\min(\sigma\sqrt{n}, \pi(n))}\right)
\end{aligned} \tag{VIII.16}$$

□

Secondly, in Theorem VIII.2.2, we also show that the influence of an individual's demand on market-clearing price in the secondary stage is $O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)$ as $n \to +\infty$. As the proof is very similar to Theorem VIII.2.1, we provide the proof sketch here.

Theorem VIII.2.2. *For a player i, assume that Q_{-i} $(i \in N)$ are independent and identically distributed with finite mean μ and variance σ^2. For any q' and q'' s.t. $0 \leq q', q'' \leq q_{max}$, we have $\left|\mathbb{E}_{Q_{-i}}[P(\mathbf{Q},\mathbf{x})|(Q_i = q')] - \mathbb{E}_{Q_{-i}}[P(\mathbf{Q},\mathbf{x})|(Q_i = q'')]\right| = O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)$ as $n \to +\infty$, given $\{x_1, x_2 \ldots, x_n\}$ have been determined.*

Proof (Sketch). Let $z_n = \sum_{j=1}^{n} x_j$, and we assume it has been determined. Firstly, we note that given any realization of Q_i (i.e. $Q_1, \ldots, Q_{i-1}, Q_{i+1}, \ldots, Q_n$), the market-clearing price p in the secondary market is *monotonically non-decreasing* with the increasing of q_i. Take the expectation of Q_{-i}, the expected price is also monotonically non-decreasing with the increasing the q_i. So we get

$$\begin{aligned} &\left|\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=q')}\right] - \mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=q'')}\right]\right| \\ =&\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=q_{max})}\right] - \mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=0)}\right] \end{aligned} \tag{VIII.17}$$

Let $p_1 = P(\mathbf{Q},\mathbf{x})|_{(Q_i=0)}$, $p_2 = P(\mathbf{Q},\mathbf{x})|_{(Q_i=q_{max})}$, and $\Delta p = p_2 - p_1$, and we consider different realizations of $\mathbf{Q}$ as follows,

- When $\sum_{j\neq i} q_j \leq z_n - q_{max}$, no matter how much player i purchased in the primary market, the supply is always over (or equal to) demand. So the market-clearing price is always 0, i.e.

$$\Delta p = 0 \tag{VIII.18}$$

- When $\sum_{j\neq i} q_j > z_n$, the demand is always over supplies. When q_i changes from 0 to q_{max}, the demand in the secondary market will be increased by q_{max}, which will lead to an increase of the price. Similar to the proof of Theorem VIII.2.1, we can show

$$\Delta p = O\left(\frac{q_{max}}{\pi(n)}\right) \tag{VIII.19}$$

- When $z_n - q_{max} < \sum_{j\neq i} q_j \leq z_n$, the demand can be either over or under the supply. Similar as the proof of Theorem VIII.2.1, we could get

$$\Delta p = \Theta(1) \tag{VIII.20}$$

Based on Lemma VIII.2.1, we have $\mathbb{P}\left(z_n - q_{max} < \sum_{j\neq i} q_j \leq z_n\right) = O\left(\frac{q_{max}}{\sigma\sqrt{n}}\right)$.

Combining the three cases of the realization of $\mathbf{Q}$, we conclude:

$$\left|\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=q')}\right]-\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|_{(Q_i=q'')}\right]\right|=O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right) \tag{VIII.21}$$

□

Together Theorem VIII.2.1 and VIII.2.2, we conclude that an individual's influence on the price in the secondary market is $O\left(\frac{\max(x_{max},q_{max})}{\min(\sigma\sqrt{n},\pi(n))}\right)$ as $n\to+\infty$. In other words, for player $i\in N$, for any x_i,x_{-i} and Q_i, we have:

$$\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})\middle|Q_i\right]=\mathbb{E}_{Q_{-i}}\left[P(Q_{-i},x_{-i})\right]\pm O\left(\frac{\max(x_{max},q_{max})}{\min(\sigma\sqrt{n},\pi(n))}\right). \tag{VIII.22}$$

VIII.2.3 Existence and Characteristics of the asymptotic Nash Equilibrium

In this subsection, we firstly analyze the utility functions and best responses of players, and then discuss the existence and characteristics of the asymptotic Nash equilibrium defined in Definition VIII.2.1. We find that an asymptotic Nash equilibrium among players always exists, and the expected price in the secondary market equal to the price in the primary market, as the number of players approaches to infinity. It implies that the secondary market can mitigate the demand uncertainty of the first stage. If a player is in short of resources, she may buy resources at the *same* (expected) price as c in the secondary market to reduce the penalty; if a player has redundant resources, she may also sell resources in (expected) price c to reduce the waste.

Before introducing the main result, we provide two lemmas that are helpful in our analysis.

Lemma VIII.2.2. *Assume that* $\{Q_i\}$ $(i\in N)$ *are independent and identically distributed with finite mean* μ *and variance* σ^2*. For a player* $i\in N$*, for any* x' *and* x''*, s.t.* $0\le x',x''\le$

x_{max}, *we have*

$$\begin{aligned}&\left|\mathbb{E}_{\mathbf{Q}}\left[\min(P(\mathbf{Q},\mathbf{x}),c_i')|(x_i=x')\right]-\mathbb{E}_{\mathbf{Q}}\left[\min(P(\mathbf{Q},\mathbf{x}),c_i')|(x_i=x'')\right]\right|\\=&O\left(\frac{x_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)\end{aligned}\tag{VIII.23}$$

as $n\to+\infty$, *given* $\{x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_n\}$ *have been determined.*

Proof (Sketch). Based on Theorem VIII.2.1,

$$\left|\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i=x')\right]-\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q},\mathbf{x})|(x_i=x'')\right]\right|=O\left(\frac{x_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)\tag{VIII.24}$$

As $\left|\mathbb{E}_{\mathbf{Q}}\left[c_i'|(x_i=x')\right]-\mathbb{E}_{\mathbf{Q}}\left[c_i'|(x_i=x'')\right]\right|=0$, we coud get the result. □

Lemma VIII.2.3. *For player* $i\in N$, *assume that* $\{Q_{-i}\}$ *are independent and identically distributed with finite mean* μ *and variance* σ^2. *For any* q' *and* q'' *s.t.* $0\le q',q''\le q_{max}$, *we have*

$$\begin{aligned}&\left|\mathbb{E}_{Q_{-i}}\left[\min(P(\mathbf{Q},\mathbf{x}),c_i')|(Q_i=q')\right]-\mathbb{E}_{Q_{-i}}\left[\min(P(\mathbf{Q},\mathbf{x}),c_i')|(Q_i=q'')\right]\right|\\=&O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)\end{aligned}\tag{VIII.25}$$

as $n\to+\infty$, *given* $\{x_1,x_2\ldots,x_n\}$ *have been determined.*

Proof (Sketch). Based on Theorem VIII.2.2,

$$\left|\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|(Q_i=q')\right]-\mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})|(Q_i=q'')\right]\right|=O\left(\frac{q_{max}}{\min(\sigma\sqrt{n},\pi(n))}\right)\tag{VIII.26}$$

As $\left|\mathbb{E}_{Q_{-i}}\left[c_i'|(Q_i=q')\right]-\mathbb{E}_{Q_{-i}}\left[c_i'|(Q_i=q'')\right]\right|=0$, we could get the result. □

Intuitively, the two lemmas show that player i's influence to $\min(P(\mathbf{Q},\mathbf{x}),c_i')$ is also $O\left(\frac{1}{\min(\sqrt{n},\pi(n))}\right)$ as $n\to+\infty$. Similarly as Equation VIII.22, for player $i\in N$, for any x_i,x_{-i}

and Q_i, we have

$$\mathbb{E}_{Q_{-i}}\left[\min(P(\mathbf{Q},\mathbf{x}),c_i')\middle|Q_i\right] = \mathbb{E}_{Q_{-i}}\left[\min(P(Q_{-i},x_{-i}),c_i')\right] \pm O\left(\frac{\max(x_{max},q_{max})}{\min(\sigma\sqrt{n},\pi(n))}\right) \tag{VIII.27}$$

We now reorganize the (expected) utility function in Equation VIII.7. For easiness of notation, let $\psi(n) = \frac{\max(x_{max},q_{max})}{\min(\sigma\sqrt{n},\pi(n))}$, and $\mathscr{P} = P(Q_{-i},x_{-i})$ in the following discussion.

$$\begin{aligned}
&\mathbf{U}_i(x_i,x_{-i}) \\
=&\mathbb{E}_{\mathbf{Q}}\left[-cx_i + (x_i-Q_i)^+ P(\mathbf{Q},\mathbf{x}) - (Q_i-x_i)^+ \min\left(P(\mathbf{Q},\mathbf{x}),c_i'\right)\right] \\
=&\mathbb{E}_{Q_i}\Big\{-cx_i + (x_i-Q_i)^+ \mathbb{E}_{Q_{-i}}\left[P(\mathbf{Q},\mathbf{x})\middle|Q_i\right] \\
&-(Q_i-x_i)^+ \mathbb{E}_{Q_{-i}}\left[\min\left(P(\mathbf{Q},\mathbf{x}),c_i'\right)\middle|Q_i\right]\Big\} \\
=&\mathbb{E}_{Q_i}\Big\{-cx_i + (x_i-Q_i)^+ \mathbb{E}_{Q_{-i}}\left[P(Q_{-i},x_{-i})\right] \\
&-(Q_i-x_i)^+ \mathbb{E}_{Q_{-i}}\left[\min\left(P(Q_{-i},x_{-i}),c_i'\right)\right]\Big\} \pm O\left(\psi(n)\right) \\
=&\mathbb{E}_{Q_i}\Big\{\left[\mathbb{E}_{Q_{-i}}(\mathscr{P})-c\right]x_i + (Q_i-x_i)^+\left[\mathbb{E}_{Q_{-i}}(\mathscr{P})) - \mathbb{E}_{Q_{-i}}\left[\min\left(\mathscr{P},c_i'\right)\right]\right] \\
&-Q_i\mathbb{E}_{Q_{-i}}(\mathscr{P})\Big\} \pm O\left(\psi(n)\right)
\end{aligned} \tag{VIII.28}$$

It is easy to see that $\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}\left[\min\left(\mathscr{P},c_i'\right)\right] \geq 0$, and $\mathbf{U}_i(x_i,x_{-i})$ is convex for x_i, as $n \to \infty$. Then the best response of player i can be got only when $x_i = 0$ or $x_i = x_{max}$.

- When $x_i = 0$, then $\mathbf{U}_i(x_i,x_{-i}) = -\mathbb{E}_{Q_i}(Q_i)\,\mathbb{E}_{Q_{-i}}\left[\min(\mathscr{P},c_i')\right] \pm O\left(\psi(n)\right)$;
- When $x_i = x_{max}$, then $\mathbf{U}_i(x_i,x_{-i}) = \left[\mathbb{E}_{Q_{-i}}(\mathscr{P})-c\right]x_{max} - \mathbb{E}_{Q_i}(Q_i)\,\mathbb{E}_{Q_{-i}}(\mathscr{P}) \pm O\left(\psi(n)\right)$.

Let $A = \mathbb{E}_{Q_i}(Q_i)\left\{\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}\left[\min(\mathscr{P},c_i')\right]\right\}$, and $B = \left[\mathbb{E}_{Q_{-i}}(\mathscr{P})-c\right]x_{max}$, then the best response of player i (as $n \to \infty$) is

$$x_i^* = \begin{cases} 0, & A > B \\ 0 \text{ or } x_{max}, & A = B \\ x_{max}, & A < B \end{cases} \tag{VIII.29}$$

Note that when $A = B$, player i is indifferent between 0 and x_{max}.

We now discuss the existence and characteristics of asymptotic Nash equilibrium among players, as shown in Theorem VIII.2.3.

Theorem VIII.2.3. *Asymptotic Nash equilibrium always exists, and* $\mathbf{x}^* = \{x_1^*, x_2^*, \ldots, x_n^*\}$ *is an asymptotic Nash equilibrium* if and only if *the following conditions holds,*

1. *For any player* $i \in N$, *there exists a threshold* θ^*, *in which*

$$x_i^* = \begin{cases} 0, & c_i' < \theta^* \\ 0 \text{ or } x_{max}, & c_i' = \theta^* \\ x_{max}, & c_i' > \theta^* \end{cases} \tag{VIII.30}$$

2. $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = c$ *as* $n \to +\infty$.

Proof. Firstly we prove that if a strategy profile $\mathbf{x}^*$ satisfies the two conditions above, then it is an asymptotic Nash equilibrium. Assume that there is a threshold θ^*, such that players in $N_+ = \{i \in N | c_i' > \theta^*\}$ order x_{max} and players in $N_- = \{i \in N | c_i' < \theta^*\}$ order 0, and for corresponding $\mathbf{x}^*$, $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = c$ holds as $n \to \infty$, we now show that $\mathbf{x}^*$ is an asymptotic Nash equilibrium.

When $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = c$, as $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = \mathbb{E}_{Q_{-i}}\left[P(Q_{-i}, x_{-i}^*)\right] \pm O(\psi(n))$, then we could get $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = \mathbb{E}_{Q_{-i}}\left[P(Q_{-i}, x_{-i}^*)\right]$ as $n \to +\infty$. And Equation VIII.28 can be simplified as follows (let $\mathscr{P} = P(Q_{-i}, x_{-i}^*)$),

$$\begin{aligned} &\mathbf{U}_i(x_i^*, x_{-i}^*) \\ =&\mathbb{E}_{Q_i}\left\{(Q_i - x_i^*)^+ \left\{\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}\left[\min(\mathscr{P}, c_i')\right]\right\} - Q_i c\right\} \pm O(\psi(n)) \end{aligned} \tag{VIII.31}$$

As $\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}[\min(\mathscr{P}, c_i')] \geq 0$, we could get that $\mathbf{U}_i(x_i^*, x_{-i}^*)$ is a *non-increasing* function with x_i^* (as $n \to \infty$).

- For any player i that in $N_- = \{i \in N | c_i' < \theta^*\}$, her best response is $x_i^* = 0$. As $\mathbf{U}_i(x_i^*, x_{-i}^*)$ is *non-increasing*, her expected utility cannot be improved.

- For any player i that in $N_+ = \{i \in N | c_i' > \theta^*\}$, her best response is $x_i^* = x_{max}$. We have known that $c_i' > \theta^* > \mathscr{P}$ (based on Equation VIII.6), so $\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}[\min(\mathscr{P}, c_i')] = 0$. So we could get

$$\mathbf{U}_i(x_i^*, x_{-i}^*) = -\mathbb{E}_{Q_i}(Q_i)c \pm O(\psi(n)) \tag{VIII.32}$$

  Then player i cannot gain by deviating.

So x^* is an asymptotic Nash equilibrium.

Then we show that if $\mathbf{x}^*$ is an asymptotic Nash equilibrium, then condition 1 holds. It is easy to see that all players order 0 or all players order x_{max} cannot be an asymptotic Nash equilibrium. Otherwise, all players will become buyers (or sellers) in the secondary market, and some players can deviate from their strategies to gain profit.

From Equation VIII.29, we observe that: If $x_i^* = x_{max}$, then for any player j with $c_j' > c_i'$, $x_j^* = x_{max}$; If $x_i^* = 0$, then for any player j with $c_j' < c_i'$, $x_j^* = 0$. Then we could know that the condition 1 holds on $\mathbf{x}^*$.

Then we show that if strategy profile $\mathbf{x}^*$ is an asymptotic Nash equilibrium, then we could get $\mathbb{E}_{\mathbf{Q}}\left[P(\mathbf{Q}, \mathbf{x}^*)\right] = c$ as $n \to +\infty$. We prove it by contradiction:

- If $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] < c$, we claim that players in $N_+ = \{i \in N | c_i' > \theta^*\}$ (i.e. players those order x_{max}) have incentive to deviate. When $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] < c$, for player i in $N_+ = \{i \in N | c_i' > \theta^*\}$, we have $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = \mathbb{E}_{Q_{-i}}(\mathscr{P})$ as $n \to +\infty$, and

$$\mathbf{U}_i(x_i^*, x_{-i}^*) = \mathbb{E}_{Q_i}\left\{\left[\mathbb{E}_{Q_{-i}}(\mathscr{P}) - c\right] x_i^* - Q_i \mathbb{E}_{Q_{-i}}(\mathscr{P})\right\} \pm O(\psi(n)) \tag{VIII.33}$$

  As $\mathbb{E}_{Q_{-i}}(\mathscr{P}) - c < 0$, any player $i \in N_+$ has incentive to decrease the x_i. So it cannot be an asymptotic Nash equilibrium.

- If $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] > c$, we claim that there exists player $i \in N_- = \{i \in N | c_i' < \theta^*\}$ (i.e.

players those order 0 unit of resources) has incentive to deviate. Intuitively, player i with c_i' that is very closed to θ^* may have incentive to deviate. Based on Assumption VIII.2.1 and definition of market clearing price, for any small δ, there exists ε such that for player i with $c_i' = \theta^* - \varepsilon$, and $c_i' \geq \mathscr{P} + \delta$. So we could get

$$\begin{aligned}
&\mathbf{U}_i(x_i = 0, x_{-i}^*) \\
=&\mathbb{E}_{Q_i}(Q_i)\left[\mathbb{E}_{Q_{-i}}(\mathscr{P}) - \mathbb{E}_{Q_{-i}}(\min(\mathscr{P}, c_i'))\right] - \mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}(\mathscr{P}) \pm O(\psi(n)) \\
=&0 - \mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}(\mathscr{P}) \pm O(\psi(n))
\end{aligned} \tag{VIII.34}$$

If $\mathbb{E}_{Q_{-i}}(\mathscr{P}) - c > 0$ as $n \to \infty$, then

$$\begin{aligned}
&\mathbf{U}_i(x_i = x_{max}, x_{-i}^*) \\
=&\mathbb{E}_{Q_i}(Q_i)\left[\mathbb{E}_{Q_{-i}}(\mathscr{P}) - c\right] x_{max} - \mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}(\mathscr{P}) \pm O(\psi(n)) \\
>&0 - \mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}(\mathscr{P}) \pm O(\psi(n)) \\
=&\mathbf{U}_i(x_i = 0, x_{-i}^*)
\end{aligned} \tag{VIII.35}$$

Then player i has incentive to deviate and $\mathbf{x}^*$ cannot be an asymptotic Nash equilibrium.

So we conclude that $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = c$ as $n \to +\infty$.

Recall that we have assumed that $x_{max} \geq q_{max}$. If all players order either 0 or x_{max} in the primary market, then we can always adjust the proportion of players those order 0 or x_{max} to make $\mathbb{E}_{\mathbf{Q}}[P(\mathbf{Q}, \mathbf{x}^*)] = c$ hold as $n \to \infty$. So the existence of the equilibrium is straightforward.

To sum up, asymptotic Nash equilibrium exists and $\mathbf{x}^*$ is an asymptotic Nash equilibrium *if and only if* the two conditions hold. □

Intuitively, Theorem VIII.2.3 warrants the existence of the asymptotic Nash equilibrium and characterizes the equilibrium. In the equilibrium, players are divided into two

sets based on the marginal penalty of shortage c'. Those players with relatively high c' would order as many resources as they can in the primary market and become sellers in the secondary market, and players with relatively low c' would not order resources and wait to buy resources in the secondary market. More importantly, in equilibrium, the expected price in the secondary market is equal to the price in the primary market, as the number of players approaches infinity. The properties of the asymptotic equilibrium imply that secondary market could mitigate the demand uncertainty in the primary market.

VIII.2.4 Social Welfare and Aggregated Orders

Based on the existence and characteristics of asymptotic Nash equilibrium in Theorem VIII.2.3, we discuss the social welfare (i.e. $\sum_{i\in N}\mathbf{U}(x_i,x_{-i})$) and the aggregated orders (i.e. $\sum_{i\in N}x_i$) in the game with a secondary market, comparing these to the social welfare and aggregated orders *without* a secondary market. We find that both social welfare and aggregated orders *with* a secondary market are significantly better than those *without* when the number of players is approaching infinity.

We begin by discussing the *optimal* social welfare and aggregated orders. The *optimal* social welfare and aggregated orders in this case can be obtained when we treat all players as a *single* player, in which the (stochastic) demand is $\sum_{i\in N}Q_i$. Assume players' demands are independent and identically distributed, let Q denote the random demand of a player, and $q_1,q_2,\ldots,q_n$ be a sequence of realization of random variable Q. Based on the Law of Large Numbers, as $n\to+\infty$, $\frac{1}{n}\sum_{i\in N}q_i\to\mathbb{E}_Q(Q)$. Consequently, the authority can simply order $n\mathbb{E}_Q(Q)$ resources and distribute resources to players according to their needs at price c. Let $SW_\#$ and $OD_\#$ denote the *optimal* social welfare and aggregated orders, respectively. Then we can easily obtain that

$$SW_\# = -cn\mathbb{E}_Q(Q) \tag{VIII.36}$$

and

$$OD_\# = n\mathbb{E}_Q(Q). \tag{VIII.37}$$

Let $SW_{\overline{2nd}}$ denote the social welfare in the case *without* a secondary market, where each player makes an independent choice about the optimal amount of resource to order (note that this is also the optimal social welfare in the setting without a secondary market, since all player decisions are entirely decoupled). Correspondingly, let SW_{2nd} denote the social welfare *in the equilibrium* of the game *with* a secondary market. Similarly, let $OD_{\overline{2nd}}$ and OD_{2nd} denote the aggregated order of players in cases of *without* and *with* a secondary market (in the latter case, in equilibrium), respectively.

Our main results in this section (Theorems VIII.2.4 and VIII.2.5) show that the difference between the optimal and equilibrium outcomes is bounded by $O(\max(\sqrt{n}, \frac{n}{\pi(n)}))$, a significant improvement over the $\Theta(n)$ difference between the optimum and the outcome with only the newsvendor model (i.e., without the secondary market).

Theorem VIII.2.4. *The difference between optimal social welfare and the social welfare* without *secondary market is* $\Theta(n)$*, i.e.* $\left|SW_{\#} - SW_{\overline{2nd}}\right| = \Theta(n)$*, as* $n \to +\infty$.

Proof. Based on Equation VIII.2, without secondary market, for player i, the (expected) utility is $\mathbf{U}_i(x_i^*) = \mathbb{E}_Q\left[-cx_i^* - c_i'(Q - x_i^*)^+\right]$, in which $x_i^* = F_Q^{-1}\left(\frac{c_i'-c}{c'}\right)$. And it can be also denoted as $\mathbf{U}_i(x_i^*) = -c\mathbb{E}_Q(Q) \pm \Theta(1)$. As players are independent, the social welfare is $SW_{\overline{2nd}} = \sum_{i\in N}\mathbf{U}_i(x_i^*) = -cn\mathbb{E}_Q(Q) \pm \Theta(n)$ as $n \to +\infty$. So $\left|SW_{\#} - SW_{\overline{2nd}}\right| = \Theta(n)$ as $n \to +\infty$. □

Theorem VIII.2.5. *The difference between optimal social welfare and the social welfare in the equilibrium* with *secondary market is upper bounded by* $O(n \cdot \psi(n))$*, i.e.* $\left|SW_{\#} - SW_{2nd}\right| = O(n \cdot \psi(n))$ *as* $n \to +\infty$*, where* $\psi(n) = \frac{\max(x_{max}, q_{max})}{\min(\sigma\sqrt{n}, \pi(n))}$.

Proof. With secondary market, based on Theorem VIII.2.3, in equilibrium, $\mathbb{E}_{\mathbf{Q}}(P(\mathbf{Q}, \mathbf{x})) = c$ as $n \to +\infty$. For any player i, $\mathbb{E}_{Q_{-i}}[P(Q_{-i}, x_{-i})] = c \pm O(\psi(n))$.

We have known that in equilibrium, players's strategy is either 0 or x_{max},

- If $x_i = 0$, then

$$\begin{aligned}\mathbf{U}_i(x_i = 0, x_{-i}) &= -\mathbb{E}_{Q_i}\left\{Q_i \mathbb{E}_{Q_{-i}}\left[\min(P(Q_{-i}, x_{-i}), c_i')\right]\right\} \pm O(\psi(n)) \\ &= -\mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}\left[\min(P(Q_{-i}, x_{-i}), c_i')\right] \pm O(\psi(n)) \\ &\geq -\mathbb{E}_{Q_i}(Q_i)\mathbb{E}_{Q_{-i}}\left[P(Q_{-i}, x_{-i})\right] \pm O(\psi(n)) \\ &= -\mathbb{E}_{Q_i}(Q_i)\, c \pm O(\psi(n))\end{aligned} \tag{VIII.38}$$

- If $x_i = x_{max}$,

$$\begin{aligned}\mathbf{U}_i(x_i = x_{max}, x_{-i}) =& \mathbb{E}_{Q_i}\left\{-c x_{max} + (x_{max} - Q_i)\mathbb{E}_{Q_{-i}}\left[P(Q_{-i}, x_{-i})\right]\right\} \\ & \pm O(\psi(n)) \\ =& \mathbb{E}_{Q_i}\left[-c x_{max} + (x_{max} - Q_i)\, c\right] \pm O(\psi(n)) \\ =& -\mathbb{E}_{Q_i}(Q_i)\, c \pm O(\psi(n))\end{aligned} \tag{VIII.39}$$

To sum them up, then (ignoring the constant factor):

$$SW_{2nd} = \sum_{i \in N} \mathbf{U}_i(x_i, x_{-i}) \geq -cn\,\mathbb{E}_Q(Q) \pm O(n \cdot \psi(n)) \tag{VIII.40}$$

So we could get that:

$$|SW_{\#} - SW_{2nd}| = O(n \cdot \psi(n)) \tag{VIII.41}$$

□

Similarly, the sum of orders in the equilibrium of the game *with* a secondary market is also much closer to the optimal sum of orders than the case *without* a secondary market.

Theorem VIII.2.6. *The difference between the aggregated orders* without *secondary market and the optimal aggregated orders is* $\Theta(n)$*, i.e.* $\left|OD_{\overline{2nd}} - OD_{\#}\right| = \Theta(n)$*, as* $n \to +\infty$*.*

Proof. Based on Equation VIII.2, without secondary market, the optimal order for player i is $x_i^* = F_Q^{-1}\left(1 - \frac{c}{c_i'}\right)$. It can be also denoted as $x_i^* = \mathbb{E}_Q(Q) \pm \Theta(1)$. So the aggregated

order is $OD_{\overline{2nd}} = n\mathbb{E}_Q(Q) \pm \Theta(n)$, and $\left|OD_{\overline{2nd}} - OD_{\#}\right| = \Theta(n)$ as $n \to +\infty$. □

Theorem VIII.2.7. *The difference between the aggregated orders* with *secondary market and the optimal aggregated orders is upper bounded by* $O(\frac{1}{c}n \cdot \psi(n))$*, i.e.* $\left|SW_{\#} - SW_{2nd}\right| = O(\frac{1}{c}n \cdot \psi(n))$ *as* $n \to +\infty$*, where* $\psi(n) = \frac{\max(x_{max}, q_{max})}{\min(\sigma\sqrt{n}, \pi(n))}$.

Proof. Based on Equation VIII.40 in the proof of Theorem VIII.2.5, we could know that $SW_{2nd} = -cn\mathbb{E}_Q(Q) \pm O(n \cdot \psi(n))$. As the price in the primary market is c and the expected price in a secondary market is also c. So the overall quantity of resources in the two markets is $OD_{2nd} = \frac{|SW_{2nd}|}{c} = n\mathbb{E}_Q(Q) \pm O(\frac{1}{c}n \cdot \psi(n))$. So we could get $|OD_{2nd} - OD_{\#}| = O(\frac{1}{c}n \cdot \psi(n))$ as $n \to +\infty$. □

VIII.3 Small Markets: Two Players

In this section, I discuss another side of the problem, where there are only two players. In this case, the secondary market takes the form of a trade of the resource between the two players. It is therefore natural to use the bargaining framework to model this market. In particular, we use the Nash bargaining solution concept to characterize the price in the secondary market [146].

Building on the Nash bargaining price in the secondary market, I firstly provide a sufficient condition that guarantees the existence of a pure strategy Nash equilibrium, and then characterize the symmetric equilibrium in a symmetric game by comparing the aggregated order and social welfare under equilibrium with those *without* a secondary market. I find that, in the two-player case, the aggregated order under equilibrium *with* a secondary market is *not necessarily* lower than that *without* a secondary market, but always more "closer" to the optimal aggregated order. Furthermore, the social welfare under equilibrium *with* a secondary market is always larger or equal to that *without* a secondary market.

Let the player set be $N = \{i, j\}$. Without loss of generality, we focus on player i - the argument for player j is symmetric. For player i, recall that her demand follows the random variable Q_i, the realized demand is q_i, the marginal penalty of resource shortage is c_i', and

the amount of ordered resources (i.e. strategy in the game) from the primary market is x_i. In this case $\mathbf{Q} = \{Q_i, Q_j\}$, and $\mathbf{x} = \{x_i, x_j\}$.

Assuming the price in the secondary market is $P(\mathbf{Q}, \mathbf{x})$, the expected utility of player i can be expressed as follows,

$$\begin{aligned}\mathbf{U}_i(x_i, x_j) = \mathbb{E}_{\mathbf{Q}}\big\{ & -cx_i + \min\left[(x_i - Q_i)^+, (Q_j - x_j)^+\right] P(\mathbf{Q}, \mathbf{x}) \\ & - \left[(Q_i - x_i)^+ - (x_j - Q_j)^+\right]^+ c_i' \\ & - \min\left[(Q_i - x_i)^+, (x_j - Q_j)^+\right] P(\mathbf{Q}, \mathbf{x})\big\}.\end{aligned} \tag{VIII.42}$$

Intuitively, in the secondary market, when $x_i > Q_i$ and $x_j < Q_j$, player i could sell $\min\left[(x_i - Q_i)^+, (Q_j - x_j)^+\right]$ units of resources to player j; when $x_i < Q_i$ and $x_j > Q_j$, player i could buy $\min\left[(Q_i - x_i)^+, (x_j - Q_j)^+\right]$ units of resources from player j, and the marginal penalty of shortage is c_i' if she is still in short of resources after the trading in the market; finally when $x_i < q_i$ and $x_j < q_j$, both players are in short of resources, and player i has to pay the penalty c_i' for each unit of shortage.

VIII.3.1 Nash Bargaining Price

When there are two players, the price in the secondary market can be any value between 0 and c_i', assuming player i is the buyer and j is the seller in the market. It is also possible that both two players have redundant or shortage of resources, and there is no trading in the secondary market. We assume that the two players are *bargaining* about the price in the secondary market when one player has an excess supply (i.e. the seller) and the other has excess demand. Specifically, we consider the *Nash bargaining* [146] solution concept to identify the price in the secondary market.

The Nash bargaining solution is the unique solution to a two-player bargaining problem that satisfies the axioms of scale invariance, symmetry, efficiency, and independence of irrelevant alternatives [146]. In our case, we find that the Nash bargaining price only depends on the marginal penalty of shortage for the buyer in the market. Indeed, the resulting

solution takes a particularly simple form, as the following theorem attests.

Theorem VIII.3.1. *When there are two players i and j, assume i is the buyer and j is the seller in the secondary market. Then the Nash bargaining price in the secondary market is* $\frac{1}{2}c_i'$.

Proof. Assume the price is p, and the amount of resources traded is y. Then the utility *gained* by accepting the price for player i is $(c_i' - p)y$, and the utility *gained* for player j is py. Based on Nash bargaining solution in [146], we maximize the product of $(c_i' - p)y \cdot py$, so the Nash bargaining price is $p = \frac{1}{2}c_i'$. □

Based on Theorem VIII.3.1, the price in the secondary market (whenever trade is possible) is characterized as follows:

$$P(\mathbf{Q},\mathbf{x}) = \begin{cases} \frac{1}{2}c_i', & x_i < Q_i \text{ and } x_j > Q_j \\ \frac{1}{2}c_j', & x_i > Q_i \text{ and } x_j < Q_j. \end{cases} \tag{VIII.43}$$

VIII.3.2 Existence of Pure Strategy Nash Equilibrium between Two Players

We will show a sufficient condition that guarantees the existence of pure strategy Nash equilibrium, and also show that symmetric equilibrium always exists when the two players are symmetric in the game.

Given the Nash bargaining price in the secondary market and a realization of demands $\{q_i, q_j\}$, the utility function of player i writes as follows,

$$\begin{aligned} U_i(x_i,x_j) = &- cx_i + \min\left[(x_i - q_i)^+, (q_j - x_j)^+\right] \frac{1}{2}c_j' \\ &- \left[(q_i - x_i)^+ - (x_j - q_j)^+\right]^+ c_i' \\ &- \min\left[(q_i - x_i)^+, (x_j - q_j)^+\right] \frac{1}{2}c_i' \end{aligned} \tag{VIII.44}$$

And the expected utility function of player i is

$$\mathbf{U}_i(x_i,x_j) = \mathbb{E}_{\mathbf{Q}}\left[U_i(x_i,x_j)\right] \tag{VIII.45}$$

Equation VIII.44 can be divided into two cases: (1) $x_j - q_j > 0$, (2) $x_j - q_j \leq 0$. Taking gradient w.r.t. x_i for the two cases, we obtain the following:

- When $x_j - q_j > 0$,

$$\frac{\partial \left[U_i(x_i, x_j)\right]}{\partial x_i} = \begin{cases} -c + c_i', & x_i - q_i \leq q_j - x_j \\ -c + \frac{1}{2}c_i', & q_j - x_j < x_i - q_i \leq 0 \\ -c, & x_i - q_i > 0 \end{cases} \tag{VIII.46}$$

- When $x_j - q_j \leq 0$,

$$\frac{\partial \left[U_i(x_i, x_j)\right]}{\partial x_i} = \begin{cases} -c + c_i', & x_i - q_i \leq 0 \\ -c + \frac{1}{2}c_j', & 0 < x_i - q_i \leq q_j - x_j \\ -c, & x_i - q_i > q_j - x_j \end{cases} \tag{VIII.47}$$

Now we show a sufficient condition that guarantees the existence of pure strategy Nash equilibrium, which based on the concavity of $U_i(x_i, x_j)$ derived from equation VIII.46 and VIII.47.

Theorem VIII.3.2. *Pure strategy equilibrium exists in the two player game if* $\frac{1}{2} \leq \frac{c_i'}{c_j'} \leq 2$.

Proof. Based on Equation VIII.46 and VIII.47, $U_i(x_i, x_j)$ is *concave* when $\frac{c_i'}{c_j'} \geq \frac{1}{2}$. Similarly, $U_j(x_j, x_i)$ is also *concave* when $\frac{c_i'}{c_j'} \leq 2$. So when $\frac{1}{2} \leq \frac{c_i'}{c_j'} \leq 2$, both $U_i(x_i, x_j)$ and $U_j(x_j, x_i)$ are *concave*. After taking expectation, $\mathbf{U}_i(x_i, x_j)$ and $\mathbf{U}_j(x_j, x_i)$ both remain being *concave*. As the strategy spaces of players are continuous and compact, based on [149], pure strategy Nash equilibrium exists in the two player game. □

For detailed characterization, we will focus on the *symmetric game*, in which players have and identical distribution of (stochastic) demand and identical marginal penalty of resource shortage. And we mainly consider *symmetric equilibrium* in the game.

Definition VIII.3.1. *(x_i, x_j) is a symmetric pure strategy Nash equilibrium if it is pure strategy Nash equilibrium and $x_i = x_j$.*

Again we show that symmetric equilibrium exists in the symmetric game.

Theorem VIII.3.3. *Symmetric pure strategy Nash equilibrium exists in the two-player symmetric game.*

Proof. Based on Equation VIII.46 and VIII.47, when two players are symmetric, $c_i' = c_j'$, both $\mathbf{U}_i(x_i, x_j)$ and $\mathbf{U}_j(x_j, x_i)$ are *concave* (similar as the argument in proof of Theorem VIII.3.2). As the strategy spaces of players are continuous and compact, based on Theorem 3 in [150], symmetric equilibrium always exists in the two-player symmetric game. □

VIII.3.3 Characteristics of a Symmetric Pure Strategy Nash Equilibrium

In this subsection, we characterize the symmetric pure strategy Nash equilibrium between two players when the game is symmetric. We will compare aggregated order and social welfare under symmetric equilibrium with those *without* a secondary market. Note that there are only two players, therefore the aggregated order is simply *twice* of the number of resources ordered by each player. The aggregated order in equilibrium *with* a secondary market will be shown to be not necessarily less than the one *without* a secondary market, it is however always "closer" to the optimal aggregated order. And we also prove that the social welfare in equilibrium *with* a secondary market is always larger or equal to that *without* a secondary market.

Recall that, without a secondary market, the order that maximizes the expected utility of player *i* is simply given by $x^* = F_{Q_i}^{-1}\left(1 - \frac{c}{c_i'}\right)$, where $F_{Q_i}^{-1}(\cdot)$ denotes the inverse cumulative distribution function of Q_i. Before introducing the main results, we first analyze the *optimal* aggregated order and social welfare in a *centralized* mechanism. The *optimal* aggregated order and social welfare can be got when we treat two players as a single player, in which the (stochastic) demand is $Q_i + Q_j$, and the marginal penalty of resource shortage is c_i'. Let

$2x^{\#}$ denote the *optimal* aggregated order, then

$$2x^{\#} = F^{-1}_{Q_i+Q_j}\left(1-\frac{c}{c_i'}\right), \tag{VIII.48}$$

where $F^{-1}_{Q_i+Q_j}(\cdot)$ is the inverse cumulative distribution function of Q_i+Q_j. Reorganizing Equation VIII.48 yields

$$\mathbb{P}(Q_i+Q_j \geq 2x^{\#}) = \frac{c}{c_i'} \tag{VIII.49}$$

For ease of exposition, slightly abusing our previous notation, let $SW_{\overline{2nd}}(x)$ denote the social welfare for the case *without* a secondary market when each player orders x in the primary market. Similarly, let $SW_{2nd}(x)$ and $SW_{\#}(x)$ denote social welfare *with* a secondary market and optimal social welfare when each player orders x, respectively. We observe that, for any given x, we have

$$SW_{\#}(x) \geq SW_{2nd}(x) \geq SW_{\overline{2nd}}(x), \text{ for any } x \geq 0 \tag{VIII.50}$$

Intuitively, $SW_{\#}(x)$ can be seen as a centralized mechanism, and it is easy to see that $SW_{\#}(x) \geq SW_{\overline{2nd}}(x)$ and $SW_{\#}(x) \geq SW_{2nd}(x)$; given a specific amount of resources, with secondary market, resources can be traded between players to *improve* the social welfare after the realization of demands, so $SW_{2nd}(x) \geq SW_{\overline{2nd}}(x)$.

Assuming the symmetric equilibrium in the two-player game is (x^e, x^e), we are ready to compare x^*, x^e, and $x^{\#}$ for the aggregated order, and compare $SW_{\overline{2nd}}(x^*)$, $SW_{2nd}(x^e)$, and $SW_{\#}(x^{\#})$ for social welfare.

Note that the aggregated order is $\sum_{k\in N} x_k = x_i + x_j$. To compare the aggregated order *with* and *without* a secondary market, we need only compare x^* with x^e.

Before presenting the main results, we first derive the best responds of players. As $\mathbf{U}_i(x_i, x_j)$ is concave, the best responds of player i obtains when $\frac{\partial \mathbf{U}_i(x_i,x_j)}{\partial x_i} = 0$. Interchange the order of expectation and gradient, we get,

$$\frac{\partial \mathbf{U}_i(x_i,x_j)}{\partial x_i} = \frac{\partial \mathrm{E}_{\mathbf{Q}}\left[U_i(x_i,x_j)\right]}{\partial x_i} = \mathrm{E}_{\mathbf{Q}}\frac{\partial\left[U_i(x_i,x_j)\right]}{\partial x_i} \tag{VIII.51}$$

We consider 4 disjoint parts of domain of Q_i and Q_j. (also illustrated in Figure VIII.1)

1. $A(x) = \{(Q_i,Q_j)\big|\, 0 \le Q_j \le x, Q_i + Q_j \ge 2x\}$,
2. $B(x) = \{(Q_i,Q_j)\big|\, 0 \le Q_j \le x, Q_i \ge x, Q_i + Q_j \le 2x\}$,
3. $C(x) = \{(Q_i,Q_j)\big|\, Q_j \ge x, Q_i \ge x\}$,
4. $D(x) = \{(Q_i,Q_j)\big|\, Q_j \ge x, 0 \le Q_i \le x, Q_i + Q_j \ge 2x\}$.

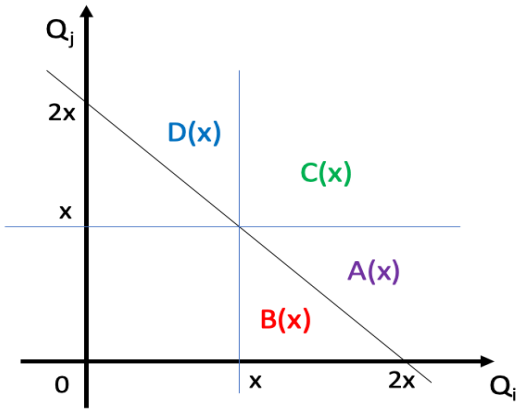

Figure VIII.1: Illustration of $A(x)$, $B(x)$, $C(x)$, and $D(x)$

In the above four parts, $A(x)$ and $B(x)$ correspond to the first two cases in Equation VIII.46; $C(x)$ and $D(x)$ correspond to the first two cases in Equation VIII.47. It easy to verify that Equation VIII.51 can be denoted as follows,

$$\frac{\partial \mathbf{U}_i(x_i,x_j)}{\partial x_i}\Big|_{x_j=x_i} = -c + c_i'\mathbb{P}(A(x_i)) + \frac{1}{2}c_i'\mathbb{P}(B(x_i)) + c_i'\mathbb{P}(C(x_i)) + \frac{1}{2}c_i'\mathbb{P}(D(x_i)) \tag{VIII.52}$$

Let $\frac{\partial \mathbf{U}_i(x_i,x_j)}{\partial x_i}\Big|_{x_j=x_i} = 0$, assume the equilibrium is (x^e,x^e), then we have

$$\mathbb{P}(A(x^e)) + \frac{1}{2}\mathbb{P}(B(x^e)) + \mathbb{P}(C(x^e)) + \frac{1}{2}\mathbb{P}(D(x^e)) = \frac{c}{c_i'} \tag{VIII.53}$$

Let $G(x) = \mathbb{P}(A(x)) + \frac{1}{2}\mathbb{P}(B(x)) + \mathbb{P}(C(x)) + \frac{1}{2}\mathbb{P}(D(x))$, then we could get the following lemma:

Lemma VIII.3.1. *For any given distribution of Q_i, $G(x)$ is monotonically* non-increasing *as the increasing of x.*

Proof. Assume $0 < x_k < x_i$, then the following statements are easy to verify,

- If $(Q_i, Q_j) \in A(x_i)$, then $(Q_i, Q_j) \in A(x_k) \cup C(x_k)$;
- If $(Q_i, Q_j) \in B(x_i)$, then $(Q_i, Q_j) \in A(x_k) \cup B(x_k) \cup C(x_k)$;
- If $(Q_i, Q_j) \in C(x_i)$, then $(Q_i, Q_j) \in C(x_k)$;
- If $(Q_i, Q_j) \in D(x_i)$, then $(Q_i, Q_j) \in C(x_k) \cup D(x_k)$.

In $G(x)$, the coefficient of $\mathbb{P}(A(x))$ and $\mathbb{P}(C(x))$ is 1, and the coefficient of $\mathbb{P}(B(x))$ and $\mathbb{P}(D(x))$ is $\frac{1}{2}$. For any (Q_i, Q_j), the coefficients of the corresponding item in $G(x)$ are *non-increasing* when we increase x_k to x_i. So we could get that $G(x_k) \geq G(x_i)$, and $G(x)$ is monotonically non-increasing with the increasing of x. □

We now show a *sufficient* and *necessary* condition under which the aggregated order in equilibrium (i.e. $2x^e$) is less or equal to that without a secondary market (i.e. $2x^*$):

Lemma VIII.3.2. *In symmetric equilibrium (x^e, x^e) of two-player symmetric game, $x^e \leq x^*$ if and only if $\mathbb{P}\left(Q_i + Q_j \geq 2x^*\right) \leq \mathbb{P}(Q_i \geq x^*)$, in which $x^* = F_{Q_i}^{-1}\left(1 - \frac{c}{c_i'}\right)$.*

Proof. As $\mathbb{P}(A(x)) + \mathbb{P}(C(x)) + \mathbb{P}(D(x)) = \mathbb{P}((Q_i, Q_j) | Q_i + Q_j \geq x)$, let $x = x^{\#}$, then

$$\mathbb{P}(A(x^{\#})) + \mathbb{P}(C(x^{\#})) + \mathbb{P}(D(x^{\#})) = \mathbb{P}((Q_i, Q_j) | Q_i + Q_j \geq x^{\#}) = \frac{c}{c'} \tag{VIII.54}$$

in which $2x^{\#} = F_{Q_i+Q_j}^{-1}\left(\frac{c_i' - c}{c_i'}\right)$.

Under equilibrium (x^e, x^e), we know Equation VIII.53 holds, therefore

$$\begin{aligned}
&\mathbb{P}(A(x^e)) + \frac{1}{2}\mathbb{P}(B(x^e)) + \mathbb{P}(C(x^e)) + \frac{1}{2}\mathbb{P}(D(x^e)) \\
=& \frac{c}{c_i'} \\
=& \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#))
\end{aligned} \tag{VIII.55}$$

As $G(x)$ is monotonically non-increasing with the increasing of x, so

$$\begin{aligned}
& x^e \geq x^\# \\
\Longleftrightarrow\ & G(x^\#) \geq G(x^e) = \frac{c}{c_i'} = \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(A(x^\#)) + \frac{1}{2}\mathbb{P}(B(x^\#)) + \mathbb{P}(C(x^\#)) + \frac{1}{2}\mathbb{P}(D(x^\#)) \geq \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(B(x^\#)) \geq \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(A(x^\#)) + \mathbb{P}(B(x^\#)) + \mathbb{P}(C(x^\#)) \geq \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(Q_i + Q_j \geq 2x^\#) \leq \mathbb{P}(Q_i \geq x^\#)
\end{aligned} \tag{VIII.56}$$

completing the proof. □

We further show another *sufficient* and *necessary* condition, under which the aggregated order in equilibrium (i.e. $2x^e$) is greater or equal to the *optimal* aggregated order (i.e. $2x^\#$).

Lemma VIII.3.3. *In symmetric equilibrium (x^e, x^e) of two-player symmetric game, $x^e \geq x^\#$ if and only if $\mathbb{P}\left(Q_i + Q_j \geq 2x^\#\right) \leq \mathbb{P}(Q_i \geq x^\#)$, in which $2x^\# = F^{-1}_{Q_i+Q_j}\left(\frac{c_i'-c}{c_i'}\right)$.*

Proof. As $\mathbb{P}(A(x)) + \mathbb{P}(C(x)) + \mathbb{P}(D(x)) = \mathbb{P}((Q_i, Q_j) | Q_i + Q_j \geq x)$, let $x = x^\#$, then

$$\mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) = \mathbb{P}((Q_i, Q_j) | Q_i + Q_j \geq x^\#) = \frac{c}{c'} \tag{VIII.57}$$

in which $2x^\# = F^{-1}_{Q_i+Q_j}\left(\frac{c_i'-c}{c_i'}\right)$.

Under equilibrium (x^e, x^e), we know Equation VIII.53 holds, therefore

$$\begin{aligned}
&\mathbb{P}(A(x^e)) + \frac{1}{2}\mathbb{P}(B(x^e)) + \mathbb{P}(C(x^e)) + \frac{1}{2}\mathbb{P}(D(x^e)) \\
=& \frac{c}{c_i'} \\
=& \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#))
\end{aligned} \tag{VIII.58}$$

As $G(x)$ is monotonically non-increasing with the increasing of x, so

$$\begin{aligned}
& x^e \geq x^\# \\
\Longleftrightarrow\ & G(x^\#) \geq G(x^e) = \frac{c}{c_i'} = \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(A(x^\#)) + \frac{1}{2}\mathbb{P}(B(x^\#)) + \mathbb{P}(C(x^\#)) + \frac{1}{2}\mathbb{P}(D(x^\#)) \geq \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(B(x^\#)) \geq \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(A(x^\#)) + \mathbb{P}(B(x^\#)) + \mathbb{P}(C(x^\#)) \geq \mathbb{P}(A(x^\#)) + \mathbb{P}(C(x^\#)) + \mathbb{P}(D(x^\#)) \\
\Longleftrightarrow\ & \mathbb{P}(Q_i + Q_j \geq 2x^\#) \leq \mathbb{P}(Q_i \geq x^\#)
\end{aligned} \tag{VIII.59}$$

completing the proof. □

The above two lemmas help us prepare Theorem VIII.3.4. The theorem states that x^e is always between $x^\#$ and x^*. Intuitively, when the marginal penalty is very large, then players have a strong incentive to *over-request* resources *without* a secondary market, i.e. $x^* > x^\#$. In this case, introducing a secondary market could decrease the total order of resources (comparing with the case *without* a secondary market). On the other hand, when the marginal penalty is very close to the unit cost of the resource, then players have incentive to *under-request* resources *without* a secondary market, i.e. $x^* < x^\#$. Introducing a secondary market could increase the total order of resources (comparing with the case *without* a secondary market).

Theorem VIII.3.4. *In symmetric equilibrium (x^e, x^e) of two-player symmetric game,*

- *If $x^\# \leq x^*$, then $x^\# \leq x^e \leq x^*$;*

- *If $x^{\#} \geq x^*$, then $x^{\#} \geq x^e \geq x^*$.*

Proof. The proof proceeds by cases:

- If $x^{\#} \leq x^*$, we have

$$\mathbb{P}(Q_i + Q_j \geq 2x^*) \leq \mathbb{P}(Q_i + Q_j \geq 2x^{\#}) = \frac{c}{c_i'} = \mathbb{P}(Q_i \geq x^*) \tag{VIII.60}$$

$$\mathbb{P}(Q_i + Q_j \geq 2x^{\#}) = \frac{c}{c_i'} = \mathbb{P}(Q_i \geq x^*) \leq \mathbb{P}(Q_i \geq x^{\#}) \tag{VIII.61}$$

  Based on Lemma VIII.3.2 and VIII.3.3, we could get that $x^{\#} \leq x^e \leq x^*$

- If $x^{\#} \geq x^*$, we have

$$\mathbb{P}(Q_i + Q_j \geq 2x^*) \geq \mathbb{P}(Q_i + Q_j \geq 2x^{\#}) = \frac{c}{c_i'} = \mathbb{P}(Q_i \geq x^*) \tag{VIII.62}$$

$$\mathbb{P}(Q_i + Q_j \geq 2x^{\#}) = \frac{c}{c_i'} = \mathbb{P}(Q_i \geq x^*) \geq \mathbb{P}(Q_i \geq x^{\#}) \tag{VIII.63}$$

  Based on Lemma VIII.3.2 and VIII.3.3, we could get that $x^{\#} \geq x^e \geq x^*$

□

Now we are ready to show that the social welfare in equilibrium with secondary market is always at least that without a secondary market when there are two players.

Theorem VIII.3.5. *In symmetric equilibrium (x^e, x^e) of two-player symmetric game, we always have $SW_{\#}(x^{\#}) \geq SW_{2nd}(x^e) \geq SW_{\overline{2nd}}(x^*)$.*

Proof. Note that $SW_{\#}(x)$ and $SW_{\overline{2nd}}(x)$ are both *single-peaked* with peak point at $x = x^{\#}$ and $x = x^*$, respectively. We argue in two cases,

- When $x^{\#} \leq x^*$, Theorem VIII.3.4 implies that $x^{\#} \leq x^e \leq x^*$. Based on the single-peakedness of $SW_{\#}(x)$ and $SW_{\overline{2nd}}(x)$, we further have

$$SW_{\#}(x^{\#}) \geq SW_{\#}(x^e) \geq SW_{\#}(x^*) \tag{VIII.64}$$

$$SW_{\overline{2nd}}(x^*) \geq SW_{\overline{2nd}}(x^e) \geq SW_{\overline{2nd}}(x^\#) \tag{VIII.65}$$

Based on Equation VIII.50, we have $SW_\#(x^e) \geq SW_{2nd}(x^e) \geq SW_{\overline{2nd}}(x^e)$. So we conclude $SW_\#(x^\#) \geq SW_{2nd}(x^e) \geq SW_{\overline{2nd}}(x^*)$.

- When $x^\# \geq x^*$, Lemma VIII.3.4 implies that $x^\# \geq x^e \geq x^*$. Based on the single-peakedness of $SW_\#(x)$ and $SW_1(x)$, we have

$$SW_\#(x^\#) \geq SW_\#(x^e) \geq SW_\#(x^*) \tag{VIII.66}$$

$$SW_{\overline{2nd}}(x^*) \geq SW_{\overline{2nd}}(x^e) \geq SW_{\overline{2nd}}(x^\#) \tag{VIII.67}$$

Similar as the argument in the first case, we again conclude $SW_\#(x^\#) \geq SW_{2nd}(x^e) \geq SW_{\overline{2nd}}(x^*)$.

□

VIII.4 Conclusion

This chapter investigates the impact of the secondary market to demand uncertainty. I develop a two-stage model with multiple players. In the first stage, players with uncertain demands order and obtain resources from the authority; in the second stage, after the realization of players' demands, resources can be traded among them. I derive the optimal decisions and pure strategy Nash equilibrium for players in the first stage, along with the price of resources in the secondary market, in both a large market with an infinite number of players and a small market with two players. In the large market, with asymptotic analysis, I find that the expected price in the secondary market is the same as the price in the primary market, and the aggregated order and social welfare under equilibrium *with* the secondary market are both better than those *without*. In the small market, I show that the aggregated order under equilibrium *with* the secondary market is always closer to the optimal aggregated order than that *without* the secondary market. And the social welfare under equilibrium *with* the secondary market is guaranteed to be no worse than the social

welfare *without*.

In this chapter, I assumed zero trading cost in the secondary market. In a large market, if a non-zero trading cost is considered, characterizing players' best response can be highly non-trivial, and the corresponding equilibrium analysis will become more involved. This merits more efforts to understand in the future. I also assume that players' demands are independent of each other. In reality, players' demands are likely to be correlated. A more elaborate model incorporating demand correlation will be a future direction that worths to pursue. Analysis of the game with more than two but not infinite players is also a challenging but meaningful work to do.

CHAPTER IX

Conclusion

In this thesis, I study situations which lead to efficiency/inefficiency in equilibrium outcomes, as well as approaches for mitigating the inefficiency.

I firstly examined a non-cooperative multi-defender security game in which defenders may protect multiple targets, offering complete characterization Average-case Stackelberg Equilibrium (or equivalently, Nash equilibrium among defenders) and approximate equilibria, socially optimal solutions, and price of anarchy. The results show that defenders generally over-protect the targets and efficiency may degrade due to players' selfish behavior. Furthermore, I also study the multi-defender security game model in spear-phishing attacks and dynamic traffic light control, respectively.

Then I consider the mechanism design in coalition formation problem, and I mainly use computational methods to design mechanisms which have good theoretical and empirical properties. I study both general coalition formation and roommates problem. For coalition formation mechanisms, I mainly address the computational challenges in implementing *rotating proposer mechanism*, which implements a subgame perfect Nash equilibrium in the corresponding rotating proposer game, and evaluate the mechanism by empirical methods. To address the challenges, I introduce preprocessing and pruning, as well as approximate versions of RPM, one tailored to the roommate problem (with coalitions of at most two), and another for coalitions of arbitrary size. The experiments show that even the approximate versions of RPM significantly outperform several alternative mechanisms for coalition formation in terms of social welfare and fairness, are do not introduce significant incentives to misreport preferences. Then I present the first treatment of the roommates problem from an automated mechanism design (AMD). The proposed approaches and analysis are mainly based on two restricted incentive compatibility: promotion-one incentive compatibility, in

which the manipulation consists of promoting a single player to the top position in the manipulator's preference list; and permutation incentive compatibility, in which the manipulation is a permutation of the preference. I also propose several automated mechanism design approaches to the roommates problem. I prove that the first approach is ε-permutation IC. While the other two don't have theoretical guarantees, empirical results show that they have even better incentive properties.

Finally, I investigate the impact of the secondary market to demand uncertainty in a resource allocation problem. I develop a two-stage model with multiple players. In the first stage, players with uncertain demands order and obtain resources from the authority; in the second stage, after the realization of players' demands, resources can be traded among them. I derive the optimal decisions and pure strategy Nash equilibrium for players in the first stage, along with the price of resources in the secondary market, in both a large market with an infinite number of players and a small market with two players. In the large market, with asymptotic analysis, I find that aggregated order and social welfare under equilibrium *with* the secondary market are both better than those *without*. In the small market, I show that the aggregated order under equilibrium *with* the secondary market is always closer to the optimal aggregated order than that *without* the secondary market. And the social welfare under equilibrium *with* the secondary market is guaranteed to be no worse than the social welfare *without*.

Below, I describe some other possible works I will do in the future.

- Mechanism Design in Computational Resource Allocation: In chapter VIII, I introduce a secondary market to make the resource allocation inside a company more efficient. We can also treat the problem as a mechanism design problem. In the problem, the authority is *benevolent*, i.e. its goal is improving the efficiency of a system instead of gaining more profits. The work in the thesis is an option to improve the efficiency of the system. But it is worthwhile to consider some other ways to deal with the mechanism design problem of allocating computational resources inside a

big company.

- Optimal Pricing in Cloud Services: I will try to leverage the Stackelberg model to optimize the pricing scheme in cloud services, especially spot instance in elastic computing. I will try to differentiate the different level of cloud service quality and decide how to set price for different levels of quality. We can consider it as a two-stage sequential game, in which the seller set the prices for different services in the first stage, and the buyers choose the most favored service to buy. The challenge of the problem is that we do not know users' preferences, even the distribution of the preferences. I will try to combine game theory, robust optimization and online learning to deal with the problem.

- Optimal Ad Auction Design: I will also do some research on mechanism design in sponsored search auction. Generalized second price (GSP) auction is widely used and well studied in sponsored search auction. One advantage of GSP is that it is very simple and does not rely on the prior knowledge of users utility functions. However, A straightforward question is whether we can get higher revenue if we can get more prior knowledge of users' utility functions. One possible option is Myerson's mechanism, which has very high revenue but is also very sensitive to the error of estimation of users' utility function. In the future, I will try to design more robust ad auction mechanism that could deal with the error of estimation better.

BIBLIOGRAPHY

[1] Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos. Optimal personalized filtering against spear-phishing attacks. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 958–964, Jan 2015.

[2] Mengchen Zhao, Bo An, and Christopher Kiekintveld. An initial study on personalized filtering thresholds in defending sequential spear phishing attacks. In *Proceedings of the 2015 IJCAI Workshop on Behavioral, Economic and Computational Intelligence for Security*, Jul 2015.

[3] Manish Jain, James Pita, Milind Tambe, Fernando Ordóñez, Praveen Paruchuri, and Sarit Kraus. Bayesian stackelberg games and their application for security at los angeles international airport. *SIGecom Exch.*, 7(2):10:1–10:3, June 2008.

[4] James Pita, Manish Jain, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Using game theory for los angeles airport security. *AI Magazine*, 30(1):43–57, 2009.

[5] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, pages 13–20, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[6] Security games with arbitrary schedules: A branch and price approach. In Maria Fox and David Poole, editors, *AAAI*. AAAI Press, 2010.

[7] Manish Jain, Jason Tsai, James Pita, Christopher Kiekintveld, Shyamsunder Rathi, Milind Tambe, and Fernando Ordóòez. Software assistants for randomized patrol

planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4):267–290, July 2010.

[8] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, pages 689–696, Richland, SC, 2009. International Foundation for Autonomous Agents and Multiagent Systems.

[9] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, EC '06, pages 82–90, New York, NY, USA, 2006. ACM.

[10] Yoram Bachrach, Moez Draief, and Sanjeev Goyal. Contagion and observability in security domains. Allerton Conference, 2013.

[11] Gerry Smith. Massive Target hack traced back to phishing email. Huffpost Business, http://www.huffingtonpost.com/2014/02/12/target-hack_n_4775640.html, Feb 2014.

[12] Robin Sidel. Target to settle claims over data breach. The Wall Street Journal, http://www.wsj.com/articles/target-reaches-settlement-with-visa-over-2013-data-breach-1439912013, Aug 2015.

[13] Kim Zetter. A cyberattack has caused confirmed physical damage for the second time ever. Wired, http://www.wired.com/2015/01/german-steel-mill-hack-destruction/, Jan 2015.

[14] Declan McCullagh. White House confirms 'spearphishing' intrusion. CNET, http://www.cnet.com/news/white-house-confirms-spearphishing-intrusion/, October 2012.

[15] James Rogers. Hackers attack Nuclear Regulatory Commission 3 times in 3 years. FOX News, http://www.foxnews.com/tech/2014/08/20/hackers-attack-nuclear-regulatory-commission/, August 2014.

[16] Kim Zetter. Top federal lab hacked in spear-phishing attack. Wired, http://www.wired.com/2011/04/oak-ridge-lab-hack/, April 2011.

[17] Jason Hong. The state of phishing attacks. *Communications of the ACM*, 55(1):74–81, 2012.

[18] Cormac Herley and Dinei Florêncio. A profitless endeavor: Phishing as tragedy of the commons. In *Proceedings of the 2008 New Security Paradigms Workshop (NSPW)*, pages 59–70. ACM, 2008.

[19] Aron Laszka, Mark Felegyhazi, and Levente Buttyan. A survey of interdependent information security games. *ACM Computing Surveys*, 47(2):23:1–23:38, Aug 2014.

[20] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, 2003.

[21] Hau Chan, Michael Ceyko, and Luis E Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attacks. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 152–162, Aug 2012.

[22] Jian Lou and Yevgeniy Vorobeychik. Equilibrium analysis of multi-defender security games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 596–602, 2015.

[23] Dirk Helbing. Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.*, 73:1067–1141, Dec 2001.

[24] Takashi Nagatani. The physics of traffic jams. *Reports on Progress in Physics*, 65(9):1331, 2002.

[25] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Rev. E 62, Issue*, 62:2000, 2000.

[26] Dirk Helbing, Jan Siegmeier, and Stefan Lmmer. Self-organized network flows. *NHM*, 2(2):193–210, 2007.

[27] M. Ebrahim Fouladvand, M. Reza Shaebani, and Zeinab Sadjadi. Intelligent controlling simulation of traffic flow in a small city network. *Journal of the Physical Society of Japan*, 73(11):3209–3214, 2004.

[28] Carlos Gershenson and David A. Rosenblueth. Self-organizing traffic lights at multiple-street intersections. *Complexity*, 17(4):23–39, 2012.

[29] S. Lammer, H. Kori, K. Peters, and D. Helbing. Decentralised control of material or traffic flows in networks using phase-synchronisation. *Physica A: Statistical Mechanics and its Applications*, 363(1):39–47, April 2006.

[30] Ming Zhong, Satish Sharma, and Pawan Lingras. Letter genetically-designed time delay neural networks for multiple-interval urban freeway traffic flow forecasting, 2006.

[31] S. Mikami and Y. Kakazu. Genetic reinforcement learning for cooperative traffic signal control. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 223–228 vol.1, Jun 1994.

[32] Tahere Royani, Javad Haddadnia, and Mohammad Alipoor. Traffic signal control for isolated intersections based on fuzzy neural network and genetic algorithm. In *Proceedings of the 10th WSEAS International Conference on Signal Processing, Computational Geometry and Artificial Vision*, ISCGAV'10, pages 87–91, Stevens Point, Wisconsin, USA, 2010. World Scientific and Engineering Academy and Society (WSEAS).

[33] R. Hoar, J. Penner, and C. Jacob. Evolutionary swarm traffic: if ant roads had traffic lights. In *Evolutionary Computation, 2002. CEC '02. Proceedings of the 2002 Congress on*, volume 2, pages 1910–1915, 2002.

[34] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201 – 230, 2002.

[35] José Alcalde and Pablo Revilla. Researching with whom? stability and manipulation. *Journal of Mathematical Economics*, 40(8):869 – 887, 2004.

[36] Haris Aziz, Felix Brandt, and Paul Harrenstein. Fractional hedonic games. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 5–12. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[37] A. E. Roth. The economics of matching: Stability and efficiency. *Mathematics of Operations Research*, 92:617–628, 1982.

[38] José Alcalde and Salvador Barberà. Top dominance and the possibility of strategy-proof stable solutions to matching problems. *Economic Theory*, 4(3):417–435, 1994.

[39] Greg Leo, Jian Lou, Martin Van der Linden, Yevgeniy, and Myrna Wooders. Matching soulmates. Working paper, Available at SSRN: https://ssrn.com/abstract=2833553, February 2017.

[40] Atila Abdulkadiroglu and Tayfun Sonmez. School choice: A mechanism design approach. In *American Economic Review*, volume 93(3), pages 729–747, 2003.

[41] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[42] Robert W Irving. An efficient algorithm for the "stable roommates" problem. In *Journal of Algorithms*, volume 6(4), pages 577–595, 1985.

[43] Sophie Bade. Serial dictatorship: The unique optimal allocation rule when information is endogenous. *Theoretical Economics*, 10:385–410, 2015.

[44] Mason Wright and Yevgeniy Vorobeychik. Mechanism design for team formation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1050–1056, 2015.

[45] Matthew Chamber, Chen Hajaj, Greg Leo, Jian Lou, Martin Van der Linden, Yevgeniy Vorobeychik, and Myrna Wooders. *Non-Cooperative Team Formation and a Team Formation Mechanism*, 2017.

[46] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[47] Katarina Cechlarova and Sona Ferkova. The stable crews problem. *Discrete Applied Mathematics*, 140(1):1 – 17, 2004.

[48] Francesca Forno and Roberta Garibaldi. Sharing economy in travel and tourism: The case of home-swapping in italy. *Journal of Quality Assurance in Hospitality & Tourism*, 16(2):202–220, 2015.

[49] Alvin E. Roth, Tayfun Sonmez, and M. Utku Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151 – 188, 2005.

[50] V. Conitzer and T Sandholm. An algorithm for automatically designing deterministic mechanisms without payments. In *International Conference on Autonomous Agents and Multiagent Systems*, 2004.

[51] I. A. Kash, R. Murty, and D. C. Parkes. Enabling spectrum sharing in secondary

market auctions. *IEEE Transactions on Mobile Computing*, 13(3):556–568, March 2014.

[52] Peter Cramton and Suzi Kerr. Tradeable carbon permit auctions: How and why to auction not grandfather. *Energy Policy*, 30(4):333 – 345, 2002.

[53] David B. Spence. Can law manage competitive energy markets. *Cornell Law Review*, 93:765–818, 2008.

[54] Heinrich Freiherr von Stackelberg. *Marktform und Gleichgewicht ((Market Structure and Equilibrium)*. Vienna, 1934.

[55] Antoine Augustin Cournot. *Recherches sur les principes math ematiques de la th eorie des richesses (Researches into the Mathematical Principles of the Theory of Wealth)*. Hachette, Paris, 1938.

[56] Praveen Paruchuri, Jonathan P. Pearce, Milind Tambe, Fernando Ordonez, and Sarit Kraus. An efficient heuristic approach for security against multiple adversaries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '07, pages 181:1–181:8, New York, NY, USA, 2007. ACM.

[57] Praveen Paruchuri, Jonathan P. Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '08, pages 895–902, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

[58] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed armor

protection: The application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track*, AAMAS '08, pages 125–132, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

[59] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, pages 689–696, Richland, SC, 2009. International Foundation for Autonomous Agents and Multiagent Systems.

[60] Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordóñez, and Milind Tambe. Security and game theory: Iris- a tool for strategic security allocation in transportation networks. In *AAMAS 2011*, 2011.

[61] William B. Haskell, Debarun Kar, Fei Fang, Milind Tamb, Sam Cheung, and Lt. Elizabeth Denicola. Robust protection of fisheries with compass. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI'14, pages 2978–2983. AAAI Press, 2014.

[62] Fei Fang, Albert Xin Jiang, and Milind Tambe. Protecting moving targets with multiple mobile resources. *Journal of Artificial Intelligence Research, 48:583-634*, 2013.

[63] Matthew P. Johnson, Fei Fang, and Milind Tambe. Patrol strategies to maximize pristine forest area. In *AAAI*, 2012.

[64] Fei Fang, Thanh Nguyen, Rob Pickles, Wai Lam, Gopalasamy Clements, Bo An,

Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying paws: Field optimization of the protection assistant for wildlife security, 2016.

[65] Bo Li and Yevgeniy Vorobeychik. Feature cross-substitution in adversarial classification. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2087–2095. Curran Associates, Inc., 2014.

[66] Weiyi Xia Ellen Wright Clayton Murat Kantarcioglu Ranjit Ganta Ray-mond Heatherly Bradley A. Malin Zhiyu Wan, Yevgeniy Vorobeychik. A game theoretic framework for analyzing re-identification risk. *PLoS ONE*, 2015.

[67] Swetasudha Panda and Yevgeniy Vorobeychik. Stackelberg games for vaccine design. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1391–1399, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[68] Albert Xin Jiang, Ariel D. Procaccia, Yundi Qian, Nisarg Shah, and Milind Tambe. Defender (mis)coordination in security games. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 220–226. AAAI Press, 2013.

[69] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, 2003.

[70] Eric Shieh, Albert Xin Jiang, Amulya Yadav, Pradeep Varakantham, and Milind Tambe. Unleasing dec-mdps in security games: Enabling effective defender teamwork. *European Conference on Artificial Intelligence*, 2014.

[71] Hau Chan, Michael Ceyko, and Luis E. Ortiz. Interdependent defense games: Modeling interdependent security under deliberate attack. In *Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 152–162, 2012.

[72] David L. Alderson, Gerald G. Brown, W. Matthew Carlyle, and R. Kevin Wood. Solving defender-attacker-defender models for infrastructure defense. *INFORMS Computing Society Conference*, 2011.

[73] Yoram Bachrach, Moez Draief, and Sanjeev Goyal. Contagion and observability in security domains. In *Allerton Conference*, 2013.

[74] Daron Acemoglu, Azarakhsh Malekian, and Asu Ozdaglar. Network security and contagion, 2013. Working paper.

[75] Yevgeniy Vorobeychik, Jackson Mayo, Robert Armstrong, and Joseph Ruthruff. Noncooperatively optimized tolerance: Decentralized strategic optimization in complex systems. *Physical Review Letters*, 107(10):108702, 2011.

[76] Diego Cerdeiro, Marcin Dziubinski, and Sanjeev Goyal. Individual security and network design. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*, 2014.

[77] Diego A. Cerdeiro, Marcin Dziubiński, and Sanjeev Goyal. Individual security, contagion, and network design. *Journal of Economic Theory*, 170:182 – 226, 2017.

[78] J. Lou, A. M. Smith, and Y. Vorobeychik. Multidefender security games. *IEEE Intelligent Systems*, 32(1):50–60, Jan 2017.

[79] Mengchen Zhao, Bo An, and Christopher Kiekintveld. Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 658–664. AAAI Press, 2016.

[80] Enrico Blanzieri and Anton Bryl. A survey of learning-based techniques of email spam filtering. *Artificial Intelligence Review*, 29(1):63–92, 2008.

[81] Ian Fette, Norman Sadeh, and Anthony Tomasic. Learning to detect phishing emails. In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, pages 649–656. ACM, May 2007.

[82] André Bergholz, Jan De Beer, Sebastian Glahn, Marie-Francine Moens, Gerhard Paaß, and Siehyun Strobel. New filtering approaches for phishing email. *Journal of Computer Security*, 18(1):7–35, 2010.

[83] Aron Laszka, Waseem Abbas, S. Shankar Sastry, Yevgeniy Vorobeychik, and Xenofon Koutsoukos. Optimal thresholds for intrusion detection systems. In *Proceedings of the Symposium and Bootcamp on the Science of Security*, HotSos '16, pages 72–81, New York, NY, USA, 2016. ACM.

[84] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. Deceiving cyber adversaries: A game theoretic approach. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, pages 892–900, Richland, SC, 2018. International Foundation for Autonomous Agents and Multiagent Systems.

[85] Shuxin Li, Xiaohong Li, Jianye Hao, Bo An, Zhiyong Feng, Kangjie Chen, and Chengwei Zhang. Defending against man-in-the-middle attack in repeated games. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, pages 3742–3748. AAAI Press, 2017.

[86] Haris Aziz and Rahul Savani. Hedonic games. In *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[87] Gilad Zlotkin and Jeffrey S. Rosenschein. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *AAAI-94 Proceedings*, 1994.

[88] Onn Shehory and Sarit Kraus. Task allocation via coalition formation among autonomous agents. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'95, pages 655–661, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.

[89] Tuomas W. Sandholm and Victor R. Lesser. Coalition formation among bounded rational agents. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'95, pages 662–669, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.

[90] Tuomas Sandholm, Sandeep Sikka, and Samphel Norden. Algorithms for optimizing leveled commitment contracts. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 535–540, 1999.

[91] Matthew E Gaston and Marie desJardins. Agent-organized networks for dynamic team formation. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 230–237. ACM, 2005.

[92] Leandro Soriano Marcolino, Albert Xin Jiang, and Milind Tambe. Multi-agent team formation: diversity beats strength? In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

[93] Nadia Burani and William S. Zwicker. Coalition formation games with separable preferences. *Mathematical Social Sciences*, 45(1):27–52, February 2003.

[94] Francis Bloch and Effrosyni Diamantoudi. Noncooperative formation of coalitions in hedonic games. *International Journal of Game Theory*, 40(2):263–280, May 2010.

[95] Szilvia Pápai. Unique stability in simple coalition formation games. *Games and Economic Behavior*, 48(2):337–354, August 2004.

[96] Carmelo Rodriguez-Alvarez. Strategy-proof coalition formation. *International Journal of Game Theory*, 38(3):431–452, 2009.

[97] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18(1):135–153, January 2001.

[98] H. Aziz, F. Brandt, and P. Harrenstein. Pareto optimality in coalition formation. *Games and Economic Behavior*, 82:562–581, 2013.

[99] Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–702, 1998.

[100] Wanyuan Wang, Zhanpeng He, Peng Shi, Weiwei Wu, and Yichuan Jiang. Truthful team formation for crowdsourcing in social networks (extended abstract). In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 1327–1328, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

[101] Szilvia Pápai. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5(1):91–105, March 2000.

[102] Michele Flammini, Gianpiero Monaco, and Qiang Zhang. Strategyproof mechanisms for additively separable hedonic games and fractional hedonic games. *CoRR*, abs/1706.09007, 2017.

[103] AE Roth. The Economics of Matching : Stability and Incentives Stable. *Mathematics of Operations Research*, 7(4):617–628, 1982.

[104] Jay Sethuraman, Chung-Piaw Teo, and Liwen Qian. Many-to-One Stable Matching: Geometry and Fairness. *Mathematics of Operations Research*, 31(3):581–596, August 2006.

[105] Federico Echenique and Jorge Oviedo. A Theory of Stability in Many-to-many Matching Markets. SSRN Scholarly Paper ID 691443, Social Science Research Network, Rochester, NY, October 2004.

[106] Fuhito Kojima and Parag A. Pathak. Incentives and Stability in Large Two-Sided Matching Markets. *The American Economic Review*, 99(3):608–627, June 2009.

[107] Kim-Sau Chung. On the existence of stable roommate matchings. *Games and Economic Behavior*, 33(2):206 – 230, 2000.

[108] Effrosyni Diamantoudi, Eiichi Miyagawa, and Licun Xue. Random paths to stability in the roommate problem. *Games and Economic Behavior*, 48(1):18 – 28, 2004.

[109] Bettina Klaus and Flip Klijn. Smith and rawls share a room: stability and medians. *Social Choice and Welfare*, 35(4):647–667, Oct 2010.

[110] Jens Gudmundsson. When do stable roommate matchings exist? a review. *Review of Economic Design*, 18(2):151–161, Jun 2014.

[111] Boris G. Pittel and Robert W. Irving. An upper bound for the solvability probability of a random stable roommates instance. *Random Structures & Algorithms*, 5(3):465–486, 1994.

[112] Péter Biró, Elena Iñarra, and Elena Molis. A new solution concept for the roommate problem: Q-stable matchings. *Mathematical Social Sciences*, 79:74 – 82, 2016.

[113] Alvin E. Roth and Uriel G. Rothblum. Truncation Strategies in Matching Markets—in Search of Advice for Participants. *Econometrica*, 67(1):21–43, 1999.

[114] Rohit Vaish and Dinesh Garg. Manipulating gale-shapley algorithm: Preserving stability and remaining inconspicuous. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 437–443, 2017.

[115] John G. Riley and William F. Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981.

[116] Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

[117] Peter Cramton. Spectrum auctions. *Handbook of Telecommunications Economics*, 2002.

[118] A. Gopinathan, Z. Li, and C. Wu. Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets. In *2011 Proceedings IEEE INFOCOM*, pages 3020–3028, April 2011.

[119] Martin Hoefer, Thomas Kesselheim, and Berthold Vöcking. Approximation algorithms for secondary spectrum auctions. *ACM Trans. Internet Technol.*, 14(2-3):16:1–16:24, October 2014.

[120] Yuefei Zhu, Baochun Li, and Z. Li. Truthful spectrum auction design for secondary networks. In *2012 Proceedings IEEE INFOCOM*, pages 873–881, March 2012.

[121] Robert N. Stavins. Experience with market-based environmental policy instruments. Technical report, Resources for the Future, Resources for the Future, 1616 P Street, NW Washington, D.C. 20036, 2001.

[122] Anita Talberg and Kai Swoboda. Emissions trading schemes around the world. Technical report, Parliament of Australia, 2013.

[123] Center for Climate and Energy Solutions. Climate change 101– cap and trade. Technical report, Center for Climate and Energy Solutions, 2011.

[124] Tom Tietenberg. The Tradable-Permits Approach to Protecting the Commons: Lessons for Climate Change. *Oxford Review of Economic Policy*, 19(3):400–419, 09 2003.

[125] Mahmut Parlar. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Research Logistics (NRL)*, 35(3):397–409, 1988.

[126] Fernando Bernstein and Awi Federgruen. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science*, 51(1):18–29, 2005.

[127] Bruce C. Hartman, Moshe Dror, and Moshe Shaked. Cores of inventory centralization games. *Games and Economic Behavior*, 31(1):26 – 49, 2000.

[128] Alfred Muller, Marco Scarsini, and Moshe Shaked. The newsvendor game has a nonempty core. *Games and Economic Behavior*, 38(1):118 – 126, 2002.

[129] Gérard P. Cachon and Serguei Netessine. *Game Theory in Supply Chain Analysis*, pages 13–65. Springer US, Boston, MA, 2004.

[130] M.G. Fiestras-Janeiro, I. Garcia-Jurado, A. Meca, and M.A. Mosquera. Cooperative game theory and inventory management. *European Journal of Operational Research*, 210(3):459 – 466, 2011.

[131] Hau Lee and Seungjin Whang. The impact of the secondary market on the supply chain. *Management Science*, 48(6):719–731, 2002.

[132] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordonez, and Milind Tambe. Computing optimal randomized resource allocations for massive security games. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*, 2009.

[133] Joshua Letchford and Yevgeniy Vorobeychik. Computing optimal security strategies for interdependent assets. In *Conference on Uncertainty in Artificial Intelligence*, pages 459–468, 2012.

[134] Aron Laszka, Jian Lou, and Yevgeniy Vorobeychik. Multi-defender strategic filtering against spear-phishing attacks. In *30th AAAI Conference on Artificial Intelligence (AAAI)*, February 2016.

[135] Jakob Erdmann Daniel Krajzewicz Michael Behrisch, Laura Bieker. Sumo – simulation of urban mobility sumo. In *The Third International Conference on Advances in System Simulation*, 2011.

[136] Julia Lima Fleck and Christos G. Cassandras. Infinitesimal perturbation analysis for quasi-dynamic traffic light controllers. In *12th International Workshop on Discrete Event Systems, WODES 2014, Cachan, France, May 14-16, 2014.*, pages 235–240, 2014.

[137] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47–97, Jan 2002.

[138] W.W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33:452–473, 1977.

[139] P. Nordlie. *A longitudinal study of interpersonal attraction in a natural group setting*. PhD thesis, University of Michigan, 1958.

[140] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*, IJCAI'11, pages 73–78. AAAI Press, 2011.

[141] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Transactions on Economics and Computation*, 1(4):18:1–18:26, 2013.

[142] James Schummer. Almost-dominant strategy implementation: exchange economies. *Games and Economic Behavior*, 48(1):154–170, 2004.

[143] S. Sanghvi and David Parkes. Hard-to-manipulate combinatorial auctions. Technical report, Harvard University, 2004.

[144] Alvin E Roth and No Jul. The Economist as Engineer : Game Theory , Experimentation , and Computation as Tools for Design Economics. 70(4):1341–1378, 2002.

[145] P. Erdös and A. Rényi. On the evolution of random graphs. In *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, pages 17–61, 1960.

[146] John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[147] K. J. Arrow, T. Harris, and Jacob Marshak. Optimal inventory policy. *Econometrica*, 1951.

[148] Ken Littlewood. Special issue papers: Forecasting and control of passenger bookings. *Journal of Revenue and Pricing Management*, 4(2):111–123, Apr 2005.

[149] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.

[150] Shih-Fen Cheng, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman. Notes on equilibria in symmetric games. In *In Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents (GTDT)*, pages 71–78, 2004.