

Online Platforms in Networked Markets: Transparency, Anticipation and Demand Management

Thesis by
John Zhen Fu Pang

In Partial Fulfillment of the Requirements for the
Degree of
Ph.D. in Computational and Mathematical Sciences

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2019
Defended May 9, 2019

© 2019

John Zhen Fu Pang
ORCID: 0000-0002-6485-7922

All rights reserved

ACKNOWLEDGEMENTS

I am deeply indebted to my advisor and friend, Adam Wierman, who has been integral in my development throughout graduate school. Adam has always supported my pursuits in research, and offered immense help towards my career. Even in his busiest moment, Adam found time for his graduate students.

Secondly, I would like to thank Steven Low, who was never officially my co-advisor, but had willingly played that role in the last five years. Steven is an amazing resource for students, and always found time for us despite his schedule. Steven's positive attitude and drive towards research is truly infectious.

Thirdly, I would also like to thank others who played important roles in this thesis. First on that list are my committee members, Professors Venkat Chandrasekaran and Laura Doval, for their help and guidance throughout the entire process. Next on that list are my collaborators: Eilyan Bitar, Subhonmesh Bose, Minghua Chen, Hu Fu, Linqi Guo, Weixuan Lin, Qiulin Lin, Zhaojian Wang, Hanling Yi and Pengcheng You. It was a great privilege interacting and working with each one of them. I am also fortunate to share an office with Xiaoqi Ren and Zachary Lee, both of whom have made long days in graduate school significantly more manageable.

I am also extremely grateful and humbled for the opportunity to be in Caltech, which truly is a place unlike any other. I believe there is no where else one can be surrounded by so many brilliant, yet approachable, scholars. I am also thankful for the support staff in the Computing and Mathematical Sciences Department who allowed us to focus on research, especially Maria Lopez, Christine Ortega, and Sheila Shull. Outside CMS, I am also grateful to have learnt much from Cassandra Horii and Jennifer Weaver from the Center for Teaching, Learning and Outreach.

Life in Pasadena extended beyond Caltech. I am blessed to have been cared for by Missio Church in the past two years. I would like to specially acknowledge Len and Amy Tang for their pastoral care and love, both for Odelia and myself.

Finally, I would like to thank my family. I would certainly not be who I am without my parents. Thanks for teaching me to learn from failure and reminding me you will always be there for me. This thesis would not be possible without the support, understanding and love from my wife, Odelia, who has held my hand throughout every difficult moment and never let go. Last but not least, I thank God for all His provisions in my life.

ABSTRACT

The global economy has been transformed by the introduction of online platforms in the past two decades. These companies, such as Uber and Amazon, have benefited and undergone massive growth, and are a critical part of the world economy today. Understanding these online platforms, their designs and how participation change with anticipation and uncertainty can help us identify the necessary ingredients for successful implementation of online platforms in the future, especially for those with underlying network constraints, e.g., the electricity grid.

This thesis makes three main contributions. First, we identify and compare common access and allocation control designs for online platforms, and highlight their trade-offs between transparency and control. We make these comparisons under a networked Cournot competition model and consider three popular designs: (i) open access, (ii) discriminatory access, and (iii) controlled allocation. Our findings reveal that designs that control over access are more efficient than designs that control over allocations, but open access designs are susceptible to substantial search costs. Next, we study the impact of demand management in a networked Stackelberg model considering network constraints and producer anticipation. We provide insights on limiting manipulation under these constrained networked marketplaces with nodal prices, and show that demand management mechanisms that traditionally aid system stability also help play a vital role economically. In particular, we show that demand management empower consumers and give them “market power” to counter that of producers, limiting the impact of their anticipation and their potential for manipulation. Lastly, we study how participants (e.g., drivers on Uber) make competitive real-time production (driving) decisions. To that end, we design a novel pursuit algorithm for making online optimization under limited inventory constraints. Our analysis yields an algorithm that is competitive and applicable to achieve optimal results in the well known one-way trading problem, and new variants of the original problem.

PUBLISHED CONTENT AND CONTRIBUTIONS

Lin, Qiulin et al. (2019). “Competitive Online Optimization under Inventory Constraints”. In: *arXiv:1901.09161*.

J.Z.F.P participated in the conception of the project, analyzed the problem, recommend and propose algorithms, introduced the generalizations and participated in the writing of the manuscript.

Pang, John ZF et al. (2019). “Transparency and Control in Networked Markets”. In: *arXiv:1903.04574*.

J.Z.F.P participated in the conception of the project, analyzed the problem, constructed the examples and participated in the writing of the manuscript.

Lin, Weixuan et al. (2017). “Networked cournot competition in platform markets: Access control and efficiency loss”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*.

J.Z.F.P participated in the conception of the problem, analyzed the problem, and participated in the writing of the manuscript. DOI: 10.1109/CDC.2017.8264340.

Pang, John ZF et al. (2017). “The Efficiency of Open Access in Platforms for Networked Cournot Markets”. In: *IEEE International Conference on Computer Communications*.

J.Z.F.P participated in the conception of the project, analyzed and solved the problem, and participated in the writing of the manuscript. DOI: 10.1109/INFOCOM.2017.8057125.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	vi
List of Illustrations	vii
List of Tables	viii
Chapter I: Introduction	1
1.1 Online Platform Design	4
1.2 Design Challenges	7
1.3 Main Contributions and Overview	9
Chapter II: Networked Cournot Competition for Online Platforms	11
2.1 The Networked Cournot Competition Model	12
Chapter III: Access and Allocation Control Designs for Online Platforms	20
3.1 Literature Review	23
3.2 Model and Preliminaries	25
3.3 Efficiency of Open Access	28
3.4 Inefficiency of Controlled Allocation Designs	34
3.5 Efficiency of Discriminatory Access	38
3.6 Search Costs and its Implications	42
Chapter IV: Demand Management in Online Platforms	45
4.1 Literature Review	49
4.2 Model & Preliminaries	50
4.3 Anticipation in Networked Markets	54
4.4 Economic Value of Demand Management	61
Chapter V: Towards Dynamic Participation in Online Platforms	68
5.1 Literature Review	71
5.2 Problem Formulation	73
5.3 CR-Pursuit Algorithmic Framework	77
5.4 Competitive Analysis of CR-Pursuit	81
5.5 Application to One-way Trading	82
Chapter VI: Conclusion and Future Work	86
6.1 Lessons and Insights	89
Bibliography	91
Appendix A: Proofs for Chapter III	100
Appendix B: Proofs for Chapter IV	126
Appendix C: Proofs for Chapter V	135
Afterward	
About the Author	

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
4.1 Two-node illustrative example defined as $\mathcal{G}(\alpha, \beta, \mathbf{c}, f)$	54
4.2 A plot of best-response dynamics for market demand fulfilled of the game $\mathcal{G}([1, 1]^T, [1, 1]^T, [2q^2, 1q^2]^T, f)$ for $f = 1/15$ (left) and $f = 1/20$ (right), where the bounds of f for non-existence of equilibria are $1/19 < f < 4/59$ respectively. The x-axis (resp. y-axis) represents the demand fulfilled at the first market d_1 (resp. d_2).	59
4.3 (Left) Profit or utility function of the producer in the two-node illustrative game $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1+\epsilon)q]^T, \infty)$ in red with axis on left.	60
4.4 Profit or utility function of the producer in the two-node illustrative game $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1+\epsilon)q]^T, \infty)$ in red with axis on left.	64

LIST OF TABLES

<i>Number</i>		<i>Page</i>
2.1	Equilibrium Computation Complexity (Abolhassani et al., 2014) . . .	16
5.1	Summary of Notations used in Chapter 5.	74

Chapter 1

INTRODUCTION

Before the insurgence of platforms, the world for decades was dominated primarily by “pipeline businesses” with linear value chains. Manufacturers in the past bought raw materials, added value by inventively combining them, and sold the end product at a significantly higher cost than the constituent raw materials. The value chain adds value and creates a competitive advantage through five primary activities: (i) inbound logistics, (ii) operations, (iii) outbound logistics, (iv) marketing and sales, and (v) service. Traditional firms carefully managed their production quantity and prices to maximize profits.

A lot has changed since then. The emergence of platforms has led to a new paradigm of online marketplaces. These platforms, unlike traditional retailers or providers, often do not need to produce on their own or store physical goods, but instead, play the role of a matchmaker. The “raw materials” that these platforms bring together are the groups of participants that they bring together, and not something they can necessarily buy or put through a value chain. Unlike their earlier counterparts, platforms cannot just manage production and prices. Instead, they need to consider *network effects*, i.e., an idea that the incremental benefit gained by an existing user grows with each new user that joins the network. As an extreme example, if no one else in your social circle had telephones, then neither owning a phone nor the telephone network bring any value. Only when at least one of your contacts is on the telephone network is there a benefit in you being on the system, and this benefit grows as more of your social circle is on the system.

Example 1.1. *As a ridesharing platform, Uber does not employ any of their drivers but instead aims to matchmake drivers and riders and help price these rides. A large rider pool means that drivers get ride requests faster, take a shorter time to get to their matched riders, and consequently earn more money. A significant driver pool implies that passengers have minimal waiting time for a ride to be served. According to a recent Bloomberg article (Newcomer, 2017), Uber is reporting that revenue growth is outpacing losses, despite its investors paying 2 billion dollars annually to subsidize rides.*

A recent estimate of the market value of some of these platforms alone in 2016 is US\$4.3 trillion as reported in (P. C. Evans and Gawer, 2016), representing a significant proportion of the US\$56.8 trillion market value of the Top 2000 Valued Companies globally according to a recent Forbes report (Forbes, 2018). However, starting a platform is not straightforward. First, it requires a solution to the platform design “chicken-and-egg” problem: How do you convince one side of the market to join if the other side is empty? The platform needs to reach a critical mass to benefit from the so-called network effects. Platforms need to subsidize at least one side of the market to achieve this critical mass. Evans and Schlamensee dedicate their popular science book *Matchmakers: The New Economics of Multisided Platforms* to those who may have failed in the process:

“To all the pioneers who tried to cross the critical mass frontier.” (D. S. Evans and Schmalensee, 2016)

These matchmakers or multisided platforms are not new. The oldest recorded auction house is the *Stockholms Auktionsverk*, or Stockholm’s Auction House, founded in 1694. On the other hand, popular auction platform eBay was founded about three centuries after in 1995. Meanwhile, newspaper advertisements started in the 18th Century, while Google introduced Search in 1998, and began monetizing it through ads in 2000. While these multisided matchmaking platforms have existed in different forms for a long time, dramatic advances in technology have driven the platform’s recent popularity and profitability. Online retail platforms have almost zero marginal cost for display and storage. The zero marginal cost means that items which were once less in demand and therefore not displayed or sold in traditional stores can now be displayed, sold, and delivered. In other words, the long-tail of demand can be met by online platforms through technology.

Example 1.2. *The popular online platform eBay has the reach of over 100 million active users per quarter. However, antique sellers today may still rather pay a hefty 10-30% seller’s premium to have it auctioned at an Auction House, which professionally appraises, advertises and ship your item. They do so because the number of items sold on eBay also means that objects in such niche areas may not have the chance to meet the right buyer. eBay, on the other hand, thrives on small and newer items, which often have both ample supply and demand. Unlike traditional auctioneers, eBay can have multiple auctions occurring at the same time and for a more extended period.*

Example 1.3. *The matchmaking platform Tinder today have an average of about 4 million active users. Unlike its pre-technology predecessor, usually known as Singles Club, Tinder reduces (i) the barrier for information, (ii) the cost of interaction, and (iii) the cost of an actual physical location. Tinder can also easily segment markets by age ranges, distance, and sexual orientation. Recently, Tinder also introduced Tinder U, a new feature that makes it easier to connect with other students around you in 4-year, accredited, not-for-profit schools. Tinder U is a good example where access control, i.e., further limiting the matches on each side, can improve social welfare. In particular, Tinder U allows for safer interactions between students from the same college, and can also potentially serve a purpose beyond dating, e.g., study groups.*

Example 1.4. *Amazon, the most valuable company in the world, and coincidentally also a platform, started as a single-sided firm. It began in 1994 as a bookseller, buying books from publishers and reselling them to consumers. There were no connections between publishers and consumers, and no network effects. Even so, Amazon did manage to provide a vast selection with more than 2.5 million titles as of 1997. Consumers then already no longer needed to leave their homes and had access to more titles than a bookstore could offer. Shifting to a platform in 2000 with the launch of Amazon Marketplace, Amazon exercises an exact amount of control, imposing rules against bad behavior. One such way Amazon does this is through its Buybox, which is the default seller for an item. This default seller is picked based on item price, customer reviews, stock availability, and return policy.*

The rise of the platform economy brings with it a wide variety of engineering, economic and social challenges. Historically, markets have been slow to evolve and finding the “right” trading partners has been a daunting task. However, the integration of networks and information technology into marketplaces has led to complicated platforms that facilitate matches among participants. There is now an unprecedented level of control over the operation of these markets. Companies are engineering platforms to control the flow of information, recommend matches, and enforce prices and terms of trade. As such, the decisions made in the design of the platforms create complex and subtle interactions between computational constraints, network constraints, and market outcomes.

1.1 Online Platform Design

In a seminal work to understand two-sided markets, (Rochet and Tirole, 2003) built a model of platform competition with two-sided markets. In that work, they unveil the determinants of price allocation under a two-sided market model. They also compare the outcome from different governance structures, e.g., profit-maximizing versus not-for-profit organizations. This groundbreaking work contributed heavily towards Tirole's Nobel Memorial Prize in Economic Sciences in 2014.

Besides different governance structures, there are many other factors to consider in designing a successful platform, including matching, access control, pricing, trust and transaction costs. The online marketplace platform Amazon carefully studies each design component before becoming what it is today. In allowing buyers to purchase from any seller on the platform, Amazon displays its transparency and fairness to all sellers. However, because open access may incur substantial search costs, Amazon works around it by its Buybox, which is a product's default seller. Some guidelines that Amazon has for Buybox winners are professional selling accounts, good reputation and transaction history, refund flexibility and competitive prices. By doing so, Amazon limits access and promote competitive pricing, reducing search costs and building trust. About 80% of consumers on Amazon do not look beyond the default seller in Buybox.

To perform well on Amazon, a seller needs to win the Buybox. To win the Buybox, sellers have started employing an automated mechanism termed as *algorithmic pricing*, which set prices dynamically based on current conditions, e.g., competitor's selling price and inventory levels. For a single item, prices can change up to hundreds of times a day (L. Chen, Mislove, and Wilson, 2016). Since prices have to be valid for multiple consumers with varying valuations over a different subset of items, obtaining optimal pricing is hard (Guruswami et al., 2005), but good approximations to the problem are actively studied (Chawla, Hartline, and Kleinberg, 2007).

Today, driven by the desire to understand what makes them successful, there is an increasingly large literature on online platforms and multi-sided market design. These works consider individual design components, e.g., pricing by (L. Chen, Mislove, and Wilson, 2016), matching by (Akbarpour, S. Li, and Oveis Gharan, 2017), and access control by (Banerjee, Gollapudi, et al., 2017). Recent studies have also moved towards balancing trade-offs between conflicting designs. An example of this is (Dinerstein et al., 2018), which studies how guiding consumers to their most desired product often weaken sellers' incentives to offer low prices.

At the heart of platform design is the design of the matching algorithm that determines matches between firms (sellers) and consumers (buyers). Platforms today have a wide variety of approaches for matching. Some platforms, such as Etsy, Airbnb, and eBay follow an open access model: they provide information on all candidate matches, allowing the sellers and the buyers to make their own decisions freely. On the other extreme, platforms like Uber adopt a controlled allocation model: they provide no information about the candidate matches, only presenting a specific opportunity for a pairing. In between, there are discriminatory access platforms, such as Amazon, which impose constraints on the firms limiting which markets they can enter through Buybox, e.g., only sellers with low enough prices and excellent reviews are eligible to be shown in the Buybox, the default seller on a product's front page. Since about 80% of consumers do not look beyond the Buy Box, in effect, they are only accessible to the eligible sellers. Beyond these designs which apply widely to most platforms, there are also practical considerations specific to different industries and settings.

One such consideration is search cost. Search costs are one form of consumer transaction cost involved in searching for better alternatives. Without a thorough understanding of the trade-offs between the objectives mentioned above with considerations like search costs, and carefully balancing them, one can end up in a situation where these costs overwhelm consumers. For example, an open access platform design may not be a good idea for ridesharing platforms, since the number of producers (e.g., drivers) in such platforms is significant, resulting in substantial search costs for consumers. By contrast, successful ride-sharing platforms show only prices and do not offer alternatives, eliminating any search cost involved.

Another practical consideration of interest is network constraints, e.g., transmission line flow constraints in electricity markets. The electricity market is a classic example of a networked marketplace, where options available to individual participants are varied, complex and highly constrained. At peak levels, demand on the electricity market compels the use of generators with a more considerable marginal cost of production. To counter this, mature networked constrained systems such as the electricity market are usually already equipped with mechanisms, such as demand response, that help ensure stability or reduce the supply-demand imbalance. Since the marginal generating unit determines clearing price for all load, these mechanisms allow for a more stable and reliable network, and can also serve as a check against the exercise of market power by generators (Walawalkar et al., 2010).

Another critical component of online platform design is information. A single signal, e.g., locational marginal prices in electricity markets or surge price multipliers in ridesharing platforms, condenses the information in these massive platforms. The underlying information structure of the platform determines practical considerations such as anticipation and decision making under uncertainty. Platforms often make allocation decisions on behalf of the producing participants, but the anticipation of the platform allocations may change the competition entirely and affect the resulting efficiency. Also, sound decisions are often not made under full information, and participants often have to make decisions sequentially under tight inventory constraints. Usually, an optimal one-shot choice can be very different from a competitive decision in uncertain markets over some time.

Example 1.5. *Independent system operators (ISO) serve as a platform over the multi-sided market for electricity. Without electricity, none of the previously mentioned platform examples can function. Above and beyond attaining conditions for platform success, e.g., reaching critical mass, the ISO has to first and foremost maintain safety and security of the power network. Constraints are governing how electricity flows over the power network, e.g., Kirchhoff's Laws provide governing equations over the current and voltages across the grid. While conventional generators (or producers) can overwork beyond their set-point (frequency), significant (frequency) deviations can lead to automatic generator shutdowns, which then causes additional strain on the network. The stress on the system may consequently cause a different generator to be overworked, and cascading failure of the power network may occur, leading to large-scale blackouts.*

Typically, to ensure that the network remains stable, the ISO collect producer bids and computes locational marginal prices based on demand forecast. The locational marginal prices are often calculated based on a social welfare maximization problem constrained on physical laws governing the power network. Both consumers and generators on the system pay and paid based on this locational marginal price.

However, with large fluctuating loads and renewable generation, net demand forecast may not always be accurate. To curb with this shortfall or oversupply, ISOs design mechanisms such as demand response, where consumers can reduce their load in exchange for certain rewards. These consumers often have a limited battery capacity to utilize or flexible loads that they can shift and they may want to use them strategically to obtain maximal rewards.

1.2 Design Challenges

There are multiple challenges for an online platform designer both in starting a new platform and improving a current one. In this thesis, we consider three such problems that help shed light on recent transformation and changes to multisided platforms. Our first consideration concerns the trade-off between transparency and control in platforms, with a focus on access and allocation control. The second challenge is to analyze the negative impact of practical physical network constraints, anticipation and manipulation under networked competition, and then design mechanisms to prevent them. The final challenge for platform design we tackle is to understand how individuals with inventory constraints sell under uncertainty or incomplete information. We dive deeper into each of the problems here.

Access and Allocation Control for Online Platform Designs

Our first focus is on access and allocation control designs with regards to tradeoffs between transparency and control. Platforms have very diverse access and allocation control designs even within the same industry. For example, within ride-sharing platforms, Didi requires confirmation of pick-up and drop-off locations before allowing drivers to pick rides in an *open access* design while Uber assigns or *allocates* each trip to the nearest driver. Some marketplaces like eBay have an *open access* design and seek to lower producer entry cost and increase competition, but may end up suffering from high consumer search cost. On the other hand, Amazon combines the scale of a populated marketplace with carefully differentiated access through their Buybox designs in a *discriminatory access* design. In each example, both platforms have similar aims but apply very different models to get there.

This motivates a variety of critical open questions with regards to the design of access and allocation control. Firstly, what is the worst case efficiency loss of platforms under these designs? In other words, how do selfish decisions propagate under each platform design? Secondly, what is the impact of allocation control, and further considering recent collusion and manipulation in ridesharing platforms, are there any unintended incentives for strategic behavior that might decrease market efficiency? Lastly, is there a sweet spot between open access and controlling allocations that balances transparency and control, and limits efficiency loss? Are there any limitations or difficulties with a balanced design?

Practical Considerations for Online Platform Design

In the first part, we show that practical considerations such as search cost can be critical in online platform designs, e.g., open access designs for ridesharing might cause user dissatisfaction and cause a reduction of demand on the system. Anticipation is another consideration that is an essential tool that can instead be abused by producers to improve their profit and hurt the system as a whole. As an example of such manipulation, the work of (Ruhi et al., 2018) has shown that strategic curtailment by aggregators for renewables can improve profits by causing an increase in locational marginal prices. Whenever possible, market manipulation takes place and are especially prevalent in networked constrained marketplaces, such as the electricity market.

Platform designs that are robust to anticipation and manipulation are critical. To that end, we answer some crucial questions in the second part of this thesis. Firstly, how does network constraints over allocations affect the efficiency of an online platform? Secondly, what is the impact of producer anticipation on the efficiency of the platform and the stability of the network? Lastly, are there any economic implications to mechanisms that are already used to upkeep the physical stability and security of the system?

Competitive Participation in Real-time Uncertain Markets

Networked markets are often large and complex. As such, individual firms may have to make decisions under uncertainty or without full information, in pursuit of some form of optimality. It is critical first to learn how an individual makes decisions under uncertainty before trying to understand the entire system under uncertainty. Examples of these uncertainties include electricity markets where generators may not know the bid of other generators or advertisement companies who may not know what future demand of advertisement.

We consider the problem of participation in real-time uncertain markets under an inventory constraint. The main question in this challenge is the following: Is there an optimal (in terms of competitive ratio) way to participate in real-time uncertain markets? We find that this problem is a variant of the classical one-way trading problem. Consequently, we aim to design an algorithm that generalizes to the original problem with price elasticity.

1.3 Main Contributions and Overview

In this thesis, we aim to understand online platforms better. Our first result focuses on the trade-offs between transparency and control in online platform designs. Next, we look into practical platform designs, taking into consideration network constraints, anticipation, and the impact of existing mechanisms such as demand response. Another practical platform design consideration is that users make decisions under uncertainty, and may seek to do so in an optimal way. To that end, we study how producers make selling decisions under uncertainty, e.g., how drivers decide when to drive in a day. We make three main contributions in this thesis, summarized here.

In Chapter 3, we first analyze the worst case efficiency loss of online platform designs under a networked Cournot competition model. Inspired by some of the largest platforms today, the platform designs considered are three different ways to balance the trade-off between transparency and control. They are (i) open access, (ii) controlled allocation, and (iii) discriminatory access designs. Our results show that open access designs incentivize increased production towards perfectly competitive levels and limit efficiency loss, while controlled allocation designs lead to producer-platform incentive misalignment, resulting in low participation and unbounded efficiency loss. We show that discriminatory access designs strike a balance between transparency and control, and achieve the best of both worlds, maintaining high participation rates while limiting efficiency loss. We also study a model of consumer search cost to include this practical consideration which further distinguishes the three designs described.

In Chapter 4, we include other practical considerations in platform design such as network constraints and anticipation, and study the economic impact of demand management on strategic networked competition with producer anticipation. We consider these in a networked Stackelberg model that bears similarity to an electricity market and is a leader-follower version of the networked Cournot model where firms submit quantity bids and the platform balances supply and demand over the network subject to network constraints in a socially optimal manner. We also show that efficient anticipatory competition in the networked Stackelberg model is fragile, and known conditions, e.g., homogeneous price intercepts, are often impractical. Our main result is that demand response can play an additional role economically, and in particular, (i) have bounded efficiency loss in the absence of network constraints, (ii) is not often binding at equilibrium, and (iii) binding network constraints do not significantly worsen the efficiency of the system.

In Chapter 5, we study how individual producers make selling decision under uncertainty and inventory constraints. A practical example relating to platforms is how drivers decide how much and when to participate in ridesharing apps in a day. We formulate an online optimization problem with inventory constraints, where a producer has a fixed amount of inventory to sell, and irrevocable decisions are necessary for revenue functions which appear sequentially. While online optimization is a well-studied topic, versions with inventory constraints have proven to be more complicated. We consider a formulation of inventory-constrained optimization that is a generalization of the classic one-way trading problem and has a wide range of applications. We present a new algorithmic framework, CR-Pursuit, and prove that it achieves the optimal competitive ratio among all deterministic algorithms (up to a problem-dependent constant factor) for inventory-constrained online optimization. Our algorithm and its analysis not only simplify and unify the state-of-the-art results for the standard one-way trading problem, but they also establish novel bounds for generalizations including concave revenue functions. For example, for one-way trading with convex price elasticity, which corresponds to concave inverse demand functions, CR-Pursuit achieves a competitive ratio within a small additive constant (i.e., $1/3$) to the lower bound of $\ln \theta + 1$, where θ is the ratio between the maximum and minimum base prices.

In Chapter 6, we also include lessons and insights from each section that serves as a quick overview and conclusion of the thesis. In each of these chapters, we begin with a more detailed overview of the contributions and a review of the relevant literature. We append proofs of our results after a conclusion to this thesis to encourage a smoother read.

Our analysis throughout this thesis builds on the networked Cournot competition model, which we first will introduce in Chapter 2. We include the history and work related to the model, prior known results on existence and uniqueness of equilibria, equilibrium computational costs, and other critical relevant concepts.

*Chapter 2*NETWORKED COURNOT COMPETITION FOR ONLINE
PLATFORMS

Having observed competition in a spring water duopoly, Antoine Augustin Cournot first outlined his theory on quantity competition in his 1838 volume *Recherches sur les Principes Mathematiques de la Theorie des Richesses*. Cournot theorized that production quantity affects prices in a simple way, i.e., the more the aggregate production, the lower the price.

“The cheaper an article is, the greater ordinarily is the demand for it. The sales or the demand (for to us these two words are synonymous, and we do not see for what reason theory need take account of any demand which does not result in a sale)—the sales or the demand generally, we say, increases when the price decreases.” (Cournot, 1838)¹

Cournot’s theory for quantity competition yielded prices and quantities that are between monopolistic and competitive levels. Additionally, equilibrium prices are larger than marginal costs. It has since been generalized to consider multiple firms, and is known to retain similar results.

Almost half a decade later, Joseph Bertrand developed a model for competition based on price instead of quantity. In that model, producers report a selling price that they are willing to commit to, and the market share goes to the producer with the lowest selling price. If the market is not monopolized by a single producer, the equilibrium market price is equivalent to marginal cost. However, the Bertrand model for competition was also met with its own set of criticism. For example, the “winner-takes-all” model is unable to account for capacity constraints, which are relevant to producers that may not benefit from economics of scale.

Neither model is necessarily “better” than the other, but perhaps a model can be more appropriate to use in a particular modeling of competition. They are best used when applied to the right industry, Cournot when firms choose quantities, e.g., oil and gas markets, electricity markets, and Bertrand when firms choose prices, e.g., electronics market.

¹A translation by Nathaniel T. Bacon.

The Bertrand and Cournot competition models have both been well-used for models of competition for electricity markets, e.g., (Bunn and Oliveira, 2003; Hobbs, 2001; Oren, 1997), and energy markets, e.g., (Salant, 1982; Golombek, Gjelsvik, and Rosendahl, 1995).

Recently, with globalization, increased connectivity and technological breakthroughs, networked generalizations of both the Cournot competition (Abolhassani et al., 2014; Bimpikis, Ehsani, and Ilkic, 2014) and Bertrand competition (Chawla and Roughgarden, 2008; Guzmán, 2011; Anshelevich and Sekar, 2015) have also been introduced to meet the needs of an increasingly networked marketplace setting.

In these networked models, there is more than one market, and firms are connected via a bipartite graph to a subset of the markets. Besides these models, there are also networked generalizations of bargaining games where agents can trade via bilateral contracts over a network that determines the set of feasible trades, e.g., (Abreu and Manea, 2012; Elliott, 2015; Nava, 2015; Nguyen, 2015).

The networked Cournot competition model inherits many of the nice equilibrium results from its vanilla counterpart, e.g., if costs are convex and the inverse demand function is concave, then a unique equilibrium exists. The required computation time to find these equilibria is reported in (Abolhassani et al., 2014), and is also included in this section. The equilibrium quantities under this model is also shown to be related to network centrality (Ilkic, 2009).

Meanwhile, the networked Cournot model also have applications in many areas, e.g., electricity markets in (Bose et al., 2014) and demand-side management in smart grids, a.k.a. demand response in (Motaleb et al., 2017).

2.1 The Networked Cournot Competition Model

We describe the competition in online platforms using the *networked Cournot competition* model introduced in (Abolhassani et al., 2014) and (Bimpikis, Ehsani, and Ilkic, 2014). As a generalization of the classical model of Cournot competition, the networked Cournot model captures settings in which firms compete to produce a homogeneous good in *multiple markets*, where each market is potentially only accessible by a subset of firms. Firm's decisions are connected across markets through an aggregate (convex) production cost while decisions across firms are connected via a single price from the markets, depending on the aggregate demand. We formally develop this model in the following subsections.

Network Model and its Variants

The network specifying the connections between firms and markets is described according to a directed bipartite graph (F, M, \mathcal{E}) . Here, we denote by $F := \{1, \dots, n\}$ the set of n firms, $M := \{1, \dots, m\}$ the set of m markets, and $\mathcal{E} \subseteq F \times M$ the set of directed edges connecting firms to markets where $(i, j) \in \mathcal{E}$ if and only if firm i has access to market j .

An alternative specification more relevant to electricity markets is described according to a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and each node $i \in \mathcal{V}$ in the graph correspond to a co-located firm and market. In the electricity market example, the graph then constrains the reallocations made by the independent system operator to optimize social welfare. This alternative will be explored in Chapter 4. There, we consider an anticipatory setting where firms know how the platform reallocates, and a non-anticipatory setting where the platform competes with a reallocation.

In general, the efficiency of such marketplaces depend on the structure of the underlying graph, which restricts the set of markets to which each firm has access. A crucial role that the platform might, therefore, play in this setting is the selection of markets that are made available to each firm. An alternative would be to dictate the amount passing through each edge, essentially deciding the allocations for each firm's production. In Chapter 3, we examine three important classes of platform designs: *open access* platforms, *controlled allocation* platforms, and *discriminatory access* platforms. In Chapter 4, we further consider platforms with and without anticipation.

Producer Model

Each firm $i \in F$ decides its producing quantity $q_{ij} \in \mathbb{R}_+$ to each market $j \in M$, where $q_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. We let $q_i := (q_{i1}, \dots, q_{im}) \in \mathbb{R}_+^m$ denote the supply profile from firm i , and denote this feasible set of quantities for firm i as $Q_i(\mathcal{E})$, and $Q(\mathcal{E}) = \prod_{i \in F} Q_i(\mathcal{E})$. We let s_i be the aggregate production of firm $i \in F$, given by

$$s_i := \sum_{j=1}^m q_{ij}. \quad (2.1)$$

The resulting production cost of firm i is defined by $C_i(s_i)$. We assume that the cost function C_i is convex, differentiable on $(0, \infty)$, and satisfies $C_i(x) = 0$ for all $x \leq 0$. Since the cost is convex, decisions across markets are coupled since the marginal cost for each firm is constant across markets. Finally, we define $C := (C_1, \dots, C_n)$ as the cost function profile.

Market Model

We model price formation according to an inverse demand function in each market. Similar to (Bimpikis, Ehsani, and Ilkic, 2014), we restrict our attention to linear inverse demand functions throughout this thesis. Specifically, the price in each market $j \in M$ is determined according to $p_j(d_j) := \alpha_j - \beta_j d_j$, where d_j denotes the aggregate quantity supplied to market j , given by

$$d_j := \sum_{i=1}^n q_{ij}. \quad (2.2)$$

Here, $\alpha_j > 0$ measures consumers' maximum willingness to pay, and $\beta_j > 0$ measures the price elasticity of demand.

Social Welfare

We measure the performance (or efficiency) of a platform according to *social welfare*. For platforms, the pursuit of social welfare benefits both buyers and sellers, and in the long run, promotes their expansion. For example, Amazon (in its Buybox design) believes that welfare measures such as availability, fulfillment, and customer service ultimately lead to increased customer satisfaction, and thereby, promote its growth in the long run (L. Chen and Wilson, 2017).

We adopt the standard notion of social welfare defined as aggregate consumer utility less total production cost. Specifically, the social welfare associated with a supply profile q and a cost function profile C is defined according to

$$\text{SW}(q, C) := \sum_{j=1}^m \int_0^{d_j} p_j(z) dz - \sum_{i=1}^n C_i(s_i), \quad (2.3)$$

where s_i and d_j are defined in Eqs. (1) and (2), respectively. We define the *efficient social welfare* associated with the cost function profile C and an edge set \mathcal{E} as

$$\text{SW}^*(\mathcal{E}, C) := \sup_{q \in Q(\mathcal{E})} \text{SW}(q, C). \quad (2.4)$$

It is straightforward to check that the above supremum can be attained, and that the set of efficient supply profiles is non-empty. The *optimal social welfare* associated with the cost function profile C is the efficient social welfare at the edge set corresponding to the complete bipartite graph $F \times M$. One can show that the optimal supply profile can be attained if a social welfare maximizing platform controls both production of firms and their allocation to the different markets, while losing control of either may lead to inefficiency.

The platform designs and practical considerations in Chapters 3 and 4 aim to maximize the social welfare at the equilibrium of the resulting game. In particular, the metric of performance we consider is the *price of anarchy*, which is the worst case (multiplicative factor) efficiency loss due to selfish behavior.

Before formally defining the price of anarchy, we first define the networked Cournot game, as defined by (Abolhassani et al., 2014; Bimpikis, Ehsani, and Ilkilic, 2014). The platform designs corresponding to access control directly applies to the networked Cournot game in that the open access design corresponds to the full bipartite graph $F \times M$ while discriminatory access correspond to the bipartite graph maximizing equilibrium social welfare.

Networked Cournot Competition

With all the economic models in hand, we now describe the equilibrium of the market specified for the open access and discriminatory access designs according to Nash, where no producers can deviate to obtain a better outcome. We consider profit maximizing firms, where the profit π_i of a firm i , given the supply profiles of all other firms $q_{-i} = (q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_n)$, is given by

$$\pi_i(q_i, q_{-i}) := \sum_{j=1}^m q_{ij} p_j(d_j) - C_i(s_i), \quad (2.5)$$

where p_j is a price formed based on the total demand at node j , coupling the decisions made between different producers. We denote by $\pi := (\pi_1, \dots, \pi_n)$ the collection of payoff functions of all firms. The triple $(F, Q(\mathcal{E}), \pi)$ defines a normal-form game applicable for the open access and discriminatory access designs, which we refer to as the *networked Cournot game* associated with the edge set \mathcal{E} . Its Nash equilibrium is defined as follows.

Definition 2.1. A supply profile $q \in Q(\mathcal{E})$ constitutes a pure strategy Nash equilibrium of the game $(F, Q(\mathcal{E}), \pi)$ if for every firm $i \in F$,

$$\pi_i(q_i, q_{-i}) \geq \pi_i(\bar{q}_i, q_{-i}), \text{ for all } \bar{q}_i \in Q_i(\mathcal{E}).$$

Under the assumptions of convex cost functions and affine inverse demand functions, (Abolhassani et al., 2014) has shown that the networked Cournot game is an ordinal potential game. Additionally, it admits a unique Nash equilibrium that is the unique optimal solution to a convex program. We summarize the most related result in the following lemma, and provide time complexity results in Table 2.1.

Lemma 2.2. (Abolhassani et al., 2014) *The game $(F, Q(\mathcal{E}), \pi)$ admits a unique Nash equilibrium $q^{NE}(\mathcal{E})$ that is the unique optimal solution to the following convex program:*

$$\underset{q \in Q(\mathcal{E})}{\text{maximize}} \quad \text{SW}(q, C) - \sum_{i=1}^n \sum_{j=1}^m \frac{\beta_j q_{ij}^2}{2}. \quad (2.6)$$

Table 2.1: Equilibrium Computation Complexity (Abolhassani et al., 2014)

Cost Functions	Inverse Demand Function	Time Complexity	Technique
Convex	Linear	$O(E^3)$	Convex Optimization, Ordinal Potential Game
Convex	Strongly Monotone Marginal Revenue Function	$\text{Poly}(E)$	Reduction to Linear Complementary Program
Convex, Separable	Concave	$O(n \log^2 Q_{max})$	Supermodular Optimization

In general, the supply profile at the unique Nash equilibrium of the networked Cournot game differs from the efficient supply profile. This efficiency loss is commonly also known as the *price of anarchy* of the game first introduced in (Koutsoupias and Papadimitriou, 1999). In the rest of this thesis, we use *worst case efficiency loss* and *price of anarchy* interchangeably.

Definition 2.3. *The price of anarchy associated with the edge set \mathcal{E} , the cost function profile C , and the corresponding networked Cournot game $(F, Q(\mathcal{E}), \pi)$ is defined as*

$$\rho(\mathcal{E}, C) := \frac{\text{SW}^*(\mathcal{E}^C, C)}{\text{SW}(q^{NE}(\mathcal{E}), C)},$$

where \mathcal{E}^C refers to the complete bipartite graph and we set $\rho(\mathcal{E}^C, C) = 1$ if $\text{SW}^*(\mathcal{E}, C) = 0$ and $\text{SW}(q^{NE}(\mathcal{E}), C) = 0$. Note here that comparisons are always made against the optimal social welfare, i.e., the optimal social welfare attained with the complete bipartite graph, and control over the productions and allocations of supply.

In general, for games with possible multiplicity of Nash equilibria, the price of anarchy is defined as the ratio of the efficient social welfare over that of the Nash equilibrium with the *worst* social welfare, and is always greater than 1. The price of anarchy is defined similarly for Nash equilibrium in Stackelberg games, or Stackelberg equilibrium.

Insights on Supply Paths, Mergers and Expansions

Beyond existence and uniqueness, the analysis from (Ilklic, 2009) and (Bimpikis, Ehsani, and Ilklic, 2014) also identify connections between equilibrium outcomes and supply paths in the underlying network structure. Additionally, they study the impact of changes in competition structure (or graph structure) either due to firms merging or expanding into new markets, on their profits and consumer welfare.

First focusing on games with identical production technology and consequently same cost functions C , and markets with identical demand slopes β , a weighted adjacency matrix can be defined as follows:

$$w_{i_1 k_1, i_2 k_2} = \begin{cases} 2c & \text{if } i_1 = i_2, k_1 \neq k_2 \\ \beta & \text{if } i_1 \neq i_2, k_1 = k_2 \\ 0 & \text{otherwise.} \end{cases}$$

This weighted adjacency matrix is not on the original graph of nodes, but instead defined on its corresponding line graph, where the links from the original graph are now nodes. Links exist between nodes on the line graph if two links are either incident on a firm or a market, and is respectively weighted $2c$ or β . One way to understand the non-zero entries of W is that the values represent the change in marginal profit corresponding to a firm-market pair that results from an infinitesimal increase in the quantity corresponding to another firm-market pair.

Given the matrix W , we restate the following theorem characterizing the Nash equilibrium of such a game, which yields interesting insights on the relationship between the degree of strategic substitutability or complementarity between the actions of firms in two markets with the supply paths that connect them.

Theorem 2.4 (Theorem 1 (Bimpikis, Ehsani, and Ilklic, 2014)). *The unique Nash equilibrium of the symmetric game is given by*

$$\mathbf{q}^* = [I + \gamma W]^{-1} \gamma \bar{\alpha},$$

where $\gamma = (2(c + \beta))^{-1}$ and $\bar{\alpha}$ is a vector with the length number of edges, where each entry correspond to the α value relating to the market incident on that edge. Furthermore, if $\lambda_{\max}(\gamma W) < 1$,

$$\mathbf{q}^* = \left[\sum_{k=0}^{\infty} (\gamma W)^{2k} - \sum_{k=0}^{\infty} (\gamma W)^{2k+1} \right] \gamma \bar{\alpha}.$$

Informally, we obtain the “*the enemy of my enemy is my friend*” intuition. More formally, the even (odd) power-ed terms correspond to paths of even (odd) length. The matrix W can also be shown to be closely related to a centrality measure known as the *Katz-Bonacich centrality* of the nodes in the network, defined as follow.

Definition 2.5. *Given a weighted adjacency matrix W and a scalar ρ , the Katz-Bonacich centrality of the nodes in the network is defined as the following vector*

$$\mathbf{b}(W, \rho) = \sum_{t=0}^{\infty} (\rho W)^t \mathbf{1}.$$

It is then trivial to see that when markets have the same maximal willingness to pay α , we obtain the following corollary.

Corollary 2.6 (Corollary 1 (Bimpikis, Ehsani, and Ilkilic, 2014)). *Suppose that $\lambda_{\max}(\gamma W) < 1$, then the unique Nash equilibrium profile of the symmetric game is given by*

$$\mathbf{q}^* = \mathbf{b}(W, -\gamma)\gamma\alpha.$$

Beyond equilibrium analysis, the matrix W also have a role in the effects of changes in the structure of competition among firms on their profits and consumer welfare. The main focus in this section is on highlighting the role of the underlying network structure and the differences between traditional single market analysis to a setting with significant network effects.

Before proceeding on to the impact of the structure of competition, we first define the following $|\mathcal{E}| \times m$ matrix Λ , which can be understood as a measure of the firms’ market power.

$$\Lambda_{ik,l} = -\beta \sum_{j \in M_l} \frac{\psi_{jl,ik}}{\psi_{ik,ik}}, \quad \forall (i, k) \in \mathcal{E} \text{ and } M_l \in M,$$

where $\Psi = [I + \gamma W]^{-1}$. Essentially, entry (ik, l) of matrix Λ is the change in market m_l ’s price that results from a marginal increase in the production that firm f_i supplies to market m_k . This allows a description of how firms react to quantity shocks from other firms.

Proposition 2.7 (Proposition 1 (Bimpikis, Ehsani, and Ilkilic, 2014)). *Consider an exogeneous shock dq_{ik} . Other firms adjust their quantities according to*

$$dq_{il} = \frac{\psi_{ik,jl}}{\psi_{ik,ik}} dq_{ik}.$$

This means that a change in one firm-market quantity can have ripple effects on the entire network. Since Ψ is a symmetric, positive semidefinite matrix with positive entries, firms i and j view their actions in markets k and l respectively as strategic complements or substitutes depending on the sign of $\psi_{ik,jl}$.

Beyond understanding these equilibria and supply paths, (Bimpikis, Ehsani, and Ilkilic, 2014) also considers expansions into new markets and horizontal mergers. In this thesis, we will instead proceed with understanding platform designs with regards to access and allocation control.

*Chapter 3*ACCESS AND ALLOCATION CONTROL DESIGNS FOR
ONLINE PLATFORMS

In this chapter, we aim to study the trade-offs between transparency and control under the networked Cournot competition model. We model and analyze three different platform designs, (i) open access, (ii) controlled allocation, and (iii) discriminatory access designs, and provide worst case efficiency loss bounds for each of them. We study these designs of access and allocation control using the model of networked Cournot competition, introduced in Chapter 2. Our analysis provides new results on efficiency loss and illustrates the trade-off between transparency and control in the design of online platforms. We further distinguish the three designs through a model of search costs.

The first platform design we consider is the open access design (Section 3.3), exemplified by online marketplaces like eBay and Etsy. The advantages of an open access platform design are its fairness, transparency and low producer entry cost, and the “increased market competitiveness” driven by such attributes. This design also increases competition among producers, which usually lead to large sales volume and lower prices for consumers. We show that under the networked Cournot competition model, open access platform designs preserve a large proportion of the optimal demand fulfilled in each market. This participation plays a significant role in limiting efficiency loss in the system where we prove a $3/2$ worst case multiplicative guarantee (Theorem 3.5).

An integral piece of that result is a novel argument that *the convex cost functions that maximize efficiency loss are linear ones*, a generalization of the results from (R. Johari and J. N. Tsitsiklis, 2005) for single market Cournot competition. This fact also allows us to obtain worst case efficiency loss bounds as a function of the degree of asymmetry in producers’ cost functions. Together, these results show that open access promotes sales volume, has a small efficiency loss when firms have symmetric cost, and retains a large proportion of the system’s optimal social welfare. Yet, we do not frequently see open access platform designs in online platforms today. Instead, designs which control the flow of each quantity supplied is increasingly utilized in new platforms.

Given the efficiency loss in open access designs, one may intuitively posit that the efficiency loss in open access platforms is due to poor allocation of supply to consumers. This intuition is one of the many motivating factors towards considering a controlled allocation platform (Section 3.4). These platforms control the allocation over each transaction, to carefully allocate supply from all producers to all consumers. Many modern platforms use designs of these forms, e.g., drivers in ride-sharing platforms like Uber decide when to drive but not who to fetch. Given the prominence of these designs, one might expect them to be efficient.

In contrast, we show that *the worst case efficiency loss can grow linearly in the number of markets even when the controlled allocation platform optimizes for social welfare* (Theorem 3.11). This highlights that controlled allocation platforms may end up being inefficient in spite of their good intentions. Additionally, we show that the controlled allocation platform under the networked Cournot competition model results in a networked Stackelberg game¹, where neither the existence nor uniqueness of Nash equilibrium can be guaranteed. We conclude Section 3.4 with *a general result (Theorem 3.14) on the worst case efficiency loss under different market clearing mechanisms*. Specifically, we consider the setting in which the platform is allowed to choose its market clearing mechanism via the choice of its objective function used. We also show that any such function that is a convex combination of consumer surplus and social welfare has the same unbounded worst case efficiency loss result while allocations that are more aligned to individual producers (such as revenue maximizing allocations) result in a smaller worst case efficiency loss.

The efficiency loss under controlled allocation designs suggests that the platform's allocation incentivizes decreasing production levels to a point where the optimal allocation of produced good cannot offset the efficiency loss associated with reduced production. On the other hand, the result on incentivizing production in open access platforms motivates us to carefully restrict producer-consumer pairs in a way that strikes a balance between transparency and control to improve the allocation while maintaining large proportions of demand fulfilled. Such a design is similar to what Amazon's Buybox attempts to do by first presenting a seller determined by Amazon based on various performance measures², and then presenting the possibility of looking at all possible sellers.

¹We study the setting with constraints over transportation or rebalancing in more detail in Chapter 4.

²The performance measures include competitive pricing, excellent reviews, flexible return policies, and short response time.

Motivated by this, we propose a new platform design called discriminatory access which optimizes over the edge set to increase the social welfare at the Nash equilibrium of the resulting networked Cournot competition. We show that while the guaranteed preserved proportion of optimal demand fulfilled may not be as large as the transparent open access platforms, the improved allocations as a result of controlling the network allows for *an improved 3/4 worst case guarantee on social welfare retention for discriminatory access platforms* (Theorem 3.15). Similar to the open access platform results, we further explore the impact of cost asymmetry on the worst case efficiency loss. Specifically, when the production cost functions of firms are close to identical, it would be desirable to choose an open access design which maximizes the extent of competition between firms. On the other hand, if firms have substantially different cost functions, it would be desirable to choose a discriminatory access design, in which the platform picks out firms with lower production costs.

A careful design over the connections in the discriminatory access platform is also often not trivial, and we show that the optimal network design problem can be written as a mathematical program with equilibrium constraints, and can be difficult to solve both analytically and computationally. However, under the restriction to linear cost functions—which were shown previously to yield the largest loss of efficiency in open access platforms—we *propose and prove the optimality of a greedy algorithm that yields an optimal worst case network design* (Theorem 3.17).

One factor for platform design that we have ignored to this point is *search cost*. Search costs are one form of consumer transaction cost involved in searching for better alternatives. Without a careful understanding of the trade-offs between the objectives as mentioned above with search cost and carefully balancing them, one can end up in a situation where search costs overwhelm consumers. For example, an open access platform design for a ride-sharing platform is not ideal, since there are many drivers, which results in substantial consumer search costs. On the other hand, ride-sharing platforms such as Uber eliminate any search cost involved. To understand the impact of search costs, we introduce a simple search cost model in Section 3.6 that yields important contrasts. In particular, since welfare losses from search costs can outweigh the benefits from open access, the worst case loss of open access platforms may no longer be bounded. On the other hand, the discriminatory access platform design remains efficient in the face of search costs and can be further optimized to balance search costs with the efficiency of the matching.

3.1 Literature Review

The recent growth of online platforms has led researchers to focus on identifying design features common to successful platforms. Earlier works in this area started by introducing different models of two-sided platform markets, e.g., (D. S. Evans and Schmalensee, 2016) modeled the cross-network externalities between the two or more groups while the renowned work of (Rochet and Tirole, 2003) considered a model whereby pricing structure affects volume of transaction. They posit that price structure matters through (i) transaction costs existing between different sides, and (ii) constraints on types of transaction costs imposed by the platform. Other definitions and models are found in many other work, such as (Hagiu and Wright, 2015; Rysman, 2009), and various platform designs are suggested in other works. We focus first on access and allocation control.

Open access is touted in (Boudreau, 2010) to increase competition among participants and shown in (Parker and Van Alstyne, 2017; Schor, 2016) to be often designed to cope with issues such as fairness and openness. The classical example studied for open access is the online marketplace eBay studied in (Chircu and Kauffman, 2001), which has been shown in (Gross and Acquisti, 2003; Hui et al., 2016) to depend much on reputation and regulation. Lastly, (Heylighen, 2006) shows that open access thrives on information symmetry, openness, and transparency.

Contrastly, an intuitive way of exert control in the platform is to assume that producers make socially inefficient decisions and instead decide allocations on their behalf, e.g., Uber as studied in (Hall, Kendrick, and Nosko, 2015). Dependent on how well the producers understand allocation decisions, (Rosenblat and Stark, 2016) showed there may be various outcomes. For example, it is well known and reported in (J. Y. Chen, 2017) that drivers on ride-sharing platforms collaborate and reduce their production to cause demand spikes in the system, often resulting in better individual payoff but worse performance for the platform. (Scheiber, 2017; A. Lu, Frazier, and Kislev, 2018; Z. Fang, Huang, and Wierman, 2018; Banerjee, Riquelme, and Ramesh Johari, 2015) look at other platform means, e.g., through subsidies and incentives, to make sure that drivers participate actively and on only one platform. (Banerjee, Freund, and Lykouris, 2016) showed an approximation algorithm to find the optimal allocation and prices to a variety of objective functions. Recent work by (Afeche, Liu, and Maglaras, 2018) characterize drivers' incentive compatibility conditions for repositioning decisions, and provide new insights on the interplay between admission control and drivers' strategic repositioning decisions.

More recently, (Ma, F. Fang, and Parkes, 2018) also considered setting prices that are smooth both in space and time in a spatio-temporal pricing mechanism.

Another possible way to improve on open access is through discriminatory access, where it is studied in (Banerjee, Gollapudi, et al., 2017; Kanoria and Saban, 2017; Akbarpour, S. Li, and Oveis Gharan, 2017) that platforms can restrict access between certain producers and markets in a bid to improve market outcomes. A significant amount of effort has been placed into what is known in the literature, e.g., (Chawla, Hartline, and Kleinberg, 2007; Chawla, Hartline, Malec, et al., 2010; L. Chen and Wilson, 2017), as *algorithmic pricing*, which is most well known to be exemplified by Amazon’s Buybox as studied in (L. Chen, Mislove, and Wilson, 2016), where they highlight one seller for each item to every consumer that is looking for it. Another example of discriminatory access is Airbnb’s Superhost program studied in (Liang et al., 2017), which highlights certain renters through badges and have been shown to vastly improve revenue on the peer-to-peer rental platform.

One of the starting objectives of online marketplaces or platforms as stated in (Bakos, 1997) was to reduce search costs — that is, to reduce the amount of effort to acquire information on sellers and the quality of their goods. (Diehl, Kornish, and Lynch Jr, 2003) showed that lower search costs leads to an increase in price sensitivity. However, it has also been shown in (Branco, Sun, and Villas-Boas, 2015) that too much information revealed in online platforms can instead result in increasing search costs from the consumers’ point of view. (Nishida and Remer, 2018) showed that reducing search costs can lead to higher prices and profit while (Gamp, 2016) showed that a monopolist may impede information acquisition to improve her profits. In a more recent work, (Dinerstein et al., 2018) studies the tradeoff between guiding consumers to their most desired product while also strengthening seller incentives to offer low prices.

Besides access and allocation control, work in this area has also covered a variety of possible design factors, including pricing in (Weyl, 2010), competition in (Armstrong, 2006), reputation in (Nosko and Tadelis, 2015; Tadelis, 2016; Luca, 2017), thickness in (Ashlagi et al., 2018), and also dynamic models for kidney exchange in (N. Agarwal et al., 2018). Recent empirical studies in (Einav et al., 2015; Bimpikis, Candogan, and Saban, 2016) also reveal significant price dispersion in online and spatial marketplaces, causing platforms to differentiate products in order to create distinct consumer markets as in (Dinerstein et al., 2018). In particular, these results highlight the need to study platforms in using models of *networked competition*.

3.2 Model and Preliminaries

In this section, we will concisely present the networked Cournot competition model and the platform designs considered. The reader is referred to Chapter 2 for further details and relevant literature on networked markets.

The network specifying the connections between firms and markets is described according to a directed bipartite graph (F, M, \mathcal{E}) , where $F := \{1, \dots, n\}$ denotes the set of n firms, $M := \{1, \dots, m\}$ the set of m markets, and $\mathcal{E} \subseteq F \times M$ the set of directed edges connecting firms to markets where $(i, j) \in \mathcal{E}$ if and only if firm i has access to market j .

Producer Model

Each firm $i \in F$ decides its producing quantity $q_{ij} \in \mathbb{R}_+$ to each market $j \in M$, where $q_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. We let $q_i := (q_{i1}, \dots, q_{im}) \in \mathbb{R}_+^m$ denote the supply profile from firm i , and denote this feasible set of quantities for firm i as $Q_i(\mathcal{E})$, and $Q(\mathcal{E}) = \prod_{i \in F} Q_i(\mathcal{E})$. We let s_i be the aggregate production of firm $i \in F$, given by

$$s_i := \sum_{j=1}^m q_{ij}. \quad (3.1)$$

The resulting production cost of firm i is defined by $C_i(s_i)$. We assume that the cost function C_i is convex, differentiable on $(0, \infty)$, and satisfies $C_i(x) = 0$ for all $x \leq 0$. Since the cost is convex, decisions across markets are coupled since the marginal cost for each firm is constant across markets. Finally, we define $C := (C_1, \dots, C_n)$ as the cost function profile.

Market Model

As is standard in Cournot models of competition, we model price formation according to a linear inverse demand function in each market. The price in each market $j \in M$ is determined according to $p_j(d_j) := \alpha_j - \beta_j d_j$, where d_j denotes the aggregate quantity supplied to market j , given by

$$d_j := \sum_{i=1}^n q_{ij}. \quad (3.2)$$

Here, $\alpha_j > 0$ measures consumers' maximum willingness to pay, and $\beta_j > 0$ measures the price elasticity of demand.

Social Welfare

We measure the performance (or efficiency) of a platform according to *social welfare*. We adopt the standard notion of social welfare defined as aggregate consumer utility less total production cost. Specifically, the social welfare associated with a supply profile q and a cost function profile C is defined according to

$$SW(q, C) := \sum_{j=1}^m \int_0^{d_j} p_j(z) dz - \sum_{i=1}^n C_i(s_i), \quad (3.3)$$

where s_i and d_j are defined in Eqs. 3.1 and 3.2, respectively. We define the *efficient social welfare* associated with the cost function profile C and an edge set \mathcal{E} as

$$SW^*(\mathcal{E}, C) := \sup_{q \in Q(\mathcal{E})} SW(q, C). \quad (3.4)$$

The *optimal social welfare* associated with the cost function profile C is the efficient social welfare at the edge set corresponding to the complete bipartite graph $F \times M$.

Networked Cournot Competition

We now describe the equilibrium of the market specified for the open access and discriminatory access designs according to Nash. We consider profit maximizing firms, where the profit π_i of a firm i , given the supply profiles of all other firms $q_{-i} = (q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_n)$, is given by

$$\pi_i(q_i, q_{-i}) := \sum_{j=1}^m q_{ij} p_j(d_j) - C_i(s_i), \quad (3.5)$$

where p_j is based on the demand at node j . We denote by $\pi := (\pi_1, \dots, \pi_n)$ the collection of payoff functions of all firms. The triple $(F, Q(\mathcal{E}), \pi)$ defines a normal-form game applicable for the open access and discriminatory access designs, which we refer to as the *networked Cournot game* associated with the edge set \mathcal{E} . Its Nash equilibrium is defined as follows.

Definition 3.1. A supply profile $q \in Q(\mathcal{E})$ constitutes a pure strategy Nash equilibrium of the game $(F, Q(\mathcal{E}), \pi)$ if for every firm $i \in F$,

$$\pi_i(q_i, q_{-i}) \geq \pi_i(\bar{q}_i, q_{-i}), \text{ for all } \bar{q}_i \in Q_i(\mathcal{E}).$$

Under the assumptions of convex cost functions and affine inverse demand functions, (Abolhassani et al., 2014) has shown that the networked Cournot game is an ordinal potential game. Additionally, it admits a unique Nash equilibrium that is the unique optimal solution to a convex program. We summarize the most related result in the following lemma.

Lemma 3.2. (Abolhassani et al., 2014) *The game $(F, \mathcal{Q}(\mathcal{E}), \pi)$ admits a unique Nash equilibrium $q^{NE}(\mathcal{E})$ that is the unique optimal solution to the following convex program:*

$$\underset{q \in \mathcal{Q}(\mathcal{E})}{\text{maximize}} \quad \text{SW}(q, C) - \sum_{i=1}^n \sum_{j=1}^m \frac{\beta_j q_{ij}^2}{2}. \quad (3.6)$$

In general, the supply profile at the unique Nash equilibrium differs from the efficient supply profile. This efficiency loss is commonly also known as the *price of anarchy* of the game first introduced in (Koutsoupias and Papadimitriou, 1999). In the rest of this thesis, we use *worst case efficiency loss* and *price of anarchy* interchangeably.

Definition 3.3. *The price of anarchy associated with the edge set \mathcal{E} , the cost function profile C , and the corresponding networked Cournot game $(F, \mathcal{Q}(\mathcal{E}), \pi)$ is defined as*

$$\rho(\mathcal{E}, C) := \frac{\text{SW}^*(\mathcal{E}^C, C)}{\text{SW}(q^{NE}(\mathcal{E}), C)},$$

where \mathcal{E}^C refers to the complete bipartite graph and we set $\rho(\mathcal{E}^C, C) = 1$ if $\text{SW}^*(\mathcal{E}, C) = 0$ and $\text{SW}(q^{NE}(\mathcal{E}), C) = 0$. Note here that comparisons are always made against the optimal social welfare, i.e., that attained with the complete bipartite graph.

In general, for games with possible multiplicity of Nash equilibria, the price of anarchy is defined as the ratio of the efficient social welfare over that of the Nash equilibrium with the *worst* social welfare, and is always greater than 1.

Open and Discriminatory Access Platforms

The open and discriminatory access platforms are choices of networks on the original networked Cournot competition model. The open access platform corresponds to the complete bipartite graph $F \times M$, while the discriminatory access platform seeks the optimal set of connections between firms and markets.

Controlled Allocation Platforms

In contrast, the controlled allocation platform is a platform with anticipated clearing mechanism. Instead of producing directly to connected markets, firms produce and allow this aggregating platform to allocate the aggregate supply through a pre-defined clearing mechanism.

In Section 3.4, we focus our attention on platforms maximizing social welfare, which keeps prices across markets constant³. The set of markets becomes equivalent to an “aggregated” inverse demand function that is piecewise linear, convex⁴.

Given an allocation function A , the profit of a firm i producing s_i can be written under the platform allocation A as

$$\pi(s_i, s_{-i}, A) = s_i p_j(A_j(\sum_i s_i)) - C_i(s_i), \quad (3.7)$$

where p_j is equivalent across markets with positive demand fulfilled under a social welfare maximizing controlled allocation platform. Similar to the networked Cournot game, we denote by $\pi := (\pi_1, \dots, \pi_n)$ the collection of payoff functions of all firms. Here, the feasible set of production for each producer is just \mathbb{R}_+ . The triple (F, A, π) defines a normal-form game applicable for the controlled allocation design, which we refer to as the *networked Stackelberg game*. The equilibrium, which we term a Stackelberg equilibrium, is defined as follows.

Definition 3.4. A supply profile $s^{SE} \in \mathbb{R}_+^n$ constitutes a Stackelberg equilibrium of the game (F, A, π) if for every firm $i \in F$, $\pi(s_i^{SE}, s_{-i}^{SE}) \geq \pi(s_i, s_{-i}^{SE})$, $\forall s_i \geq 0$.

3.3 Efficiency of Open Access

Open access platforms such as eBay and Etsy thrive on transparency and offer buyers/consumers access to all possible sellers. This platform design is usually accompanied by lower entry costs that typically improve competition between producers, thereby boosting transaction volume, and oftentimes leads to socially efficient markets by fairly presenting similar opportunities to every participant.

In this section, we aim to provide tight bounds on the loss of welfare due to selfish behavior in networked Cournot games under open access platform designs. The main theorem in this section provides a (tight) efficiency loss bound of a networked Cournot game governed by an open access platform with respect to the number of firms n .

³We also extend our analysis beyond social welfare clearing mechanisms, in which case prices across markets may not be equivalent, but we assume that the platform averages prices such that it remains equivalent for all producers

⁴Consequently, the inverse demand function is no longer necessarily concave. This means that previous existential and uniqueness results from (Abolhassani et al., 2014) no longer applies, and the controlled allocation platform can have multiple equilibrium or suffer market failure, i.e., no equilibrium exists.

Theorem 3.5. *The worst case efficiency loss associated with a cost function profile C and the corresponding open access networked Cournot game $(F, \mathcal{Q}(F \times M), \pi)$ is upper bounded by*

$$\rho(F \times M, C) \leq \frac{3}{2} \left(1 - \frac{1}{3n+6} \right).$$

The bound is tight if markets have the same maximal willingness to pay, i.e., $\alpha_1 = \alpha_2 = \dots = \alpha_m$. Precisely, for any choice of n , there exists a cost function profile \bar{C} , such that

$$\rho(\bar{C}) = \frac{3}{2} \left(1 - \frac{1}{3n+6} \right).$$

This result generalizes those in (R. Johari and J. N. Tsitsiklis, 2005) from a single market case to the networked setting, succinctly presented in the following corollary to Theorem 3.5.

Corollary 3.6. *Open access platforms have worst case efficiency loss of at most $3/2$.*

In the remainder of this section, we prove Theorem 3.5 first by proving a critical result in Lemma 3.8 that the cost functions maximizing the worst case efficiency loss are linear ones, a result presented in (R. Johari and J. N. Tsitsiklis, 2005) for the case without a network. That means that given a game with any set of convex cost functions, we can find a corresponding set of linear cost functions that increase the worst case efficiency loss. Applying Lemma 3.8, bounds on efficiency loss are developed for both the symmetric and asymmetric costs setting in Proposition 3.9 (Section 3.3) and Theorem 3.5 (Section 3.3) respectively.

The result in Theorem 3.5 (for firms with asymmetric costs) is potentially counter-intuitive since an increase in the number of firms n (presumably increasing competition) results instead in a larger worst case efficiency loss. To shed more light on this, we further develop worst case efficiency loss bounds on networked Cournot games under open access platforms with linear cost functions and bounds on asymmetry, presented later in Proposition 3.10. Additionally, in the symmetric cost case, in Proposition 3.9 we see that it is indeed that as the number of perfectly competing firms increase, the worst case efficiency loss decreases.

We present in the following lemma that demand fulfilled at the Nash equilibrium of the networked Cournot game for any market is at least half of the fulfilled demand under a socially optimal allocation.

Lemma 3.7. *For any market $j \in M$, the demand fulfilled at Nash is at least half of the demand fulfilled at the social optimal allocation.*

The lemma underlines one of the critical reasons open access platforms perform well—that open access platform designs preserve a large proportion of production. However, open access platforms still suffer from selfish behavior that manifest in the misalignment between optimal allocations and Nash allocations. Contrasting to this, in the next section, we explore an aggregator platform which allocates perfectly.

Identifying the Worst-case Cost Function Profile

The following lemma reveals that the worst case efficiency loss of networked Cournot games are maximized at a cost function profile consisting of cost functions that are linear over the non-negative real numbers, demonstrating that any cost function profile C as specified in Section 2.1 can be correspondingly re-designed to construct a piecewise linear cost function profile \bar{C} , resulting in an increased worst case efficiency loss. In light of this result, in designing efficiency loss bounds guaranteed to hold for all cost functions belonging to the family specified in Section 2.1, it suffices to consider cost functions that are linear on $(0, \infty)$.

Lemma 3.8. *Given a cost function profile C , define the cost function profile $\bar{C} = (\bar{C}_1, \dots, \bar{C}_n)$ according to*

$$\bar{C}_i(s_i) = \left(\partial^+ C_i \left(\sum_{j=1}^m q_{ij}^{NE}(F \times M) \right) \cdot s_i \right)^+$$

for $i = 1, \dots, n$, where $\partial^+ C_i$ denotes the right-derivative of the function C_i . It holds that $\rho(F \times M, C) \leq \rho(F \times M, \bar{C})$.

We provide some intuition here behind the proof of Lemma 3.8, while the full proof can be found in the Appendix. The social welfare achieved at a Nash equilibrium depends on the underlying cost function profile only at the marginal cost at Nash equilibrium, which follows from the fact that each firm's profit is a concave function of her own supply profile. By replacing the original convex cost functions by linear cost functions which have the same marginal cost at Nash equilibrium, the resulting Nash equilibrium has identical production quantities and a worse-off social welfare compared to the original one. A careful reader may notice here that the efficient social welfare may also decrease, but one can show that the decrease in the efficient social welfare is guaranteed to be larger than that of the social welfare at Nash equilibrium.

Efficiency Loss in Open Access Platforms

The characterization of the worst-case or *optimal* cost function profile in Lemma 3.8 is crucial to the derivation of tight upper bounds on the worst case efficiency loss bounds for networked Cournot games. One example of this simplification is that under linear cost function profiles, the networked Cournot game can essentially be decoupled into m many different Cournot games with a single market since the marginal cost of each firm remain constant. In what follows, we examine the role played by *(a)symmetry* in linear cost functions in determining platform efficiency.

Before proving Theorem 3.5, we prove a result for the case where firms have symmetric costs. This help us show that the intuition that increased competition does indeed lead to less efficiency loss is indeed correct given the right assumptions. In proving Theorem 3.5, we devise worst case asymmetric cost functions compared to the single one utilized in the optimal case. Intuitively, since the optimal solution would only utilize the firm with the smallest marginal cost, the worst case family of cost functions should consist of one firm with a low cost and the remaining with high costs. As we increase the number of firms in this scenario, the set with high costs increase, leading to a slight increase in competition which, however, is inefficient. Another way to approach the counter-intuitive result is that the additional firm can always be designed to have arbitrarily large marginal cost, such that it does not participate, and therefore maintains the efficiency loss without its participation.

The above worst case efficiency loss bounds exhibit distinct behaviors in terms of their monotonicity in the number of firms n . We thus further explore as to how the degree of asymmetry between firms' cost functions affect efficiency loss in open access platforms. In doing so, we characterize tight worst case efficiency loss bounds when the underlying cost functions are restricted to be linear over $(0, \infty)$, and whose slope lies within a bounded interval $[c_{\min}, c_{\max}] \subseteq \mathbb{R}_+$.

Symmetric Cost Functions

Our analysis begins with the setting where firms have identical cost functions, representing producers with similar technology. We establish a tight upper bound on the worst case efficiency loss under this setting in Proposition 3.9 that *monotonically decreases* in the number of firms, and converges to 1 as the number of producers grows large. This conforms with the intuition that increasing suppliers strictly increases competition, and thereby reduces the extent to which any producer might exert market power, improving efficiency as selfish behavior becomes restricted.

Proposition 3.9. *For firms with symmetric costs, i.e., $C_i = C_j, \forall i, j \in F$, the efficiency loss associated with the corresponding open access networked Cournot game $(F, Q(F \times M), \pi)$ is bounded above by*

$$\rho(F \times M, C) \leq 1 + \frac{1}{(n+1)^2 - 1}.$$

In addition, this bound is tight, i.e., for any choice of n , there exists a symmetric cost function profile with a corresponding worst case efficiency loss equal to the upper bound.

In proving Proposition 3.9, a direct application of Theorem 3.8 reveals that the worst-case symmetric cost function profile consists of n identical cost functions that are linear on $(0, \infty)$. The details of the proof are deferred to the Appendix.

Arbitrary Asymmetric Cost Functions

We now consider the more general setting in which firms have arbitrary asymmetric cost functions satisfying the assumptions in Section 2.1, summarized in Theorem 3.5, where we establish a tight upper bound on the efficiency loss that is *monotonically increasing* in the number of firms. In this case, it may seem counterintuitive that a perceived increased competition instead leads to efficiency loss.

This seemingly counterintuitive result can occur if an expensive firm enters the market. First, note that the entry of this new firm results in an increase in aggregate supply at Nash equilibrium, because of increased ‘competition’ in the market. However, its entry takes away production from its (cheap) competitors. This manifests in a reduction in social welfare if the increase in production cost exceeds the increase in consumer utility. Such a phenomenon is known as the “excess entry theorem” in the economics literature, e.g., (Suzumura and Kiyono, 1987; Mankiw and Whinston, 1986; Lahiri and Ono, 1988), and reveals the possibility that a new firm’s entry can lead to a reduction in social welfare.

Additionally, taking the number of firms $n \rightarrow \infty$ yields a worst case efficiency loss bound that is valid for any number of firms, and any number of markets, presented in Corollary 3.6. This recovers the 3/2 worst case efficiency loss or price of anarchy bound first established by (R. Johari and J. N. Tsitsiklis, 2005) for a single market. In fact, under this generalized characterization, one can obtain both the previous result in their work, and for a networked market of any size with open access platform design.

Linear Cost Functions with Bounds on Asymmetry

The efficiency loss results in Proposition 3.9 and Theorem 3.5 appear contradictory at first sight. Namely, the efficiency loss bound is decreasing in n if producers have symmetric cost functions but it is increasing in n if producers are allowed to have asymmetric cost functions. In what follows, we explore how the efficiency loss depends on the asymmetry between firms' cost functions, providing more intuition towards this counter-intuitive result, by restricting ourselves to cost functions that are linear on $(0, \infty)$, and whose slopes lie within $[c_{\min}, c_{\max}] \subseteq \mathbb{R}_+$.

$$\mathcal{L}(c_{\min}, c_{\max}) := \left\{ C_0 : \mathbb{R} \rightarrow \mathbb{R}_+ \mid C_0(x) = (cx)^+, c \in [c_{\min}, c_{\max}] \right\}.$$

We write $C \in \mathcal{L}^n(c_{\min}, c_{\max})$ if the cost function profile C satisfies $C_i \in \mathcal{L}(c_{\min}, c_{\max})$ for each firm $i \in F$. It will be convenient to define a non-dimensional parameter γ_j , which measures the degree of (a)symmetry between firms for each market $j \in M$. Specifically, for each market $j \in M$, define

$$\gamma_j := 1 - \frac{c_{\max} - c_{\min}}{\alpha_j - c_{\min}}.$$

It holds that $\gamma_j \in (-\infty, 1]$ if $c_{\min} < \alpha_j$. Clearly, γ_j is increasing in consumers' maximum willingness to pay α_j , and decreasing in the maximum cost c_{\max} . It follows that a value of γ_j close to one implies a small degree of asymmetry between firms' cost functions relative to consumers' maximum willingness to pay in market j .

The following proposition provides a tight bound on the worst case efficiency loss bound when firms have linear cost functions with a bounded degree of asymmetry.

Proposition 3.10. *Let $C \in \mathcal{L}^n(c_{\min}, c_{\max})$, and assume that $c_{\min} < \max_{j \in M} \alpha_j$ (that at least one firm would be willing to produce to one market). The worst case efficiency loss associated with the corresponding open access networked Cournot game $(F, \mathbf{Q}(F \times M), \pi)$ is upper bounded by*

$$\rho(F \times M, C) \leq \frac{\sum_{j=1}^m \frac{((\alpha_j - c_{\min})^+)^2}{\beta_j}}{\sum_{j=1}^m \left(\frac{2n+4}{3n+5} + \delta(\gamma_j, n) \right) \frac{((\alpha_j - c_{\min})^+)^2}{\beta_j}},$$

where the function $\delta(\gamma, n)$ is defined according to

$$\delta(\gamma, n) = \begin{cases} 0 & \text{if } \gamma < \frac{2n+3}{3n+5}, \\ \frac{(n-1)(3n+5)}{(n+1)^2} \left(\gamma - \frac{2n+3}{3n+5} \right)^2 & \text{otherwise.} \end{cases}$$

The bound is tight if $\alpha_1 = \alpha_2 = \dots = \alpha_m$.

The worst case efficiency loss bound specified in Proposition 3.10 depends on the degree of (a)symmetry between firms' cost functions only through the terms $\delta(\gamma_j, n)$ for $j = 1, \dots, m$. In particular, as $\delta(\gamma, n)$ is non-decreasing in γ , a reduction in the degree of asymmetry between firms' cost functions manifests in a reduction in the worst case efficiency loss bound.

With the positive worst case efficiency loss results from open access platform designs, together with the result in Lemma 3.7 that open access platform designs preserve a good proportion of demand fulfilled, we explore in the next section the platform design that controls allocations on behalf of the participating firms, allowing them to only decide their production quantity.

3.4 Inefficiency of Controlled Allocation Designs

Despite preserving large production rates, the open access platform design suffers from some efficiency loss through misaligned allocations. We study *controlled allocation platforms* that aggregate production and perfectly allocate to various markets in hope that this control circumvents the misaligned allocations. Unlike the open access platform described in Section 3.3, a controlled allocation platform allocates the production to different markets in a socially optimal way for any given aggregate production quantity.

However, controlled allocation platforms are not as simple to optimize as they seem. In particular, the control over allocations exerted by the platform may create unintended incentives for firms to withhold production, as empirically shown in (J. Y. Chen, 2017) for the ride-sharing platform Uber. Though firms cannot strategically choose prices or matches, they retain control over their participation in the platform. In the networked Stackelberg setting, this takes the form of strategic choices of production levels, with the platform only allocating on their behalf to maximize a certain quantity, e.g., social welfare.

Our main result in this section highlights that distorted incentives in controlled allocation platforms can lead to inefficient market outcomes.

Theorem 3.11. *A controlled allocation platform maximizing social welfare can have unbounded price of anarchy. In particular, there exists a family of networks where the Stackelberg equilibrium is unique and the worst case efficiency loss is $\Omega(m)$, i.e., there exists s^{SE} and cost functions C such that:*

$$\rho(C, M) \geq \Omega(m).$$

The contrast between Theorem 3.11 for controlled allocation platforms and Theorem 3.5 for open access platforms is stark. While open access platforms preserves at least $2/3$ of the optimal social welfare regardless of market parameters, controlling allocations can lead to reactive producer behavior which drives efficiency preservation rates possibly all the way to 0. In this section, we first prove Theorem 3.11. Additionally, we provide a generalization of Theorem 3.11 in Theorem 3.14, which shows that controlled allocation platforms remain inefficient for a variety of platform objective functions in clearing the market, spanning between social welfare and consumer surplus.

Social Welfare Maximizing Controlled Allocation Platforms

Given the results from Section 3.3, one may intuitively think that perfect allocation can lead to improved efficiency loss bounds. In Theorem 3.11, we show that such a platform, even with the best intentions, i.e., maximizing social welfare, can lead instead to an unbounded efficiency loss. On the other hand, allowing for an open access design and not controlling allocations lead to a bounded worst case efficiency loss as in Section 3.3.

Note that the networks constructed to prove Theorem 3.11 are simple. They use a single firm with costless production. The construction begins with a single market and then, as markets are added one by one to the system, the parameters of each new market are such that the firm has no incentive to increase production due to the reallocation under the controlled allocation, whereas the socially optimal (non-Nash) production level does increase, as does the optimal welfare. Note that, when a firm has no production cost, the socially optimal production is always to produce until the price in every market is driven to 0.

Theorem 3.11 is proven using a construction which highlights that an efficient platform should prevent firms from withholding production, which may be achieved by modifying the market clearing mechanism appropriately. This begs the question: Can the platform do better by optimizing a different quantity? We now show that the efficiency loss in controlled allocation platforms remain unbounded when the platform chooses a family of alternative objective functions in clearing the market. In particular, such a family of objective functions are given by convex combinations of social welfare and consumer surplus.

Consumer Surplus Maximizing Controlled Allocation Platforms

Before attending to the family of objectives, we first consider the corner case in consumer surplus, which helps give some intuition as to why the worst case efficiency loss can grow large. Given a fixed supply, the consumer surplus is defined as follows.

Definition 3.12. Specifically, the consumer surplus associated with a supply profile q is defined according to

$$CS(q) := \sum_{j=1}^m \int_0^{d_j} p_j(z) dz - d_j p_j(d_j), \quad (3.8)$$

where d_j is defined in Eq. (2).

In fact, the worst case efficiency loss is unbounded under the objective of consumer surplus. For two markets with different willingness to pay α and the same price elasticity to demand β , an equal quantity in each market produce the same consumer surplus. To some extent, by considering consumer surplus, the platform becomes agnostic towards consumers' maximal willingness to pay.

Proposition 3.13. *A controlled allocation platform maximizing consumer surplus have infinitely large price of anarchy. In particular, even under a two market, one firm setting, the price of anarchy is ∞ .*

The idea behind the two node construction which proves the proposition is simple: Maximizing consumer surplus implies that the platform is agnostic to the maximal willingness to pay, whereas rational producers consider that. A simple example of this is a one firm, two market network has the firm with linear cost cs , and the markets have parameters

$$\alpha_1 = c + \epsilon, \beta_1 = \epsilon/2, \alpha_2 = c - \epsilon, \beta_2 = \epsilon,$$

which results in the platform allocating fully to the second market before accessing the first market. It is trivial to show that being forced to participate in the whole of the second market, the firm would rather not produce anything at all. On the other hand, the optimal allocation has 2 units in the first market and none in the second, resulting in a positive social welfare, and therefore an unbounded worst case efficiency loss.

A General result on Controlled Allocation Platforms.

The control over allocations is delicate, and explicitly not optimizing (as in the open access case) is more efficient than optimizing either social welfare or consumer surplus. To understand how the worst case efficiency loss is affected by different objective functions, we now extend the analysis of the social welfare case to consider an objective $\text{OBJ}(\lambda)$ to be a convex combination of consumer surplus and revenue, parametrized by λ :

$$\text{OBJ}(\lambda) = \sum_{j \in M} \left[\lambda \left(\int_0^{d_j} p_j(q) dq - d_j p_j(d_j) \right) + (1 - \lambda) (d_j p_j(d_j)) \right].$$

For example, when $\lambda = 1$, the above objective is consumer surplus whereas $\lambda = 0$ gives revenue, and when $\lambda = 0.5$, the resulting objective is half of consumer welfare, which is equivalent to a social welfare objective. In particular, we are able to provide a lower bound on the worst case efficiency loss for the controlled allocation platform maximizing the convex combination of objective functions.

We present a more general version of the previous results on consumer surplus ($\lambda = 1$) and social welfare ($\lambda = 1/2$), where worst case efficiency loss is infinitely large and growing linearly in the number of markets respectively. This result is presented in the following theorem.

Theorem 3.14. *The worst case efficiency loss $\rho(k)$ of controlled allocation platforms optimizing various objective functions (characterized by $\sigma = \frac{2-3\lambda}{1-\lambda}$) satisfies:*

$$\rho(\sigma) \geq \begin{cases} \infty, & \text{if } \sigma \leq 0. \\ \Omega(m), & \text{if } 0 < \sigma \leq 1 \\ \max\left\{ \frac{2}{3\sqrt{\frac{1}{\sigma}(1-\frac{1}{\sigma})}}, \frac{3}{2} \right\}, & \text{otherwise.} \end{cases}$$

This means that the construction developed to prove Theorem 3.11 also can be used to show that any objective between consumer surplus and social welfare leads to similar results to the social welfare case in Theorem 3.11, and that past social welfare the upper bound for worst case efficiency loss decreases towards a nice constant, which corresponds to the open access result.

When considering linear inverse demand functions and assuming $d_j \leq \alpha_j/\beta_j$, the above objective can be simplified to the following after dividing by a constant $(1 - \lambda)$ assuming $\lambda \neq 1$ (having dealt with the consumer surplus case in the previous

subsection),

$$\text{OBJ}(\lambda) = \sum_{j \in \mathcal{M}} \left[\frac{\lambda}{2(1-\lambda)} \beta_j d_j^2 + (\alpha_j d_j - \beta_j d_j^2) \right].$$

Observing the stationarity conditions with respect to d_j , we find that the platform maximizing such an objective would keep the following quantity constant across markets:

$$\frac{\delta \text{OBJ}(\lambda)}{\delta d_j} = \alpha_j - \frac{2-3\lambda}{1-\lambda} \beta_j d_j,$$

which confirms that as λ approaches 1, the emphasis on the maximal willingness to pay α_j decreases, and becomes potentially inefficient. Increasing past the point $\lambda = 2/3$, the optimization function also changes from a concave maximization problem to a convex one, i.e., the problem ends up being a (nonconvex) Quadratic Program (QP). This also reveals that whenever a new market is accessed (regardless of $\lambda < 1$), that the following relationship is preserved for any new additional supply $q_j, q_k, \beta_j q_j = \beta_k q_k$, with the parameter λ basically dictating when the next market is accessed. For example, when $\lambda = 0.5$, prices are kept constant for each participated market.

The issue with controlled allocation platforms is that the platform's market clearing mechanism is agnostic to underlying cost functions. While this action seemingly makes much sense under this setting, we show above that having the wrong objectives can yield (in the extreme) unbounded or infinitely large price of anarchy, preserving close to or none of the optimal social welfare. It should be clear through the section that such a strategy of controlling allocations potentially leads to undesirable and reactive producer behavior. As such, an alternative market clearing mechanism is necessary.

3.5 Efficiency of Discriminatory Access

Open access designs preserve large productions but suffer from misaligned allocations, while controlled allocation designs lead to undesirable and unintended incentives for the producers to curtail production. Recently, some platforms have started shifting from open access platform designs towards more delicate designs which aim to both retain the incentive to produce in open access platforms and exert some control over allocations. Results from this section show the promise of some of these changes. We design such a platform to be similar to mechanisms such as Amazon's Buybox and Airbnb's Superhost, which highlights particular sellers or renters while hiding others.

In the setting of our model, such access control corresponds to the specification of the edge set of the bipartite graph that connect firms to markets, with the goal of maximizing social welfare at the unique Nash equilibrium of the resulting networked Cournot game. This platform is referred to as a *discriminatory access platform*. Unlike the controlled allocation platform, the efficiency of discriminatory access platforms is guaranteed to be no worse than open access platforms, as open access is a valid choice of access control.

Our work on discriminatory access design culminates in the following theorem, highlighting the bounded efficiency loss from discriminatory access platforms at the optimal network design. We prove a complementary result in Theorem 3.17 that this optimal network design problem can be computed for the worst case cost functions via a greedy algorithm presented later in the section.

Theorem 3.15. *Assume that each firm's cost function is linear over $(0, \infty)$, and assume that \mathcal{E}^* is an the optimal network design for the discriminatory access platform. Let $C \in \mathcal{L}^n(c_{\min}, c_{\max})$, and assume that $c_{\min} < \max_{j \in M} \alpha_j$. The efficient social welfare associated with the edge set \mathcal{E}^* satisfies*

$$\text{SW}^*(\mathcal{E}^*, C) = \text{SW}^*(F \times M, C).$$

The efficient social welfare can be attained at the edge set \mathcal{E}^ . Moreover, the worst case efficiency loss associated with the discriminatory access networked Cournot game $(F, \mathcal{Q}(\mathcal{E}^*), \pi)$ is upper bounded by*

$$\rho(\mathcal{E}^*, C) \leq \frac{\sum_{j=1}^m \frac{((\alpha_j - c_{\min})^+)^2}{\beta_j}}{\sum_{j=1}^m \max_{k \in \{1, \dots, n\}} \left\{ \frac{2k+4}{3k+5} + \delta(\gamma_j, k) \right\} \frac{((\alpha_j - c_{\min})^+)^2}{\beta_j}},$$

where the function $\delta(\gamma, n)$ is defined in Proposition 3.10, and serves as a measure of the degree of asymmetry between firms' cost functions.

The above bound is tight if $\alpha_1 = \alpha_2 = \dots = \alpha_m$.

The functions that maximize worst case efficiency loss in networked Cournot games are linear cost functions. Therefore, the bounds in Theorem 3.15 also applies to cases where firms have general convex costs. Additionally, choosing the number of firms $n = 1$ yields a worst case efficiency loss bound of $4/3$ for (optimized) discriminatory access platforms with any number of firms and markets, improving upon the $3/2$ price of anarchy bound for open access platforms established in Corollary 3.6. The result is formally stated as follows.

Corollary 3.16. *Assume that each firm's cost function is linear over $(0, \infty)$. Discriminatory access platforms have worst case efficiency loss of at most $4/3$.*

In general, we show that choosing an optimal edge set that maximizes social welfare at Nash equilibrium amounts to a mathematical program with equilibrium constraints (MPEC), and is, in general, computationally intractable.

The Optimal Network Design Problem

The optimal network design problem amounts to the selection of an edge set \mathcal{E} , which maximizes social welfare at the unique Nash equilibrium of the resulting networked Cournot game. Formally, Lemma 3.2 provides a characterization of the supply profile at the unique Nash equilibrium of the game $(F, \mathcal{Q}(\mathcal{E}), \pi)$ as the unique optimal solution to a convex program. Therefore, the *optimal network design problem* admits a formulation as the following MPEC:

$$\begin{aligned} & \text{maximize} && \text{SW}(q, C) \\ & \text{subject to} && \mathcal{E} \subseteq F \times M \\ & && q \in \arg \max_{x \in \mathcal{Q}(\mathcal{E})} \left\{ \text{SW}(x, C) - \sum_{j=1}^m \sum_{i=1}^n \frac{\beta_j x_{ij}^2}{2} \right\}. \end{aligned} \tag{3.9}$$

Here, the decision variables are the edge set \mathcal{E} and the supply profile q . The challenge in solving the above problem stems from the equilibrium constraint on q , and the presence of the discrete decision variable \mathcal{E} . (An equilibrium constraint requires that a vector be an optimal solution to a optimization problem. In general, this leads to a nonconvex and disconnected feasible region for MPECs. See (Luo, J.-S. Pang, and Ralph, 1996) for a more detailed discussion.) In what follows, we show that, considering linear cost functions which represents the worst case efficiency loss, the above problem can be solved using a greedy algorithm.

Greedy Algorithm for Optimal Worst-case Network Design

In this section, we restrict ourselves to cost functions that are linear on $(0, \infty)$, i.e., the ones maximizing worst case efficiency loss. Specifically, we assume that the cost function of each firm $i \in F$ satisfies $C_i(s_i) = (c_i s_i)^+$, where $c_i \geq 0$. Leveraging on this assumption, we propose a greedy algorithm for solving the optimal network design problem (9) in Algorithm 1. For each market $j \in M$, the greedy algorithm visits firms in ascending order of marginal cost, and provides each firm it visits access to market j if its inclusion in that market increases social welfare.

loa 1: The Greedy Algorithm

Require: $c_1 \leq \dots \leq c_n$.

- 1: Initialize edge set $\mathcal{E} \leftarrow \emptyset$.
- 2: **for** $j = 1$ to m **do**
- 3: Initialize firm index $i \leftarrow 1$.
- 4: Initialize edge set $\tilde{\mathcal{E}} \leftarrow \mathcal{E}$.
- 5: **repeat**
- 6: Update edge set $\mathcal{E} \leftarrow \tilde{\mathcal{E}}$.
- 7: **if** $i \leq n$ **then**
- 8: Set edge set $\tilde{\mathcal{E}} \leftarrow \mathcal{E} \cup (i, j)$.
- 9: Set firm index $i \leftarrow i + 1$.
- 10: **end if**
- 11: **until** $\text{SW}(q^{NE}(\tilde{\mathcal{E}}), C) \leq \text{SW}(q^{NE}(\mathcal{E}), C)$.
- 12: **end for**
- 13: **return** \mathcal{E} .

Clearly, Algorithm 1 yields an edge set \mathcal{E}^* , whose corresponding Nash equilibrium has a social welfare that is no smaller than that of the open access platform.

Theorem 3.17. *If \mathcal{E}^* is an edge set obtained from Algorithm 1, then $(\mathcal{E}^*, q^{NE}(\mathcal{E}^*))$ is an optimal solution to (3.9), that is, the greedy algorithm achieves the optimal network design.*

Theorem 3.15 reveals the advantage discriminatory access platforms have over open access ones in reducing the efficiency loss at Nash equilibrium, while Theorem 3.17 guarantees we can find such an edge set when optimizing over the worst case network design. Namely, when the edge set is chosen to be an optimal solution of the network design problem (9), the discriminatory access platform is guaranteed to have a tight bound on the price of anarchy that is no larger than that of the open access platform. Moreover, this price of anarchy bound is guaranteed to be non-increasing in the number of firms n .

We observe under our model that sacrificing transparency and discriminating choices of firms for different markets allow for a slight improvement in the worst case social welfare preserved, or a slight improvement in efficiency. In what follows, we present a simple search cost model to study the impacts of consumer search cost on platform design, and further demonstrate the flexibility of the discriminatory access platform design to consider search costs.

3.6 Search Costs and its Implications

The model and analysis presented thus far considers production costs but not other costs which can also be on the consumer side, and may be monetary, e.g., transportation and communication, or non-monetary, e.g., search costs. Platform designs center around lowering entry costs for producers, oftentimes leading to increased production and competition, and lowering search costs for consumers. In this section, we present a simple search cost model which highlights the ability of the discriminatory access design to also consider search costs, and the distinction between open access and discriminatory access designs.

As we are working in the networked Cournot setting where demand is aggregated, we must use a simple model for search costs. We define a search cost for each consumer market j , r_j , on the consumers as a product of their consumer surplus (taking care of the simple assuring assumption that markets are rational and that search costs can never be more than their surplus) and a discount factor $f(n)$ which monotonically increases with the number of firms n , and lastly, a market-specific parameter θ highlighting potentially different market segments with differing search costs, defined as follows:

$$r_j = \theta f(n) CS_j, \quad 0 \leq \theta \leq 1,$$

where function f has the following properties:

$$f(1) = 0, \quad \lim_{n \rightarrow \infty} f(n) = 1, \quad f \text{ monotonically increasing}$$

This search cost model ensures that (i) search cost penalties can never exceed consumer surplus—an assumption that search cost can never cause the market to receive negative utility, (ii) differentiation can be made through θ for low and high search cost participants, and that (iii) search cost increases in the number of firms exposed to - the effort or cost has to increase as the number of choices increase.

Remark 3.18. *Since controlled allocation platforms essentially make all decisions on allocation, the search cost is 0 for consumers but continue to maintain its previous negative result from Section 4, implying that search costs has no effect on controlled allocation platforms.*

In light of the difficulty to analyze without a further parametric form of $f(n)$, we assume hereafter that $f(n)$ assumes the form of $f(n) = \frac{n-1}{n+1}$, fulfilling the preceding properties we listed.

Open access designs, by definition, presents all possible choices to consumers, allowing them to make production and allocation decisions, but also thereby incurring large search costs. In the following theorem, we highlight the impact of search costs on open access platforms.

Theorem 3.19. *Open Access Platforms with search costs defined with $f(n) = \frac{n-1}{n+1}$ have worst case efficiency loss $\Omega(n)$.*

On the other hand, referring to Corollary 3.16, where the price of anarchy is $\frac{4}{3}$ at $n = 1$ (which coincidentally incurs zero search cost) for discriminatory access platform designs, they maintain its worst case constant bound price of anarchy under linear cost functions and search costs.

Prior to this section, it seems that the performance guarantees attained by open access and discriminatory access platform designs are similar. This section presents a preliminary look at the impact of search cost on platform design, in particular highlighting the potentially post-search unbounded worst case efficiency loss result for open access platforms, and the flexibility of the discriminatory access platform design to consider search costs. It remains important to study this further, e.g., to find out how to optimize network design with search costs, or how computation or other costs involved may affect the performance of discriminatory access designs too.

Concluding Remarks

This chapter presents trade-offs between transparency and control in online platform design. We present three platform designs under a networked Cournot competition model that resembles some of the more popular platform designs today. They are open access, controlled allocation, and discriminatory access designs. Open access designs transparently allow full access, while controlled allocation designs make allocation decisions for producers. Discriminatory access designs seek a balance between transparency and control, by controlling over the firm-market access.

The driving force behind the significant results in this work is the insight in Lemma 3.8 that linear costs are the worst among convex cost functions for minimizing efficiency loss. This result implies that for analyzing worst case efficiency loss bounds, a networked competition model is only as complex as a competition model with a single market since linear production costs imply constant marginal cost, which decouples decisions across markets.

From that, we show that open access designs preserve a significant proportion of production and alignment to social welfare. Further, we show that open access designs have a worst case efficiency loss bound of $3/2$. Complementing that, we also include bounds for when firms have symmetric and asymmetric cost functions, which reveals a well-known result that an increase in competitors leads to improved efficiency, but what is surprising is that this does not hold in the case with asymmetric costs. Such a result is because increasing the number of competitors with asymmetric costs in that case only reduce competition in the worst case. However, when further considering a simple search cost model, we find that the worst case efficiency loss for open access designs can be large, and can grow in the number of firms.

Next, we show that controlling allocations to maximize social welfare cause sufficient misalignment in incentives with producers for them to curtail production significantly. We show that the platform keeps prices constant across active markets, and also show that the resulting game is equivalent to a Cournot game with convex inverse demand, as studied in (John N Tsitsiklis and Yunjian Xu, 2014). We show that this results in a worst case efficiency loss lower bound that grows in the number of markets. We present a generalization of that result to platform objectives or market clearing mechanisms beyond social welfare.

Lastly, we show that discriminatory access platforms can achieve improved worst case efficiency loss bounds of $4/3$. The optimal network design problem is formulated as a mathematical program with equilibrium constraints, highlighting its complexity. We find an optimal greedy algorithm for the case when firms' production costs are linear, representing an optimization over the worst case cost function. We also show that discriminatory access platforms can balance efficiency and search costs, retaining its $4/3$ bound even in the presence of search costs.

Chapter 4

DEMAND MANAGEMENT IN ONLINE PLATFORMS

We witnessed in Chapter 3 how practical considerations like search cost affect the efficiency of different platform designs. This chapter aims to study another practical implementation in platforms known as *demand management*.

Demand management refers to signals and rewards designed to incentivize customers' demand reduction from nominal consumption patterns. It has been utilized in multiple different areas, e.g., electricity markets and ridesharing platforms. It manifests as load-side demand management (a.k.a. demand response) and locational marginal prices (LMPs) in electricity markets and surge pricing in ridesharing platforms. They are known to be useful in promoting the interaction and responsiveness of consumers, guiding short term impacts for the market to improve stability, ensure supply meets demand while meeting certain network constraints.

Demand management can take on very different forms. LMPs (similarly, surge pricing) provide prices across a congested network that reflect the marginal price of demand at each location, to minimize supply-demand imbalances. Demand Response, on the other hand, rewards consumers for reducing their nominal demand. Dependent on how demand management is set up, producer manipulation can instead turn it on its head, leading to significant inefficiencies in these markets. Manipulations in the electricity market and ridesharing platforms evidence these inefficiencies, e.g., Enron during the California Energy Crisis in 2000 costing US\$40-45 Billion (Weare, 2003), JPMorgan Ventures Energy Corp in 2012, coordinated logging off from Uber applications by drivers to induce surge pricing. By anticipating these demand management tools, producers have the power to manipulate prices in these markets for their gains.

The markets that require strict governance are networked, firstly in that there are multiple producers and consumers, but also in a sense that there are often constraints on how demand can be allocated or transported across the network. For example, drivers on ridesharing platforms usually prefer to drive within their geographical location if prices and demand allow for it. Meanwhile, transferring electricity from one node to another on the electrical grid depends on power and line flow constraints, driven by the physical laws of electricity transportation.

Beyond considering networked marketplaces, a complete study on the impact of demand management must also include the potential for anticipation and manipulation. To that end, we introduce a novel model of networked competition that takes into account producer anticipation on platform allocation, designed to be a leader-follower version of the networked Cournot competition model with a platform, where producers choose production quantity to maximize individual profits while the platform rebalances across the network to maximize social welfare. In this networked Stackelberg game, producers compete against each other knowing (in anticipation of) the platform market clearing mechanism, i.e., all producers know how the platform rebalances across the network, as in Section 3.4. In contrast, without knowledge (and anticipation) of the clearing mechanism, the producers compete against each other and the platform simultaneously in the networked Cournot competition together. Under these models, we can gain new and interesting insights on anticipation and demand management in networked markets.

Example: Electricity Markets

Behind the success of the grid is an electricity market governed by an independent system operator (ISO), who ensures that the *right prices* across the grid are set such that (i) supply meets demand, (ii) social welfare is maximized, and (iii) network flow physics and constraints are considered. To do this, the ISO usually receives quantity bids from producers and estimates demand over the grid, and optimizes rebalancing constrained on network flow and physics, resulting in prices that serve as *optimal incentives* for producers at each node to generate the right amount.

If producers cannot anticipate allocations of the ISO, then this set-up is indeed optimal given truthful quantity bids and demand. However, market clearing mechanisms and network constraints may be known to producers. This knowledge can lead to price manipulation when entities obtain sufficient “market power”, e.g., Enron in the California Energy Crisis. As such, it is critical that we consider producer anticipation when studying electricity markets under a governing platform.

Thankfully, mature electricity markets have developed mechanisms to circumvent some of these problems. One of them is known as load-side demand management, a.k.a. demand response. It helps alleviate binding constraints on the grid, consequently aiding with network stability and can help with the increasing amount of renewable energy on the grid. We show in this chapter that demand management mechanisms can also help curb manipulation caused by producer anticipation.

Example: Ridesharing Platforms

Ridesharing platforms have revolutionized transportation in cities and change the ways car ownership is viewed today. Behind the success of ridesharing are platforms who carefully set prices for rides and match riders to drivers. These platforms (i) provide sufficient incentives for driver participation, (ii) ensure reasonable waiting time for passengers, and (iii) carefully subsidize the short side to improve engagement further and hopefully thrive on *network effects* in these networked marketplaces.

Unlike electricity, there is a cost involved for the drivers to be relocated without a passenger, in terms of opportunity and movement cost. To circumvent this problem, recent changes in *surge pricing* promise an additive bonus for re-locations, priced carefully to cover the mentioned costs.

Information through ridesharing platforms is explicitly designed so that individual drivers (along with their corresponding production constraints) do not have sufficient “market power” individually. This design is also the reason why manipulation on ridesharing platforms require coordinated cooperation, e.g., drivers synchronizing their logging out of ride-sharing applications to cause a supply-demand mismatch, activating surge pricing. This manipulation again compels studying anticipation in these networked marketplaces.

As ridesharing platforms mature, it will be critical that they prevent such collusion and manipulation due to the anticipation of market clearing and pricing mechanisms. We propose that like demand response for electricity markets, demand management reduces the market power and impact of manipulation for drivers.

Our Contributions

In this chapter, we show that demand management is a powerful tool in preventing producer manipulation. We show that beyond what others usually associate demand management with, it further limits manipulation, but does so in a unique manner. In particular, we posit that demand management mechanisms empower consumers and allow them to provide back “non-strategically”, essentially handing “market power” to consumers to counter that which producers may have. We first introduce a leader-follower or Stackelberg version of the well-studied networked Cournot competition model. We call this model the *networked Stackelberg model*, where producers jointly act as the leader, and a controlling platform who rebalances quantity over the network follows. Alternatively, producers without anticipation participate in a simultaneous networked Cournot game with the platform.

In Section 4.2, we present our competition models¹ and define the corresponding equilibrium concepts. We also describe the constraints over rebalancing allocations and present a characterization of platform allocations. We design the network such that each producer is co-located with a consumer, where nodal pricing results from binding constraints over allocations.

Our first result seeks to understand the impact of anticipation in these networked marketplaces. We find in Proposition 4.2 that the networked competition *where producer do not understand the market clearing mechanism leads to inefficiency even in simple networks*. On the other hand, we find in Section 4.3 conditions for successful and efficient anticipatory competition in networked marketplaces. If price intercepts are homogeneous across an unconstrained marketplace, then the networked Stackelberg game yields *a unique equilibrium with bounded multiplicative worst case efficiency loss of 3/2*. However, we show by construction that the conditions listed are “almost necessary”, i.e., taking away either will remove the existence of equilibria or the bounded worst case efficiency loss.

Our aim in the final part of this work is to highlight the potential of load-side demand management beyond its traditional benefits in also improving economic efficiency. We develop a simple model of demand response under the networked Stackelberg model and show that under the absence of network constraints, a unique equilibrium exists, again with a worst case efficiency loss bound of 3/2. To supplement this result, we further provide conditions such that demand response serves as a threat that does not get activated at equilibrium. We exhibit the efficacy of demand response by applying it to one of the previous examples and highlight its ability to discourage equilibria with lower production levels. Lastly, we show that equilibria with binding network constraints do not suffer significant losses when consumers participate in demand management. More specifically, we show that the worst case loss can be split into two parts: (i) the original multiplicative 3/2 bound, and (ii) an additive loss bounded by the revenue passing through constrained lines, which we see as a *price of congestion*. This result means that if constrained lines do not carry too much revenue, then the equilibrium remains efficient. More notably, if none of the constraints are binding, then we retrieve the original 3/2 multiplicative efficiency loss bound.

¹The main difference between the controlled allocation platform in Chapter 3 and the one considered here is the co-location of a firm with each market, and the network constraints over the rebalancing. The co-location allows for a manifestation of nodal prices. In the controlled allocation platform design, there were no constraints over rebalancing.

4.1 Literature Review

Our focus in this work is the impact of demand management mechanisms in networked marketplaces where producers anticipate the actions of the governing platform or market maker. This work lies in the intersection of platform design and demand management in networked competition, but such work has traditionally also been mostly studied for applications in electricity market analysis. We present here related literature from electricity market load-side demand management, while related literature for networked competition and platform design can be found in Chapters 2 and 3 respectively.

The idea of demand management is not new, and have been widely used in electricity markets, and studied as early as the 1980s (Gellings, 1985). Electricity market demand management promotes the interaction and responsiveness of customers, and may offer a broad range of potential benefits on system operation and expansion and on market efficiency (Siano, 2014). Additionally, by improving the reliability of the power system and, in the long term, lowering peak demand (L. P. Qian et al., 2013), demand response reduces overall plant and capital cost investments and postpones the need for network upgrades. Today, innovative enabling technologies and systems, such as smart meters, energy controllers, communication systems, decisive to facilitate the coordination of efficiency and demand response in a smart grid. (Huang, Walrand, and Ramchandran, 2012) also devised lightweight algorithms to solve for optimal energy and load-side demand management.

Beyond stability, load-side demand management strategies can also help with delaying investment on new generation capacity from renewable resources and improve operation of existing installed capacity (Pina, Silva, and Ferrão, 2012). An implication of this is that increased demand response can help us utilize an increased amount of renewable energy, which can be cheaper, but more critically, less harmful to our environment. A study on the practical implementation of secure and private load-side demand management can be found in (Palensky and Dietrich, 2011), and a study on real-time implementation can be found in (Conejo, Morales, and Baringo, 2010).

To the best of our knowledge, this work is the first to consider economic impact in terms of economic stability and reduction in market power.

4.2 Model & Preliminaries

In this chapter, we focus on the impact of demand management in a networked marketplace. We model the interaction of a collection of strategic firms and a governing platform \mathcal{P} that can transport a commodity over a network. We begin by describing the marketplace and the competition model below.

Modeling the marketplace

Consider M markets labeled $1, \dots, M$ that are connected via a network. A single strategic firm F_m supplies at market m . A firm in our model can only supply to its local market. The platform \mathcal{P} provides transport, thereby reallocating goods across markets. Firm F_m decides its supply quantity $q_m \geq 0$ for which it incurs a cost of $c_m(q_m)$. Assume that supply costs of all firms are continuously differentiable increasing convex functions with zero investment costs ($c_m(0) = 0$). Denote the supply profile across the markets by $\mathbf{q} := (q_1, \dots, q_M)^\top$. Define \mathbf{q}_{-i} as the same profile save the i -th one.

Consider an aggregate price-taking consumer at market m characterized by an inverse demand function of the form

$$p_m(d_m) := \alpha_m - \beta_m d_m.$$

Here, $\alpha_m > 0$ denotes the maximal price the consumers are willing to pay and β_m denotes the price elasticity in market m .² Let $\mathbf{d} \in \mathbb{R}^M$ denote the demands across all markets.

Platform \mathcal{P} provides transportation of the commodity over a network joining the markets. Denote by r_m , the amount that \mathcal{P} allocates to market m . Collect them across all markets in

$$\mathbf{r} = \mathbf{d} - \mathbf{q} \in \mathbb{R}^M.$$

The transport is constrained by the capabilities of the network, modeled as a set of line constraints in the sequel with $\mathbf{1}^\top \mathbf{r} = 0$, where $\mathbf{1} \in \mathbb{R}^M$ is a vector of all ones. Such a model ensures that \mathcal{P} may choose not to reallocate ($\mathbf{r} = 0$ is feasible) and no supply is wasted during the reallocation process. This model for the network is inspired by electricity markets where generators produce and loads consume power at their respective buses of an electric power network. A system operator clears the market in a way that the power injections across the network respect the engineering constraints of the grid. Decisions in electricity markets are often made based on a

²Our results mostly generalize to the case with general concave decreasing inverse demands.

linearized lossless power flow model known as the DC-approximation Stoft (2002) and Purchala et al. (2005). Under this popular approximation, (directed) power flows on N transmission lines become $\mathbf{H}(-\mathbf{r})$ for power injections $-\mathbf{r}$ across the network, which is assumed to be no larger than the production \mathbf{q} . The matrix $\mathbf{H} \in \mathbb{R}^{2N \times M}$ is the so-called injection shift-factor matrix. The directed power flows respect the transfer capabilities of the transmission lines, modeled as $\mathbf{H}(-\mathbf{r}) \leq \mathbf{f}$ with $\mathbf{f} \in \mathbb{R}_+^{2N}$.

Modeling Competition

Present models of networked competition study simultaneous games between all participants, firms and platform alike, as in the networked Cournot competition in (Cai, Bose, and Wierman, 2017). Our focus on demand management under a governing social planner necessitates the consideration of producer anticipation on these platform rebalancing allocations. We do this under a model of networked Stackelberg competition, where firms act as joint leaders and platform \mathcal{P} follows. We formally introduce both models after we define each participant's payoff functions. Given rebalancing quantities \mathbf{r} , the utility for producer m is its profit π_m , i.e.,

$$\pi_m(q_m, r_m) := q_m p_m(q_m + r_m) - c_m(q_m),$$

which it aims to maximize over its production constraints $q_m \geq 0$. The platform \mathcal{P} plays the role of a socially benevolent planner, seeking to maximize social welfare of the system. It is a well-studied objective, e.g., (Ramesh Johari and John N Tsitsiklis, 2004), and is also used for the optimization of electricity markets. Further, optimizing social welfare minimize variance of prices across the distinct markets in a network. Social welfare is defined as aggregate consumer welfare less producers' total cost, and equivalently, the sum of consumer, producer and merchandising surplus, first defined by

$$\begin{aligned} \text{CS}(\mathbf{q}, \mathbf{r}) &:= \sum_{m=1}^M \left(\int_0^{q_m+r_m} p_m(w_m) dw_m - (q_m + r_m) p_m(q_m + r_m) \right), \\ \text{PS}(\mathbf{q}, \mathbf{r}) &:= \sum_{m=1}^M (p_m(q_m + r_m) - c_m(q_m)), \quad \text{MS}(\mathbf{q}, \mathbf{r}) := \sum_{m=1}^M (r_m p_m(q_m + r_m)). \end{aligned}$$

By maximizing the platform payoff function $\Pi_{\mathcal{P}}(\mathbf{q}, \mathbf{r})$, i.e., social welfare, or

$$\Pi_{\mathcal{P}}(\mathbf{q}, \mathbf{r}) := \text{SW}(\mathbf{q}, \mathbf{r}) = \sum_{m=1}^M \left(\int_0^{q_m+r_m} p_m(w_m) dw_m - c_m(q_m) \right),$$

over its constraint set $\mathbf{r} \in \mathbb{X}$, the platform aim to balance objectives of the various parties in the marketplace. Such a market clearing mechanism defines a reaction

function $\rho : \mathbb{R}_+^m \rightarrow \mathbb{X}$, which determines the platform's reallocation \mathbf{r} for each set of supply quantity profile \mathbf{q} submitted by the producers. This reaction function ρ is determined by the following optimization:

$$\rho(\mathbf{q}) = \arg \max_{\mathbf{r}} \sum_{m=1}^M \left(\int_0^{q_m+r_m} p_m(w_m) dw_m - c_m(q_m) \right),$$

subject to $\mathbf{r} \in \mathcal{X}(\mathbf{q})$,

which admits an efficient (convex programming) solution for common choices of cost functions, e.g., linear and quadratic production cost.

The critical distinction between the two competition models is whether the producers consider this reaction function, or equivalently, how the platform reallocation \mathbf{r} changes, with a change in production quantity \mathbf{q} . Since the reallocation \mathbf{r} is decided based on the supply profile of all firms, firms in the networked Stackelberg game are actually aware of the competition they are in with other firms in the system.

Non-anticipatory competition (Cournot)

In the non-anticipatory competition, producers and the platform \mathcal{P} participate simultaneously. A profile $(\mathbf{q}^C, \mathbf{r}^C)$ is a generalized Nash-Cournot (or non-anticipatory) equilibrium if

$$\pi_m(q_m^C, \mathbf{q}_{-m}^C, \mathbf{r}^C) \geq \pi_m(q_m, \mathbf{q}_{-m}^C, \mathbf{r}^C), \quad \forall q_m \geq 0,$$

for each producer m , and

$$\Pi_P(\mathbf{q}^C, \mathbf{r}^C) \geq \Pi_P(\mathbf{q}^C, \mathbf{r}), \quad \forall \mathbf{r} \text{ such that } \mathbf{r} \in \mathbb{X}(\mathbf{q}^C).$$

for the platform \mathcal{P} whose strategy set depends on the firms' productions. A careful reader may notice that the profit of each firm can be written in the form $\pi_m(q_m, r_m)$, i.e., each firm plays a game only with the platform \mathbf{P} , since each firm does not anticipate reallocations of the platform. At the generalized Nash equilibrium of the Cournot competition, neither the producers nor platform \mathcal{P} can unilaterally deviate to obtain improved utility. By showing that these are concave games, one can prove the existence of this (generalized Nash) equilibrium by applying (Rosen, 1965).

Anticipatory Competition (Stackelberg)

When all producers anticipate platform rebalancing, they participate in a simultaneous game only against the other producers. The platform \mathcal{P} follows with a

rebalancing allocation that is known to all producers which depends on the quantity bids from each of them, effectively making firms aware of the coupling across production quantities at different nodes.. In other words, the producers account for the response of \mathcal{P} 's actions in its quantity offer in this game, as in a Stackelberg (leader-follower) game. We model competition based on complete information, and while this represents idealized conditions in real-life settings, this model is rich enough to bear insights on the impact of anticipatory behavior of producers in such marketplaces.

We define a production and rebalancing allocation profile $(\mathbf{q}^S, \mathbf{r}^S)$ to be a *Stackelberg equilibrium* or an equilibrium in the networked Stackelberg competition if,

$$\pi_m(\mathbf{q}^S, \rho(\mathbf{q}^S)) \geq \pi_m(q_m, \mathbf{q}_{-m}^S, \rho(q_m, \mathbf{q}_{-m}^S)), \forall q_m \geq 0,$$

i.e., no producer can unilaterally deviate with anticipation of rebalancing allocations.

Modeling Demand Management

In contrast to the original consumer model, we allow in Section 4.4 for fulfilled demand d_m at some markets to be negative, which corresponds to load-side demand management. Concave quadratic consumer utilities yield linear inverse demand, indicating prices one is willing to pay for demand. In particular, $p_m(d_m) = \alpha_m - \beta_m d_m$ corresponds to utility

$$U_m(d_m) = \int_0^{d_m} \alpha_m - \beta_m x \, dx = \alpha_m d_m - \frac{\beta_m}{2} d_m^2. \quad (4.1)$$

We assume that there is already nominal demand D_m fulfilled at each node, and that the above utility functions model that of residual demand. Extending to the negative domain of p_m implies that consumers are willing to curtail its nominal demand D_m at a particular payment rate, incurring a utility $U(d)$ which is negative when d is as well. In return, these consumers will be paid $-q_m p_m(d_m)$ as a reward.

A Two-node Illustrative Game

While networked marketplaces with platforms as described in our paper are abundant, we provide here a two-node illustrative game, with a producer and consumer in each and a link joining them with flow constraints. We do not constrain this example game to be for a particular application, but use it to illustrate results that we will prove in this work.




Figure 4.1: Two-node illustrative example defined as $\mathcal{G}(\alpha, \beta, \mathbf{c}, f)$.

In this two-node network, we simplify the line constraints by a simple line capacity f . At each node $m = \{1, 2\}$, let $p_m(d_m) = \alpha_m - \beta_m d_m$ with cost $c_m(q_m)$. Whenever we present an example of this two-node game, we use $\mathcal{G}(\alpha, \beta, \mathbf{c}, f)$ to describe it. ³

4.3 Anticipation in Networked Markets

We model producer anticipation in these markets by an understanding of how the platform allocates demand over the network, e.g., power flow allocations in electricity markets. Recent events on anticipation and manipulation in electricity and ridesharing markets motivate our model and analysis including anticipation.

In this section, we prove the following result on conditions under which competition in these networked Stackelberg setting can be efficient and inefficient.

Theorem 4.1. *Assume that price intercepts of markets are homogeneous, i.e., $\alpha_{m_1} = \alpha_{m_2}$ for all markets m_1, m_2 , with no network restrictions for the platform \mathcal{P} on reallocation, i.e., $\mathbf{f} = \infty$, then there is a unique equilibrium $(\mathbf{q}^S, \mathbf{r}^S)$ to the networked Stackelberg game, and its worst case efficiency loss is bounded, i.e.,*

$$SW(\mathbf{q}^*, \mathbf{r}^*) \leq \frac{3}{2} SW(\mathbf{q}^S, \mathbf{r}^S).$$

Otherwise, the networked Stackelberg game, in general, can also have either (i) no equilibrium or (ii) multiple equilibria. Even when an equilibrium exists, the networked Stackelberg game have unbounded worst case efficiency loss, i.e., there exists a family of networked Stackelberg games such that $(\mathbf{q}^S, \mathbf{r}^S)$,

$$SW(\mathbf{q}^*, \mathbf{r}^*) \geq \Omega(M)(SW(\mathbf{q}^S, \mathbf{r}^S)),$$

i.e., worst case efficiency loss grows linearly in the number of nodes M in the network.

³We will find throughout this chapter that while simple, this two-node game will be useful for illustrating and contrasting the different settings with regards to the existence of equilibria and its corresponding efficiency with/without anticipation and/or demand management.

To contrast the efficiency in the first part of this result, we first show using the two-node game $\mathcal{G}([1, 1]^T, [1, \gamma]^T, [\gamma q^2, q^2]^T, \infty)$ that the Cournot equilibrium can be arbitrarily inefficient as γ grows, summarized concisely in the following proposition.

Proposition 4.2. *The networked Cournot game have worst case efficiency loss arbitrarily large, i.e., for any $K \in \mathbb{R}_+$, there exists a $\gamma > 0$ such that the generalized Nash equilibrium $(\mathbf{q}^C, \mathbf{r}^C)$ of the game $\mathcal{G}([1, 1]^T, [1, \gamma]^T, [\gamma q^2, q^2]^T, \infty)$ has efficiency loss at least K , i.e.,*

$$SW(\mathbf{q}^*, \mathbf{r}^*) \geq K(SW(\mathbf{q}^C, \mathbf{r}^C)).$$

Note that the unbounded line flow is not actually necessary. In fact, any bounded line flow larger than $1/3$ is sufficient for the result above to hold. Having the unbounded line flow allow us to easily contrast against results from the first part of Theorem 4.1.

In the rest of this section, we first discuss the potential of anticipation in improving efficiency when networked markets are “hard to manipulate”, providing intuition to the efficiency loss bounds in the theorem. In contrast to the arbitrarily large efficiency loss in networked Cournot competition (Proposition 4.2), we show that a large class of games remain efficient in the presence of producer anticipation. In fact, the illustration for Proposition 4.2 would be efficient if the producers anticipate. An essential part of the proof involves showing the equivalence of the networked Stackelberg game with conditions as described in the first part of the theorem to a single market Cournot competition with linear inverse demand and applying known results to the single market Cournot competition yield the results on existence and uniqueness of equilibrium and the corresponding efficiency loss.

However, we next show that the conditions listed, i.e., homogeneous price intercepts and no network restrictions on transportation, in the first part of Theorem 4.1 are “almost necessary”. These networked games with anticipation can easily lead to either non-existence of equilibria or arbitrarily large losses and can be caused by dropping either of the two conditions, highlighting the “almost necessary” nature of these conditions. Without network constraints, we show that the networked Stackelberg game is, in general, also equivalent to a single market Cournot competition but with convex, piecewise linear, inverse demand.

Efficient Anticipation in Networked Markets

By anticipating, producers understand the impact of their decisions better, and also become aware of how their actions couple with the production quantities of the other firms. The first part of Theorem 4.1 is an example where anticipation in these games only serve as the knowledge that platform \mathcal{P} couple the markets in the network, consequently reducing producer “market power”.

A key ingredient in proving the first result in Theorem 4.1 is to show that under such properties, the platform is allowed to keep prices constant across nodes in the network perpetually. This effectively couples the price elasticity to demand at every node, reducing the impact of manipulation on prices, and anticipation in this setting only serves to help producers understand their inability to price manipulate. To be precise, suppose the platform’s allocation is not constrained and that price intercepts across markets are homogeneous, i.e., $\alpha_j = \alpha_k$ for all nodes j, k , then the networked Stackelberg game is equivalent to a single Cournot competition with the same price intercept α but with a price elasticity to demand $\beta = (\sum_j \beta_j^{-1})^{-1}$.

It is hard not to notice the semblance of this harmonic mean to the effective resistance of resistors placed in parallel. This means that in the absence of network constraints, the homogeneous price intercepts allow for the platform to increase the “effective resistance” to price manipulate in these markets.

Such an “aggregation” also implies that participating in such a networked market is equivalent to participating in a single Cournot competition, with linear inverse demand, as studied previously in (Ramesh Johari and John N Tsitsiklis, 2004). We state an important result from the literature that is directly applicable to the networked Stackelberg game described in the first part of Theorem 4.1.

Theorem 4.3. *Consider a Cournot oligopoly with linear inverse demand functions and where producers have convex cost. The corresponding Cournot game admits a unique Cournot equilibrium \mathbf{q}^C with bounded worst case efficiency loss of $3/2$ as compared to the social optimal profile \mathbf{q}^* , i.e.,*

$$SW(\mathbf{q}^*) \leq \frac{3}{2} SW(\mathbf{q}^C).$$

Additionally, the worst case cost functions when optimizing for worst case efficiency loss are linear. Precisely, for any game with arbitrary convex costs, there exists a Cournot game except with linear cost functions that have efficiency loss no less than the original game.

Since the networked Stackelberg game in question is equivalent to a single market Cournot oligopoly with linear inverse demand, a straightforward application of Theorem 4.3 yields the uniqueness and existence of equilibria result, and the desired efficiency loss bound.

Additionally, since the worst case producer cost functions under this setting are linear ones, analysis over linear cost functions also extend to all convex cost functions. Also, our results derived in Chapter 3, e.g., the worst case efficiency loss bounds for (a)symmetric cost functions, can also be applied under these settings.

When platform actions are easy to understand, like in this setting, anticipation does little to no harm. In fact, when platforms retain sufficient “market power”, then anticipation only serve to remind the producers of it. By contrast, Proposition 4.2 showed that when producers fail to understand platform allocations and incentives, the networked market can instead be arbitrarily inefficient.

We show in the next subsection that the above results on bounded worst case efficiency loss do not extend beyond this setting, i.e., there are many settings where platforms no longer hold on to sufficient “market power”, and opportunities for anticipation and manipulation arise. The conditions listed in the first part of Theorem 4.1 are “almost necessary” in the sense that dropping either conditions will lead to results that are significantly different.

To be precise, we show in the following subsection that if there are network constraints, then neither existence nor uniqueness of equilibria necessarily holds. On the other hand, if the price intercepts across markets are not homogeneous, then the worst case efficiency loss is unboundedly large⁴. We will illustrate these in two-node examples of networks without either of the two conditions. These games will show the “almost necessary” nature of the conditions in the first part of Theorem 4.1.

Inefficient Anticipation in Networked Markets

In this part, we show that the results obtained from the first part of Theorem 4.1 in the previous section do not hold for networked competition with anticipation in general. In fact, these networked Stackelberg games have unbounded worst case efficiency losses. Further, as compared to networked Cournot games studied by (Bimpikis, Ehsani, and Ilklic, 2014; Abolhassani et al., 2014) which exhibits tractable analytic properties in uniqueness and existence of equilibria and bounded efficiency loss as shown in Chapter 3, the networked Stackelberg games are similar to the controlled

⁴We use a set-up similar to that in Section 3.4.

allocation platforms in Section 3.4, where there may not be an equilibrium, and is significantly more challenging to analyze and understand.

Even in cases where an equilibrium exists and is unique, the worst case efficiency loss can still be unboundedly large, as in Section 3.4. Further, anticipatory competition can also lead to non-existence of equilibrium. Our main result in this part is that the equilibrium of these games can be fragile, and that worst case efficiency losses are unboundedly large. We carefully construct the following two-node illustrative games to simultaneously illustrate the “almost necessary” nature of the conditions in the first part of Theorem 4.1 and the possibility of no equilibria, multiple equilibria, and unbounded worst case efficiency loss. In the case with no equilibria, we observe cycles under simple best-response dynamics. We adopt a simple best response dynamics where at each round, the firms first make decisions simultaneously under a fixed platform reallocation quantity, and the platform responds.

Consider the two-node illustrative game $\mathcal{G}([1, 1]^T, [1, 1]^T, [c_1 q^2, c_2 q^2]^T, f)$ further illustrated in Figure 4.2. We show that if flow constraint f is constrained within the following lower and upper bounds (corresponding to thresholds for (non-)existence of equilibria on a binding or non-binding constraint respectively), i.e.,

$$\frac{c_1 - c_2}{4c_1c_2 + 3c_1 + 3c_2 + 2} < f < \frac{4(c_1 - c_2)}{16c_1c_2 + 8c_1 + 8c_2 + 3},$$

then no equilibrium exists in the corresponding networked Stackelberg game. Note that the left plot in Figure 4.2 depict the scenario where ($f = 1/15$) and no equilibria exists, leading to a cycle in the graph. Such behavior is common when equilibria does not exist. The right plot has an equilibrium with flow constraints, where ($f = 1/10$). This illustration shows that even when price intercepts are homogeneous, the presence of network constraints can potentially cause a non-existence of equilibrium.

One may perhaps then suggest that if there are no network constraints, that networked competition is efficient with anticipation. We show this is not the case. We first show that the networked Stackelberg game without network constraints but not necessarily with homogeneous price intercepts, is equivalent to a single market Cournot oligopoly, except with convex (piecewise linear) inverse demand function, previously studied in (J. N. Tsitsiklis and Y. Xu, 2012). We formally prove the following result.

Proposition 4.4. *A networked Stackelberg game with linear inverse demands $\{\alpha_m - \beta_m d_m\}_{m=1}^M$ with $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_M$ is equivalent to a piecewise linear (convex)*




Figure 4.2: A plot of best-response dynamics for market demand fulfilled of the game $\mathcal{G}([1, 1]^T, [1, 1]^T, [2q^2, 1q^2]^T, f)$ for $f = 1/15$ (left) and $f = 1/20$ (right), where the bounds of f for non-existence of equilibria are $1/19 < f < 4/59$ respectively. The x-axis (resp. y-axis) represents the demand fulfilled at the first market d_1 (resp. d_2).

inverse demand function that links the following points: $(0, \alpha_1), (\frac{\alpha_1 - \alpha_2}{\beta_1^}, \alpha_2), (\frac{\alpha_1 - \alpha_2}{\beta_1^*} + \frac{\alpha_2 - \alpha_3}{\beta_2^*}, \alpha_3), \dots, (\sum_{j=1}^{m-1} \frac{\alpha_j - \alpha_{j+1}}{\beta_j^*}, \alpha_m), \dots, (\sum_{j=1}^{M-1} \frac{\alpha_j - \alpha_{j+1}}{\beta_j^* + \frac{\alpha_M}{\beta_M^*}}, 0)$, where β_m^* is the harmonic mean of all prior price elasticity β including its own, i.e.,*

$$\beta_m^* = \frac{1}{\sum_{j=1}^m \frac{1}{\beta_j}}.$$

Despite showing the equivalence to convex inverse demand functions, one cannot simply apply bounds derived in (J. N. Tsitsiklis and Y. Xu, 2012) since they are upper bounds to efficiency loss, and are not necessarily tight. Unlike the bounds in that work which are for Cournot candidates and therefore not necessarily tight for equilibrium, our results on efficiency loss are constructive. This implies that any of the losses we claim come with a precise construction of an equilibrium exhibiting that loss. In particular, when we consider a networked Stackelberg game without network constraints, the bounds in (J. N. Tsitsiklis and Y. Xu, 2012) say that “the game cannot be too inefficient”, while our results show that “the game is *at least* this inefficient in the worst case”.

Inspired by the controlled allocation platforms studied earlier, we construct a family of networked Stackelberg games to exhibit an efficiency loss that similarly increases linearly in the number of nodes. We present a two-node illustrative game that




Figure 4.3: (Left) Profit or utility function of the producer in the two-node illustrative game $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1+\epsilon)q]^T, \infty)$ in red with axis on left.

exhibits the necessary features, and an extension to a larger system with greater efficiency losses dependent on the number of nodes will also be described.

Consider the two-node illustration $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1+\epsilon)q]^T, \infty)$ for any $\lambda \in (0, 1/2)$, further illustrated in Figure 4.3. The zero-cost producer is indifferent between participating such that only one market or both are served because the market parameters are selected such that profit is $1/4$ in both cases. On the other hand, since there is a zero-cost producer, the optimal social welfare is obtained when all demand is fulfilled. As markets are added, the optimal social welfare increases, whereas the one at equilibrium remains the same. Contrasting to that, the $(1+\epsilon)q$ cost producer will never participate in any of the equilibrium nor in the social optimal since its marginal cost is always greater than that of the maximal willingness to pay in any of the markets. We design the market parameters in a particular manner such that the trapezoids under each piecewise linear region yield the same area (social welfare) even as we add more markets.

Observe in Figure 4.3 that the profits at both peaks in the left plot are equivalent at $1/4$, and the producer is indifferent between the two. Resulting inverse demand function that is piecewise linear, and convex, of the game, in blue with axis on right. Notice that the homogeneous intercepts scenario on the right plot is equivalent to a single market Cournot competition model with linear inverse demand. One may also make comparisons with the Cournot competition under this setting; in that case, the equilibrium depends on the position of the generator in the network, i.e., its co-located market. For example, if the generator is co-located with the market with largest (resp. smallest) willingness to pay, then the equilibrium is at the first (resp. last) peak.

In this section, we provide intuition on how we prove Theorem 4.1. We first show that under the conditions listed in the theorem, the networked Stackelberg game is equivalent to the well studied single market Cournot competition, providing us the existence and uniqueness of equilibria, and efficiency loss results. The conditions in Theorem 4.1 are idealized and inhibit producer “market power”, and anticipation serves as a reminder to producers of this limitation. However, these conditions fail to capture the setting of interest where network constraints affect prices and allocations. Further, we show through precise constructions in the second part of Theorem 4.1 that networked competition with anticipation is open to manipulation. This highlights the “almost necessary” nature of the conditions listed for efficient networked competition with anticipation.

In the next section, we show the impact of demand management both in limiting efficiency loss and in mitigating manipulations. It does so by allowing consumers to produce to other consumers, giving them “market power” to counter that of producers, revealing a new side to demand management.

4.4 Economic Value of Demand Management

In the previous section, we show how producer anticipation of a controlling platform or social planner can lead to various outcomes, dependent on how much “market power” is held by the platform. In this section, we explore the impact of load-side demand management, a.k.a. demand response in electricity markets, under platforms for networked marketplaces with producer anticipation. Recall that when consumers participate in demand management, the platform is no longer bound by the $q_m + r_m \geq 0$ constraint in each market m .

In this section, we prove the following theorem, which is the key result in this chapter, essentially bounding the worst case efficiency loss independent of price intercepts, even in the presence of binding network constraints, as long as consumers participate in demand management. It also provides conditions under which these reduction in demand is not actually necessary at equilibrium, but serve purely only as a threat to counter the “market power” of the producers.

Theorem 4.5. *Consider a networked Stackelberg game with demand management. Denote the set of edges with binding constraints at a Stackelberg equilibrium (q^S, r^S) by I^S . For each edge $(i, j) \in I^S$, let $f_{i,j}$ denote its edge constraint. The worst case*

efficiency loss at the Stackelberg equilibrium is bounded, i.e.,

$$SW(\mathbf{q}^*, \mathbf{r}^*) \leq \frac{3}{2} \left(SW(\mathbf{q}^S, \mathbf{r}^S) + \sum_{(i,j) \in \mathbf{I}} p_i^S f_{i,j} \right),$$

where p_i^S is the price at the outgoing node i at the Stackelberg equilibrium.

If the network is unconstrained, a unique equilibrium exists, and if weighted pairwise differences between consumer price intercepts α_m and $\alpha_M = \min_m \alpha_m$ are smaller than aggregate demand, i.e.,

$$\frac{1}{\sum_m \frac{1}{\beta_m}} \sum_{m=1}^M \frac{\alpha_m - \alpha_M}{\beta_m} \leq \sum_{m=1}^M q_m^S + r_m^S,$$

then demand management is unnecessary at the Stackelberg equilibrium, i.e., $q_m^S + r_m^S \geq 0, \forall j$.

We first show that when there are no network constraints on the platform allocation, the networked Stackelberg game with demand management has small efficiency loss. Essentially, this result follows the spirit to the first part of Theorem 4.1, first revealing its equivalence to a single market Cournot oligopoly, where the platform retains the bulk of the “market power”, manifested in the decreased and coupled price elasticity to demand across markets in the network. From there, results for existence and uniqueness of equilibria, and its worst case efficiency losses follow similarly as in Section 4.3. In addition to that, we provide conditions such that demand management plays the role of a threat but is not necessary at equilibrium, and highlight the impact on a two-node illustrative game.

Lastly, we conclude this chapter by showing that demand management in these networked games with anticipation is robust against network constraints. In particular, we show that these constraints at equilibrium have limited impact on its worst case efficiency loss. This also means that if network constraints exist but are not binding at equilibrium, then we retrieve the original 3/2 worst case efficiency loss bound. To show this, we first convert the networked Stackelberg game to a set of smaller networked Stackelberg games with properties that are easier to analyze. This conversion comes at a cost, manifested in the second term of the efficiency loss bound, which can be thought of as a *price of congestion*. Post-conversion, the smaller networked Stackelberg games, albeit with binding constraints, can be shown to each have a worst case efficiency loss of 3/2. We include another two-node illustrative game with an equilibrium that has binding line constraints to illustrate this.

Efficiency of Demand Management

It is well known that demand management tools like demand response in electricity markets help reduce supply-demand imbalance and allow for additional flexible generation or production in the network. In this section, we show that demand management has an important and critical role to play economically, both in terms of improving efficiency but also to discourage manipulations. In particular, we show that without network constraints, networked competition with anticipation and demand management have small worst case efficiency loss. This is because platforms again can increase the "resistance" of price manipulation in markets. Further, we provide conditions where demand management may be unnecessary at the equilibrium, and only serves as a threat against strategic curtailment.

Our first step is to show that networked competition with anticipation and demand management is again equivalent to a single market Cournot oligopoly. Suppose that the platform is social welfare maximizing with no network constraints. Without loss of generality, let $\alpha_1 = \max_m \alpha_m$. If consumers participate in demand management, then the system is equivalent to a single Cournot competition with maximal willingness to pay $\alpha = \left(\alpha_1 - \sum_m \frac{\alpha_1 - \alpha_m}{\beta_m} \right)$, and price elasticity to demand $\beta = \frac{1}{\sum_m \frac{1}{\beta_m}}$. The maximal willingness to pay has decreased, and the ability to couple the price elasticity to demand across markets leaves little room for anticipation and manipulation, with anticipation again serving as knowledge to the power of the platform.

Here, this implies that engaging Theorem 4.3 again provides similar results for uniqueness and existence of equilibria, and efficiency loss. The underlying idea is that demand management, in the absence of network constraints, couples the markets' price elasticity to demand, and reduces the impact and incentives for manipulation. Demand management is just a different way to allow for prices to remain constant across the markets when there are no network constraints. The case with constraints binding will be discussed in the following subsection, where we also develop worst case bounds there.

Oftentimes, platforms often would rather not have to engage demand management. The second part of Theorem 4.5 provide conditions under which demand management remains a threat that is not necessary at equilibrium. We show that if weighted pairwise comparisons of price intercepts α_j with the minimum over price intercepts α_M is not larger than aggregate demand, then demand management is not needed at equilibrium, i.e., $d_j \geq 0, \forall j$.

To illustrate this, we apply demand management on the construction in Section 4.3

with multiple equilibria, which initially suffered unbounded loss without demand management. Keeping all else constant, demand management essentially forces the producer to opt for the “highest production” equilibrium since the once “lowest production” equilibrium now yields a lower revenue and will no longer be an equilibrium. In particular, instead of a loss that grows linearly in the number of nodes, the loss here with demand management is $1 - \frac{1}{m}$, which decreases in m . One can observe that at the extreme ($m \rightarrow \infty$), the original networked Stackelberg game with anticipation incurs an unbounded loss, while the same game with demand management attains optimal social welfare.




Figure 4.4: Profit or utility function of the producer in the two-node illustrative game $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1+\epsilon)q]^T, \infty)$ in red with axis on left.

In Figure 4.4, observe that the original profits at both peaks are equivalent at $1/4$, and the producer is indifferent between the two. Similarly, the resulting piecewise linear, and convex inverse demand function of the game is in blue with the axis on the right. With demand management (dotted lines), the producer is now forced into the good equilibrium. One can compute that demand management is not active but serve as a threat against manipulation at the equilibrium.

This hints towards another impact of demand management—demand management reduces the number of (less favorable) equilibria. We posit that it does so by giving more “market power” to consumers, resulting in an effective decrease in “market power” on the producer side. Recently, two ridesharing giants Uber and Lyft also agreed to allow Google to display a price comparison on its Maps applications. Besides the fact that they rely on the Google Maps API heavily, we believe this comparison also serves as a form of “demand management” and a check against the “market power” of the drivers, consequently limiting their ability to manipulate in these markets.

Congestion in these systems though manifested in binding network constraints, can potentially lead to inefficiencies. Recall through the example of the non-existence of equilibrium that one can show that even in the homogeneous price intercept case, cases with network constraints can be hard to study. Can we then similarly quantify the efficiency loss due to binding network constraints with consumers participating in demand management? Can demand management also curb manipulation and limit efficiency loss in the presence of these network constraints?

Impact of Network Constraints

The positive results of demand response thus far assume that there are no network constraints binding. However, in practical scenarios, network constraints bind regularly, as is the case in electricity markets. In our previous two-node illustration with the non-existence of equilibrium, the lower bound dictates whether an equilibrium exists where constraints are binding. To deal with network constraints efficiently, platforms often resort to nodal pricing, e.g., locational marginal pricing, which further incentivizes manipulation and exacerbates the loss. Additionally, none of our results thus far provide satisfactory performance with binding network constraints.

To prove the second part of Theorem 4.5, we reconstruct a set of subsystems (networks of smaller sizes each with constraints only on edges incident on one leaf node) that are equivalent to the system with these binding constraints. The set of subsystems have both an equilibrium and an optimal social profile with welfare slightly larger (the difference is bounded by the revenue passing through the constrained line.) but only have binding edges on leaf nodes. We show that these edges are bound for both the Stackelberg equilibrium and the optimal social profile, and the multiplicative worst case efficiency loss for each subsystem is $3/2$. The additive loss is due to the change in welfare when we break the system down into subsystems and can be seen as a *price of congestion*. For the Stackelberg equilibrium, we show that the previous profile must remain an equilibrium, while such conversions only lead to an increased optimal social welfare.

This result implies that the market remains efficient when network constraints are not binding on large transmission lines. Alternatively, if the social welfare obtained in the entire system is large relative to the revenue flowing through constrained lines, then there is room to accommodate some loss from these binding network constraints.

To illustrate this, consider again the two-node game $\mathcal{G}([1, 1]^T, [1, 1]^T, [c_1, c_2]^T, f)$

used to illustrate the non-existence of equilibria. One can show that if the flow constraint f is constrained within

$$f < \frac{\frac{1}{4} \left(\frac{(c_1+1)-(c_2+1)}{(c_1+1)(c_2+1)} \right)}{1 - \frac{1}{4} \left(\frac{1}{c_1+1} + \frac{1}{c_2+1} \right)},$$

then an equilibrium exists in the networked Stackelberg game, but is constrained. For concreteness, consider $c_1 = 2, c_2 = 1$, then the bound above is $1/19$.

Suppose $f = 1/20$, then we can compute the equilibrium to be $q_1^S = \frac{19}{120}$, $q_2^S = \frac{21}{80}$, and $r_1^S = \frac{1}{20} = -r_2^S$, and prices at the equilibrium are $p_1 = \frac{19}{24}$, $p_2 = \frac{63}{80}$.

Our new result allow us to bound the loss in such situations by

$$SW(\mathbf{q}^*, \mathbf{r}^*) \leq \frac{3}{2} \left(SW(\mathbf{q}^S, \mathbf{r}^S) + \frac{63}{3200} \right).$$

A nice and potentially important corollary to the theorem though is that if network constraints exist but is not binding at an equilibrium, then it preserves the efficiency results from the no network constraint case. The intuition behind the above result is that such an equilibrium would also be an equilibrium in the case whereby the network constraints \mathbb{X} is removed. If that is the case, then in the case without network constraints (where the socially optimal solution can only be larger), the price of anarchy is $3/2$. Since the optimal social value can only decrease with network constraints, the price of anarchy without any of the network constraints binding is also $3/2$.

We have seen that demand management circumvents many of the shortcomings faced when competition is purely anticipatory. By offering "market power" to consumers, demand management mechanisms effectively reduce the "market power" of producers. Beyond securing and stabilizing the power network, we show that demand management can also play a big role in terms of economic efficiency, reducing the multiplicity of equilibria and potentially serving as a threat against manipulation and a check against "market power". Additionally, efficiency loss remains bounded in networked Stackelberg games with demand management even in the presence of binding network constraints.

Concluding Remarks

Demand management has been extensively used in electricity markets and is known to decrease variance in locational marginal prices (or shadow prices), improve stability of the electricity grid by reducing supply-demand imbalance and accommodate

increasing uncertainty from greater penetration of renewable energy. There are also multiple benefits for engaging in demand management, both for the participant, and the network at large.

In this chapter, we study the economic impact of demand management in networked marketplaces with governing platforms or market makers, e.g., electricity markets and its independent system operators. We explore this using a networked competition model that accounts for producer anticipation on the platform's allocations and network constraints, known as the networked Stackelberg competition.

This anticipation can be helpful when the platform has sufficient power and control but can be harmful when producers have sufficient "market power", leading to price manipulation. We first show conditions when the networked Stackelberg game have bounded efficiency loss. In particular, when the price intercepts are homogeneous across networked markets without flow constraints, we show that a unique equilibrium exists, and have bounded multiplicative loss of $3/2$. The conditions listed allow for the platform to couple and aggregate the networked marketplace, decreasing the opportunities and incentives for manipulation, and limiting the impact of anticipation. However, there exists a general class of networked Stackelberg games which do not fall under this setting and can either have non-existence of equilibria or unboundedly large efficiency loss. As a result, we find the conditions to be "almost necessary".

We show that demand management mechanisms in these networked games with anticipation shares the "market power" of producers with consumers, thereby reducing opportunities for price manipulation and incentives for collusion. Without network constraints, games with demand management have bounded multiplicative loss of $3/2$, while binding network constraints incur an additional additive loss bounded by the revenue that flows through these constrained lines. Additionally, we show that demand management may often be used as a threat but may not be necessary at an equilibrium.

*Chapter 5*TOWARDS DYNAMIC PARTICIPATION IN ONLINE
PLATFORMS

Thus far, Chapters 3 and 4 consider settings with full or perfect information, where participants make decisions privy to cost and demand at various nodes in a static game. In this chapter, we take the first step towards understanding uncertainty in markets. To do so, we consider an online optimization problem where the decision maker has a fixed amount of inventory to sell over markets with uncertain future conditions.

We show that this is a generalization of a classical problem known as the one-way trading problem. In that problem, a user has a fixed amount of inventory that he needs to sell in markets with different prices at each time period and is concerned with obtaining a conservative amount of revenue in the presence of uncertain (but bounded) prices.

“In the one-way trading problem a trader is given the task of trading dollars to yen. Each day, a new exchange rate is announced and the trader must decide how many dollars to convert to yen according to the current rate. The game ends when the trader trades his entire dollar wealth to yen and his payoff is the number of yen acquired.” - (El-Yaniv et al., 2001)

We focus more generally on an important class of online optimization problems that we term *online optimization under inventory (budget) constraints (OOIC)*. In these problems, a decision maker has a fixed amount of inventory, e.g., drivers on ridesharing markets with limit on driving hours or battery owners participating in power contingency reserves market, and must make an irrevocable decision in each of T rounds with the goal of optimizing aggregate revenue. The challenge is that the decision maker does not have knowledge of future revenue functions nor when the final round will occur, i.e., T . Further, the strict inventory constraint means that action now has consequences for future rounds. As a result of this entanglement, positive results have only been possible for inventory constrained online optimization

in special cases to this point, e.g., the one-way trading problem (El-Yaniv et al., 2001).

Beyond the one-way trading problem, the general formulation of OOIC also captures a variety of other practical applications with regards to real-time selling. Two examples that have motivated our interest in OOIC are (i) power contingency reserve markets (Shafiee et al., 2018; Akhavan-Hejazi and Mohsenian-Rad, 2014) and (ii) network spectrum trading (Bogucka et al., 2012; L. Qian et al., 2011). We delve deeper into each of the examples here.

In power contingency reserve markets, the system operator faces a contingency, e.g., shortfall of supply that may lead to cascading blackouts, and communicates this need to either supplement the power system using battery or cut down large scale power supply. Consider the perspective of a battery supply owner¹ that is deciding when to take part in a contingency. A contingency may be solved immediately, or it may instead cause a larger contingency whereby the system operator is perhaps willing to pay more at a later time. In preparation to participate in these contingencies, batteries are charged earlier and therefore the marginal cost of participation manifests as an opportunity cost against future participation in the day. These situations highlight the need for the online properties considered in our work: (i) the unknown ending time T , (ii) future revenue functions are not known, and (iii) a costless, strict inventory constraint.

Similarly, in spectrum trading, the owner of a spectrum band sells bandwidth to make sure that profit or revenue is maximized given the investments that have already been made to procure the particular bandwidth. This means that any cost with regards to sales only appears as opportunity cost against future possible sales. Similarly, a potential buyer who is turned down may seek bandwidth from a different provider, and may never return, or situations may change between time periods, highlighting the same three properties as before: (i) the unknown ending time T , (ii) future revenue functions are not known, and (iii) a costless and strict inventory constraint.

In this chapter, we develop a new algorithmic framework, and apply it to develop online algorithms for the OOIC problem with an optimal competitive ratio (up to a problem-dependent constant factor). Further, we prove that the algorithm provides the first positive results for a generalization of the classical one-way trading problem with concave revenue functions and price elasticity, which coincidentally aligns with

¹An alternative is flexible loads which may be curtailed to respond to supply shortfalls. Similarly, responding at the current moment may be an opportunity cost to later participation.

the price formation in the Cournot competition models considered earlier. In more detail, we summarize our contributions as follows.

Firstly, we introduce a new algorithmic framework, CR-Pursuit, in Section 5.3. The framework is based on the idea of “pursuing” an (in some cases, almost) optimal competitive ratio. The framework is parameterized by a tight upper bound on the competitive ratio, which is then “pursued” with the actions in each round. We apply the framework to OOIC and generalizations of the classic one-way trading problem in this paper, but the framework has the potential for broad applicability beyond these settings as well. For online platform design, the CR-Pursuit framework presents an almost optimal online optimization for selling a limited inventory under uncertainty.

Next, in Section 5.4, we apply CR-Pursuit to the OOIC problem to achieve the optimal competitive ratio among all deterministic algorithms (up to a problem-dependent constant factor). To obtain these bounds we use two technical ideas that may be of general interest beyond the OOIC problem. First, we prove that it suffices to focus on the single-parametric CR-Pursuit algorithm for achieving optimal competitive ratio, thus significantly reducing the search space of optimal online algorithms. We also identify a “critical” input sequence that highlights an important structural property of the space of input sequences. By applying CR-Pursuit to this sequence, we characterize a lower bound on the optimal competitive ratio as $\ln \theta + 1$ where θ is the ratio between the highest and lowest possible maximal willingness to pay. For any other input, the performance ratio achieved by CR-Pursuit is upper bounded by the product of a problem-dependent factor and the stated lower bound. This structure not only suggests a principled approach to characterizing the optimal competitive ratio, but also immediately shows that CR-Pursuit achieves the optimal competitive ratio (up to a problem-dependent factor) among all deterministic algorithms.

Lastly, we apply CR-Pursuit to various one-way trading problems in Sec. 5.5. The novel framework simplifies and unifies the state-of-the-art results of classic one-way trading problems. In particular, the critical input discussed above is indeed the worst-case one for classical one-way trading; hence, CR-Pursuit achieves the optimal competitive ratio $\ln \theta + 1$. Further, we show that CR-Pursuit performs well for generalizations of one-way trading where no positive results were previously known. Specifically, for one-way trading with price elasticity and concave revenue functions, CR-Pursuit achieves a competitive ratio that is within a small additive constant (i.e., $1/3$) to the general lower bound of $\ln \theta + 1$.

5.1 Literature Review

Online optimization is a large and rich research area and excellent surveys can be found in (Albers, 2003; Fiat, 1998). Well-known problems in the online optimization paradigm include the age-old secretary problem (Chow et al., 1964), the ski rental problem (Karlin et al., 1988), the one-way trading problem (El-Yaniv et al., 2001), and the k -server problem (Fiat, Rabani, and Ravid, 1990). Our results represent the most general results to date for a situation where actions are subject to a fixed inventory constraint.

The problem considered here is a generalization of the classical one-way trading problem, which has received considerable attention, e.g., (El-Yaniv et al., 2001; Chin et al., 2015; Damaschke, Ha, and Tsigas, 2009; Fujiwara, Iwama, and Sekiguchi, 2011; Lorenz, Panagiotou, and Steger, 2009; W. Zhang et al., 2012). In the one-way trading problem an online decision maker is sequentially presented with exchange rates within a bounded region, and she desires to trade all her assets to another. The amount of assets traded in a single time period is assumed to be small enough to not affect the eventual price. El-Yaniv et al. propose a threshold-based online algorithm with competitive ratio $O(\ln \theta)$ (El-Yaniv et al., 2001). Any remaining items must be sold at the last period as that is revenue maximizing. On the other hand, our analysis allows for leftover inventory (since selling all assets at the last time step may not be the revenue maximizing solution for the last time step in the presence of price elasticity or concave revenue functions) and an unknown stopping time, while retaining the attained competitive ratio.

Variants of the one-way trading problem have been studied in the literature. The one-way trading problem has been studied with unbounded prices and time-varying price bounds, respectively in (Chin et al., 2015) and (Damaschke, Ha, and Tsigas, 2009). It has also been studied when every two consecutive prices are interrelated (W. Zhang et al., 2012). Average-case competitive analysis under the assumption that the distribution of the maximum exchange rate is known has also been studied for the same problem (Fujiwara, Iwama, and Sekiguchi, 2011). (Kakade et al., 2004) incorporate market volume information and study another one-way trading model in stock market, called the price-volume trading problem. While the classical one-way trading problem mostly deals with linear revenue functions, we note that in our problem we consider general concave revenue functions, which allow us to capture a broader class of interesting settings, e.g., one-way trading with price elasticity or markets with linear inverse demand.

Beyond the one-way trading problem, OOIC is also highly related to generalizations of the secretary problem and prophet inequalities, e.g., (Rubinstein, 2016; Feldman and Zenklus, 2018; Babaioff et al., 2008). Strong positive results have been obtained for these problems; however the analytic setting considered differs dramatically from our work. Specifically, we consider a worst-case analysis whereas analysis of the secretary problem and prophet inequalities focus on stochastic instances. Under the stochastic setting, so-called “thresholding” algorithms are effective; however such algorithms have unbounded competitive ratios in the worst-case setting, even under the simplest assumptions.

Prior to this work, the most general results known for online problems with inventory constraints are for the class of problems termed online optimization with packing constraints, e.g., (Buchbinder and J. Naor, 2005; Buchbinder, J. S. Naor, et al., 2009; Azar et al., 2016; Arora, Hazan, and Kale, 2012; Bansal, Buchbinder, and J. S. Naor, 2012). This stream of work developed an interesting algorithmic framework based on a primal-dual or multiplicative weights update approach, which centers around maintaining a dual variable for each constraint, understood as a shadow (or pseudo) price for the constraint given the information thus far. While the inventory constraints we consider are packing constraints, our formulation is fundamentally different than the formulation considered in these papers. In these papers, the constraints come in an online fashion; whereas in our work, the revenue functions are arriving in an online fashion.

Another related online optimization problem is the k -search problem, where a player is searching for the k highest prices in a sequence that is revealed to her sequentially. When $k \rightarrow \infty$, the k -max search problem becomes the one-way trading problem (Lorenz, Panagiotou, and Steger, 2009). Lorenz et. al. propose optimal deterministic and randomized online algorithms for both the k -max search and k -min search problem (Lorenz, Panagiotou, and Steger, 2009).

That is different from the well-known k -server problem, where an online algorithm must control the movement of k servers in a metric space to minimize the movement (or latency involved) in serving future requests. A popular algorithmic framework for the k -server problem is the *potential function framework*. In contrast to our CR-Pursuit approach, the potential function approach requires a bound between the offline optimal cost and the online cost at each time period with respect to the potential. Instead, the CR-Pursuit maintains a bound between the current offline optimal revenue and current total online revenue.

Finally, it is important to distinguish our work from the literature studying *regret* in online optimization, e.g., (Hazan, A. Agarwal, and Kale, 2007; N. Chen et al., 2015). While regret is a natural measure for many online optimization problems, when inventory constraints are present it is no longer appropriate to compare against the best static action, as is done by regret. Static actions are poor choices when optimizing revenue subject to inventory constraints. Instead, competitive ratio is the most appropriate measure. Further, note that there is a fundamental algorithmic trade-off between optimizing regret and competitive ratio, even when inventory constraints are not present. In particular, (L. Andrew et al., 2013) shows that no algorithm can obtain both sub-linear regret and constant competitive ratio.

5.2 Problem Formulation

The key notations used in this chapter are summarized in Table 5.1. We study an online optimization problem where functions appear discretely and the stopping time T is unknown to the decision maker. The functions $g_t(\cdot)$ are revealed sequentially in a particular period t , and are unknown beforehand. Additionally, there is an inventory Δ that constrains the sum of the decision maker's actions, which is fixed and given in advance as a constraint. We emphasize here that in contrast to the line of online optimization work with packing constraints (Buchbinder and J. Naor, 2005; Buchbinder, J. S. Naor, et al., 2009; Azar et al., 2016; Arora, Hazan, and Kale, 2012; Bansal, Buchbinder, and J. S. Naor, 2012), the uncertainty in our optimization problem is not on the inventory constraints but the incoming revenue functions which are arriving in an online manner.

More specifically, at time $t \in [T]$ ², upon observing the revenue function at period t , the decision maker has to make an irrevocable decision on an action (quantity) v_t , with the objective of maximizing the aggregate revenue, while respecting the inventory constraint $\sum_{t \in [T]} v_t \leq \Delta$.³ Upon choosing v_t the decision maker receives a revenue of $g_t(v_t)$, where $g_t, \forall t \in [T]$ satisfy the following conditions:

²Throughout this chapter, we use $[n]$ to denote the set $\{1, 2, \dots, n\}$.

³The assumption that the action is chosen *after* observing the function differs from the classical online convex optimization literature (Hazan, A. Agarwal, and Kale, 2007; J. Li et al., 2012), but matches the literature on online convex optimization with switching costs (M. Lin et al., 2012; Bansal, Gupta, et al., 2015; Y. Li, Qu, and N. Li, 2018) and the literature on competitive algorithm design for online algorithm, including those on buy-or-rent decision making problems (Karlin et al., 1988; T. Lu, M. Chen, and L. L. Andrew, 2013; Ying Zhang et al., 2018) and metrical task systems (Borodin, Linial, and Saks, 1992; Fiat and Mendel, 2003; L. Lu et al., 2013). It allows an isolation of the inefficiency resulting from inventory constraints rather than also including the inefficiency resulting from the lack of knowledge of the function.

Table 5.1: Summary of Notations used in Chapter 5.

T	Total time slots
$g_t(v)$	Revenue function at time t
$\sigma^{[1:t]}$	Input sequence up to time t , i.e., $\{g_1, g_2, \dots, g_t\}$
$p(t)$	Base price at time t , i.e., $g'_t(0)$
m, M	lower/upper bound for $p(t)$
θ	fluctuation ratio M/m
λ	Dual solution
v_t	Output of online algorithm at time t
v_t^*	Output of optimal offline algorithm at time t
\hat{v}_t	The maximizer of $g_t(v)$
$\Phi_\Delta(\pi)$	Worst case (maximal) inventory over all possible sequences of inputs needed to maintain the competitive ratio $\pi \geq 1$ for CR-Pursuit(π) and a fixed inventory Δ

- $g_t(v)$ is concave, continuous and differentiable on $v \in [0, \Delta]$;
- $g_t(0) = 0$;
- $p(t) \triangleq g'_t(0) > 0$ and $p(t) \in [m, M]$.

The first condition is a smoothness condition on the revenue function and a natural decreasing marginal revenue assumption. The second condition implies that selling nothing yields zero revenue while the third condition limits the marginal revenue at the origin (named maximum willingness to pay hereafter) and ensures that it is beneficial to sell, since the maximum willingness to pay is positive. Denote the family of all possible revenue functions at time t as \mathcal{G} . We assume m and M are known beforehand in the online setting, as is standard in one-way trading. We denote $\theta = M/m$ as the ratio between the highest and lowest maximal willingness to pay. We note that in this paper, we only consider the domain when $g_t(v)$ is increasing in v , which intuitively means that selling more can never decrease revenue. For example, if $g_t(v) = (p_t - \alpha_t v)v$ as in the case of the one-way trading problem with linear price elasticity, then we only consider $v \in [0, \frac{p_t}{2\alpha_t}]$. This is because selling any more than the optimal amount at that time, e.g., $\frac{p_t}{2\alpha_t}$ in this example, will consume more resources while at the same time decrease aggregate revenue. For the original one-way trading problem example, these functions g are linear functions that passes through the origin.

In summary, an instance of online optimization under inventory constraints (OOIC) is formulated as follows:

$$\text{OOIC : } \max_{\mathbf{v}} \sum_{t=1}^T g_t(v_t) \quad (5.1)$$

$$s.t. \quad \sum_{t=1}^T v_t \leq \Delta, \quad (5.2)$$

$$v_t \geq 0, \forall t \in [T]. \quad (5.3)$$

Note that we can interpret the inventory constraint (5.2) in an OOIC in a parallel way to the inventory constraint in a one-way trading problem. In particular, in the one-way trading problem the trader has to decide in each slot the selling quantity v_t to maximize the total revenue at the stopping time T . In fact, when setting the family of functions \mathcal{G} to be the family of revenue functions of the form $g_t(v_t) = p(t)v_t$ we obtain the classical one-way trading problem. Additionally, when addressing revenue functions of the form $g_t(v_t) = v(t)(p(t) - f_t(v_t))$ where f_t is a convex function representing price elasticity, we obtain the generalized one-way trading problem with price elasticity.

To complete the specification of an OOIC, we state the Lagrangian of the OOIC problem here:

$$L(\mathbf{v}, \lambda, \mu) = \sum_{t=1}^T g_t(v_t) + \lambda(\Delta - \sum_{t=1}^T v_t) + \sum_{t=1}^T v_t \mu(t), \quad (5.4)$$

where $\lambda \geq 0$ and $\mu(t) \geq 0, \forall t \in [T]$ are Lagrangian multipliers.

To study the performance of an algorithm for OOIC we use the *competitive ratio* as the metric of interest. Note that many papers in the online optimization literature, e.g., (Hazan, A. Agarwal, and Kale, 2007), focus on *regret* instead of competitive ratio, but regret is not an appropriate measure when inventory constraints are considered since static actions are no longer appropriate⁴. Formally it can be defined as follows. Denote a deterministic online algorithm as \mathcal{A} , then \mathcal{A} is called *π -competitive* if

$$\pi = \max_{\sigma \in \Sigma} \frac{\eta_{OPT}(\sigma)}{\eta_{\mathcal{A}}(\sigma)}, \quad (5.5)$$

⁴Our focus on competitive ratio matches that of the literature on secretary problems e.g., (Babaioff et al., 2008; Rubinstein, 2016), prophet inequalities e.g., (Hajiaghayi, Kleinberg, and Sandholm, 2007; Rubinstein, 2016), online optimization with switching costs e.g., (M. Lin et al., 2012; L.Lu et al., 2013; T. Lu, M. Chen, and L. L. Andrew, 2013; Bansal, Gupta, et al., 2015), etc.

where Σ is the set of all possible inputs $(T, g_t(\cdot), t \in [T])$, and $\eta_{OPT}(\sigma)$ and $\eta_{\mathcal{A}}(\sigma)$ are the revenues generated by the optimal offline algorithm OPT and the online algorithm \mathcal{A} , respectively. The value π is the competitive ratio of algorithm \mathcal{A} .

The comparison against an offline optimal solution necessitates the analysis and understanding of the offline optimal solution of the OOIC problem, where the input sequence of functions σ and stopping time T are known beforehand to the decision maker.

Understanding the Offline Optimal solution

It is easy to check that the offline version of an OOIC is a convex optimization problem; thus it can be efficiently solved. Beyond this observation, we can also give a more efficient solution. More specifically, in the offline setting, both T and $g_t(\cdot), \forall t \in [T]$ are known in advance to the decision maker. By investigating the KKT conditions of the problem and exploring the “water-filling” structure of the optimal solution, we propose a binary-search based algorithm to obtain the dual solution. For ease of presentation, we denote λ^* as the optimal dual solution of problem OOIC and denote

$$V_t(\lambda^*) \triangleq \{v | g'_t(v) = \lambda^*, v \in [0, \Delta]\}, \quad (5.6)$$

the corresponding set of values of v in each time t that matches the optimal dual solution. The optimality of the primal and dual solutions are stated in the following proposition.

Proposition 5.1. *The optimal dual solution λ^* can be obtained by the binary search algorithm and the optimal primal solution satisfies*

$$v_t^* \in V_t(\lambda^*) \quad (5.7)$$

$$\sum_{t=1}^T v_t^* = \Delta \quad (5.8)$$

Note that in (5.7), if $V_t(\lambda^*) = \emptyset$ we set $v_t^* = 0$. Here, the Lagrange multiplier λ^* can be interpreted as the marginal cost (shadow price) of the inventory. In essence, we should only sell in the slots in which the base price, i.e., $p(t)$, exceeds the marginal cost determined by the optimal dual variable. Note that when $v_t > 0$, the marginal revenue at epoch t is $g'_t(v_t)$. Thus (5.7) essentially implies that the marginal revenue equals to the marginal cost (i.e., λ^*) in slots that the selling quantity is positive.

However, in practice, when a decision is made, the future revenue functions and stopping time T are not known, and robust decisions have to be made. One consequence is that these decisions are made essentially assuming that there is a possibility that the current time period is the stopping time T . Thus we are interested in the online setting where we do not assume any further distributional information on the revenue functions and decisions are irrevocable.

5.3 CR-Pursuit Algorithmic Framework

The class of online algorithms that make up the CR-Pursuit framework can be described as follows, where we denote $\sigma^{[1:t]} \triangleq \{g_1, g_2, \dots, g_t\}$ as the input up to time t and, under $\sigma^{[1:t]}$, we denote the optimal offline revenue as $OPT(\sigma^{[1:t]})$. Given any $\pi \geq 1$, at time t , the online algorithm (called CR-Pursuit(π)) chooses an action v_t that satisfies:

$$\frac{OPT(\sigma^{[1:t]})}{\pi} = \eta^{t-1} + g_t(v_t), \quad (5.9)$$

where η^{t-1} is the revenue of the online algorithm CR-Pursuit(π) up to time $t - 1$. Clearly, we have $\eta^0 = 0$ and

$$\eta^t = \eta^{t-1} + g_t(v_t). \quad (5.10)$$

Essentially, (5.9) and (5.10) imply that the online algorithm CR-Pursuit(π) tries to keep the *offline-to-online revenue ratio* at each slot to be π , i.e., we have $\forall t \in [T]$, $OPT(\sigma^{[1:t]})/\eta^t = \pi$.

While CR-Pursuit is defined for any arbitrary competitive ratio bound π , it may not be feasible for all bounds⁵. This limitation motivates the following definition.

Definition 5.2. For CR-Pursuit(π) with $\pi \geq 1$, if for any T and $\sigma^{[1:T]}$, we have $\sum_{t=1}^T v_t \leq \Delta$, then we say CR-Pursuit(π) is feasible. Otherwise, it is infeasible.

If CR-Pursuit(π) is feasible, i.e., it can maintain the ratio π under all possible sequence and inventory constraint inputs, then we know that it is at least π -competitive⁶. Intuitively, if π is large, the left-hand-side in (5.9) (i.e., $OPT(\sigma^{[1:t]})/\pi$) is small. In this case, CR-Pursuit(π) only needs to sell a small amount of inventory to maintain the revenue ratio. (Recall that we only consider the domain when $g_t(v)$ is increasing in v .) In order to maintain the revenue ratio to be π , at time $t \in [T]$, the online algorithm need to consume inventory v_t .

⁵One may not be able to pursue an over-optimistic competitive ratio.

⁶We note that there always exists a π (large enough) that can guarantee the feasibility of CR-Pursuit(π).

A careful reader may note that the inventory v_t at a time period t may not suffice to pursue the competitive ratio, then causing another potential infeasibility. Here we present the following proposition to handle this case and allow us to focus on feasibility based on sufficiency of total inventory in the remainder of the paper.

Proposition 5.3. *We have $OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \leq g_t(\hat{v}_t)$, where $\hat{v}_t \in [0, \Delta]$ is the maximizer of $g_t(v)$.*

Having defined the algorithmic framework and assured that, given sufficient inventory, it is always possible to pursue the competitive ratio, we now present the main theorem of this work, which states that a characteristic equation exists for the optimality of the competitive ratio among all deterministic online algorithms.

Theorem 5.4. *For CR-Pursuit(π) with $\pi \geq 1$ and for a fixed inventory Δ , let $\Phi_\Delta(\pi) \triangleq \max_{\sigma^{[1:T]}} \sum_{t=1}^T v_t$ be the worst case (maximal) inventory over all possible sequences of inputs needed to maintain the competitive ratio π , where v_t is the output of CR-Pursuit(π) at time t under input $\sigma^{[1:T]}$. The characteristic equation for optimality is $\Phi_\Delta(\pi^*) = \Delta$, and π^* , the solution to the above equation, is the optimal competitive ratio among all the deterministic online algorithms for OOIC problem.*

Theorem 5.4 implies that it suffices to focus on the single-parametric CR-Pursuit algorithm for achieving optimal competitive ratio. In other words, the problem of designing an optimal deterministic online algorithm for OOIC problem reduces to finding a $\pi^* \geq 1$ that satisfies the characteristic function $\Phi_\Delta(\pi^*) = \Delta$ ⁷. We first present a useful observation that given a fixed inventory constraint Δ , among all CR-Pursuit algorithms with feasible competitive ratio, algorithms with a (strictly) smaller competitive ratio requires a (strictly) larger output at each time period.

Lemma 5.5. *For any fixed input $\sigma^{[1:T]}$, the output of a feasible CR-Pursuit(π) at time $t \in [T]$, is decreasing in π . Consequently, $\Phi_\Delta(\pi)$ is decreasing in $\pi \geq 1$.*

Lemma 5.5 follows almost naturally since the offline solution remains the same, and attempting to preserve a smaller competitive ratio requires you to always commit to a larger inventory output⁸.

⁷By definition, if $\Phi_\Delta(\pi) \leq \Delta$, then CR-Pursuit(π) is feasible and π -competitive while $\Phi_\Delta(\pi) > \Delta$ implies that CR-Pursuit(π) may be infeasible under certain input sequences.

⁸It also implies that if CR-Pursuit(π_1) is feasible, then any online algorithm CR-Pursuit(π) with $\pi \geq \pi_1$ is also feasible. Thus an upper bound on the optimal competitive ratio in this case will give us a feasible competitive online algorithm.

If we can obtain a closed-form expression of $\Phi_\Delta(\pi)$, then by setting $\Phi_\Delta(\pi) = \Delta$, we can obtain the minimum competitive ratio π^* such that $\text{CR-Pursuit}(\pi^*)$ is feasible and thus it is π^* -competitive. Additionally, as shown in Theorem 5.4, π^* is the best competitive ratio among all the deterministic online algorithms. To see this, we first present the following lemma, setting bounds on the increase in the offline optimal solution with any new sequence input, based on the shadow prices or marginal benefits at each time, manifested in the dual variables of the optimization problem.

Lemma 5.6. *Let λ_t and λ_{t-1} be the optimal dual variable (or current offline marginal prices) under $\sigma^{[1:t]}$ and $\sigma^{[1:t-1]}$, respectively, and let v_t^* be the optimal offline solution⁹ at slot t under $\sigma^{[1:t]}$. Given any input $\sigma^{[1:T]}$, at any time $t \in [T]$, we have the following inequalities:*

$$OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \geq g_t(v_t^*) - \lambda_t v_t^*,$$

and

$$OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \leq g_t(v_t^*) - \lambda_{t-1} v_t^*,$$

essentially bounding the difference in the offline optimal given a new input in the sequence.

Before we proceed to prove Theorem 5.4, we first present the following lemma revealing a useful structure of the worst-case input for CR-Pursuit. It implies that for any inventory Δ and for any pursuit algorithm based on a competitive ratio $\pi \geq 1$, the worst case input sequence causes CR-Pursuit to sell at non-decreasing marginal prices across the time periods. The structure will be used in the proof of Theorem 5.4.

Lemma 5.7. *There exists an input sequence $\sigma \in \arg \max_\sigma \sum_t v_t$ such that $g'_t(v_t)$ is non-decreasing in t , where g_t is the revenue function and v_t is the output of $\text{CR-Pursuit}(\pi)$ under σ .*

Lemma 5.7 suggests that such an input sequence that forces the pursuit algorithm to face non-decreasing marginal prices $g'_t(v_t)$ exists. As such, it suffices to focus on the input space where $g'_t(v_t)$ is non-decreasing in t ¹⁰.

⁹Note that this solution v_t^* is obtained by taking into consideration the entire input sequence thus far and should not be confused with the maximizer of the function g_t .

¹⁰Intuitively, these sequences require the pursuit algorithm to make inventory commitments earlier on in the time periods, and at lower prices, which result in significant inventory commitment at these non-decreasing marginal prices.

The proof of Lemma 5.7 is provided in the Appendix, but the idea of the proof, roughly speaking, is that if the input sequence is not as stated, then there exists a change of positions of revenue functions in the input sequence that eventually yield a new input sequence such that $g'_t(v_t)$ is non-decreasing in t and it has to be more worst-case than before. An important result we require to prove Lemma 5.7 is the following lemma which describes how the offline optimal solution changes when we interchange the position of g_τ and $g_{\tau+1}$ under a fixed input sequence.

Lemma 5.8. *Let $\tilde{\sigma}$ be an input sequence. $\bar{\sigma}$ is another input sequence constructed by interchanging g_τ and $g_{\tau+1}$ in $\tilde{\sigma}$. We claim that*

$$OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}) \geq OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau]}). \quad (5.11)$$

Lemma 5.8 basically states that regardless of the input sequence thus far, the impact or improvement in the offline optimal that it brings at that point has "diminishing returns" in time. This helps us identify the worst case input sequence in Lemma 5.7 where the marginal benefits is non-decreasing across time periods. Putting the above lemmas together, we are now ready to prove our main result in Theorem 5.4.

Proof of Theorem 5.4. Consider an arbitrary deterministic online algorithm, denoted as \mathcal{A} . We show that \mathcal{A} cannot achieve a ratio smaller than π^* .

For \mathcal{A} and CR-Pursuit(π^*), denote the output at time t as $v_t^{\mathcal{A}}$ and v_t , respectively. For ease of presentation, denote $\tilde{\sigma}^{[1:T]} = \{\tilde{g}_1, \tilde{g}_2, \dots, \tilde{g}_T\}$ as the worst case input sequence for CR-Pursuit(π^*), i.e., under this input, we have $\sum_{\tau=1}^T v_\tau = \Phi_\Delta(\pi^*) = \Delta$.

From Lemma 5.7, we claim that $\tilde{g}'_t(v_t)$ is non-decreasing in t , i.e., $\tilde{g}'_t(v_t) \leq \tilde{g}'_{t+1}(v_{t+1})$, $\forall t$. We present \tilde{g} sequentially to \mathcal{A} . If \mathcal{A} commits at any time a larger inventory than CR-Pursuit, then we calculate the offline-to-online ratio that \mathcal{A} is now pursuing. Since CR-Pursuit cannot attain that offline-to-online optimal ratio, it means that at some point along the sequence, \mathcal{A} has to have at least offline-to-online ratio equivalent to π^* . On the other hand, if \mathcal{A} reaches an offline-to-online ratio larger than π^* at any time, then we set that as the stop time T , and we are done.

Thus the competitive ratio of \mathcal{A} should at least be π^* . It follows that \mathcal{A} must coincide with CR-Pursuit(π^*), achieving a ratio of π^* , or otherwise \mathcal{A} incurs a higher ratio on $\tilde{\sigma}^{[1:T]}$. \square

5.4 Competitive Analysis of CR-Pursuit

The results in the previous section highlight that the key to applying the CR-Pursuit algorithmic framework is to mathematically characterize $\Phi_\Delta(\pi)$. For special cases such as the one-way trading problem (El-Yaniv et al., 2001) where $g_t(v) = p_t v$ is a linear function on v , we can obtain close-form expressions of $\Phi_\Delta(\pi)$ and compute the optimal competitive ratio (as demonstrated in Section 5.5). However, it is difficult in general to obtain such closed-form expressions for more general families of concave revenue functions. Instead, we characterize an upper bound on $\Phi_\Delta(\pi)$, which gives an upper bound on the optimal competitive ratio, and consequently a feasible online CR-Pursuit algorithm.

The critical input also plays an important role in establishing the upper bound for $\Phi_\Delta(\pi)$. It turns out that for any other inputs, the performance ratio achieved by CR-Pursuit is upper bounded by the product of a problem-dependent factor and the lower bound achieved under the critical input. This insight leads to the following results.

Theorem 5.9. *For OOIC, the optimal competitive ratio π^* is upper bounded by $c(\ln \theta + 1)$, i.e., we have $\pi^* \leq c(\ln \theta + 1)$, where $c = \sup_{g \in \mathcal{G}} \frac{\hat{v} g'(0)}{g(\hat{v})}$, $\hat{v} = \arg \max_{v \geq 0} g(v)$.*

Theorem 5.9 characterizes an upper bound on the optimal competitive ratio in the case for general revenue functions g_t , and also implies that CR-Pursuit($c(\ln \theta + 1)$) is feasible and its competitive ratio is upper bounded $c(\ln \theta + 1)$. Note that c is a constant that depends on the gradient properties (a.k.a. base price) and the maximizers of the revenue functions¹¹. For many interesting problems, this c is bounded and small. For example, for the one-way trading problem where the revenue functions are linear, i.e., $g_t(v) = p(t)v, \forall t \in [T]$, we have $c = 1$. For another example, for the one-way trading with linear price elasticity where the revenue functions are quadratic, i.e., $g_t(v) = (p(t) - \alpha_t v)v, \forall t \in [T]$, we have $c = 2$. The proof of Theorem 5.9 is in the Appendix.

¹¹While c is a constant when the family of revenue functions are fixed, it is indeed true that c could presumably be driven to be infinitely large, e.g., with revenue functions that are concave and increasing. This parameter c can be seen in an economical sense as a comparison between the base price and the average price at the maximizer of the function. Since the former is already bounded in $[m, M]$, we look at the case when the latter is small. These situations are hard to derive any interesting online optimization as the functions require too much commitment even in bad time epochs, and have low average prices. This results in low committed average prices while the offline optimal may eventually not have to participate in these time epochs.

5.5 Application to One-way Trading

In this section, we apply the CR-Pursuit algorithmic framework to standard one-way trading (El-Yaniv et al., 2001) and its generalizations, illustrating that the framework can both match state-of-the-art results for the classic setting and provide new results for generalizations that have previously resisted analysis. In particular, using the CR-Pursuit framework, we obtain an online algorithm matching the optimal $(\ln \theta + 1)$ competitive ratio for the classic one-way trading problem (Theorem 5.11) and a near-optimal $(\ln \theta + 4/3)$ result for the case with linear price elasticity (Theorem 5.13)¹².

This section also provides an illustration of how the framework can be applied to specific problem domains to obtain tighter results that are possible for the general OOIC problem studied in the previous section. In particular, bounds may be obtained in general given gradient properties at the origin and optimal solutions of revenue functions, but tighter bounds can be obtained given a more specific family of revenue functions, e.g., for one-way trading with price elasticity ($c = 2$), the upper bound derived from Section 5.3 is $2(\ln \theta + 1)$ while the bound obtained in this section is $\ln \theta + 4/3$.

The one-way trading problem is a special case of the OOIC problem with $g_t(v_t) = p(t)v_t$ for all $t \in [T]$ and the input at time t can be simplified as $p(t)$. In Section 5.5, we obtain the close-form expression of $\Phi_\Delta(\pi)$ and compute the optimal π^* in this special case. Additionally, in Sec. We also show the ease of generalizing the one-way trading problem, to cases where price formation include price elasticity, an aspect that has been left out but desired in various communities.

Classic One-way Trading

As a direct application, one can obtain from Section 5.3 that the upper bound for the one-way trading problem is $\ln \theta + 1$, which matches the lower bound. Thus, we immediately know that the optimal competitive ratio for one-way trading is $\ln \theta + 1$ and CR-Pursuit($\ln \theta + 1$) obtains the best guarantee possible among deterministic online algorithms. In this section, we apply the other approach presented in Section 5.3, with the aim of demonstrating the possibility of mathematically characterizing $\Phi_\Delta(\pi)$ under specific conditions. In particular, in the following, we first compute the closed-form expression of $\Phi_\Delta(\pi)$, then proceed with similar analysis as in Section 5.3 to obtain the optimal competitive ratio.

¹²The algorithmic framework also extends to any convex price elasticity, and yield online algorithms with near-optimal competitive ratio in these cases.

In the one-way trading problem, given any input up to time t , denoted as $\sigma^{[1:t]} \triangleq \{p(1), p(2), \dots, p(t)\}$, the optimal offline revenue can be expressed as $OPT(\sigma^{[1:t]}) = \Delta \cdot \max \sigma^{[1:t]}$. Given any $\pi \geq 1$, consider the class of online algorithm CR-Pursuit(π) defined in Sec. 5.3. At time t , CR-Pursuit(π) sells the amount v_t that satisfies:

$$\frac{OPT(\sigma^{[1:t]})}{\pi} = \eta^{t-1} + p(t)v_t, \quad (5.12)$$

where η^{t-1} is the revenue of the online algorithm CR-Pursuit(π) up to time $t-1$. Clearly, we have $\eta^0 = 0$ and

$$\eta^t = \eta^{t-1} + v_t p(t). \quad (5.13)$$

Essentially, (5.12) and (5.13) imply that the online algorithm CR-Pursuit(π) tries to keep the offline-to-online revenue ratio at each slot to be π , i.e., we have $\forall t \in [T]$, $OPT(\sigma^{[1:t]})/\eta^t = \pi$.

In the following, our goal is to compute the close-form expression of $\Phi_\Delta(\pi)$. Observe that at slot t , the selling decision of CR-Pursuit(π^*) can be simplified as

$$v_t = \frac{(\max \sigma^{[1:t]} - \max \sigma^{[1:t-1]})\Delta}{\pi p(t)}.$$

This suggests that CR-Pursuit(π^*) will sell only when the current price is larger than the highest price in history. With this observation, we have the following lemma¹³.

Lemma 5.10. *For CR-Pursuit(π) with $\pi \geq 1$, given any input $\sigma^{[1:T]}$, to compute $\Phi_\Delta(\pi)$, it is sufficient to consider increasing price sequence.*

From Lemma 5.10, we know that it is sufficient to consider the following increasing price sequence with length $n \leq T$:

$$m \leq p_1 < p_2 < \dots < p_n \leq M. \quad (5.14)$$

Under the given increasing price sequence, the optimal offline revenue at time $t \in [n]$ can be simplified as $OPT(\sigma^{[1:t]}) = p_t \Delta$. According to (5.12), the output of CR-Pursuit(π) at time $t \in [n]$ is

$$v_t = \frac{OPT(\sigma^{[1:t]}) - \eta^{t-1} \pi}{\pi p_t} = \frac{(p_t - p_{t-1})\Delta}{\pi p_t},$$

¹³It can actually be seen as a corollary of Lemma 5.7.

where $p_0 = 0$. Then we have

$$\begin{aligned}\Phi_{\Delta}(\pi) &= \max_{p_1, p_2, \dots, p_n} \sum_{t=1}^n v_t \\ &= \max_{p_1, p_2, \dots, p_n} \frac{\Delta}{\pi} \left(1 + \frac{p_2 - p_1}{p_2} + \dots + \frac{p_n - p_{n-1}}{p_n}\right) \\ &\stackrel{(a)}{=} \frac{\Delta}{\pi} \left(1 + \int_m^M \frac{1}{x} dx\right) = \frac{\Delta}{\pi} (1 + \ln \theta),\end{aligned}$$

where (a) holds when the input sequence in (5.14) satisfies $n \rightarrow \infty$ and $p_i \rightarrow p_{i+1}, \forall i \in [n-1]$. Indeed, this is the worst-case input sequence for one-way trading problem, also known as the "critical" input sequence. Thus for $\text{CR-Pursuit}(\pi)$ to be feasible, the minimum possible π should satisfy $\Phi_{\Delta}(\pi) = \Delta$, which yields the solution that $\pi^* = \ln \theta + 1$. Consequently, we have the following result:

Theorem 5.11. *Let $\pi^* = \ln \theta + 1$, $\text{CR-Pursuit}(\pi^*)$ is $(\ln \theta + 1)$ -competitive.*

Using a similar technique as used in Section 5.3 and in (El-Yaniv et al., 2001), we can show that π^* is the optimal competitive ratio for one-way trading problem. We present the result in the following for completeness.

Theorem 5.12. *Any deterministic online algorithm for one-way trading problem has a competitive ratio that is no smaller than π^* .*

One-way Trading with Price Elasticity

In this subsection, we consider the one-way trading problem in a generalized setting with an additional flexibility on the price model playing the role of *price elasticity*. We assume that price is affected by the total quantity sold at each slot, implying that the decision of how much to sell affects the trading price, usually known in the economics literature as *price elasticity*.

Specifically, we assume that at each slot $t \in [T]$, the price elasticity ($\triangleq f_t(v)$) is a convex non-negative function of the selling quantity with $f(0) = 0$. Under this setting, the revenue function at time t becomes $g_t(v) = (p(t) - f_t(v))v$. This setting can be considered as a special case of *OOIC* and the input at time t can be simplified as $p(t), f_t(v)$. Here we have $g'(0) = p(t) \in [m, M]$ and $f_t(v) \in [0, +\infty), \forall v \in [0, \Delta], f_t(0) = 0$. Namely, the set of all possible revenue functions can be expressed as

$$\begin{aligned}\mathcal{G} &= \{g_t(v) | g_t(v) = (p(t) - f_t(v))v, p(t) \in [m, M], \\ &\quad f_t(v) \in [0, +\infty), \forall v \in [0, \Delta], f_t(0) = 0\}.\end{aligned}$$

Note that when $f_t(v) = 0, \forall t \in [T]$, the problem reduces to one-way trading problem considered in Section 5.5. Thus we note that any deterministic online algorithm in one-way trading with price elasticity has a competitive ratio of at least $\ln \theta + 1$. When $f_t(v) = \alpha_t v, \alpha_t \geq 0, \forall t \in [T]$, the problem becomes a one-way trading problem with linear price elasticity, which is the same market model as in Chapters 3 and 4.

Consider the online algorithm CR-Pursuit(π) defined in (5.9). In this case, it is difficult to obtain the closed-form expression of $\Phi_\Delta(\pi)$. Instead, we follow the analysis in Section 5.3 to obtain an upper bound on $\Phi_\Delta(\pi)$, and obtain the following result.

Theorem 5.13. *Let $\bar{\pi} = \frac{z^2}{2z-1}$, where $z = 2(1 + \ln \theta)$, then online algorithm CR-Pursuit($\bar{\pi}$) is $\bar{\pi}$ -competitive. In particular, that means that the competitive ratio of CR-Pursuit is $O(\ln \theta)$ as well.*

Proof. When $\bar{\pi} = \frac{z^2}{2z-1}$, we have $\bar{\Phi}(\bar{\pi}) = \Delta$. From Lemma C.5, we know that $\Phi_\Delta(\bar{\pi}) \leq \bar{\Phi}(\bar{\pi}) = \Delta$. Thus CR-Pursuit($\bar{\pi}$) is $\bar{\pi}$ -competitive. \square

Note that $\bar{\pi} = \ln \theta + 1 + \frac{\ln \theta + 1}{4 \ln \theta + 3} \in [\ln \theta + \frac{5}{4}, \ln \theta + \frac{4}{3}]$, thus the competitive ratio of CR-Pursuit($\bar{\pi}$) is very close to the lower bound. Further, this competitive ratio is better than the result in Sec. 5.3 which yields a $2(\ln \theta + 1)$ upper bound, since we have a tighter bound through Lemma C.4.

Concluding Remarks

We consider the problem of firms' optimal decision making under uncertain future information. To ground our problem, we consider a firm with a fixed inventory selling in uncertain markets in that revenue functions at each time is only revealed at that time. We show first that this generalizes the classical one-way trading problem.

We design a novel framework for online optimization called CR-pursuit, where algorithms pursue pre-defined competitive ratios. We show that such an algorithm presents an optimal algorithm in terms of minimizing competitive ratio in the class of problems that we consider.

In particular, our CR-Pursuit algorithm matches the state-of-the-art solutions for the one-way trading problem and is applicable for a wider class of problems. With regards to optimal decision making, we find that firms are both wary and hopeful of the future, hedging against it by making safe decisions at each time step.

Chapter 6

CONCLUSION AND FUTURE WORK

Online platforms are catalysts for a large bulk of social and economic interactions today. Platforms are not new, and have existed for a long time, but recent advances in communication technology allow for platforms to thrive under online settings. Driven by this success, studies of multi-sided platforms, networked competition and successful design features started surfacing alongside. Beginning with the seminal work of (Rochet and Tirole, 2003), research in this area have grown, in breadth but also in depth, e.g., designing fine-grained control in ride-sharing platforms through surge pricing in (A. Lu, Frazier, and Kislev, 2018). Many online platforms today still fail to attain the *critical mass* required for a successful platform that takes advantage of *network effects*, an idea that increase in scale on one side have benefits to both sides of the market. On the other hand, some older online platforms which are once successful are now on the decline. These platforms usually do not understand platform design well, nor identify loopholes for manipulation, or the impact of constraints due to regulation. *The goal of this thesis is to identify aspects of platform design that is critical to the success of online platforms.*

In order to accomplish this goal, we have identified three observations that can help us understand online platforms better, namely, (i) understanding successful online platform design features and how and when they balance transparency and control, (ii) recognizing potential abuse of market power and designing checks against such manipulations, and (iii) considering the impact of dynamic decision making by participants under limited information on the future. We elaborate on each of these observations here.

1. **Platforms make trade-offs between transparency and control.**

For example, platforms like eBay have open access designs and are fully transparent, while platforms like Uber exercise significant control and make allocation decisions on behalf of its participants. Meanwhile, platforms like Amazon seek a sweet spot between these two designs. Even with the same purpose in mind, platform designs can contrast drastically, e.g., eBay vs. Amazon. Obtaining the right balance depend on competition structures, e.g., transaction costs involved and heterogeneity of product.

2. **Platforms prevent manipulation by producers with market power.**

Platforms often have designs that allow them to retain most of the “market power”. For example, a study showed that the impact of producer anticipation in Uber relies heavily on its power and information asymmetry with the drivers. Further, in recent years, we have witnessed manipulation both in electricity markets, e.g., Enron in California, and ridesharing markets, e.g., coordinated manipulation to cause surge pricing. This happens because participants collude and exercise their joint market power in a negative manner, and network constraints over allocations exacerbate the effect of these manipulations. In designing platforms that control over allocations, preventing or having checks against such manipulation is critical.

3. **Platforms make better decisions when their participants do.**

Platforms often assume that their participants have complete information and make fully rational decisions. However, participants may have limited information of the future as compared to the platform. As such, participants in platforms often make decisions that are unexpected. When they do, then platform designs may turn out to be suboptimal too.

In this thesis, we provide a number of results in Chapters 3, 4 and 5 respectively, studying some of these observations, shedding light on the following issues:

1. **Trade-offs between transparency and control in platform design**

Designs of successful platforms can be very diverse—understanding when certain design features are efficient is important. We study three platform designs on a networked Cournot competition model, highlighting the trade-off between transparency and control. The first one is an open access design, exemplified by eBay, where all firms have access to all markets, and firms decide their production quantity to the markets jointly. The second design, exemplified by ridesharing platforms, is a controlled allocation design, where producers decide their aggregate production, but not its allocation. The final design, discriminatory access, seeks to balance transparency and control, providing restricted firm-market access, but firms retain control over the allocation over their connected markets. We show that open access can incur large search costs, while controlled allocation designs cause misaligned platform-producer incentives towards production. Lastly, we show that discriminatory access suffers from small inefficiencies but may be hard to design.

2. Demand management as a check against market power

“Market power” and anticipation often allows for producer manipulation which results in increased individual profits, but decreased system performance. Network constraints over the platform allocations exacerbate the “potential” of these manipulations, since nodal or marginal prices over the network can be controlled easily. We study the impact of demand management under a networked Stackelberg model, which accounts for producer anticipation but is also amenable to consider practical network constraints pertaining to the electricity grid. We first provide conditions for efficient anticipatory competition, i.e., if price intercepts are homogeneous over markets that are connected in an unconstrained manner, then a unique equilibrium exists and efficiency loss is bounded. Demand management serves as a powerful tool against anticipation, essentially giving consumers “market power” to counter that of the producers, limiting their ability for manipulations. In fact, we show that when consumers participate in demand management, then efficiency loss is bounded multiplicatively by $3/2$ and binding network constraints further incur an additional additive loss. When markets are not too heterogeneous, then demand management may serve as a threat that is not necessary at equilibrium.

3. Online decision making under inventory constraints

We have witnessed the impact of anticipation and the power of information. It is also possible that firms have to make decisions under uncertainty, and one source of doubt comes from the online nature of participation in platforms. With uncertain prices and demand in the future, firms may hedge against this by selling less in each market, being cautiously optimistic of the future. We show that the problem of online optimization to sell in markets under an inventory constraint is a generalization of the classical one-way trading problem. We develop the CR-Pursuit framework to consider this class of online optimization problems, showing that a simple pursuit of a competitive ratio is possible—and the problem simplifies to finding the optimal or a feasible competitive ratio to pursue. We show that our CR-Pursuit algorithm achieves state-of-the-art results for the original one-way trading problem and is applicable to generalizations that we consider. The analysis bears insights to the larger picture in this thesis, in that the optimal solution of firms with uncertainty can be far from the platform’s expectation. This will only be exacerbated in the presence of multiple firms.

Since we cannot present a complete overview of every section in this thesis, we provide a summary of the important lessons and interesting insights that we learn in each section.

6.1 Lessons and Insights

This thesis extends over multiple topics, and in each topic/section our results have provided us interesting insights about online platform design and considerations. In this section, we aim to conclude this thesis by providing high-level insights from each topic covered in this thesis. As we expound on each topic, we also provide reference to the corresponding section to obtain a fuller and more complete picture.

Open access designs perform well when consumer search costs are low.

Open access designs are transparent and easy to implement in practice. Through our analysis of open access platform designs under the networked Cournot competition, we find that they promote participation, i.e., the Nash demand fulfilled at each market is at least half that of the socially optimal profile (Lemma 3.7). When search costs are negligible, then the worst case efficiency loss of open access designs are small (Theorem 3.5). However, if search costs are high, then worst case efficiency loss can be arbitrarily large (Theorem 3.19).

Additional control does not necessarily lead to improved performance.

Control is delicate. We show that if producers can anticipate these control over allocations that maximize social welfare, then they tend to manipulate prices by under-producing for increased profits, which can lead to arbitrarily large efficiency loss (Theorem 3.11). We extend our analysis to consider different market clearing mechanisms (Theorem 3.14), showing how outcomes depend on them.

If you have to increase control, do it over access whenever possible.

Controlling over access is not easy. In our work, we show that it is equivalent to a Mathematical Program with Equilibrium Constraints (MPEC) which are, in general, NP-hard. However, the discriminatory access design has bounded worst case efficiency loss (Theorem 3.15), and amenable to balance other factors such as search costs. While the network optimization is difficult, applying greedy algorithms and heuristics may at times lead to optimality (Theorem 3.17).

Anticipation is a necessary evil.

Anticipation drives a platform. For incentives to be effective in promoting efficiency, the producers need to know them. Without anticipation, efficiency loss can be arbitrarily large even in small networks, regardless of the market clearing mechanism (Proposition 4.2). On the other hand, we provide conditions that jointly inhibit negative effects of anticipation. When we are able to restrict the “market power” of the producers, we can guarantee a bounded loss (Theorem 4.1). However, we also show that these conditions for efficiency under anticipation are “almost necessary”, in that removing either condition leads to potentially an absence of equilibrium or arbitrarily large losses.

Demand management give consumers “market power” to counter producers’.

Beyond stabilizing supply-demand imbalance across the market, demand management schemes like demand response have a large role to play economically too. We show that demand management restricts the “market power” of producers by actually empowering and giving consumers “market power” to counter that of the producers. By doing so, we show that demand management in the absence of network constraints have bounded loss, while it suffers a small *price of congestion* when network constraints are binding (Theorem 4.5).

Sometimes, simple pursuit can be optimal.

Making decisions under uncertainty can be difficult. For example, an Uber driver needs to make a decision when he should drive in the day. We show that these online optimization problems with inventory (driving time) constraints have online optimal (competitive ratio) solutions that are easy to follow (Theorem 5.4). For the Uber scenario, it involves finding what you would have earned up till this point and aim for at least a fraction (the target competitive ratio) of it. We show what these ratios are for different classes of revenue functions (Theorem 5.9).

BIBLIOGRAPHY

- Abolhassani, Melika et al. (2014). “Network Cournot competition”. In: *International Conference on Web and Internet Economics*. Springer, pp. 15–29.
- Abreu, Dilip and Mihai Manea (2012). “Bargaining and efficiency in networks”. In: *Journal of Economic Theory* 147.1, pp. 43–70.
- Afeche, Philipp, Zhe Liu, and Costis Maglaras (2018). “Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance”. In:
- Agarwal, Nikhil et al. (2018). *Market failure in kidney exchange*. Tech. rep. National Bureau of Economic Research.
- Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan (2017). “Thickness and information in dynamic matching markets”. In:
- Akhavan-Hejazi, H. and H. Mohsenian-Rad (Mar. 2014). “Optimal Operation of Independent Storage Systems in Energy and Reserve Markets With High Wind Penetration”. In: *IEEE Transactions on Smart Grid* 5.2, pp. 1088–1097.
- Albers, Susanne (2003). “Online algorithms: a survey”. In: *Mathematical Programming* 97.1-2, pp. 3–26.
- Andrew, Lachlan et al. (2013). “A tale of two metrics: Simultaneous bounds on competitiveness and regret”. In: *Conference on Learning Theory*, pp. 741–763.
- Anshelevich, Elliot and Shreyas Sekar (2015). “Price Competition in Networked Markets: How Do Monopolies Impact Social Welfare?” In: *Web and Internet Economics*. Springer, pp. 16–30.
- Armstrong, Mark (2006). “Competition in two-sided markets”. In: *The RAND Journal of Economics* 37.3, pp. 668–691.
- Arora, Sanjeev, Elad Hazan, and Satyen Kale (2012). “The Multiplicative Weights Update Method: a Meta-Algorithm and Applications.” In: *Theory of Computing* 8.1, pp. 121–164.
- Ashlagi, Itai et al. (2018). “On matching and thickness in heterogeneous dynamic markets”. In: *Available at SSRN 3067596*.
- Azar, Y. et al. (Oct. 2016). “Online Algorithms for Covering and Packing Problems with Convex Objectives”. In: *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 148–157.
- Babaioff, Moshe et al. (2008). “Online auctions and generalized secretary problems”. In: *ACM SIGecom Exchanges* 7.2.
- Bakos, J Yannis (1997). “Reducing buyer search costs: Implications for electronic marketplaces”. In: *Management science* 43.12, pp. 1676–1692.

- Banerjee, Siddhartha, Daniel Freund, and Thodoris Lykouris (2016). “Multi-objective pricing for shared vehicle systems”. In: *CoRR* abs/1608.06819. arXiv: 1608.06819. URL: <http://arxiv.org/abs/1608.06819>.
- Banerjee, Siddhartha, Sreenivas Gollapudi, et al. (2017). “Segmenting two-sided markets”. In: *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pp. 63–72.
- Banerjee, Siddhartha, Carlos Riquelme, and Ramesh Johari (2015). “Pricing in ride-share platforms: A queueing-theoretic approach”. In:
- Bansal, Nikhil, Niv Buchbinder, and Joseph Seffi Naor (2012). “A primal-dual randomized algorithm for weighted paging”. In: *Journal of the ACM (JACM)* 59.4.
- Bansal, Nikhil, Anupam Gupta, et al. (2015). “A 2-competitive algorithm for on-line convex optimization with switching costs”. In: *LIPICs-Leibniz International Proceedings in Informatics*. Vol. 40. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Bimpikis, Kostas, Ozan Candogan, and Daniela Saban (2016). “Spatial pricing in ride-sharing networks”. In:
- Bimpikis, Kostas, Shayan Ehsani, and Rahmi Ilkiliç (2014). “Cournot competition in networked markets”. In: *EC*, p. 733.
- Bogucka, H. et al. (Nov. 2012). “Secondary spectrum trading in TV white spaces”. In: *IEEE Communications Magazine* 50.11, pp. 121–129.
- Borodin, Allan, Nathan Linial, and Michael E Saks (1992). “An optimal on-line algorithm for metrical task system”. In: *Journal of the ACM (JACM)* 39.4, pp. 745–763.
- Bose, Subhonmesh et al. (2014). “The role of a market maker in networked cournot competition”. In: *53rd IEEE Conference on Decision and Control*. IEEE, pp. 4479–4484.
- Boudreau, Kevin (2010). “Open platform strategies and innovation: Granting access vs. devolving control”. In: *Management Science* 56.10, pp. 1849–1872.
- Branco, Fernando, Monic Sun, and J Miguel Villas-Boas (2015). “Too much information? Information provision and search costs”. In: *Marketing Science* 35.4, pp. 605–618.
- Buchbinder, Niv and Joseph Naor (2005). “Online primal-dual algorithms for covering and packing problems”. In: *European Symposium on Algorithms*. Springer, pp. 689–701.
- Buchbinder, Niv, Joseph Seffi Naor, et al. (2009). “The design of competitive on-line algorithms via a primal–dual approach”. In: *Foundations and Trends® in Theoretical Computer Science* 3.2–3, pp. 93–263.

- Bunn, D. W. and F. S. Oliveira (2003). “Evaluating individual market power in electricity markets via agent-based simulation”. In: *Annals of Operations Research* 121.1-4, pp. 57–77.
- Cai, Desmond, Subhonmesh Bose, and Adam Wierman (2017). “On the role of a market maker in networked cournot competition”. In: *arXiv preprint arXiv:1701.08896*.
- Chawla, Shuchi, Jason D Hartline, and Robert Kleinberg (2007). “Algorithmic pricing via virtual valuations”. In: *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, pp. 243–251.
- Chawla, Shuchi, Jason D Hartline, David L Malec, et al. (2010). “Multi-parameter mechanism design and sequential posted pricing”. In: *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, pp. 311–320.
- Chawla, Shuchi and Tim Roughgarden (2008). “Bertrand competition in networks”. In: *Algorithmic Game Theory*. Springer, pp. 70–82.
- Chen, Julie Yujie (2017). “Thrown under the bus and outrunning it! The logic of Didi and taxi drivers’ labour and activism in the on-demand economy”. In: *New Media & Society*, p. 1461444817729149.
- Chen, Le, Alan Mislove, and Christo Wilson (2016). “An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace”. In: *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pp. 1339–1349.
- Chen, Le and Christo Wilson (2017). “Observing algorithmic marketplaces in-the-wild”. In: *ACM SIGecom Exchanges* 15.2, pp. 34–39.
- Chen, Nianguan et al. (2015). “Online convex optimization using predictions”. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 43. 1. ACM, pp. 191–204.
- Chin, Francis YL et al. (2015). “Competitive algorithms for unbounded one-way trading”. In: *Theoretical Computer Science* 607, pp. 35–48.
- Chircu, AM and Robert J Kauffman (2001). “Digital intermediation in electronic commerce—the eBay model”. In: *E-commerce and V-business*, pp. 43–66.
- Chow, YS et al. (1964). “Optimal selection based on relative rank (the “secretary problem”)”. In: *Israel Journal of mathematics* 2.2, pp. 81–90.
- Conejo, Antonio J, Juan M Morales, and Luis Baringo (2010). “Real-time demand response model”. In: *IEEE Transactions on Smart Grid* 1.3, pp. 236–242.
- Cournot, Antoine-Augustin (1838). *Research on the Mathematical Principles of the Thrift of Wealth by Augustin Cournot*. L. Hachette.
- Damaschke, Peter, Phuong Hoai Ha, and Philippas Tsigas (2009). “Online search with time-varying price bounds”. In: *Algorithmica* 55.4, pp. 619–642.

- Diehl, Kristin, Laura J Kornish, and John G Lynch Jr (2003). “Smart agents: When lower search costs for quality information increase price sensitivity”. In: *Journal of Consumer Research* 30.1, pp. 56–71.
- Dinerstein, Michael et al. (2018). “Consumer price search and platform design in internet commerce”. In: *American Economic Review* 108.7, pp. 1820–59.
- Einav, Liran et al. (2015). “Assessing sale strategies in online markets using matched listings”. In: *American Economic Journal: Microeconomics* 7.2, pp. 215–247.
- Elliott, Matthew (2015). “Inefficiencies in Networked Markets”. In: *American Economic Journal: Microeconomics* 7.4, pp. 43–82. DOI: 10.1257/mic.20130098.
- Evans, David S and Richard Schmalensee (2016). *Matchmakers: the new economics of multisided platforms*. Harvard Business Review Press.
- Evans, Peter C and Annabelle Gawer (2016). “The rise of the platform enterprise: a global survey”. In:
- Fang, Zhixuan, Longbo Huang, and Adam Wierman (2018). “Loyalty Programs in the Sharing Economy: Optimality and Competition”. In: *arXiv preprint arXiv:1805.03581*.
- Feldman, Moran and Rico Zenklusen (2018). “The submodular secretary problem goes linear”. In: *SIAM Journal on Computing* 47.2, pp. 330–366.
- Fiat, Amos (1998). “Online Algorithms: The State of the Art (Lecture Notes in Computer Science)”. In:
- Fiat, Amos and Manor Mendel (2003). “Better algorithms for unfair metrical task systems and applications”. In: *SIAM Journal on Computing* 32.6, pp. 1403–1422.
- Fiat, Amos, Yuval Rabani, and Yiftach Ravid (1990). “Competitive k-server algorithms”. In:
- Forbes (2018). *GLOBAL 2000: THE WORLD’S LARGEST PUBLIC COMPANIES*. URL: %5Cur1%7Bhttps://www.forbes.com/global2000/#50eb23d5335d%7D.
- Fujiwara, Hiroshi, Kazuo Iwama, and Yoshiyuki Sekiguchi (2011). “Average-case competitive analyses for one-way trading”. In: *Journal of Combinatorial Optimization* 21.1, pp. 83–107.
- Gamp, Tobias (2016). *Guided search*.
- Gellings, Clark W (1985). “The concept of demand-side management for electric utilities”. In: *Proceedings of the IEEE* 73.10, pp. 1468–1470.
- Golombek, Rolf, Eystein Gjelsvik, and Knut Einar Rosendahl (1995). “Effects of liberalizing the natural gas markets in Western Europe”. In: *The Energy Journal*, pp. 85–111.
- Gross, Ben and Alessandro Acquisti (2003). “Balances of power on ebay: Peers or unequals”. In: *Workshop on Economics of Peer-to-peer systems*. Citeseer.

- Guo, Linqi, John Pang, and Anwar Walid (2016). “Dynamic service function chaining in sdn-enabled networks with middleboxes”. In: *2016 IEEE 24th International Conference on Network Protocols (ICNP)*. IEEE, pp. 1–10.
- (2018). “Joint Placement and Routing of Network Function Chains in Data Centers”. In: *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, pp. 612–620.
- Guruswami, Venkatesan et al. (2005). “On profit-maximizing envy-free pricing”. In: *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, pp. 1164–1173.
- Guzmán, Carlos Lever (2011). *Price competition on network*. Tech. rep.
- Hagiu, Andrei and Julian Wright (2015). “Multi-sided platforms”. In: *International Journal of Industrial Organization* 43, pp. 162–174.
- Hajiaghayi, Mohammad Taghi, Robert Kleinberg, and Tuomas Sandholm (2007). “Automated online mechanism design and prophet inequalities”. In: *AAAI*. Vol. 7, pp. 58–65.
- Hall, Jonathan, Cory Kendrick, and Chris Nosko (2015). “The effects of Uber’s surge pricing: A case study”. In: *The University of Chicago Booth School of Business*.
- Hazan, Elad, Amit Agarwal, and Satyen Kale (2007). “Logarithmic regret algorithms for online convex optimization”. In: *Machine Learning* 69.2-3, pp. 169–192.
- Heylighen, Francis (2006). “Why is Open Access Development so Successful? Stigmergic organization and the economics of information”. In: *arXiv preprint cs/0612071*.
- Hobbs, Benjamin F (2001). “Linear complementarity models of Nash-Cournot competition in bilateral and POOLCO power markets”. In: *Power Systems, IEEE Trans. on* 16.2, pp. 194–202.
- Huang, Longbo, Jean Walrand, and Kannan Ramchandran (2012). “Optimal demand response with energy storage management”. In: *2012 IEEE third international conference on smart grid communications (SmartGridComm)*. IEEE, pp. 61–66.
- Hui, Xiang et al. (2016). “Reputation and regulations: evidence from ebay”. In: *Management Science* 62.12, pp. 3604–3616.
- Ilkilic, Rahmi (2009). “Cournot competition on a network of markets and firms”. In:
- Johari, R. and J. N. Tsitsiklis (2005). “Efficiency loss in Cournot games”. In: *Technical Report*. URL: <http://web.mit.edu/jnt/www/Papers/R-05-cournot-tr.pdf>.
- Johari, Ramesh and John N Tsitsiklis (2004). “Efficiency loss in a network resource allocation game”. In: *Mathematics of Operations Research* 29.3, pp. 407–435.

- Kakade, Sham M et al. (2004). “Competitive algorithms for VWAP and limit order trading”. In: *Proceedings of ACM EC*, pp. 189–198.
- Kanoria, Yash and Daniela Saban (2017). “Facilitating the Search for Partners on Matching Platforms: Restricting Agents’ Actions”. In:
- Karlin, Anna R et al. (1988). “Competitive snoopy caching”. In: *Algorithmica* 3.1-4, pp. 79–119.
- Koutsoupias, Elias and Christos H. Papadimitriou (1999). “Worst-case Equilibria”. In: *Symposium on Theoretical Aspects of Computer Science*, pp. 404–413.
- L.Lu et al. (2013). “Online Energy Generation Scheduling for Microgrids with Intermittent Energy Sources and Co-Generation”. In: *Proceedings of ACM Sigmetrics*.
- Lahiri, Sajal and Yoshiyasu Ono (1988). “Helping minor firms reduces welfare”. In: *The economic journal* 98.393, pp. 1199–1202.
- Li, Jiayin et al. (2012). “Online optimization for scheduling preemptable tasks on IaaS cloud systems”. In: *Journal of Parallel and Distributed Computing* 72.5, pp. 666–677.
- Li, Yingying, Guannan Qu, and Na Li (2018). “Using Predictions in Online Optimization with Switching Costs: A Fast Algorithm and A Fundamental Limit”. In: *2018 Annual American Control Conference (ACC)*. IEEE, pp. 3008–3013.
- Liang, Sai et al. (2017). “Be a “Superhost”: The importance of badge systems for peer-to-peer rental accommodations”. In: *Tourism management* 60, pp. 454–465.
- Lin, Minghong et al. (2012). “Online optimization with switching cost”. In: *ACM SIGMETRICS Performance Evaluation Review* 40.3, pp. 98–100.
- Lin, Weixuan et al. (2017). “Access Control in Online Platforms: Competition and Efficiency Loss”. In: *Proceedings of the IEEE Conference for Decision and Control*.
- Lorenz, Julian, Konstantinos Panagiotou, and Angelika Steger (2009). “Optimal algorithms for k-search with application in option pricing”. In: *Algorithmica* 55.2, pp. 311–328.
- Lu, Alice, Peter I Frazier, and Oren Kislev (2018). “Surge Pricing Moves Uber’s Driver-Partners”. In: *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, pp. 3–3.
- Lu, Tan, Minghua Chen, and Lachlan LH Andrew (2013). “Simple and effective dynamic provisioning for power-proportional data centers”. In: *IEEE Transactions on Parallel and Distributed Systems* 24.6, pp. 1161–1171.
- Luca, Michael (2017). “Designing online marketplaces: Trust and reputation mechanisms”. In: *Innovation Policy and the Economy* 17.1, pp. 77–93.
- Luo, Zhi-Quan, Jong-Shi Pang, and Daniel Ralph (1996). *Mathematical programs with equilibrium constraints*. Cambridge University Press.

- Ma, Hongyao, Fei Fang, and David C Parkes (2018). “Spatio-temporal pricing for ridesharing platforms”. In: *arXiv preprint arXiv:1801.04015*.
- Mankiw, N Gregory and Michael D Whinston (1986). “Free entry and social inefficiency”. In: *The RAND Journal of Economics*, pp. 48–58.
- Motalleb, Mahdi et al. (2017). “A game-theoretic demand response market with networked competition model”. In: *Power Symposium (NAPS), 2017 North American*. IEEE, pp. 1–6.
- Nava, Francesco (2015). “Efficiency in decentralized oligopolistic markets”. In: *Journal of Economic Theory* 157, pp. 315–348.
- Newcomer, Eric (Apr. 2017). *Uber, Lifting Financial Veil, Says Sales Growth Outpaces Losses*. URL: <https://www.bloomberg.com/news/articles/2017-04-14/embattled-uber-reports-strong-sales-growth-as-losses-continue>.
- Nguyen, Thanh (2015). “Coalitional bargaining in networks”. In: *Operations Research* 63.3, pp. 501–511.
- Nishida, Mitsukuni and Marc Remer (2018). “Lowering consumer search costs can lead to higher prices”. In: *Economics Letters* 162, pp. 1–4.
- Nosko, Chris and Steven Tadelis (2015). *The limits of reputation in platform markets: An empirical analysis and field experiment*. Tech. rep. National Bureau of Economic Research.
- Oren, Shmuel S (1997). “Economic inefficiency of passive transmission rights in congested electricity systems with competitive generation”. In: *The Energy Journal*, pp. 63–83.
- Palensky, Peter and Dietmar Dietrich (2011). “Demand side management: Demand response, intelligent energy systems, and smart loads”. In: *IEEE transactions on industrial informatics* 7.3, pp. 381–388.
- Pang, John ZF et al. (2017). “The Efficiency of Open Access in Platforms for Networked Cournot Markets”. In: *Proceedings of IEEE INFOCOM*.
- Pang, John ZF, Linqi Guo, and Steven H Low (2017). “Optimal load control for frequency regulation under limited control coverage”. In: *IREP Symposium*, pp. 1–7.
- Parker, Geoffrey and Marshall Van Alstyne (2017). “Innovation, openness, and platform control”. In: *Management Science*.
- Pina, André, Carlos Silva, and Paulo Ferrão (2012). “The impact of demand side management strategies in the penetration of renewable electricity”. In: *Energy* 41.1, pp. 128–137.
- Purchala, K. et al. (2005). “Usefulness of DC power flow for active power flow analysis”. In: *Proc. of IEEE PES General Meeting*. IEEE, pp. 2457–2462.

- Qian, L. et al. (Nov. 2011). “Spectrum Trading in Cognitive Radio Networks: An Agent-Based Model under Demand Uncertainty”. In: *IEEE Transactions on Communications* 59.11, pp. 3192–3203.
- Qian, Li Ping et al. (2013). “Demand response management via real-time electricity price control in smart grids”. In: *IEEE Journal on Selected areas in Communications* 31.7, pp. 1268–1280.
- Rochet, Jean-Charles and Jean Tirole (2003). “Platform competition in two-sided markets”. In: *Journal of the european economic association* 1.4, pp. 990–1029.
- Rosen, J Ben (1965). “Existence and uniqueness of equilibrium points for concave n-person games”. In: *Econometrica: Journal of the Econometric Society*, pp. 520–534.
- Rosenblat, Alex and Luke Stark (2016). “Algorithmic labor and information asymmetries: A case study of Uber’s drivers”. In:
- Rubinstein, Aviad (2016). “Beyond matroids: Secretary problem and prophet inequality with general constraints”. In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, pp. 324–332.
- Ruhi, Navid Azizan et al. (2018). “Opportunities for price manipulation by aggregators in electricity markets”. In: *IEEE Transactions on Smart Grid* 9.6, pp. 5687–5698.
- Rysman, Marc (2009). “The economics of two-sided markets”. In: *Journal of Economic Perspectives* 23.3, pp. 125–43.
- Salant, Stephen W (1982). “Imperfect competition in the international energy market: a computerized Nash-Cournot model”. In: *Operations research* 30.2, pp. 252–280.
- Scheiber, Noam (2017). “How Uber uses psychological tricks to push its drivers’ buttons”. In: *The New York Times* 2.
- Schor, Juliet (2016). “DEBATING THE SHARING ECONOMY.” In: *Journal of Self-Governance & Management Economics* 4.3.
- Shafiee, S. et al. (Apr. 2018). “Economic Assessment of Energy Storage Systems in Alberta’s Energy and Operating Reserve Markets”. In: *2018 IEEE/PES Transmission and Distribution Conference and Exposition*, pp. 1–5.
- Siano, Pierluigi (2014). “Demand response and smart grids—A survey”. In: *Renewable and sustainable energy reviews* 30, pp. 461–478.
- Stoft, S. (2002). *Power System Economics: Designing Market for Power*. Piscataway, NJ: IEEE Press.
- Suzumura, Kotaro and Kazuharu Kiyono (1987). “Entry barriers and economic welfare”. In: *The Review of Economic Studies* 54.1, pp. 157–167.

- Tadelis, Steven (2016). “Reputation and feedback systems in online platform markets”. In: *Annual Review of Economics* 8, pp. 321–340.
- Tsitsiklis, J. N. and Y. Xu (2012). “Efficiency Loss in a Cournot Oligopoly with Convex Market Demand”. In: *Game Theory for Networks*. Springer, pp. 63–76.
- Tsitsiklis, John N and Yunjian Xu (2014). “Efficiency loss in a cournot oligopoly with convex market demand”. In: *Journal of Mathematical Economics* 53, pp. 46–58.
- Walawalkar, Rahul et al. (2010). “Evolution and current status of demand response (DR) in electricity markets: Insights from PJM and NYISO”. In: *Energy* 35.4, pp. 1553–1560.
- Wang, Zhaojian et al. (2018). “Distributed optimal frequency control considering a nonlinear network-preserving model”. In: *IEEE Transactions on Power Systems*.
- Weare, Christopher (2003). *The California electricity crisis: causes and policy options*. Public Policy Instit. of CA.
- Weyl, E Glen (2010). “A price theory of multi-sided platforms”. In: *The American Economic Review* 100.4, pp. 1642–1672.
- Xu, Yunjian et al. (2017). “On the efficiency of networked Stackelberg competition”. In: *Information Sciences and Systems (CISS), 2017 51st Annual Conference on*. IEEE, pp. 1–6.
- El-Yaniv, Ran et al. (2001). “Optimal search and one-way trading online algorithms”. In: *Algorithmica* 30.1, pp. 101–139.
- You, Pengcheng, Peng Cheng, et al. (2018). “Efficient Online Station Assignment for EV Battery Swapping”. In:
- You, Pengcheng, John ZF Pang, Minghua Chen, et al. (2017). “Battery swapping assignment for electric vehicles: A bipartite matching approach”. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, pp. 1421–1426.
- You, Pengcheng, John ZF Pang, and Enoch Yeung (2018a). “Deep Koopman Controller Synthesis for Cyber-Resilient Market-Based Frequency Regulation”. In: *IFAC-PapersOnLine* 51.28, pp. 720–725.
- (2018b). “Stabilization of Power Networks via Market Dynamics”. In: *Proceedings of the Ninth International Conference on Future Energy Systems*. ACM, pp. 139–145.
- Zhang, Wenming et al. (2012). “Optimal algorithms for online time series search and one-way trading with interrelated prices”. In: *Journal of combinatorial optimization* 23.2, pp. 159–166.
- Zhang, Ying et al. (2018). “Peak-aware online economic dispatching for microgrids”. In: *IEEE Transactions on Smart Grid* 9.1, pp. 323–335.

Appendix A

PROOFS FOR CHAPTER III

Proof of Lemma 3.8

Let $q^{\text{NE}}(F \times M) \in Q(F \times M)$ be the unique Nash equilibrium of the game $(F, Q(F \times M), \pi)$ associated with an arbitrary cost function profile c . Throughout the proof, we always consider a networked Cournot game associated with the edge set $F \times M$. Thus, for notational simplicity we use q^{NE} instead of $q^{\text{NE}}(F \times M)$ for the remainder of the proof. For each $i \in F$, we define the non-negative scalar λ_i according to

$$\lambda_i := \partial^+ C_i \left(\sum_{j=1}^m q_{ij}^{\text{NE}} \right). \quad (\text{A.1})$$

Here, λ_i is the marginal cost of firm i at the unique Nash equilibrium of the game $(F, Q(F \times M), \pi)$. We define a (piecewise) linear cost function profile $\bar{c} = (\bar{c}_1, \dots, \bar{c}_n)$ according to

$$\bar{c}_i(s_i) := \max\{\lambda_i s_i, 0\}, \quad i = 1, \dots, n.$$

Clearly, \bar{c}_i is convex, and differentiable on $(0, \infty)$ for each $i \in \{1, \dots, n\}$. Additionally, the stationarity conditions of firms' profit maximization problem under the cost function profiles C and \bar{C} are identical at q^{NE} . The combination of these two facts shows that q^{NE} is the unique Nash equilibrium of the game $(F, Q(F \times M), \bar{\pi})$ associated with the cost function profile \bar{C} . Our objective is to show that $\rho(F \times M, C) \leq \rho(F \times M, \bar{C})$, i.e.,

$$\frac{\text{SW}(q^{\text{NE}}, C)}{\text{SW}^*(F \times M, C)} \geq \frac{\text{SW}(q^{\text{NE}}, \bar{C})}{\text{SW}^*(F \times M, \bar{C})}.$$

In showing this, we first define the scalar μ_i for each firm $i \in F$ according to

$$\mu_i := \partial^+ C_i \left(\sum_{j=1}^m q_{ij}^{\text{NE}} \right) \cdot \left(\sum_{j=1}^m q_{ij}^{\text{NE}} \right) - C_i \left(\sum_{j=1}^m q_{ij}^{\text{NE}} \right).$$

For each firm i , μ_i equals the absolute difference in his production cost at Nash equilibrium associated with the cost function profiles C and \bar{C} .

With the definition of μ_i in hand, we define an ‘‘intermediate’’ cost function profile $\tilde{c} = (\tilde{c}_1, \dots, \tilde{c}_n)$ according to:

$$\tilde{c}_i(s_i) := \max\{\lambda_i s_i - \mu_i, 0\}, \quad i = 1, \dots, n.$$

The cost function profile \tilde{C} makes the connection between the cost function profiles C and \bar{C} . On one hand, for each firm $i \in F$, the cost function \tilde{C}_i can be regarded as a linearization of the original cost function C_i around the Nash equilibrium of the game $(F, Q(F \times M), \pi)$. On the other hand, it can be interpreted as a translation of the cost function \bar{c}_i downwards along the y-axis of length μ_i , while keeping the resulting cost function non-negative for all real numbers. Note that the stationarity conditions of firms' profit maximization problem under the cost function profiles C and \tilde{C} are identical at q^{NE} . It follows that q^{NE} is the unique Nash equilibrium of the game $(F, Q(F \times M), \tilde{\pi})$ associated with the cost function profile \tilde{C} .

Since firms' production costs at q^{NE} are equal for the cost function profiles C and \tilde{C} , we have that $\text{SW}(q^{\text{NE}}, C) = \text{SW}(q^{\text{NE}}, \tilde{C})$. Moreover, since $C_i(s_i) \geq \tilde{C}_i(s_i)$ for all $s_i \in \mathbb{R}_+$, we have that $\text{SW}^*(F \times M, C) \leq \text{SW}^*(F \times M, \tilde{C})$. It follows that

$$\frac{\text{SW}(q^{\text{NE}}, C)}{\text{SW}^*(F \times M, C)} \geq \frac{\text{SW}(q^{\text{NE}}, \tilde{C})}{\text{SW}^*(F \times M, \tilde{C})}. \quad (\text{A.2})$$

Additionally, the previous translation from \bar{C} to \tilde{C} implies that $\text{SW}(q^{\text{NE}}, \tilde{C})$ and $\text{SW}(q^{\text{NE}}, \bar{C})$ are related according to:

$$\text{SW}(q^{\text{NE}}, \bar{C}) = \text{SW}(q^{\text{NE}}, \tilde{C}) - \sum_{i=1}^n \mu_i \geq 0. \quad (\text{A.3})$$

We claim that the following inequality holds for the efficient social welfare $\text{SW}^*(F \times M, \bar{C})$ associated with the cost function profile \bar{C} :

$$\text{SW}^*(F \times M, \bar{C}) \geq \text{SW}^*(F \times M, \tilde{C}) - \sum_{i=1}^n \mu_i.$$

To see this, let q^* be an efficient supply profile under the cost function profile \tilde{C} . For each $i \in F$, we have

$$\bar{C}_i \left(\sum_{j=1}^m q_{ij}^* \right) \leq \tilde{C}_i \left(\sum_{j=1}^m q_{ij}^* \right) + \mu_i.$$

This inequality implies that

$$\begin{aligned} \text{SW}^*(F \times M, \bar{C}) &\geq \text{SW}(q^*, \bar{C}) \geq \text{SW}(q^*, \tilde{C}) - \sum_{i=1}^n \mu_i \\ &= \text{SW}^*(F \times M, \tilde{C}) - \sum_{i=1}^n \mu_i, \end{aligned}$$

The above inequality, in combination with inequalities Equations A.2 and A.3, shows that

$$\begin{aligned} \frac{SW(q^{\text{NE}}, C)}{SW^*(F \times M, C)} &\geq \frac{SW(q^{\text{NE}}, \tilde{C})}{SW^*(F \times M, \tilde{C})} \\ &\geq \frac{SW(q^{\text{NE}}, \tilde{C}) - \sum_{i=1}^n \mu_i}{SW^*(F \times M, \tilde{C}) - \sum_{i=1}^n \mu_i} \geq \frac{SW(q^{\text{NE}}, \bar{C})}{SW^*(F \times M, \bar{C})}, \end{aligned}$$

as needed to be shown.

Proof of Proposition 3.9

It follows from Lemma 3.8 that the worst symmetric cost function profile that maximizes $\rho(F \times M, C)$ consists of n identical cost functions that are linear on $(0, \infty)$. Thus, to upper bound the PoA of the networked Cournot game, it suffices to consider a cost function profile \bar{C} that satisfies

$$\bar{C}_i(x) = cx \quad \text{for } x \geq 0, \quad \text{for all } i \in F, \quad (\text{A.4})$$

for a finite positive constant $c > 0$ that is independent of i .

Given the assumption on the (piecewise) linearity of cost, it is straightforward to show that the unique Nash equilibrium and an efficient supply profile of the corresponding networked Cournot game are given by

$$q_{ij}^{\text{NE}} = \frac{(\alpha_j - c)^+}{\beta_j(n+1)}, \quad \text{and} \quad q_{ij}^* = \frac{(\alpha_j - c)^+}{\beta_j n},$$

respectively, for each $i \in F, j \in M$. It follows that the social welfare at the unique Nash equilibrium q^{NE} of this Cournot game is given by

$$SW^*(q^{\text{NE}}, \bar{C}) = \sum_{j=1}^m \frac{((\alpha_j - c)^+)^2}{2\beta_j} \left(1 - \frac{1}{(n+1)^2}\right).$$

And the efficient social welfare is given by

$$SW^*(F \times M, \bar{C}) = \sum_{j=1}^m \frac{((\alpha_j - c)^+)^2}{2\beta_j}.$$

Hence, the price of anarchy associated with the cost function profile \bar{C} is given by

$$\rho(F \times M, \bar{C}) = \begin{cases} 1 + \frac{1}{(n+1)^2 - 1} & \text{if } \max_{j \in M} \alpha_j > c, \\ 1 & \text{otherwise} \end{cases}.$$

Choosing $c < \max_{j \in M} \alpha_j$ gives the worst-case cost function profile that maximizes the price of anarchy over symmetric cost function profiles. This completes the proof.

Proof of Theorem 3.5

It follows from Lemma 3.8 that the worst cost function profile that maximizes $\rho(F \times M, C)$ consists of cost functions that are linear on $(0, \infty)$. We thus assume that each firm's cost function satisfies $\bar{C}_i(s_i) = (c_i s_i)^+$, for $i = 1, \dots, n$, where we assume without loss of generality that $c_1 \leq \dots \leq c_n$.

Note that, for the case in which $n = 1$, a direct application of Proposition 3.9 provides a tight price of anarchy bound of $\rho(F \times M, C) \leq 4/3$. For the remainder of the proof, we restrict ourselves to $n \geq 2$.

We first consider the simple setting where the number of markets $m = 1$. That is, the set of markets $M = \{1\}$. Without loss of generality, we assume that $c_1 < \alpha_1$. Since cost functions are linear, there exists an efficient supply profile q^* that assigns all production to firm 1, i.e., $q_{i1}^* = 0$ for $i = 2, \dots, n$. The supply from firm 1 and the corresponding efficient social welfare are given by:

$$q_{11}^* = \frac{\alpha_1 - c_1}{\beta_1}, \quad \text{and} \quad SW^*(F \times \{1\}, \bar{C}) = \frac{(\alpha_1 - c_1)^2}{2\beta_1}.$$

Fixing α_1, β_1, c_1 , we can optimize over c_2, \dots, c_n in order to minimize the social welfare at the unique Nash equilibrium of the Cournot game. Using similar arguments as in the proof of (R. Johari and J. N. Tsitsiklis, 2005, Theorem 12), one can formulate this problem as a symmetric convex program over the production quantities at Nash equilibrium of the remaining firms. The optimal value of c_2, \dots, c_n is given by

$$c_i^* = \alpha_1 - \frac{2n+3}{3n+5}(\alpha_1 - c_1), \quad \text{for } i = 2, \dots, n.$$

And the production quantity of each producer is given by

$$q_{i1} = \begin{cases} \frac{\alpha_1 - c_1}{\beta_1} \cdot \frac{n+3}{3n+5} & \text{if } i = 1. \\ \frac{\alpha_1 - c_1}{\beta_1} \cdot \frac{1}{3n+5}, & \text{if } i = 2, \dots, n. \end{cases}$$

Define the cost function profile $\bar{C}^* = (\bar{C}_1^*, \dots, \bar{C}_n^*)$ according to $\bar{C}_1^*(s_1) = (c_1 s_1)^+$ and $\bar{C}_i^*(s_i) = (c_i^* s_i)^+$ for $i = 2, \dots, n$. Thus, for the fixed parameters α_1, β_1, c_1 , the minimum social welfare at Nash equilibrium is given by

$$SW(q^{NE}(F \times \{1\}, \bar{C}^*)) = \frac{(n+2)(\alpha_1 - c_1)^2}{(3n+5)\beta_1}.$$

It follows that for any linear cost function profile \bar{C} , we have

$$\rho(F \times \{1\}, \bar{C}) \leq \rho(F \times \{1\}, \bar{C}^*) = \frac{3}{2} \left(1 - \frac{1}{3n+6} \right).$$

We now consider the slightly more complicated setting in which $m > 1$. The efficient social welfare associated with a linear cost function profile \bar{C} is given by

$$SW^*(F \times M, \bar{C}) = \sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{2\beta_j}.$$

Given the linearity of firms' cost functions over $(0, \infty)$, the networked Cournot game decouples across markets. Thus, the social welfare at Nash equilibrium satisfies

$$\begin{aligned} SW(q^{NE}(F \times M), \bar{C}) &= \sum_{j=1}^m SW(q^{NE}(F \times \{j\}), \bar{C}) \\ &\geq \sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{3\beta_j \left(1 - \frac{1}{3n+6}\right)} = \frac{SW^*(F \times M, \bar{C})}{\frac{3}{2} \left(1 - \frac{1}{3n+6}\right)}. \end{aligned}$$

It follows that the price of anarchy $\rho(F \times M, \bar{C})$ satisfies

$$\rho(F \times M, \bar{C}) \leq \frac{3}{2} \left(1 - \frac{1}{3n+6}\right).$$

Additionally, it is straightforward to check that this price of anarchy bound is achieved by the cost function profile \bar{C}^* if $\alpha_1 = \alpha_2 = \dots = \alpha_m$. This completes the proof.

Proof of Proposition 3.10

We only provide a proof of the price of anarchy bound, since the proof on its tightness is straightforward. The key intuition of the proof is that, under a linear cost function profile, the networked Cournot game decouples across markets. The proof proceeds in two parts.

Part 1: We provide a price of anarchy bound for the case in which the number of markets $m = 1$. Namely, for any cost function profile $C \in \mathcal{L}^n(c_{\min}, c_{\max})$, we have

$$\rho(F \times \{j\}, C) \leq \frac{1}{\frac{2n+4}{3n+5} + \delta(\gamma_j, n)}, \quad j = 1, \dots, m.$$

We omit the proof of this price of anarchy bound, as it follows from similar arguments as in the proof of Theorem 3.5.

Part 2: We provide a price of anarchy bound for the case in which the number of market $m > 1$. Let the cost function of each firm $i \in F$ be given by

$$C_i(s_i) = (c_i s_i)^+, \quad \text{where } c_i \in [c_{\min}, c_{\max}].$$

Without loss of generality, we assume that $c_1 \leq \dots \leq c_n$. Thus, the efficient social welfare associated with the edge set $F \times M$ and the cost function profile C is given by

$$SW^*(F \times M, C) = \sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{2\beta_j}.$$

Since firms' cost functions are linear on $(0, \infty)$, the networked Cournot game decouples across markets. It follows that the social welfare at the unique Nash equilibrium of the game $(F, Q(F \times M), \pi)$ satisfies

$$SW(q^{NE}(F \times M), C) = \sum_{j=1}^m SW(q^{NE}(F \times \{j\}), C). \quad (A.5)$$

Additionally, the term $SW(q^{NE}(F \times \{j\}), C)$ satisfies

$$SW(q^{NE}(F \times \{j\}), C) = \frac{SW^*(F \times \{j\}, C)}{\rho(F \times \{j\}, C)} \quad (A.6)$$

$$\geq \sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{2\beta_j} \left(\frac{2n+4}{3n+5} + \delta(\gamma_j, n) \right). \quad (A.7)$$

Here, equation A.6 follows from the definition of the price of anarchy, and inequality A.7 follows from the price of anarchy bound in Step 1. A combination of Eq. (A.5) and inequality (A.7) provides the following lower bound on the reciprocal of the price of anarchy $\rho(F \times M, C)$:

$$\frac{1}{\rho(F \times M, C)} \geq \frac{\sum_{j=1}^m \left(\frac{2n+4}{3n+5} + \delta(\gamma_j, n) \right) \frac{((\alpha_j - c_1)^+)^2}{\beta_j}}{\sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{\beta_j}}.$$

One can verify that the partial derivative of the right-hand-side (RHS) of the above inequality with respect to c_1 is non-negative for $c_1 \in [c_{\min}, c_{\max}]$. Hence, choosing $c_1 = c_{\min}$ minimizes the RHS of the above inequality. This completes the proof.

Lemma A.1. *There exists firm i , such that $s_i^N \leq s_i^*$.*

Proof. Proof: Suppose note. Then $\forall i, s_i^N > s_i^*$.

Case 1: $S_N \subseteq S_*$:

$$s_i^N > s_i^* \Rightarrow \exists j \in S_N \text{ s.t. } d_j^N > d_j^* \Rightarrow p_j^* > p_j^N,$$

but considering any firm i such that $q_{ij}^N > 0$,

$$\alpha_j > p_j^N = c'_i(s_i^N) + \beta_j q_{ij}^N > c'_i(s_i^N) > c'_i(s_i^*)$$

but by the optimality of $*$,

$$c'_i(s_i^*) \geq p_j^* \Rightarrow p_j^N > p_j^*,$$

a contradiction.

Case 2: \exists market j s.t. $j \in S_N \setminus S_*$.

For any i such that $q_{ij}^N > 0$, $c'_i(s_i^N) < \alpha_j \leq c'_i(s_i^*)$, a contradiction to $s_i^N > s_i^*$. \square

Proof. Proof of Lemma 3.7: Suppose on the other hand that there is a market j such that $d_j^N < d_j^*$. We show that this contradicts the preceding lemma. For any firm i such that $q_{ij}^N = 0$, we have $c'_i(s_i^N) \geq p_j^N \geq p_j^* = c'_i(s_i^*)$, and by monotonicity, $s_i^N > s_i^*$. On the other hand, if $q_{ij}^N > 0$,

$$c'_i(s_i^N) = \alpha_j - \beta_j d_j^N - \beta_j q_{ij}^N > \alpha_j - \beta_j \frac{d_j^*}{2} - \beta_j \frac{d_j^*}{2} = \alpha_j - \beta_j d_j^* = c'_i(s_i^*),$$

again implying that $s_i^N > s_i^*$ for all i , a contradiction to the preceding lemma. \square

Before providing a proof of Theorem 3.11 we first state and prove the key structural lemma we use in the proof.

Lemma A.2. *For a firm with costless production, a set of linear demand markets under controlled allocation is equivalent to a single market with a convex, piecewise linear demand curve. Conversely, any convex, decreasing, piecewise linear demand curve with finitely many linear segments can be realized by a set of linear demand markets under controlled access.*

Proof. Proof of Lemma A.2. The characterization of the socially optimal production, with s fixed highlights that the platform will reallocate this amount to d_1, \dots, d_m such that $\sum_j d_j = s$, and for each market j where $d_j > 0$, p_j is equivalent to a fixed price p across markets; for each market j where $d_j = 0$, it must be that $\alpha_j \leq p$. This shows that, as s increases, the allocation will enter the markets one by one in the order in which α_j decreases.

We say a market becomes *active* when supply starts entering it. For a set of active markets, before the next market becomes active, the marginal increase in supply will be allocated in proportion to $1/\beta_j$ (in order to keep the prices the same). This fully describes the behavior of the platform.

Without loss of generality, assume the markets are ordered such that $\alpha_1 \geq \dots \geq \alpha_m$. From the firm's point of view, the platform is equivalent to a single market with a piecewise linear demand curve: when the price is between α_1 and α_2 , the rate at which price drops when s increases is β_1 ; for $p \in [\alpha_2, \alpha_3]$, the rate is $1/(\frac{1}{\beta_1} + \frac{1}{\beta_2})$. In general, when the first k markets are active, prices drop at the rate of $(\sum_{j=1}^k \frac{1}{\beta_j})^{-1}$. We call this single demand curve the *aggregate demand curve*.

For a given production level, the area under the aggregate demand curve is equal to the welfare in the original markets. The aggregate demand curve fully characterizes the set of markets under controlled access. Note that whenever a new market joins, the rate at which price drops becomes slower, therefore the aggregate demand curve is always convex.

Conversely, one can show that any convex, decreasing, piecewise linear demand curve consisting of finitely many linear segments is equivalent to a set of linear demand markets under controlled access.

□

Proof. Proof of Theorem 3.11. By Theorem A.2, we can focus on constructing an aggregate demand curve. Fix a constant $\lambda \in (0, \frac{1}{2})$. The aggregate demand curve we construct for m markets, ($m \geq 2$), is

$$p(d) = \max_{0 \leq k < m} (\lambda^k - \lambda^{2k} d).$$

It is not hard to verify that this is the piecewise linear function that connects the following points: $(0, 1)$, $(\frac{1}{1+\lambda}, \frac{\lambda}{1+\lambda})$, $(\frac{1}{\lambda(1+\lambda)}, \frac{\lambda^2}{1+\lambda})$, \dots , $(\frac{1}{\lambda^{m-2}(1+\lambda)}, \frac{\lambda^{m-1}}{1+\lambda})$, $(\frac{1}{\lambda^{m-1}}, 0)$. Being the maximum of a family of decreasing linear functions, $p(d)$ is obviously a convex decreasing function. We first calculate the optimal social welfare, the area under $p(d)$. The trapezoid whose vertices are $(\frac{1}{\lambda^{k-1}(1+\lambda)}, 0)$, $(\frac{1}{\lambda^{k-1}(1+\lambda)}, \frac{\lambda^k}{1+\lambda})$, $(\frac{1}{\lambda^k(1+\lambda)}, 0)$, $(\frac{1}{\lambda^k(1+\lambda)}, \frac{\lambda^{k+1}}{1+\lambda})$ has area

$$\frac{1}{2} \left(\frac{1}{\lambda^k(1+\lambda)} - \frac{1}{\lambda^{k-1}(1+\lambda)} \right) \left(\frac{\lambda^k}{1+\lambda} + \frac{\lambda^{k+1}}{1+\lambda} \right) = \frac{1-\lambda}{1+\lambda}.$$

There are $m - 2$ such trapezoids under $p(d)$, and therefore the socially optimal welfare is $\Omega(m)$. On the other hand, the linear components of $p(d)$ are designed so that producing on any of the linear segment gives a maximal profit of $\frac{1}{4}$ (for the k -th segment, the profit maximizing production level is $\frac{1}{2\lambda^{k-1}}$). The firm is indifferent to best responding to any of the linear segments, and all the production levels $\frac{1}{2\lambda^k}$

for $k = 0, \dots, m - 1$ are Nash equilibria. If we push the starting point of $p(d)$ from $(0, 1)$ to $(0, 1 + \epsilon)$ for some $\epsilon > 0$, then producing $\frac{1+\epsilon}{2}$ (by responding only to the first market) will be the unique equilibrium, resulting in a social welfare of only $3(1 + \epsilon)^2/8$. Therefore the price of anarchy is $\Omega(m)$. \square

Proof. Proof of Theorem 3.14:

In the $k \leq 0$ regime, a similar example to the consumer surplus case would work. For example, if $\alpha_1 = \alpha_2$, with $\beta_1 = 0$ and $\beta_2 = \epsilon$, then the platform will assign all of the allocations to the second market. Consider a linear cost function $c = (\alpha_1 - \epsilon)x$, then the optimal is to fulfill all of market 1 which gives us infinitely large social welfare, whereas the platform forces the firm to go to market 2, which results in a constant social welfare.

On the other hand, when $k \in (0, 1]$, we can design a series of markets for a costless firm with a resulting price of anarchy that grows in the number of markets accessed, similar to the social welfare maximizing platform ($\lambda = 0.5$ case). Lastly, in the $k > 1$ regime, our price of anarchy bound decreases gracefully as the market maker objective approaches revenue maximizing, up till a point where it is worst to set up an open access platform, explaining the max operator.

We first show that regardless of the platform objective, the firms produce on aggregate a quantity equivalent to the quantity that maximizes revenue:

Lemma A.3. *Regardless of the parameter k for one costless firm, optimal aggregate production in the platform including markets M maximizes social welfare, i.e.*
 $q^* = \sum_{j \in M} \frac{\alpha_j}{2\beta_j}$.

Proof. Proof of Lemma A.3: Since market m is included, then the markets are allocated at least the following amount of quantity before the last market is accessed:

$$d_j = \frac{\alpha_j - \alpha_m}{k\beta_j}.$$

Further, we know that any additional quantity introduced into the platform will be assigned e_j to the respective markets respecting the following condition:

$$\beta_j e_j = \beta_m e_m \Rightarrow e_j = \frac{\frac{1}{\beta_j}}{\sum_k \frac{1}{\beta_k}} Q$$

Optimizing over the additional quantity e_j , the firm's profit can be written:

$$\sum_{j=1}^m \left(\frac{\alpha_j - \alpha_m}{k\beta_j} + e_j \right) \left(\alpha_j - \beta_j \left(\frac{\alpha_j - \alpha_m}{k\beta_j} + e_j \right) \right)$$

but with e_j defined above and letting $A = \sum_k \frac{1}{\beta_k}$, this can be rewritten as:

$$\sum_{j=1}^m \left(\frac{\alpha_j - \alpha_m}{k\beta_j} + \frac{Q}{A\beta_j} \right) \left(\alpha_j - \frac{\alpha_j - \alpha_m}{k} - \frac{Q}{A} \right)$$

Differentiating with respect to Q , we have:

$$\sum_{j=1}^m \left[\frac{1}{A\beta_j} \left(\alpha_j - \frac{\alpha_j - \alpha_m}{k} - \frac{Q}{A} \right) + \left(\frac{\alpha_j - \alpha_m}{k\beta_j} + \frac{Q}{A\beta_j} \right) \left(-\frac{1}{A} \right) \right] = 0$$

which gives us that:

$$Q = \sum_{j=1}^m \left(\frac{\alpha_j}{2\beta_j} - \frac{\alpha_j - \alpha_m}{k\beta_j} \right),$$

which implies that the optimal aggregate production is indeed $\sum_{j=1}^m \frac{\alpha_j}{2\beta_j}$.

□

We propose a set of parameters for each market as in the social welfare maximizing case and prove that the inefficiency in these markets are reflected as in Theorem 3.14.

Market parameters are designed in the following form, similar to that of the social welfare maximizing scenario:

$$\alpha_1 = 1, \beta_1 = 1, \text{ \& } \alpha_j = \frac{a^{j-1}\theta^{j-1}}{1+\theta}, \beta_j = \frac{\theta^{2j-2}}{1-\theta^2}, \forall j \geq 2,$$

then if we can show that the (single costless) firm wants to stay in the first market only region, then the price of anarchy grows as a function of the sum of a geometric progression with ratio a , and thereby can be written:

$$\rho = \frac{4}{3} \left(1 + \frac{1-\theta}{1+\theta} (a^2 + a^4 + \dots) \right) = \frac{4}{3} \left(1 + \frac{(1-\theta)a^2}{(1+\theta)(1-a^2)} \right) \rightarrow_{\theta \rightarrow 0} \frac{4}{3} \left(\frac{1}{1-a^2} \right)$$

The optimal revenue considering only the first market is $\frac{\alpha_1^2}{4\beta_1} = \frac{1}{4}$. After some algebra, the amount allocated to market j in the presence of m many markets of the form defined above, or $d_{j,m}$, can be expressed as:

$$\left(\left(\frac{1}{2} - \frac{1}{k} \right) \sum_{i=1}^m \frac{\alpha_i}{\beta_i} \right) \frac{\frac{1}{\beta_j}}{\sum_{i=1}^m \frac{1}{\beta_i}} + \frac{\alpha_j}{k\beta_j}$$

and the prices at the markets then can be computed via $\alpha_j - \beta_j d_{j,m}$:

$$\left(1 - \frac{1}{k}\right) \alpha_j - \frac{1}{\sum_{i=1}^m \frac{1}{\beta_i}} \left(\left(\frac{1}{2} - \frac{1}{k}\right) \sum_{i=1}^m \frac{\alpha_i}{\beta_i} \right)$$

and thereby, revenue from market j can be written:

$$\left(\frac{1}{2} - \frac{1}{k}\right) \frac{\sum_{i=1}^m \frac{\alpha_i}{\beta_i}}{\sum_{i=1}^m \frac{1}{\beta_i}} \left(1 - \frac{2}{k}\right) \frac{\alpha_j}{\beta_j} + \left(1 - \frac{1}{k}\right) \frac{\alpha_j^2}{k\beta_j} - \frac{1}{\beta_j} \left(\frac{1}{\sum_{i=1}^m \frac{1}{\beta_i}}\right)^2 \left(\frac{1}{2} - \frac{1}{k}\right)^2 \left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2$$

and revenue from all markets can be written:

$$\left(\frac{1}{2} - \frac{1}{k}\right)^2 \frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} + \frac{1}{k} \left(1 - \frac{1}{k}\right) \sum_{i=1}^m \frac{\alpha_i^2}{\beta_i}$$

and we wish to show that $\forall k$, we can find a and θ such that the above aggregate revenue is less than the revenue at only the first market. (Recall that when $k = 1$, the second term is gone and the first term is the revenue from the social welfare case, and taking $a = 1$, we retrieve the same righthand side as the social welfare case too, which equates to the revenue at the first market, i.e., $1/4$.)

Lemma A.4. *When $a = 1$ and $k = 1$, then for any $\lambda \in (0, 0.5)$,*

$$\left(\frac{1}{2} - \frac{1}{k}\right)^2 \frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} = \frac{1}{4} \iff \frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} = 1.$$

Proof. Proof of Lemma A.4:

This is proved under the social welfare maximizing platform, and restated here for convenience. \square

Another important and interesting point to note is that the two coefficients in the revenue term sum up to $1/4$, i.e.,

$$\left(\frac{1}{2} - \frac{1}{k}\right)^2 + \frac{1}{k} \left(1 - \frac{1}{k}\right) = \frac{1}{4}.$$

We use this identity to prove the following lemma:

Lemma A.5. *For $k \in (0, 1]$, take $a = 1$, and $\theta < 0.5$, then*

$$\left(\frac{1}{2} - \frac{1}{k}\right)^2 \frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} + \frac{1}{k} \left(1 - \frac{1}{k}\right) \sum_{i=1}^m \frac{\alpha_i^2}{\beta_i} < \frac{1}{4}$$

Proof. Proof of Lemma A.5: From above, it remains to show that $\sum_{i=1}^m \frac{\alpha_i^2}{\beta_i} \geq 1$, but the first term in the sum is already 1 while the rest are nonnegative so we are done. \square

We now move to prove the case whereby $k > 1$, through the following lemma:

Lemma A.6. *For $k > 1$, and taking θ small enough,*

$$\frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} \leq a^2.$$

This can be seen by considering the $\sum_{i=1}^m \frac{\alpha_i}{\beta_i}$ part and extracting a from each summand. With this result, we can show the following:

Lemma A.7. *For $k > 1$, $a \leq \sqrt{1 - 2\sqrt{\frac{1}{k}\left(1 - \frac{1}{k}\right)}}$ and taking λ small enough, then*

$$\left(\frac{1}{2} - \frac{1}{k}\right)^2 \frac{\left(\sum_{i=1}^m \frac{\alpha_i}{\beta_i}\right)^2}{\sum_{i=1}^m \frac{1}{\beta_i}} + \frac{1}{k} \left(1 - \frac{1}{k}\right) \sum_{i=1}^m \frac{\alpha_i^2}{\beta_i} < \frac{1}{4}.$$

The proof of Theorem 3.14 follows from the above lemmas, and the PoA bounds follow.

Proof of Theorem 3.17

Without loss of generality, we assume that $c_1 \leq \dots \leq c_n$. We only provide a proof for the case in which the number of market $m = 1$, as it is straightforward to generalize our proof to the case in which $m > 1$ under the assumption on the linearity of firms' cost functions.

We denote by $F_1(\mathcal{E})$ the set of firms that have access to market 1. It is defined according to

$$F_1(\mathcal{E}) = \{i \in F \mid (i, 1) \in \mathcal{E}\}.$$

We first introduce the concept of a contiguous set of firms, which plays a central role in the remainder of the proof. We have the following definition.

Definition A.8 (Contiguous Set). *The set $F_1(\mathcal{E})$ is contiguous if*

$$F_1(\mathcal{E}) = \{1, 2, \dots, |F_1(\mathcal{E})|\}, \quad \text{or} \quad F_1(\mathcal{E}) = \emptyset.$$

Here, $|F_1(\mathcal{E})|$ denotes the cardinality of the set $F_1(\mathcal{E})$. Qualitatively, the set $F_1(\mathcal{E}) \subseteq F$ is contiguous, if it consists of consecutive elements of the set F . Clearly, for the edge set \mathcal{E}^* generated by the greedy algorithm, the set $F_1(\mathcal{E}^*)$ is contiguous.

The rest of the proof consists of two parts. In Part 1, we show that if the set $F_1(\mathcal{E})$ is contiguous, then the social welfare at the Nash equilibrium associated with edge set \mathcal{E} is guaranteed to be no larger than that of the edge set \mathcal{E}^* . In Part 2, we consider the case in which the set $F_1(\mathcal{E})$ is not contiguous. We show that there exists an edge set $\tilde{\mathcal{E}}$ that yields a contiguous set $F_1(\tilde{\mathcal{E}})$, and has social welfare at Nash equilibrium that is no smaller than that of the edge set \mathcal{E} .

Part 1: In this part, we assume that the set $F_1(\mathcal{E})$ is contiguous, and show that $\text{SW}(q^{\text{NE}}(\mathcal{E}), C) \leq \text{SW}(q^{\text{NE}}(\mathcal{E}^*), C)$. We first define a sequence of edge sets according to

$$\mathcal{E}_k = \bigcup_{i=1}^k \{(i, 1)\}, \quad k = 0, \dots, n. \quad (\text{A.8})$$

In particular, we have that $\mathcal{E}_0 = \emptyset$. Let $k^* = |\mathcal{E}^*|$. To show that $\text{SW}(q^{\text{NE}}(\mathcal{E}), C) \leq \text{SW}(q^{\text{NE}}(\mathcal{E}^*), C)$ if $F_1(\mathcal{E})$ is contiguous, it suffices to show that the sequence $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C)$ is strictly increasing in k over $k = 0, \dots, k^*$, and monotonically non-increasing in k over $k = k^*, \dots, n$.

We assume without loss of generality that $c_i \leq \alpha_1$ for all i . If this is not the case, one can work with an alternative cost function profile $\tilde{C} = (\tilde{C}_1, \dots, \tilde{C}_n)$ that is defined according to

$$\tilde{C}_i(s_i) = (\min\{c_i, \alpha_1\} \cdot s_i)^+$$

for $i = 1, \dots, n$. Clearly, $\min\{c_i, \alpha_1\} \leq \alpha_1$ for all i . Additionally, it is straightforward to show that

$$\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) = \text{SW}(q^{\text{NE}}(\mathcal{E}_k), \tilde{C})$$

for all $k = 0, \dots, n$.

The proof of this claim on monotonicity relies on the following lemma. Its proof is deferred to Appendix A.

Lemma A.9. *Assume that each firm i 's cost function is of the form $C_i(s_i) = (c_i s_i)^+$, where $c_1 \leq \dots \leq c_n$. Let the number of markets $m = 1$, and the edge set \mathcal{E}_k be defined according to Eq. (A.8). For each $k = 1, \dots, n$, we have that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) >$*

$SW(q^{NE}(\mathcal{E}_{k-1}), C)$ if and only if

$$\alpha_1 - c_k > \frac{1}{k} \left(1 + \frac{1}{k - \frac{1}{2(k+1)}} \right) \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right). \quad (\text{A.9})$$

It follows from the description of the greedy algorithm that $SW(q^{NE}(\mathcal{E}_k), C)$ is strictly increasing in k for $0 \leq k \leq k^*$. Additionally, we have that

$$SW(q^{NE}(\mathcal{E}_{k^*+1}), C) \leq SW(q^{NE}(\mathcal{E}_{k^*}), C).$$

It follows from Lemma A.9 that for $k = k^* + 1$, the following inequality is satisfied:

$$\alpha_1 - c_k \leq \frac{1}{k} \left(1 + \frac{1}{k - \frac{1}{2(k+1)}} \right) \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right). \quad (\text{A.10})$$

To complete Part 1 of the proof, we have the following lemma. Its proof can be found in Appendix A.

Lemma A.10. *Let $k^* \in \{0, \dots, n\}$, and assume that $c_1 \leq \dots \leq c_n \leq \alpha_1$. If inequality (A.10) is satisfied for $k = k^* + 1$, then it is satisfied for $k = k^* + 1, \dots, n$.*

A combination of Lemma A.9 and A.10 reveals that

$$SW(q^{NE}(\mathcal{E}_k), C) \leq SW(q^{NE}(\mathcal{E}_{k-1}), C)$$

for $k = k^* + 1, \dots, n$. This completes Part 1 of the proof.

Part 2: In this part, we assume that the set $F_1(\mathcal{E})$ is not contiguous. We show that there exists an edge set $\tilde{\mathcal{E}} \subseteq F \times \{1\}$, such that the set $F_1(\tilde{\mathcal{E}})$ is contiguous, and $SW(q^{NE}(\mathcal{E}), C) \leq SW(q^{NE}(\tilde{\mathcal{E}}), C)$.

Our proof of the above claim is constructive. Given an edge set $\mathcal{E} \subseteq F \times \{1\}$, we construct a sequence of $(n^2 + 1)$ edge sets according to the following procedure

1. Set $k = 0$, and $\mathcal{E}_k = \mathcal{E}$.
2. If the set $F_1(\mathcal{E}_k)$ is contiguous, then set $\mathcal{E}_{k+1} = \mathcal{E}_k$, and go to Step 5. If not, go to Step 3.
3. Define the edge set $\tilde{\mathcal{E}}_k$ according to

$$\tilde{\mathcal{E}}_k = \mathcal{E}_k \setminus \{(\max F_1(\mathcal{E}_k), 1)\}. \quad (\text{A.11})$$

4. If $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) \leq \text{SW}(q^{\text{NE}}(\tilde{\mathcal{E}}_k), C)$, then set $\mathcal{E}_{k+1} = \tilde{\mathcal{E}}_k$. If not, set

$$\mathcal{E}_{k+1} = \tilde{\mathcal{E}}_k \cup \{(\min(F \setminus F_1(\mathcal{E}_k)), 1)\}. \quad (\text{A.12})$$

5. If $k < n^2$, update $k = k + 1$, and go to Step 2. If $k \geq n^2$, terminate the procedure.

We will show that both of the following claims are true

(i) The set $F_1(\mathcal{E}_{n^2})$ is contiguous.

(ii) For $k = 0, \dots, n^2 - 1$, we have

$$\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) \leq \text{SW}(q^{\text{NE}}(\mathcal{E}_{k+1}), C). \quad (\text{A.13})$$

We note that the second claim implies that $\text{SW}(q^{\text{NE}}(\mathcal{E}), C) \leq \text{SW}(q^{\text{NE}}(\mathcal{E}_{n^2}), C)$. In what follows, we show that Claim (i) and (ii) are true in Part 2.1 and 2.2 of the proof, respectively.

Part 2.1: Proof of Claim (i). We show that Claim (i) is true according to a ‘‘potential function’’ argument. Namely, we define a potential function on the edge set $\mathcal{E}_k \subseteq F \times \{1\}$ as follows

$$\Phi(\mathcal{E}_k) = |F_1(\mathcal{E}_k)|(\max F_1(\mathcal{E}_k) - |F_1(\mathcal{E}_k)|)$$

for $k = 0, \dots, n^2$. For all $\mathcal{E}_k \subseteq F \times \{1\}$, we have that $\Phi(\mathcal{E}_k) \geq 0$. It is straightforward to show $\Phi(\mathcal{E}_k) = 0$ if and only if the set $F_1(\mathcal{E}_k)$ is contiguous. It follows that if $\Phi(\mathcal{E}_k) > 0$, then \mathcal{E}_{k+1} is specified according to either $\mathcal{E}_{k+1} = \tilde{\mathcal{E}}_k$ or Eq. (A.12). If $\mathcal{E}_{k+1} = \tilde{\mathcal{E}}_k$, we have

$$\Phi(\mathcal{E}_{k+1}) = |F_1(\tilde{\mathcal{E}}_k)|(\max F_1(\tilde{\mathcal{E}}_k) - |F_1(\tilde{\mathcal{E}}_k)|) \quad (\text{A.14})$$

$$= |F_1(\tilde{\mathcal{E}}_k)|(\max F_1(\tilde{\mathcal{E}}_k) - (|F_1(\mathcal{E}_k)| - 1)) \quad (\text{A.15})$$

$$\leq |F_1(\tilde{\mathcal{E}}_k)|(\max F_1(\mathcal{E}_k) - 1 - (|F_1(\mathcal{E}_k)| - 1)) \quad (\text{A.16})$$

$$= (|F_1(\mathcal{E}_k)| - 1)(\max F_1(\mathcal{E}_k) - |F_1(\mathcal{E}_k)|) \quad (\text{A.17})$$

$$\leq \Phi(\mathcal{E}_k) - 1, \quad (\text{A.18})$$

where inequality (A.16) follows from the inequality $\max F_1(\tilde{\mathcal{E}}_k) \leq \max F_1(\mathcal{E}_k) - 1$, Eq. (A.17) follows from the fact that $|F_1(\tilde{\mathcal{E}}_k)| = |F_1(\mathcal{E}_k)| - 1$, and inequality (A.18) follows from the fact that $F_1(\mathcal{E}_k)$ is not contiguous.

On the other hand, if \mathcal{E}_{k+1} is specified according to Eq. (A.12), then we have

$$\Phi(\mathcal{E}_{k+1}) = |F_1(\mathcal{E}_k)| (\max F_1(\mathcal{E}_{k+1}) - |F_1(\mathcal{E}_k)|) \quad (\text{A.19})$$

$$\leq |F_1(\mathcal{E}_k)| (\max F_1(\mathcal{E}_k) - 1 - |F_1(\mathcal{E}_k)|) \quad (\text{A.20})$$

$$= \Phi(\mathcal{E}_k) - |F_1(\mathcal{E}_k)| \leq \Phi(\mathcal{E}_k) - 1, \quad (\text{A.21})$$

where Eq. (A.19) follows from the fact that $|F_1(\mathcal{E}_k)| = |F_1(\mathcal{E}_{k+1})|$, and inequality (A.20) follows from the fact that $\max F_1(\mathcal{E}_{k+1}) \leq \max F_1(\mathcal{E}_k) - 1$. In both cases, we have that if $\Phi(\mathcal{E}_k) > 0$, then

$$\Phi(\mathcal{E}_k) - \Phi(\mathcal{E}_{k+1}) \geq 1.$$

It is straightforward to show that $\Phi(\mathcal{E}_0) < n^2$. It immediately follows that $\Phi(\mathcal{E}_{n^2}) = 0$. This finishes the proof of Claim (i).

Part 2.2: Proof of Claim (ii). According to the procedure we use in generating the sequence of edge sets, inequality (A.13) is trivially satisfied if $\mathcal{E}_{k+1} = \mathcal{E}_k$ or $\mathcal{E}_{k+1} = \tilde{\mathcal{E}}_k$. For the remainder of this part, we show that inequality (A.13) is satisfied if \mathcal{E}_{k+1} is specified according to Eq. (A.12). The key idea in this proof is to show that if the removal of the most expensive producer leads to a strict decrease in social welfare, then a unilateral decrease in the marginal cost of said producer will lead to an increase in social welfare.

We first introduce some notation pertinent to the remainder of the proof. Define the indices g_k and h_k according to

$$g_k = \max F_1(\mathcal{E}_k) \quad \text{and} \quad h_k = \min (F \setminus F_1(\mathcal{E}_k)).$$

Since the set $F_1(\mathcal{E}_k)$ is not contiguous, we must have that $h_k < g_k$. We define a new cost function profile $C^\theta = (C_1^\theta, \dots, C_n^\theta) \in \mathcal{L}^n(c_{\min}, c_{\max})$ according to

$$C_i^\theta(s_i) = \begin{cases} (c_i s_i)^+ & \text{if } i \neq g_k \\ (\theta s_i)^+ & \text{if } i = g_k, \end{cases} \quad (\text{A.22})$$

where θ is a scalar parameter.

The unique Nash equilibrium of the networked Cournot game depends on firms' cost function profiles. With a slight abuse of notation, we denote by $q^{\text{NE}}(\mathcal{E}, C)$ the unique Nash equilibrium of the networked Cournot game associated with an edge set $\mathcal{E} \subseteq F \times \{1\}$, and a cost function profile C .

Our proof relies on the following technical lemma stating the monotonicity of $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C^\theta), C^\theta)$ in the scalar θ . Its proof is deferred to Appendix A.

Lemma A.11. *Let the edge set $\tilde{\mathcal{E}}_k$ and the cost function profile C^θ be specified according to Eq. (A.11) and (A.22), respectively. If $SW(q^{NE}(\mathcal{E}_k, C), C) > SW(q^{NE}(\tilde{\mathcal{E}}_k, C), C)$, then $SW(q^{NE}(\mathcal{E}_k, C^\theta), C^\theta)$ is monotonically decreasing in θ for $c_{\min} \leq \theta \leq c_{g_k}$.*

Note that, by specifying \mathcal{E}_{k+1} according to Eq. (A.12), we essentially replace the most expensive producer with another producer with a cheaper cost. In particular, if \mathcal{E}_{k+1} is specified according to Eq. (A.12), then $SW(q^{NE}(\mathcal{E}_k, C), C)$ and $SW(q^{NE}(\mathcal{E}_{k+1}, C), C)$ are related according to

$$SW(q^{NE}(\mathcal{E}_k, C), C) = SW(q^{NE}(\mathcal{E}_k, C^{c_{g_k}}), C^{c_{g_k}}), \quad (\text{A.23})$$

$$SW(q^{NE}(\mathcal{E}_{k+1}, C), C) = SW(q^{NE}(\mathcal{E}_k, C^{c_{h_k}}), C^{c_{h_k}}). \quad (\text{A.24})$$

Recall that $h_k < g_k$. It follows that $c_{\min} \leq c_{h_k} \leq c_{g_k}$. An application of Lemma A.11 shows that

$$SW(q^{NE}(\mathcal{E}_k, C^{c_{g_k}}), C^{c_{g_k}}) \leq SW(q^{NE}(\mathcal{E}_k, C^{c_{h_k}}), C^{c_{h_k}}).$$

The above inequality, in combination with Eq. (A.23) and (A.24), shows that when \mathcal{E}_{k+1} is specified according to Eq. (A.12), inequality (A.13) is satisfied. This completes the proof.

Proof of Theorem 3.15

Without loss of generality, we assume that $c_1 \leq \dots \leq c_n$. It follows from the description of the greedy algorithm that for each market $j \in M$, $(1, j) \in \mathcal{E}^*$ if and only if $c_1 < \alpha_j$. It immediately follows that the efficient social welfare associated with the edge set \mathcal{E}^* satisfies

$$SW^*(\mathcal{E}^*, C) = \sum_{j=1}^m \frac{((\alpha_j - c_1)^+)^2}{2\beta_j} = SW^*(F \times M, C).$$

For the second part of the theorem, we only provide a proof for the case in which the number of markets $m = 1$. The generalization to the case in which $m > 1$ can be carried out following similar steps as in the proof of Proposition 3.10.

Given the restriction that $c_1 \leq \dots \leq c_n$, we define the subset of ordered linear cost function profiles in $\mathcal{L}^n(c_{\min}, c_{\max})$ according to

$$\mathcal{O}^n(c_{\min}, c_{\max}) = \left\{ C \in \mathcal{L}^n(c_{\min}, c_{\max}) \mid \begin{aligned} C_i(s_i) &= (c_i s_i)^+, \quad i = 1, \dots, n, \\ c_1 &\leq \dots \leq c_n \end{aligned} \right\}.$$

Thus, we have the following chain of inequalities

$$\rho(\mathcal{E}^*, C) = \inf_{\mathcal{E} \subseteq F \times M} \frac{SW^*(F \times M, C)}{SW(q^{NE}(\mathcal{E}), C)} \quad (\text{A.25})$$

$$\leq \sup_{C \in \mathcal{L}^n(c_{\min}, c_{\max})} \inf_{\mathcal{E} \subseteq F \times M} \frac{SW^*(F \times M, C)}{SW(q^{NE}(\mathcal{E}), C)} \quad (\text{A.26})$$

$$= \sup_{C \in \mathcal{O}^n(c_{\min}, c_{\max})} \inf_{\{(1,1)\} \subseteq \mathcal{E} \subseteq F \times M} \frac{SW^*(F \times M, C)}{SW(q^{NE}(\mathcal{E}), C)} \quad (\text{A.27})$$

$$\leq \inf_{\{(1,1)\} \subseteq \mathcal{E} \subseteq F \times M} \sup_{C \in \mathcal{O}^n(c_{\min}, c_{\max})} \frac{SW^*(F \times M, C)}{SW(q^{NE}(\mathcal{E}), C)} \quad (\text{A.28})$$

$$\leq \inf_{\{(1,1)\} \subseteq \mathcal{E} \subseteq F \times M} \frac{1}{\frac{2|\mathcal{E}|+4}{3|\mathcal{E}+5|} + \delta(\gamma_1, |\mathcal{E}|)} \quad (\text{A.29})$$

$$= \frac{1}{\max_{k \in \{1, \dots, n\}} \left\{ \frac{2k+4}{3k+5} + \delta(\gamma_1, k) \right\}}. \quad (\text{A.30})$$

Here, Eq. (A.25) follows from the fact that $SW^*(\mathcal{E}^*, C) = SW^*(F \times M, C)$, Eq. (A.27) follows from our restriction that $c_1 \leq \dots \leq c_n$, inequality (A.28) follows from the min-max inequality, and inequality (A.29) is a direct application of Proposition 3.10. This completes the proof.

Proof of Lemma A.9

The proof proceeds in two parts. In Part 1, we provide a necessary and sufficient condition for the production quantity of firm k to be strictly positive at Nash equilibrium, when the edge set is given by \mathcal{E}_k . In Part 2, we leverage on the intermediary result in Part 1 to show that $SW(q^{NE}(\mathcal{E}_k), C) > SW(q^{NE}(\mathcal{E}_{k-1}), C)$ if and only if inequality (A.9) is satisfied.

Part 1: We show that for the game $(F, Q(\mathcal{E}_k), \pi)$, the production quantity of firm k at Nash equilibrium $q_{k1}^{NE}(\mathcal{E}_k)$ is strictly positive, if and only if

$$\alpha_1 - c_k > \frac{1}{k} \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right). \quad (\text{A.31})$$

We first prove the “if” part of the desired claim. First note that inequality (A.31) implies that the following inequality is satisfied:

$$(k+1)(\alpha_1 - c_k) > \sum_{i=1}^k (\alpha_1 - c_i). \quad (\text{A.32})$$

One can check that if inequality (A.31) is satisfied, then the unique Nash equilibrium of the game $(F, Q(\mathcal{E}_k), \pi)$ is given by

$$q_{i1}^{NE}(\mathcal{E}_k) = \frac{(k+1)(\alpha_1 - c_i) - \sum_{\ell=1}^k (\alpha_1 - c_\ell)}{(k+1)\beta_1}, \quad i = 1, \dots, k,$$

and $q_{i1}^{NE}(\mathcal{E}_k) = 0$ for $i = k+1, \dots, n$. It follows from inequality (A.32) that $q_{k1}^{NE}(\mathcal{E}_k) > 0$.

Next, we prove the “only if” part of the desired claim. First note that $c_i \leq c_k$ for $i = 1, \dots, k$. Recall that for the game $(F, Q(\mathcal{E}_k), \pi)$, firm k 's production quantity $q_{k1}^{NE}(\mathcal{E}_k)$ at Nash equilibrium is strictly positive. It follows that $q_{i1}^{NE}(\mathcal{E}_k) > 0$ for $i = 1, \dots, k$. The first order optimality condition for Nash equilibrium of the game $(F, Q(\mathcal{E}_k), \pi)$ implies that

$$\alpha_1 - \beta_1 \left(\sum_{\ell=1}^k q_{\ell 1}^{NE}(\mathcal{E}_k) \right) - \beta_1 q_{i1}^{NE}(\mathcal{E}_k) - c_i = 0, \quad (\text{A.33})$$

for $i = 1, \dots, k$. Consequently, we have that

$$q_{k1}^{NE}(\mathcal{E}_k) = \frac{(k+1)(\alpha_1 - c_k) - \sum_{\ell=1}^k (\alpha_1 - c_\ell)}{(k+1)\beta_1} > 0.$$

This implies that inequality (A.31) is satisfied.

Part 2: We show $\text{SW}(q^{NE}(\mathcal{E}_k), C) > \text{SW}(q^{NE}(\mathcal{E}_{k-1}), C)$ if and only if inequality (A.9) is satisfied. First note when $k = 1$, it is straightforward to see that $\text{SW}(q^{NE}(\mathcal{E}_1), C) > 0$ if and only if $\alpha_1 - c_1 > 0$. Thus, for the remainder of the proof, we assume that $k \geq 2$.

We only provide a proof for the “only if” part of the claim, as the “if” part of the claim can be proved using similar arguments. First note that $\text{SW}(q^{NE}(\mathcal{E}_k), C) > \text{SW}(q^{NE}(\mathcal{E}_{k-1}), C)$ implies that $q_{k1}^{NE}(\mathcal{E}_k) > 0$. If this is not the case, then we have that $q^{NE}(\mathcal{E}_k) = q^{NE}(\mathcal{E}_{k-1})$, which clearly leads to a contradiction. Note that $c_i \leq c_k$ for $i = 1, \dots, k$. It follows that

$$q_{i1}^{NE}(\mathcal{E}_k) \geq q_{k1}^{NE}(\mathcal{E}_k) > 0, \quad \text{for } i = 1, \dots, k.$$

One can show that the unique Nash equilibrium of the game $(F, Q(\mathcal{E}_k), \pi)$ is given by

$$q_{i1}^{NE}(\mathcal{E}_k) = \frac{(k+1)(\alpha_1 - c_i) - \sum_{\ell=1}^k (\alpha_1 - c_\ell)}{(k+1)\beta_1}, \quad i = 1, \dots, k,$$

and $q_{i1}^{NE}(\mathcal{E}_k) = 0$ for $i = k + 1, \dots, n$. We denote the aggregate supply in market 1 at the unique Nash equilibrium of the game $(F, Q(\mathcal{E}_k), \pi)$ by $d_1^{NE}(\mathcal{E}_k)$. It is given by

$$d_1^{NE}(\mathcal{E}_k) = \sum_{i=1}^n q_{i1}^{NE}(\mathcal{E}_k) = \frac{\sum_{i=1}^k (\alpha_1 - c_i)}{(k+1)\beta_1}$$

Additionally, the social welfare at the Nash equilibrium of the game $(F, Q(\mathcal{E}_k), \pi)$ satisfies

$$\begin{aligned} & SW(q^{NE}(\mathcal{E}_k), C) \\ &= \alpha_1 d_1^{NE}(\mathcal{E}_k) - \frac{1}{2} \beta_1 d_1^{NE2}(\mathcal{E}_k) - \sum_{i=1}^k c_i q_{i1}^{NE}(\mathcal{E}_k) \\ &= \alpha_1 d_1^{NE}(\mathcal{E}_k) - \frac{1}{2} \beta_1 d_1^{NE2}(\mathcal{E}_k) - \sum_{i=1}^k c_i \left(\frac{\alpha_1 - c_i}{\beta_1} - d_1^{NE}(\mathcal{E}_k) \right) \\ &= (k+1)\alpha_1 d_1^{NE}(\mathcal{E}_k) - \frac{1}{2} \beta_1 d_1^{NE2}(\mathcal{E}_k) - \sum_{i=1}^k \left(\frac{c_i(\alpha_1 - c_i)}{\beta_1} \right) + \left(\sum_{i=1}^k c_i - k\alpha_1 \right) d_1^{NE}(\mathcal{E}_k) \\ &= \frac{\alpha_1 \sum_{i=1}^k (\alpha_1 - c_i)}{\beta_1} - \frac{2k+3}{2} \beta_1 d_1^{NE2}(\mathcal{E}_k) - \sum_{i=1}^k \frac{c_i(\alpha_1 - c_i)}{\beta_1} \\ &= \frac{\sum_{i=1}^k (\alpha_1 - c_i)^2}{\beta_1} - \frac{2k+3}{2} \beta_1 d_1^{NE2}(\mathcal{E}_k) \\ &= \frac{\sum_{i=1}^k (\alpha_1 - c_i)^2}{\beta_1} - \frac{2k+3}{2} \frac{\left(\sum_{i=1}^k (\alpha_1 - c_i) \right)^2}{(k+1)^2 \beta_1}. \end{aligned}$$

Next, we provide a closed-form expression for $SW(q^{NE}(\mathcal{E}_{k-1}), C)$. Recall that $q_{k1}^{NE}(\mathcal{E}_k) > 0$. As we showed in Part 1, this implies that inequality (A.31) is satisfied. We thus have

$$\alpha_1 - c_{k-1} \geq \alpha_1 - c_k > \frac{1}{k} \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right),$$

which further implies that

$$\alpha_1 - c_{k-1} > \frac{1}{k-1} \left(\sum_{i=1}^{k-2} (\alpha_1 - c_i) \right).$$

It follows from our result in Part 1 that for the game $(F, Q(\mathcal{E}_{k-1}), \pi)$, producer $k-1$'s production quantity at Nash equilibrium is strictly positive. Using similar arguments

as in our derivation on the closed-form expression for $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C)$, we have the following closed-form expression for $\text{SW}(q^{\text{NE}}(\mathcal{E}_{k-1}), C)$

$$\begin{aligned} & \text{SW}\left(q^{\text{NE}}(\mathcal{E}_{k-1}), C\right) \\ &= \frac{\sum_{i=1}^{k-1}(\alpha_1 - c_i)^2}{\beta_1} - \frac{2k+1}{2} \frac{\left(\sum_{i=1}^{k-1}(\alpha_1 - c_i)\right)^2}{k^2\beta_1}. \end{aligned}$$

Thus, the difference between $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C)$ and $\text{SW}(q^{\text{NE}}(\mathcal{E}_{k-1}), C)$ is given by

$$\begin{aligned} & \text{SW}\left(q^{\text{NE}}(\mathcal{E}_k), C\right) - \text{SW}\left(q^{\text{NE}}(\mathcal{E}_{k-1}), C\right) \\ &= \frac{(\alpha_1 - c_k)^2}{\beta_1} + \frac{2k+1}{2} \frac{\left(\sum_{i=1}^{k-1}(\alpha_1 - c_i)\right)^2}{k^2\beta_1} \\ & \quad - \frac{2k+3}{2} \frac{\left(\sum_{i=1}^k(\alpha_1 - c_i)\right)^2}{(k+1)^2\beta_1}. \end{aligned}$$

An algebraic calculation reveals that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) > \text{SW}(q^{\text{NE}}(\mathcal{E}_{k-1}), C)$ if and only if

$$\frac{\alpha_1 - c_k}{\sum_{i=1}^{k-1}(\alpha_1 - c_i)} < \frac{1}{k}, \quad \text{or} \quad \frac{\alpha_1 - c_k}{\sum_{i=1}^{k-1}(\alpha_1 - c_i)} > \frac{k+2 + \frac{1}{2k}}{k^2 + k - \frac{1}{2}}.$$

Recall that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) > \text{SW}(q^{\text{NE}}(\mathcal{E}_{k-1}), C)$ implies that inequality (A.31) is satisfied. It follows that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C) > \text{SW}(q^{\text{NE}}(\mathcal{E}_{k-1}), C)$ implies inequality (A.9) is satisfied.

Proof of Lemma A.10

We prove this lemma by induction in k .

Base Step: For $k = k^* + 1$, inequality (A.10) is satisfied by the assumption of this lemma.

Induction Step: Assume that inequality (A.10) is satisfied for $k \geq 1$. We show that it is satisfied for $k + 1$ by showing that

$$\alpha_1 - c_{k+1} \leq \frac{1}{k+1} \left(1 + \frac{1}{k+1 - \frac{1}{2(k+2)}} \right) \left(\sum_{i=1}^k (\alpha_1 - c_i) \right). \quad (\text{A.34})$$

Since $c_{k+1} \geq c_k$, inequality (A.34) is satisfied if the following inequality holds

$$(k+1)(\alpha_1 - c_k) \leq \left(1 + \frac{1}{k+1 - \frac{1}{2(k+2)}} \right) \left(\sum_{i=1}^k (\alpha_1 - c_i) \right). \quad (\text{A.35})$$

An algebraic calculation reveals that inequality (A.35) is satisfied if and only if

$$\frac{k^3 + 3k^2 + \frac{1}{2}k - 2}{k^2 + 3k + \frac{3}{2}}(\alpha_1 - c_k) \leq \frac{k^2 + 4k + \frac{7}{2}}{k^2 + 3k + \frac{3}{2}} \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right).$$

Given that $k \geq 1$, the above inequality is satisfied if and only if

$$\alpha_1 - c_k \leq \frac{k^2 + 4k + \frac{7}{2}}{k^3 + 3k^2 + \frac{1}{2}k - 2} \left(\sum_{i=1}^{k-1} (\alpha_1 - c_i) \right). \quad (\text{A.36})$$

The induction hypothesis implies that inequality (A.10) is satisfied for k . Given that $c_i \leq \alpha_1$ for $i = 1, \dots, n$, we have that inequality (A.36) is satisfied if

$$\frac{1}{k} \left(1 + \frac{1}{k - \frac{1}{2(k+1)}} \right) \leq \frac{k^2 + 4k + \frac{7}{2}}{k^3 + 3k^2 + \frac{1}{2}k - 2}. \quad (\text{A.37})$$

Given that $k \geq 1$, inequality (A.37) holds if and only if

$$\begin{aligned} & \left(k^2 + 2k + \frac{1}{2} \right) \left(k^3 + 3k^2 + \frac{1}{2}k - 2 \right) \\ & \leq k \left(k^2 + k - \frac{1}{2} \right) \left(k^2 + 4k + \frac{7}{2} \right). \end{aligned}$$

And the above inequality holds if and only if $(k+1)^2 \geq 0$. Thus, (A.37) is satisfied if $k \geq 1$. This further implies that inequalities (A.34)-(A.36) are all satisfied. Hence, inequality (A.10) also holds for $k+1$. This completes the proof by induction.

Proof of Lemma A.11

Let $n_k = |F_1(\mathcal{E}_k)|$. For the ease of exposition, we assume that the set $F_1(\mathcal{E}_k)$ is given by

$$F_1(\mathcal{E}_k) = \{1, 2, \dots, n_k - 1, g_k\}.$$

However, it is straightforward to generalize the proof to the case in which $F_1(\mathcal{E}_k)$ is any subset of $\{1, 2, \dots, g_k\}$ satisfying $n_k = |F_1(\mathcal{E}_k)|$ and $\max F_1(\mathcal{E}_k) = g_k$.

The remainder proof proceeds in two parts. In Part 1, we provide a closed-form expression for $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C^\theta), C^\theta)$, and show that it is piecewise quadratic in θ . In Part 2, the assumption stated in this lemma implies that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C^\theta), C^\theta)$ is strictly decreasing in θ for $c_{\min} \leq \theta \leq c_{g_k}$.

Part 1: We first show that when $\theta = c_{g_k}$, we have that $q_{i1}^{\text{NE}}(\mathcal{E}, C^\theta) > 0$ for all $i \in F_1(\mathcal{E}_k)$. First recall that

$$\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C), C) > \text{SW}(q^{\text{NE}}(\tilde{\mathcal{E}}_k, C), C).$$

It follows that $q_{g_k 1}^{NE}(\mathcal{E}_k, C) > 0$. If this is not the case, then Nash equilibrium remains unchanged after the removal of producer g_k . This implies that $\text{SW}(q^{NE}(\mathcal{E}_k, C), C) = \text{SW}(q^{NE}(\tilde{\mathcal{E}}_k, C), C)$, which is a contradiction. Recall that $c_i \leq c_{g_k}$ for all $i \in F_1(\mathcal{E}_k)$. In combination with the fact that $q_{g_k 1}^{NE}(\mathcal{E}_k, C) > 0$, this implies that $q_{i 1}^{NE}(\mathcal{E}, C^\theta) > 0$ for all $i \in F_1(\mathcal{E}_k)$.

Assume that when $\theta = c_{\min}$, the vector $q^{NE}(\mathcal{E}_k, C^\theta)$ includes n_{\min} strictly positive entries. Since $|\mathcal{E}_k| = n_k$, we must have that $n_{\min} \leq n_k$. We define a collection of subsets $\Theta_{n_{\min}}, \dots, \Theta_{n_k}$ of the set $[c_{\min}, c_{g_k}]$ according to

$$\Theta_\ell = \left[\alpha_1 + \sum_{r=1}^{\ell-2} (\alpha_1 - c_r) - \ell(\alpha_1 - c_{\ell-1}), \right. \\ \left. \alpha_1 + \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) - (\ell+1)(\alpha_1 - c_\ell) \right] \cap [c_{\min}, c_{g_k}],$$

for $\ell = n_{\min}, \dots, n_k - 1$, and

$$\Theta_{n_k} = [c_{\min}, c_{g_k}] \setminus \bigcup_{\ell=n_{\min}}^{n_k-1} \Theta_\ell.$$

One can check that for any $\theta \in \text{int}(\Theta_\ell)$, $\ell \in \{n_{\min}, \dots, n_k\}$, we have

$$q_{i 1}^{NE}(\mathcal{E}_k, C^\theta) > 0, \text{ for } i = 1, \dots, \ell - 1, \text{ and } i = g_k.$$

That is, the vector $q^{NE}(\mathcal{E}_k, C^\theta)$ contains ℓ strictly positive entries. Additionally, given that the vector $q^{NE}(\mathcal{E}_k, C^\theta)$ includes n_{\min} strictly positive entries when $\theta = c_{\min}$, we can show that

$$\bigcup_{\ell=n_{\min}}^{n_k} \Theta_\ell = [c_{\min}, c_{g_k}], \quad \text{and} \quad \Theta_{\ell_1} \cap \Theta_{\ell_2} = \emptyset$$

for any $\ell_1, \ell_2 \in \{n_{\min}, \dots, n_k\}$ satisfying $\ell_1 \neq \ell_2$.¹

It follows from similar arguments² that $\text{SW}(q^{NE}(\mathcal{E}_k, C^\theta), C^\theta)$ admits the following closed-form expression

$$\text{SW}(q^{NE}(\mathcal{E}_k, C^\theta), C^\theta) = \frac{(\alpha_1 - \theta)^2 + \sum_{r=1}^{\ell-1} (\alpha_1 - c_r)^2}{\beta_1} \\ - \frac{2\ell + 3}{2} \frac{\left(\alpha_1 - \theta + \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) \right)^2}{(\ell + 1)^2 \beta_1}, \quad \text{for } \theta \in \Theta_\ell.$$

¹The collection of sets $\{\Theta_{n_{\min}}, \dots, \Theta_{n_k}\}$ might not be a partition of $[c_{\min}, c_{g_k}]$, as some of these sets can be empty.

²as in the proof of Lemma A.9

We remark that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C^\theta), C^\theta)$ is a piecewise quadratic function of θ that is continuous in θ for $\theta \in [c_{\min}, c_{g_k}]$, and continuously differentiable in θ for $\theta \in \text{int}(\Theta_\ell)$, $\ell \in \{n_{\min}, \dots, n_k\}$.

Part 2: We show that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C^\theta), C^\theta)$ is strictly monotonically decreasing in θ for $\theta \in [c_{\min}, c_{g_k}]$. First recall that the union of the intervals $\Theta_{n_{\min}}, \dots, \Theta_{n_k}$ satisfies

$$\bigcup_{\ell=n_{\min}}^{n_k} \Theta_\ell = [c_{\min}, c_{g_k}].$$

Additionally, $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta)$ is continuous in θ for $\theta \in [c_{\min}, c_{g_k}]$. Thus, in order to show that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta)$ is strictly decreasing in θ on $[c_{\min}, c_{g_k}]$, it suffices to show that

$$\frac{\partial}{\partial \theta} \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) < 0,$$

for $\theta \in \text{int}(\Theta_\ell)$, $\ell = n_{\min}, \dots, n_k$.

For $\theta \in \text{int}(\Theta_\ell)$, we have the following closed-form expression for $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta$:

$$\begin{aligned} & \frac{\partial}{\partial \theta} \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) \\ &= \frac{(2\ell + 3) \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) - (2\ell^2 + 2\ell - 1)(\alpha_1 - \theta)}{(\ell + 1)^2 \beta_1}. \end{aligned}$$

Thus, for each $\theta \in \text{int}(\Theta_\ell)$, $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta < 0$ if the following inequality is satisfied:

$$\alpha_1 - \theta > \frac{2\ell + 3}{2\ell^2 + 2\ell - 1} \sum_{r=1}^{\ell-1} (\alpha_1 - c_r). \quad (\text{A.38})$$

We first show that $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta < 0$ for $\theta \in \text{int}(\Theta_{n_k})$. Recall that $\text{SW}(q^{\text{NE}}(\mathcal{E}_k, C), C) > \text{SW}(q^{\text{NE}}(\tilde{\mathcal{E}}_k, C), C)$. It follows from Lemma A.9 that the following inequality is satisfied

$$\alpha_1 - c_{g_k} > \frac{1}{n_k} \left(1 + \frac{1}{n_k - \frac{1}{2(n_k+1)}} \right) \sum_{r=1}^{n_k-1} (\alpha_1 - c_r). \quad (\text{A.39})$$

It follows from inequality (A.39) that for each $\theta \in \text{int}(\Theta_{n_k})$, we have

$$\begin{aligned} \alpha_1 - \theta &> \alpha_1 - c_{g_k} > \frac{1}{n_k} \left(1 + \frac{1}{n_k - \frac{1}{2(n_k+1)}} \right) \sum_{r=1}^{n_k-1} (\alpha_1 - c_r) \\ &> \frac{2n_k + 3}{2n_k^2 + 2n_k - 1} \sum_{r=1}^{n_k-1} (\alpha_1 - c_r). \end{aligned}$$

Thus, $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta < 0$ for $\theta \in \text{int}(\Theta_{n_k})$.

Next, we show that $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta < 0$ for $\theta \in \text{int}(\Theta_\ell)$, $\ell = n_{\min}, \dots, n_k - 1$. Recall that $c_1 \leq \dots \leq c_{n_k-1} \leq c_{g_k} < \alpha_1$. It follows from a combination of Lemma A.10 and inequality (A.39) that

$$\alpha_1 - c_\ell > \frac{1}{\ell} \left(1 + \frac{1}{\ell - \frac{1}{2(\ell+1)}} \right) \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) \quad (\text{A.40})$$

for $\ell = n_{\min}, \dots, n_k - 1$. Inequality (A.40) implies that

$$\begin{aligned} \ell(\alpha_1 - c_\ell) &> \left(1 + \frac{1}{\ell - \frac{1}{2(\ell+1)}} \right) \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) \\ &> \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) \end{aligned} \quad (\text{A.41})$$

for $\ell = n_{\min}, \dots, n_k - 1$. It follows from inequality (A.41) that the following chain of inequalities are satisfied for $\theta \in \text{int}(\Theta_\ell)$, $\ell = n_{\min}, \dots, n_k - 1$:

$$\begin{aligned} \theta &< \alpha_1 + \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) - (\ell + 1)(\alpha_1 - c_\ell) \\ &= c_\ell + \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) - \ell(\alpha_1 - c_\ell) < c_\ell. \end{aligned}$$

The above inequality, in combination with inequality (A.40), provides the following lower bound on $\alpha_1 - \theta$ for $\theta \in \text{int}(\Theta_\ell)$:

$$\begin{aligned} \alpha_1 - \theta &> \alpha_1 - c_\ell > \frac{1}{\ell} \left(1 + \frac{1}{\ell - \frac{1}{2(\ell+1)}} \right) \sum_{r=1}^{\ell-1} (\alpha_1 - c_r) \\ &> \frac{2\ell + 3}{2\ell^2 + 2\ell - 1} \sum_{r=1}^{\ell-1} (\alpha_1 - c_r). \end{aligned}$$

Thus, $\partial \text{SW}(q^{\text{NE}}(\mathcal{E}_k), C^\theta) / \partial \theta < 0$ for $\theta \in \text{int}(\Theta_\ell)$, $\ell = n_{\min}, \dots, n_k - 1$. This completes the proof.

Proof. Proof of Theorem 3.19

Example A.12. In the case of symmetric linear costs from the previous section, recall that consumer surplus can be written:

$$CS_j^{\text{NE}} = \frac{(\alpha_j - c_1)^2 n^2}{2\beta_j(n+1)^2}$$

and in this case, the overall penalty with n firms can be written:

$$p_j = \theta \frac{(\alpha_j - c_1)^2 n^2}{2\beta_j(n+1)^2} f(n), \quad 0 \leq \theta \leq 1$$

In this scenario, we can again make the (tight) price of anarchy explicit for the case with search costs where $\max_j \alpha_j > c$, i.e. when the solution is non-degenerate:

$$\rho(F \times M, \bar{C}) \leq \frac{(n+1)^2}{n^2(1 - \theta f(n)) + 2n}$$

We can already see in the above example that search costs can bring significant changes to the price of anarchy bounds we obtain. For example, when $\theta = 0$, then we obtain the original bound for open access with symmetric costs firms, suggesting that the larger n is, the better the bound. On the other hand, when $\theta = 1$, as n grows large, $\theta f(n) \rightarrow 1$, and our price of anarchy bound becomes linear in n , the number of firms.

□

Appendix B

PROOFS FOR CHAPTER IV

Proof. Proof of Proposition 4.2: The existence of a unique equilibrium follows an application of Rosen's theorem for concave games (Rosen, 1965). Assume the following parameters for the two-node game.

$$p_1(d_1) = 1 - d_1, \quad p_2(d_2) = 1 - \gamma d_2, \quad c_1(q_1) = \gamma q_1^2, \quad c_2(q_2) = q_2^2,$$

with an infinite transmission capacity on the link between nodes 1 and 2, and with $\gamma > 0$.

We denote the reallocation quantity chosen by the platform to be $r = r_1 = -r_2$, and it is obvious the following stationarity conditions hold at optimal supply profile:

$$q_1^*(1 + 2\gamma) + r^* = 1, \quad q_2^*(2 + \gamma) - \gamma r^* = 1, \quad q_1^* - \gamma q_2^* + (1 + \gamma)r^* = 0,$$

which follows from maximizing social welfare Π with respect to q_1, q_2, r together. We show that

$$q_1^* + r^* = \frac{1 - r^*}{1 + 2\gamma} + r^* = \frac{1 + 2\gamma r^*}{1 + 2\gamma}, \quad q_2^* - r^* = \frac{1 + \gamma r^*}{2 + \gamma} - r^* = \frac{1 - 2r^*}{2 + \gamma}.$$

Now, we can also write

$$r^* = \frac{2\gamma^2 - 2}{6\gamma + 6\gamma^2} \rightarrow \frac{1}{3},$$

and therefore, we have that the demand at both markets are non-negative. The social welfare obtained at the solution corresponding to the above set of equations can be written:

$$\Pi(q_1^*, q_2^*, r^*) = \frac{\gamma + 1}{6\gamma}.$$

On the other hand, stationarity conditions for an equilibrium are the following:

$$q_1^C(2 + 2\gamma) + r^C = 1, \quad q_2^C(2 + 2\gamma) - \gamma r^C = 1, \quad q_1^C - \gamma q_2^C + (1 + \gamma)r^C = 0,$$

which corresponds to a unique solution (q^C, r^C) , which has the following social welfare:

$$\Pi(q^C, r^C) = \frac{5(2\gamma^3 + 7\gamma^2 + 7\gamma + 2)}{8(\gamma^2 + 4\gamma + 1)^2}.$$

Once again, we can show that

$$q_1^C + r^C = \frac{1 - r^C}{2 + 2\gamma} + r^C = \frac{1 + (1 + 2\gamma)r^C}{2 + 2\gamma}, \quad q_2^C - r^C = \frac{1 + \gamma r^C}{2 + 2\gamma} - r^C = \frac{1 - (2 + \gamma)r^C}{2 + 2\gamma}.$$

Similarly, we can show

$$r^C = \frac{\gamma - 1}{1 + 3\gamma + \gamma^2},$$

and it is easy to observe that the demand at both markets are non-negative. Further, it becomes trivial to note that the following limit holds:

$$\lim_{\gamma \rightarrow \infty} \frac{\Pi(q^*, r^*)}{\Pi(q^C, r^C)} = \infty,$$

which shows that the two-node Cournot game can be unboundedly inefficient.

*Since this example is one where the price intercepts are homogeneous, and where there are no network constraints, then the platform is essentially identical irregardless of the objective function between consumer surplus and social welfare. This is because while the stationarity conditions for the social welfare case is

$$\alpha_j - \beta_j = \alpha_m - \beta_m d_m,$$

and the stationarity conditions for other market clearing mechanisms are

$$\alpha_j - k\beta_j d_j = \alpha_m - k\beta_m d_m,$$

for some value $k \in \mathbb{R}$. However, if $\alpha_j = \alpha_m$, then the conditions are essentially the same. This means that the same example with a different platform market clearing mechanism would still result in an infinitely large worst case efficiency loss. \square

Proof. Proof of Theorem 4.1 We prove Theorem 4.1 in three steps. Our first step proves that the networked Stackelberg game is equivalent to a single market Cournot market with linear inverse demand. The results on uniqueness and existence of equilibrium, and efficiency loss is then an application of Theorem 4.3. Next, we show the ‘‘almost necessary’’ nature of this result. In particular, we provide proofs that indeed no equilibrium exists in the $\mathcal{G}([1, 1]^T, [1, 1]^T, [2q^2, q^2]^T, f)$ illustrative game. Next, we show proof of the construction exhibiting multiple equilibrium and efficiency loss growing in the number of nodes (and its generalization).

Consider now the following networked Stackelberg game where markets have homogeneous price intercepts:

$$p_j(d_j) = \alpha - \beta_j d_j,$$

i.e., $\alpha_j = \alpha_k$, for all markets j, k (homogeneous price intercepts). We claim this networked Stackelberg game is equivalent to a Cournot game with one market and linear inverse demand, parameterized by:

$$p(d) = \alpha - \frac{1}{\sum_j \frac{1}{\beta_j}} d,$$

where β_j is defined from the set of markets in the networked Stackelberg game. We first show these two games are identical in the sense that, given a fixed aggregate production q , the resulting prices obtained through the following optimization problem:

$$\max_{d_1, \dots, d_m} \sum_{j=1}^m \int_0^{d_j} \alpha - \beta_j x \, dx - \sum_{i=1}^n c_i(s_i) \quad (\text{B.1})$$

$$s.t. \sum_{j=1}^m d_j = q \quad (\text{B.2})$$

is the same as the price obtained at the single market with quantity q , i.e., $p = \frac{1}{\sum_j \frac{1}{\beta_j}} q$.

By stationarity conditions, we find that

$$\alpha - \beta_j d_j = \alpha - \beta_k d_k \iff \beta_j d_j = \beta_k d_k.$$

For all markets j , aggregate demand q is allocated to markets j in the following manner

$$d_j(q) = \frac{\frac{1}{\beta_j}}{\sum_k \frac{1}{\beta_k}} q \iff p_j(q) = \alpha - \beta_j d_j(q) = \alpha - \frac{1}{\sum_k \frac{1}{\beta_k}} q,$$

which is equivalent to price p at the single market system. Since the prices faced at all aggregate demand are exactly the same, then the two games are essentially equivalent. Invoking Theorem 4.3, we obtain the existence and uniqueness of equilibria, and the corresponding 3/2 worst case efficiency loss.

To prove Proposition 4.4, one can use a similar argument as above to show the equivalence of a set of markets $\{\alpha_k - \beta_k d_k\}_{k=1}^n$ with $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$ is equivalent to a piecewise linear (convex) inverse demand function that links the following points: $(0, \alpha_1), (\frac{\alpha_1 - \alpha_2}{\beta_1^*}, \alpha_2), (\frac{\alpha_1 - \alpha_2}{\beta_1^*} + \frac{\alpha_2 - \alpha_3}{\beta_2^*}, \alpha_3), \dots, (\sum_{j=1}^{k-1} \frac{\alpha_j - \alpha_{j+1}}{\beta_j^*}, \alpha_k), \dots, ((\sum_{j=1}^{n-1} \frac{\alpha_j - \alpha_{j+1}}{\beta_j^* + \frac{\alpha_n}{\beta_n^*}}), 0)$, where β_j^* is the harmonic mean of all prior β_j including itself, i.e.,

$$\beta_k^* = \frac{1}{\sum_{j=1}^k \frac{1}{\beta_j}}.$$

Essentially, the platform allocate starting from markets with largest willingness to pay and access new markets as the current price matches the maximal willingness to pay of these incoming markets. Once a new market is added, prices are constant between all active markets, resulting in the modified price elasticity.

Next, we show that the two examples (and their generalization) we present are indeed as claimed. The following examples are (i) the generalization of the game with multiple equilibria and efficiency loss growing linearly in the number of nodes, and (ii) the example with no equilibria.

Example B.1. *The following example is a generalization of the two node illustrative game $\mathcal{G}([1, \frac{1}{1+\lambda}]^T, [1, \frac{\lambda^2}{1-\lambda^2}]^T, [0, (1 + \epsilon)q]^T, \infty)$ and has multiple equilibria and has unbounded worst case efficiency loss. We consider n nodes in a fully connected graph without constraints. There is a costless firm on node 1, and every other node has a firm with linear cost with marginal cost larger than 1, the largest willingness to pay. Optimally, the costless firm will produce to fulfill all the demand from the consumers. We design market parameters such that the area under each trapezoid is fixed. Recall that since the producer has zero cost, then the social welfare is just the area under the inverse demand curve. As such, the increase of optimal social welfare as we add nodes becomes clear.*

At any Stackelberg equilibrium, only the costless producer will produce. The markets are designed such that there will be an additional Stackelberg equilibrium including the additional market, but the additional markets never affect its profit. Therefore, the worst case is when it chooses to participate only in the first market. As such, the worst case social welfare at a Stackelberg equilibrium remains the same as we add nodes.

Fix $\lambda \in (0, 1/2)$. We design the markets such that the aggregate (convex, piecewise linear) demand curve is as follows:

$$p(d) = \max_{0 \leq k < m} (\lambda^k - \lambda^{2k} d)$$

One can show that this is convex and piecewise linear, combining the points $(0, 1)$, $(\frac{1}{1+\lambda}, \frac{\lambda}{1+\lambda})$, \dots , $(\frac{1}{\lambda^{m-1}}, 0)$.¹ The maximal profit along each piecewise linear component is $1/4$ and at the first segment, the corresponding social welfare is $3/8$. On the other hand, the optimal social welfare is the area under the function

¹Note that these set of markets can be also described by their price intercept and their price elasticity. In particular, the price intercept of the j -th market is $\frac{\lambda^{j-1}}{1+\lambda}$ and the price elasticity is $\frac{\lambda^{2(j-1)}}{1-\lambda^2}$.

$p(d)$. Each trapezoid has an area of $\frac{1-\lambda}{1+\lambda}$, and it is obvious that the worst case multiplicative efficiency loss is $\Omega(m)$.

□

Example B.2. Consider the two-node network with a finite line capacity, and demand and cost functions at nodes $k = 1, 2$ are given by

$$p_1(d_1) = 1 - d_1, \quad p_2(d_2) = 1 - d_2, \quad c_k(q_k) = c_k q_k^2,$$

and assume without loss that $c_1 > c_2 > 0$ and $m > 1$. Together, for these set of parameters, we find that no equilibrium exists if the flow constraint f satisfies the following bound:

$$\frac{\frac{1}{4} \left(\frac{(c_1+1)-(c_2+1)}{(c_1+1)(c_2+1)} \right)}{1 - \frac{1}{4} \left(\frac{1}{c_1+1} + \frac{1}{c_2+1} \right)} < f < \frac{4c_1 - 4c_2}{16c_1c_2 + 8c_1 + 8c_2 + 3},$$

Proof. Proof of bounds in Example B.2 Since we assume $c_1 > c_2 > 0$, then the second producer always produces more and the platform allocates from node 2 to node 1. Now, the profit of each firm then can be written as the following:

$$\Pi_1(q_1, q_2) = \begin{cases} q_1(1 - \frac{q_1+q_2}{2}) - c_1 q_1^2, & \text{if } \frac{q_2-q_1}{2} < f \\ q_1(1 - q_1 - f) - c_1 q_1^2, & \text{otherwise} \end{cases}$$

$$\Pi_2(q_2, q_1) = \begin{cases} q_2(1 - \frac{q_1+q_2}{2}) - c_2 q_2^2, & \text{if } \frac{q_2-q_1}{2} < f \\ q_2(1 - q_2 + f) - c_2 q_2^2, & \text{otherwise} \end{cases}$$

From the profit functions of the firms, we can then look at the stationarity conditions, given the different cases.

$$\frac{\partial \Pi_1}{\partial q_1} = \begin{cases} 1 - \frac{2q_1+q_2}{2} - 2c_1 q_1, & \text{if } \frac{q_2-q_1}{2} < f \\ 1 - 2q_1 - f - 2c_1 q_1, & \text{otherwise} \end{cases}$$

$$\frac{\partial \Pi_2}{\partial q_2} = \begin{cases} 1 - \frac{q_1+2q_2}{2} - 2c_2 q_2, & \text{if } \frac{q_2-q_1}{2} < f \\ 1 - 2q_2 + f - 2c_2 q_2, & \text{otherwise} \end{cases}$$

In this example, a Stackelberg equilibrium exists if $\frac{\partial \Pi_1}{\partial q_1} = \frac{\partial \Pi_2}{\partial q_2} = 0$ and either $\frac{q_2-q_1}{2} < f$ (case 1) or $\frac{q_2-q_1}{2} > f$ (case 2), or $\frac{\partial \Pi_1}{\partial q_1} \leq 0, \frac{\partial \Pi_2}{\partial q_2} \leq 0$ and $\frac{q_2-q_1}{2} = f$ (case

3). We first examine the first case, where $\frac{q_2 - q_1}{2} < f$, i.e., the flow constraints allow for the desired movement of the platform.

Under the first scenario, from the stationary conditions, we have that productions are related by

$$(4c_1 + 1)q_1 = (4c_2 + 1)q_2,$$

and that the sum of productions can be written in the following two manner:

$$q_1 + q_2 = \left(1 + \frac{4c_1 + 1}{4c_2 + 1}\right)q_1 = \left(1 + \frac{4c_2 + 1}{4c_1 + 1}\right)q_2$$

From the stationarity conditions, we can find the production values can be written:

$$q_1 = \frac{2(4c_2 + 1)}{(4c_2 + 1) + (4c_2 + 1) + (4c_1 + 1)(4c_2 + 1)}, \quad q_2 = \frac{2(4c_1 + 1)}{(4c_2 + 1) + (4c_2 + 1) + (4c_1 + 1)(4c_2 + 1)},$$

so no equilibrium of this case can exist if we have the following equation to hold:

$$\frac{4c_1 - 4c_2}{(4c_2 + 1) + (4c_2 + 1) + (4c_1 + 1)(4c_2 + 1)} > f$$

Under the second scenario, from stationarity conditions, we have that

$$q_1 = \frac{1 - f}{2 + 2c_1}, \quad q_2 = \frac{1 + f}{2 + 2c_2},$$

and therefore no such equilibrium exists if the following constraints on f hold:

$$f > \frac{c_1 - c_2}{4(1 + c_1)(1 + c_2) - (2 + c_1 + c_2)}$$

We get our final result by combining the two results. *For the interested reader, one can obtain similar bounds for other objectives that are a convex combination of social welfare and revenue. In particular, there are examples with no equilibria as well for these platform objectives. \square

Proof. Proof of Theorem 4.5

Theorem 4.5 is proved in three steps. The first two steps focuses on the case without network constraints, proving first that the efficiency loss is indeed $3/2$ and that there are conditions under which demand management remains a threat but is not needed at equilibrium. Essentially, the first step again follows the spirit of the proofs from before for efficiency loss—that under these conditions, the game is equivalent to a single market Cournot game with a linear inverse demand. The third step is longer, where we show that when network conditions bind, we can essentially form sets of

smaller networks that each have efficiency loss bounded by $3/2$. The process of forming these sets of smaller networks come at a cost, which is exactly the second term in the efficiency loss value.

Having demand management implies that we drop the $q_j + r_j \geq 0$ constraint. Similar to before, we can show that without network constraints, prices at markets are equivalent, i.e.,

$$\alpha_j - \beta_j d_j = \alpha_k - \beta_k d_k,$$

where d'_j 's may be negative. Again, one can show that this is equivalent to a single-market Cournot competition with the following price function (where without loss we assume $\alpha_1 = \max_j \alpha_j$).

$$p(d) = \left(\alpha_1 - \sum_j \frac{\alpha_1 - \alpha_j}{\beta_j} \right) - \frac{1}{\sum_j \frac{1}{\beta_j}} d$$

Similar to before, we can show that these are equivalent, resulting in a single market, and applying (Ramesh Johari and John N Tsitsiklis, 2004) yields a $3/2$ worst case bound. Now we find conditions under which demand response remains a threat but is not activated at equilibrium. In particular, without loss, let $\alpha_m = \min_j \alpha_j$, then we require:

$$\left(\alpha_1 - \sum_j \frac{\alpha_1 - \alpha_j}{\beta_j} \right) - \frac{1}{\sum_j \frac{1}{\beta_j}} d < \alpha_m$$

Equivalently, we have

$$d > \frac{1}{\sum_j \frac{1}{\beta_j}} \sum_j \frac{\alpha_j - \alpha_m}{\beta_j}.$$

Now we consider the case with network constraints. Assume there exists an equilibrium to the networked Stackelberg game. Let q^{NE} be the production at each location, and $f_{i,j}^S$ the flow at equilibrium between nodes i and j at the Stackelberg equilibrium. We prove the case for one constrained edge but the proof generalizes to multiple edges in a straightforward manner. Let the constrained edge be $\{j_1, j_2\}$ where the constraint lies in the direction from node j_1 to node j_2 . Assume there are no other paths (unconstrained) from j_2 to j_2 , otherwise either there exists an equilibrium that does not bind constraints, or this cannot be an equilibrium.

We now make cuts on the graph to consider smaller graphs with particular kinds of constrained edges. Our aim is to get graphs of smaller size and with specific type of constraints and only on lines that are incident on one "leaf" node. These "leaf" nodes will either have a market fully fulfilled or a production fully utilized at

both equilibrium and the optimal solution. Under these settings, we can then easily compare the smaller graphs and obtain the $3/2$ loss easily in each.

We replace nodes j_1, j_2 by j_1^*, j_2^* in each subgraph, such that j_1 originally is now linked to j_2^* and j_1^* is now linked to j_2 . The original graph with m nodes is now two graphs with a total of $m + 2$ nodes. The replacement outgoing node now has a costless producer and no co-located market. The replacement incoming node now has a market with price at the equilibrium μ_{j_1} and no producer (or an infinitely costly one).

We prove the following lemma, essentially summarizing the impact of these cuts and replacements of nodes on the equilibrium and optimal quantity profile.

Lemma B.3. *With cuts and replacement nodes described above, the following statements are true.*

1. $SW_{new}(q^S, r^S) = SW(q^S, r^S) + \mu_{j_1} f_{j_1, j_2}^S$
2. $SW_{new}(q^*, r^*) \geq SW(q^*, r^*)$
3. $SW_{new}(q^*, r^*) \leq \frac{3}{2} SW_{new}(q^S, r^S)$

Proof. Proof of Lemma B.3: We first prove statement 1. By KKT conditions, the quantity profile is exactly the same in the two games. The difference in social welfare between the new and old cases is the revenue passing through that line.

The second statement follows since we have decreased cost and kept everything else essentially the same. In particular, the optimal profile should also be an optimal profile under the new set-up.

Lastly, one can show that these are “almost equivalent” to the game with no constraints, since both the optimal and equilibria cases use up the flow constraints that the equilibrium profile uses. For example, in the case with the zero cost consumer, then the game is similar to an unconstrained case but with the convex cost piecewise defined to be zero below constraint and linear with parameter greater than the willingness to pay in any market. For the constrained demand that is fulfilled by both optimal and Nash, then we end up with a concave piecewise linear inverse demand, except that the Nash is at the edge of concavity, while the social optimal may produce more potentially. However, even when no such concavity exists, the efficiency loss is $3/2$, so the decreased prices cannot help optimal social welfare.

□

With the above lemma in place, we are now ready to prove the main part of Theorem 4.5. In particular.

$$\begin{aligned}
 SW(q^*, r^*) &\leq SW_{new}(q^*, r^*) \\
 &\leq \frac{3}{2} SW_{new}(q^S, r^S) \\
 &\leq \frac{3}{2} \left[SW(q^S, r^S) + \sum_{(j_1, j_2) \in \mathcal{I}} \mu_{j_1} f_{j_1, j_2}^S \right]
 \end{aligned}$$

Appendix C

PROOFS FOR CHAPTER V

Proof of Proposition 5.1

Proof. We prove this theorem by investigating the KKT conditions of problem *OOIC* and exploring the structure of the optimal solution.

The Lagrangian for problem *OOIC* is defined as

$$L(\mathbf{v}, \lambda, \mu) = \sum_{t=1}^T g_t(v_t) + \lambda(\Delta - \sum_{t=1}^T v_t) + \sum_{t=1}^T v_t \mu(t), \quad (\text{C.1})$$

where $\lambda \geq 0$ and $\mu(t) \geq 0, \forall t \in [T]$ are the Lagrangian multipliers. The following KKT conditions give us sufficient and necessary conditions for optimality:

$$g'_t(v_t) - \lambda + \mu(t) = 0 \quad \forall t \in [T], \quad (\text{C.2})$$

$$\sum_{t=1}^T v_t \leq \Delta, \quad (\text{C.3})$$

$$v_t \geq 0 \quad \forall t \in [T], \quad (\text{C.4})$$

$$\mu(t) \geq 0 \quad \forall t \in [T], \quad (\text{C.5})$$

$$\lambda \geq 0, \quad (\text{C.6})$$

$$v_t \mu(t) = 0 \quad \forall t \in [T], \quad (\text{C.7})$$

$$\lambda \left(\sum_{t=1}^T v_t - \Delta \right) = 0. \quad (\text{C.8})$$

Suppose v^*, μ^* and λ^* are the optimal solutions that satisfy the KKT conditions. Denote the set $\mathcal{T}_0 = \{t | v^*(t) > 0, \forall t \in [T]\}$, then according to the KKT conditions,

we have

$$\mu^*(t) = 0, \quad \forall t \in \mathcal{T}_0, \quad (\text{C.9})$$

$$\lambda^* \left(\sum_{t \in \mathcal{T}_0} v^*(t) - \Delta \right) = 0, \quad (\text{C.10})$$

$$g'_t(v^*(t)) - \lambda^* = 0, \quad \forall t \in \mathcal{T}_0, \quad (\text{C.11})$$

Since g'_t is concave, $g'_t(v^*(t))$ is non-increasing in v_t . According to (C.11) we have

$$g'_t(0) \geq g'_t(v^*(t)) = \lambda^*, \quad \forall t \in \mathcal{T}_0, \quad (\text{C.12})$$

namely,

$$g'_t(0) \geq \lambda^* \quad \forall t \in \mathcal{T}_0. \quad (\text{C.13})$$

Thus given a λ^* , we can use (C.13) to determine the set \mathcal{T}_0 .

For ease of presentation, we denote

$$V_t(\lambda) = \{v | g'_t(v) = \lambda, v \in [0, \Delta]\}.$$

Now consider the following two cases:

(1) $\Delta \geq \max_{v_t \in V_t(0)} \sum_{t=1}^T v_t$. In this case, observe that the solution

$$v^*(t) \in V_t(0), \forall t \in [T],$$

$$\lambda^* = 0,$$

$$\mu^*(t) = 0, \forall t \in [T],$$

satisfies the KKT conditions, thus it is the optimal solution.

(2) $\Delta < \max_{v_t \in V_t(0)} \sum_{t=1}^T v_t$. In this case, we must have $\lambda^* > 0$. According to (C.10) and (C.11), we have

$$v^*(t) \in V_t(\lambda^*) \quad (\text{C.14})$$

$$\sum_{t=1}^T v^*(t) = \Delta \quad (\text{C.15})$$

It is straightforward to check that $v_t, \forall t \in \mathcal{T}_0$ is non-increasing w.r.t. λ . Meanwhile, according to (C.13), we know that the size of set \mathcal{T}_0 is non-increasing w.r.t. λ . Putting together these two observations, we conclude that $\sum_{t \in \mathcal{T}_0} v_t$ is non-increasing w.r.t. λ . Thus given $\Delta > 0$, there exists a unique $\lambda = \lambda^*$ that satisfies $\sum_{t \in \mathcal{T}_0} v^*(t) = \Delta$. Since KKT conditions are sufficient and necessary for optimality of convex problem, we can conclude that λ^* is the optimal dual solution. \square

Proof of Lemma 5.6

Proof of Lemma 5.6. We prove this lemma in the following two steps:

Step I, we prove that $OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \geq g_t(v_t^*) - \lambda_t v_t^*$. To see this, we denote optimal solution at time $\tau \in [t]$ under input $\sigma_1^{[1:t]}$ as v_τ^* . Note that $v_\tau^* \in V_\tau(\lambda_t), \tau \in [t]$ or $v_\tau^* = 0$ if $V_\tau(\lambda_t) = \emptyset$. Similarly, denote optimal solution at time $\tau \in [t-1]$ under input $\sigma_1^{[1:t-1]}$ as \bar{v}_τ . Note that $\bar{v}_\tau \in V_\tau(\lambda_{t-1}), \tau \in [t-1]$ or $\bar{v}_\tau = 0$ if $V_\tau(\lambda_{t-1}) = \emptyset$. Also $v_\tau^* \leq \bar{v}_\tau, \tau \in [t-1]$ (by the non-increasing of $g'_i(v)$ and $\lambda_t \geq \lambda_{t-1}$). Then we have

$$\begin{aligned}
& OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \\
&= \sum_{\tau=1}^t g_\tau(v_\tau^*) - \sum_{\tau=1}^{t-1} g_\tau(\bar{v}_\tau) \\
&= g_t(v_t^*) + \sum_{\tau=1}^{t-1} (g_\tau(v_\tau^*) - g_\tau(\bar{v}_\tau)) \\
&\stackrel{(a)}{\geq} g_t(v_t^*) + \sum_{\tau=1}^{t-1} \lambda_t (v_\tau^* - \bar{v}_\tau) \\
&\stackrel{(b)}{\geq} g_t(v_t^*) - \lambda_t v_t^*
\end{aligned}$$

For (a), it comes from the concavity of $g_\tau(v)$ and $v_\tau^* \leq \bar{v}_\tau, \tau \in [t-1]$. For (b), we claim that $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t v_\tau^*$. To see this, when $\lambda_t = 0$, we must have $\lambda_{t-1} = 0$. In this case, $v_\tau^* = \bar{v}_\tau, \forall \tau \in [t-1]$ and thus we have $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t v_\tau^*$. When $\lambda_t > 0$, from the KKT conditions in (C.10), we have $\sum_{\tau=1}^t v_\tau^* = \Delta \geq \sum_{\tau=1}^{t-1} \bar{v}_\tau$. Then we conclude that $\sum_{\tau=1}^{t-1} \bar{v}_\tau \leq \sum_{\tau=1}^t v_\tau^*$ and consequently, we have $\sum_{\tau=1}^{t-1} (v_\tau^* - \bar{v}_\tau) \geq -v_t^*$.

Step II, we prove that $OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \leq g_t(v_t^*) - \lambda_{t-1} v_t^*$. Similarly, we have

$$\begin{aligned}
& OPT(\sigma^{[1:t]}) - OPT(\sigma^{[1:t-1]}) \\
&= g_t(v_t^*) + \sum_{\tau=1}^{t-1} (g_\tau(v_\tau^*) - g_\tau(\bar{v}_\tau)) \\
&\stackrel{(a)}{\leq} g_t(v_t^*) + \sum_{\tau=1}^{t-1} \lambda_{t-1} (v_\tau^* - \bar{v}_\tau) \\
&\stackrel{(b)}{=} g_t(v_t^*) - \lambda_{t-1} v_t^*
\end{aligned}$$

For (a), it is by the concavity of g_τ : $g_\tau(v_\tau^*) \leq g_\tau(\bar{v}_\tau) + \lambda_{t-1}(v_\tau^* - \bar{v}_\tau)$ (Note that $\lambda_\tau = g'_\tau(\bar{v}_\tau)$) and $\lambda_{t-1} \geq \lambda_\tau, \forall \tau \in [t-1]$. For (b), when $\lambda_{t-1} = 0$, it holds straightly; when $\lambda_{t-1} > 0$, we have $\sum_{\tau=1}^t v_\tau^* = \Delta = \sum_{\tau=1}^{t-1} \bar{v}_\tau$, which implies $\sum_{\tau=1}^{t-1} (v_\tau^* - \bar{v}_\tau) = -v_t^*$ \square

Proof of Lemma. 5.8

Proof. Denote the input under σ as g_t . Denote the input under $\bar{\sigma}$ as \bar{g}_t , The optimal dual variable under $\tilde{\sigma}^{[1:t]}$ (resp. $\bar{\sigma}^{[1:t]}$) as λ_t (resp. $\bar{\lambda}_t$). We have,

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2. \quad (\text{C.16})$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$.

1) If $\lambda_\tau \leq \bar{\lambda}_\tau$, then

$$\begin{aligned} & OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}) \\ & \stackrel{(a)}{\geq} g_\tau(v_\tau) - \lambda_\tau v_\tau \\ & \stackrel{(b)}{\geq} g_\tau(\bar{v}_{\tau+1}) - \lambda_\tau \bar{v}_{\tau+1} \\ & \stackrel{(c)}{\geq} g_\tau(\bar{v}_{\tau+1}) - \bar{\lambda}_\tau \bar{v}_{\tau+1} \\ & \stackrel{(a)}{\geq} OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau]}) \end{aligned}$$

For (a), it is by lemma 5.6. For (b), it is by the concavity of g_t and for (c), it by $\lambda_\tau \leq \bar{\lambda}_\tau$.

2) If $\lambda_\tau \geq \bar{\lambda}_\tau$, then similarly

$$\begin{aligned} & OPT(\bar{\sigma}^{[1:\tau]}) - OPT(\bar{\sigma}^{[1:\tau-1]}) \\ & \stackrel{(a)}{\geq} g_{\tau+1}(\bar{v}_\tau) - \bar{\lambda}_\tau \bar{v}_\tau \\ & \stackrel{(b)}{\geq} g_{\tau+1}(\bar{v}_\tau) - \lambda_\tau \bar{v}_\tau \\ & \stackrel{(c)}{\geq} g_{\tau+1}(v_{\tau+1}) - \lambda_\tau v_{\tau+1} \\ & \stackrel{(a)}{\geq} OPT(\tilde{\sigma}^{[1:\tau+1]}) - OPT(\tilde{\sigma}^{[1:\tau]}) \end{aligned}$$

For (a), it is by lemma 5.6. For (b), it is by $\lambda_\tau \geq \bar{\lambda}_\tau$. For (c), it is by the concavity of g_t . Also, with

$$\begin{aligned} & OPT(\bar{\sigma}^{[1:\tau]}) - OPT(\bar{\sigma}^{[1:\tau-1]}) + OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau]}) \\ & = OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau-1]}) \\ & = OPT(\tilde{\sigma}^{[1:\tau+1]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}) \\ & = OPT(\tilde{\sigma}^{[1:\tau+1]}) - OPT(\tilde{\sigma}^{[1:\tau]}) + OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}), \end{aligned}$$

we can have

$$OPT(\tilde{\sigma}^{[1:\tau+1]}) - OPT(\tilde{\sigma}^{[1:\tau]}) \leq OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}) \quad (\text{C.17})$$

□

Proof of Lemma 5.7

Proof of Lemma 5.7. Suppose an arbitrary $\tilde{\sigma} \in \arg \max_{\sigma} \sum_t v_t$, under which $g'_t(v_t)$ is not non-decreasing in t . That is, exist a τ , $g'_\tau(v_\tau) > g'_{\tau+1}(v_{\tau+1})$. Denote the optimal dual variables under $\tilde{\sigma}^{[1:t]}$ as λ_t . Note that λ_t is non-decreasing in t . Without loss of generality, we assume that $\lambda_t < \lambda_{t+1}$ or $\lambda_t = \lambda_{t+1} = 0$, $\forall t$. We construct a new input sequence $\bar{\sigma}$ by interchanging g_τ and $g_{\tau+1}$ in $\tilde{\sigma}$ and denote the input under $\bar{\sigma}$ as \bar{g}_t , the output of $CRP(\pi^*)$ under $\bar{\sigma}$ as \bar{v}_t . The optimal dual variable under $\bar{\sigma}^{[1:t]}$ as $\bar{\lambda}_t$. By definition, we can easily observe that,

$$OPT(\tilde{\sigma}^t) = OPT(\bar{\sigma}^t), \forall t \leq \tau - 1 \vee t \geq \tau + 1; \quad (\text{C.18})$$

$$v_t = \bar{v}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2; \quad (\text{C.19})$$

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2. \quad (\text{C.20})$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$. We claim that $\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t$ and $\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1})$. To see this, consider the following two cases:

Case I, $\lambda_\tau = \lambda_{\tau+1} = 0$. Under this case, we have $OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}) = OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau]}) = g_\tau(\hat{v}_\tau)$, where $\hat{v}_\tau = \arg \max_v g_\tau(v)$. Then $v_\tau = \bar{v}_{\tau+1}$. Similarly, we have $v_{\tau+1} = \bar{v}_\tau$. $\sum_t v_t = \sum_t \bar{v}_t$. We conclude that $\bar{\sigma} \in \arg \max_{\sigma} \sum_t v_t$ and $\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1})$.

Case II, $0 \leq \lambda_\tau < \lambda_{\tau+1}$. First, we have

$$\begin{aligned} g_\tau(v_\tau) + g_{\tau+1}(v_{\tau+1}) &= \frac{OPT(\tilde{\sigma}^{[\tau+1]}) - OPT(\tilde{\sigma}^{[\tau-1]})}{\pi^*} \\ &= \frac{OPT(\bar{\sigma}^{[\tau+1]}) - OPT(\bar{\sigma}^{[\tau-1]})}{\pi^*} \\ &= g_{\tau+1}(\bar{v}_\tau) + g_\tau(\bar{v}_{\tau+1}), \end{aligned}$$

which implies

$$g_\tau(v_\tau) - g_\tau(\bar{v}_{\tau+1}) = g_{\tau+1}(\bar{v}_\tau) - g_{\tau+1}(v_{\tau+1}). \quad (\text{C.21})$$

Second, we claim that $\bar{v}_{\tau+1} \leq v_\tau$. From Lemma 5.8, we have

$$OPT(\bar{\sigma}^{[1:\tau+1]}) - OPT(\bar{\sigma}^{[1:\tau]}) \leq OPT(\tilde{\sigma}^{[1:\tau]}) - OPT(\tilde{\sigma}^{[1:\tau-1]}).$$

Then $g_\tau(v_\tau) \geq g_\tau(\bar{v}_{\tau+1})$ and $\bar{v}_{\tau+1} \leq v_\tau$ are straightforward.

Third, we show $g_\tau(v_\tau) = g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} = v_\tau$ by contradiction. Suppose $g_\tau(v_\tau) > g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} < v_\tau$. we show that $\sum_t v_t < \sum_t \bar{v}_t$ which contradict the fact that $\tilde{\sigma} \in \arg \max_\sigma \sum_t v_t$. To see this, observe that we have

$$\begin{aligned}
& g'_{\tau+1}(v_{\tau+1})(v_\tau - \bar{v}_{\tau+1}) \\
& \stackrel{(a)}{<} -g'_\tau(v_\tau)(\bar{v}_{\tau+1} - v_\tau) \\
& \stackrel{(b)}{\leq} g_\tau(v_\tau) - g_\tau(\bar{v}_{\tau+1}) \\
& = g_{\tau+1}(\bar{v}_\tau) - g_{\tau+1}(v_{\tau+1}) \\
& \stackrel{(b)}{\leq} g'_{\tau+1}(v_{\tau+1})(\bar{v}_\tau - v_{\tau+1})
\end{aligned}$$

For (a), it is by $g'_\tau(v_\tau) > g'_{\tau+1}(v_{\tau+1}) \geq \lambda_{t+1} > 0$ and $\bar{v}_{\tau+1} < v_\tau$. For (b), it is from the concavity of g_τ . As $g'_{\tau+1}(v_{\tau+1}) \geq \lambda_{t+1} > 0$, we have

$$v_\tau + v_{\tau+1} < \bar{v}_\tau + \bar{v}_{\tau+1} \tag{C.22}$$

, which leads to $\sum_t v_t < \sum_t \bar{v}_t$.

So we conclude that $g_\tau(v_\tau) = g_\tau(\bar{v}_{\tau+1})$ and thus $\bar{v}_{\tau+1} = v_\tau$. Consequently, $g_\tau(v_{\tau+1}) = g_{\tau+1}(\bar{v}_\tau)$ and thus $\bar{v}_\tau = v_{\tau+1}$. It is then straightforward that

$$\bar{\sigma} \in \arg \max_\sigma \sum_t v_t; \tag{C.23}$$

$$\bar{g}'_\tau(\bar{v}_\tau) = g'_{\tau+1}(v_{\tau+1}) < g'_\tau(v_\tau) = \bar{g}'_{\tau+1}(\bar{v}_{\tau+1}) \tag{C.24}$$

By continuously interchanging g_τ and $g_{\tau+1}$ which fails to satisfy $g'_{\tau+1}(v_\tau) \leq g'_\tau(v_{\tau+1})$, we finally attain a sequence $\in \arg \max_\sigma \sum_t v_t$ such that $g'_t(v_t)$ is non-decreasing in t . \square

To prove this theorem, we use the following sequence of lemmas. For ease of presentation, we first define $\tilde{p}(t) = OPT(\sigma^{[1:t]})/\Delta$. Here $\tilde{p}(t)$ can be interpreted as the weighted averaged trading price for the offline algorithm when the input sequence is $\sigma^{[1:t]}$. Lemma C.1 limits the amount of inventory required to pursue the competitive ratio, in terms of the problem-dependent factor c , the competitive ratio π , the inventory size Δ , and the difference between the subsequent weighted averaged trading prices of the offline optimal solution. Lemma C.2 seeks to bound the subsequent differences, which help in the proof of Lemma C.3 that serves as an

upper bound relying only on θ and not the actual sequence of functions or prices. According to (5.9) and (5.10), the output of the algorithm CR-Pursuit(π) at slot t is v_t that satisfies

$$\frac{\Delta(\tilde{p}(t) - \tilde{p}(t-1))}{\pi} = g_t(v_t). \quad (\text{C.25})$$

We begin with the following lemma, which gives an upper bound on the amount of inventory the online algorithm consumes at each time step to maintain the offline-to-online revenue ratio.

Lemma C.1. *For any input sequence $\sigma^{[1:T]}$, we have*

$$v_t \leq c \frac{\Delta(\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)},$$

where $p(t) = g'_t(0)$.

We prove Lemma C.1 in the appendix. The main idea of the proof is to show that the condition on c can lead to the inequality $g_t(v_t) \leq g_t(c \frac{\Delta(\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)})$, and the result then follows by the fact that $g_t(v)$ is concave and increasing in v . Next, we prove a lemma that allows us to bound the sum over all subsequent optimal average prices, which then helps us limit the impact of price changes and thereby achieve the competitive ratio bound.

Lemma C.2. *For any input sequence $\sigma^{[1:T]}$, $\forall p \in [m, M]$, we have $x_p \leq p$, where $x_p \triangleq \sum_{t, p(t) \leq p} (\tilde{p}(t) - \tilde{p}(t-1))$.*

To prove Lemma C.2, we first observe that if all the $p(t) \in \sigma^{[1:T]}$ is less than p , the result is immediate. As for general cases, based on Lemma 5.8, we can construct new input sequences by moving forward the slots with $p(t) \leq p$ in σ , while increasing x_p . At last, we can attain an input sequence with larger x_p , while the slots with $p(t) \leq p$ of it is all at the front and following the first observation, its x_p is bounded by p . Lemma C.2 then allows us to bound a component of the worst case competitive ratio in the following lemma, eventually used to prove Theorem 5.9.

Lemma C.3. *For any input sequence $\sigma^{[1:T]}$, we have*

$$\sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)} \leq \ln \theta + 1, \quad (\text{C.26})$$

where $\theta = M/m$.

The idea to prove Lemma C.3 is to construct an optimization problem that maximizes the left hand size in (C.26), subjected to the constraint from Lemma C.2 and show its maximum value is equal to the right hand size in (C.26).

We are now ready to prove the upper bound on $\Phi_\Delta(\pi)$, i.e., the result in Theorem 5.9.

Proof of Theorem 5.9. From Lemma C.1 and C.3, we have

$$\begin{aligned}\Phi_\Delta(\pi) &= \max_{\sigma^{[1:T]}} \sum_{t=1}^T v_t \leq \sum_{t=1}^T c \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)} \\ &\leq c \frac{\Delta}{\pi} (\ln \theta + 1).\end{aligned}$$

By solving $c \frac{\Delta}{\pi} (\ln \theta + 1) = \Delta$, we get that $\bar{\pi} = c(\ln \theta + 1)$. Thus by setting $\pi = \bar{\pi}$, we have $\Phi_\Delta(\bar{\pi}) \leq \Delta$ and CR-Pursuit($\bar{\pi}$) is $\bar{\pi}$ -competitive. It then immediately follows that $\bar{\pi}$ is an upper bound for the optimal competitive ratio π^* . \square

Recall that since CR-Pursuit is at best $(\ln \theta + 1)$ -competitive, then the result in Theorem 5.9 implies that CR-Pursuit achieves the optimal competitive ratio (up to a problem-dependent factor c) among all deterministic online algorithms.

Proof of Lemma C.1

Proof. First, from Proposition 5.3, we easily conclude that $v_t \leq \hat{v}_t$. So if $c \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)} \geq \hat{v}_t$, it's trivial. We now assume that $c \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)} \leq \hat{v}_t$. As $g_t(v)$ is a concave increasing function in $[0, \hat{v}_t]$ and thus it's equivalent to show that $g_t(v_t) \leq g_t(c \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)})$. For ease of presentation, we denote $k = g_t(v_t) = \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi}$. Then we have the following sequence of equivalent (or consequent) statements,

$$\begin{aligned}g_t(v_t) &\leq g_t\left(c \frac{\Delta (\tilde{p}(t) - \tilde{p}(t-1))}{\pi p(t)}\right) \\ \iff &k \leq g_t\left(\frac{ck}{p(t)}\right) \\ \stackrel{(a)}{\iff} &k \leq \frac{ck}{p(t)\hat{v}_t} g_t(\hat{v}_t) \\ \iff &c \geq \frac{g_t'(0)}{g_t(\hat{v}_t)/\hat{v}_t},\end{aligned}$$

where the last inequality holds by the definition of c . For (a), by the concavity of g_t , $g_t(0) = 0$, and $\frac{ck}{p(t)} \leq \hat{v}_t$, we have

$$g_t\left(\frac{ck}{p(t)}\right) = g_t\left(\frac{ck}{p(t)\hat{v}_t} \hat{v}_t + \left(1 - \frac{ck}{p(t)\hat{v}_t}\right) 0\right) \geq \frac{ck}{p(t)\hat{v}_t} g_t(\hat{v}_t).$$

\square

Proof of Lemma C.2

Proof. Denote $T_1 \triangleq \min\{t : p(\tau) > p, \forall \tau \geq t\} - 1$, i.e., for any $t > T_1$, we have $p(t) > p$, or equivalently if $p(t) \leq p$, then $t \leq T_1$. By definition, x_p is determined by $\sigma^{[1:T_1]}$ only. Thus, in this proof, we only focus on the input horizon $t \in [T_1]$.

First, we consider a special case when $p(t) \leq p, \forall t \in [T_1]$. In this case, we have

$$\begin{aligned} x_p &= \sum_{t \in [T_1], p(t) \leq p} (\tilde{p}(t) - \tilde{p}(t-1)) \\ &= \sum_{t=1}^{T_1} \tilde{p}(t) - \tilde{p}(t-1) \\ &= \tilde{p}(T_1) \\ &= \frac{OPT(\sigma^{[1:T_1]})}{\Delta} \\ &\leq \frac{\sum_{t=1}^{T_1} p(t)v_t^*}{\Delta}, \end{aligned}$$

where $v_t^*, t \in [T_1]$ are the solution of the optimal offline algorithm under input $\sigma^{[1:T_1]}$. Here, the last inequality follows from the fact that $g_t(v), \forall t \in [T_1]$ are concave functions and we have

$$OPT(\sigma^{[1:T_1]}) = \sum_{t=1}^{T_1} g_t(v_t^*) \leq \sum_{t=1}^{T_1} (g_t(0) + g_t'(0)v_t^*) = \sum_{t=1}^{T_1} p(t)v_t^*.$$

Further, since $\sum_{t=1}^{T_1} v_t^* \leq \Delta$ and $p(t) \leq p, \forall t \in [T_1]$, we have $x_p \leq \frac{\sum_{t=1}^{T_1} p(t)v_t^*}{\Delta} \leq p$.

Second, we now consider the general cases. Suppose exist a slot $\tau (\tau \leq T_1)$ such that $p(\tau) > p$. Then we construct a new input sequence $\bar{\sigma}$ by interchange g_τ and $g_{\tau+1}$ in σ . Denote the input under $\bar{\sigma}$ as \bar{g}_t . Let $\bar{x}_p, \bar{p}(t)$ be the corresponding variables under $\bar{\sigma}$. To show that $x_p \leq p$, we first show $x_p \leq \bar{x}_p$. We then prove $x_p \leq p$ as follows: We continuously interchange with $p(\tau) > p$ with the input at its next slot until all the slots with $p(t) \leq p$ is at the front of it. At the meantime, x_p keeps on non-decreasing. Finally, we get a σ' , in which the price at each slot in $[T_1']$ (T_1' is corresponding to T_1 but defined under σ') is less or equal to p , and $x_p \leq x_p'$. Since in $\sigma', p \geq p(t), \forall t$, from our analysis in the first part (special case), we have $x_p' \leq p$. It then follows that $x_p \leq p$.

We now prove $x_p \leq \bar{x}_p$. By definition, we can easily observe that,

$$OPT(\sigma^t) = OPT(\bar{\sigma}^t), \forall t \leq \tau - 1 \vee t \geq \tau + 1; \quad (\text{C.27})$$

$$g_t = \bar{g}_t, \forall t \leq \tau - 1 \vee t \geq \tau + 2. \quad (\text{C.28})$$

Besides, $g_\tau = \bar{g}_{\tau+1}, g_{\tau+1} = \bar{g}_\tau$.

1)if $p(\tau + 1) > p$, it is easy to see that $\bar{x}_p = x_p$.

2)if $p(\tau + 1) \leq p$, we have

$$\begin{aligned} x_p - \bar{x}_p &= \tilde{p}(\tau + 1) - \tilde{p}(\tau) - (\bar{p}(\tau) - \bar{p}(\tau - 1)) \\ &= \frac{OPT(\sigma^{[1:\tau+1]}) - OPT(\sigma^{[1:\tau]})}{\Delta} \\ &\quad - \frac{OPT(\bar{\sigma}^{[1:\tau]}) - OPT(\bar{\sigma}^{[1:\tau-1]})}{\Delta} \\ &\stackrel{(a)}{\leq} 0 \end{aligned}$$

(a) is by lemma. 5.8.

□

Proof of Lemma C.3

Proof. Suppose in $\sigma^{[1:T]}$, $p(t)$ takes n different values, which are denoted as $m \leq p_1 \leq p_2 \leq \dots \leq p_n \leq M$. And define $y_i \triangleq \sum_{t, p(t)=p_i} (\tilde{p}(t) - \tilde{p}(t-1))$. Note that we have

$$\sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)} = \sum_{i=1}^n \frac{y_i}{p_i}. \quad (\text{C.29})$$

From Lemma C.2, we have $\sum_{j=1}^i y_j = x_{p_i} \leq p_i$.

Consider the following optimization problem:

$$\begin{aligned} \max \quad & \sum_{i=1}^n \frac{y_i}{p_i} \\ \text{s.t.} \quad & \sum_{j=1}^i y_j \leq p_i, i \in [n] \\ & y_i \geq 0, i \in [n]. \end{aligned}$$

The KKT conditions are sufficient and necessary conditions for optimality for the above convex problem. Denote $\mu_i \geq 0, i \in [n]$ as the dual variables, then the KKT

conditions can be expressed as:

$$\frac{1}{p_i} - \sum_{j=1}^{n+1-i} \mu_j = 0, \forall i \in [n], \quad (\text{C.30})$$

$$\mu_i(p_i - \sum_{j=1}^i y_j) = 0, \forall i \in [n], \quad (\text{C.31})$$

$$\mu_i \geq 0, \forall i \in [n],$$

$$y_i \geq 0, \forall i \in [n].$$

From (C.30), we know that $\mu_i > 0$ for all $i \in [n]$. Thus from (C.31), we have

$$p_i - \sum_{j=1}^i y_j = 0, \forall i \in [n]. \quad (\text{C.32})$$

Thus we know the optimal primal solution is

$$y_i = p_i - p_{i-1}, \forall i \in [n], \quad (\text{C.33})$$

where $p_0 = 0$. And the optimal objective value equals to $\sum_{i=1}^n \frac{p_i - p_{i-1}}{p_i}$.

So

$$\begin{aligned} \sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)} &= \sum_{i=1}^n \frac{y_i}{p_i} \\ &\leq \sum_{i=1}^n \frac{p_i - p_{i-1}}{p_i} \\ &= \frac{p_1}{p_1} + \sum_{i=2}^n \frac{p_i - p_{i-1}}{p_i} \\ &\leq 1 + \int_{p_1}^{p_n} \frac{1}{x} dx \\ &\leq 1 + \ln \theta. \end{aligned}$$

This completes our proof. □

Proof of Lemma 5.10

Proof. We show that any input $\sigma^{[1:T]}$ is equivalent to an increasing price sequence as the following:

$$m \leq p_1 < p_2 < \cdots < p_n \leq M, \quad (\text{C.34})$$

where $n \leq T$. According to (5.12), $CRP(\pi)$ will sell only when the current price is larger than the highest price in history. Thus for any input $\sigma^{[1:T]}$, we can delete the

slots when $CRP(\pi)$ does not sell, and the outputs of $CRP(\pi)$ is then equivalent to the resulting increasing price sequence. \square

Proof of Theorem 5.12

Proof. Consider an arbitrary deterministic online algorithm different from $CRP(\pi^*)$, denoted as \mathcal{A} . Using an adversary argument we show that \mathcal{A} cannot achieve a ratio smaller than π^* .

For \mathcal{A} and $CRP(\pi^*)$, denote the output at time t as $v_{\mathcal{A}}(t)$ and v_t , respectively. For ease of presentation, denote

$$\tilde{\sigma}^{[1:T]} = \{\tilde{p}(1), \tilde{p}(2), \dots, \tilde{p}(T)\}$$

as the worst case input for $CRP(\pi^*)$, i.e., under this input, we have $\sum_{\tau=1}^T v(\tau) = \Phi_{\Delta}(\pi^*) = \Delta$. According to Lemma 5.10, we must have $m \leq \tilde{p}(1) < \tilde{p}(2) < \dots < \tilde{p}(T) \leq M$.

We present $\tilde{p}(1)$ to \mathcal{A} at the first slot. If $v_{\mathcal{A}}(1) \leq v(1)$, then we end the trading period, i.e., $T = 1$. In this case, we have $\tilde{p}(1)v_{\mathcal{A}}(1) \leq \tilde{p}(1)v(1) = OPT(\tilde{\sigma}^{[1:1]})/\pi^*$, thus the competitive ratio of \mathcal{A} is at least π^* . Otherwise, if $v_{\mathcal{A}}(1) > v(1)$, we continue to present $\tilde{p}(2)$ to \mathcal{A} . In general, if at time t the total amount of inventory \mathcal{A} consumed is no larger than $\sum_{\tau=1}^t v(\tau)$, we immediately end the trading period. Otherwise, we continue and present \mathcal{A} with the next input. Let t' be the minimum t such that at the end of time t , the total amount of inventory \mathcal{A} consumed is less than $\sum_{\tau=1}^t v(\tau)$. We note that t' must exist, otherwise we have at time T , $\sum_{\tau=1}^T v_{\mathcal{A}}(\tau) > \sum_{\tau=1}^T v(\tau) = c$ is a contradiction. Then we have

$$\begin{aligned} v_{\mathcal{A}}(1) &> v(1) \\ \sum_{\tau=1}^2 v_{\mathcal{A}}(\tau) &> \sum_{\tau=1}^2 v(\tau) \\ &\dots \\ \sum_{\tau=1}^{t'-1} v_{\mathcal{A}}(\tau) &> \sum_{\tau=1}^{t'-1} v(\tau) \end{aligned}$$

and by the definition of t' , we have

$$\sum_{\tau=1}^{t'} v_{\mathcal{A}}(\tau) \leq \sum_{\tau=1}^{t'} v(\tau). \quad (\text{C.35})$$

Since $\tilde{p}(\tau)$ are increasing in τ , \mathcal{A} would have achieved a larger revenue by selling exactly $v(\tau)$ for any $\tau \in [t' - 1]$ and by selling $v_{\mathcal{A}}^*(t') = v_{\mathcal{A}}(t') + \sum_{\tau=1}^{t'-1} v_{\mathcal{A}}(\tau) - \sum_{\tau=1}^{t'-1} v(\tau)$ at time t' . Namely, we have

$$\eta_{\mathcal{A}}^{t'} \leq \sum_{\tau=1}^{t'-1} \tilde{p}(\tau)v(\tau) + \tilde{p}(t')v_{\mathcal{A}}^*(t'), \quad (\text{C.36})$$

where $\eta_{\mathcal{A}}^{t'}$ is the revenue of \mathcal{A} up to time t' . However, from (C.35), we know $v_{\mathcal{A}}^*(t') \leq v(t')$ and thus we have

$$\eta_{\mathcal{A}}^{t'} \leq \sum_{\tau=1}^{t'} \tilde{p}(\tau)v(\tau) = OPT(\tilde{\sigma}^{[1:t']})/\pi^*. \quad (\text{C.37})$$

Thus the competitive ratio of \mathcal{A} should at least be π^* .

It follows that \mathcal{A} must coincide with $CRP(\pi^*)$, achieving a ratio of π^* , or otherwise \mathcal{A} incurs a higher ratio on $\tilde{\sigma}^{[1:T]}$. \square

Proof of Lemma C.4

Proof. First, by Proposition 5.3, we know that $v_t \leq \hat{v}_t$, where \hat{v}_t as the optimizer of $g_t(v_t)$. So if $c \frac{\Delta(\tilde{p}(t)-\tilde{p}(t-1))}{p(t)} \geq \hat{v}_t$, it's trivial. We now assume $c \frac{\Delta(\tilde{p}(t)-\tilde{p}(t-1))}{p(t)} \leq \hat{v}_t$. As $g_t(v)$ is a concave function, $g_t(v)$ is increasing in $[0, \hat{v}_t]$ and thus it's equivalent to show that $g_t(v_t) \leq g_t(c \frac{\Delta(\tilde{p}(t)-\tilde{p}(t-1))}{p(t)})$. To simplify the explanation, let $k = g_t(v_t) = \frac{\Delta(\tilde{p}(t)-\tilde{p}(t-1))}{\pi}$. By simple calculation, it's equivalent to show the following inequalities:

$$\begin{aligned} g_t(v_t) &\leq g_t\left(\frac{2\Delta}{\pi(1+\sqrt{1-\frac{1}{\pi}})} \frac{(\tilde{p}(t)-\tilde{p}(t-1))}{p(t)}\right) \\ \iff k &\leq (p(t) - f_t\left(\frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)}\right)) \frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)} \\ \iff f_t\left(\frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)}\right) &\leq \frac{1-\sqrt{1-\frac{1}{\pi}}}{2} p(t) \\ \stackrel{(a)}{\iff} f_t(\hat{v}_t) \frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)} \frac{1}{\hat{v}_t} &\leq \frac{1-\sqrt{1-\frac{1}{\pi}}}{2} p(t) \\ \stackrel{(b)}{\iff} f_t(\hat{v}_t) \frac{2}{\pi(1+\sqrt{1-\frac{1}{\pi}})} \frac{g_t(v_t^*)}{p(t)} \frac{1}{\hat{v}_t} &\leq \frac{1-\sqrt{1-\frac{1}{\pi}}}{2} p(t) \\ \iff f_t(\hat{v}_t) \frac{g_t(\hat{v}_t)}{\hat{v}_t} &\leq \frac{p^2(t)}{4} \\ \iff f_t(\hat{v}_t)(p_t - f_t(\hat{v}_t)) &\leq \left(\frac{p(t)-f_t(\hat{v}_t)+f_t(\hat{v}_t)}{2}\right)^2 \end{aligned}$$

The last inequality follows by inequality of arithmetic and geometric means. For (a), by the convexity of f_t , $f_t(0) = 0$, and $\frac{2k}{\pi(1+\sqrt{1-\frac{1}{\pi}})p(t)} \leq \hat{v}_t$, we have

$$f_t\left(\frac{2k}{\pi(1+\sqrt{1-\frac{1}{\pi}})p(t)}\right) \leq f_t(\hat{v}_t) \frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)} \frac{1}{\hat{v}_t}.$$

For (b), by lemma 5.3, we have

$$f_t(\hat{v}_t) \frac{2k}{(1+\sqrt{1-\frac{1}{\pi}})p(t)} \frac{1}{\hat{v}_t} \leq f_t(\hat{v}_t) \frac{2}{\pi(1+\sqrt{1-\frac{1}{\pi}})} \frac{g_t(\hat{v}_t)}{p(t)} \frac{1}{\hat{v}_t}$$

□

Proof of Lemma C.5

Proof. From Lemma C.4, we have

$$\Phi_{\Delta}(\pi) = \max_{\sigma^{[1:T]}} \sum_{t=1}^T v_t \leq \frac{2\Delta}{\pi \left(1 + \sqrt{1 - \frac{1}{\pi}}\right)} \sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)}.$$

By Lemma C.3, we know that

$$\sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)} \leq 1 + \ln \theta. \quad (\text{C.38})$$

Then we can bound $\Phi_{\Delta}(\pi)$ as

$$\begin{aligned} \Phi_{\Delta}(\pi) &\leq \frac{2\Delta}{\pi(1+\sqrt{1-\frac{1}{\pi}})} \left(\sum_{t=1}^T \frac{\tilde{p}(t) - \tilde{p}(t-1)}{p(t)} \right) \\ &\leq \frac{2\Delta}{\pi(1+\sqrt{1-\frac{1}{\pi}})} (1 + \ln \theta) = \bar{\Phi}(\pi). \end{aligned}$$

This completes our proof. □

Recall that $\tilde{p}(t) \triangleq \frac{OPT(\sigma^{[1:t]})}{\Delta}$, and we have $\tilde{p}(t) \in [0, M]$ and $\tilde{p}(t)$ is non-decreasing in t . The output of the algorithm CR-Pursuit(π) at time t satisfies

$$\frac{\Delta(\tilde{p}(t) - \tilde{p}(t-1))}{\pi} = v_t(p(t) - f_t(v_t)). \quad (\text{C.39})$$

Restating Lemma C.1 under the parametric assumptions of $g_t(v_t)$ in the one-way trading with price elasticity setting, we can upper bound the selling quantity of CR-Pursuit(π) at each slot using a more precise value of c , reflected in the following lemma.

Lemma C.4. For any input sequence $\sigma^{[1:T]}$, we have

$$v_t \leq c \frac{\Delta \tilde{p}(t) - \tilde{p}(t-1)}{p(t)},$$

where $c = 2 / \left(1 + \sqrt{1 - \frac{1}{\pi}}\right)$.

We note that the bound in Lemma C.4 ($c = 2 / \left(1 + \sqrt{1 - \frac{1}{\pi}}\right)$) is related to π and is always tighter than in Lemma C.1 ($c = 2$). The idea of the proof is similar to that of Lemma C.1, but we further utilize the special structure of $g_t(\cdot)$ here (i.e., the convexity of $f_t(\cdot)$). A proof is in the Appendix. The tighter bound allows us to develop an online algorithm with better competitive ratio as compared to the one obtained as a result of Sec. 5.3.

Lemma C.5. For CR-Pursuit(π) with $\pi \geq 1$, we have $\Phi_\Delta(\pi) \leq \bar{\Phi}(\pi)$, where $\bar{\Phi}(\pi) \triangleq \frac{2\Delta}{\pi(1+\sqrt{1-\frac{1}{\pi}})}(1 + \ln \theta)$.

Lemma C.5 shows that $\Phi_\Delta(\pi)$ is upper bounded by $\bar{\Phi}(\pi)$. It is easy to show that $\bar{\Phi}(\pi)$ is decreasing in $\pi \geq 1$. Thus by setting $\bar{\Phi}(\bar{\pi}) = \Delta$, we can guarantee that CR-Pursuit($\bar{\pi}$) is feasible. Then we have the following result, which shows that the competitive ratio of CR-Pursuit($\bar{\pi}$) is $\ln \theta + \Omega(1)$.

AFTERWARD

I joined the Computing and Mathematical Sciences Department Ph.D. Program at California Institute of Technology in the fall of 2014. Soon after the preliminary examination held in the spring of 2014, I began working with my advisor Adam Wierman.

The first piece of my thesis (Chapter 3) started from the earlier half of my second year. Together with Adam and his postdoctoral scholar Hu Fu, we analyzed the efficiency of open access platforms under a networked Cournot competition model (J. Z. Pang et al., 2017). We also showed that controlling allocations can disincentivize production. This work appeared at the International Conference on Computer Communications (INFOCOM) in 2017, where I gave my first conference talk from my Ph.D. career. This paper served as the starting point for our work on online platform designs. It was during one of the breaks between sessions at the Simons Institute for the Theory of Computing that we explained this work to Weixuan (Sam) Lin and his advisor, Eilyan Bitar, who was previously a postdoctoral scholar at Caltech, and then, an Assistant Professor at Cornell. Together, we found an important result in Lemma 3.8, which was pivotal to the proof of many new results in Chapter 3 (W. Lin et al., 2017), which also previously appeared in the Conference for Decision and Control that same year. Finally, work on search cost started through Adam's visiting SURF student Jack Kleeman from University of Cambridge. A culmination of this work was recently submitted to Operations Research.

I started getting involved in the second piece of my thesis (Chapter 4) spring of 2018, building on work that was done previously in the group (Yunjian Xu et al., 2017). The connection between controlling allocation platforms and networked Stackelberg games considered in was stark, and new insights to the Stackelberg model followed naturally. Anticipation led to manipulation, and conditions for efficiency under anticipation were "almost necessary". Instead, inspired by a demand-side load management mechanism in electricity markets, we develop a demand response model, and show that demand management can lead to efficient markets, even with network constraints binding. Discussions online and at the INFORMS general meeting with Subhonmesh Bose helped push this work forward. This work moved forward with much advise from Laura Doval, and concluded with a research visit to the University of Illinois, Urbana-Champaign a year later in April 2019. Together with previous results, the work was recently submitted to Management Science.

The final part of my thesis (Chapter 5) started in the fall of 2017, when I had wonderful discussions with a visiting student Hanling Yi from the Chinese University of Hong Kong (CUHK). I was invited the following spring to CUHK by his advisor Minghua Chen, and together with another student Qiulin Lin, we made big progress in development of competitive online algorithms for generalizations of the classic one-way trading problem. Much of the work took place within two weeks in CUHK, but Qiulin (with minimal further guidance) tidied up our photographed whiteboard proofs and ironed out details. This work will be presented in Conference of the ACM Special Interest Group for Computer Systems Performance Evaluation (Sigmetrics) in 2019. The research visit to Hong Kong was exceptionally memorable as I also spent three days after the trip in Singapore, where I officially married my wife, Odelia.

In preparation of this thesis, I made the difficult choice to focus on my work relating to online platform design. As a result, a significant part of work done over the course of my graduate studies was not reported as part of this thesis prior to this. The following serves as a summary similar to an extended abstract of each work and will not include much technical details or results. More importantly, it serves to show my gratitude for the opportunities in collaboration, and the beautiful friendships made over these years.

The first project outside the online platform design regime is on load-side frequency regulation (J. Z. Pang, L. Guo, and Low, 2017). This work bore much fruits in multiple extensions and generalizations, e.g., (Wang et al., 2018; You, J. Z. Pang, and Yeung, 2018b; You, J. Z. Pang, and Yeung, 2018a). Another project covers the study of scheduling for battery swapping, where we derive solutions for both the offline (You, J. Z. Pang, M. Chen, et al., 2017) and online (You, Cheng, et al., 2018) setting by first reformulating the problem as an bipartite graph optimization problem. The last project looked at the joint optimization of placement and routing of network function chains in networks (L. Guo, J. Pang, and Walid, 2016; L. Guo, J. Pang, and Walid, 2018).

Besides these research projects, I was also glad to participate in two summer internships at the Software, Technology and Innovation Center of Schlumberger Limited.

ABOUT THE AUTHOR

John Pang received a BS with first class honors in Mathematical Sciences with a focus on Pure Mathematics from Nanyang Technological University in 2013. He was a Research Engineer at the Institute of High Performance Computing under the Agency for Science, Technology and Research (A*STAR) of Singapore until 2014. He then continued to the Computing and Mathematical Sciences (CMS) Department in California Institute of Technology and received his Ph.D. in 2019. He served as a visiting researcher to the Chinese University of Hong Kong in May 2018 and the University of Illinois, Urbana-Champaign in April 2019. He also served as the inaugural CMS Teaching Fellow from 2018 to 2019. He is a recipient of the A*STAR Undergraduate Scholarship and Singapore's National Science Scholarship.

His research lies in the intersection of applied mathematics, computer science and operations research, with applications in economics and power systems. His focus is on understanding how to design online platforms under practical trade-offs and considerations.