

University of Nevada, Reno

Decentralized Auction Solutions for Dynamic, Networked Markets

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in Computer Science and Engineering

by

Jordana Katherine Blazek

Dr. Frederick C. Harris Jr., Advisor

December, 2025

© by Jordana Katherine Blazek
All Rights Reserved



THE GRADUATE SCHOOL

We recommend that the dissertation
prepared under our supervision by

Jordana Katherine Blazek

Entitled

Decentralized Auction Solutions for Dynamic, Networked Markets

be accepted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

Dr. Frederick C. Harris, Jr.,
Advisor

Dr. Christopher Herald
Committee Member

Dr. David Feil-Seifer
Committee Member

Dr. Ankita Shukla
Committee Member

Dr. Eric Olson
Graduate School Representative

Markus Kemmelmeier, Ph.D., Dean
Graduate School

December 2025

Abstract

This work is a focused contribution to decentralized auction mechanisms in markets driven by autonomous, dynamically interacting nodes. These markets are characterized by strategic interactions, limited information, and fluctuating competition. Traditional auction models, such as the Second Price Auction, encourage truthful bidding but often assume centralized control, which is impractical in decentralized settings. To address this gap, we propose some extension to the Progressive Second-Price (PSP) auction, a mechanism that allows nodes to iteratively adjust bids based on local information from neighboring nodes.

This work develops a dynamic framework for decision-making processes in decentralized environments, utilizing a game-theoretic approach. By process, we define a subset of right-continuous, left-limited (*cadlag*) valuation functions used in order to model deterministic events.

Our framework represents strategies as a finite set of feasible actions, formalized through buyer–seller interactions on the bipartite graph representing participation, or the set of active bids, capturing the interdependencies between players within the network. We introduce a set of mixed strategies defined by probability distributions over these feasible actions, allowing for the modeling of intelligent decision-making within dynamic, competitive and alternatively, cooperative environments.

Key contributions include the development of influence sets to capture direct and indirect network effects on bidding behavior, an opt-out mechanism for strategic exit based on utility gain, derivations and proofs realizing and validating our extensions of the original theorems. Finally, we design and implement simulations to support our claims, allowing for further investigation of the solution space.

Dedication

I dedicate this thesis to my advisor, Dr. Frederick C. Harris, Jr., for his expertise in guiding the progression of an endogenous work. There is no other advisor in the world who could have led me to complete my research.

To my parents, my son.

And to Dr. Olson, for inspiring me with his intelligence and aptitude.

Acknowledgments

This dissertation represents the culmination of a long journey that began when I joined the University of Nevada, Reno in 2009, when Dr. Kevin Manley wrote to Dr. Thomas Quint to secure my place in his Calculus III summer course. I was working at a call center and dealing black jack at the Cal Neva Casino.

The University decided to accept my unusual enrollment, by unanimous vote in committee, and I later had the honor of contributing to both the Mathematics and Computer Science departments. During those years, I had the privilege of collaborating with some of the kindest and most intelligent people I have ever met—lab partners, mentors, and colleagues whose curiosity and generosity shaped how I think about research and discovery. I studied under professors whose fields have since evolved or disappeared, yet their influence continues to shape my understanding of mathematical rigor and creative inquiry.

I would like to thank my committee for their invaluable guidance throughout my graduate studies, and for their time generously given to complete this work.

I would like to thank Dr. Yakov Varol; I am so grateful for your invitation to continue my studies, and to Dr. Kevin Manley, you remain a source of confidence and inspiration.

To Dr. Frederick Harris and Dr. Eric Olson, thank you for your support, patience and kindness. You taught me so much.

Contents

Abstract	i
Dedication	ii
Acknowledgments	iii
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Motivation	1
1.1.1 Overview of Market Dynamics and Distributed Auctions	1
1.1.2 Limitations of Current Auction Mechanisms	2
1.2 Objectives and Outline of the Dissertation	2
2 Background and Review of the Literature	6
2.1 Introduction and Definitions	6
2.1.1 Transform from Message-Passing Network Constraints to Optimal Solutions	6
2.1.2 Criteria for Optimal Message Transmission	7
2.1.3 Other Network Optimization Problems	8
2.1.4 Understanding Network Equilibrium	10
2.2 Auction Theory Foundations	13
2.2.1 Auction Mechanisms and Their Evolution	13
2.2.2 Introduction to Auction Theory	13
2.3 Auction Mechanisms in Decentralized and Networked Markets	14
2.3.1 Challenges in Decentralized Markets	14
2.3.2 Auction Design in Networked Markets	14
2.3.3 Introduction to the Progressive Second-Price (PSP) Mechanism	15
2.4 Markets and Pricing in Network Auctions	17
2.4.1 Market Clearing and Equilibrium Pricing	17
2.4.2 Dynamic Pricing, Queueing, and Resource Allocation	18
2.5 Representations and Models	19
2.5.1 Time Series and Event Ordering	19

2.5.2	Dynamic Models and Difference Equations	19
2.6	Game Theory and Strategic Decision-Making in Auctions	21
2.6.1	Game-Theoretic Properties Associated with Network Constraints	21
2.6.2	Nash Equilibrium in Auction Models	23
2.6.3	Mean-Field Theory and Potential Games	24
2.6.4	Application of Mean-Field Theory and Potential Games in Market Models	24
2.6.5	Fixed-Point Theory and Applications in Network Optimization	25
2.6.6	Influence of Network Structure on Nash Stability	26
2.6.7	From Mechanisms to Matrix Representations	27
2.7	Graph Theory and Optimization in Auction and Network Models . .	27
2.7.1	Fundamentals of Graph Theory in Networked Systems	28
2.7.2	Properties and Metrics for Stability in Auction Networks	28
2.7.3	Centrality Measures and Their Role in Auction Stability	29
2.7.4	Applications of Graph Theory in Auction Optimization	30
2.8	Sphere of Influence Graphs (SIG) and Advanced Influence Metrics .	31
2.8.1	Sphere of Influence Graphs in PSP Auctions	32
2.8.2	Applications of SIGs in Auction Models	33
2.8.3	Influence Graphs and Equilibrium Convergence	34
2.9	Decision Theory and Network Models	35
2.9.1	Foundations of Decision Theory	35
2.9.2	Probabilistic Decision Models in Ecology	35
2.9.3	Linking Decision Theory and Ecological Models	36
2.9.4	Ecology and Game Theory Connections	36
2.9.5	Applications of Decision Theory in Ecological and Networked Systems	37
2.10	Conclusion	38
3	Challenges and Motivations	39
3.1	Motivation	39
3.2	Challenges	40
3.3	Contributions	42
3.4	Road Map	43
4	Strategic Bidding and Opt-Out Mechanism in the PSP Auction	46
4.1	Introduction	47
4.2	Related Work	49
4.3	The Market Mechanism	50
4.3.1	Market incentive.	52
4.4	User Strategy	54
4.4.1	User valuation (strategic incentive).	54
4.4.2	User behavior.	57
4.4.3	A simple example.	59

4.5	Conclusion and Future Work	64
5	An Equilibrium Analysis of a Secondary Mobile Data-Share Market	66
5.1	Introduction	67
5.2	Related Work	69
5.3	Market Formulation and Definitions	72
5.3.1	The Market Mechanism	72
5.3.2	Market Incentive	73
5.3.3	The Anonymity Problem	75
5.4	Strategic Framework	78
5.4.1	User Valuation (Strategic Incentive).	78
5.4.2	User Behavior.	83
5.4.3	A Simple Example.	89
5.5	Equilibrium Analysis	96
5.6	Conclusion and Future Work	110
6	Bipartiteness in Progressive Second-Price Multi-Auction Networks with Perfect Substitute	112
6.1	Introduction	113
6.2	Background and Related Work	114
6.3	The PSP Auction Mechanism	116
6.3.1	Bounded Participation	117
6.3.2	Residual Quantity and Allocation	118
6.3.3	Exclusion–Compensation	120
6.3.4	Valuation and Utility	121
6.4	Influence Sets	122
6.4.1	Primary (Direct) Influence Sets	123
6.4.2	Expanded (Indirect) Influence Sets	124
6.4.3	Projection Domains and Influence Operations	126
6.4.4	Projection-Based Influence Propagation	126
6.4.5	Partial Ordering and Market Shifts	128
6.5	Influence Shells and Local Saturation	129
6.5.1	Ordering and Influence Propagation	131
6.5.2	Asynchronous Sellers and Coupled Buyers	134
6.6	Simulation Framework and Implementation	137
6.6.1	Event-Driven Algorithm and Asynchronous Updates	138
6.6.2	Price Ladder Verification	141
6.6.3	Connectivity	143
6.7	Conclusion and Future Work	149
	Appendix: Market Shift Revealed by Partial Ordering	151
7	The Effects of Latency in Progressive Second-Price Auctions	154
7.1	Introduction	154

7.2	Preliminaries	156
7.3	Zero-Revenue Equilibria	160
7.4	Latency and Asynchronicity	164
7.5	Different Latencies	168
7.6	Conclusions	169
8	An Analysis of a Progressive Second-Price Multi-Auction Market with Perfect Substitute and Perfect Information	171
8.1	Introduction	171
8.2	PSP Auction Rules	175
8.2.1	Bid Structure and Strategy Space	175
8.2.2	Residual Quantity and Allocation	177
8.3	Exclusion–Compensation	178
8.3.1	(Local) Externality Cost.	179
8.3.2	(Global) Externality Cost	181
8.3.3	Valuation and Utility	183
8.4	Buyer Best Response	184
8.5	Simulations	196
8.6	Conclusion and Future Work	200
	Appendix: Proving Price and Cost Equality	203
9	Conclusions and Future Work	207
9.1	Conclusions	207
9.2	Future Work	209
	References	212
	Appendices	221
	A List of Publications by Jordana Katherine Blazek	221

List of Tables

6.1	Basic sets and notation for a bundle of J independent PSP auctions	117
6.2	Buyer regimes and their economic interpretation.	147
6.3	Interpretation of ridges in the buyer's utility surface.	147
7.1	The effects of reserve price on the bid price, total value, utility and revenue in the ϵ -Nash equilibria obtained from Algorithm 7.2 averaged over 100 different random initial bids. Except for the revenue corresponding to a zero reserve price, the standard deviations—not shown—were less than 1 percent of the averages.	163
8.1	Basic sets and notation for a bundle of J independent PSP auctions	176

List of Figures

2.1	Spheres of Influence for a Set of Edges and Vertices. Modified from Toussaint's original figure [81].	31
5.1	Product / Quotient (step) Space	78
6.1	3D matrix view: rows (buyers), columns (sellers), and z encodes price tiers. The colored surface shows the buyer price; filled markers are active bids; open circles show marginal winners. The right panel shows a transition where a new high-tier participant appears at $j=1$, a demand shortfall removes a low cell, and a reconfiguration shifts activity at $j=2$	130
6.2	Adjacency structure showing market connectivity between buyers and sellers.	142
6.3	Buyer 0 valuation curve and marginal diagnostics.	143
6.4	Adjacency and market connectivity for the 8×2 experiment. Connectivity is set at 50%.	144
6.5	Single buyer-seller utility surface for buyer 6 at seller 0. The surface plots $u_i(z_i^j, w_i^j) = \theta_i(z_i^j) - z_i^j w_i^j$ over quantity z_i^j and unit price w_i^j , holding the opposing bids fixed at the snapshot.	145
6.6	Buyer-level diagnostics under the Progressive Second Price (PSP) joint best response. Each panel shows the valuation $\theta_i(z)$ with the realized point $(Z_i, \theta_i(Z_i))$ and the marginal curve $\theta'_i(z)$ with a dashed line at p^* , illustrating the transition from constraint-limited to price-limited behavior along improvement paths.	146
6.7	Shared-seller utility surface where buyer 6's utility is a function of total requested quantity $Z_i = z_0 + z_1$; $u_i(z_0, z_1) = \theta_i(Z_i) - w(z_0, z_1)(Z_i)$ and $w = \theta'(Z_i)$: the feasible participation surface.	148
7.1	Comparison of the probability density functions governing the time between evaluation of bids and the communication latency to transmit a bid to the auction.	165
7.2	Left shows that changing the scale λ_c of the communication latency has minimal effect on the ensemble-averaged price and total utility received by all buyers. Right shows the average value and cost for each individual buyer in the case $\lambda_c = 1$	167

7.3	The outcomes for lazy buyers who evaluate the market 17 times less frequently and experience 17 times the latency in their bid messages compared to an equal number of industrious buyers with identical valuations.	169
8.1	Average bid price versus the percentage of buyers participating in multiple auctions. The resource in each auction is given by $Q^1 = 1000$ and $Q^2 = 2000$. The reserve price is $P^j = 6$ in both actions.	197
8.2	A buyer-seller graph for a market with four auctions, 400 buyers and a 6-edge distance between the first and last auctions. Here p is the percentage of buyers allowed the bid in the neighboring auctions. Shown is $p = 20$ where $ \mathcal{I}_1 = 120$, $ \mathcal{I}_2 = 140$, $ \mathcal{I}_3 = 140$ and $ \mathcal{I}_4 = 120$. Note that when $p = 300$ all bidders bid in all auctions.	198
8.3	Average bid price at an ϵ -Nash equilibrium for a market consisting of four auctions versus the percentage of buyers who can bid in the neighboring auction—see Figure 8.2. The shadow illustrates the standard deviation. The resource in each auction is given by $Q^1 = 500$, $Q^2 = 2000$, $Q^3 = 500$ and $Q^4 = 2000$. The reserve price is $P^j = 10$ in all actions.	199
8.4	The same buyers as in Figure 8.3 except with $Q^1 = 500$, $Q^2 = 2500$, $Q^3 = 500$ and $Q^4 = 1500$. The reserve price is $P^j = 10$ in all actions.	199

*In the dense undergrowth of a forest, the true slime mold *Physarum polycephalum* searches for food. Despite its simple structure, this organism can solve complex network optimization problems without centralized control. As the slime mold extends its tubular network (the ectoplasm), it assesses multiple pathways, balancing efficiency and risk to ensure optimal resource allocation [61].*

Chapter 1

Introduction

1.1 Motivation

This introductory chapter establishes the foundational motivation and mechanism design principles for the Progressive Second-Price (PSP) auction framework. It connects the early theoretical groundwork to the later developments in network modeling, latency analysis, and dynamic PSP systems. We outline the scope of the dissertation and demonstrate continuity in analysis from static equilibrium theory to fully dynamic decentralized markets.

1.1.1 Overview of Market Dynamics and Distributed Auctions

Decentralized systems such as mobile data sharing markets, bandwidth auctions, and distributed digital markets present unique challenges in terms of resource allocation. Nodes dynamically enter and exit these markets, creating continuously changing conditions. Decentralized, agent-based market models often struggle in such dynamic environments because they assume static participation and global information access.

In these distributed markets, auctions are a preferred mechanism for efficiently allocating resources. Auctions assign resources to nodes who value them most, optimizing overall utility. The Progressive Second-Price (PSP) auction is a strong candidate for these dynamic environments due to its iterative bidding and pricing rules.

1.1.2 Limitations of Current Auction Mechanisms

Traditional auction models, such as first-price, second-price, and Vickrey-Clarke-Groves (VCG) auctions, work well in static environments. These models ensure truthful bidding and efficient resource allocation. However, in decentralized and dynamic markets, these models face significant limitations:

- **Lack of Iterative Feedback:** Traditional auctions occur in a single round or in isolated instances, whereas decentralized markets require continuous feedback loops for nodes to adapt their strategies.
- **Fixed Participation:** Classical models assume all nodes are present throughout the auction process, which does not reflect real-world decentralized environments where nodes may join or leave the market.
- **Real-Time Adjustments:** Traditional models struggle to account for real-time changes in market conditions, leading to inefficient resource allocation in fast-moving markets.

These limitations necessitate an auction mechanism that not only accommodates dynamic participation but also enables strategic decision-making based on local, evolving information.

1.2 Objectives and Outline of the Dissertation

The main objective of this dissertation is to enhance the Progressive Second-Price (PSP) auction mechanism by introducing a new, innovative framework, thus addressing the limitations of traditional auction models in dynamic, decentralized markets. These innovations aim to improve efficiency, stability, and strategic depth in auction environments where autonomous nodes make rational decisions using information from the dynamically changing environment, i.e. environments where autonomous nodes make decisions based on limited, time-sensitive information.

We demonstrate the ability to incorporate complex mathematical tools to extend the capabilities of auction models. The enhanced PSP auction will be applied to various contexts, such as bandwidth allocation, decentralized finance, and peer-to-peer data markets, while incorporating game theory, network economics, and iterative feedback mechanisms.

Chapter 2 provides a comprehensive overview and establish theoretical foundations that align with the subsequent [12], [10], and [11] results. The foundational work on PSP mechanisms gives the structure and support for our transition to multi-auction and latency analyses. We provide a comprehensive overview of the theories foundational to this research, beginning with core concepts in network theory, including network structures, message transmission, and utility optimization. These basics set the stage for a deeper exploration of auction theory, from traditional auction formats (English, Dutch, and second-price auctions) to advanced mechanisms like Vickrey-Clarke-Groves (VCG) auctions, emphasizing their role in promoting incentive compatibility and addressing computational challenges.

We introduce our focus: progressive, iterative auction models. A detailed review of the Progressive Second-Price (PSP) auction by Lazar and Semret, first presented in 1999, is given, particularly as it pertains to the real-time, decentralized contexts explored in this dissertation. By examining market dynamics and pricing strategies—such as equilibrium pricing and demand-based pricing models. The chapter illustrates how PSP auctions manage adaptive bidding and resource allocation in changing environments.

Additional sections delve into various network representations, including time series, graph theory, and ecological models, which provide tools for analyzing dynamic interactions within auction settings. Applications of mean-field theory and potential games allow for stability analysis in large systems. Game theory concepts, particularly Nash equilibrium, potential games, and fixed-point theorems, are discussed as frameworks for understanding auction stability and strategic behavior among participants. Graph theoretical tools and sphere of influence (SOI) graphs are introduced

as methods for analyzing network stability and interdependencies. Finally, decision theory in ecological models inspire our modeling of adaptive behavior in networked markets, offering a perspective on how participants optimize resource allocation under uncertainty.

Chapter 3 will outline the motivation for developing adaptive auction mechanisms in decentralized markets and the main challenges faced.

Chapter 4 will discuss the opt-out function and the persistence of the existing Nash equilibrium in the context of a real-world scenario, the Hong Kong Mobile Data Exchange Market. We will test the PSP mechanism in this networked environment. We will analyze the outcomes of the simulations, focusing on metrics such as convergence relative to network connectivity and the role of network topology.

Chapter 5 will refine and define the Nash Equilibrium from Chapter 4, expanding the ITNG (2021) [12] analytical framework that modeled fully connected, noncooperative markets. This chapter revisits the secondary data exchange market example, illustrating how elasticity of supply and demand shapes the PSP mechanism and supports equilibrium formation. It provides a bridge from the foundational single-auction environment to the multi-auction, graph-theoretic interpretation explored in later chapters.

Chapter 6 transitions from single-market equilibrium models to network-based analysis. It introduces a bipartite structure that connects buyers and sellers through multiple overlapping auctions and presents the projection-based influence framework on the active bid set. This section defines the graph-theoretic concept of partial orders, saturation over influence sets that emerge naturally from the PSP formulation, linking the earlier equilibrium results to partial orders on bid prices that govern allocation, market shifts, and demonstrates the emergence of saturated one-edge shells (BFS-sets of bipartite connectivity).

Chapter 7 examines how message delays and random initial bids influence the formation of ϵ -Nash equilibria under truthful ϵ -best replies. It introduces asynchronous bid updates and bounded delay mechanisms that explain how latency and initializa-

tion noise shape convergence and predictability in decentralized PSP markets. This chapter refines the understanding of temporal asynchronicity.

Chapter 8 extends the PSP mechanism into a fully dynamic multi-auction system. It introduces bounded participation across concurrent auctions and generalizes the exclusion-compensation principle to a global setting, revealing how aggregate externalities can be expressed as price-ordered compositions of local market staircases.

Chapter 9 will present our conclusions, highlighting key findings on market stability, influence sets, and dynamic participation, and demonstrate our test case evaluations. We determine the potential applications of the PSP model in wireless and vehicular networks, mobile data and other consumable resource exchanges, and the associated decentralized financial markets and distributed systems, emphasizing the interdisciplinary potential of the enhanced PSP mechanism, and highlight the implications of our results for efficient resource allocation in decentralized and dynamic markets.

Finally, we outline future research directions to extend this work, increasing the resilience of auction models by using adaptive techniques responsive to noise and imperfect environments.

Chapter 2

Background and Review of the Literature

This chapter provides the necessary background to understand the theories and models that inform the research in this dissertation. We explore foundational auction theory, game theory, network economics, and various modeling approaches essential for understanding decentralized auction mechanisms.

2.1 Introduction and Definitions

2.1.1 Transform from Message-Passing Network Constraints to Optimal Solutions

This subsection establishes the link between classical message-passing optimization and the Progressive Second-Price (PSP) auction framework. In PSP markets, iterative updates to bids, prices, and allocations correspond directly to variable updates in message-passing algorithms. Each buyer–seller interaction can be viewed as a step towards equilibrium, where local information exchange replaces global coordination. Convergence of the iterative updates represents the formation of network equilibrium under the PSP mechanism, showing how distributed optimization techniques naturally extend to decentralized market behavior.

To transform from message-passing network constraints to optimal solutions, we first identify the problem. Familiar network problems that have seen a considerable amount of research are capacity planning and scheduling; each problem type requires a

different type of optimal solution, and therefore optimization algorithms will perform a search in the associated space of possible solutions. We implement the algorithm by iteratively updating variables and ensuring updates respect network constraints. In this way, we solve for the constraints, and our transform reveals our optimal solution.

The game-theoretic implications of auction mechanisms allow for robustness, efficiency, and scalability in achieving optimal solutions. With the advancement of technology, network links are more intelligent, and able to act independently, reducing the need for centralized control and so considerable research has been towards the goal of designing a mechanism suitable for real-time applications in modern networked systems. Networks are growing in their availability, complexity and we are dealing with an astonishing decrease in the heterogeneity of link type and ability.

2.1.2 Criteria for Optimal Message Transmission

Optimal message transmission in networks is about ensuring that messages reach their destination in the most effective way. The criteria can vary depending on the network's goals and constraints.

Latency Minimization: Latency minimization in communication networks reflects responsiveness in decentralized PSP auctions, where faster bid propagation and information exchange improve market convergence and stability. Efficient communication supports convergence to equilibrium in PSP markets.

Latency minimization is concerned with reducing end-to-end delay; critical in applications where real-time responsiveness is essential, such as video conferencing, online gaming, and automated control systems. Latency minimization ensures that messages are transmitted as quickly as possible, which can be particularly challenging in decentralized networks.

Energy-Efficiency: In decentralized PSP networks, energy-efficiency can be interpreted as a cost-efficiency tradeoff. Nodes seek to conserve their computational and communication resources while continuing to participate strategically in the market.

In networks with limited power resources (e.g., wireless sensor networks), minimizing energy consumption during message transmission is vital. Energy-efficient routing protocols reduce the number of transmissions, use lower-power nodes, or adjust transmission power to conserve energy without compromising reliability.

Reliability Measures: Reliability focuses on ensuring that messages are delivered accurately, even in the presence of potential failures. Fault tolerance is particularly crucial in networks with high churn or intermittent connectivity (e.g., ad hoc or peer-to-peer networks), where optimal message transmission includes mechanisms for retransmission, redundancy, or error correction.

Load Balancing: Distributing traffic evenly across the network helps prevent congestion and ensures no single node or link is overloaded. This is particularly relevant in mesh and peer-to-peer networks where nodes may have different capacities and routes can be dynamically adjusted to balance the load.

Scalability and Performance: In large-scale networks, optimal message transmission should not only work well for a small number of nodes but should scale effectively as more nodes join. Scalability considerations ensure that transmission performance (like latency and throughput) remains consistent as the network grows.

Cost Minimization: For networks with associated costs (e.g., cellular networks where users pay for data usage), cost minimization is an important criterion. Optimal message transmission, in this case, aims to achieve desired performance levels (e.g., low latency or high reliability) while keeping costs as low as possible.

2.1.3 Other Network Optimization Problems

Network optimization encompasses a wide range of problems aimed at improving the performance and efficiency of communication systems. While optimal message transmission focuses on the real-time movement of data through a network, other

areas such as capacity planning and scheduling address broader infrastructure and resource management challenges. These complementary problems operate at different layers of the network, collectively ensuring efficient and reliable communication.

Capacity Planning for Network Resources: Capacity planning focuses on ensuring the network has sufficient resources (e.g., bandwidth, storage, processing power) to handle anticipated traffic volumes. While capacity planning is about provisioning resources in advance, optimal message transmission is more concerned with efficiently using the existing capacity at any given moment. Capacity planning deals with the broader infrastructure setup, whereas message transmission tackles real-time or near-real-time data flow within those constraints. Optimal message transmission depends on effective capacity planning, as having insufficient capacity can lead to network congestion and degraded performance. However, they are distinct because capacity planning does not involve the actual routing or transmission protocols, focusing instead on high-level resource allocation.

Scheduling Problems in Networks: Scheduling ensures that network resources (e.g., CPU time, bandwidth) are allocated to different tasks or data flows in a way that meets specific performance goals, such as fairness or priority handling. Scheduling determines the order and timing of data transmission but doesn't directly address the methods for transmitting messages across the network. Optimal message transmission focuses on the mechanics of data delivery, such as selecting the best routes or adjusting transmission parameters, rather than deciding which messages go first. Effective scheduling supports optimal message transmission by organizing data flows to reduce contention and improve access to shared resources. However, they operate at different layers: scheduling often works at the data-link or transport layer, while message transmission can involve network-layer routing and link-layer transmission protocols.

2.1.4 Understanding Network Equilibrium

In networked environments, equilibrium represents a state where nodes (or participants) achieve a stable balance between their objectives, constraints, and available resources. Network equilibrium emerges from the interaction of several key factors: utility, cost, phase, and valuation. Each node seeks to maximize its utility, which is often influenced by its valuation of network resources, while managing the costs associated with resource usage. Phases within network operations introduce temporal dynamics, adding complexity to decision-making processes and impacting how equilibrium is maintained over time. Together, these elements create a dynamic environment where nodes adjust strategies to reach a state of optimal resource allocation, balancing individual needs with network-wide efficiency.

Utility Measure: Utility represents the level of satisfaction or benefit that a node (or participant) gains from participation. It is a measure of how valuable a certain outcome or resource allocation is to an individual participant. Utility is central to determining equilibrium in networked environments. Each node's goal is typically to maximize its utility, which might involve achieving lower latency, greater bandwidth, or better reliability.

Cost of Network Resources: Cost refers to the resources that a node expends to achieve its desired utility. This could include monetary costs, energy consumption, processing power, or other resources needed for transmission and reception. Cost is often what constrains a node's actions within the network. In equilibrium, nodes seek a balance between maximizing utility and minimizing cost, leading to an efficient allocation of resources. High costs can discourage certain network behaviors, pushing nodes toward an equilibrium where only beneficial or resource-efficient activities occur. Lowering cost often requires sacrificing some utility.

Phase in Networked Auctions: Each phase can be explicitly associated with the auction iteration variable, such as time step t or iteration count n , in the PSP mechanism. Here, a *phase* represents a dynamic bid–response cycle in which buyers and sellers update their strategies, prices, and allocations. This connection between temporal phases and discrete iterations clarifies how equilibrium evolves through successive PSP updates, making each phase a measurable component of convergence behavior.

The balance between utility and cost shapes the stability and sustainability of network behavior over time. The market will experience shifts as nodes adjust their strategies, attempting to sustain equilibrium. We may use the term phase to refer to this dynamic behavior. In networked auctions or dynamic networks, phase can refer to different stages or time intervals within the network’s operation, often linked to bidding rounds, decision-making points, or transmission cycles. Phase is crucial for understanding how equilibrium evolves over time in dynamic networks. Each phase may represent a temporary equilibrium point or an adjustment period as nodes respond to new information, alter strategies, or update bids.

Equilibrium Shifts in Multi-Phase Environments: Equilibrium can shift as nodes transition from one phase to another, particularly in auctions where each phase might introduce new bids or changes in network topology. In multi-phase, or iterative systems, nodes face trade-offs related to timing and decision-making. For example, acting early in a phase might lead to favorable bids or access to resources, while waiting until a later phase could offer more information but at a higher cost or reduced utility. Phased interactions introduce temporal dynamics that can affect strategy and equilibrium stability, as nodes must consider not only their immediate utility and cost but also may construct potential strategies, and remember previous strategies for themselves, and other players.

An equilibrium is achieved when nodes’ bids, reflecting their valuations, balance resource supply and demand. Nodes must balance their valuation against their

willingness to incur costs, potentially adjusting their valuations based on resource availability, competition, or budget constraints.

Valuation of Networked Nodes: In PSP frameworks, valuation can be formalized through elastic demand functions θ_i and their derivatives θ'_i , which represent each node's marginal valuation of resources. These functions define how a node evaluates additional allocation or cost and directly informs its bidding behavior and strategy updates. By grounding valuation in θ_i and θ'_i , we connect the perceived worth of network resources to a quantitative measure that governs PSP dynamics, making valuation the core of each node's strategic decision-making.

Finally, valuation refers to the perceived worth or value that a node assigns to a specific resource or service within the network. In auctions, valuation typically reflects how much a node is willing to pay or trade for a resource, such as bandwidth, processing power, or data storage. Valuation is fundamental to the bidding process and directly influences equilibrium outcomes in auction-based networks. Each node's valuation affects its bid and strategy, as nodes aim to secure resources at prices that align with their valuations. High valuation for a scarce resource might lead to aggressive bidding, while low valuation might result in a node opting out of certain auctions or negotiations. Valuation can change dynamically.

Reliability Protocols: Reliability protocols in networked environments are often designed to tolerate faults probabilistically, balancing between high availability and resource constraints [75, 84]. Vosoughi et al. (2016) explore fault tolerance in large-scale networks, noting that protocol reliability can be augmented through redundancy and error-checking, though at the expense of additional resources. These considerations are particularly relevant to dynamic market environments, where network conditions and participant availability may change rapidly.

2.2 Auction Theory Foundations

2.2.1 Auction Mechanisms and Their Evolution

Auction mechanisms have evolved from traditional formats, such as English and Dutch auctions, to more complex models designed for specific strategic and economic outcomes. In an **English auction**, bids increase incrementally, with the highest bid winning, while a **Dutch auction** decreases the price until a bidder accepts, both offering straightforward bid placement and winner determination strategies. These types are widely used but are limited by their centralized and synchronous nature.

2.2.2 Introduction to Auction Theory

First-Price and Second-Price Auctions In a first-price auction, the highest bidder wins and pays their bid. In a second-price auction, the highest bidder also wins, but pays the second-highest bid. This crucial distinction impacts bidding strategies: while first-price auctions encourage strategic underbidding.

Second-price auctions, pioneered by Vickrey [83], allow the highest bidder to pay the second-highest bid, incentivizing truthful bidding, as paying the second-highest price minimizes the need for strategic manipulation. This auction type supports fair pricing but faces challenges in dynamic and decentralized environments where information and timing constraints can hinder optimal bidding [19, 51].

Vickrey-Clarke-Groves (VCG) Auctions VCG auctions generalize the second-price auction to multi-unit or multi-item settings, ensuring truthful bidding by maximizing social welfare. The incentive compatibility of VCG auctions arises from the design that aligns individual interests with group efficiency, making them optimal for resource allocation in complex environments.

The VCG auction extends second-price principles to multi-item auctions, optimizing social welfare by encouraging truthful bidding across participants [24, 32]. However, VCG mechanisms can be computationally intensive and vulnerable to col-

lusion, especially in large, decentralized networks. These limitations have led to the development of alternative mechanisms better suited to dynamic and distributed markets.

These challenges motivated the development of distributed mechanisms such as the Progressive Second-Price (PSP) auction [45, 70], which retains the incentive-compatible structure of Vickrey-type mechanisms while enabling iterative updates through local communication. PSP provides a bridge between static VCG principles and dynamic networked markets, forming the foundation for the decentralized model explored in this dissertation.

2.3 Auction Mechanisms in Decentralized and Networked Markets

2.3.1 Challenges in Decentralized Markets

In decentralized systems, such as mobile data sharing or bandwidth auctions, participants lack access to global information and must base decisions on local, often incomplete, data. This introduces challenges in designing auction mechanisms that can effectively aggregate dispersed information.

2.3.2 Auction Design in Networked Markets

Auction design in networked markets involves creating mechanisms that allocate resources efficiently while accounting for the decentralized and dynamic nature of these environments. Unlike traditional auctions, networked markets require mechanisms that adapt to changing supply and demand, manage interdependencies among participants, and ensure fairness and efficiency under constraints such as limited information and asynchronous decision-making. Progressive auction mechanisms have addressed these challenges by enabling iterative bidding and local adaptation. This section explores key principles of auction design, emphasizing the role of network structure, strategic interactions, and stability in creating robust and efficient market

systems.

2.3.3 Introduction to the Progressive Second-Price (PSP) Mechanism

Shivkumar Kalyanaraman Semret's doctoral thesis, "Market-Based Resource Allocation and Pricing for Dynamic Spectrum Access in Wireless Networks", published in 1999, was a pioneering contribution to the field of decentralized resource allocation. His work introduced the Progressive Second-Price (PSP) auction as a mechanism for achieving efficient and fair allocation of resources in dynamic, distributed networks. This thesis laid the foundation for integrating market-based principles into network optimization, emphasizing decentralized decision-making and adaptability, which have become crucial in modern auction and network models [73].

Progressive auctions allow participants to adjust their bids iteratively over multiple rounds. This iterative process is crucial in environments where nodes progressively learn about other bids, adjusting their strategies based on emerging information [45]. Lazar and Semret's Progressive Second-Price (PSP) auction extends the second-price auction to an iterative, decentralized framework, emerging as a distributed extension of the Vickrey and VCG principles that form the foundation of modern mechanism design. Designed for decentralized networks, the PSP auction allocates an infinitely divisible resource efficiently and fairly among multiple users by allowing each participant to submit bids independently at each network link [45, 51]. This structure is particularly suitable for networks with variations in demand and availability [73].

Valuation and Elastic Demand in the PSP Mechanism: Elastic demand refers to the idea that a user's valuation of a resource decreases as the quantity of the resource they receive increases. This property encourages users to adjust bids according to their diminishing returns, fostering stable bidding behavior in PSP auctions [1].

In the PSP auction, each player i submits a bid $s_i = (q_i, p_i)$, where, q_i is the

quantity of the resource the player desires; p_i is the unit price the player is willing to pay. This decentralized bidding approach allows each participant to react to changes in other bids dynamically, promoting a stable equilibrium through iterative adjustments [45].

Allocation of Resources in the PSP Mechanism: The allocation rule defines how resources are distributed based on bids. For a given bid profile $s = (s_1, s_2, \dots, s_I)$, the allocation $a_i(s)$ for player i is given by:

$$a_i(s) = q_i \wedge Q_i(p_i; s_{-i})$$

where $Q_i(p_i; s_{-i})$ represents the maximum available quantity at a bid price of p_i , considering the bids of other players [45].

Cost and Utility in the PSP Mechanism: The cost to player i is:

$$c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})]$$

and their utility is given by the quasi-linear function:

$$u_i(s) = \phi_i(a_i(s)) - c_i(s)$$

where $\phi_i(a_i(s))$ is the valuation function for the allocated resource.

Each expression introduced above has a direct interpretation: the cost function represents the social opportunity cost of a participant's actions, while the utility function is quasi-linear in allocation, illustrating how each bidder's valuation translates into measurable surplus.

Equilibrium of the PSP Mechanism: A Nash equilibrium in the PSP auction is a bid profile s^* such that no player can unilaterally improve their utility by changing their bid [73]. The PSP mechanism promotes incentive compatibility by aligning each player's strategy with their true valuation, ensuring truthful bidding and stability [1].

At equilibrium, the allocation maximizes the total user value: $\sum_{i=1}^I \phi_i(a_i(s^*))$, ensuring that resources are allocated efficiently to those who value them the most.

The PSP auction mechanism addresses the unique challenges of decentralized resource allocation by promoting truthful bidding, efficient allocation, and stability. Its design leverages elastic demand, allowing the allocation and price to adjust dynamically to reflect each participant's marginal valuation of resources. The PSP auction's iterative nature and incentive-compatible structure make it ideal for networks with variable demand and resource constraints, ensuring both individual utility maximization and overall efficiency [45, 51, 73].

These foundational PSP rules provide the basis for the extended mechanism developed in this dissertation. The introduction of influence sets, dynamic participation, and graph-based connectivity generalizes PSP to networked markets where interactions evolve iteratively over time, linking auction equilibrium to network equilibrium.

2.4 Markets and Pricing in Network Auctions

2.4.1 Market Clearing and Equilibrium Pricing

In network auctions, achieving market-clearing, where the quantity of resources supplied matches the quantity demanded, can be essential for reaching equilibrium. Progressive auction models, such as the PSP mechanism, enable these market-clearing prices by iteratively adjusting bids based on real-time demand, given by a numerical study by Maillé et al. [51]. By allowing prices to adapt dynamically, the iteration converges to an efficient allocation of resources even in unstable market situations. Similar principles apply in supermarket games, where equilibrium is reached as users strategically select resources to balance load and minimize costs in a decentralized environment [89].

These are demand-based pricing models, and support this equilibrium by using elasticity to adjust prices according to demand shifts. For example, as discussed by Morris and Semret, elasticity captures how a user's valuation changes with the

quantity of resource received, encouraging fairer distribution in decentralized networks where resource needs vary [59, 73].

Delenda, Maillé and Tuffin [26] build directly on Lazar and Semret’s 1999 PSP framework by addressing the one remaining free parameter in the model — the reserve price. They demonstrate that while PSP guarantees convergence, efficiency, and incentive compatibility, the seller’s minimum acceptable price crucially determines revenue and market clearing. Their key contribution is a concavity proof showing that the expected revenue is concave under mild assumptions on user demand. This means the reserve price can be optimized by simple numerical methods, allowing PSP markets to balance efficiency with revenue maximization.

2.4.2 Dynamic Pricing, Queueing, and Resource Allocation

In PSP markets, queueing and dynamic pricing behaviors map naturally onto iterative bid-response cycles. Each PSP iteration can be viewed as a queue update, where latency represents the delay in bid propagation and response. As buyers adjust bids and sellers update reserve prices, the system behaves like a queue reaching steady state—we see how congestion or waiting time corresponds to delayed convergence.

Dynamic pricing and queueing theory work hand-in-hand in network auctions, particularly for optimizing resource allocation in environments where demand fluctuates. Queueing models allow researchers to examine the effects of resource congestion, waiting times, and the availability of resources in real-time, which can directly impact auction outcomes. For instance, by modeling queue lengths and service times, queueing theory helps in determining how pricing adjustments can alleviate congestion, thus ensuring that resources are distributed efficiently in high-demand scenarios [40].

Simulation-based approaches to dynamic pricing allow researchers to model complex, decentralized interactions, as Morris demonstrates in the context of evolving market conditions [59]. These simulations are particularly relevant in multiuser cognitive radio networks, where game-theoretic models such as supermarket games have been used to optimize resource allocation and address the challenges of real-time

competition for limited bandwidth [1, 49, 89].

Competitive pricing models in network auctions also consider inefficiencies that can arise from Nash equilibria, as described by Niyato and Tasnadi [63, 80]. These models analyze how competitive dynamics influence price stability and resource allocation, highlighting that equilibrium prices may not always align with optimal outcomes. This misalignment led to research towards the design of pricing mechanisms that minimize inefficiencies in decentralized networks, and a large field of study that considers the transformation between decentralized network constraints and global optimization; for example, mean-field theory. In large-scale PSP systems, mean-field theory approximates the aggregate behavior of numerous agents as individual influence diminishes, allowing macro-level analysis of equilibrium and convergence dynamics.

2.5 Representations and Models

2.5.1 Time Series and Event Ordering

Time series models are crucial for understanding how information propagates across a network, which directly impacts the convergence of distributed algorithms. In systems where rigid intervals or discrete events create discontinuities, time series models enable precise tracking of these sequences, as shown in Lamport’s work on logical clocks [44]. Effective time series analysis can guide the design of algorithms for network auctions by accounting for delays in information propagation, ultimately influencing convergence behavior.

2.5.2 Dynamic Models and Difference Equations

Population Dynamics: Dynamic models provide a foundation for analyzing interactions within networked systems, where population dynamics and ecological models can serve as analogs for competitive market systems. In network auctions, models such as the Lotka-Volterra equations illustrate competition and resource allocation,

framing interactions between bidders as analogous to predator-prey relationships [15]. These dynamics capture how strategies within an auction adapt over time, where competition for resources mirrors ecological interdependencies, and understanding these patterns supports efficient market design.

Building on population dynamics, phase dynamics allow us to design algorithms around binary analogs—such as supply and demand or cost and utility. By modeling systems as converging to fixed points along 45° lines, phase dynamics highlight the underlying stability in competitive interactions, much like the equilibrium between opposing forces in markets [78]. This analogy extends to autonomous decision-making, where nodes strategically adjust behavior to reach optimal outcomes, effectively creating a balance in market-like environments.

Diffusion Models: The diffusion models from physics provide further insights into networked markets, particularly in how resources and information spread. Helbing’s diffusion models describe the movement of particles in physical systems, analogous to the flow of information or resources within a dynamic market [34]. Here, diffusion can represent the spread of bidding information or resource availability, with equilibrium acting as a mixed strategy solution in game-theoretic terms. As information diffuses through the network, strategies stabilize, leading to an equilibrium where resources are allocated based on collective demand and individual utility maximization.

Clustering Models: Clustering models based on simplex structures provide frameworks for understanding phase transformations and equilibrium states in networks [88]. These models are particularly applicable in auction settings with multiple interacting agents where phase shifts occur as strategies converge. Clustering allows us to observe equilibrium states that emerge as similar strategies group together, forming clusters of optimal decisions [88, 96]. These models are especially valuable in auctions with multiple bidders, where equilibrium emerges through phase shifts as strategies converge. The deterministic nature of simplex structures offers researchers a mecha-

nism to drive systems toward desired outcomes; clustering creates naturally optimal allocations.

In cognitive radio ad-hoc networks, flow-based power control enables efficient resource allocation by adjusting power levels according to demand [43]. This deterministic strategy mirrors the allocation rules in PSP auctions, and reflects game-theoretic principles where resource allocation adapts efficiently to current needs, providing a practical analog for designing network auctions that continuously optimize resource usage.

Each of these dynamic models provides a lens for analyzing various aspects of network auctions, from equilibrium pricing to resource allocation stability. Their applications extend beyond traditional networking to capture complex, evolving interactions within auction-based systems.

2.6 Game Theory and Strategic Decision-Making in Auctions

2.6.1 Game-Theoretic Properties Associated with Network Constraints

Game-theoretic properties associated with network constraints are essential in understanding individual strategies, and how solutions may be designed where the actions of autonomous nodes have an impact on overall network performance. In network equilibrium, the modeled constraints interact as nodes strive to maximize utility while managing costs within a time-phased structure. The equilibrium state is reached when each node optimally balances its utility and cost in a given phase, with minimal incentive to change its strategy unless conditions or phases shift. The interactions form our simple model real-world network dynamics, where nodes continuously adapt to achieve efficient and sustainable outcomes in the face of changing resources and competitive pressures.

Nash Equilibrium: The equilibrium state, defined by Nash [62], is known as a Nash Equilibrium, is particularly relevant for networks where each node optimizes its own performance (e.g., minimizing delay, cost, or maximizing throughput) in response to others' actions.

Pareto Efficiency: Another important game-theoretic concept is Pareto efficiency. In networked settings, Pareto efficiency often aligns with optimal resource allocation and balancing network traffic. It implies a configuration where performance improvements for one user don't degrade service for others

Efficiency in decentralized decision-making is a measure of strategic selfishness of individual nodes, and is sometimes characterized as the Price of Anarchy (PoA). For example, in a routing game, nodes may choose paths that minimize their individual travel times, leading to congestion. The PoA metric demonstrates the need for incentives in decentralized auctions, which can help align individual strategies with overall network performance.

Incentive-Compatibility (Truthfulness): In decentralized auction networks, mechanisms like the Progressive Second-Price (PSP) auction incentivize nodes to bid truthfully, ensuring that each node's behavior aligns with the system's goals without centralized enforcement. Thus, truthfulness, or incentive-compatibility is an important goal in networked games. A system is incentive-compatible if every participant maximizes their utility by truthfully reporting their preferences or constraints.

Truthfulness allows for nodes to exhibit best-response dynamics. In networks with auction mechanisms, best-response dynamics describe how each node adapts its strategy based on others' strategies. We strive for a socially optimal solution, and must improve our collective performance. Mechanisms without incentive-compatibility built-in often require coordination or incentive alignment to achieve similar goals.

As researchers, we are able to measure the aggregate utility of all nodes in the network. We desire to maximize the social welfare of all participants; our mechanism

is carefully designed, utilizing these game-theoretic properties which allow nodes to converge iteratively to an equilibrium by continuously adjusting to local information. In a decentralized setting, these best-response dynamics allow optimization to occur organically, with each node contributing to equilibrium stability as best it can given the information available. In games of partial information, trustfulness is key in determining dynamic equilibrium, and allows researchers to increase the network complexity based on the decision-making abilities of intelligent nodes, and the uncertainty inherent in all message-passing schemes.

Feasibility Constraints: Feasibility constraints (e.g., bandwidth limits, processing power) restrict nodes' possible actions. For auction mechanisms, these resource constraints are key in finding an equilibrium state. Additional constraints may be interpreted and added to the financial model; for example, a budget constraint contributing to the equilibrium stability. The number and complexity of the constraints defines the network problem, and so defines the solution state; the mechanism; the auction, or game, will converge to the equilibrium solution. These constraints will determine the existence of a solution, along with the nodes' individual strategies, determining how stable the solution will be.

Solution Stability: Stability is especially relevant in dynamic network environments where nodes may frequently join or leave the network. A stable equilibrium can handle such fluctuations without significant degradation in performance. Different network types will experience different difficulties in finding and maintaining equilibrium. We must be careful to determine the actions of our intelligent players, the strategic decisions made towards the goal of the mechanisms implementation. The solution of the truthful embedding of the network game is a solution to the game [73].

2.6.2 Nash Equilibrium in Auction Models

Nash equilibrium is a fundamental concept in game theory, representing a state where no player can benefit by unilaterally changing their strategy given the strategies of

others [62]. In auction models, Nash equilibrium provides a basis for stable bidding strategies that align with the overall market conditions. Xu's work on equilibrium dynamics in networked systems illustrates how participants adjust bids in response to neighboring nodes, showing how equilibrium states evolve dynamically in network auctions as participants seek optimal responses [89]. Best response dynamics are central to these interactions, as players iteratively adjust their bids based on observed changes in the network, ultimately converging to a stable equilibrium.

2.6.3 Mean-Field Theory and Potential Games

Introduction to Mean-Field Theory: Mean-field theory approximates interactions in large systems by focusing on a global measure of the behavior of the population, rather than on individual actions. In large auction markets or networked environments, mean-field theory allows complex systems to be simplified, as the behavior of each participant becomes part of an aggregate distribution [29]. This approach is particularly useful in large auctions, where the actions of single bidders have minimal impact on the collective outcome. In these scenarios, participants' strategies often converge to a mean-field equilibrium, representing the system's overall balance [64].

2.6.4 Application of Mean-Field Theory and Potential Games in Market Models

Applications of Mean-Field Theory in Auctions and Networked Systems: Mean-field theory provides a framework for analyzing interactions in large markets by approximating the aggregate effect of individual behaviors. Supermarket games apply this theory effectively in decentralized settings, modeling how individual user strategies impact system-wide resource availability and equilibrium [89]. In spectrum-sharing scenarios where many users compete for limited resources, Maharjan and Wang show that mean-field approximations reduce computational complexity by focusing on average behavior, making it feasible to model decentralized decision-making at scale [49, 85]. In network resource allocation, this approach allows researchers to

model complex systems by focusing on the average interaction effect rather than detailed individual strategies. Boualem and Nutz discuss how mean-field models are especially effective in large-scale decentralized settings, where they simplify the analysis of market-clearing prices and resource allocation [29, 64]. This theory applies well to dynamic networks, where decentralized systems benefit from the aggregate effects captured by mean-field models, providing a framework for predicting market behavior with less computational overhead.

Potential Games and Auction Stability: Potential games further contribute to stability in network auctions by aligning individual incentives with an overall potential function, as explained by Manshaei and Xu [52, 89]. This concept is highly applicable in cognitive radio and other resource-constrained environments, where mean-field theory combined with potential games helps to understand how collective behaviors influence network-wide resource availability and efficiency; efficient resource allocation is achieved in [85] for spectrum optimization, where consistent allocation results in better performance. In these games, players' incentives are structured to naturally lead towards convergence, with each strategy adjustment contributing to a stable state [52].

Potential games are particularly applicable to networked auction settings; Xu's work on spectrum allocation illustrates how potential games enhance stability in auctions, as each bidder's strategy aligns with collective efficiency goals [89].

2.6.5 Fixed-Point Theory and Applications in Network Optimization

Fixed-Point Theorems in Optimization: Fixed-point theorems like Brouwer's provide the mathematical basis for ensuring equilibrium in optimization problems, asserting that under specific conditions, a stable allocation will emerge [56, 71]. In network auctions, these theorems help confirm that equilibrium solutions are achievable, supporting the stability of flow and pricing configurations where participants'

strategies balance with resource constraints.

Methods for Fixed-Point Solutions: Fixed points in networked systems are often difficult to compute directly, but tools such as the Krawczyk operator offer an interval-based approach for locating fixed points within bounded regions [77]. Additionally, projection-like retractions on matrix manifolds, as outlined by Absil, facilitate efficient optimization for high-dimensional fixed-point problems, supporting stable configurations in network auctions [2].

Applications to Dynamical Systems and Stability: Iterative methods for finding fixed points play a critical role in analyzing stability in dynamical systems. Domes and Montanher's research on iterative methods for feedback systems is relevant here, as networked auctions often feature feedback loops where each participant's bid influences others' strategies [30, 58]. Smale's contributions to differentiable dynamical systems provide further insights into stability in competitive environments, revealing how equilibria evolve in complex auction settings [79].

2.6.6 Influence of Network Structure on Nash Stability

The structure of a network heavily influences the stability and robustness of Nash equilibria in auction settings. In highly connected networks, equilibrium can be reached more quickly as strategies and information spread efficiently. Methods such as contraction mapping and supermodular games contribute to stabilizing equilibrium in interconnected environments by leveraging the interdependence among participants [93]. For instance, Zhang's application of Stackelberg games shows how leader-follower dynamics can enhance stability in network optimization, while Maharjan's use of potential games and mean-field models highlights the balancing act between competition and cooperation necessary for achieving robust outcomes in decentralized auctions [49]. Highly connected networks often promote stability by ensuring that strategic deviations propagate quickly, realigning incentives across participants. Quint and Shubik's study illustrates that when vertices interact mainly with local

neighbors, stability relies on the density of local connections and the structure of these interactions within the network [72].

2.6.7 From Mechanisms to Matrix Representations

Analyzing stability within networked auctions may also involve representing these mechanisms within a matrix framework, where methods like the Krawczyk operator can provide insights into fixed-point stability [77]. While not central to auction design, stability analysis using operators helps assess equilibrium configurations, but a more tangible approach involves matrix representations that facilitate decentralized solutions.

Linear programming techniques, such as the Danzig-Wolfe method, solve for multi-hop routing problems by optimizing allocation based on matrix fitness ratios [46]. These matrix-based approaches translate auction mechanisms into a structured format, supporting decentralized solutions that allow each node to optimize local objectives within a global framework. By combining linear programming with matrix representations, decentralized systems can achieve optimal allocations, effectively bridging the gap between individual strategies and network-wide equilibrium.

2.7 Graph Theory and Optimization in Auction and Network Models

Graph theory offers powerful tools for modeling and optimizing auction and networked systems. By representing participants as vertices and their interactions as edges, graph structures provide a foundation for analyzing how connectivity, stability, and disruptions impact resource distribution and equilibrium in auction environments. This section explores key graph-theoretical concepts and their relevance to decentralized auctions, particularly progressive second-price (PSP) markets.

2.7.1 Fundamentals of Graph Theory in Networked Systems

Graph-theoretic models capture the complexity of auction networks by representing interactions and dependencies through edges and vertices. These models help to visualize and analyze resource flows, dependencies, and stability in PSP markets.

2.7.1.1 Directed Acyclic Graphs (DAGs)

DAGs simplify scheduling and dependency management by ensuring that no cycles exist, enabling clear bid flow organization. Barret's work on DAGs illustrates their utility in traffic management and routing, which translates directly to PSP auctions for managing bid dependencies and ensuring efficient allocation across iterative bidding rounds [8]. By leveraging DAG structures, PSP markets can reduce conflicts, optimize scheduling, and maintain transparency in resource allocation.

Minimizing Crossing Numbers: Graph visualization and efficiency improve significantly when crossing numbers—the number of edge crossings in a graph layout—are minimized. DeVos explores methods for reducing crossing numbers, which is particularly relevant in auction environments where bid flows and resource allocations must be optimized for clarity and efficiency [27]. For progressive allocation mechanisms, minimizing crossing numbers facilitates faster convergence by reducing complexity in iterative processes.

2.7.2 Properties and Metrics for Stability in Auction Networks

Graph stability metrics assess the resilience and robustness of auction networks, influencing convergence rates and equilibrium outcomes. These metrics, including reachability, resistance distance, and centrality, are particularly important for modeling interactions in PSP auctions.

Reachability Constraints: Reachability defines which vertices can be accessed from a given starting point, shaping influence paths within auction networks. In PSP

auctions, reachability constraints determine influence sets, where participants impact one another's bids based on network connections. Algorithms like depth-first and breadth-first search systematically evaluate reachability, revealing strategic influence patterns and ensuring robust resource distribution in decentralized systems.

Resistance Distance: Resistance distance, adapted from electrical circuit theory, quantifies the ease of information or resource flow between vertices. Osvaldo demonstrates how resistance distance captures temporal interactions, modeling long-term stability by tracking resource diffusion across dynamic networks [4]. This measure directly applies to PSP auctions, where resource propagation drives equilibrium adjustments and influences bid dynamics. As resistance decreases, influence spreads more effectively, stabilizing competitive auction environments.

Centrality Metrics: Centrality measures, including degree, betweenness, and eigenvector centrality, identify influential vertices in a network. High-centrality vertices streamline resource flow, stabilize bidding behavior, and reduce congestion in auction systems. Borgatti and Everett show that centrality amplifies stability by enhancing connectivity and ensuring resilience in decentralized networks [17]. These measures are essential for designing PSP markets where strategically placed vertices can accelerate convergence and maintain equilibrium.

2.7.3 Centrality Measures and Their Role in Auction Stability

In auction networks, centrality metrics play a crucial role in stabilizing decentralized systems. High-centrality vertices facilitate information flow, align strategies, and ensure resource allocation efficiency.

Connectivity and Stability: Increased connectivity aids in achieving equilibrium by enhancing information flow and aligning participant strategies. Jackson and Wolinsky's work highlights how central vertices in decentralized systems stabilize auction

dynamics by reducing delays and enabling faster convergence to equilibrium [37]. These properties are especially valuable in PSP auctions, where stable bidding interactions depend on seamless communication across the network.

Strategic Influence of Central Vertices: Freeman’s foundational research on centrality demonstrates how strategically positioned vertices gain competitive advantages by influencing adjacent nodes [31]. In PSP auctions, high-centrality vertices stabilize bidding dynamics by coordinating strategies and facilitating resource allocation, ensuring robust equilibrium outcomes. These vertices act as stabilizing forces in competitive markets, mitigating volatility and promoting efficiency.

2.7.4 Applications of Graph Theory in Auction Optimization

Graph theory provides practical tools for optimizing auction networks, transforming theoretical concepts into actionable strategies that enhance stability, efficiency, and adaptability in PSP markets.

Spectral Properties for Stability and Convergence: Spectral graph theory offers insights into network robustness through the spectral gap of the Laplacian matrix—the difference between its largest and second-largest eigenvalues. A larger spectral gap indicates faster convergence to equilibrium and greater resilience to disruptions. Mohar and Chung’s research highlights how spectral gaps reduce instability, ensuring stable resource allocation in PSP auctions [57, 23]. These properties support the design of auction networks that adapt effectively to changing market conditions.

Resistance Distance in Dynamic Networks: Resistance distance models the diffusion of influence and resources in PSP auctions, capturing long-term stability in dynamic settings. Osvoldo’s research demonstrates how resistance-based metrics predict resource flow and influence propagation, shaping equilibrium outcomes in decentralized markets [4]. By understanding these dynamics, auction designers can optimize network interactions to sustain stability.

Message-Passing Algorithms and Optimization: Yedidia’s work on message-passing algorithms provides a framework for inference and optimization in decentralized systems, offering insights into autonomous decision-making in PSP auctions [90]. By leveraging message-passing strategies, auction networks can achieve efficient resource allocation and align individual strategies with global objectives.

2.8 Sphere of Influence Graphs (SIG) and Advanced Influence Metrics

Sphere of Influence Graphs (SIGs) provide a powerful framework for modeling local interactions in decentralized systems, including PSP auctions. By connecting vertices based on proximity within a metric space, SIGs capture localized dependencies and strategic influences, offering insights into how bidding behavior propagates and stabilizes in networked markets. This section explores the properties, applications, and strategic implications of SIGs, emphasizing their role in achieving equilibrium in PSP auctions.

Definition: For a set $X = \{X_1, X_2, \dots, X_n\}$ in a metric space M , each point X_i has an influence region (or sphere) determined by the closest point distance r_i . An edge between X_i and X_j exists if their spheres of influence intersect.

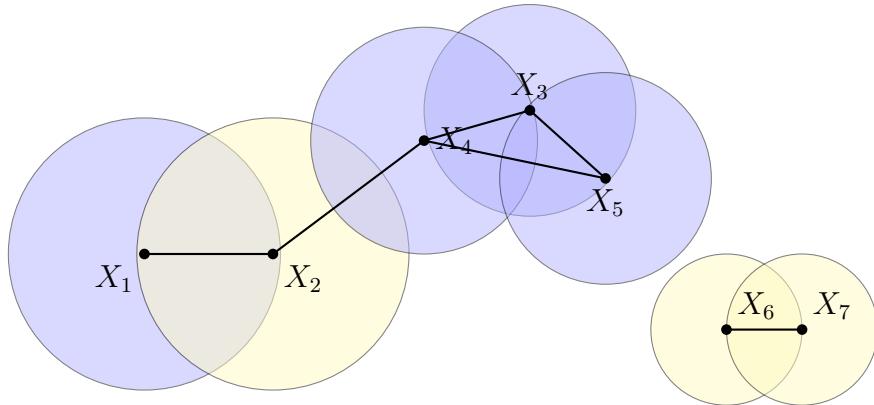


Figure 2.1: Spheres of Influence for a Set of Edges and Vertices. Modified from Toussaint’s original figure [81].

Real-World Example - Ecological Zones: Consider an ecosystem where various plant species grow in specific areas, each species representing a node. Each plant’s “sphere of influence” is the surrounding area where it competes for resources like sunlight, water, and soil nutrients. If two plants’ spheres overlap, an edge forms, indicating competition or interaction. SIGs model these relationships, showing which species affect each other, aiding in the study of competition, biodiversity, and ecosystem stability in fragmented habitats. This example highlights SIGs’ utility in environmental and biological research.

2.8.1 Sphere of Influence Graphs in PSP Auctions

SIGs model localized interactions where participants’ bids influence their neighbors, forming clusters of aligned strategies. Introduced by Michael and Quint for spatial data modeling, SIGs have been widely applied to networked systems to represent dynamic dependencies and influence patterns [55, 56]. In PSP auctions, SIGs provide a natural representation of influence zones, where local interactions ripple through the network, driving adjustments in bids and stabilizing market outcomes.

Metrics for Stability in SIGs: Key metrics for analyzing SIGs include *edge density* and *clique size*. High edge density indicates tightly connected regions where participants’ strategies are synchronized, promoting stability in bidding behavior. Barrett’s analysis shows that cliques within SIGs foster equilibrium by aligning bids within connected groups, reducing volatility and enhancing market efficiency [8]. These metrics are particularly relevant in PSP markets, where local strategic alignment drives convergence to stable equilibria.

Dynamic Properties and Micro-Level Modeling: Dynamic SIG properties enable the modeling of adaptive strategies at a micro level, where each participant’s behavior evolves based on local interactions. Barrett and Osvaldo highlight SIGs as tools for visualizing influence dynamics in PSP auctions, showing how bid dependencies evolve over time [8, 4]. This approach aligns with the iterative nature of

PSP mechanisms, where bid adjustments are guided by feedback from neighboring participants, creating a layered and adaptive market structure.

Influence Radius and Bid Dependencies: The *influence radius* of a vertex in a SIG determines the extent of its impact on neighboring vertices, capturing the localized effects of bidding behavior. Influence graphs based on SIG structures model these dependencies, showing how local interactions propagate adjustments across the network. Quint and Shubik’s work demonstrates how influence paths within SIGs drive equilibrium by cascading strategic shifts throughout the network [72].

2.8.2 Applications of SIGs in Auction Models

SIGs offer practical insights for designing and optimizing decentralized auctions by representing local interdependencies and strategic influences.

Weighted Influence Graphs: In PSP auctions, weighted influence graphs extend the SIG framework by assigning weights to edges based on the strength of influence. These weights capture bid intensity and resource demand, allowing for dynamic adjustments that align with changing market conditions. Lazar and Semret’s PSP mechanism leverages this adaptability to ensure efficient resource allocation in decentralized environments [45].

Centrifugal Number and Network Stability: The *centrifugal number* in SIGs measures the maximum number of non-overlapping spheres around a central vertex, providing insights into clustering behavior and stability. Quint’s research shows that high centrifugal numbers indicate robust clustering, reducing disruptive interactions and facilitating smoother convergence to equilibrium [55]. This property is particularly useful in PSP markets, where stable clustering minimizes volatility and supports iterative bid adjustments.

Dynamic Interactions in Multi-Phase Systems: In multi-phase PSP markets, SIGs model the progression of bidding dynamics across iterative rounds. Each phase represents a temporary equilibrium adjusted as bids evolve, with influence graphs capturing real-time dependencies and interactions. Quint and Shubik’s work on multi-phase systems highlights how SIGs stabilize PSP auctions by structuring interactions over multiple iterations, ensuring consistent convergence to equilibrium [72, 81].

2.8.3 Influence Graphs and Equilibrium Convergence

Influence graphs represent the interdependencies among bidders in PSP auctions, where vertices denote participants and edges indicate bid influence. These graphs provide a framework for analyzing strategic dependencies, capturing both local and global effects on equilibrium behavior.

Impact on Convergence Rates: The structure of influence graphs directly affects the rate of convergence to equilibrium. Denser influence graphs with larger spectral gaps promote faster convergence by ensuring that bid adjustments propagate efficiently across the network. Quint and Shubik’s analysis demonstrates how localized interactions within influence graphs accelerate equilibrium stabilization by realigning incentives across connected participants [72].

Resilience to Bid Disruptions: Influence graphs enhance resilience by modeling how strategic deviations are absorbed and corrected within the network. Weighted influence graphs adapt to disruptions by adjusting edge weights based on changing bid intensity and demand, creating a flexible structure that maintains equilibrium even under fluctuating conditions [45].

Applications to Resource Allocation: Influence graphs support efficient resource allocation by modeling dependencies and aligning bidding strategies with market conditions. Barrett’s work on SIGs illustrates how influence-based clustering improves resource distribution in PSP auctions, minimizing contention and ensuring

stable outcomes [8].

2.9 Decision Theory and Network Models

Decision theory offers a framework for understanding and modeling choices under uncertainty, making it essential for analyzing behavior in both ecological systems and networked environments. This section explores the foundational principles of decision theory, its connections to ecological dynamics, and applications to networked systems, highlighting how these models inform the stability and optimization of decentralized auctions.

2.9.1 Foundations of Decision Theory

Decision theory addresses the principles of making optimal choices in uncertain environments. The foundational work of Boole [16] established deductive reasoning as a basis for decision-making, introducing early ideas on uncertainty and logical deduction. Building on these principles, expected utility theory became a cornerstone of decision theory, providing a framework for evaluating choices based on their expected outcomes. Diecidue's rank-dependent utility model [28] extends this concept, offering a robust approach for incorporating risk preferences and addressing discrepancies in decision-making under uncertainty.

In dynamic environments, decision-making often occurs under time constraints or incomplete information. Ordonez's study on decision-making under time pressure [66] demonstrates how agents balance immediate actions against long-term optimization. These insights are directly applicable to auction systems and network models, where participants must adapt strategies rapidly in response to fluctuating conditions.

2.9.2 Probabilistic Decision Models in Ecology

Probabilistic decision models are widely used to predict ecological outcomes, where uncertainty arises from environmental variability and species interactions. Chiou's work on robust traffic control [22] demonstrates the utility of probabilistic models

in managing flows, a concept that translates to ecological networks for resource and population management. Bayesian methods, as outlined by Zhou [96], provide a framework for incorporating uncertainty into decision-making processes, enabling robust clustering in dynamic ecological settings. These methods help bridge the gap between theoretical models and practical applications in systems where uncertainty shapes outcomes.

2.9.3 Linking Decision Theory and Ecological Models

Ecological systems offer natural analogs for decision-theoretic processes, where species behavior can be modeled as a series of decisions influenced by environmental conditions and interspecies dynamics. Bomze's application of the Lotka-Volterra predator-prey model [15] demonstrates how decision-making can reflect competition and adaptation in ecological settings, framing species interactions as strategic processes. May's exploration of bifurcation phenomena [53] extends this framework, showing how small environmental changes can lead to significant shifts in system dynamics.

Adaptive decision-making under environmental uncertainty is another critical area where decision theory informs ecological models. Oki's biologically-inspired frameworks [65] and Meyer's work on noise-driven dynamics [54] illustrate how organisms and systems adapt to fluctuating conditions, offering insights into resilience and optimization in decentralized networks.

2.9.4 Ecology and Game Theory Connections

The interplay between ecology and game theory reveals deep parallels between population dynamics and strategic interactions. Replicator dynamics, as explored by Bomze [15], model ecological equilibria by treating population fractions as mixed strategies. These models capture the evolution of competitive interactions, with phase portraits identifying equilibria as sources, sinks, or other critical points.

Evolutionary game theory offers algorithmic tools for understanding these dynamics. Chastain's study on multiplicative weight updates [21] connects evolutionary

games to frameworks like the experts algorithm, balancing cumulative utility and entropy maximization. These methods have direct implications for auction systems, where strategies evolve over time, reflecting a balance between short-term gains and long-term equilibrium. The Karush-Kuhn-Tucker conditions provide a mathematical foundation for ensuring optimization in these games, emphasizing the importance of maximizing efficiency in dynamic environments.

2.9.5 Applications of Decision Theory in Ecological and Networked Systems

Decision theory informs a wide range of applications in ecological and networked systems, particularly in managing resources and optimizing outcomes in uncertain environments. Behavioral models guide population management, with Hansen's interval-based optimization techniques [33] and Liu's chaotic time series methods [47] providing tools for analyzing complex dynamics. These approaches are crucial for understanding how decentralized networks, such as ecological systems or auction markets, adapt to changing conditions.

In networked systems, decision theory integrates with resource allocation models to optimize equilibrium and stability. Shary's Krawczyk operator [77] offers a method for interval-based analysis, supporting robust optimization under uncertainty. Padilla's cognitive frameworks [67] extend this approach by incorporating decision-making under incomplete information, emphasizing resilience and adaptability. These techniques align with concepts from Lamport's study on event ordering [44], addressing the challenges of time-based decision-making in distributed networks.

Resource allocation in ecological systems often parallels strategies in auction markets, where network equilibrium reflects balanced resource distribution. Quint's work on balancedness in partitioning games [71] highlights how equilibrium concepts can inform resource-constrained networks, ensuring efficient allocation despite competing demands.

2.10 Conclusion

This chapter provides a comprehensive review of the theoretical and mathematical foundations of decentralized auction mechanisms, with a particular focus on the Progressive Second-Price (PSP) auction. Through an exploration of auction theory, game theory, and network economics, we lay the groundwork for understanding how decentralized auctions operate and how participants interact strategically within such systems.

Key concepts from auction theory, such as equilibrium pricing, incentive compatibility, and resource allocation, were examined in the context of dynamic and networked environments. Game-theoretic principles, including Nash equilibria, potential games, and mean-field theory, were highlighted as essential tools for analyzing strategic interactions and achieving stable outcomes. The intersection of these disciplines with network economics emphasizes the role of connectivity, influence, and competition in shaping market dynamics.

The chapter also introduces the foundational role of graph theory in modeling and optimizing PSP auctions. Graph structures, metrics, and algorithms—such as Sphere of Influence Graphs (SIGs), centrality measures, resistance distance, and reachability—provide critical insights into how information propagates across networks. These tools enable the design of robust, efficient auction mechanisms that adapt to changing market conditions.

In summary, this chapter has established the theoretical and methodological basis for studying PSP auctions in decentralized markets. By integrating insights from multiple disciplines, it has paved the way for the analysis and design of mechanisms that leverage local interactions, network connectivity, and strategic decision-making to achieve optimal resource allocation in complex, networked environments. This synthesis will guide the development of advanced models and algorithms in subsequent chapters, building on the foundations outlined here.

Chapter 3

Challenges and Motivations

3.1 Motivation

This research addresses the complexities of decentralized auctions and market mechanisms, particularly in environments where strategic interactions are influenced by network structure and information constraints. In modern markets, nodes often face uncertainty, incomplete information, and dynamic competition, all of which pose challenges to efficient resource allocation. Auctions, and specifically second-price auction mechanisms, provide a framework for truthful bidding. However, in decentralized settings, the lack of centralized control and real-time feedback necessitates mechanisms that adapt to evolving market conditions.

Our focus is on developing adaptive auction mechanisms, like the Progressive Second-Price (PSP) auction, that respond to market dynamics by allowing nodes to adjust their bids based on local information gathered from their network neighbors. This motivates the study of influence sets, dynamic participation, and the role of network effects in shaping bidding behavior. In these settings, nodes lack full market information, and are affected by network dependencies.

Picture a bustling urban environment during rush hour, with hundreds of autonomous vehicles navigating the city streets. Each vehicle is equipped with sensors and communication systems, forming a decentralized vehicular network. These vehicles must make decisions on route selection, speed adjustments, and lane changes to minimize travel time and fuel consumption, all while avoiding congestion and colli-

sions.

However, each vehicle only has local information—data from its immediate surroundings and neighboring vehicles—making it difficult to predict global traffic patterns. To optimize their routes, vehicles engage in a decentralized bidding process for access to high-priority lanes or faster routes. Each vehicle’s bid depends on its urgency, destination, and local traffic conditions.

The network faces challenges similar to those in auctions: as vehicles adjust their bids based on changing local conditions, traffic flow becomes unpredictable. Some vehicles drop out of the bidding for high-priority lanes when costs outweigh benefits, while others strategically bid higher to secure faster routes. Over time, the system self-organizes, with vehicles reaching an equilibrium where resources like road space and time are efficiently allocated.

Just as vehicles in the network balance individual utility with overall traffic optimization, nodes in decentralized markets leverage dynamic participation and network effects to stabilize and achieve optimal outcomes. The dynamic nature of decentralized systems requires auction mechanisms that can adapt to shifting market conditions and heterogeneous network structures. Our implementation of a real-world network faces significant obstacles; it is a game of partial information played in a web of interconnected decisions, dynamic participation, and evolving market constraints.

To address these complexities, we must identify and overcome key challenges that impact both the stability and efficiency of decentralized markets. The next section outlines these challenges, forming the basis for the innovative solutions proposed in this work.

3.2 Challenges

Designing auction mechanisms in decentralized environments poses significant challenges due to the complexity and variability of networked systems. These challenges arise from a combination of incomplete information, dynamic behaviors, and the interconnected nature of decision-making among participants. Addressing these issues

requires a comprehensive understanding of the following key factors:

Information Asymmetry: In decentralized markets, nodes lack access to global market information and must rely solely on local signals obtained from their immediate neighbors. This limited perspective increases the likelihood of suboptimal bidding strategies and strategic miscoordination, where decisions made based on partial information fail to align with the overall market dynamics. For instance, consider a bidder in a networked auction who overestimates demand due to limited local signals, leading to inflated bids. These inflated bids ripple across the network, prompting neighboring participants to adjust their own strategies, often exacerbating inefficiencies. Mitigating information asymmetry requires mechanisms that enable nodes to infer broader market trends from local interactions without centralized oversight.

Dynamic Participation: Participants in decentralized auctions, whether buyers or sellers, often enter or exit the market based on their perceptions of value and competition. This fluid participation introduces variations in bidding dynamics, disrupting stability and complicating the prediction of equilibrium outcomes. To address this, adaptive mechanisms must accommodate changing participation patterns and interconnected decisions while ensuring market stability and efficiency.

Network Effects: The interdependence of decisions among participants amplifies the complexity of decentralized markets. A node's bid impacts its direct neighbors and propagates through secondary connections, making prediction and coordination complex. To address this, adaptive mechanisms must accommodate changing participation patterns and interconnected decisions while ensuring market stability and efficiency. Effective auction design must account for both direct and indirect influences on bidding strategies.

Market Saturation: As decentralized markets mature, the incremental effects of additional interactions diminish, leading to saturation. In such conditions, price variability decreases, and influence sets converge, limiting the scope for strategic adjustments. Overcoming market saturation requires mechanisms that maintain dynamic bidding opportunities and adapt to the diminishing returns of further interac-

tions.

Bid Adjustment Mechanisms: In dynamic auctions, participants must continuously adjust their bids in response to local changes and new information. However, determining optimal bid adjustment strategies is inherently challenging. Effective strategies must balance responsiveness to immediate conditions with the need for long-term stability, particularly in environments where bids are shaped by both primary interactions (direct connections) and secondary interactions (indirect influences). Auction mechanisms must enable nodes to navigate these complexities while preserving market efficiency.

These challenges lay the groundwork for the development of our extended Progressive Second-Price (PSP) mechanism. By integrating adaptive features, our mechanism enables dynamic bid adjustments, fosters strategic flexibility, and enhances resilience against information asymmetry and network effects. In traffic systems, the concept of Wardrop's Principle describes an equilibrium state where no driver can improve their travel time by unilaterally changing routes [87]. Our proposed framework is designed to maintain market efficiency and stability, even under constraints evolving from dynamic node type and the resulting heterogeneous connections. By design, these challenges form our extended PSP mechanism, ensuring that decentralized markets remain robust, scalable, and capable of meeting the demands of modern networked environments.

3.3 Contributions

Building upon the challenges outlined in the previous section, this chapter transitions toward the design logic of the extended PSP mechanism. The challenges of information asymmetry, dynamic participation, and network interdependence directly inform the adaptive strategies and opt-out mechanisms introduced next, ensuring continuity between problem definition and methodological development.

This dissertation proposes an innovative solution to the challenges of overloaded networked systems, focusing on extending the Progressive Second-Price (PSP) mech-

anism to meet the demands of decentralized markets. By leveraging concepts from physical and biological systems, the extended PSP mechanism enhances the decision-making capabilities of autonomous nodes, enabling them to adapt to dynamic environments and heterogeneity in network structure. Acting as rational agents within distributed systems, these nodes create a dynamical system that bridges classical auction theory and real-world decentralized markets, illustrating how local interactions and network effects influence global outcomes. The adaptive behavior of nodes in this framework mirrors natural systems, where individual agents respond to changing conditions by exploring alternative configurations. These systems thrive under uncertainty, leveraging variation as a means to optimize outcomes and avoid suboptimal equilibria. By incorporating adaptive mechanisms and utilizing perturbations constructively, this work demonstrates how decentralized markets can achieve robust and efficient resource allocation, even in the presence of limited information and dynamic conditions.

Our extended PSP mechanism incorporates several key innovations: dynamic bid adjustments, the integration of influence sets, and the inclusion of opt-out conditions. Together, these components create a stable and adaptive auction environment that balances local decision-making with global efficiency. This dissertation outlines a framework for formalizing utility functions, modeling network interactions, and analyzing game-theoretic implications, ensuring that the proposed solution is robust, flexible, and widely applicable.

3.4 Road Map

This section discusses the intellectual trajectory of the dissertation and provides context for what follows. Each chapter is motivated by a set of theoretical or practical questions that arise naturally from the limitations of the previous stage. The goal is not only to describe what is to come, but to explain why each transition is necessary for understanding the emergence of stability, influence, and adaptation in PSP markets.

The dissertation begins by confronting a central question: “How do truthful, efficient, and adaptive market mechanism operate without centralized control, what motivates the participants to act? How can we model adaptive behavior under uncertainty?” Each subsequent chapter contributes a different level of resolution to this question, expanding the theoretical foundation while connecting it to real-world networked markets.

The early chapters establish the mechanism and its initial application. Chapter 4 presents the PSP auction within the Hong Kong Mobile Data Exchange Market, introducing the opt-out function and demonstrating how local participation rules preserve equilibrium in decentralized settings. Chapter 5 refines the theoretical framework, extending the equilibrium analysis to include elasticity and reserve pricing. This chapter provides a formal bridge from single-auction equilibrium models to the multi-auction interpretation developed later.

The middle chapters shift from equilibrium analysis to network and temporal structure. Chapter 6 introduces a bipartite graph model of multi-auction markets, where buyers and sellers are connected through overlapping influence sets. It formalizes the projection-based influence framework, defining primary and secondary influence sets and demonstrating how local interactions produce market-wide coordination and saturation effects. Chapter 7 examines how asynchronous bid updates and latency affect convergence. By introducing bounded delay mechanisms and measuring the effects of initialization noise, it refines our understanding of temporal asynchronicity and stability in distributed auctions.

The final chapters integrate these results into a cohesive dynamic model. Chapter 8 presents the dynamic multi-auction PSP mechanism, generalizing exclusion-compensation and bounded participation principles into a continuous decision process. It unifies valuation, allocation, and cost under a shared temporal framework, showing convergence toward an absorbing ϵ -Nash region. Chapter 9 concludes by reflecting on how adaptive PSP mechanisms address the core challenges of asymmetry, participation, and interdependence. It identifies opportunities for future applications

in wireless communication, vehicular coordination, and decentralized resource allocation.

It is our intent that each chapter builds conceptually and methodologically upon the last, culminating in a coherent theory of decentralized market dynamics that connects equilibrium reasoning, network structure, and temporal adaptation.

Chapter 4

Strategic Bidding and Opt-Out Mechanism in the PSP Auction

This chapter first appeared in the proceedings of ITNG 2020 [14].

Jordan Blocher, Frederick C. Harris, Jr. An Optimization Algorithm for the Sale of OverageData in Hong Kong’s Mobile Data Exchange Market, in Latifi, S. (eds.), 17th International Conference on Information Technology : New Generations (ITNG 2020) Advances in Intelligent Systems and Computing, Volume 1134, Chapter 73, pp 553-561. April 6-8, Las Vegas, NV. Springer, DOI https://doi.org/10.1007/978-3-030-43020-7_73

Abstract

Internet service providers are offering shared data plans where multiple users may purchase and share a single pool of data. In the Chinese economy, users have the ability to sell unused data on the Hong Kong Exchange Market, called “2cm”, currently maintained by AT&T internet services. We propose a software-defined network for modeling this wireless data exchange market; a fully connected, pure “point of sale” market. A game-theoretical analysis identifies and defines rules for a progressive second-price (PSP) auction, which adheres to the underlying market structure. We allow for a single degree of statistical freedom – the reserve price – and show that data exchange markets allow for greater flexibility in acquisition decision-making and mechanism design with an emphasis on optimization of software-defined networks.

We have designed a framework to optimize this strategy space using the inherent elasticity of supply and demand. Using a game theoretic analysis, we derive a buyer-response strategy for wireless users based on second-price market dynamics and prove the existence of a balanced pricing scheme. We examine shifts in the market price function and prove that the desired properties for optimization to a Nash equilibrium hold.

Keywords: software-defined networks, mobile share, game theory, second-price auction

4.1 Introduction

Mobile data usage is quickly outpacing voice and SMS in wireless networks. Multi-device ownership has led to the introduction of the shared data plan [7]. Using an account service, users are able to keep track of data usage in real time across all their devices. The shared data service plan requires that users hold an a priori knowledge of demand and supply with respect to their data plan in order to form a strategy, meaning that a user must *plan* to either buy or sell thier overage data. In our formulation, we address several topics: data as a product in the real-monetary market, and data as network resource in a wireless topology.

Many new services are found exclusively on mobile devices. Companies are moving their software from (wired) grid-based to node-based communication. For example, the move from a standard website to a mobile phone app. Software-defined networking (SDN) addresses the new environment of wireless communication devices, allowing for a programmable network architecture. The account services that manage wireless shared data plans decentralize network management, and mobility becomes a factor in SDN design. Individual mobile devices provide flexibility, and may make decisions regarding local network infrastructure. There is a clear need for algorithms designed for optimization in this space. In many cases, the direct communication between mobile devices allows for a simple mutation of classic optimization models. Auctions are key in SDN for the fair allocation of resources. For this work, we focus on

mobile data, an infinitely divisible and distributable quantity. Mobile data represents online data accessed using a wireless network. In [45], Lazar and Semret introduced the Distributed Progressive Second-Price Mechanism (PSP) for bandwidth allocation. Such an auction is (1) easily distributed, and (2) allocates an infinitely divisible resource. A PSP auction is defined as distributed when the allocations at any element depend only on local state; no single entity holds a global market knowledge. We consider the multi-auction: where each auctioneer is a user selling data to their peers.

The model for data exchange was recently adopted by China Mobile Hong Kong (CMHK), who released a platform, called 2cm (secondary exchange market), creating a secondary market where users can buy and sell data from each other. CMHK owns and moderates 2cm, where CMHK the only auctioneer, and computes allocations of mobile data based on bids submitted to the platform. We focus on providing users with an incentive framework so rational users will choose a collaborative exchange. This collaborative exchange is the (built-in) transformation from the direct-revelation mechanism (truthful bidding) to the desired message space (actual bids).

We describe our auction mechanism as a pure-strategy progressive game with incomplete, but perfect information. The market strategy is determined by the impact of user behavior on market dynamics. The optimal objective is defined as a rational user's valuation of digital property. In classic mechanism design, with multiple user types, there is no single way to design the transformation from the direct revelation mechanism to its corresponding computational design. As in [45], our incentive for a user to truthfully reveal its type is built into the user strategies. We determine (at least one) local equilibrium is a result of incentive compatibility (truthfulness) in strategic bidding, and so our formulation holds the desired PSP qualities. Our derivation of strategies depend on the ratio of supply and demand, and consequently, on the ratio of buyers to sellers.

This is the first work to provide a comprehensive derivation of an auction mechanism with respect to the CMHK platform. The rest of this paper is structured as

follows: Section 4.2 presents the related work on auction theory and resulting policy software. Section 4.3 details the mathematical structure of the data-exchange market, which we present as an extension of the market in [45]. The analysis of user behavior and the resulting algorithms are presented in Section 4.4 along with a simple example. Conclusions and Future Work follow in Section 4.5.

4.2 Related Work

Progressive second-price auctions are used for optimal allocation in a variety of scenarios, and for different reasons. Different definitions of social welfare define different strategies. Typical goals of optimization are the maximization of revenue, and optimal allocation. Other papers focus, taken from auction theory, optimize seller’s reserve prices, or market price. Results derived from game theory focus on player strategy, as in this work. In [70], user strategy gives a “quantized” version of PSP, improving the rate of convergence of the game. Modifications to the mechanism that result in improved convergence also appear in [50], which relies on an approximation of market demand. Another mechanism derived from game theory [74], derives optimal strategies for buyers and brokers (sellers), and further shows the existence of network-wide market equilibria by representing the market dynamics as a system of equations.

Allowing a user preference to, loosely, represent a policy, we may interpret the rules of the data exchange market as a policy scheme, where the ISP is assumed to enforce the rules and the market dynamics play out as a game among “users” of the game. So in a distributed system, users are allowed to set their own policies, and the ISP is responsible for implementing the framework to support their preferences. Trusted management systems are based on the Common Information Model (CIM), and focus on policy-based management, for example the “Policy-Maker” toolkit. In general, the translation of policy-based management systems to SDN focuses on combining the simplicity of policy-based implementation with the flexibility of SDN, as in the meta-policy system, CIM-SDN [68].

Game-theoretical analysis of mobile data has been presented in [92] as a framework for mobile-data offloading. In our analysis, the stability of the game is expressed as the set of equilibria, or fixed points, of the system. When considering the distributed and decentralized allocation of resources, a variety of equilibria exist for heterogeneous and homogeneous services once a certain set of conditions is met, one of which is truthfulness.

4.3 The Market Mechanism

In a distributed PSP auction, the design must meet a certain set of known criteria: (1) *truthfulness* (incentive compatibility), (2) *individual rationality/selfishness*, and (3) *social welfare maximization (exclusion-compensation)*. We examine the PSP auction as the constraints are able to attain the desirable property of truthfulness through incentive compatibility, meaning that an user has more of an incentive to tell the truth. This is because in second-price markets, the winning bid does not pay the winning bid price, but the price from *next lowest bid*. The pricing mechanism also upholds the exclusion-compensation principle, or Pareto criterion, where any change to the system would make at least one user worse-off. We construct the model for a PSP data auction for mobile users participating in secondary mobile data exchange market.

Let the set of all wireless users to be labeled by the index set $\mathcal{I} = \{1, \dots, I\}$. In our current formulation, we do not allow a seller to host multiple auctions, thus we may identify each local auction with the index of the seller $j \in \mathcal{I}$. The bid profiles of the users are given as, $s \equiv [s_i^j]$ where $(i, j) \in \mathcal{I} \times \mathcal{I}$. Now, this is a single bid, where we fill the space by submitting zero bids to all non-active users, meaning that if there is no interaction between two players i and j , then $(i, j) = 0$. One may think of it as an $\mathcal{I} \times \mathcal{I}$ matrix, with each element of the matrix representing a pair-interaction. However this matrix is just one projective representation of the space. A single snapshot of a static system, all quantities and prices are fixed may be represented by this matrix. Once users begin to bid, then we must consider all

possible interactions, which is done by fixing one index in \mathcal{I} at a time, allowing all other quantities to vary. So the strategy space in fact includes all the possibilities for an user in \mathcal{I} ; another dimension to the problem is added with each possible variation. We call this space S , the (full) strategy space for buyer i as all possible bids at all auctions (where i 's bid changes with respect to the variation of all other bids): $S_i = \Pi_{j \in \mathcal{I}} S_i^j$, and $S_{-i} = \Pi_{j \in \mathcal{I}} (\Pi_{k \neq i \in \mathcal{I}} S_k^j)$ as the associated opponent profiles, as in standard game-theoretic notation.

The grid(s) of bid profiles, s , represents the uncertain state of the distributed PSP auction mechanism in the secondary market, where we take uncertain to mean the statistical distribution of player types and corresponding actions. In general, we will not reference the full grid s . We will also use the context of the bid to indicate the user type. To further clarify our analysis, we adopt the following notational conventions: a seller's profile is denoted by $s^j = [s_i^j]_{i \in \mathcal{I}}$, and $s_i = [s_i^j]_{j \in \mathcal{I}}$ denotes a buyer's profile, where $s_{-i} \equiv [s_1^j, \dots, s_{i-1}^j, s_{i+1}^j, \dots, s_I^j]_{j \in \mathcal{I}}$ as the profile of user i 's opponents. Furthermore, noting that this is a simplification for ease of notation, we let $Q^j = \sum_{i \in \mathcal{I}} q_i^j$ be the total amount of data j has to sell, and $Q_i = \sum_{j \in \mathcal{I}} q_i^j$ represent the total amount of data desired by buyer i .

We assume a public platform, published by the ISP, that allows sellers to advertise their auctions. Buyers may submit bids directly to sellers over the wireless network. We also assume that a buyer's budget is sufficient, as the alternative would be to pay a higher price to the ISP. We describe the rules as follows:

- The bid is represented by $s_i^j = (q_i^j, p_i^j)$, meaning i would like to buy from j a quantity q_i^j and is willing to pay a unit price p_i^j .
- The seller takes responsibility for notifying i of opponent bid profiles s_{-i} , and updates the bid profile when buyer i joins the auction.
- $s_i^j > 0$ represents a buyer-seller pair in s , with bid, $s_i^j = (q_i^j, p_i^j)$, where quantity $q_i^j \in d^j$ is an element of $[q_i^j]_{i \in \mathcal{I}}$, with reserve unit price $p_i^j \in p^j$, an element of $[p_i^j]_{i \in \mathcal{I}}$.

- If a buyer does not submit a bid to a seller, then this implies $s_i^j = 0$. A buyer that does not submit a bid will not receive opponent profiles from seller j .
- A user who does not submit a bid is holding to the previous bid, either zero or nonzero.

We emphasize that buyers are consistently referenced using the index i as a subscript, and sellers using the index j as a superscript, as in [73].

4.3.1 Market incentive.

We examine the role of buyers, who are able to directly influence global market dynamics, and assume that the sellers take a reactionary role. Each buyer i will have information from each seller j , as well as opponent profiles s_{-i} , from each auction in which it is participating. In the extreme case, where i submits bids to all auctions $j \in \mathcal{I}$, buyer i gains access all buyer profiles, $[s_1, \dots, s_I]$. However, sellers can only gain information about the market by observing buyer behavior in their local auction. Buyers, on the other hand, can see all the sellers reserve prices, although they can only see their opponent bid profiles.

Define the set of sellers chosen by buyer $i \in \mathcal{I}$ as,

$$\mathcal{I}_i(n) = \arg \max_{\mathcal{I}' \subset \mathcal{I}, |\mathcal{I}'|=n} \sum_{j \in \mathcal{I}'} Q^j,$$

and similarly, for a seller $j \in \mathcal{I}$, we define the set of buyers participating in auction j as,

$$\mathcal{I}^j(m) = \arg \max_{\mathcal{I}' \subset \mathcal{I}, |\mathcal{I}'|=m} \sum_{i \in \mathcal{I}'} p_i^j,$$

where $m, n \in \mathcal{I}$.

The PSP auction given in [45] is a set of simple and symmetric rules that closely follow market theory. We now formally define a PSP auction, which determines the actions buyers and sellers in the secondary market. We define an **opt-out function**, σ_i , associated with a buyer i as part of its type. Buyer i , when determining how to

acquire a possible allocation a , will determine its bid quantities by,

$$\sigma_i(a) = [\sigma_i^j(a)]_{j \in \mathcal{I}}. \quad (4.1)$$

In a general sense, σ_i applies our user strategy to the PSP rules. The rules presented here incorporate of the opt-out function with the auction mechanism, and closely follows the work presented in [45]. The market price function, P_i , for a buyer in the secondary market can be described as follows:

$$\begin{aligned} P_i(z, s_{-i}) &= \sum_{j \in \mathcal{I}} \sigma_i^j \circ p_i^j(z_i^j, s_{-i}^j) \\ &= \sum_{j \in \mathcal{I}} \left(\inf \left\{ y \geq 0 : q_i^j(y, s_{-i}^j) \geq \sigma_i^j(z) \right\} \right), \end{aligned} \quad (4.2)$$

and is interpreted as the aggregate of minimum prices that buyer i bids in order to obtain data amount z given opponent profile s_{-i} . We note that in the following analysis the total minimum price for the buyer cannot be an aggregation of the *individual* prices of the buyers, as it is possible that the reserve prices of the sellers may vary. The maximum available quantity of data in auction j at unit price y given s_{-i}^j is:

$$q_i^j(y, s_{-i}^j) = \sigma_i^j \circ q_i^j(y, s_{-i}^j) = \left[Q^j - \sum_{p_k^j > y} \sigma_k^j(a) \right]^+. \quad (4.3)$$

It follows from the upper-semicontinuity of Q_i^j that for s_{-i}^j fixed, $\forall y, z \geq 0$,

$$\sigma_i^j(z) \leq \sigma_i^j \circ q_i^j(y, s_{-i}^j) \Leftrightarrow y \geq \sigma_i^j \circ p_i^j(z, s_{-i}^j). \quad (4.4)$$

The resulting data allocation rule is a function of the local market interactions between buyers and sellers over all local auctions, as is composed with i 's opt-out value, so that for each $i \in \mathcal{I}$, the allocation from auction j is,

$$\begin{aligned} a_i^j(s) &= \sigma_i^j \circ a_i^j(s) \\ &= \min \left\{ \sigma_i^j(a), \sum_{\substack{p_k^j = p_i^j \\ p_k^j < \sigma_i^j(a)}} \sigma_k^j(a) q_i^j(p_i^j, s_{-i}^j) \right\}, \end{aligned} \quad (4.5)$$

noting that for the full allocation from all auctions we may simply aggregate over the seller pool.

Remark: The bid quantity $\sigma_i^j(a)$ and the allocation a_i^j are complementary. In fact, the buyer strategy is the first term in the minimum, the second term being owned by the seller.

Finally, we must have that the cost to the buyer adheres to the second-price rule for each local auction, with total cost to buyer i ,

$$c_i(s) = \sum_{j \in \mathcal{I}} p_i^j \left(a_i^j(0; s_{-i}^j) - a_i^j(s_i^j; s_{-i}^j) \right). \quad (4.6)$$

The cost to buyer i adds up the willingness of all buyers excluded by player i to pay for quantity a_i^j . i.e.

$$c_i^j(s) = \int_0^{a_i^j} p_i^j(z, s_{-i}) dz.$$

This is the “social opportunity cost” of the PSP pricing rule.

4.4 User Strategy

In any market, a buyer or seller would like to obtain the maximum amount of utility possible while staying within budget. The buyer’s utility maximizes the amount of data allocated by the seller, while the seller’s utility maximizes the cost of the data sold. Clearly, the cost is the product of the unit price and the desired allocation. We examine cases where the buyer has found an allocation that satisfies its demand AND price constraints, and define a strategic bid to a move to a better market position.

4.4.1 User valuation (strategic incentive).

We define each buyer as a user $i \in \mathcal{I}$ with quasi-linear utility function $u_i = [u_i^j]_{j \in \mathcal{I}}$. A buyers’ utility function is of the form,

$$u_i = \theta_i \circ \sigma_i(a) - c_i, \quad (4.7)$$

where the composition of the elastic valuation function θ_i with σ_i distributes a buyers’ valuation of the desired allocation a across local markets, submitting the strategic bid

to multiple seller's auctions. The composition map represents the codomain of $\theta_i(\sigma_i)$, which is the same as the domain of $\sigma_i(a)$, and performs the function of restricting the buyer's domain to minimize $d^j p^j - c_i$, i.e., maximize u_i . Using this rule, we extend the PSP rules described in [73] in order to find equilibria in subsets of local data-exchange markets.

The sellers, $j \in \mathcal{I}$ are not associated with an opt-out function. We consider their valuation to be a functional extension of the buyers, where θ^j is constructed from buyer demand. We adopt the definition for an elastic valuation function as in [45], which allows for continuity of constraints imposed by the user strategies.

Definition 4.1. (*Elastic demand*) [45] *A real valued function, $\theta(\cdot) : [0, \infty) \rightarrow [0, \infty)$, is an (elastic) valuation function on $[0, D]$ if*

- $\theta(0) = 0$,
- θ is differentiable,
- $\theta' \geq 0$, and θ' is non-increasing and continuous,
- There exists $\gamma > 0$, such that for all $z \in [0, D]$, $\theta'(z) > 0$ implies that for all $\eta \in [0, z)$, $\theta'(z) \leq \theta'(\eta) - \gamma(z - \eta)$.

We begin our analysis with buyer valuation θ_i . A buyers' valuation of an amount of data represents how much a buyer is willing to pay for a unit of data (bandwidth). This is equivalent to the bid price when given a fixed amount of data. The buyers' utility-maximizing bid (fixing the desired allocation $z \geq 0$) is a mapping to the lowest possible unit price,

$$f_i(z) \triangleq \inf \{y \geq 0 : \rho_i(y) \geq z, \forall j \in \mathcal{I}\}, \quad (4.8)$$

where $\rho_i(y)$ represents the demand function of buyer i at bid price $y \geq 0$. The market supply function is the extreme case of possible buyer demand, and acts as an “inverse” function of f_i . We have, for bid price $y \geq 0$, $\rho_i(y) = \sum_{j \in \mathcal{I}: p_i^j \geq y} Q^j$. The utility-maximizing bid price is the lowest unit cost for the buyer to be able participate

in all the auctions in \mathcal{I}_i , and corresponds to the maximum reserve price amongst the sellers.

From the perspective of the seller we have a more direct interpretation of valuation as revenue. The demand function of seller j at reserve price $y \geq 0$ is $\rho^j(y) = \sum_{i \in \mathcal{I}: p_i^j \geq y} \sigma_i^j(a)$. We define the “inverse” of the buyer demand function for seller j as potential revenue at unit price y , we have,

$$f^j(z) \triangleq \sup \{y \geq 0 : \rho^j(y) \geq z, \forall i \in \mathcal{I}\}. \quad (4.9)$$

Unsurprisingly, f^j maps quantity z to the highest possible unit data price.

We show that user valuation satisfies the conditions for an elastic demand function, based on (4.9). We first note that, in general (and so we omit the subscript/superscript notation), the valuation of data quantity $x \geq 0$ is given by, $\theta(x) = \int_0^x f(z) dz$. We propose the following Lemma,

Lemma 4.2. *(user valuation) For any buyer $i \in \mathcal{I}$, the valuation of a potential allocation a is,*

$$\theta_i \circ \sigma_i(a) = \sum_{j \in \mathcal{I}} \int_0^{\sigma_i^j(a)} f_i(z) dz. \quad (4.10)$$

Now, we may define seller j ’s valuation in terms of revenue,

$$\theta^j = \sum_{i \in \mathcal{I}} \theta^j \circ \sigma_i^j(a) = \sum_{i \in \mathcal{I}} \int_0^{\sigma_i^j(a)} f^j(z) dz. \quad (4.11)$$

We have that θ_i and θ^j are elastic valuation functions, with derivatives θ_i' and θ^j' satisfying the conditions of elastic demand.

Proof. Let ξ be a unit of data from buyer bid quantity $\sigma_i^j(a)$. If ξ decreases by incremental amount x , then seller bid q_i^j must similarly decrease. The lost potential revenue for seller j is the price of the unit times the quantity decreased, by definition, $f^j(\xi)x$, and so, $\theta^j(\xi) - \theta^j(\xi - x) = f^j(\xi)x$, and (4.11) holds. As we may use the same argument for (4.10), as such, we will denote $f_i = f^j = f$ for the remainder of the proof. We observe that the function f is the first derivative of the valuation function

with respect to quantity. Letting $\theta_i = \theta^j = \theta$, the existence of the derivative implies θ is continuous, and therefore, in this context, f represents the marginal valuation of the user, θ' . Also, clearly $\theta(0) = \theta(\sigma(0)) = 0$. Now, as we consider data to be an infinitely divisible resource, we have a continuous interval between allocations a and b , where $a \leq b$. Now, as θ is differentiable, for some $c \in [a, b]$,

$$\theta'(c) = \lim_{x \rightarrow c} \frac{\theta(x) - \theta(c)}{x - c} = f(c),$$

and so $f = \theta'$ is differentiable at $c \in [a, b]$, and so as $a \geq 0$, $\theta' \geq 0$. Finally, we have that concavity follows from the demand function. Then, as θ' is non-increasing, we may denote its derivative $\gamma \leq 0$, and taking the derivative of the Taylor approximation, we have, $\theta'(z) \leq \theta'(\eta) + \gamma(z - \eta)$. \square

Finally, it is worth mention that the analysis of the auction as a game only assumes some form of demand and supply, in order to derive properties. The mechanism itself does not require any knowledge of user demand or valuation.

4.4.2 User behavior.

Buyers and sellers are able to change their bid strategies asynchronously. A user's local strategy space is therefore non-deterministic as the preferences of users are subject to change. Although it is possible for a seller to fully satisfy a buyer i 's demand, it is also reasonable to expect that a seller's overage data may not satisfy even a single buyer's demand. In this case, a buyer must split its bid among multiple sellers. The buyer strategy bids in auctions with the highest quantities first, a natural result of the demand curve.

The buyer strategy tends towards equal valuation of all local markets, and therefore similar prices. Buyer i 's seller pool is determined by minimizing n , the smallest set of sellers that satisfy its demand Q_i : $\min \{n \in \mathcal{I} \mid nD^n \geq Q_i\}$. Similarly, seller j determines the minimal set of buyers that maximizes revenue and sells all of its data, Q^j , i.e. $\min \{m \in \mathcal{I} \mid \sum_{i \in \mathcal{I}^j(m)} q_i^j \geq Q^j\}$. We use $j^* = n \leq I$ to represent the seller with the least amount of data in \mathcal{I}_i , i.e. $D^{j^*} \leq Q^j, \forall j \in \mathcal{I}^j$.

Define the composition,

$$\sigma_i^j \circ a = \sigma_i^j(a) = \frac{a_i^j}{|\mathcal{I}_i|}, \quad (4.12)$$

to be the buyer strategy with respect to quantity for all sellers $j \in \mathcal{I}_i$. We propose the following scheme:

Definition 4.3. (*Opt-out buyer strategy*) Let $i \in \mathcal{I}$ be a buyer and fix all other buyers' bids s_{-i} at time $t > 0$, and let a be i 's desired allocation. Define,

$$\sigma_i^j(a) \triangleq \begin{cases} \sigma_i^{j^*}(a), & j \in \mathcal{I}^j, \\ 0, & j \not\in \mathcal{I}^j. \end{cases} \quad (4.13)$$

and bid price $p_i^j = \theta'_i(\sigma_i^j(a))$.

Let the reserve price for seller j be,

$$p_*^j = p_{i^*}^j + \epsilon, \quad (4.14)$$

where i^* is the bidder with the highest "losing" bid price. A truthful bid implies that the new bid price differs from the last bid price by at least ϵ .

We will show that sellers are able to maximize revenue in a restricted subset of buyers in \mathcal{I} , and as such will attempt to facilitate a local market equilibrium for this subset. A local auction j converges when $s_i^{j(t+1)} = s_i^{j(t)} \forall i \in \mathcal{I}$, at which point the allocation is stable, the data is sold, and the auction ends. We propose a strategy to maximize (local) seller revenue.

Lemma 4.4. (*Localized seller strategy (i.e. progressive allocation)*) For any seller j , fix all other bids $[s_i^k]_{i,k \neq j \in \mathcal{I}}$ at time $t > 0 \in \tau$. For each $t \in \tau$, let $\omega(t)$ be define the winner at time t , and perform the update,

$$D^{j(t+1)} = D^{j(t)} - \sigma_{\omega(t)}^{j(t)}(a). \quad (4.15)$$

Allowing t to range over τ , we have that $Q^j = 0$, and a local market equilibrium.

We omit the proof, and provide a simple example.

4.4.3 A simple example.

We give a simple example of convergence to a local market equilibrium, where the buyers are assumed to respond according to (4.5).

Name	Bid total	Unit price
A	50	1
B	40	1.2
C	26	1.5
D	20	2
E	14	2.2

Let $s^{(1)} = [(65, \epsilon)]_{i \in \mathcal{I}}$ and $s^{(2)} = [(85, \epsilon)]_{i \in \mathcal{I}}$. The buyer bids are as follows:

$$s_A = [(0, 0), (50, 1)],$$

$$s_B = [(0, 0), (40, 1.2)],$$

$$s_C = [(0, 0), (26, 1.5)],$$

$$s_D = [(0, 0), (20, 2)],$$

$$s_E = [(0, 0), (14, 2.2)].$$

Then at $t = 1$, $s^{(2)} = [(0, p^{(2)}), (20, p^{(2)}), (26, p^{(2)}), (20, p^{(2)}), (14, p^{(2)})]$, and so $(D^{(2)}, p^{(2)}) = (85, 1 + \epsilon)$, The buyer response is,

$$s_A = [(50, 1), (0, 0)],$$

$$s_B = [(40, 1.2), (0, 0)],$$

$$s_C = [(0, 0), (26, p^{(2)})],$$

$$s_D = [(0, 0), (20, p^{(2)})],$$

$$s_E = [(0, 0), (14, p^{(2)})].$$

At $t = 2$, $(D^{(1)}, p^{(1)}) = (65, 1 + \epsilon)$, with bid vector $s^{(1)} = [(25, p^{(1)}), (40, p^{(1)}), (0, 0), (0, 0), (0, 0)]$. $(D^{(2)}, p^{(2)}) = (25, 1 + \epsilon)$. Then,

$$s_A = [(25, p^{(1)}), (25, p^{(2)})],$$

$$s_B = [(40, p^{(1)}), (0, 0)],$$

where we have removed bids to indicate winner(s) with a tentative allocation. At $t = 3$, $(D^{(1)}, p^{(1)}) = (50, 1 + \epsilon)$, with bid vector $s^{(1)} = [(25, p^{(1)}), (40, p^{(1)}), (0, 0), (0, 0), (0, 0)]$. $(D^{(2)}, p^{(2)}) = (0, 1 + \epsilon)$ and $s^{(2)} = [(25, p^{(1)}), (0, 0), (26, p^{(2)}), (20, p^{(2)}), (14, p^{(2)})]$. Then,

$$s_A = [(25, p^{(1)}), (0, 0)].$$

At $t = 4$ the auction ends.

4.4.3.1 Individual rationality/selfishness.

Value is modeled as a function of the entire marketplace: a buyer's valuation is aggregated over all the auctions, and the seller's valuation is aggregated over its own auction. We must ensure that a user's private action satisfies the conditions of a direct-revelation mechanism, as well as adheres to the collective goals. We show that, from (4.4) and (4.3), an individual user will contribute to local stability, given global market dynamics S .

We model the impact of the dynamics of S of the data-exchange market on a local auction j . As we have shown, the seller behavior is a reaction of buyer behavior, and have presented some rules. The market fluctuations from S give auctioneer j the chance to infer information about the global market. We demonstrate that the symmetry between buyer and seller behavior stretches across subsets of local auctions. Additionally, we identify a clear bound restricting the range of influence that local auctions have on each other. Consider a single iteration of the auction, where a seller updates bid vector s^j , and the buyers' response s_i , to comprise a single time step. We have the following Proposition,

Proposition 4.5. *(Valuation across local auctions) For any $i, j \in \mathcal{I}$,*

$$j \in \mathcal{I}_i \Leftrightarrow i \in \mathcal{I}^j. \tag{4.16}$$

Fix an auction $j \in \mathcal{I}$ with duration τ and define the influence sets of users. The

primary and secondary influencing sets are given as,

$$\Lambda = \bigcup_{i \in \mathcal{I}^j} \mathcal{I}_i, \quad \text{and} \quad \lambda = \bigcup_{i \in \mathcal{I}^j} \left(\bigcup_{k \in \mathcal{I}_i} \mathcal{I}^k \right). \quad (4.17)$$

Define $\Delta = \Lambda \cup \lambda$. Fixing all other bids $s_i^j \in \mathcal{I}$, and time $t > 0 \in \tau$, we have that,

$$\sum_{j \in \Lambda} \theta_i^j = \sum_{i \in \lambda} \theta_i^j. \quad (4.18)$$

Proof. As this is our main result, we provide an outline of the (exhaustive) proof, illustrating the most important case, when a market shifts affect auction j , and the direct influence of the shift on the connected subset of local markets.

A local auction $j \in \mathcal{I}$, is determined by the collection of buyer bid profiles. Using Proposition 4.4 and (4.16), we have that,

$$i \in \mathcal{I}^j \Leftrightarrow p_i^j > p_{i^*}^j, \quad (4.19)$$

where we define i^* as the losing buyer with the highest bid price in auction j . By (4.8) $p_i^j \geq p_{i^*}^j + \epsilon$, thus $p_i^j < p_{i^*}^j$ can only happen during a market shift. Consider $k \in \mathcal{I}^j$ at time t where, for example, some buyer(s) enter the auction, and so (4.19) implies that $\sum_{i \in \mathcal{I}^j} \sigma_i^j(a) > Q^j$. Now, $p_i^j < p_{i^*}^j \Rightarrow k \in \mathcal{I}^j$ and $s_k^j > 0$ will cause k to initiate a shift. By definition 4.3, k will set $s_k^j = 0$, and begin to add sellers to its pool. Suppose that at time t , j 's market is at equilibrium. Unless k adds a seller with a higher reserve price within $|\mathcal{I}^j|$ time steps, by (4.15), the auction ends. We have that, $\forall i \in \mathcal{I}^j, \nexists s_i^j > 0$ where $i \in \mathcal{I}^j$, and (4.16) holds.

Now, the subset $\mathcal{I}^j \subset \mathcal{I}$ determines j 's reserve price $p_{i^*}^j$. We will assume the buyer submits a coordinated bid, using (4.5). The reserve price (4.14) of seller j is determined at each shift, and is the lowest price that j will accept to perform any allocation. Let p_*^j denote the reserve price of auction j and p_i^* denote the bid price of buyer i , i.e. $p_i^k = p_i^*$, $\forall k \in \mathcal{I}_i$. Using Proposition 4.4, for each $i \in \mathcal{I}^j$, we have from (4.8), (4.9), that $p_i^* \geq p_*^j$, $\forall k \in \mathcal{I}_i$. In the simplest case, consider a disjoint local market j , where $\forall i \in \mathcal{I}^j, s_i^k = 0$, $\forall k \neq j \in \mathcal{I}_i \Rightarrow \Lambda = \{j\}$ and $\lambda = \mathcal{I}^j$. Again

using (4.8) and (4.9), it is clear that $\theta_i = \theta^j$, $\forall i \in \mathcal{I}^j$. In all other cases, the sellers $\in \Lambda$ are competing to sell their respective resources to buyers whose valuations are distributed across multiple auctions. The bid price of buyer $i \in \mathcal{I}^j$ is determined by, $p_i^* = \max_{k \in \mathcal{I}_i}(p_k^*)$. Λ is the set of sellers directly influencing the bids of buyers in auction j . Now, the reserve price for auction j is such that, $p_*^j \leq \min_{i \in \mathcal{I}^j}(p_i^*) - \epsilon$. From (4.17), Λ is defined by a seller $j \in \mathcal{I}$, where each user $k \in \lambda$ has some direct or indirect influence on j . Denote $\Delta^j = \Lambda^j \cup \lambda^j$.

Consider the set λ^j . For some buyer $i \in \mathcal{I}^j$, and then for some seller $k \in \mathcal{I}_i$, we have a buyer $l \in \mathcal{I}^k$. By (4.16), $i, l \in \mathcal{I}^k$, and so the reserve price $p_*^k \leq \min(p_l^*, p_i^*)$, and $k, j \in \mathcal{I}_i \Rightarrow p_i^* \geq \max(p_*^k, p_*^j)$. Suppose that $l \in \mathcal{I}^j \Leftrightarrow j \in \mathcal{I}_l$, so that $p_l^* < p_*^j$, and the valuation of buyer l does not impact auction j and vice versa, i.e. $\theta_l^j = 0$. Since $l \in \mathcal{I}^k$, $p_l^* \geq p_*^k \Rightarrow p_*^k < p_*^j$, and $i \in \mathcal{I}^j \Rightarrow p_i^* \geq p_*^j$. Therefore, we have that the ordering implied by (4.17) holds, and,

$$p_*^k \leq p_l^* < p_*^j \leq p_i^*, \quad (4.20)$$

for any buyer $l \in \lambda^j$ such that $l \in \mathcal{I}^j$. We use a similar argument for a secondary user $q \in \mathcal{I}_l$.

Finally, consider the subset Δ^j ; a shift occurs in 2 cases. (1) If $i \in \mathcal{I}^j$ decreases its bid quantity so that $\sum_{a \in \mathcal{I}^j} \sigma_i^j(a) < Q^j$, and (2) if buyer i^* , defined in Proposition 4.4, increases its valuation so that $p_{i^*}^j < p_*^j$. Fixing all other bids, a decrease in q 's demand will directly impact buyer i . If at the end of the bid iteration, we still have that i is the buyer with the lowest bid price, then (4.9) holds and j 's valuation does not change. Otherwise a new i^* will be chosen upon recomputing \mathcal{I}^j , as in Lemma 4.3, and the market will attempt to regain equilibrium. We determine the influence of Δ^{k^*} on Δ^j by (4.19).

In each case we have that (4.8) and (4.9) hold for some fixed time t , and so, $\forall i \in \mathcal{I}^j$, any bid outside of our construction has a zero valuation, with respect to buyers $\in \lambda$ and sellers $\in \Lambda$, and therefore cannot cause shifts to occur except through a shared buyer, e.g. some $l \in \mathcal{I}^k$. Thus, in all cases, (4.8) and (4.9) hold. Fixing all

bids in any auction where $q \ni \Lambda^j, \forall i \in \mathcal{I}^j, \forall k \in \mathcal{I}_i, \forall l \in \mathcal{I}^k$,

$$\int_0^{\sigma_i^k(a)} f_i(z) dz = \int_0^{\sigma_i^k(a)} f^k(z) dz, \quad (4.21)$$

and

$$\int_0^{\sigma_l^k(a)} f^k(z) dz = \int_0^{\sigma_l^k(a)} f_l(z) dz. \quad (4.22)$$

Thus, with a slight abuse of notation for clarity,

$$\sum_{\lambda} \int_0^{\sigma(a)} f^{\Lambda}(z) dz = \sum_{\Lambda} \int_0^{\sigma(a)} f_{\lambda}(z) dz, \quad (4.23)$$

where the result follows by construction, and the continuity of θ' . \square

4.4.3.2 Truthfulness (incentive compatibility).

We will prove that the dominant strategy for buyers is to submit coordinated bids, where all bids the buyer submits are equal. Our motivation for coordinated bids comes from the idea of potential games. In potential games, the incentive of all users to change strategy can be expressed as a single global function. The necessary condition of an ϵ -best reply is that the new bid price must differ from the last by at least ϵ . Thus, our strategic bid is an ϵ -best response. Now, an ϵ -best reply for user i is $p_i^* = \theta'_i(\sigma_i(a)) + \epsilon$, for a given opponent profile s_{-i} , and for each $j \in \mathcal{I}_i$. Now, as ϵ is the bid fee, we have that p_i^j is equal to the marginal valuation of player i in auction j , and so incentive compatibility holds.

4.4.3.3 Social welfare maximization (exclusion-compensation).

We define an optimal state of social welfare to be when valuations are equal across a subset of local auctions. Then, $\Delta \subset \mathcal{I}$ is the subset of users where social welfare is achieved. We finally have:

Corollary 4.6. (Δ -Pareto efficiency) *The subset $\Delta \subset \mathcal{I}$ is Pareto efficient, in that no user can make a strategic move without making any other user worse off.*

Proof. Define $s_* = (z_*, \theta'_*(z_*)$ as the set of truthful ϵ -best replies for user i given opponent bid profile S_{-i} , where $\forall j \in \mathcal{I}_i, s_*^j = s_*$. Since θ'_i is continuous, as was shown in Lemma 4.2, and as $s|_{\Delta} = \{[s_i^j] \in \lambda^j \times \Lambda^j\}$ is continuous in s on $S_k = \Pi_{k \in \lambda^j} S_k^j$, then given that $s_* = s^* = (f^*(p^*), p^*) = (z^*, \theta'(z^*))$, we have that s^* is truthful. The result now follows directly from the result of Proposition 4.5. \square

4.5 Conclusion and Future Work

We take these results as evidence of (at least one) fixed point, and conjecture that an optimal solution exists, where all users will receive the desired amount of data (negative or positive), at a fair price.

The PSP auction is a natural data-pricing scheme for consumers accessing a data-exchange market in their wireless network, and that the desired properties of (1) *truthfulness*, (2) *individual rationality/ selfishness*, and (3) *social welfare maximization* are met. We conclude that there is a need for better management of data on the consumer level; an advanced implementation such as the PSP auction presented here would ensure that the consumers in any such exchange market benefit from their participation. It is clear that there is profit to be made by supplying data to the data-driven consumer. However, customer care is necessary to hold the “lifetime consumer”. Consumers, when allowed to manage their own overage data, are able to do so fairly and efficiently. It is not unreasonable to allow them to manage their own data; this benefits all wireless users.

Mathematically, we have shown that if truthfulness holds locally for both buyers and sellers, i.e. $p_i = \theta'_i, \forall j \in \mathcal{I}_i$ and $p^j = \theta^{j'}$, $\forall i \in \mathcal{I}^j$, then, in the absence of market shifts, there exists an ϵ -Nash equilibrium extending over a subset of connected local markets. Observing the symmetric, natural topology of the strategy space, we conjecture that a unique subspace limit exists for connected Δ . A study of this space and the design of the necessary framework is the direction of our future work.

In future work, we intend to show that $s|_{\Delta}$ represents a continuous mapping

$[0, \sum_{k \in \lambda^j} Q^k]_{i \in \Lambda^j}$ onto itself, and show that the continuous mapping of the convex compact set s_* into itself (s^*) has at least one fixed point, i.e., \exists some $k \neq i$ such that $z^* = \sigma^*(z) \in [0, D_k]_{i \in \Lambda^j}$. We want to show that the symmetry built into strategy space provides built-in conditions for convergence and stability, indicating a network Nash equilibrium (NE). Wireless users are modeled as a distribution of buyers and sellers with normal incentives.

Finally, as a result of user behavior and subsequent strategies, we determine that the data-exchange market behaves in a predictable way. However, each auction may be played on the same or on different scales in valuation, time, and quantity; therefore the rate at which market fluctuations occur is impossible to predict. Nonetheless, we have shown that our bidding strategy results in (at least one), Nash equilibrium, where again the reserve prices are fixed by the seller at bid time.

Chapter 5

An Equilibrium Analysis of a Secondary Mobile Data-Share Market

An Equilibrium Analysis of a Secondary Mobile Data-Share Market

This chapter first appeared as a Journal publication in Information [13].

Jordan Blocher and Frederick C Harris, Jr. An Equilibrium Analysis of a Secondary Mobile Data-Share Market, *Information*, Vol 12, No 11, pg 434, October 2021.
<https://doi.org/10.3390/info12110434>

Abstract

Internet service providers are offering shared data plans where multiple users may buy and sell their overage data in a secondary market managed by the ISP. We propose game-theoretic approach to a software-defined network for modeling this wireless data exchange market: a fully connected, noncooperative network. We identify and define the rules for the underlying progressive second-price (PSP) auction for the respective network and market structure. We allow for a single degree of statistical freedom – the reserve price – and show that the secondary data exchange market allow for greater flexibility in the acquisition decision-making of mechanism design. We have designed a framework to optimize the strategy space using elasticity of supply and demand. Wireless users are modeled as a distribution of buyers and sellers with normal incentives. Our derivation of a buyer-response strategy for wireless users

based on second-price market dynamics leads us to prove the existence of a balanced pricing scheme. We examine shifts in the market price function and prove that our network upholds the desired properties for optimization with respect to software-defined networks and prove the existence of a Nash equilibrium in the overlying noncooperative game.

Keywords: software-defined networks; mobile share; game theory; second-price auction

5.1 Introduction

We study an evolving noncooperative real time market. The changing demand of consumers operating on a wireless network has led companies to change their strategy towards data management and networking among wireless users. Mobile phone data experiences dynamically changing demand. Companies are beginning to measure, customize prices, and serve customers according to their individual demand in real-time. Recently, AT&T has introduced a mobile data sharing plan [7], which allows consumers to manage their own data usage in a real time auction on a public platform managed by AT&T. In this work, we derive the Nash equilibrium by examining the resulting structure of this secondary data market. We address several topics: data as a product in the real-monetary market, and data as a network resource in a wireless topology.

As older softwares move from (wired) grid-based to node-based communication, a new paradigm of programmable network architectures is becoming the standard for wireless communication. These decentralized software-centric networks manage services modularly and allow for flexibility of demand; a decentralized design better suited to managing mobile networks, and mobility becomes a factor in network design. The relationship between the provider and the consumer is shifting; it is becoming a necessity to discretize consumer valuation. This is the new paradigm of digital economics.

There is a clear need for algorithms designed for optimization in this mobile space.

Mobile devices provide flexibility, and may make individual decisions regarding local network infrastructure. In many cases, the direct communication between mobile devices allows for a simple mutation of classic optimization models. In online (web) resource allocation algorithms, auctions are used to implement the fair distribution of resources in network allocation algorithms. Mobile data is individually requested by each consumer as demand is unique to each mobile device. For the customer, they perceive a higher level of value for the service. For the ISP, progressive pricing allows for all customers to actively control, create, and share value. As stated by Izaret and Schürmann, "Making progressive pricing a ... reality can happen only if firms change how they create, define, and measure value so that they can share it fairly." [36]

An auction mechanism is defined as distributed when the allocations at any element depend only on local state, i.e. no single entity holds a global market knowledge. In [45], Lazar and Semret introduced the Distributed Progressive Second-Price (PSP) Mechanism for bandwidth allocation, auction mechanisms that are (1) easily distributed, and (2) allocate an infinitely divisible resource. In classic mechanism design, with multiple user types, there is no single way to design the transformation from the direct revelation mechanism to its corresponding computational design. We apply a modifier to the PSP mechanism in order to mutate the strategy space, following the dynamics of user correspondence. As in [45], we obtain our result by design in composition with the PSP rules.

The secondary market provides a unique opportunity for social equilibrium, as it allows users to share data without sharing the same data plan, a restriction in most ISPs, such as [7]. We address the need for privacy in the bidding market; bid privacy is a concern for two reasons: (1) Sellers may use a buyer's valuation to discriminate against a specific buyer(s), (2) an auctioneer might create a fake second highest bid slightly below the highest bid in order to increase his revenue. In general, the buyer does not trust the auctioneer. We therefore determine that our mechanism must be locally privacy-preserving. By privacy-preserving, we mean anonymous. That is, the winning bid in each local auction maintains anonymity, although the bid itself may

be public information. In this way, consumers are able to trust the pricing scheme of the local auction and the resulting allocation.

To the best of our knowledge, this is the first work to provide a comprehensive derivation of a truthful mechanism that is self-contained within this specific dynamic market topology, the second-price auction platform pioneered by AT&T. We focus on providing the existing auction platform with an incentive framework, and so rational users choose a collaborative exchange. In other words, adhering to the second-price rule, where price is derived from autonomous demand, we build "order" within the dynamic network of shifting demand and supply based on noncooperative, autonomous consumers. This is the (built-in) transformation from the direct-revelation mechanism to the desired message space. Our auction mechanism may be described in game-theoretical terms as a pure-strategy progressive game with incomplete, but perfect, information.

The rest of this paper is structured as follows: Section 5.2 presents the related work on auction theory and resulting policy software. Section 5.3 details the market derivation and mathematical form, which we present as an extension of the formulation found in [45]. The analysis of user behavior and the resulting algorithms that drive the noncooperative game are presented in Section 5.4 along with a simple example. The analysis of the VCG properties and the network Nash equilibria are given in Section 5.5. Conclusions and Future Work follow in Section 5.6.

5.2 Related Work

Different definitions of social welfare define different auction strategies. Typical goals of optimization are the maximization of revenue and optimal allocation. Google AdWords allows advertisers to set their own prices by using an auction system where advertisers bid on keywords to get their ads placed in Google search results [36]. According to [36], progressive pricing, when used in combination with an auction platform, is a fairer way to determine prices. This is the "smart" pricing rule. We are used to understanding prices in units of data, or rates such as units of data per hour.

If we can instead see prices in terms of a unit of value, then the price the customer pays can scale in proportion to the value demanded; this is our dynamic reserve price. Progressive second-price auctions strive for this dynamic price in different ways, some try to optimize the sellers' reserve price, or market price, as in this work. Research has been done for PSP auctions, and improvements have been made to the original work from Lazar and Semret. In [70], user strategy gives a “quantized” version of PSP, improving the rate of convergence of the game by shifting the bid price based on some threshold. Modifications to the mechanism that result in improved convergence also appear in [50], which relies on a global approximation function of demand.

Approximation of demand is a popular avenue of research for the division of data in a PSP auction. The complexity and amount of data inherent in digital data sharing creates a natural necessity; platforms must take advantage of the continuity of auctions restricted to simple sellers and buyers, as well as grid-based platforms. An approximation of the global demand function that uses a statistical approximation of the state space is derived from the theory of potential games. Potential games make use of a global strategy defined by the potential function; many companies use this as a mathematical tool to gather user data in order to further shape their market space. The idea of potential games in PSP markets was used by Zhong to coordinate the fair charging of electric vehicles in [91]; the potential game modeled the distribution of the load variance in electric vehicle charging, minimizing it as a global function with some constraints. The benefit of potential games is that, under some conditions, we are guaranteed convergence to a Nash equilibrium.

The analysis and interpretation of the data exchanged between the data-serve platform and the user quantify user value individually and strive to understand the decision process with increasing granularity. Usually, companies only know their marginal costs and can only infer user value by sample estimation. Today, with pervasive data and increasingly precise analytical capabilities, companies can derive a more precise estimate of value per user, and still maintain a zero marginal cost. Recently, using a mobile app as a platform, [94] assigns electric vehicles to charging

stations, replacing a first-come-first-serve system, and ensuring that there is enough power assigned to keep each user satisfied without overloading the supply. Importantly, [94] considers the individual valuations of the users. More generally, [74], defines an optimal strategy for buyers and brokers using a game-theoretic derivation, and further shows the existence of network-wide market equilibria obtained by the specific provisioning of this network given specific network dynamics. Indeed, various equilibria may be derived by designing the structure of provisioning in distributed systems, motivated by individual users interests. The continuity of data supplied by the customer demand allows for the continuity of change to the pricing system of the data-serve platform. By allowing a user preference to, loosely, represent a policy, we may interpret a user preference from the data exchange market as a plan to allow users to set their own policies, and rely on the existing framework to implement their preferences. Game-theoretical analysis of mobile data has been presented in [92] as a framework for mobile-data offloading.

This new type of provisioning is described as virtual elasticity in [60], a paper that has recently provided an innovative way to estimate future prices, an avenue of research that is of great interest to the marketing community. The estimation of future prices is difficult in PSP auctions, due to the dynamically changing bids. However, an estimation of future demand (price) can greatly affect the efficiency of the provisioning system, particularly with static resources, such as computer memory.

Finally, we mention that the concern of privacy is addressed in [35]. This paper addresses the problem of privacy between the web user and the advertiser, which is beyond the scope of this paper. However, we mention that the issue of privacy can change based on the platform, and the extent to which we have implemented data privacy, through anonymous bids, may not be sufficient. Our analysis is largely based on the work of [73] and his examination of second-price auctions in networked settings. In particular, we make use of the assumption of consistent bids, in the special case where consistent bidding is an optimal solution, i.e. a Nash equilibrium.

5.3 Market Formulation and Definitions

5.3.1 The Market Mechanism

We aim to design a distributed PSP auction, operating within a strategic framework that determines the bidding behavior of users in a wireless network. The auction design must meet a certain set of known criteria: (1) *truthfulness*, (2) *individual rationality/ selfishness*, (3) *social welfare maximization*, and (4) *an anonymous winning bid*. For the secondary data exchange market, we determine that the strategy space must meet additional criteria: (5) *privacy and independence from the ISP*, (6) *locally fair division*, and (7) *minimize crossover in buyer/ seller pools*. In second-price markets, the winning bid does not pay the winning bid price, but the price from *next lowest bid*. This provides the market with the property of truthfulness through incentive compatibility, meaning that a bidder will truthfully reveal its valuation of the resource. The exclusion-compensation principle, or Pareto criterion, is built into the pricing mechanism, and guarantees that at equilibrium, any change to the system would make at least one user worse-off.

Let the set of all wireless users be labeled by the index set $\mathcal{I} = \{1, \dots, I\}$. In our current formulation, we do not allow a seller to host multiple auctions, and so we may assume that data is a unary resource belonging to the seller, and identify each local auction with the index of the seller $j \in \mathcal{I}$. The bid profiles of the users are given as, $s \equiv [s_i^j]$ where $(i, j) \in \mathcal{I} \times \mathcal{I}$. We assume that all inactive bids are zeroed, i.e. if there is no interaction between two players i and j , then $(i, j) = 0$. Then, $\mathcal{I} \times \mathcal{I}$ is a matrix, with each element of the matrix representing a single buyer-seller interaction, one projective representation of the space. The matrix allows for ease in our analysis by vectorizing the space, and represents a single snapshot of a static system, all quantities and prices are fixed. We call this space S , as in standard game-theoretic notation, and so the (full) strategy space for buyer i as all possible bids at all auctions: $S_i = \Pi_{j \in \mathcal{I}} S_i^j$, and $S_{-i} = \Pi_{j \in \mathcal{I}} (\Pi_{k \neq i \in \mathcal{I}} S_k^j)$ as the associated opponent profiles.

The grid(s) of bid profiles, s , represents the statistical distribution of player types and corresponding actions, the userwise distributed state of actions in the secondary market. We will use the context of the bid to indicate the user type as well as the notation: s_i is a buyer's bid, $s_{-i} \equiv [s_1^j, \dots, s_{i-1}^j, s_{i+1}^j, \dots, s_I^j]_{j \in \mathcal{I}}$ is the profile of user i 's opponents, and s^j a seller's local auction. We describe the rules as follows:

- The bid is represented by $s_i^j = (q_i^j, p_i^j)$, indicating that i would like to buy from j a quantity q_i^j and is willing to pay a unit price p_i^j .
- The auction platform maintains and updates all bids.
- $s_i^j > 0$ represents an active bid in s , with bid, $s_i^j = (q_i^j, p_i^j)$.
- A buyer that does not submit a bid, i.e. $s_i^j = 0$, will not receive opponent profiles from seller j .
- A seller's profile $[s_i]$ is comprised of buyers in auction j , and a buyer's profile $[s^j]$ a set of auctions in which the buyer holds active bids.

We will assume that a buyer's budget is sufficient, as the alternative would be to pay a higher price to the ISP.

5.3.2 Market Incentive

As a market with perfect with incomplete information, we determine that sellers can only gain information about the market by observing buyer behavior in their local auction. Buyers are able to see the sellers reserve price for each market in which they bid. The reserve price of each auction is determined by the valuations of the buyers, that is, the seller takes a passive, or reactive role, and modifies its reserve price according to market demand. Thus the buyer is able to determine the state of the market through the reserve price of its active auctions, and so may be able to infer some behavior resulting from opponent bid profiles. Our analysis focuses the role of buyers, who are able to directly influence global market dynamics. Each buyer i will have information from each seller j , as well as opponent profiles s_{-i} , from each

auction in which it is participating. In the extreme case, where i submits bids to all auctions $j \in \mathcal{I}$, buyer i gains access all buyer profiles, $[s_1, \dots, s_I]$ for each auction j .

Define the set of sellers chosen by buyer $i \in \mathcal{I}$ as,

$$\mathcal{I}_i(n) = \arg \max_{\mathcal{I}' \subset \mathcal{I}, |\mathcal{I}'|=n} \sum_{j \in \mathcal{I}'} Q^j,$$

and similarly, for a seller $j \in \mathcal{I}$, we define the set of buyers participating in auction j as,

$$\mathcal{I}^j(m) = \arg \max_{\mathcal{I}' \subset \mathcal{I}, |\mathcal{I}'|=m} \sum_{i \in \mathcal{I}'} p_i^j,$$

where $m, n \in \mathcal{I}$.

We now must determine a function which regulates the behavior of the buyers in our dynamic market. We define an **opt-out function**, σ_i , associated with a buyer i as part of its type. Buyer i , when determining how to acquire a possible allocation a , will determine its bid quantities by,

$$\sigma_i(a) = [\sigma_i^j(a)]_{j \in \mathcal{I}}. \quad (5.1)$$

In a general sense, σ_i applies the PSP rules to our user strategy. The rules presented here incorporate the opt-out function with the auction mechanism. The market price function, P_i , for a buyer in the secondary market is:

$$\begin{aligned} P_i(z, s_{-i}) &= \sum_{j \in \mathcal{I}} \sigma_i^j \circ p_i^j(z_i^j, s_{-i}^j) \\ &= \sum_{j \in \mathcal{I}} \left(\inf \left\{ y \geq 0 : q_i^j(y, s_{-i}^j) \geq \sigma_i^j(z) \right\} \right), \end{aligned} \quad (5.2)$$

which we interpret as the aggregate of minimum prices that buyer i bids in order to obtain data amount z given opponent profile s_{-i} . The maximum available quantity of data in auction j at unit price y given s_{-i}^j is given as:

$$q_i^j(y, s_{-i}^j) = \sigma_i^j \circ q_i^j(y, s_{-i}^j) = \left[Q^j - \sum_{p_k^j > y} \sigma_k^j(a) \right]^+. \quad (5.3)$$

It follows from the upper-semicontinuity of D_i^j that for s_{-i}^j fixed, $\forall y, z \geq 0$,

$$\sigma_i^j(z) \leq \sigma_i^j \circ q_i^j(y, s_{-i}^j) \Leftrightarrow y \geq \sigma_i^j \circ p_i^j(z, s_{-i}^j). \quad (5.4)$$

The data allocation rule is a function of the local market interactions between buyers and sellers over all local auctions, as is composed with i 's opt-out value, so that for each $i \in \mathcal{I}$, the allocation from auction j is,

$$\begin{aligned} a_i^j(s) &= \sigma_i^j \circ a_i^j(s) \\ &= \min \left\{ \sigma_i^j(a), \frac{\sigma_i^j(a)}{\sum_{\substack{k \in \mathcal{I} \\ p_k^j = p_i^j}} \sigma_k^j(a)} q_i^j(p_i^j, s_{-i}^j) \right\}. \end{aligned} \quad (5.5)$$

Remark 5.1. *The bid quantity $\sigma_i^j(a)$ and the allocation a_i^j are complementary.*

Finally, we must have that the cost function.

$$c_i(s) = \sum_{j \in \mathcal{I}} p_i^j \left(a_i^j(0; s_{-i}^j) - a_i^j(s_i^j; s_{-i}^j) \right). \quad (5.6)$$

The cost to buyer i adds up the willingness of all buyers excluded by player i to pay for quantity a_i^j . i.e.

$$c_i^j(s) = \int_0^{a_i^j} p_i^j(z, s_{-i}^j) dz.$$

This is the “social opportunity cost” of the PSP pricing rule.

5.3.3 The Anonymity Problem

The PSP auction given in [45] is comprised of a set of simple and symmetric rules that closely follow market theory, and as it is distributed we require privacy to be computed on an individual basis, each user must be able to confirm its own anonymity. We describe the process as given in [19]. In general, a distributed computation, where buyer i is part of a coalition comprising auction j , is as follows:

Denoting $m_{-i} = [(s_i^j, r_i), m_1, \dots, m_n]_{k \neq i \in \mathcal{I}}$, buyer sends a message to each of its opponents, where s_i^j is i 's bid, r_i is an independent random value, and m_1, \dots, m_n the messages i has received so far. Then, all buyers are able to confirm the winning

bid s_i^* . It was proven in [19] that full privacy is not possible in a second-price auction, even if we allow partial revelation and weak coalitions. We propose anonymity, in the winning bid only.

Buyers are anonymous

In our secondary market, we have that any local auction is anonymous by definition, as a permutation of the valuations results in a permutation of allocations and prices, equivalently, exchanging the bids of two losing buyers does not change the auction's result. Formally,

Definition 5.2. (*Anonymous auction*) [19] *Given an auction j and buyers $i \in \mathcal{I}$, a protocol for computing $\max\{i \in \mathcal{I} : p_i^j \geq p_k^j \forall k \in \mathcal{I}\}$ if for all coalitions $T \subset \mathcal{I}$, any pair of inputs $x = [s_1^j, \dots, s_I^j], \xi$, so that ξ is a permutation of x , $\forall i \in T : x_i = \xi_i$, and $\max()$, and any choice of random inputs $\{r_i\}_{i \in T}$. Let $\bar{T} = T \times \mathcal{I} \setminus T$,*

$$\begin{aligned} & \Pr([x, \{r_i\}_{i \in T}]_{x \in \{r_i\}_{i \in T}}) \\ &= \Pr([\xi, \{r_i\}_{i \in T}]_{\xi \in T \times \mathcal{I} \setminus T} \mid \{r_i\}_{i \in T}), \end{aligned}$$

which states that any two inputs, the messages seen by coalition T are identically distributed.

The winning bid is trusted (anonymous) information

We claim that a buyer's trust in a local auction is fulfilled when the outcome of the auction is guaranteed to be correct, and if the winner's identity remains private information. For each local auction, we define a coalition to be the participating buyers. The winning bidder is chosen by distributed computation via homomorphic encryption. We present the Lemma in its general form,

Lemma 5.3. (*Benaloh 1987*) $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n x_i \bmod p$ is privately computable.

Thus, it is possible to anonymously compute,

$$\omega = \max([p_i^j]_{i \in \mathcal{I}^j}) = (p_1^j, p_2^j, \dots, p_n^j, \arg \max(p^j)). \quad (5.7)$$

We have the following procedure determining the winner of auction j for some fixed time in the bid progression, where $Q^j > 0$.

Algorithm 5.1 (Max bid private computation)

```

1:  $\omega \leftarrow 1, e \leftarrow 0$ 
2: while  $e \leq 1$  do
3:   for  $i \in \mathcal{I}^j$  do
4:     if  $p_i^j \leq \omega$  then
5:        $p_i^j \leftarrow 1$ 
6:     else
7:        $p_i^j \leftarrow 0$ 
8:     end if
9:   end for
10:   $e = \sum_{i \in \mathcal{I}^j} p_i^j \bmod (n + 1)$  (Lemma 5.3)
11:  for  $i \in \mathcal{I}^j$  do
12:    if  $p_i^j \geq e$  then return  $i$  (winner)
13:    end if
14:  end for
15: end while
16:

```

The winning buyer then leaves the auction, and so we have that the privacy of the winning buyer is persistent. We note that it is possible for a winner to anonymously rejoin an auction, however this does not alter our result. At time $t = 0$, a seller j entering the market will submit bid $s_k^j = (Q^j, \epsilon)$ to the public data exchange platform, and so the initial bid s_k^j , is public knowledge. The auction begins at time $t > 0$, and at $t = 0$, j will initializing its reserve price by executing a single bid iteration.

We will assume that the cost of participating in the secondary market is absorbed by the bid fee, which could represent data used in submitting bids, or a fee charged per unit of data, or a flat rate charged at the completion of the purchase. We do not model ISP revenue, but assume it may be extracted from the bid fee at $t = 0$.

The formulation is inspired to the thinnest allocation route for bandwidth given in [45]. We note that if a single seller j can satisfy i 's demand, then (5.8) reduces to the original form, defined in [73] as "a simple buyer at a single resource element".

Truthfulness (incentive compatibility)

We will prove that the dominant strategy for buyers is to submit coordinated bids, where all bids the buyer submits are equal. Our motivation for coordinated bids comes from the idea of potential games. In potential games, the incentive of all users to change strategy can be expressed as a single global function. We map the incentive of a buyer over all auctions $j \in \mathcal{I}$ to a single potential function. This is a standard method that is used often, as it simplifies the analysis of both strategy and auction design. Thus, our strategic bid is an ϵ -best response. The necessary condition of an ϵ -best reply is that the new bid price must differ from the last by at least ϵ . Now, an ϵ -best reply for user i is $p_i^* = \theta'_i(\sigma_i(a)) + \epsilon$, for a given opponent profile s_{-i} , and for each $j \in \mathcal{I}_i$. Now, as ϵ is the bid fee, we have that p_i^j is equal to the marginal valuation of player i in auction j , and so incentive compatibility holds.

5.4 Strategic Framework

5.4.1 User Valuation (Strategic Incentive).

In any market, a buyer or seller would like to obtain the maximum amount of utility possible while staying within budget. The buyer's utility maximizes the amount of data allocated by the seller, while the seller's utility maximizes the cost of the data sold. We may illustrate the resulting product space for the buyer:

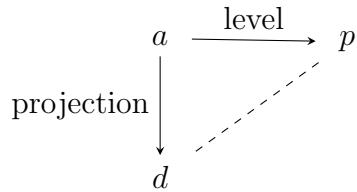


Figure 5.1: Product / Quotient (step) Space

The level, or price associated with the buyer's bid may be projected onto a line. As we will show, this holds when buyers use the same bid price for all nonzero bids. The projection, or the amount of data requested, is a plane, since buyers may bid in more

than one auction. Clearly, the allocation is the product of the price and the data request. The resulting step function is convex. We define a move to a better market position to be synonymous with a strategic bid.

Remark 5.4. *The terms “bid” and “strategy” are often interchangeable, from auction design and game theory, respectively.*

Our mechanism allows a buyer to *opt-out* of auctions by submitting zero bids. This strategy maximizes utility while minimizing the number of positive bids submitted to the overall market. We define each buyer as a user $i \in \mathcal{I}$ with quasi-linear utility function $u_i = [u_i^j]_{j \in \mathcal{I}}$, a buyers’ utility function is of the form,

$$u_i = \theta_i \circ \sigma_i(a) - c_i, \quad (5.8)$$

where the composition of the elastic valuation function θ_i with σ_i distributes a buyers’ valuation of data allocation a across local markets (and thus multiple sellers). The composition map (the codomain of $\theta_i(\sigma_i)$ is the same as the domain of $\sigma_i(a)$) and restricts the buyer’s domain to minimize $q^j p^j - c_i$ and so maximize u_i . We formally extend the PSP rules described in [73] to determine the presence of equilibria across fully connected subsets of local data-exchange markets. By fully connected, we mean that the market subset maintains its own equilibrium without the influence of any other data-exchange (any other auction). The sellers, $j \in \mathcal{I}$ are not associated with an opt-out function. The sellers’ strategy can only be to determine the reserve price of their local auction, using only information from buyers who have not opted out.

Remark 5.5. *It is possible that a seller would be able to derive information about other auctions by examining buyer bids over time, particularly if the seller had knowledge of the buyer strategy. In this work, we assume sellers are unable to derive opponent information from buyer bids.*

Elastic valuation functions allow for even infinitesimal changes in the market dynamics to be modeled. We give the definition for an elastic valuation function as in [45].

Definition 5.6. (*Elastic demand*) [45] A real valued function, $\theta(\cdot) : [0, \infty) \rightarrow [0, \infty)$, is an (elastic) valuation function on $[0, D]$ if

- $\theta(0) = 0$,
- θ is differentiable,
- $\theta' \geq 0$, and θ'_i is non-increasing and continuous,
- There exists $\gamma > 0$, such that for all $z \in [0, D]$, $\theta'(z) > 0$ implies that for all $\eta \in [0, z]$, $\theta'(z) \leq \theta'(\eta) - \gamma(z - \eta)$.

The elastic valuation of users and homogeneous nature of data in the secondary market allows for continuity in the constraints imposed by the user strategies. We begin our analysis with buyer valuation θ_i . A buyers' valuation of an amount of data represents how much a buyer is willing to pay for that amount. This is equivalent to the bid price, given a fixed amount of data, satisfying θ_i . We determine the buyers' utility-maximizing bid given quantity $z \geq 0$ to be a mapping to the lowest possible unit price. We have,

$$f_i(z) \triangleq \inf \{y \geq 0 : \rho_i(y) \geq z, \forall j \in \mathcal{I}\}, \quad (5.9)$$

where $\rho_i(y)$ represents the demand function of buyer i at bid price $y \geq 0$, and gives the quantity that buyer i would buy at a given price. We determine that the market supply function corresponds to an extreme of possible buyer demand, and acts as an “inverse” function of f_i . We have, for bid price $y \geq 0$,

$$\rho_i(y) = \sum_{j \in \mathcal{I}: p_i^j \geq y} Q^j. \quad (5.10)$$

We note that f_i is such that i could still bid in *any* auction $j \in \mathcal{I}$. Therefore, the utility-maximizing bid price is the lowest unit cost of the buyer to participate in all auctions, and corresponds to the maximum reserve price amongst the sellers.

From the perspective of the seller we have a more direct interpretation of valuation as revenue. We determine the demand function of seller j at reserve price $y \geq 0$

to be,

$$\rho^j(y) = \sum_{i \in \mathcal{I}: p_i^j \geq y} \sigma_i^j(a), \quad (5.11)$$

and define the “inverse” of the buyer demand function for seller j as potential revenue at unit price y , we have,

$$f^j(z) \triangleq \sup \{y \geq 0 : \rho^j(y) \geq z, \forall i \in \mathcal{I}\}, \quad (5.12)$$

and, unsurprisingly, f^j maps quantity z to the highest possible unit data price.

The valuation of any user must be modeled as a function of the entire market-place. Naturally, a buyers’ valuation is aggregated over local markets, and the sellers’ valuation is aggregated over its own auction. We have already introduced the composition $\theta_i \circ \sigma_i$ as the valuation of the buyers. We further show that user valuation satisfies the conditions for an elastic demand function, with valuations based on (5.11) and (5.12). We first note that, in general (and so we omit the subscript/superscript notation), the valuation of data quantity $x \geq 0$ is given by,

$$\theta(x) = \int_0^x f(z) dz,$$

as in [73]. Now, we have the following Lemma,

Lemma 5.7. (*User valuation*) For any buyer $i \in \mathcal{I}$, the valuation of a potential allocation a is,

$$\theta_i \circ \sigma_i(a) = \sum_{j \in \mathcal{I}} \int_0^{\sigma_i^j(a)} f_i(z) dz. \quad (5.13)$$

Now, we may define seller j ’s valuation in terms of revenue,

$$\theta^j = \sum_{i \in \mathcal{I}} \theta^j \circ \sigma_i^j(a) = \sum_{i \in \mathcal{I}} \int_0^{\sigma_i^j(a)} f^j(z) dz. \quad (5.14)$$

We have that θ_i and θ^j are elastic valuation functions, with derivatives θ_i' and θ^j' satisfying the conditions of elastic demand.

Proof. Let ξ be a unit of data from buyer bid quantity $\sigma_i^j(a)$. If ξ decreases by incremental amount x , then seller bid q_i^j must similarly decrease. The lost potential revenue for seller j is the price of the unit times the quantity decreased, by definition, $f^j(\xi)x$, and so,

$$\theta^j(\xi) - \theta^j(\xi - x) = f^j(\xi)x.$$

Thus (5.14) holds. As we may use the same argument for (5.13), as such, we will denote $f_i = f^j = f$ for the remainder of the proof. We observe that the function f is the first derivative of the valuation function with respect to quantity. Letting $\theta_i = \theta^j = \theta$, the existence of the derivative implies θ is continuous, and therefore, in this context, f represents the marginal valuation of the user, θ' . Also, clearly $\theta(0) = \theta(\sigma(0)) = 0$. Now, as we consider data to be an infinitely divisible resource, we have a continuous interval between allocations a and b , where $a \leq b$. Now, as θ is continuous, for some $c \in [a, b]$,

$$\theta'(c) = \lim_{x \rightarrow c} \frac{\theta(x) - \theta(c)}{x - c} = f(c),$$

and so $f = \theta'$ is continuous at $c \in [a, b]$, and so as $a \geq 0$, $\theta' \geq 0$. Finally, we have that concavity follows from the demand function. Then, as θ' is non-increasing, we may denote its derivative $\gamma \leq 0$, and taking the derivative of the Taylor approximation, we have, $\theta'(z) \leq \theta'(\eta) + \gamma(z - \eta)$. \square

The sellers' natural utility is the potential profit, or simply $u^j = \theta^j$, where we have chosen to omit the original cost of the data paid to the ISP, as it is not a component of our mechanism, and as a discussion of mobile data plans is outside the scope of this paper. Now, a rational user will try to maximize its utility, thus, user incentive manifests as a response to market dynamics. A buyer has the choice to opt-out of any auction, and as a seller will try to sell the maximum amount of data, the highest possible reserve price is conditioned by "natural" constraints. Utility-maximization acts as revenue maximization for a rational seller, and as cost minimization for a rational buyer. Thus, for each user $p_i^j \geq \min(p_i^j)$ and $p_i^j \leq \max(p_i^j)$, which holds

$\forall i, j \in \mathcal{I}$ such that $s_i^j > 0$. Now, rational buyer does not want to purchase extra data, as this would be equivalent to overpaying, however i submits positive bids to a set of sellers, and a rational seller will attempt to maximize profit, and so will try and sell all of its data. Therefore,

$$\sum_{i \in \mathcal{I}} \sigma_i^j(a) \geq Q^j \quad \text{and} \quad \sum_{j \in \mathcal{I}} q_i^j \geq Q_i, \quad (5.15)$$

which holds $\forall i, j \in \mathcal{I}$. We will assume that buyers and sellers do not overbid, and so omit this constraint from our formulation. Thus, at equilibrium all users are satisfied, and $Q^j = Q_i$, although we observe that this result does *not* imply that $s_i = s^j$.

Finally, it is worth mention that the analysis of the auction as a game assumes some forms of demand and supply, in order to derive properties. The mechanism itself does not require any knowledge of user demand or valuation.

5.4.2 User Behavior.

The user local strategy space is non-deterministic: the preferences of users are subject to change, determinations and predictions are based on the binary dependence of the variables. Arrow's Theorem states that no deterministic strategy can provide a mapping of the preferences of users into a market-wide (complete and transitive) strategy. As individual bids cannot map to a general objective, a better market position can only be determined by an adaptive strategy. We will address the market risks and securities in our secondary data exchange market, and provide a game-theoretic model of a real market progression, which we then use to derive, and then define, adaptive variables.

Assuming equal bandwidth for all users, we derive a globally optimal strategy suited for users with local information in a distributed data-sharing model. In a multi-auction market, each auction a buyer joins has the possibility of decreasing the potential cost of its data. However, increasing the size of the auction implies a certain risk, which we may interpret as a definite liability. Increasing the number of transactions causes additional messaging overhead, fees, and increased competition

from other buyers. A transaction also causes potential indirect costs, which may be considered work done to find sellers, or effort of communication from participation. A seller has the potential for greater profit with each new buyer in its auction, taking the same risk. To simplify our analysis, here the liability of any user is naturally absorbed into the bid fee ϵ , as in [73]. Therefore, according to our interpretation, the bid fee is dependent on the association between two users and their market positions, in addition to the underlying network structure. Now, both sellers and buyers must consider the cost of adding additional users to their subsequent pools.

Buyer i 's seller pool is determined by minimizing n , and is the smallest set of sellers that allows for a coordinated bid, and the aggregate bids satisfy its demand, Q_i .

$$\min \{n \in \mathcal{I} \mid nD^n \geq Q_i\}. \quad (5.16)$$

Similarly, seller j determines the minimal set of buyers that maximizes revenue and sells all of its data, Q^j .

$$\min \left\{ m \in \mathcal{I} \mid \sum_{i \in \mathcal{I}^j(m)} q_i^j \geq Q^j \right\}, \quad (5.17)$$

We further determine that the set of buyers and sellers participating in a single equilibrium is bounded by the potential indirect costs of participation. We will denote this individual cost to each user as ϱ . The indirect cost is the portion of the bid fee ϵ that is dependent on the underlying network and the individual. Observing that ϱ indirectly effects user utility, and therefore acts to establish a natural budget for each user. We define this constraint as,

$$u \leq \varrho, \quad (5.18)$$

which may be interpreted as the effort a rational user is willing to expend on its message space, and serves to limit the size of the buyer/seller pools. This information may be collected from a specific device's configuration, i.e. enabled roaming, daily data restrictions. It is clear that an unconstrained market, even with a finite number

of users, could suffer from the expense of many local auctions trading an infinitely divisible resource, thus ϱ is interpreted as the "liability" component of ϵ , and attempts to regulate network congestion.

Buyer Strategy

Although it is possible for a seller to fully satisfy a buyer i 's demand, it is also reasonable to expect that a seller may come close to using their entire data cap, and only sell the fractional overage. In this case, a buyer must split its bid among multiple sellers. The buyer strategy bids in auctions with the highest quantities first, a natural exploitation of the demand curve. A new seller entering the market with a large quantity of data will be in high demand. This behavior contributes to market price stability, as seller valuation is determined by buyer demand, the buyer strategy tends towards equal valuation of all local markets, and therefore similar prices. If a buyers' demand is not satisfied, they will need to bid in markets with smaller data quantities, and so will bid on a larger portion of the sellers' bid quantity, increasing their unit price. We define $j^* = n \leq I$ represent the seller with the least amount of data $\in \mathcal{I}_i$, i.e. $D^{j^*} \leq Q^j, \forall j \in \mathcal{I}^j$. We define the composition,

$$\sigma_i^j \circ a = \sigma_i^j(a) = \frac{a_i^j}{|\mathcal{I}_i|},$$

to be the buyer strategy with respect to quantity for all sellers $j \in \mathcal{I}_i$. We propose the following strategy.

Lemma 5.8. (*Opt-out buyer strategy*) *Let $i \in \mathcal{I}$ be a buyer and fix all other buyers' bids s_{-i} at time $t > 0$, and let a be i 's desired allocation. Define,*

$$\sigma_i^j(a) \triangleq \begin{cases} \sigma_i^{j^*}(a), & j \in \mathcal{I}^j, \\ 0, & j \not\in \mathcal{I}^j. \end{cases} \quad (5.19)$$

and bid price $p_i^j = \theta_i'(\sigma_i^j(a))$. Now, (5.19) holds $\forall j \in \mathcal{I}$.

Each time step, s^j , the vector of bids held by auction j , is updated it is shared with all participating buyers. At this point buyers have the opportunity to bid again,

where a buyer that does not bid again is assumed to hold the same bid, since a buyer dropping out of the auction will set their bid to $s_i^j = (0, 0)$.

Proof. We assume that a buyer will try and fill their data requirement. In the case that there exists a seller who can completely satisfy a buyers' demand, $j^* = 1$, $|\mathcal{I}_i| = 1$ and (5.16) holds. If such a buyer does not exist, as the set \mathcal{I}_i is ordered by the quantity of the sellers' bids, i may discover j^* by computing \mathcal{I}_i . Suppose that $Q_i > \sum_{j \in \mathcal{I}} Q^j$, then $j^* > I$ and $\mathcal{I}_i = \emptyset$. We model the ISP at time $t > 0$ as a seller κ with bid $s^\kappa = (d^\kappa, p^\kappa)$, where $d^\kappa > Q^j$, $\forall j \in \mathcal{I}_i$, and p^κ represents the price for data set by the ISP, which we note is also the upper bound of the sellers' pricing function. We note that in [95] this cost is the data overage fee. Consider some $k \neq i \in \mathcal{I}$ where $p_i^j = p_k^j$. The allocation rule (5.5) determines that the data will be split proportionally between all buyers with the same unit price. It is possible that the resulting partial allocation of data to i and k would not satisfy some demand. As the two cases i and k are the same, we will only consider one. Suppose seller j updates its bid to reflect the new data quantity, where $q_i^{j(t+1)} < \sigma_i^{j(t)}(a)$. First, i sets its bid to $s_i^j = 0$, and from the new subset \mathcal{I}_i , submits bids until $\sum_{j \in \mathcal{I}_i} \sigma(a)_i^j \geq Q_i$, by (5.15). Now, we consider the case where a new buyer k with bid price $p_k^j > p_i^j$ for some $j \in \mathcal{I}_i$, in other words, a new buyer k may enter the market with a better price, decreasing the value of i 's bid for $j \in \mathcal{I}_i$. In this case, by (5.16), i will choose \mathcal{I}_i so that, $\sigma_i^{j(t+1)}(a) = \sigma_i^{j(t)}(a) - \sigma_k^{j(t)}(a)$, and so \mathcal{I}_i is large enough to balance the additional demand from k . Finally, we consider the case where $|\mathcal{I}^j| = I$, where the demand of buyer i exceeds the supply, and the case where $\sigma_i(\varrho) > \theta_i(\sigma_i(a))$, where the overhead exceeds the current valuation of the data. Then, by (5.9), the valuation of the data increases until either the demand is satisfied, the debit from the overhead costs are balanced (5.18), or the upper bound of the sellers' reserve price p^κ is reached. Thus, as in each case we have that i is able to satisfy their demand, and we determine that the opt-out strategy is optimal. \square

Algorithm 5.2 (Buyer response)

```

1:  $p_{i(0)} \leftarrow \epsilon$ ,  $s_{i(0)} \leftarrow (p_i, Q_i)$ ,  $D_t \leftarrow Q_i$ , compute  $\mathcal{I}_{i(0)}$ 
2: Update  $s_i$ 
3: while  $Q_i(t) > 0$  do
4:    $D_{i(t+1)}^j \leftarrow \sum_{j \in \mathcal{I}_i} \sigma_i^{j(t)}(a)$ 
5:   if  $D_{i(t+1)}^j < D_t$  then
6:     Compute  $\mathcal{I}_{i(t)}$ 
7:      $p_i \leftarrow \theta_i(\sigma_i(a))$ 
8:   end if
9:    $s_{i(t+1)} \leftarrow (\sigma_i(a), p_i)$ 
10:  Update  $s_i$ 
11:   $D_{i(t+1)}^j \leftarrow D_{i(t)}^j$ 
12:   $t \leftarrow t + 1$ 
13: end while

```

Finally, we note that \mathcal{I}_i is not the only possible minimum subset $\in \mathcal{I}$ able to satisfy i 's demand, in fact, by restricting the size of the set \mathcal{I}_i , we would be able to improve the computation time of buyer i , at the cost of increasing the price.

Seller Strategy

We define the reserve price for seller j as,

$$p_*^j = p_{i^*}^j + \epsilon, \quad (5.20)$$

where i^* is the highest losing bidder with respect to bid price. We claim that the choice of reserve price p_*^j does not force any buyers out of the local auction. In order to maximize revenue, the seller must also be able to respond dynamically to strategic bids. In order to do this, we determine that the seller may modify its reserve price in response to the changing market dynamics.

Define any auction duration to be $\tau \in [0, \infty)$. We will show that sellers are able to maximize revenue in restricted subset of buyers in \mathcal{I} , and as such will attempt to facilitate a local market equilibrium for this subset. A local auction j converges when $\forall i \in \mathcal{I}$, $s_i^{j(t+1)} = s_i^{j(t)}$, at which point the allocation is stable, the data is sold, and the auction ends. In the sellers' local environment, we determine that the best

course of action is to maximize revenue, and then try to keep its buyer pool stable until convergence occurs.

Lemma 5.9. (*Localized seller strategy (i.e. progressive allocation)*) For any seller j , fix all other bids $[s_i^k]_{i,k \in \mathcal{I}}$ at time $t > 0 \in \tau$. For each $t \in \tau$, let $\omega(t)$ be given by (5.7), and perform the update,

$$D^{j(t+1)} = D^{j(t)} - \sigma_{\omega(t)}^{j(t)}(a). \quad (5.21)$$

Allowing t to range over τ , we have that $Q^j = 0$, and a local market equilibrium.

Consider a user purchasing data from a subset of other network users. The sellers' auction will function as follows: at each bid iteration all buyers submit bids, and the winning bid is the buyer i that has the highest price p_i^j . The seller allocates data to this winner, at which point all other buyers are able to bid again, and the winner leaves the auction (or equivalently, maintain their bid). The auction progresses as such until all the sellers' data has been allocated.

Proof. We assume that the seller will try to maximize its revenue. In the case where $|\mathcal{I}^j| = 1$, then if $\sigma_i^j(a) = Q^j$, then j 's market is at equilibrium. Otherwise, we arrive at the case of multiple buyers, which we note includes the case where $\sigma_i^j(a) < Q^j$, which is reflected trivially here.

For auction j with multiple buyers, i^* is the *losing* buyer with the highest unit price offer, determined by (5.17). Suppose that for some $i \in \mathcal{I}^j$, buyer demand is not met. In this case, by (5.15) the seller must notify i of a partial allocation by changing the bid vector at index i . With this caveat, and Proposition 5.8, we have that the aggregate demand of subset \mathcal{I}^j is satisfied by seller j . Although the buyers' valuation θ_i is not known to the seller, we will assume that buyers are bidding truthfully, and so the new reserve price $p_{i^*}^j + \epsilon = \theta_{i^*}' + \epsilon$. For clarity, let the reserve price be denoted by p_*^j . Now, by the elasticity of (5.9) and (5.12), we have that, $\forall z \geq 0$, $f_{i^*}(z) < f^j(z) \leq f_i(z)$, which holds $\forall i \in \mathcal{I}^j$, and $\forall j \in \mathcal{I}_i$. We claim that the choice of reserve price p_*^j does not force any buyers out of the local auction. To show this,

we use the assumption of truthful bids, and the fact that since the auction begins at time $t > 0$, buyers will bid at least once. As will be addressed in further analysis, we assume that a new bid price differs from the last bid price by at least ϵ . Suppose the auction starts at equilibrium, so $\sum_{i \in \mathcal{I}^j} \sigma_i^j(a) = Q^j$ at time $t = 0$. The reserve price p_*^j set at time $t = 0$ begins the auction with the first bid iteration, and so at $t > 0$, $\forall i \in \mathcal{I}^j$, we have that $p_i^j - p_*^j \geq \epsilon$. Now, in the case where at $t = 0$, $\sum_{i \in \mathcal{I}^j} \sigma_i^j(a) > Q^j$, by (5.5), the seller notifies (any) buyer k with the lowest bid price of a partial allocation by changing d_k^j thus by Proposition 5.8, k either decreases its demand or increases its valuation until $\sigma_i^j(a) \leq q_i^j$. Then, as the seller computes the set \mathcal{I}^j at each time step, a new i^* may be chosen and the buyers bid again. Suppose $\exists k \in \mathcal{I}^j$ such that $\forall l \in \mathcal{I}_k$, $i \in \mathcal{I}^l \forall i \neq k \in \mathcal{I}^j$. That is, k is disconnected from all other buyers $i \in \mathcal{I}^j$, and suppose that d_k^j is partial allocation at $t > 0$, and further suppose that there are many $l \in \mathcal{I}_k$ where $|\mathcal{I}^l| > |\mathcal{I}^j|$. The more buyers an auction has, the more likely that cases will occur that cause buyers to rebid, particularly if auctions $l \in \mathcal{I}_k$ have overlapping buyers, then k may opt-out of auction j , i.e. $s_k^{j(t)} \neq s_k^{j(t+1)} = 0$, then the seller may simply return the tentatively allocated data to Q^j . Finally, we note that if for some $i \in \mathcal{I}^j \exists k \in \mathcal{I}^j$ such that $p_i^j = p_k^j$, then the seller again notifies the buyers of a partial allocation by changing q_i^j and d_k^j by (5.5). Thus we determine the valuation between seller j and buyer i is well-posed, the reserve price (5.20) is justified, and we have a local equilibrium at time τ . \square

We provide a simple example.

5.4.3 A Simple Example.

Example 5.10. *Finally, we give an additional simple example of convergence to a local market equilibrium, where the buyers are assumed to respond with their truthful, ϵ -best replies.*

Name	Bid total	Unit price
A	50	1
B	40	1.2
C	26	1.5
D	20	2
E	14	2.2

Let $s^{(1)} = [(65, \epsilon)]_{i \in \mathcal{I}}$ and $s^{(2)} = [(85, \epsilon)]_{i \in \mathcal{I}}$. The buyer bids are as follows:

$$s_A = [(0, 0), (50, 1)],$$

$$s_B = [(0, 0), (40, 1.2)],$$

$$s_C = [(0, 0), (26, 1.5)],$$

$$s_D = [(0, 0), (20, 2)],$$

$$s_E = [(0, 0), (14, 2.2)].$$

Then at $t = 1$, we have bid vector $s^{(2)} = [(0, p^{(2)}), (20, p^{(2)}), (26, p^{(2)}), (20, p^{(2)}), (14, p^{(2)})]$, and so $(D^{(2)}, p^{(2)}) = (85, 1 + \epsilon)$, The buyer response is,

$$s_A = [(50, 1), (0, 0)],$$

$$s_B = [(40, 1.2), (0, 0)],$$

$$s_C = [(0, 0), (26, p^{(2)})],$$

$$s_D = [(0, 0), (20, p^{(2)})],$$

$$s_E = [(0, 0), (14, p^{(2)})].$$

At $t = 2$, $(D^{(1)}, p^{(1)}) = (65, 1 + \epsilon)$, with bid vector $s^{(1)} = [(25, p^{(1)}), (40, p^{(1)}), (0, 0), (0, 0), (0, 0)]$. $(D^{(2)}, p^{(2)}) = (25, 1 + \epsilon)$. Then,

$$s_A = [(25, p^{(1)}), (25, p^{(2)})],$$

$$s_B = [(40, p^{(1)}), (0, 0)],$$

where we have removed bids to indicate winner(s) with a tentative allocation. At $t = 3$, $(D^{(1)}, p^{(1)}) = (50, 1 + \epsilon)$, with bid vector $s^{(1)} = [(25, p^{(1)}), (40, p^{(1)}), (0, 0), (0, 0), (0, 0)]$.

$(D^{(2)}, p^{(2)}) = (0, 1 + \epsilon)$ and $s^{(2)} = [(25, p^{(1)}), (0, 0), (26, p^{(2)}), (20, p^{(2)}), (14, p^{(2)})]$. Then,

$$s_A = [(25, p^{(1)}), (0, 0)].$$

At $t = 4$ the auction ends.

Remark: In the case where market resources do not satisfy (5.15), however as this constraint is not restricted in time, we reason that in the case of insufficient data in the market buyers may wait for additional sellers or purchase from the ISP, κ , as a monopoly sale. Similarly, in the case of insufficient demand, where we may assume that data is held at time $t = 0$ by κ at bid price ϵ .

Algorithm 5.3 (Seller progressive allocation)

```

1:  $p^{j(0)} \leftarrow \epsilon$ ,  $s^{j(0)} \leftarrow (p^j, Q^j)$ ,  $\bar{\mathcal{I}} = \emptyset$ , compute  $\mathcal{I}^{j(0)}$ 
2: Update  $s^j$ 
3: while  $Q^j(t) > 0$  do
4:    $\bar{i} \leftarrow \max_{i \in I^j} \sum_{i \in I^j} p_i^j$ 
5:    $D^{j(t+1)} \leftarrow D^{j(t)} - \sigma_{\bar{i}}^{j(t)}(a)$ 
6:    $p^j \leftarrow p_{i^*}^j + \epsilon$  and  $q^j \leftarrow D^{j(t+1)}$ 
7:    $s^{j(t+1)} \leftarrow (q^j, p^j)$ 
8:   Update  $s^j$ 
9:    $\bar{\mathcal{I}} \leftarrow \bar{\mathcal{I}} \cup \bar{i}$ 
10:  for  $k \in \bar{\mathcal{I}}$  do
11:    if  $p_k^j < p_{i^*}^j$  then
12:       $D^{j(t+1)} = d_k^j$ 
13:       $\bar{\mathcal{I}} \leftarrow \bar{\mathcal{I}} \setminus \{k\}$ 
14:    end if
15:  end for
16:  Compute  $\mathcal{I}^{j(t)}$ 
17:   $\mathcal{I}^{j(t+1)} = \mathcal{I}^{j(t)} \setminus \bar{\mathcal{I}}$ 
18:   $t \leftarrow t + 1$ 
19: end while

```

Individual rationality/selfishness.

We conclude this portion by examining the relationship between the strategies of buyers and sellers in local auctions. We have proven that a buyer cannot have a

negative utility. Our strategic framework creates an incentive for the seller to maintain a local equilibrium, where supply equals demand. A truthful bid implies that the new bid price differs from the last bid price by at least ϵ . As a seller must distribute bid vectors to all buyers in its auction, we reason that the seller may employ a strategic caveat. The seller will notify a buyer who is subject to a market shift by changing its bid at the appropriate index. As we have shown, the seller is a functional extension of the buyer, with rules determined by the buyers' behavior. This gives an auction j a natural logical extension into the global market through its buyers. We demonstrate that the symmetry between buyer and seller behavior, consequently strategies, stretches into a symmetry across subsets of local auctions. Value is modeled as a function of the entire marketplace: a buyer's valuation is aggregated over all the auctions, and the seller's valuation is aggregated over its own auction. We must ensure that a user's private action satisfies the conditions of a direct-revelation mechanism, as well as adheres to the collective goals. We show that, from Lemma 5.9 and Definition 5.8, an individual user will contribute to local stability, given global market dynamics S .

We model the impact of the dynamics of S of the data-exchange market on a local auction j . The market fluctuations from S give auctioneer j the chance to infer information about the global market. We identify a clear bound restricting the range of influence that local auctions have on each other. Consider a single iteration of the auction, where a seller updates bid vector s^j , and the buyers' response s_i , to comprise a single time step. We have the following Proposition,

Proposition 5.11. *(Valuation across local auctions) For any $i, j \in \mathcal{I}$,*

$$j \in \mathcal{I}_i \Leftrightarrow i \in \mathcal{I}^j. \quad (5.22)$$

Fix an auction $j \in \mathcal{I}$ with duration τ and define the influence sets of users. The primary and secondary influencing sets are given as,

$$\Lambda = \bigcup_{i \in \mathcal{I}^j} \mathcal{I}_i, \quad \text{and} \quad \lambda = \bigcup_{i \in \mathcal{I}^j} \left(\bigcup_{k \in \mathcal{I}_i} \mathcal{I}^k \right). \quad (5.23)$$

Define $\Delta = \Lambda \cup \lambda$. Fixing all other bids $s_i^j \in \mathcal{I}$, and time $t > 0 \in \tau$, we have that,

$$\sum_{j \in \Lambda} \theta_i^j = \sum_{i \in \lambda} \theta_i^j. \quad (5.24)$$

Proof. As this is our main result, we provide an outline of the (exhaustive) proof, illustrating the most important case, when a market shifts affect auction j , and the direct influence of the shift on the connected subset of local markets.

A local auction $j \in \mathcal{I}$, is determined by the collection of buyer bid profiles. Using Lemma 5.9 and (5.22), we have that,

$$i \in \mathcal{I}^j \Leftrightarrow p_i^j > p_{i^*}^j, \quad (5.25)$$

where we define i^* as the losing buyer with the highest bid price in auction j . By (5.9) $p_i^j \geq p_{i^*}^j + \epsilon$, thus $p_i^j < p_{i^*}^j$ can only happen during a market shift. Consider $k \in \mathcal{I}^j$ at time t where, for example, some buyer(s) enter the auction, and so (5.25) implies that $\sum_{i \in \mathcal{I}^j} \sigma_i^j(a) > Q^j$. Now, $p_i^j < p_{i^*}^j \Rightarrow k \in \mathcal{I}^j$ and $s_k^j > 0$ will cause k to initiate a shift. By Definition 5.8, k will set $s_k^j = 0$, and begin to add sellers to its pool. Suppose that at time t , j 's market is at equilibrium. Unless k adds a seller with a higher reserve price within $|\mathcal{I}^j|$ time steps, by (5.21), the auction ends. We have that, $\forall i \in \mathcal{I}^j, \nexists s_i^j > 0$ where $i \in \mathcal{I}^j$, and (5.22) holds.

Now, the subset $\mathcal{I}^j \subset \mathcal{I}$ determines j 's reserve price $p_{i^*}^j$. We will assume the buyer submits a coordinated bid, using (5.5). The reserve price (5.20) of seller j is determined at each shift, and is the lowest price that j will accept to perform any allocation. Let p_*^j denote the reserve price of auction j and p_i^* denote the bid price of buyer i , i.e. $p_i^k = p_i^*$, $\forall k \in \mathcal{I}_i$. Using Lemma 5.9, for each $i \in \mathcal{I}^j$, we have from (5.9), (5.12), that $p_i^* \geq p_*^j$, $\forall k \in \mathcal{I}_i$. In the simplest case, consider a disjoint local market j , where $\forall i \in \mathcal{I}^j, s_i^k = 0, \forall k \neq j \in \mathcal{I}_i \Rightarrow \Lambda = \{j\}$ and $\lambda = \mathcal{I}^j$. Again using (5.9) and (5.12), it is clear that $\theta_i = \theta^j, \forall i \in \mathcal{I}^j$. In all other cases, the sellers $\in \Lambda$ are competing to sell their respective resources to buyers whose valuations are distributed across multiple auctions. The bid price of buyer $i \in \mathcal{I}^j$ is determined by,

$p_i^* = \max_{k \in \mathcal{I}_i} (p_*^k)$. Λ is the set of sellers directly influencing the bids of buyers in auction j . Now, the reserve price for auction j is such that, $p_*^j \leq \min_{i \in \mathcal{I}^j} (p_i^*) - \epsilon$. From (5.23), Λ is defined by a seller $j \in \mathcal{I}$, where each user $k \in \lambda$ has some direct or indirect influence on j . Denote $\Delta^j = \Lambda^j \cup \lambda^j$.

Consider the set λ^j . For some buyer $i \in \mathcal{I}^j$, and then for some seller $k \in \mathcal{I}_i$, we have a buyer $l \in \mathcal{I}^k$. By (5.22), $i, l \in \mathcal{I}^k$, and so the reserve price $p_*^k \leq \min(p_l^*, p_i^*)$, and $k, j \in \mathcal{I}_i \Rightarrow p_i^* \geq \max(p_*^k, p_*^j)$. Suppose that $l \ni \mathcal{I}^j \Leftrightarrow j \ni \mathcal{I}_l$, so that $p_l^* < p_*^j$, and the valuation of buyer l does not impact auction j and vice versa, i.e. $\theta_l^j = 0$. Since $l \in \mathcal{I}^k$, $p_l^* \geq p_*^k \Rightarrow p_*^k < p_*^j$, and $i \in \mathcal{I}^j \Rightarrow p_i^* \geq p_*^j$. Therefore, we have that the ordering implied by (5.23) holds, and,

$$p_*^k \leq p_l^* < p_*^j \leq p_i^*, \quad (5.26)$$

for any buyer $l \in \lambda^j$ such that $l \ni \mathcal{I}^j$. We use a similar argument for a secondary user $q \in \mathcal{I}_l$.

Finally, consider the subset Δ^j ; a shift occurs in 2 cases. (1) If $i \in \mathcal{I}^j$ decreases its bid quantity so that $\sum_{a \in \mathcal{I}^j} \sigma_i^j(a) < Q^j$, and (2) if buyer i^* , defined in Lemma 5.9, increases its valuation so that $p_{i^*}^j < p_*^j$. Fixing all other bids, a decrease in q 's demand will directly impact buyer i . If at the end of the bid iteration, we still have that i is the buyer with the lowest bid price, then (5.12) holds and j 's valuation does not change. Otherwise a new i^* will be chosen upon recomputing \mathcal{I}^j , as a consequence of Definition 5.8 and Lemma 5.9, and the market will attempt to regain equilibrium. We determine the influence of Δ^{k^*} on Δ^j by (5.25).

In each case we have that (5.9) and (5.12) hold for some fixed time t , and so, $\forall i \in \mathcal{I}^j$, any bid outside of our construction has a zero valuation, with respect to buyers $\in \lambda$ and sellers $\in \Lambda$, and therefore cannot cause shifts to occur except through a shared buyer, e.g. some $l \in \mathcal{I}^k$. Thus, in all cases, (5.9) and (5.12) hold. Fixing all bids in any auction where $q \ni \Delta^j$, $\forall i \in \mathcal{I}^j, \forall k \in \mathcal{I}_i, \forall l \in \mathcal{I}^k$,

$$\int_0^{\sigma_i^k(a)} f_i(z) dz = \int_0^{\sigma_i^k(a)} f^k(z) dz, \quad (5.27)$$

and

$$\int_0^{\sigma_l^k(a)} f^k(z) dz = \int_0^{\sigma_l^k(a)} f_l(z) dz. \quad (5.28)$$

Thus, with a slight abuse of notation for clarity,

$$\sum_{\lambda} \int_0^{\sigma(a)} f^{\Lambda}(z) dz = \sum_{\Lambda} \int_0^{\sigma(a)} f_{\lambda}(z) dz, \quad (5.29)$$

where the result follows by construction, and the continuity of θ' . \square

For completeness, in the case where the ISP κ does not adhere to the market dynamics, so $p^{\kappa} > p^j + \epsilon$, $\forall j \in \mathcal{I}$, then we may absorb the overage (difference) as part of the bid fee.

Locally fair division.

We claim that the allocation a by seller j for a local auction at equilibrium is an *equitable division*, a fair division where each buyer equally values their valuation. We have that equitable division holds from (5.27) Proposition 5.11.

Social welfare maximization.

We define an optimal state of social welfare to be when valuations are equal across a subset of local auctions. Then, $\Delta \subset \mathcal{I}$ to be a subset of users where an optimal social welfare is achieved.

Social welfare maximization (exclusion-compensation)

We define an optimal state of social welfare to be when valuations are equal across a subset of local auctions. Then, $\Delta \subset \mathcal{I}$ is the subset of users where social welfare is achieved. We finally have:

Corollary 5.12. (Δ -Pareto efficiency) *The subset $\Delta \subset \mathcal{I}$ is Pareto efficient, in that no user can make a strategic move without making any other user worse off.*

Proof. Define $s_* = (z_*, \theta'_*(z_*))$ as the set of truthful ϵ -best replies for user i given opponent bid profile S_{-i} , where $\forall j \in \mathcal{I}_i$, $s_*^j = s_*$. Since θ'_i is continuous, as was shown

in Lemma 5.7, and as $s|_{\Delta} = \{[s_i^j] \in \lambda^j \times \Lambda^j\}$ is continuous in s on $S_k = \Pi_{k \in \lambda^j} S_k^j$, then given that $s_* = s^* = (f^*(p^*), p^*) = (z^*, \theta'(z^*))$, we have that s^* is truthful. The result now follows directly from the result of Proposition 5.11. \square

5.5 Equilibrium Analysis

We intend to show evidence shared network optima (a global optimum). A buyer $i \in \mathcal{I}$ will have incentive to change its bid quantity if it increases its opt-out value σ_i , and therefore its utility (5.8). We will show that, without loss of utility, buyer i may use a “consistent” bid strategy within its seller pool, i.e. $q_i^j = q_i^k, \forall j, k \in \mathcal{I}_i$, and as such, Proposition 5.8 supports an optimal strategy with respect to (5.8). Our result shows that a buyer may select \mathcal{I}_i in order to maximize its utility while maintaining a coordinated bid strategy. Reasonably, if $j^* < I$, a buyer may increase the size of its seller pool \mathcal{I}_i , thereby lowering its coordinated bid quantity while obtaining the same (potential) allocation a_i . As buyer i submits identical bids to multiple auctions, the bid price must be as high as the highest reserve price $p_i^j \in \mathcal{I}_i$. Buyer i ’s bid then has identical bid price $p_i^j \forall j \in \mathcal{I}_i$. We further note that i ’s optimal strategy does not require reducing its bid price to a minimum in each auction, where the bid quantity $\sigma_i^j(a)$ is still fulfilled. The pricing rule of the PSP auction dictates that a buyer i will pay the cost of excluding other players from the auction, and as i ’s bid price reflects its valuation of its data requirement Q_i across all local markets, we have identical bid prices in each auction where $s_i^j > 0$. Obviously, if $j \notin \mathcal{I}_i$, then $\theta_i^j = 0$.

Lemma 5.13. *(Opt-out buyer coordination) Let $i \in \mathcal{I}$ be a opt-out buyer and fix all sellers’ profiles s^j . For any profile $S_i = (Q_i, P_i)$, let $a_i \equiv \sum_j a_i^j(s)$ be a tentative data allocation. For any fixed S_{-i} , a better reply for i in any auction is $x_i = \sigma_i \circ (z_i, y_i)$, where $\forall j \in \mathcal{I}_i$,*

$$z_i^j = \sigma_i^j(a),$$

$$y_i^j = \theta_i'(z_i^j).$$

Furthermore,

$$a_i^j(z_i, y_i) = z_i^j, \quad (5.30)$$

and

$$c_i^j(z_i, y_i) = y_i^j, \quad (5.31)$$

where i 's strategy is as in Proposition 5.8.

The proof follows closely the work in [73].

Proof. As s_{-i} is fixed, we omit it, in addition, we will use $u \equiv u_i \equiv u_i(s_i) \equiv u_i(s_i; s_{-i})$.

In full notation, we intend to show

$$u_i((q_i, p_i); s) \leq u_i((z_i, y_i); s_{-i}).$$

Now, if there exists a seller who can fully satisfy i 's demand, then $|\mathcal{I}_i| = 1$, and the case is trivial as no coordination is necessary for a single bid. Otherwise, buyer i 's demand can only be satisfied by purchasing data from multiple sellers. We will show that i may increase $|\mathcal{I}_i|$, and so decreasing q_i^j , $\forall j \in \mathcal{I}_i$, without decreasing $\sum_{j \in \mathcal{I}_i} u_i^j$. Buyer i maintains ordered set \mathcal{I}_i where the sellers with the largest bid quantities are considered first; the index of seller j^* defines a minimal subset \mathcal{I}_i , satisfying (5.16). By construction, $q_i^{j^*}$ is the minimum quantity bid offered by any $j \in \mathcal{I}_i$. Thus by (5.16) and (5.19), $\forall j \in \mathcal{I}_i$, $k \ni \mathcal{I}_i$, $\sigma_i^k(a) \leq z_i^j = \sigma_i^j(a)$, and so, using (5.24),

$$\sigma_i^j(a) \leq \left[D^j - \sum_{k \in \mathcal{I}^j: p_k^j > y_i^j} d_k^j \right]^+. \quad (5.32)$$

The buyer valuation function (5.13), guarantees that $\forall j \in \mathcal{I}_i$, $y_i^j \geq p_{i^*}^j$, where $p_{i^*}^j$ is the reserve price of seller j , defined in Proposition 5.9, and is by definition the minimum price for a buyer bid to be accepted. As \bar{D}_i^j is non-decreasing, $\forall j \in \mathcal{I}_i$, $k \ni \mathcal{I}_i$,

$$Q_i^j(y_i^j) \geq Q_i^j(p_i^j) \geq Q_i^j(p_i^k).$$

Thus (5.32) holds and so, by (5.5),

$$\begin{aligned} a_i^j(z_i, p_i) &= \min_{i \in \mathcal{I}^j} \left(z_i^j, \left[D^j - \sum_{p_k^j > y_i^j} d_k^j \right]^+ \right) \\ &= z_i^j = \sigma_i^j(a) \end{aligned}$$

where the last equality is by definition, and so (5.30) is proven. From (5.3), $\bar{D}_i^j(y, s_{-i}) = 0 \forall y < p_i^j$, and $\bar{D}_i^j(y, s_{-i}) = 0 \leq \epsilon \Rightarrow \sigma_i^j(a) = 0 \Rightarrow z_i^k = 0, \forall k \in \mathcal{I}_i$, and therefore,

$$\sum_{j \in \mathcal{I}_i} c_i^j(z_i, y_i) = \sum_{j \in \mathcal{I}_i} c_i^j(z_i, p_i),$$

thus (5.31) simply shows that changing the price p_i^j to y_i^j does not exclude any additional buyers, as the bid p_i^j was already above the reserve price of any seller $j \in \mathcal{I}_i$. We proceed to show that x_i does not result in a loss of utility for buyer i , that is,

$$u_i \leq u_i(z_i, y_i).$$

From (5.30), we have $a_i^j(z_i, y_i) = z_i^j = \sigma_i^j(a(z_i, y_i))$, and so,

$$\theta_i \circ \sigma_i^j(a(z_i, y_i)) = \theta_i \circ \sigma_i^j(a),$$

which holds $\forall j \in \mathcal{I}_i$. Therefore, by the definition of utility (5.8), and the buyers' valuation (5.13),

$$\begin{aligned} &\theta_i \circ \sigma_i(a(z_i, y_i)) - \theta_i(a) \circ \sigma_i(a) \\ &= u_i(z_i, y_i) - u_i = \sum_{j \in \mathcal{I}_i} c_i^j - c_i^j(z_i, y_i) \\ &= \sum_{j \in \mathcal{I}_i} \int_{a_i^j(z_i, p_i)}^{a_i^j} f_i(q_i^j - x) dx. \end{aligned}$$

Then, as $a_i(z_i, p_i) \leq z_i^j \leq a_i^j$, and noting that $z_i^j > 0 \Rightarrow \theta_i \geq 0 \Rightarrow f_i \geq 0$, we have $u_i(z_i, y_i) - u_i \geq 0, \forall j \in \mathcal{I}_i$. \square

The property of truthfulness is an essential component of equilibrium in second-price markets. The strategies described in this paper have removed the necessity for

a user to determine its own valuation function, we intend to show that the market dynamics resulting from the construction of the user strategy space results in truthful bids that are optimal for all users, i.e. bid prices are to the marginal value as determined by market dynamics. To achieve incentive compatibility, we find that the opt-out buyer must choose We have so far only made the assumption of truthful bids throughout our analysis. As was shown in Proposition 5.11, a buyer only has incentive to change its bid as a result of a market shift or partial allocation. In a truthful reply, the term $\epsilon/\theta'_i(0)$ ensures that a new bid price differs from the last bid price by at least ϵ , thereby ensuring that a buyer does not change its bid without correcting the effects of unstable shifts. For any buyer i , it suffices to show the continuity of the set of truthful ϵ -best replies in the set of opponent bid profiles. So, for a buyer i , define the set of possible ϵ -best replies,

$$\begin{aligned} S^\epsilon(s) &= \{s_i \in S_i(s_{-i}) : u(s_i; s_{-i}) \\ &\geq u_i(s'_i; s_{-i}) - \epsilon, \forall s'_i \in S_i(s_{-i})\}, \end{aligned} \tag{5.33}$$

and the set of *truthful* bids,

$$T_i = \{s_i \in S_i(s_i) : z = \sum_{j \in \mathcal{I}_i} \sigma_i^j(a) \wedge p_i = \theta'_i(z)\}, \tag{5.34}$$

where \wedge denotes the logical "and" operator. We note that the "strategic" set T_i is restricted by Proposition 5.8. We have the following Proposition,

Proposition 5.14. *(Incentive compatibility across local auctions)* Let Λ, λ be defined as in Proposition (5.11), and fix time $t > 0 \in \tau$, and fix $s^j, \forall j \in \Lambda$, and for some buyer $i \in \mathcal{I}^j$, let s_l also be fixed $\forall l \ni i \in \lambda$. Define,

$$\chi_i = \left\{ x \in [0, Q_i] : \theta'_i(x) > \max_{j \in \Lambda} P_i^j(x) \right\}, \tag{5.35}$$

and $z = \sup(\chi_i - \epsilon/\theta'_i(0))^+$, and for each $j \in \Lambda$,

$$v_i^j = \sigma_i^j(z),$$

and

$$w_i^j = \theta'_i(z).$$

Then a (coordinated) ϵ -best reply for the opt-out buyer is $t_i = (v_i, w_i) \in T_i \cap S_i^\epsilon(s_{-i})$, i.e., $\forall s_i, u_i(t_i; s_{-i}) + \epsilon \geq u_i(s_i; s_{-i})$. With reserve prices $p^j > 0$, there exists a "truthful" strategy game embedded in Δ . Therefore, a fixed point in Δ is a fixed point in the multi-auction game.

Proof. We claim that t_i is an ϵ -best reply for buyer i . That is,

$$u_i(t_i; s_{-i}) + \epsilon \geq u_i(s_i; s_{-i}).$$

As a result of auction initialization, a seller j 's valuation defines its reserve price to be determined by a buyer $i \in \lambda$, even if this price is zero, we have that $p^j = \epsilon \geq 0 \forall j \in \Lambda$. Let $z = \sup(\chi_i^j)$, and again let $p_*^j = f^j \circ \sigma_i^j(a)$ denote the reserve price of auction j , and $p_i^* = f_i \circ \sigma_i^j(a)$ denote the (coordinated) bid price of buyer i . We have that $i \in \mathcal{I}^j$, and (5.9) defines $\theta'_i(z)$ as being max of the reserve prices $p_*^j, \forall j \in \mathcal{I}_i$, therefore (5.35) is such that,

$$\theta'_i(z) > \max_{j \in \Lambda} P_i^j(v_i^j),$$

which implies, as θ'_i is non-increasing and $P_i^j \geq 0$, we have $\forall j \in \mathcal{I}_i$,

$$\begin{aligned} w_i^j &> P_i^j(v_i^j) \\ \Rightarrow v_i^j &\leq Q_i^j(w_i^j) = Q^j - \rho^j(w_i^j). \end{aligned}$$

And so, by (5.5),

$$\begin{aligned} a_i^j(t_i; s_{-i}) &= v_i^j \\ \Rightarrow \sum_{j \in \Lambda} a_i^j(t_i; s_{-i}) &= z. \end{aligned}$$

Therefore, $\forall j \in \Lambda$ and $\forall i \in \lambda$ such that (5.27) and (5.28) hold,

$$\int_0^{v_i^j} \bar{P}_i(x) dx = \sum_{j \in \Lambda} \int_0^{\sigma_i^j(z)} P_i^j(x) dx.$$

It follows that,

$$u_i(t_i; s_{-i}) = \int_0^z \theta'_i(x) dx - \sigma_i \circ \int_0^z \bar{P}_i(x) dx.$$

Suppose $\exists s_i = (q_i, p_i)$ such that $u_i^j(s_i; s_{-i}) > u_i^j(t_i; s_{-i}) + \epsilon$. Propositions 5.13 and 5.8, define the coordinated bid, $\nu_i = (\zeta_i, p_i)$, using (5.27) and (5.28), for each $j \in \Lambda$, $\sigma_i^j(a_i^j(\nu_i; s_{-i})) = \zeta_i^j$, then clearly $u_i(\nu_i, s_{-i}) \geq u_i(s_i, s_{-i}) \Rightarrow u_i(t_i; s_{-i}) - u_i(s_i; s_{-i}) > \epsilon$. Denoting ζ_i^j (fixed) as ζ ,

$$\int_z^\zeta \theta'_i(x) dx - \int_z^\zeta \bar{P}_i(x) dx > \epsilon.$$

For concave valuation functions, the first-order derivative of θ at point 0 gives the maximum slope of the valuation function, and so the factor $\epsilon/\theta'_i(0)$ guarantees that new bids will differ by at least ϵ , and as such, buyer i will remain in any local auction with reserve price determined by (5.20). We therefore verify that,

$$\int_z^{z+\epsilon/\theta'_i(0)} \theta'_i(x) dx \leq \epsilon,$$

and as $P_i^j \geq 0$, we have that, from the construction of ζ ,

$$\int_{z+\epsilon/\theta'_i(0)}^\zeta \theta'_i(x) dx - \int_{z+\epsilon/\theta'_i(0)}^\zeta \bar{P}_i(x) dx > 0.$$

If $\zeta > z + \epsilon/\theta'_i(0)$, then for some $\delta > 0$, $\theta_i(z + \epsilon/\theta'_i(0) + \delta) > P_i^j(z + \epsilon/\theta'_i(0) + \delta)$, contradicting (5.35). Now, if $\zeta \leq z$, then $\theta'_i(z + \epsilon/\theta'_i(0)) < P_i^j(z + \epsilon/\theta'_i(0))$, also a contradiction of (5.35), and so buyer s_i cannot exist. Finally, as we may consider $\Delta \subset \mathcal{I}$ to be a multi-auction game, our user strategies form a "truthful" local game with strategy space restricted to ϵ -best replies from buyers $\in \lambda$. Therefore we have that a fixed point in the "truthful" game is a fixed point for the auction. \square

In linear analysis, we may determine a Nash equilibrium by finding a local optima of the potential function. Additionally, as the potential function also iterates, it may be used in an analysis of convergence. The convergence of a Nash equilibrium results from the progression of ϵ -best replies, where each subsequent bid is a unilateral improvement, provided that t_i is continuous in opponent profiles. From the original

proof by [45], we observe that the collection of unconstrained truthful bids may be a subset of the collection of ϵ -best replies, i.e. $T_i \subset S_i^\epsilon$.

Now, the strategy space is comprised of a collection of bid, or "strategy", vectors that together, may be represented as a collection of potential functions, where change in buyer i 's utility, resulting from a change in strategy, equals the change in the local market objective of each seller $j \in \mathcal{I}_i$. These local objectives are known as potential functions, and are formulated by mapping the incentives of all users in a local auction to a single function. The goal of our analysis is to therefore construct a global potential function that encompasses all local markets, and show that this space adheres to the construction described in our proof. The conditions of convexity, connectedness and continuity must apply to the global market space in order for a global equilibrium to exist.

In order to address continuity in a global sense, we must again demonstrate continuity in the construction of our model. We will show that our global market holds a differentiable topology, where our opt-out function σ extends to an injective, differentiable map. We show that locally, the connectivity of our market subspace Δ provides a linearization, or approximation of a linear map, and so continuity holds in the global sense. We construct an extension and determine the existence and uniqueness of a global market objective by mathematical correspondence. We begin with the definition of correspondence,

Definition 5.15. *(Correspondence) A correspondence is mathematically defined as an ordered triple (X, Y, R) , where R is a relation from X to Y , i.e. any subset of the Cartesian product $X \times Y$.*

In an economic model, a correspondence (S_i, S_{-i}, R) defines a map from S_i to the power set S_{-i} , where R is a binary relation, i.e. $R \subset S_i \times S_{-i}$. The classic example of a correspondence in our model is the buyers' best response B_i^ϵ , where, for the multi-auction, S_i and S_{-i} are built by repeatedly using the Cartesian product over bid profiles. The power set $S_{-i} = \Pi_j(\Pi_{k \neq i} S_i^j)$ arises naturally from the product of ordered

sets. The binary equality relation $i \sim j$ naturally occurs in the strategy space, and is both an equivalence relation and a partial order, and therefore is reflexive, transitive, symmetric and anti-symmetric. We use the the axiom of set equality based on first-order logic, which states that, $\forall i \in \mathcal{I}, \forall j \in \mathcal{I}, (i \in \mathcal{I}^j \Leftrightarrow j \in \mathcal{I}_i) \Rightarrow i \sim j$, and follows from (5.22). Given any set of buyers $\mathcal{I}_i \in \Lambda$ where we have an allocation from some $j \in \lambda_i$, $1_\lambda(S_i, S_{-i}) : \lambda \rightarrow \Lambda \in S / \sim$, and so is a canonical mapping as well as an inclusion map. The product topology of the strategy space is preserved, and the set of all indicator functions on S forms the power set $\mathcal{P}(S) = S_i \times S_{-i}$ on S defines a quotient space, and forms the partition $\{s^j \in S : s^j \sim s_i\}$ of S . Now, $\Delta \subset \mathcal{P}(S)$ is the result of the correspondence map, and we have that set of users in an auction is uniquely determined by its members where sellers have fixed market prices; all users who are not changing their bids are considered equal. Therefore each seller $j \in \mathcal{I}$ is equivalent to some buyer $i \in \mathcal{I}$; buyer i 's utility constraint is satisfied in auction j if and only if seller j 's utility constraint is also satisfied.

The best response is a reaction correspondence defined by the mixed-strategy game. Denoting $T_i^\epsilon = T_i \cap B_i^\epsilon$, we have the set of truthful ϵ -best replies in opponent bid profiles S_{-i} . A natural induced topology of this space is the product topology, e.g. the canonical map $S_i \rightarrow \prod_{j \in \mathcal{I}} S^j$. Now, in order to find a fixed point in the mixed strategy space, we must have a continuous mapping, i.e. $\Lambda^j \Leftrightarrow \lambda_i, \forall i, j \in \Delta$. The data-sharing market consists of inter-dependent sets of multi-auction games around possible fixed points. Clearly, the union of all possible sets $\bigcup_{j \in \mathcal{I}} \Delta^j$ covers \mathcal{I} . We claim that the shared buyers between the different subsets Δ form a sufficiently connected set, so that Proposition 5.11 holds. We have the following Lemma.

Lemma 5.16. *(Continuity of ϵ -best reply on Δ) Let Δ be defined as in Proposition (5.11). For any buyer $i \in \lambda^j$, the collection of bids B_i is continuous in S_{-i}*

Proof. Define $\sigma_i \circ \bar{P} = \max_{i \in \mathcal{I}^j} \theta'_i(0)$, and $\bar{P}_i(z, s_i) = \underline{P} = \epsilon - \varrho$, where ϵ is the bid fee, and ϱ is i 's liability estimate for auction $j \in \mathcal{I}$. We observe that $\sigma_i \circ B_i^\epsilon$ is simply B_i^ϵ restricted to seller pool \mathcal{I}_i , i.e. $\sigma_i \circ B_i^\epsilon \equiv B_i^\epsilon|_{\mathcal{I}_i}$. Thus, we have $\sigma_i \circ T_i =$

$([0, Q^k]_{k \in \mathcal{I}^j} \times [0, \sigma_i \circ \bar{P}]^{|\mathcal{I}^j|})$ is a product of closed subsets of compact sets. Now, we have that a closed subset of a compact set is compact and the resulting product topology gives Tychonoff's theorem, i.e. every product of a compact space is compact, we have $\sigma_i \circ B_i^\epsilon$ is compact subset of B_i^ϵ . Now, letting $\bar{P} = \max_{i \in \Lambda^j} \theta'_i(0)$, and we have by definition of Δ and the product,

$$\begin{aligned} \sigma_i \circ S_i(s_{-i}) &\equiv \sigma_i|_{\Lambda^j} : S_i \mapsto T_i \subset S_i \\ &\Rightarrow \left(\bigcup_{i \in \mathcal{I}^j} [0, Q^k]_{k \in \mathcal{I}_i}, [0, \bar{P}] \right) = \bigcup_{i \in \mathcal{I}^j} \left([0, Q^k]_{k \in \mathcal{I}_i} \times [0, \bar{P}] \right) \\ &= ([0, Q^k]_{k \in \Lambda^j}, [0, \bar{P}]) \in \Lambda^j \times \lambda^j \subset T. \end{aligned}$$

The result follows from the fact that t_i is continuous in s_i , as was proven in [73], and as a finite union of compact sets is a compact set. \square

We have proven that buyers will submit bids according to their marginal valuations. We have that all bids represent ϵ -best replies, and, as was proven in [45]. The sellers' positive reserve price implies that bids are truthful. Finally, by properties determined by the construction of a mixed strategy symmetric game with a 2-dimensional message space, we may now restrict our analysis to the set of continuous, truthful, ϵ -best replies, T^ϵ . In mathematics, the notion of the continuity of functions is not immediately extensible to multivalued mappings; we show the correspondences between the two sets λ and Λ . The correspondence between i and j forms the set $\lambda \sim \Lambda$. We note that due to the binary relation, the set of all possible ϵ -best replies,

$$\Delta^\epsilon = \{(i, j) \in \lambda \times \Lambda\}|_{T^\epsilon},$$

is well-posed by Hadamard [1923], definition (4.3) and Corollary (4.4). We show that our bidding strategy results in (at least one), Nash equilibrium, where again the sellers reserve prices are fixed.

Lemma 5.17. (Δ^ϵ -Nash Equilibrium) *Let Δ be defined as in Proposition (5.11), and suppose that auction $j \in \Delta$ is not in a transient state, e.g. $t = \tau^j$. Fix all $s_i^k, k \neq j$.*

Using the rules of the data auction mechanism, along with type-based strategic moves, j converges to an ϵ -Nash equilibrium. The proof follows closely that of [73].

Proof. As auction j is at equilibrium, and since θ'_i is continuous, as was shown in Lemma 5.7, and $t = \{[t_i^j] \in \lambda^j \times \Lambda^j\}$ is continuous in s on $T_k = \Pi_{k \in \lambda^j} T_k^j$. Now, t represents a continuous mapping of $[0, \sum_{k \in \lambda^j} Q^k]_{i \in \Lambda^j}$ onto itself and we may use Brouwer's fixed point theorem, as in [73], which states that the continuous mapping of a convex compact set into itself has at least one fixed point. Therefore, \exists some $k \neq i$ such that $z^* = \sigma^*(z) \in [0, D_k]_{i \in \Lambda^j}$. Then, given that $s^* = (z^*, \theta'(z^*))$, we have that $s^* = t(s^*) \in T$. \square

The rules of the PSP multi-auction drive market mutations that evolve and are regulated by the user strategies. As a result of user behavior, and subsequent strategies, we determine that the data-exchange market behaves in a predictable way. We point out the need for better management of data on the consumer level. It is obvious that there is profit to be made by supplying data to the data-driven consumer. Mathematically, we have shown that if truthfulness holds locally for both buyers and sellers, i.e. $p_i = \theta'_i, \forall j \in \mathcal{I}_i$ and $p^j = \theta^{j'}$, $\forall i \in \mathcal{I}^j$, then, in the absence of market shifts, there exists an ϵ -Nash equilibrium extending over a subset of connected local markets. However, each auction may be played on the same or on a different scale in valuation, time and quantity, and so the rate at which market fluctuations occur is impossible to predict. This presents a problem, as in our linear analysis we rely on the stability of the market equilibrium at a fixed time to find a convergent sequence of ϵ -best replies within any auction j , whereas in the global market discontinuities may occur when we have $\mathcal{I}_i \cap \mathcal{I}^j = \emptyset$. In this case, we must address the market using a non-linear analysis. Up to this point, we have constructed our proofs around connected local markets, such as in 5.11, where we defined connectivity via a set of influencing users Λ and λ . The result was the existence of a sequence of vector-valued functions on the union of the influencing sets, Δ , allowing for the requirements of differentiability and therefore continuity to hold, resulting in a fully connected subset

of local auctions.

For each Δ , the Kuhn-Tucker optimality conditions imply that s_i is the optimal response of player i to s_{-i} if and only if there exists a Lagrange multiplier ρ_i such that:

$$\rho_i = \theta'(s_i^j), \text{ if } s_i^j > 0, \quad j \in \Delta \quad (5.36)$$

$$\rho_i \leq \theta'(s_i^j), \text{ if } s_i^j = 0, \quad j \in \Delta \quad (5.37)$$

$$\sum_{j \in \Delta} s_i^j = a_i, \quad s_i^j \geq 0, \quad i \in \Delta. \quad (5.38)$$

where ρ turns out, in fact, to be the (stable) marginal price [41]. Without loss of generality, we will consider adding data in such a way that the ordering of the sets λ and Λ is preserved. That is, the buyers bids are such that $s_1 \geq s_2 \geq \dots \geq s_I$. We may even define a function $\mathcal{N} : \Delta \rightarrow \mathbb{R}^S$ that assigns each $s \in S_\Delta$ the Nash equilibrium $\mathcal{N}(s)$ of its respective game. Each assignment induces a game with a unique Nash equilibria. We consider disjoint sets $\{\Delta\}$, we construct an extension so that for any influencer $i \in \Delta^e$, there is an extension such that for $k \in \Delta^k$ such that the dominant strategy for $i \in \Delta^{j \times k}$.

We construct our extension in the form of a new user type, a *broker* type. A broker type fills the space between Δ^j and Δ^k by purchasing data from one auction subset and selling it in the other. This user performs the function of connecting two fully-connected auction subsets Δ^j and Δ^k by supplying additional data from one auction (Δ^j) to the ordered set \mathcal{I}^j of each seller in each auction $j \in \Delta^k$. We show that this additional broker type preserves the optimality of the set $\Delta^j \cup \Delta^k$. Suppose that the broker assumes that there is an infinite amount of data available to buy and creates bids on the assumption that a market Δ will be available to fill the data request. The broker may create orders that are not feasible in the actual market, causing buyers

in Δ to shift their bids according to the (presumed) available data. In this way, the broker may actually plan the data requirements, overage or underage, based on some finite scheduling scheme. This concept is beyond the scope of our current research, but merits some future consideration. In this case, the broker can only add data δ to Δ from the set $\{d \in \mathbb{R}_+^I : s_1 \geq \dots \geq s_I; (\delta_i, \cdot) \geq (q_i, \cdot); \sum(q_i - \delta_i) \leq \delta\}$. We are presented with the problem of how to add additional data to Δ that is optimal with respect to its Nash equilibrium \mathcal{N}_Δ , as in [42]. For each optimal strategy s , the unique Nash equilibria $\mathcal{N}_\Delta(s)$ describes the allocations of data Q^j from each seller j in Δ .

We begin by examining the addition of data to a single market subset. We will show that this strategy, which we will call s^* , is therefore userwise price optimal for the entire space $\mathcal{N}(s) \forall s \in S$. We show that, under certain conditions, the transfer of data happens in an "ordered" way, so that the natural price ordering of the space is preserved and thus, the Nash equilibrium. In particular, there exists a Lagrange multiplier ρ^* such that (5.36) - (5.38) hold.

Theorem 5.18. (Δ_ϵ^* -Nash Equilibrium)

Let $\Delta^* = \Delta^j \cup \Delta^k$ be a union of market subsets at equilibria. Let $\mathcal{A}_\Delta^*(s)$ be an allocation where a bids $\hat{s}^j = (p^j, q^j)$, are augmented such that $\hat{q}^j = q^j - \delta_i$ and δ_i is such that $\sum_{j \in \Delta^j} \hat{a}_i^j = \sum_{k \in \Delta^k} \hat{a}_i^k$. Then,

$$\hat{p}_*^j < p_*^j \in \Lambda, \quad (5.39)$$

and,

$$\hat{p}_i^* > p_i^* \text{ for all } i \in \lambda. \quad (5.40)$$

The transfer of data δ from influencing set Λ^j to Λ^k preserves the ordering $\hat{p}_*^j \leq \hat{p}_l^* < \hat{p}_*^k \leq \hat{p}_i^*$ for all users $k, j \in \mathcal{I}_i$, $i, l \in \mathcal{I}^j$, $i \in \mathcal{I}^k$ such that $l \in \Lambda^k, l \in \mathcal{I}^k$.

Proof. We speculate that two things are going to happen with this change in allocation. 1. Buyers in Δ^j will no longer have enough data to satisfy the bid requirement of their current auction, and will increase their reserve price. 2. Sellers in Δ^k will

have additional data to sell, and so will lower their reserve price. We determine the influence of the modified allocation by (5.25) and (5.20). Without loss of generality, suppose that a broker purchases data from auction $j \in \Delta^j$, selling it in as auction $k \in \Delta^k$. This may happen for any pair of auctions in which the sellers reserve prices differ by more than ϵ , that is $p_*^k \geq p_*^j + \epsilon$ for some $k \in \Delta^k$ and for some $j \in \Delta^j$. In this way the broker may make a profit. We have that $\hat{p}_*^j > p_*^j \in \Lambda^j$ as the losing buyer with the highest bid price changes; less data in auction j will push buyers out of the auction, as seller j increases its price according to (5.20). We will call this set of buyers \hat{I} . Now buyers in \hat{I} are making consistent bids, and so must increase their bid uniformly. The set Λ^j does not have enough data to satisfy the demands of all the sellers λ^j , and so bids must be made that include sellers from Λ^k , making Δ^* a connected set (noting the allocation restriction of two sets in the premise). The sets Δ^j and Δ^k must be connected through auction \hat{k} , as buyers in $\hat{\mathcal{I}}$ will need to bid in auction \hat{k} in order to satisfy their data requirement. This, in effect, adds auction \hat{k} to Δ^j , along with the corresponding influencers. As buyers bid consistently, buyers from $\hat{\mathcal{I}}$ will bid the same in all auctions from Δ^j , now including auction \hat{k} .

We address the buyer side. The bid price of buyer $i \in \mathcal{I}^j$ is determined by $p_i^* = \max_{j \in \mathcal{I}_i}(p_*^j)$. For some buyer $i \in \mathcal{I}^j$, and then for some seller $k \in \mathcal{I}_i$, we have a buyer $\hat{l} \in \mathcal{I}^k$. By (5.22), $i, \hat{l} \in \mathcal{I}^k$, and so the reserve price $p_*^k \leq \min(p_{\hat{l}}^*, p_i^*)$, and $k, j \in \mathcal{I}_i \Rightarrow p_i^* \geq \max(p_*^k, p_*^j)$. With the addition data to auction k , we now have that $\hat{p}_i^* \geq \max(\hat{p}_*^k, \hat{p}_*^j)$, and $\hat{p}_i^* > p_i^*$ by (5.39). Seller k , now adding data to Δ^k , must bring new buyers in, changing its reserve price according to (5.20). Seller k will choose its reserve price to compete with sellers Λ^k . As Δ^k is a fully connected set at equilibrium, as defined in Proposition 5.11, the reserve price for auction \hat{k} must be set to $\max_{j \in \Lambda^k}(p_{i^*}^j \mid \sum_{i \in I^k} a_i^k = Q^k)$. As we assume that the broker would not act without gaining a profit, we must have that the reserve price of auction \hat{k} is at least ϵ higher than that of auction j . Again, for all $i \in \hat{I}$, $\hat{p}_i^* > p_i^*$.

Now suppose that $\hat{l} \in \hat{\mathcal{I}} \Leftrightarrow \hat{k} \in \mathcal{I}_l$, the valuation of buyer \hat{l} does not impact auction \hat{k} and vice versa, i.e. $\theta_{\hat{l}}^k = 0$. We must have that $\hat{p}_l^* \geq p_j^*$ and $\hat{p}_l^* < \hat{p}_*^k$. Since

$l \in \mathcal{I}^j$, $p_l^* \geq p_*^j \Rightarrow p_*^j < \hat{p}_*^k$, and $i \in \mathcal{I}^k \Rightarrow p_i^* \geq \hat{p}_*^k$. Therefore, we have that the ordering implied by (5.23) holds, and,

$$\hat{p}_*^j \leq \hat{p}_l^* < \hat{p}_*^k \leq \hat{p}_i^*, \quad (5.41)$$

for any buyer $\hat{l} \in \lambda^k$ such that $\hat{l} \ni \mathcal{I}^k$. Market shifts will occur due to the new reserve prices chosen in auction \hat{k} according to Proposition 5.11 until the market reaches equilibrium, and so (5.41) holds \Rightarrow (5.39) holds. In effect, the transfer of data has forced the reserve prices of Λ^j and Λ^k to "squeeze" together, with the broker profiting off of the difference. Sellers in Λ^j lower their reserve prices, according to the lower demand. At the same time, buyers in λ^k lower their bid prices according to the increased supply of data. The sets Λ^k and Λ^j are connected via auction \hat{k} , and so the reserve price determined in auction \hat{k} will affect Λ^j through the bid price of buyers in $\hat{\mathcal{I}}$. We have, by transfer of data, for all $i \in \Delta^*$,

$$\int_0^{\sigma_i^j(a)} \theta_i'(x) \, dx - \int_0^{\sigma_i^j(a) - \hat{\sigma}_i^k(a)} \bar{P}_i(x) \, dx < \epsilon,$$

and,

$$\int_0^{\sigma_i^k(a)} \theta_i'(x) \, dx - \int_0^{\sigma_i^k(a) + \hat{\sigma}_i^k(a)} \bar{P}_i(x) \, dx < \epsilon.$$

Using Proposition 5.11, we may conclude that equilibrium is achieved and so (5.36) - (5.38) hold for the sets Δ^j and Δ^k connected through auction \hat{k} , and we have the existence of at least one Nash equilibrium in Δ^* . \square

In this scenario, the presence of a broker causes the two sets Δ^j and Δ^k to become connected, arriving at a Nash equilibrium for $\Delta^* = \Delta^k \cup \Delta^j$. These connected sets are defined by the influencing users around the auctions k and j , and are dynamically defined as such. This complicates the analysis and makes it difficult to determine stability in time. By restricting the transfer of data to two auctions within two auction subsets, we manage to create some sort of structure in the underlying market dynamics that is intuitively simple, however analytically difficult to describe.

5.6 Conclusion and Future Work

Mathematically, we have shown that if truthfulness holds locally for both buyers and sellers, i.e. $p_i = \theta'_i$, $\forall j \in \mathcal{I}_i$ and $p^j = \theta^{j'}$, $\forall i \in \mathcal{I}^j$, then, in the absence of market shifts, there exists an ϵ -Nash equilibrium extending over a subset of connected local markets. We have provided the analysis for a network that is operated according to a game theoretic paradigm, so that its Nash equilibrium upholds the requirements of a second-price auction, showing characteristics of efficiency, truthfulness and rationality with respect to certain system-wide criteria. We have focused on Nash equilibria whose uniqueness has been established, such as those for users with consistent bids [73]. We show that $s|_{\Delta}$ represents a continuous mapping $[0, \sum_{k \in \Lambda^j} Q^k]_{i \in \Lambda^j}$ onto itself, and that the continuous mapping of the convex compact set s_* into itself (s^*) has at least one fixed point. We show that the symmetry built into strategy space provides built-in conditions for convergence and stability of a ϵ -Nash equilibrium over pairwise connected subsets Δ^* .

The dynamics of the system with the inclusion of brokers provides an interesting direction for future research. We speculate that certain network-wide objectives may be achieved (such as stability, bandwidth regulation, throttling) through the use of brokers. The brokers would exercise a type of feedback control, both a priori (in the static analysis) or in time, in order to maintain a desired network topology; one with a stable network-wide equilibrium. Here we have examined the relation between two connected auction subsets under allocation constraints. We expect that with the addition of more players the strategy space will begin to suffer "the curse of dimensionality", rendering the analytic techniques we have used here (based on order and continuity) ineffective.

We would like to begin real-world simulations of the results presented in this paper, and extend our theory to include some practical, statistical analysis. In particular, we would like to add queueing theory and an underlying Poisson arrival process or Brownian motion [73], in order to add practical structure and timing to our

game. This practical experience may give us insight as to the breadth and variation of the game according to the user strategy of truthfulness. An alternative method to determining Nash equilibria in the higher dimensional strategy space (with more users) could be found, and the nature of the space could be described using simulated results.

Chapter 6

Bipartiteness in Progressive Second-Price Multi-Auction Networks with Perfect Substitute

Abstract

We consider a bipartite network of buyers and sellers, where the sellers run locally independent Progressive Second-Price (PSP) auctions, and buyers may participate in multiple auctions, forming a multi-auction market with perfect substitute. The paper develops a projection-based influence framework for decentralized PSP auctions. We formalize primary and expanded influence sets using projections on the active bid index set and show how partial orders on bid prices govern allocation, market shifts, and the emergence of saturated one-hop shells. Our results highlight the robustness of PSP auctions in decentralized environments by introducing saturated components and a structured framework for phase transitions in multi-auction dynamics. This structure ensures deterministic coverage of the strategy space, enabling stable and truthful embedding in the larger game. We further model intra-round dynamics using an index τ_k to capture coordinated asynchronous seller updates coupled through buyers' joint constraints. Together, these constructions explain how local interactions propagate across auctions and gives premise for coherent equilibria—without requiring global information or centralized control.

6.1 Introduction

The Progressive Second Price (PSP) auction, introduced by Lazar and Semret [45], and later expanded upon in Semret’s dissertation [73], presented a full theoretical framework for distributed resource pricing and demonstrated the linkage between PSP and VCG-type efficiency results. The PSP auction is a decentralized mechanism characterized by truthfulness, individual rationality, and social welfare maximization. Unlike traditional centralized auctions, PSP allows buyers and sellers to iteratively interact through local bidding rounds, dynamically allocating consumable resources such as network bandwidth and other communicative and computational resources. In PSP auctions, winners pay a cost determined by the externality that is imposed on others, calculated from the distribution of allocations and the bid, This ensures truthful reporting of valuation through incentive compatibility as was shown in the foundational work of Vickrey, Clarke, and Groves [83, 24, 32]. The resulting equilibria adhere to the exclusion-compensation principle, preventing unilateral improvement without harming another participant.

Our focus is on developing adaptive auction mechanisms, like the Progressive Second Price (PSP) auction, that respond to market dynamics by allowing agents to adjust their bids based on local information gathered from their network neighbors. This motivates the study of influence sets, dynamic participation, and the role of network effects in shaping bidding behavior. In these settings, agents lack full market information and are affected by network dependencies.

Maillé et al. [26] build directly on Lazar and Semret’s 1999 PSP framework by addressing the one remaining free parameter in the model — the reserve price. They demonstrate that while PSP guarantees convergence, efficiency, and incentive compatibility, the seller’s reserve price can be optimized by simple numerical methods, allowing PSP markets to balance efficiency with revenue maximization.

These iterative updates operate as strategic interactions in a decentralized framework, where the PSP auction converges, perhaps astonishingly, deterministically to

an ε -Nash equilibrium. This has been shown to be true on the networks of 20 years ago, when bandwidth and bandwidth allocation was perhaps a different game. A real-world, modern network faces significant obstacles; it is a game of partial information played in a web of interconnected decisions, dynamic participation, and evolving market constraints. This motivates a graph-theoretic treatment of information flow and motivates the introduction of the concept of market saturation.

The structure of this paper is as follows. Section 6.2 introduces the foundations necessary to model influence propagation in decentralized auctions. In Section 6.3, we present the Progressive Second Price (PSP) auction mechanism, outlining its bidding rules, participation logic, and price allocation behavior. Section 6.4 defines and explores the dynamics of influence sets, establishing a framework for analyzing how strategy updates propagate across the market. Our approach adopts and extends these concepts through graph-based methods, specifically leveraging the bipartite graph to systematically represent buyer–seller interactions. Section 6.5 introduces the concept of saturation as the limit of influence propagation, characterizing a locally evolving equilibrium structure. The simulation framework and implementation are discussed in Section 6.6, and this paper’s conclusion and future work are presented in Section 6.7.

6.2 Background and Related Work

This paper introduces a graph-based analytical framework to examine the dynamics of Progressive Second Price (PSP) auctions within decentralized market structures. Our approach builds on foundational concepts in auction theory, network influence propagation, and graph analysis, while situating the PSP model among several related domains.

The Progressive Second Price (PSP) auction, initially proposed by Lazar and Semret [45], extends classical second-price mechanisms [83, 24, 32] into decentralized contexts. Earlier studies such as Maille and Tuffin [51] and Semret’s dissertation [73] provided a full system-level model of distributed market control and the theoretical grounding for the PSP auction mechanism, analyzing network-based PSP equilibria

and pricing strategies. Subsequent formal analyses such as Qu, Jia, and Caines [70] presented key results on the Uniformly Quantized PSP (UQ-PSP) mechanism, showing that it guarantees convergence to a unique limit price independent of initial conditions, achieves γ -incentive compatibility, and extends naturally to network topologies where equilibria depend on local information exchange. Their framework provided the first rigorous quantized extension of the PSP model, establishing discrete convergence proofs that later generalizations such as those developed in this work, Qu, Jia, and Caines [69] further extended these results to networked PSP convergence, introducing asynchronous coordination and bounded-delay convergence. Subsequent work has investigated distributed or multi-resource variants, including privacy-preserving and differential frameworks in data and spectrum markets [19, 86], expanding PSP-like mechanisms to new computational settings.

Local coordination rules, when combined with bounded delays and limited information exchange, can achieve global properties similar to those in consensus and averaging protocols. Beyond traditional equilibrium analysis, distributed consensus and coordination models offer insight into asynchronous bidding and update rules. Aguilera and Toueg [3] and Lynch [48] describe protocols ensuring eventual consistency under partial information, concepts that are applicable to asynchronous PSP updates. These works demonstrate that convergent systems operating under bounded delay result in deterministic convergence guarantees in decentralized markets.

In decentralized markets, agents' strategies depend on local interactions but propagate indirectly through shared participation and local coordination. This connects PSP analysis to the broader literature on influence diffusion and cascading behavior, as in Kleinberg [39], Oki et al. [65], and Osvaldo and Queen [4], which examine network-driven contagion and adaptive decision processes. The theoretical foundation of influence sets also aligns with the study of sphere-of-influence graphs [55, 81] and dynamic graph structures that represent iterative strategic dependencies.

Graph-based approaches are central to understanding multi-agent optimization. Baur [9], Barrett [8], and related work on planar and dynamic graphs illustrate how

reachability, closure, and resistance distance can capture evolving connectivity. In the PSP context, our use of projection operators extends these methods by linking graph reachability to economic stability, enabling a deterministic interpretation of market influence propagation.

This work examines decentralized auction theory, distributed coordination, influence propagation, and graph-theoretic modeling to provide a coherent analytical framework for PSP auctions. This expanded foundation motivates the later sections on local saturation and asynchronous market dynamics.

6.3 The PSP Auction Mechanism

The Progressive Second Price (PSP) auction is a decentralized mechanism in which buyers iteratively submit bids to sellers, and sellers update reserve prices based on received bids. Each auction operates locally, and coordination emerges through repeated interactions across the market graph. The mechanism rules first appears in [45], defining the bid structure, auction dynamics, pricing rules, allocation strategies, and participation behavior. In what follows, we define the bid structure, auction dynamics, pricing rules, allocation strategies, and participation behavior that govern the PSP mechanism. Let $\mathcal{I} = \mathcal{B} \cup \mathcal{L}$ denote the set of all agents, partitioned into buyers and sellers. Each seller $j \in \mathcal{L}$ manages a local auction for a divisible resource, and each buyer $i \in \mathcal{B}$ may submit bids to a subset of sellers. The bid profile of auction j is given by the column vector s^j with entries s_i^j , where $(i, j) \in \mathcal{B} \times \mathcal{L}$. A bid

$$s_i^j = (q_i^j, p_i^j) \in S_i^j = [0, Q^j] \times [0, \infty)$$

represents a single interaction between buyer i and seller j , where q_i^j is the quantity requested by the buyer and p_i^j is the unit price offered.

In decentralized markets governed by distributed Progressive Second Price (PSP) auctions, agents submit bids in the form of price-quantity pairs at discrete time steps. These bids are locally observable: buyers receive feedback from auctions in which they participate, and sellers observe aggregate demand over time. However,

the global structure of the market—including overlapping buyer influence, competition externalities, and inferred network effects—must be reconstructed from these partial, temporally indexed signals.

Table 6.1: Basic sets and notation for a bundle of J independent PSP auctions

Object	Single auction j	Across all auctions
quantity	Q^j	(Q^1, \dots, Q^J)
Player i 's bid pair	$s_i^j = (q_i^j, p_i^j)$	$s_i = (s_i^1, \dots, s_i^J)$
Strategy space of player i	$S_i^j = [0, Q^j] \times [0, \infty)$	$S_i = \prod_{j=1}^J S_i^j$
<i>Opposing bids w.r.t. player i</i>	$s_{-i}^j = (s_1^j, \dots, s_{i-1}^j, s_{i+1}^j, \dots, s_n^j)$	$s_{-i} = (s_{-i}^1, \dots, s_{-i}^J)$
Profile in auction j	$s^j = (s_1^j, \dots, s_n^j)$	$s = (s^1, \dots, s^J)$
Grand strategy space	$S^j = \prod_{i=1}^n S_i^j$	$S = \prod_{j=1}^J S^j$

6.3.1 Bounded Participation

Each buyer will know the available quantity for each market in which they bid. Buyers act strategically by selecting sellers, adjusting bid quantities, and choosing whether to participate based on their expected ability to satisfy demand. In the PSP framework buyers cannot reveal their entire valuation functions in a single step; instead they must request allocations iteratively. To regulate this behavior we introduce a bounded participation rule, which endogenously limits the set of sellers a buyer engages with, and can be seen as an analogue of the opt-out behavior given in [12].

Fix buyer i at time t and let p^* denote the common marginal price identified from opponents' bids. For each seller j let $c_j = \text{cap}_j(p^*)$ be the residual quantity available to i at price p^* . Define the desired total quantity

$$z_i^* = \min \left\{ \bar{q}_i(t), \sum_j c_j \right\}. \quad (6.1)$$

Definition 6.1 (Bounded participation rule). *Buyer i selects a minimal-cost subset*

of sellers $\mathcal{L}_i(t) \subseteq \mathcal{L}$, ordered by nondecreasing price $p_{(n)}^j(t)$, such that

$$\sum_{j \in \mathcal{L}_i(t)} c_j(t) \geq z_i^*. \quad (6.2)$$

The buyer allocates requests sequentially to the least expensive sellers until the desired total quantity z_i^* is reached, subject to residual capacities $c_j(t)$. For $j \notin \mathcal{L}_i(t)$, set $q_i^j = 0$.

This rule formalizes bounded participation at fixed t : each buyer interacts only with the fewest necessary sellers to realize z^* , in an attempt to minimize the cost of participation. The resulting allocation targets allocations at a common marginal price $p^*(t)$ under residual quantity constraints.

6.3.2 Residual Quantity and Allocation

As a market with perfect but incomplete information, sellers can only gain information about demand by observing buyer behavior, determined by the connectivity of the auction graph. In each iteration, every seller completes one update of its local auction.

For each seller j , the reserve price $p_*^j(t)$ is the price at which seller j is indifferent between selling her final unit of resource and keeping it. Equivalently, the seller may be viewed as submitting an internal bid $(Q^j, p_*^j(t))$ on her own auction. At the end of each round t , the reserve price is updated with information from the set of active bids, where $\mathcal{B}^j(t)$ is the set of buyers who win strictly positive allocations at seller j in round t , and $\epsilon > 0$.

We define the clearing price at seller j to be the smallest price at which aggregate awarded quantity meets available quantity:

$$\chi^j(t) = \min \left\{ y : \sum_{k: p_k^j(t) > y} a_k^j(t) \geq Q^j(t) \right\}. \quad (6.3)$$

Any residual supply must therefore be allocated among bids that tie at prices just above $\chi^j(t)$, after higher-priced bids are filled. Let

$$\underline{p}^j(t) := \min \{ p_i^j(t) : i \in \mathcal{B}^j(t) \}, \quad \bar{p}^j(t) := \max \{ p_i^j(t) : i \notin \mathcal{B}^j(t) \}, \quad (6.4)$$

be the lowest winning and highest losing bid prices at seller j , and where buyers *not* in $\mathcal{B}^j(t)$ receive zero allocation at seller j . The clearing price satisfies

$$\bar{p}^j(t) < \bar{p}^j(t) + \epsilon \leq \chi^j(t) \leq \underline{p}^j(t) - \epsilon < \underline{p}^j(t)$$

whenever there is at least one winning and one losing bidder at seller j . In particular, $\chi^j(t)$ lies in the open interval between the highest losing and lowest winning bid. At equilibrium, the reserve price $p_*^j(t)$ coincides with the clearing price at seller j , i.e., the clearing price implied by the PSP allocation rule.

Buyers at higher prices are therefore always served in full, whereas buyers at the threshold price may be rationed. At each price level y , the residual quantity is given by

$$R^j(y, t) = \left[Q^j(t) - \sum_{k: p_k^j(t) > y} a_k^j(t) \right]^+. \quad (6.5)$$

When multiple buyers tie at $p_i^j(t) = y$, the awarded allocation respects both the buyer's request and the residual supply. We refer to the tie-splitting rule originated in the analysis of quantized PSP auctions by Qu, Jia, and Caines [70],

$$a_i^j(s(t)) = \min \left\{ q_i^j(t), \frac{q_i^j(t)}{\sum_{\ell: p_\ell^j(t) = y} q_\ell^j(t)} R^j(y, t) \right\}. \quad (6.6)$$

The bid quantity $q_i^j(t)$ and the allocation $a_i^j(t)$ are complementary. In fact, the buyer strategy is the first term in the minimum, the second term being owned by the seller. For each buyer-seller pair (i, j) at time t , $a_i^j(t)$ is the *awarded* amount that seller j allocates to buyer i once the allocation rule has been applied. By construction,

$$a_i^j(t) \leq q_i^j(t), \quad (6.7)$$

with equality holding when residual supply at the buyer's price suffices to satisfy all tied requests. The mechanism therefore never awards more than requested and may award less when quantity is limited.

We remark that the reserve price $p_*^j(t)$ that lies in the margin interval determined by the bids

$$\bar{p}^j(t) < p_*^j(t) < \underline{p}^j(t), \quad (6.8)$$

whenever both $\bar{p}^j(t)$ and $\underline{p}^j(t)$ are defined, and we deliberately leave the precise rule for selecting $p_*^j(t)$ within the interval (6.8) unspecified. In particular, admissible choices include

$$p_*^j(t) = \chi^j(t), \quad p_*^j(t) = \bar{p}^j(t) + \epsilon, \quad p_*^j(t) = \underline{p}^j(t) - \epsilon,$$

provided that reserve price updates lie within ϵ and the resulting sequence $\{p_*^j(t)\}_t$ is nondecreasing.

6.3.3 Exclusion–Compensation

Each buyer's payment follows a second–price externality principle, this is the “social opportunity cost” of the PSP pricing rule. The exclusion–compensation payment to buyer i equals the loss imposed on other buyers at that seller when i participates. For a fixed auction j we use the opposing buyers' piecewise–constant marginal price function $P^j(\cdot, s_{-i}^j)$ built from s_{-i}^j ,

$$c_i^j(s) = \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz, \quad (6.9)$$

which holds true locally at each auction, where the opposing bids are calculated against the allocated resource to buyer i . The amount of resource available at price $p_{(n)}^j$ is $\xi_{n-1}^j - \xi_n^j \geq 0$. The local inverse price function is then

$$P^j(z, s_{-i}^j) = p_{(n)}^j \quad \text{for } z \in (\xi_n^j, \xi_{n-1}^j]. \quad (6.10)$$

For each ordered price y , we have that $P_i(z, s_{-i})$ is defined for the range of z corresponding to the total resource available from all sellers at that price, i.e.,

$$z \in \left(\sum_{p_{(n)}^j > y} (\xi_{n-1}^j - \xi_n^j), \sum_{p_{(n)}^j \geq y} (\xi_{n-1}^j - \xi_n^j) \right]. \quad (6.11)$$

Define the aggregate residual quantity

$$Q_i(y, s_{-i}) = \sum_{j=1}^{\mathcal{L}} Q_i^j(y, s_{-i}^j), \quad P_i(z, s_{-i}) = \inf\{y \geq 0 : Q_i(y, s_{-i}) \geq z\}, \quad (6.12)$$

where because $Q_i(y, s_{-i})$ is a right-continuous, nondecreasing step function with finitely many jumps at $\{p_{(m)}^j\}$, the infimum is attained.

6.3.4 Valuation and Utility

Each buyer i has an elastic valuation function $\theta_i : [0, Q_i] \rightarrow [0, \infty)$ with strictly decreasing derivative θ'_i . The valuation depends on the total awarded quantity across all sellers:

$$V_i(a) = \theta_i \left(\sum_{j=1}^J a_i^j(t) \right) = \int_0^{\sum_{j=1}^J a_i^j(t)} \theta'_i(z) dz. \quad (6.13)$$

Given a strategy profile s , the utility of buyer i for potential allocation a is dependent on the cost, $c_i(s)$, where the cost to buyer i as a function of the entire strategy profile s .

In the dynamic setting this profile evolves with iteration t , where $c_i(s)$ may represent total participation costs, including membership fees, per-round overhead, and per-auction message costs. Utility is given by

$$u_i(s) = V_i(a) - c_i(s), \quad (6.14)$$

where $c_i(s)$ is a dynamic cost function that evolves over time with bid updates.

The buyers' utility functions implicitly define a potential over the allocation space, as buyers seek to maximize their utility through strategic allocation requests. We note that a uniform (coordinated) bid price from buyers across active sellers upholds strategic simplicity and second-price incentives, which are rational under quasi-linear utilities, as shown in the original PSP framework [45].

Following Lazar and Semret [45], updates occur only when the buyer's utility improvement exceeds a small positive threshold, ensuring asynchronous convergence under bounded delay. In a single-auction market, buyer i accepts a new bid s'_i only

if $u_i(s'_i; s_{-i}) - u_i(s_i; s_{-i}) > \varepsilon$. In the multi-auction setting, buyer i posts a vector of bids that share a common marginal price p_i^* across all connected sellers. The utility comparison therefore becomes an aggregate test, where, in terms of the opposing bid vector s_{-i} , any gain in utility at time t depends on the current state of play. Information propagation across the market affects how the vector of opposing bids s_{-i} is formed, and thus how externalities are computed. The realized utility improvement $\Delta u_i(t)$ is evaluated relative to the previous round to determine if a new bid exceeds the cost of participation.

The discussion of externality under multiple auctions running asynchronously and a formal convergence analysis of the PSP mechanism to a single, unique, global ε -Nash network equilibrium, as was given in [73], is outside the scope of this paper. We instead focus on the iterative application of a uniform marginal price and the localized pricing structure resulting from progressive bid updates on connected network components consisting of multiple sellers sharing multiple buyers under an assumed bipartite structure. A formal analysis of the effects of latency on a PSP auction is given in [11].

6.4 Influence Sets

We model the behavior of vertices (buyers and sellers) in this bipartite structure using influence sets. Each vertex's strategy space is influenced by neighboring vertices and evolves over time. Influence sets restrict the strategy space of buyers and sellers within bounded regions, stabilizing auction dynamics and supporting predictable market equilibria. As rational agents, buyers and sellers do not optimize perfectly but instead operate within acceptable thresholds of cost and utility. Influence propagation determines the flow of information, where bidding saturation occurs once influence sets stabilize. At saturation, there is no vertex that, upon calculating his measure of utility, suffers a changing set of opponent bids, and all subsequent bid updates are calculated on locally stable subgraphs of the market.

Thus, influence sets are subsets of the auction graph that represent the scope of

influence a particular vertex (buyer or seller) has on others over a finite number of auction iterations. These sets structure interactions into subsets of the auction graph where local equilibria form dense regions where bid updates are stabilized.

6.4.1 Primary (Direct) Influence Sets

Following the original definition from [12], the *primary influence set*, denoted Λ , for a given seller j at time t , is defined set-theoretically,

$$\Lambda_{\mathcal{L}}(j, t) = \bigcup_{i \in \mathcal{B}^j(t)} \mathcal{L}_i(t), \quad (6.15)$$

where $\mathcal{B}^j(t)$ is the set of buyers bidding on seller j , and $\mathcal{L}_i(t)$ is the set of sellers that buyer i bids on. Thus, $\Lambda_{\mathcal{L}}^{(1)}(j, t)$ represents all sellers directly connected to auction j via shared buyers at time t . This definition captures the notion that influence propagates across the auction graph through buyer–seller connections. The superscript (1) denotes the first layer of influence anchored at seller j expanding through buyer-mediated connections.

To illustrate, for a buyer i , the relevant bids s_i^j flow *from* the buyer to sellers. For a seller j , we reverse this; bids flow *into* the seller from buyers. The base case captures this directionality, which we get from market theory: buyers have positive demand, and sellers have negative demand (otherwise known as surplus). The direct influence set for buyer i , denoted $\Lambda_{\mathcal{B}}^{(1)}(i, t)$, includes buyers directly connected to i through shared sellers,

$$\Lambda_{\mathcal{B}}(i, t) = \bigcup_{j \in \mathcal{L}_i(t)} \mathcal{B}^j(t). \quad (6.16)$$

We now have the first layer of *buyer-to-buyer* influence induced by common seller participation. It serves as a foundation for constructing buyer–buyer influence graphs and identifying bid coordination structures within the network. This expression gathers the buyers indirectly connected to buyer i through shared sellers, filtered by active bids at the given iteration. It provides a way to trace buyer–buyer influence mediated through seller auctions.

We extend the definition from [12] in a theoretical and practical sense, defining the base case explicitly as the vertex itself,

$$\Lambda^{(0)}(x, t) = \{x\}, \quad (6.17)$$

emphasizing that at the zeroth level, the influence set represents only the vertex itself. This represents a measure of “self-influence”, such as reserve prices (for sellers) or initial valuations (for buyers), economically aligning with the idea that a seller starts from a reserve price reflecting their own valuation, while a buyer’s self-valuation corresponds to their initial maximum willingness-to-pay.

6.4.2 Expanded (Indirect) Influence Sets

For any vertex x in the auction graph (buyer or seller), and for $n \geq 1$, the primary influence set is expanded from the $(n - 1)$ -step influence set by aggregating direct neighbors at the next layer:

$$\Lambda^{(n)}(x, t) = \bigcup_{y \in \Lambda^{(n-1)}(x, t)} \Lambda^{(1)}(y, t). \quad (6.18)$$

Where we define a two-edge projection operator

$$\Lambda^{(1)}(y, t) := \begin{cases} \bigcup_{i \in \mathcal{B}^y(t)} \mathcal{L}_i(t), & \text{if } y \in \mathcal{L}, \\ \bigcup_{j \in \mathcal{L}_y(t)} \mathcal{B}^j(t), & \text{if } y \in \mathcal{B}, \end{cases}$$

which always returns vertices of the *same type* as y after one buyer–seller alternation.

For sellers, the n -step influence set may be computed recursively,

$$\Lambda_{\mathcal{L}}^{(n)}(j, t) = \bigcup_{i \in \Lambda_{\mathcal{B}}^{(n-1)}(j, t)} \mathcal{L}_i(t),$$

where $\Lambda_{\mathcal{B}}^{(n-1)}(j, t)$ is the set of buyers reachable from seller j in $n - 1$ steps. This returns the set of sellers that receive nonzero bids from buyers who are indirectly connected to auction j via shared bidding activity across n rounds of the PSP auction.

It describes how seller j 's influence propagates through buyer behavior across seller neighborhoods. For buyers,

$$\Lambda_{\mathcal{B}}^{(n)}(i, t) = \bigcup_{j \in \Lambda_{\mathcal{L}}^{(n-1)}(i, t)} \mathcal{B}^j(t),$$

where $\Lambda_{\mathcal{L}}^{(n-1)}(i, t)$ collects sellers indirectly connected to buyer i .

Each new layer $\Lambda^{(n)}(x, t)$ therefore adds the direct neighbors of all vertices in the previous layer, producing a breadth-first expansion in the auction graph. This recursive expansion therefore builds a “growing influence ball” centered at x , where the secondary set acts as a generalized neighborhood closure or hull around the initial primary set. At each step n , the influence set $\Lambda^{(n)}(x, t)$ forms an outer boundary surrounding the influence set $\Lambda^{(1)}(x, t)$, recursively aggregating direct neighborhoods around previously identified influence vertices.

Pathwise Characterization. In graph theory, this structure parallels the n -edge neighborhood closure or a breadth-first expansion of distance- n shells. We characterize $\Lambda^{(n)}(x, t)$ pathwise as the set of all vertices reachable from x by paths alternating between buyers and sellers, of length up to $2n$. Formally, let $\mathcal{G} = (\mathcal{I}, E)$ denote the bipartite auction graph, where \mathcal{I} is the set of agents (buyers and sellers), and an edge $(i, j) \in E$ exists if buyer i bids on seller j . The graph alternates between buyers and sellers by construction: no two buyers or two sellers are directly connected. Because of the bipartiteness we have the parity rule, and consequently

$$\Lambda^{(n)}(x, t) = \{y \in \mathcal{I} \mid \text{dist}_{\mathcal{G}}(x, y) = 2n\}.$$

From a strategic perspective, the expanded influence set $\Lambda^{(n)}(x, t)$ describes the scope of anticipated externalities: the agents whose actions may not affect x directly, but may impact x 's incentives via shared neighbors. These influence chains emerge in environments with incomplete information and approximate the region of the market that affects the *expected utility gradient* of vertex x . In equilibrium analysis, these indirect sets are crucial for understanding stability, coordination potential, and susceptibility to shock propagation (e.g., strategic manipulation or correlated noise). As

noted in [12] and echoed in broader decentralized market theory (e.g., [75]), indirect influence plays a key role in shaping convergence. While $\Lambda(x, t)$ governs observed interaction, $\Lambda^{(n)}(x, t)$ governs inferred or mediated interdependence—and together, they define the full strategic visibility of a vertex.

6.4.3 Projection Domains and Influence Operations

The influence set framework captures cascading dependencies and forms the foundation for our graph-theoretic analysis. To analyze the propagation of influence in the auction network, we construct influence sets using a sequence of projection operations on the underlying bid graph. Each active bid is indexed by a pair $(i, j) \in \mathcal{B} \times \mathcal{L}$, where buyer i submits a bid to seller j . These interactions collectively form the strategy space $S(t)$, which consists of the full collection of price-quantity bids s_i^j . We extract the active subgame by identifying $\mathcal{I}_{\text{active}}(t) \subseteq \mathcal{B} \times \mathcal{L}$, the set of observed interactions between buyers and sellers at time t . These pairs serve as both an interaction graph and an index set for the time-dependent strategy array $\mathbf{s}(t)$.

6.4.4 Projection-Based Influence Propagation

Let $\mathcal{I}_{\text{active}}(t) \subseteq \mathcal{B} \times \mathcal{L}$ be the set of buyer–seller pairs that submit positive bids at time t . Each pair (i, j) indexes the strategy array $\mathbf{s}(t) \in S(t)$, so $\mathcal{I}_{\text{active}}(t)$ is both a graph on $\mathcal{B} \cup \mathcal{L}$ and an index set for the strategy space.

Projection maps.

$$\pi : \mathcal{I}_{\text{active}}(t) \longrightarrow \mathcal{B}, \quad \pi(i, j) = i, \quad \varpi : \mathcal{I}_{\text{active}}(t) \longrightarrow \mathcal{L}, \quad \varpi(i, j) = j.$$

- *Structural role.* Alternating compositions $\varpi \circ \pi^{-1} \circ \pi \circ \varpi^{-1} \dots$ trace paths through the bipartite auction graph, giving n -edge neighborhoods.
- *Strategic role.* Acting on $S(t)$, the same maps carve out partial strategy profiles (e.g. all bids of a given buyer).

Full pre-images. We take the composition of the the projections in order to restrict

and vectorize the space $S(t)$. For any buyer i and seller j we write

$$\varpi^{-1}(i, t) = \{(i, j') \in \mathcal{I}_{\text{active}}(t)\}, \quad \pi^{-1}(j, t) = \{(i', j) \in \mathcal{I}_{\text{active}}(t)\},$$

so that

$$\varpi \circ \pi^{-1}(i, t) = \{j' \mid (i, j') \in \mathcal{I}_{\text{active}}(t)\} = \mathcal{L}_i(t),$$

$$\pi \circ \varpi^{-1}(j, t) = \{i' \mid (i', j) \in \mathcal{I}_{\text{active}}(t)\} = \mathcal{B}^j(t).$$

n -step influence. Because the graph is bipartite, two successive projections always return a vertex of the *same* type. These projections serve dual purposes: structurally, they trace paths through the auction graph, alternating between buyers and sellers; strategically, they extract subspaces of $S(t)$ that represent partial strategies or responses and may evolve as patterns in the form of active bids sets.

Connected components via iterated projections Because the auction graph is bipartite, two successive projections return a vertex of the same type. Define the composition operators

$$P := \varpi \circ \pi^{-1} \quad (\text{buyer} \rightarrow \text{seller}), \quad Q := \pi \circ \varpi^{-1} \quad (\text{seller} \rightarrow \text{buyer}).$$

Starting from a seller j , one step of influence expansion is

$$\Lambda_{\mathcal{L}}^{(1)}(j, t) = P Q (j) = \varpi \circ \pi^{-1} \circ \pi \circ \varpi^{-1}(j, t),$$

which moves seller \rightarrow buyers \rightarrow sellers. Analogously, for a buyer i we set $\Lambda_{\mathcal{B}}^{(1)}(i, t) = Q P(i)$.

The n -edge neighborhoods follow by simple iteration,

$$\Lambda_{\mathcal{L}}^{(n)}(j, t) = (P Q)^n(j), \quad \Lambda_{\mathcal{B}}^{(n)}(i, t) = (Q P)^n(i), \quad n \geq 1,$$

with the base case $\Lambda^{(0)}(x, t) = \{x\}$.

Each application of PQ (or QP) adds exactly one buyer–seller alternation, so $\Lambda^{(n)}(x, t)$ is the breadth-first shell lying n edges away from x . Iterating until $\Lambda^{(n)}(x, t) =$

$\Lambda^{(n-1)}(x, t)$ closes the connected component containing x . Thus the recursive projection operator captures both direct and indirect influence flows. Thus, $\Lambda^{(n)}$ may be interpreted as the n -step neighborhood in the auction graph and as a dynamic closure of best-response behavior. Each expansion layer captures not just structural proximity but strategic influence—the transmission of incentive, information, and utility across the network. In this way, we convert local participation patterns into global influence propagation, formalized as graph-theoretic expansions over the projected structure of the strategy space.

6.4.5 Partial Ordering and Market Shifts

While the projection mappings π, ϖ and their compositions produce index sets (subsets of \mathcal{B} or \mathcal{L}), these sets are characterized by the underlying strategy space $S(t)$. These indices correspond directly to elements of the strategy space $S(t)$, which contains structured bid information $s_i^j(t)$. That is, the projection $\varpi^{-1}(i, t)$ retrieves all bid tuples (i, j) in the index set, but equivalently defines the subspace of $s_i(t)$ consisting of all bid array submitted by buyer i at time t .

Each element $s^j(t) \in S^j(t)$ is a bid array, and the collection $\pi \circ \varpi^{-1}(j, t)$. Aside from being an index set of buyers, we have a set of bid arrays $\{s_i^j(t)\}_{i \in \mathcal{B}^j(t)}$ that can be partially ordered by their prices $p_{(n)}^j(t)$. While the projection operators isolate buyers or sellers structurally, the *functional influence* between market participants is mediated through the comparison of bid prices.

This introduces a natural partial order among bidders for each seller at a fixed time, and we define a partial ordering on $S^j(t) \subset S(t)$ by $p_i^j(t) < p_k^j(t)$ if buyer i bids less than buyer k for the same seller j . Given the set of buyers $\mathcal{B}^j(t) = \pi \circ \varpi^{-1}(j, t)$, we may impose a partial order structure based on the associated price bids $\{p_i^j(t)\}$. This ordering determines which buyers are accepted by seller j (those with highest prices until the resource is exhausted), and the sensitivity of a potential equilibrium to shifts in bidding behavior.

As shown in [12], a market shift occurs precisely when a buyer outside of $\mathcal{B}^j(t)$

improves their relative position in this order, causing $\mathcal{B}^j(t)$ to be recomputed. Such shifts reflect structural changes in functional influence, as they alter the competitive hierarchy among bids and propagate through the multi-auction environment. Specifically, a market shift in auction j occurs when the partial order of bids at seller j changes in a way that affects allocation. From [12], two cases are critical:

1. **Demand Shortfall:** A buyer $i \in \mathcal{B}^j(t)$ reduces their bid quantity so that total demand falls below available supply:

$$\sum_{i \in \mathcal{B}^j(t)} a_i^j(t) < Q^j(t).$$

The auction must recompute its reserve price or reallocate supply among remaining buyers.

2. **Bid Overtake:** A buyer $i^* \notin \mathcal{B}^j(t)$ improves their valuation so that

$$p_{i^*}^j(t) < p_*^j(t),$$

where $p_*^j(t)$ is the minimum accepted bid price at auction j . The buyer i^* displaces the marginal buyer, triggering a shift in $\mathcal{B}^j(t)$.

Either case changes the minimal winning price, breaking the partial ordering within the projected sets, and forces the auction to recompute $\mathcal{B}^j(t)$. As the seller frontier $\Lambda_{\mathcal{L}}^{(1)}(j, t)$ consists of sellers connected to j through shared buyers, the reallocation thereby propagates influence through layers of the the expansion $\Lambda_{\mathcal{L}}^{(n)}(j, t) = (\varpi \circ \pi^{-1} \circ \pi \circ \varpi^{-1})^n(j)$.

6.5 Influence Shells and Local Saturation

The projection and ordering framework developed above allows us to describe how influence propagates through the auction network. We now ask under what conditions this propagation stabilizes. As buyers and sellers iteratively adjust bids, certain neighborhoods of the market reach a state in which no participant can improve

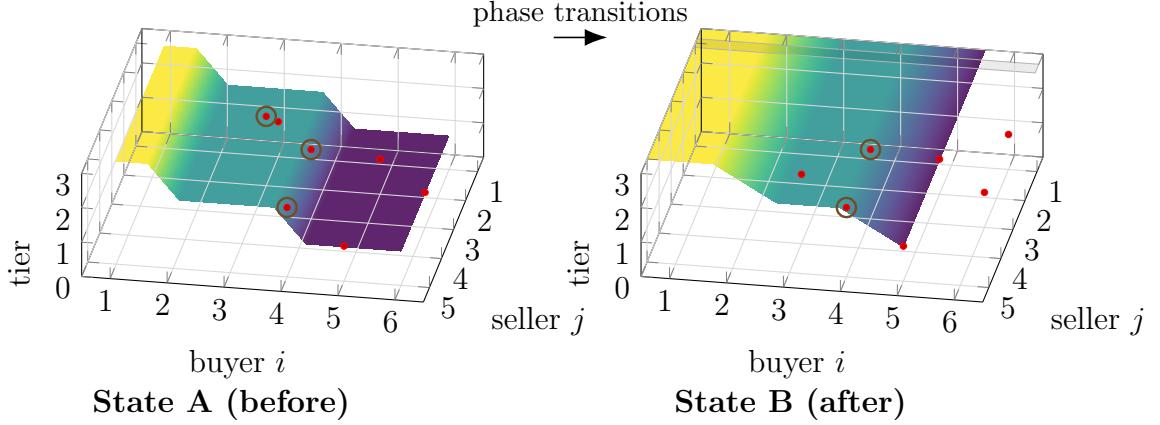


Figure 6.1: 3D matrix view: rows (buyers), columns (sellers), and z encodes price tiers. The colored surface shows the buyer price; filled markers are active bids; open circles show marginal winners. The right panel shows a transition where a new high-tier participant appears at $j=1$, a demand shortfall removes a low cell, and a reconfiguration shifts activity at $j=2$.

their utility through unilateral deviation. These locally stable regions form influence shells–bounded subgraphs within which allocations, prices, and bid updates remain consistent under further iterations. When every buyer and seller in such a region satisfies this best-response property, the shell is said to be saturated. The following gives the formal notation and enumerates the assumptions that we have made in the generalization of influence sets as were defined in [12].

Definition 6.2 (Saturated Influence Shell). *A primary influence set $\Lambda_{\mathcal{L}}^{(1)}(j, t)$ associated with seller j at time t is said to be saturated if no buyer or seller within this set can improve their utility by unilaterally altering their bids. Formally, for every buyer $i \in \mathcal{B}^j(t)$ and every seller $\ell \in \Lambda_{\mathcal{L}}^{(1)}(j, t)$, the following holds,*

$$u_i(t) \geq u'_i(t), \quad \text{for any feasible alternative strategy } s'_i(t).$$

Global market equilibrium decomposes into interconnected saturated shells, each functioning as stable subsystems. We establish conditions for the existence of a saturated influence shell.

(i.) **Countable and Locally Finite Graph.** The sets of buyers \mathcal{B} and sellers \mathcal{L} are at most countably infinite. Each participant engages in only finitely

many transactions, ensuring finite degree at every instant. This guarantees that all projection maps (π, ϖ) encounter only finite fibres and that the influence operator (6.18) perform finite unions.

Buyers and sellers participate in locally finite networks, enabling stable equilibrium convergence within compact, bounded strategy spaces. Market rules explicitly limit resources and interactions, ensuring finite dimensionality.

- (ii.) **Bounded Influence and Bids.** Influence propagation strength remains bounded, preventing divergence. Each seller's fixed endowment Q^j and each buyer's fixed demand cap Q_i are finite and time-independent. Hence every non-zero bid quantity $q_i^j(t)$ lies in the compact interval $[0, Q_i]$, and every realized allocation is in $[0, Q^j]$.
- (iii.) **Partial Ordering Stability.** The partial ordering induced by the bid structure must satisfy stable bid threshold rankings for all relevant buyers and sellers within the shell, thus establishing clear marginal price tiers.

6.5.1 Ordering and Influence Propagation

We recall the ordering relationship from [12] that holds for any seller within a saturated primary influence set $\Sigma := \Lambda_{\mathcal{L}}^{(1)}(j, t)$.

Lemma 6.3 (Local Price Ladder [12, Thm. 2.3, proof]). *Let the market be at time t with seller j and its saturated primary influence set $\Sigma := \Lambda_{\mathcal{L}}^{(1)}(j, t)$. Pick any neighbor $k \in \Sigma$ and two buyers*

$$i \in \mathcal{B}^j(t), \quad \ell \in \Lambda_{\mathcal{B}}^{(1)}(i, t) \setminus \mathcal{B}^j(t),$$

i.e., i bids on j , k is another seller reached from i , and ℓ is a buyer that bridges further to k but not to j .

If the shell $\Lambda_{\mathcal{L}}^{(1)}(j, t)$ is saturated—no profitable deviation exists for any vertex in this set—because ℓ does not bid on j while i bids on both j and k , the marginal

thresholds must nest as follows

$$\sup(\bar{p}^k(t), \underline{p}^k(t)) \leq p_\ell^*(t) < \inf(\bar{p}^j(t), \underline{p}^j(t)) \leq p_i^*(t), \quad (6.19)$$

where the marginal intervals are chosen as in (6.4), and in particular, every choice of reserve prices $p_*^k(t) \in (\bar{p}^k(t), \underline{p}^k(t))$ and $p_*^j(t) \in (\bar{p}^j(t), \underline{p}^j(t))$ satisfies

$$p_*^k(t) \leq p_\ell^*(t) < p_*^j(t) \leq p_i^*(t). \quad (6.20)$$

Proof. The argument for (6.20) is contained within the proof of [12, Thm. 2.3]. Left-to-right the chain reads

- (i.) p_*^k seller k 's reserve price;
- (ii.) p_ℓ^* buyer ℓ 's bid that clears the marginal tier in *both* auctions;
- (iii.) p_*^j seller j 's reserve price;
- (iv.) p_i^* the highest active bid of buyer i on seller j .

The left inequality holds because k clears at the minimum of the two buyers' bids; the strict middle inequality follows from ℓ placing no bid on j ; the right inequality is enforced by buyer i 's cross-auction participation. In a saturated shell, any profitable deviation by k toward j or by i away from j is ruled out. Hence buyer ℓ 's marginal valuation must lie weakly above all prices at which $newk$ can clear its final unit, while buyer i 's valuation must lie weakly below all prices at which j can still clear. This forces the margin intervals to nest as in (6.19). Since $p_*^k(t) \in (\bar{p}^k(t), \underline{p}^k(t))$ and $p_*^j(t) \in (\bar{p}^j(t), \underline{p}^j(t))$ by the definition of the reserve price p_*^j , the pointwise ladder (6.20) follows immediately. \square

Economically, this implies that higher prices at neighboring sellers prevent buyers from deviating profitably, ensuring that no participant has an incentive to alter their bidding strategy unilaterally. Hence, the ordering captures a stable distribution of resources and prices, reflecting locally optimal market conditions.

Weak (local) Monotonicity The projected influence sets, together with the induced partial orders, thus form the dynamic framework for market evolution, where projections identify which vertices are connected, and partial orders determine how influence is transmitted via price shifts. Market shifts occur when the partial order structure is perturbed beyond certain thresholds, forcing recomputation of $\mathcal{B}^j(t)$ or reserve prices.

Proposition 6.4 (Local Monotonicity). *Let $\Sigma = \Lambda_{\mathcal{L}}^{(1)}(j, t)$ be the saturated shell of seller j at time t . For each seller $k \in \Sigma$, let the marginal intervals be chosen as in (6.4), and let the reserve price $p_*^k(t)$ be selected inside the interval $(\bar{p}^k(t), \underline{p}^k(t))$.*

Now, consider the vector of seller reserves and buyer marginal valuations restricted to Σ . Under elastic, strictly decreasing valuation functions $\theta'_i(z)$ and any reserve-update rule that

(i) selects $p_^k(t+1) \in (\bar{p}^k(t+1), \underline{p}^k(t+1))$ for each $k \in \Sigma$, and*

(ii) is nondecreasing in the bids $\{p_i^k(t)\}_i$ and the previous reserve $p_^k(t)$,*

the PSP price-update map

$$p_*^k(t) \mapsto \tilde{p}_*^k(t+1) \quad \text{and} \quad p_i^k(t) \mapsto \tilde{p}_i^k(t+1) \quad (6.21)$$

is locally monotone on Σ .

Proof. Elasticity of the valuation functions implies that each buyer's marginal value $p_i^*(t) = \theta'_i(z_i(t))$ is strictly decreasing in its own allocation. Increasing any bid or reserve inside the saturated shell Σ may raise some allocations and lower others, but it cannot create a reversal of the bid ordering that defines the interval $(\bar{p}^k(t), \underline{p}^k(t))$. In particular, all ladder relations (6.20) remain invariant under such updates.

For sellers, the reserve-update rule is assumed nondecreasing in both the previous reserve $p_*^k(t)$ and the local bids $\{p_i^k(t)\}_i$. Thus any componentwise increase in $(p_*^k(t), p_i^k(t))$ cannot decrease any updated reserve $p_*^k(t+1)$. By construction, each updated reserve remains inside its interval $(\bar{p}^k(t+1), \underline{p}^k(t+1))$, and, because the

shell Σ is saturated, these intervals evolve compatibly with the ordering relations and cannot induce a downward jump that violates the ladder.

Therefore, every coordinate of the updated pair $(p_*^k(t+1), p_i^k(t+1))$ is weakly increasing in the corresponding coordinate of $(p_*^k(t), p_i^k(t))$. Hence the update rule satisfies the order-preserving property (6.21), and the PSP price–update map is locally monotone on Σ . \square

The partial ordering structure induced by bidding behavior is essential in analyzing and predicting the direction and magnitude of market shifts resulting from influence dynamics. A local allocation triggers global bid adaptation, reinforcing that while seller auctions operate independently, buyer strategy space remains tightly coupled. In integrated markets (scarce supply), the partial orders are dense and tightly coupled, making markets highly sensitive and globally coordinated. In fragmented markets (abundant supply), the partial orders become sparse and disconnected, leading to localized equilibria and insulating submarkets from external shocks. Thus, the transition from integrated to fragmented equilibrium is not just a graph phenomenon—it is a transition in the connectivity of the partial order structure induced by bidding.

Remark 6.5 (On the ordering of the PSP price map). *Although the local update map is written on the pair (p^j, p_i^j) , the PSP rule actually acts only on the first coordinate: the buyer’s bid p_i^j is simply carried forward (or set to 0 if $j \notin \mathcal{L}_i$).*

6.5.2 Asynchronous Sellers and Coupled Buyers

To represent the fine-grained dynamics of bid selection and displacement within an auction round, we introduce an internal index τ_k to describe local progression steps. Each τ_k denotes a partial-ordering resolution event—an allocation decision at auction j followed by an update to the reserve price and potentially to the projected sets. The index acts as a local time variable inside the global iteration t , allowing us to separate micro-adjustments from round-to-round evolution.

Definition 6.6 (Allocation Step τ_k). *At each τ_k within round t , seller j selects the*

highest bidder in $\pi \circ \varpi^{-1}(j)$ not yet fulfilled, allocates a feasible amount $a_i^j(\tau_k)$, updates the reserve price $p_*^j(\tau_{k+1})$, and recomputes \mathcal{B}^j and $\Lambda_{\mathcal{L}}(j)$ as needed.

Each τ_k inside a global round t is therefore a local ordering-resolution event: seller j picks the highest unfilled bidder, allocates a feasible amount $a_i^j(\tau_k)$, updates its reserve, and recomputes the bidder set $\mathcal{B}^j(t)$ as well as the seller shell $\Lambda_{\mathcal{L}}^{(1)}(j, t)$. The sequence τ_1, τ_2, \dots terminates when no remaining buyer meets the current reserve or when supply is exhausted.

Sellers operate independently: each seller's τ_k sequence proceeds without synchronization with others. However, buyers must maintain a consistent strategy across all sellers they bid on. Since the buyer's bid array $\sigma_i(t)$ is defined jointly over $\mathcal{L}_i(t)$, any change to the outcome of one auction requires coordinated updates across all components. This coupling between independent seller threads through shared buyers produces the feedback mechanism responsible for market coherence. From a game-theoretic perspective, each seller executes a local best-response process, while each buyer enforces cross-auction consistency of marginal valuation.

We define the buyer update rule as

$$q_i^j(\tau_{k+1}) = Q_i(t) - \sum_{j' \in \mathcal{L}_i(t)} a_i^{j'}(\tau_k), \quad \forall j \in \mathcal{L}_i(t), \quad (6.22)$$

indicating that the buyer updates all bids simultaneously based on observed allocations. Equation (6.22) ensures that a buyer's total requested quantity never exceeds its available resource $Q_i(t)$ and redistributes residual demand across the active seller set $\mathcal{L}_i(t)$. The rule formalizes how buyers translate local allocation feedback into revised offers, maintaining a form of budget balance across asynchronous auctions.

Because sellers run their τ_k threads asynchronously while buyers must update all bids coherently according to (6.22), local price changes propagate through the partial orders defined above, layer by layer. Each seller's reserve adjustment initiates a chain of bid updates in the neighborhoods that share its buyers. This extends the projection-ordering framework, enabling intra-round modeling of bid dynamics, bid-induced reordering, and precise tracking of influence propagation through up-

dates to the projected domains. In this sense, the τ_k sequence acts as a micro-time resolution that reveals how local saturation unfolds inside each global auction round.

Stability under asynchronous evolution The following proposition shows that, even under these independent update threads, the relative ordering of bids remains stable and the local price ladder is preserved.

Proposition 6.7 (Saturated Shell). *Let $\Sigma = \Lambda_{\mathcal{L}}^{(1)}(j, t)$ be saturated at τ_k . Assume that for every seller $k \in \Sigma$ the reserve price lies in the interval,*

$$\bar{p}^k(\tau_k) < p_*^k(\tau_k) < \underline{p}^k(\tau_k),$$

and that each τ -update preserves all ladder relations (6.20) inside Σ . If a local resolution step $\tau_k \rightarrow \tau_{k+1}$ modifies the strategy space only inside Σ , then:

- (i) *The ladder (6.20) is preserved: no ordering reversal is possible, and every affected marginal or reserve price weakly increases; a strict increase occurs whenever the winning bid or reserve at some seller in Σ rises.*
- (ii) *If every buyer or seller that first appears in $\Sigma_{n+1} \setminus \Sigma_n$ has no profitable deviation given the preserved ladder, then the expanded shell Σ_{n+1} is saturated.*

Proof. If a τ -update reallocates quantity within Σ , the clearing price or reserve at the affected seller can only move upward within its margin interval. Because all reserves and winning bids in Σ lie inside their intervals $(\bar{p}^k, \underline{p}^k)$, no update can create a reversal of the ordering that defines the ladder (6.20). Thus each affected component moves weakly upward, and whenever the winning bid or reserve at some seller in Σ increases, at least one of the four prices in the ladder strictly increases.

Saturation implies every τ -update inside Σ preserves best-response conditions given the ladder; hence extending Σ by including agents in $\Sigma_{n+1} \setminus \Sigma_n$ yields a saturated larger shell whenever those newly added agents also have no profitable deviation under the same ordering. \square

Saturation implies every τ_k update inside Σ is either a demand–shortfall or bid–overtake event; both raise the marginal price they touch, propagating weakly upward along every chain of the form (6.20). A strict increase occurs whenever the winning bid or reserve at some auction in Σ is lifted. Thus, local “saturation” is a best–response property of a one–edge influence shell.

By an inductive test we extend saturation shell–by–shell.

Corollary 6.8. *The new shell Σ_{n+1} inherits the price–ordering ladder (6.20): all its marginal prices are no smaller than those in Σ_n ; in particular, $p_*^k(\tau_{k+1}) \geq p_*^k(\tau_k)$ for all $k \in \Sigma_{n+1}$.*

Proof. Consider any buyer–seller quadruple (k, ℓ, j, i) whose seller k lies in the freshly revealed layer $\Sigma_{n+1} \setminus \Sigma_n$. Applying Proposition 6.4 to that quadruple shows that the ladder inequality (6.20) is preserved and all four prices weakly increase. Repeating this argument for every such quadruple that touches Σ_{n+1} completes the extension of the monotone ladder one edge outward. For every edge (i, k) with $k \in \Sigma_n$, Proposition 6.7 guarantees that a local τ –update cannot decrease the marginal price p_*^k . Hence $(p_*^k)_{k \in \Sigma_n}$ is component–wise non–decreasing from τ_k to τ_{k+1} . \square

The asynchronous update model developed above provides the conceptual foundation for our simulation studies. It captures the essential features of decentralized PSP dynamics: sellers acting independently on local information, buyers co-ordinating across overlapping auctions, and influence propagating through partially ordered interactions. In the following section, we use these principles to construct an event–driven simulation framework that allows us to observe how local saturation emerges in practice.

6.6 Simulation Framework and Implementation

This section summarizes the simulation code used to study the PSP markets with multiple sellers and buyers. The simulation architecture explores the practical realizations of decentralized coordination. The event-driven approach reproduces the

iterative best-response behavior implied by the mechanism and allows examination of convergence properties, price dispersion, and efficiency loss due to network coupling.

Following Semret and Lazar [45], each buyer's valuation is given by a parabolic curve of the form

$$\theta_i(z) = \kappa_i(\bar{q}_i - z/2)z \quad \text{for} \quad z \in [0, \bar{q}_i]$$

where \bar{q}_i represents the maximum quantity of goods desired and $\kappa_i = \bar{p}_i/\bar{q}_i$ has dimensions marginal price per unit where \bar{p}_i is the maximum marginal value that buyer i would ever place on the resource.

6.6.1 Event-Driven Algorithm and Asynchronous Updates

The simulation operates as a discrete-time event system. Events are scheduled and processed in a priority queue, advancing the simulation clock t to the next event. Two event types exist:

1. **Buyer Compute:** Buyer i evaluates its local state, computes updated bids (z_i^j, p_i^j) on each connected seller j , and schedules bid events when meaningful changes occur.
2. **Post Bid:** Seller j clears its auction, applying second-price allocation and updating quantities, payments, and revenues.

Buyer and seller events may reschedule each other (e.g., clears triggered after meaningful bid changes). The loop halts when no effective changes remain or a step limit is reached. The simplified pseudocode is shown in Algorithm 6.1.

Algorithm 6.1 Event-driven PSP simulation

```

1: Initialize market state  $M$ ; schedule all buyers.
2: while queue not empty and not converged do
3:    $(t, \text{type}, \text{payload}) \leftarrow \text{pop}()$ 
4:   if type = BUYER_COMPUTE then
5:     Update  $(z_i^j, p_i^j)$  for buyer  $i$  on feasible links.
6:     Schedule POST_BID events for affected sellers.
7:   else if type = POST_BID then
8:     Seller  $j$  clears auction, enforcing  $Q^j$ , opponent ordering, and payments.
9:   end if
10: end while
  
```

Each buyer i computes a uniform (or per-seller) bid price using a valuation-based update $w = \theta_i'(\sum_j z_i^j)$. Quantities are apportioned across incident sellers using a local best-response step. Buyers are sorted by descending unit price $p_{(n)}^j$. The clearing process accumulates allocations until total demand equals available resource Q^j . The threshold price

$$p_*^j = \min\{p_i^j : \sum_{k: p_k^j \geq p_i^j} q_k^j \geq Q^j\} \quad (6.23)$$

identifies the seller's marginal (clearing) price.

For each seller j , the clear routine builds a partial ordering by posted marginal prices (bid prices), serves opponents until available resource Q^j is exhausted, and charges the price incurred by the externality of participation to all served buyers for that seller. The routine updates \mathbf{a} , seller revenue, and per-buyer costs. We define the set of active buyers with positive bids as

$$\mathcal{I}_j = \{i : q_i^j > 0, a_{ij} \in \mathbf{A}\}, \quad (6.24)$$

where A represents the biadjacency matrix captures direct buyer–seller interactions:

$$\mathbf{A}(t) \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{L}|},$$

where $\mathbf{A}_{ij}(t) = 1$ if buyer i bids on seller j at time t , and $\mathbf{A}_{ij}(t) = 0$ otherwise. Rows of $\mathbf{A}(t)$ identify each buyer's active sellers; columns identify all buyers bidding on a given seller.

Experiments are conducted using randomized networks of $I = |\mathcal{B}|$ buyers and $J = |\mathcal{L}|$ sellers. The connectivity matrix \mathbf{A}_{ij} determines which buyers may interact with which sellers. Each run uses the following protocol:

1. Initialize market state M with parameters $(I, J, Q^j, \varepsilon, \text{reserve})$.
2. Assign buyer valuations (\bar{q}_i, κ_i) and budgets b_i from uniform ranges.
3. Generate random biadjacency matrices with varying density (percentage of shared buyers).
4. Execute the event-driven simulation for a fixed iteration limit.
5. Record convergence statistics, prices, allocations, and revenues.

The event scheduler supports both deterministic and stochastic updates, allowing controlled comparison between synchronous and asynchronous dynamics.

Experimental Setup. Each experiment initializes a market with I buyers and J sellers. Seller capacities are fixed at $Q_{\max} = [60.0, 40.0]$, with buyers distributed across both sellers according to a connectivity percentage that varies from 0% (fully isolated) to 100% (fully connected) in increments of 10%. For each connectivity level, a base random seed (`base_seed = 20405008`) ensures reproducibility while allowing controlled stochastic variation across runs.

Following Semret and Lazar [45], each buyer's valuation is given by a parabolic curve of the form

$$\theta_i(z) = \kappa_i(\bar{q}_i - z/2)z \quad \text{for} \quad z \in [0, \bar{q}_i]$$

where \bar{q}_i represents the maximum quantity of goods desired and $\kappa_i = \bar{p}_i/\bar{q}_i$ has dimensions marginal price per unit where \bar{p}_i is the maximum marginal value that buyer i would ever place on the resource. Note that κ_i is larger for buyers who derive more value from the resource. Now choose \bar{q}_i and \bar{p}_i independently for all i such that

$$\bar{q}_i \sim U[50, 100] \quad \text{and} \quad \bar{p}_i \sim U[10, 20], \quad (6.25)$$

where $U[a, b]$ represents the uniform distribution over the interval $[a, b]$. Noise and perturbation effects are controlled by $\epsilon = 2.5$. For each seed and connectivity level, the simulation executes until convergence, measuring clearing prices, allocations, and bid prices.

A sequence of derived seeds $\text{seed} = \text{base_seed} + s$ is used for each connectivity level s , ensuring comparable random draws while preserving independence across runs.

6.6.2 Price Ladder Verification

The simulation presented here focuses on verifying the *price ladder condition* across interconnected sellers. This experiment represents a localized instance of the broader PSP market, designed to test whether clearing prices obey a monotonic relationship when sellers share buyers through overlapping influence sets.

The experiment initializes a small market composed of two sellers ($j = 1$ and $\ell = 0$) and four buyers ($i = 0, 1, 2, 3$). The adjacency structure allows some buyers to connect to both sellers, while others remain local. Sellers have distinct capacities, $Q^1 = 8$ and $Q^0 = 15$, reflecting asymmetric market sizes. The buyer valuation and bid initialization follow:

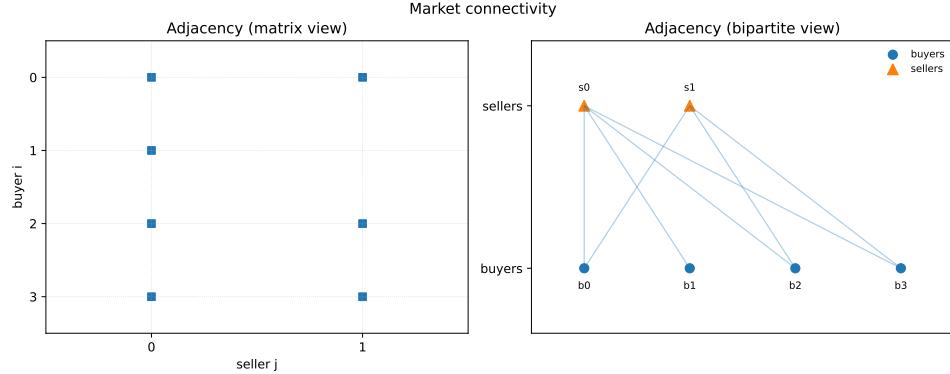
$$\begin{array}{ll} (0, 1) : q = 8, p = 40, & (0, 0) : q = 8, p = 40, \\ (1, 0) : q = 2, p = 4, & (2, 0) : q = 6, p = 1. \end{array}$$

We have the connectivity of the market, in alignment with Lemma 6.3.

The resulting market has a highly skewed valuation distribution, allowing one buyer to dominate both sellers, while others form the marginal tiers that define second-price boundaries.

The algorithm scans all sellers and their one-hop neighbors to evaluate tuples (ℓ, k, j, i) where Buyer i connects the two sellers. It tests the three inequalities defining the ladder ordering $p_\ell^* < p_k < p_j^* \leq p_i$. If these inequalities hold for all tuples, the market satisfies the monotone price ladder condition. Violations are reported with

Figure 6.2: Adjacency structure showing market connectivity between buyers and sellers.



detailed tuple traces to aid in diagnosing market inconsistencies.

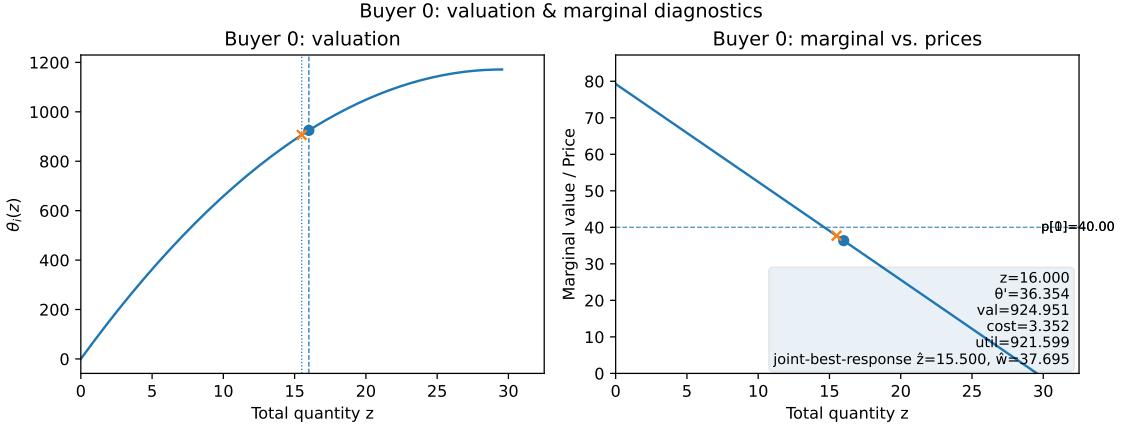
In this configuration, the ladder tuples satisfy all three inequalities, confirming a monotone relationship among clearing and bid prices. The system outputs a detailed report including, number of valid tuples and unique seller pairs (j, ℓ) , margins between successive price tiers: $(p_k - p_\ell^*)$, $(p_j^* - p_k)$, and $(p_i - p_j^*)$, and a summary of any violations detected. For this experiment, the output indicates no violations and consistent monotonicity, demonstrating that the PSP clearing mechanism maintains a globally ordered price structure when local competition and influence overlap exist.

This controlled experiment provides an analytical validation of the price ladder lemma in a simplified setting, and is intended to act as a unit test. It confirms that bid prices across connected sellers obey the expected inequalities implied Lemma 6.3. More generally, it shows that when buyers bridge multiple sellers, the second-price mechanism induces a coherent ordering of marginal prices, and provides an analytical tool for extending this verification to larger graphs. In the tuples $(\ell, k, j, i) = (0, 1, 1, 0)$ and $(0, 2, 1, 0)$ we observe

$$p_\ell^* = 1.0, \quad p_k \in \{4.0, 1.0\}, \quad p_j^* = 40.0, \quad p_i = 40.0.$$

Thus the high-tier buyer at seller j sits at the clearing price, while mid-tier competitors remain strictly below p_j^* . The reported margins $(p_k - p_\ell^*) = 0$, $(p_j^* - p_k) = 36$, and $(p_i - p_j^*) = 0$ reveal a wide central gap: a single dominant tier clears seller j ,

Figure 6.3: Buyer 0 valuation curve and marginal diagnostics.



whereas seller ℓ is anchored by low-tier participation at a much smaller price.

The monotone relationship validated here provides empirical confirmation of Lemma 6.3. The experiment illustrates how buyers bridging sellers stabilize the market through consistent price ordering, even when capacities and bid magnitudes differ substantially. The asymmetry in seller revenues and capacities demonstrates how equilibrium adapts to network structure, with high-valuation buyers dominating smaller auctions and lower-tier participants anchoring larger ones.

Our next experiment allows us to observe the propagation of equilibrium constraints across overlapping influence shells, offering empirical evidence for Propositions 6.7 of this paper.

6.6.3 Connectivity

Further experimental results are aggregated as functions of the overlap percentage between buyer–seller pairs, revealing how market interdependence affects stability, bid prices, and efficiency. The framework also enables sensitivity analysis under perturbations to parameters such as ε , budget distributions, and the structure of influence sets.

In this experiment, connectivity was gradually increased to observe how equilibrium formation and price alignment change as the market transitions from isolated

to coupled seller networks. Starting from a sparse adjacency structure, buyers were allowed to participate in multiple auctions, creating overlaps that induced cross-seller influence and coupling of price dynamics.

Figure 6.4 illustrates the adjacency structure used in the experiment. Connectivity defines the feasible market domain $I_{\text{active}}(t) \subset \mathcal{B} \times \mathcal{L}$ that bounds all strategic interactions.

Figure 6.4: Adjacency and market connectivity for the 8×2 experiment. Connectivity is set at 50%.

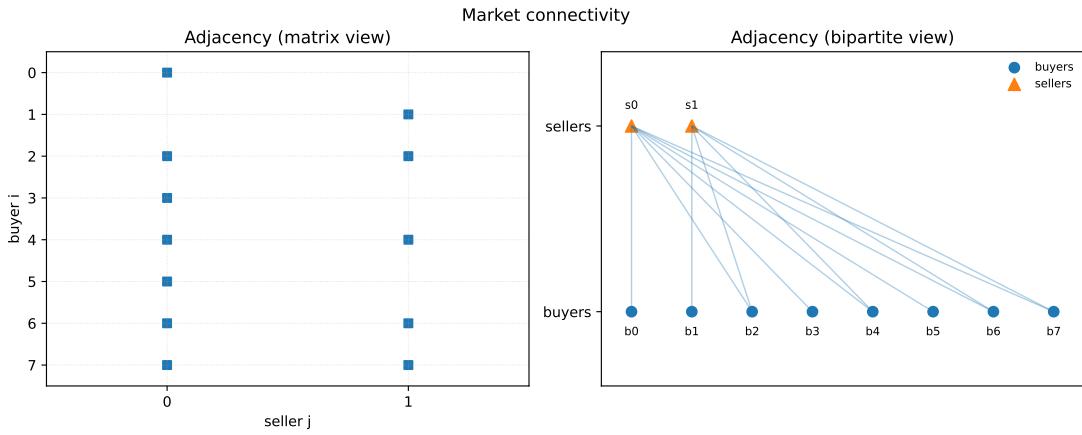
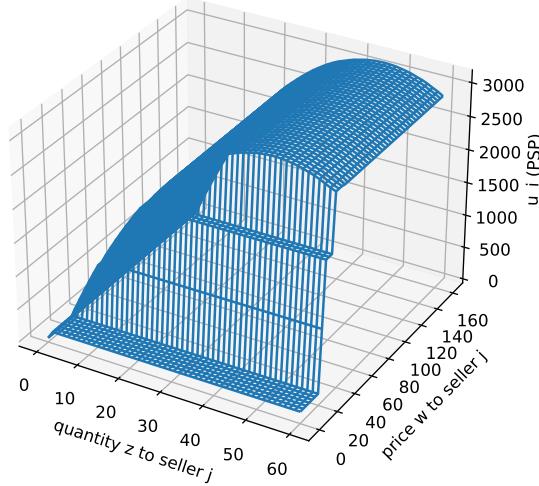


Figure 6.5 shows the utility surface for a single buyer–seller pair, as was presented in [45], here buyer 6 and seller 0, plotted over bid quantity z_i^j and price w_i^j . The surface depicts the buyer’s instantaneous utility $u_i(z_i^j, w_i^j) = \theta_i(z_i^j) - z_i^j w_i^j$ given the opponent bids and current market reserve. The concave ridge indicates the buyer’s optimal quantity at the current price level, while the lower regions show diminishing returns and cost-dominated outcomes. We see a stable interior optimum: movements along the quantity axis correspond to allocation changes, whereas movements along the price axis reflect valuation gradients.

Next, we present an algorithm: an iterative evaluation of the aggregate staircase $P_i(z, s_{-i})$. At each iteration t , buyers and sellers perform the following operations: Because each accepted update increases some buyer’s utility by a bounded discrete amount and the state space is finite, every sequence of threshold-improving updates

Figure 6.5: Single buyer–seller utility surface for buyer 6 at seller 0. The surface plots $u_i(z_i^j, w_i^j) = \theta_i(z_i^j) - z_i^j w_i^j$ over quantity z_i^j and unit price w_i^j , holding the opposing bids fixed at the snapshot.

Utility Surface (buyer 6, seller 0)



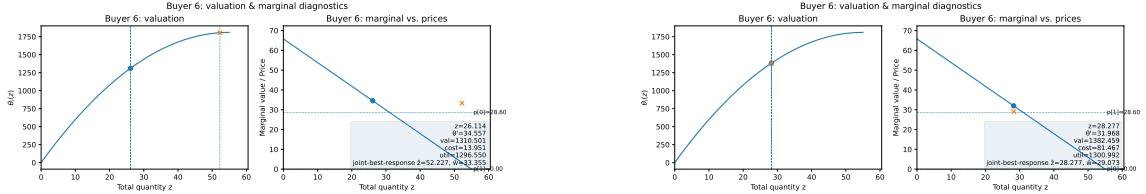
Algorithm 6.2 Buyer Update Dynamics under Bounded Participation

- 1: **Bid formation.** Each buyer i applies the opt–out map $q_i(a(s) : \mathcal{L}_i(t)) = [q_i^j(a)]_{j \in \mathcal{L}_i(t)}$ and selects the minimal–cost subset of sellers.
- 2: **Utility evaluation & Rebid.** Buyer i computes the utility increment $\Delta u_i(t)$ from its updated bids. Buyer i updates its bids iff $\Delta u_i(t) > \epsilon$.
- 3: **Allocation and clearing.** Sellers allocate proportionally at each price $p_*^j(t)$. Buyers at the cutoff price may receive partial allocations.
- 4: **Advance iteration.** Set $t \leftarrow t + 1$.

must terminate in an absorbing ϵ –NE region. We speculate that the induced dynamics are weakly acyclic: from any initial state, at least one finite improvement path leads to equilibrium.

Figure 6.6 are produced by the above construction. For Buyer 6 we evaluate the staircase $P_i(\cdot, s_{-i})$, compute (q_i, w_i) , perform the minimal–cost fill, and then read off the realized total Z_i and price $p^* := P_i(Z_i, s_{-i})$. The left panel shows $(Z_i, \theta_i(Z_i))$ on the concave valuation curve; the right panel shows θ'_i together with the dashed price level p^* . In the runs shown, the valuation is quadratic,

$$\theta_i(z) \approx az - \frac{b}{2}z^2, \quad \theta'_i(z) = a - bz, \quad a \approx 66, \quad b \approx 1.1,$$



(a) Constraint-limited: the joint best response lies on a feasibility boundary; $\theta'_i(Z_i) > p^*$ so the buyer would expand if capacity at p^* were available.

(b) Interior optimum: here $\theta'_i(Z_i) > p$, so the buyer would expand if possible. The joint-best-response is the point where $\theta'_i(z) = p$.

Figure 6.6: Buyer-level diagnostics under the Progressive Second Price (PSP) joint best response. Each panel shows the valuation $\theta_i(z)$ with the realized point $(Z_i, \theta_i(Z_i))$ and the marginal curve $\theta'_i(z)$ with a dashed line at p^* , illustrating the transition from constraint-limited to price-limited behavior along improvement paths.

so marginal value declines approximately linearly from $\theta'_i(0) \approx 66$ to near zero around $z \approx 60$. Two sellers induce a two-step staircase in P_i ; the three snapshots correspond to marginal price levels near $p^* \approx 32.1$ with $Z_i \approx 28.2$ (interior), $p^* \approx 28.6$ with $Z_i \approx 28.3$ (price-limited), and a high-availability case with $Z_i \approx 52.2$ where the buyer is constrained by feasibility at that price. These values are taken directly from the algorithm's output and no post-hoc smoothing is applied.

When the orange marker in the marginal panel lies above the dashed line, the realized point satisfies $\theta'_i(Z_i) > p^*$; the buyer would buy more at the prevailing price, but the minimal-cost fill has saturated feasible capacity at that price, so the joint best response is attained on a boundary of the feasible region rather than at marginal equality. When the marker sits on the dashed line, $\theta'_i(Z_i) = p^*$ holds and the allocation is locally efficient; here the construction returns an interior maximizer of $U_i(z) = \theta_i(z) - C_i(z)$. When the marker lies below the dashed line, $\theta'_i(Z_i) < p^*$ and any further increase in quantity would decrease utility; the best response is therefore at or near a participation boundary even though the valuation point on the left panel is well inside the curve. In every case the left panel places the realized point on $\theta_i(\cdot)$ because the algorithm maximizes value minus payment over the compact feasible set under the current price.

Figure 6.7 extends the analysis to the joint allocation space of the two sellers, each

Table 6.2: Buyer regimes and their economic interpretation.

Regime	Relation	Economic meaning
Constraint-limited	$\theta'_i(z^*) > p^*$	Supply prohibitive.
Equilibrium (interior)	$\theta'_i(z^*) = p^*$	Marginally efficient allocation.
Price-limited	$\theta'_i(z^*) < p^*$	Cost prohibitive.

point on the surface corresponds to a feasible distribution across sellers 0 and 1 under a uniform price $w = \theta'_i(Z_i)$. The height of the surface indicates utility under the split given the opposing bids present in the snapshot. The ridge along constant Z_i identifies the efficient split between sellers: solutions on the plateau indicate a local optimum; as the solution shifts below a ridge the buyer could improve utility by increasing its bid quantity, and the solution shifts toward the seller facing weaker opposing demand. Therefore, even with a uniform price tied to Z_i , the allocation decision remains two-dimensional due to how opponent demand and residual available resource shape the intersection of feasible pricing and allocations.

Table 6.3: Interpretation of ridges in the buyer's utility surface.

Ridge type	What it corresponds to
Sharp rise in u_i	A new seller step becomes available (increase in supply).
Sharp drop	Another buyer's bid dominates \rightarrow PSP second-price step kicks in.
Plateau	Both sellers saturated or prices equalize (local equilibrium).

To summarize our results, seller 0, with greater available resource ($Q_{\max} = 200$), cleared at a slightly higher price $p_0^* = 30.084$ than Seller 1 ($p_1^* = 28.597$). Despite this asymmetry, both sellers exhibited similar expected revenues ($E_0 = 44.76$, $E_1 = 31.16$) and low variance, attributed to 40% of buyers participating in both markets. The shared influence among these buyers synchronized seller behavior, leading to a nearly uniform price surface.

Buyer-level data shows that bridging buyers—particularly buyers 6 and 7—maintained bids across both sellers with marginal valuations (32.118, 33.355) close to Seller 0's clearing price. Their dual participation enforced cross-market coherence,

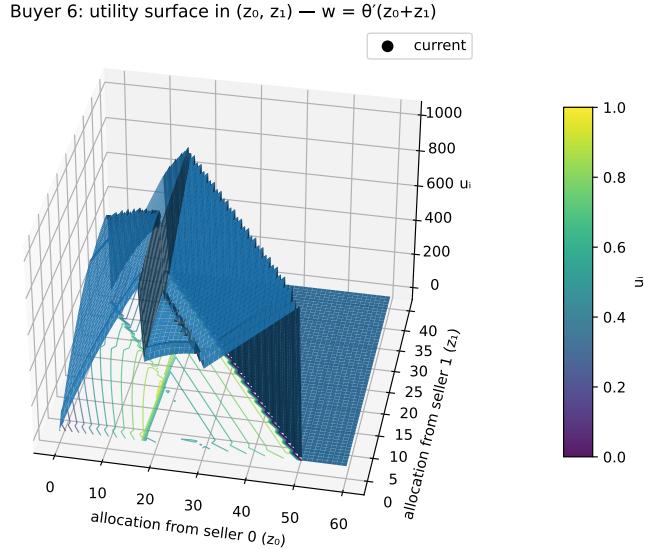


Figure 6.7: Shared-seller utility surface where buyer 6’s utility is a function of total requested quantity $Z_i = z_0 + z_1$; $u_i(z_0, z_1) = \theta_i(Z_i) - w(z_0, z_1)(Z_i)$ and $w = \theta'(Z_i)$: the feasible participation surface.

ensuring that no single auction could deviate significantly from the shared equilibrium.

Overall, increasing connectivity transforms the market from a set of independent price islands into a unified utility surface. Sparse configurations produce local equilibria with greater variance between sellers, while denser networks encourage faster convergence through influence propagation. These findings validate the theoretical expectation that shared influence sets promote global coordination and equilibrium alignment.

Statement on Supplementary Material The code for the experiments presented in this paper can be found at:

- <https://github.com/jkblazek/arXiv-2511.19225>

6.7 Conclusion and Future Work

This paper establishes a graph-theoretic framework for analyzing Progressive Second-Price (PSP) auctions, connecting decentralized market dynamics to structural properties of influence graphs. We formalized and expanded the concepts of influence sets and saturation, which together bound strategy spaces deterministically and ensure stable, truthful convergence in decentralized settings.

Our analysis relies on two levels of saturation, linked by the partial-ordering structure of bids; local saturation is a set-level best-response property. Establishing a formal fixed-point characterization of this process—perhaps using lattice or order-theoretic methods—remains an important direction for future work.

Our present analysis instead focuses on the constructive evolution of influence shells and the preservation of local monotonicity. Specifically, our approach demonstrates how recursive expansions of influence sets reveal market interactions across successive auction rounds. By introducing intra-round resolution via the τ_k steps, we provide a finer-grained analytical tool to model the internal dynamics of auction iterations, clarifying the subtle interactions between buyer strategies and seller pricing rules.

The establishment of monotonicity in bid updates via induced partial ordering ensures stable, non-oscillatory convergence under realistic, elastic valuation conditions. Our framework provides a robust method to anticipate market shifts, characterize equilibrium thresholds, and ensure consistent propagation of influence across dynamic network topologies.

Future research will explore several promising extensions in reserve price optimization. Could there be an optimal coordinated reserve vector, chosen using buyer feedback, that upholds Lemma 6.3? We start by defining an admissible reserve price region, where for fixed t , the admissible set of reserve profiles is

$$R(t) = \{p_* \in \mathbb{R}^J : p_*^j \in (\bar{p}^j(t), \underline{p}^j(t)) \ \forall j, \text{ and Lemma 6.3 holds}\}.$$

Thus, we may determine the existence of *coordinated reserve profile*, where, given

seller-side weights $w_j \geq 0$ at time t , is defined as any $p_*^{\text{coord}}(t) \in R(t)$ that maximizes

$$\Phi(p_*) = \sum_{j=1}^J w_j Q^j(t) p_*^j$$

over $R(t)$. We conjecture that at least one coordinated reserve profile $p_*^{\text{coord}}(t)$ exists for every $\epsilon > 0$. Moreover, every such profile preserves the interval ladder inequalities (6.19) and hence is consistent with local saturation of primary influence shells.

A formal proof of a best-response property is beyond the scope of this work at this time. Instead, we provide a sketch of the proof that would demonstrate the existence of a joint ϵ -best reply for a buyer participating in multiple concurrent auctions.

To form a coordinated reply at a common marginal price, we collect the sellers visible to buyer i under s_{-i} and compute their prices at a target marginal value w_i . Ordering these sellers by nondecreasing price and applying tentative allocation until the requested total is reached yields the minimal-payment split across auctions. Denote our buyer-level payment by

$$P_i(z, s_{-i}) := \inf\{y : Q_i(y, s_{-i}) \geq z\},$$

which we call an aggregate price staircase. The cumulative payment is

$$C_i(z; s_{-i}) := \int_0^z P_i(\zeta, s_{-i}) d\zeta.$$

Knowing by finite-valuation that $P_i(\cdot, s_{-i})$ is bounded, nondecreasing, and piecewise constant, this construction implements exactly the payments returned by PSP at the target marginal price and, among all feasible potential allocations with the same total quantity, attains the minimum payment.

First, we aggregate availability across auctions at a common marginal price by $P_i(z, s_{-i}) = \inf\{y : Q_i(y, s_{-i}) \geq z\}$. Finiteness in the number of bids ensures boundedness and right-continuity; the plateau condition $Q_i(y^-, s_{-i}) < z \leq Q_i(y, s_{-i})$ characterizes $P_i(z, s_{-i}) = y$. The cumulative payment $C_i(z) = \int_0^z P_i(\zeta, s_{-i}) d\zeta$ is continuous and convex. Consider

$$U_i(z; s_{-i}) := \theta_i(z) - C_i(z; s_{-i}).$$

Realize Z_i by sorting tiers by effective PSP price at level $p_i = \theta'_i(Z_i)$ and taking quantities until the sum equals Z_i ; PSP returns the same aggregate staircase, hence the same total.

Finally, integrating stochastic perturbations and noise into the PSP framework will deepen the realism of the model, allowing exploration of market resilience under uncertainty. Additionally, applying resistance distance [8] to the reserve profiles and diffusion-based influence models could yield deeper insights into influence propagation and market stability. Empirical validation through simulation and real-world decentralized applications, such as spectrum and bandwidth auctions, will be critical to validate and refine theoretical predictions and improve practical mechanism design.

Appendix: Market Shift Revealed by Partial Ordering

Example 6.9 (Market Shift Revealed by Partial Ordering). *In this example we model a simple reactive reserve update,*

$$p_*^j(t+1) = \max \left\{ p_*^j(t), \max_{i \notin \mathcal{B}^j(t)} p_i^j(t) + \epsilon \right\},$$

with $\epsilon > 0$ providing strict improvement for convergence. Thus the seller always keeps its internal bid strictly above the highest losing buyer, and the reserve price is nondecreasing in t . Alternative clearing-price rules that set p_*^j to the threshold $\chi^j(t)$ are equivalent for our results.

Consider a PSP auction market with two sellers L_1, L_2 and four buyers B_3, B_4, B_5, B_6 . Initial buyer-seller connections are represented by the adjacency matrix:

$$\mathbf{A}^{(0)} = \begin{array}{c|cc} & L_1 & L_2 \\ \hline B_3 & 1 & 0 \\ B_4 & 1 & 1 \\ B_5 & 1 & 1 \\ B_6 & 0 & 1 \end{array}$$

Auction Iteration $t = 1$

Seller L_1 initially receives bids from B_3, B_4, B_5 . Suppose initial bid prices are ordered as follows:

$$p_{B_3}^{L_1}(1) = 2.0 > p_{B_4}^{L_1}(1) = 1.8 > p_{B_5}^{L_1}(1) = 1.5.$$

Progressive allocation steps for L_1 :

τ_1 : B_3 allocated requested quantity, pays second-highest price 1.8. Seller updates reserve to $1.8 + \epsilon$.

τ_2 : B_4 allocated next available quantity, pays 1.5. Reserve updates to $1.5 + \epsilon$.

τ_3 : B_5 receives allocation, pays reserve $(1.5 + \epsilon)$.

Seller L_2 has bids from B_4, B_5, B_6 , with initial ordering:

$$p_{B_5}^{L_2}(1) = 1.9 > p_{B_4}^{L_2}(1) = 1.7 > p_{B_6}^{L_2}(1) = 1.4.$$

Progressive allocation steps for L_2 :

τ_1 : B_5 allocated, pays second-highest price 1.7. Reserve updates to $1.7 + \epsilon$.

τ_2 : B_4 allocated next, pays 1.4. Reserve updates to $1.4 + \epsilon$.

τ_3 : B_6 allocated, pays reserve $(1.4 + \epsilon)$.

Partial Ordering and Initial Influence Sets Initially, influence projections:

$$\pi \circ \varpi^{-1}(L_1) = \{B_3, B_4, B_5\}, \quad \pi \circ \varpi^{-1}(L_2) = \{B_4, B_5, B_6\}$$

Both sellers share buyers B_4, B_5 , forming an integrated influence structure.

Market Shift at $t = 2$: **Buyer** B_4 **increases bid on** L_1 Buyer B_4 increases their bid on seller L_1 to overtake B_3 :

$$p_{B_4}^{L_1}(2) = p_{B_3}^{L_1}(1) + \delta, \quad \delta > 0.$$

This triggers an immediate, asynchronous allocation decision at seller L_1 :

τ_1 : Seller L_1 allocates to the new highest bidder B_4 , who pays the second-highest bid $p_{B_3}^{L_1}(1)$. Reserve updates accordingly.

The new partial ordering on L_1 :

$$p_{B_4}^{L_1}(2) > p_{B_3}^{L_1}(2) > p_{B_5}^{L_1}(2).$$

Coupled Buyer Rebid Buyer B_4 observes this new allocation outcome and immediately updates their residual demand. Since buyers maintain consistent bid strategies across all sellers, buyer B_4 must now adjust their bid quantity for seller L_2 simultaneously:

$$\sigma_{B_4}^{L_2}(\tau_2) = Q_{B_4}(2) - a_{B_4}^{L_1}(\tau_1),$$

and submits this updated bid quantity at price $p_{B_4}^{L_2}(2)$.

Seller L_2 , asynchronously and independently from L_1 , now processes this rebid at its next local step:

τ_2 : Seller L_2 allocates quantity to buyer B_4 , charging the next-highest price among competing bidders (e.g., buyer B_5 's previous bid).

Propagation of Influence via Projection Mappings: The shift at L_1 updates the projection and partial ordering structure, immediately affecting the shared buyer set with seller L_2 . The updated projections remain:

$$\pi \circ \varpi^{-1}(L_1) = B_3, B_4, B_5,$$

$$\pi \circ \varpi^{-1}(L_2) = B_4, B_5, B_6,$$

but buyer B_4 's strategic rebid triggers a recomputation of reserve prices and rebidding decisions at L_2 , influencing buyer allocations in subsequent τ_k steps.

Thus, a local change in buyer B_4 's bid on one seller (L_1) creates a cascading effect through the partial ordering structure, inducing market shifts and influencing allocation outcomes on another seller (L_2). The explicit recomputation of partial orderings demonstrates clearly how strategic perturbations propagate through interconnected auction markets.

Chapter 7

The Effects of Latency in Progressive Second-Price Auctions

Abstract

The progressive second-price auction of Lazar and Semret is a decentralized mechanism for the allocation and real-time pricing of a divisible resource. This auction proceeds with each buyer asynchronously updating their bids by sending a message of the form (q, p) where q is the amount of resource desired and p is the marginal value the buyer places on that quantity of resource. Our focus is on how delays in the receipt of bid messages and randomness in the initial bids affect the ϵ -Nash equilibria obtained by the method of truthful ϵ -best reply. We then introduce additional features to the auction mechanism that increase the predictability of outcomes.

7.1 Introduction

Lazar and Semret [45] introduced a decentralized second-price auction for allocation of network bandwidth in which buyers asynchronously update bid prices in a way that maximizes their allocations while increasing their utility. The novelty in this auction mechanism is a buyer's individual price valuation is revealed only locally when that same buyer updates their bid with a new quantity and corresponding marginal value. Thus, the seller does not need the full price valuation curve of each buyer to allocate the resource in a way that maximizes the welfare or total value in the auction.

This progressive second-price auction was extended to the case where prices were

quantized by Jia and Caines in [38]. An opt-out selection method in the presence of multiple sellers was considered in [12]. For markets consisting of many simultaneous auctions a graph theoretic framework that characterizes influences resulting from which buyers buy from which sellers in terms of a bipartite graph is further explored in [10].

It is known provided ϵ is large enough that Algorithm 1 in [45] leads to bid updates that converge to an ϵ -Nash equilibrium—an equilibrium state where no buyer can individually increase individual their utility by more than ϵ through changing their bid. We call this the method of truthful ϵ -best reply. Proposition 3 in [45] shows further at such an equilibrium the auction is efficient, that is, the total value in the final allocation is within $\mathcal{O}(\sqrt{\epsilon\kappa})$ of optimal. Here κ is a constant related to the maximal rate of diminishing returns among all buyers in the auction.

We consider three possible sources of randomness in the progressive second-price auction: the initial bids of each buyer, the asynchronous order in which buyers update their bids and communication latency in the receipt of bid messages. Perceptions of fairness are influenced by predictability of outcomes; however, it is possible that randomness in the initialization and operation of the auction do not lead to significantly different equilibrium states. Our focus is on the degree the above sources of randomness influence both market aggregates and individual outcomes.

While aggregate quantities such as total value, utility and revenue are of interest to the seller, an individual buyer is more concerned about their individual allocations, cost and utility. Assuming the buyers do not change, then it is desirable that the allocation of resource between them also not change. Our goal is to characterize the expected outcomes and deviations in those outcomes for both the seller and individual buyers that arise from the distributed asynchronous nature of the progressive second-price auction under a realistic model of communication latency.

We begin our study by introducing an algorithm to construct ϵ -Nash equilibria consisting of truthful bids in which each buyer receives the exact quantity requested and consequently impose no externality on the other buyers in the auction. As a

result the second price for each buyer is zero and the revenue received by the seller also zero. Note that the existence of zero-revenue ϵ -Nash equilibria may be inferred from Semret [73] who observed that a buyer can obtain resource at zero cost by making a bid with a zero per-unit price valuation. On the other hand there are also ϵ -Nash equilibria which generate revenue much closer to the total value in the auction. Therefore, total revenue can vary significantly between different ϵ -Nash equilibria of the same auction. Since total utility is total value minus total cost, it follows that the total utility enjoyed by the buyers can also vary significantly between different equilibrium states.

To avoid the possibility of obtaining resource with a zero-price bid a reserve price was introduced in [73] less than which no quantity of resource is ever allocated. Under the assumption of scarcity, where the amount Q of resource available is less than what the buyers could use, there is a reserve price $P > 0$ such that all resource is allocated in the market. In this case setting the reserve price near the largest value such that all resource is allocated has the effect of reducing variability of outcomes. It follows that our algorithm for constructing a zero revenue equilibrium instead leads to an equilibrium with minimal revenue QP where each buyer pays exactly the reserve price for their allocation.

The structure of this paper is as follows: Section 7.2 introduces the formal setup, notation, and buyer–seller valuation framework. Section 7.3 analyzes zero-revenue equilibria, showing how truthful ϵ -best replies can lead to efficient but unprofitable outcomes, motivating the use of a reserve price. Section 7.4 examines the role of communication latency and asynchronicity on equilibrium convergence and variability through stochastic modeling. Finally, Section 7.6 concludes with implications for decentralized market design and directions for future research.

7.2 Preliminaries

Consider a progressive second-price auction consisting of one seller and multiple buyers. During the operation of the auction buyers place their bids asynchronously and

decentralization leads to communication latency in the receipt of bid messages. Thus, the outcome of the auction is affected by randomness in the communication delays and bid ordering. We also study what effect the initial bids have on the outcome of the auction.

Suppose a quantity Q of a divisible resource is to be allocated among a fixed set of buyers. Let \mathcal{I} be an index set such that $i \in \mathcal{I}$ represents a buyer able to participate in the auction. Each buyer has a price valuation $\theta_i(q)$ which identifies the value that buyer obtains upon receipt of a quantity q of resource. We suppose the valuation increases in q up to a maximum quantity \bar{q}_i and is concave to reflect diminishing returns. Assuming θ_i is differentiable, it follows that θ'_i is decreasing and the greatest marginal value buyer i ever places on the resource is given by $\bar{p}_i = \theta'_i(0)$.

Intuitively, $\theta'_i(q)$ represents the value of the next unit of resource after q units have been obtained. For definiteness, take θ_i as in [45] to be quadratic of the form

$$\theta_i(z) = \begin{cases} (1 - \frac{1}{2}z/\bar{q}_i)z\bar{p}_i & \text{for } z < \bar{q}_i \\ \frac{1}{2}\bar{q}_i\bar{p}_i & \text{otherwise.} \end{cases} \quad (7.1)$$

where buyer demand \bar{q}_i is sampled uniformly over $[50, 100]$ and the maximal marginal valuation \bar{p}_i is uniform on $[10, 20]$. Note that the resulting decrease in marginal value is bounded uniformly in i both above and below. Unless otherwise mentioned we consider 100 buyers that participate in the auction and keep their respective price valuations θ_i fixed as well as the amount of resource $Q = 1000$ available in the auction.

Remark 7.1. *Since $\bar{q}_i \geq 50$ for each buyer, the total quantity of resource valued by 100 buyers is at least 5000. Therefore $Q = 1000$ is guaranteed to be a condition of scarcity in which at least one buyer has the potential to increase the value of their allocation.*

Now consider the bids $(q_i, p_i) \in [0, \bar{q}_i] \times [0, \infty)$ from all buyers $i \in \mathcal{I}$ able to participate in the auction. Here q_i is the quantity requested and p_i the amount the buyer would be willing to spend to obtain an additional unit of resource, that is, the marginal value at q_i . Thus, a truthful bid always satisfies $p_i = \theta_i(q_i)$. We emphasize

that the full valuation curve θ_i for each buyer is not made available to the seller but revealed only locally through the marginal value p_i at the quantity q_i being bid on.

Denote the bid (q_i, p_i) submitted by buyer i as (i, q_i, p_i) . Then the set of all bids in the auction is

$$s = \{ (i, q_i, p_i) : i \in \mathcal{I} \}.$$

For definiteness we let the bid $(0, Q, P)$ represent the reserve price set by the seller, never update this bid and suppose $0 \in \mathcal{I}$. In this special case $\theta_0(a) = Pz$ is linear and $\theta'_0(z) = P$ constant. For notational convenience denote $\mathcal{I}_0 = \mathcal{I} \setminus \{0\}$.

If $k \neq i$ then buyer k is in competition with buyer i and the opposing bids against which buyer i must bid are

$$s_{-i} = \{ (k, q_k, p_k) \in s : k \neq i \}.$$

At a marginal price of y the resource available to buyer i is

$$Q_i(y, s_{-i}) = \max\{Q - z, 0\}$$

where $z = \sum \{ q_k : (k, q_k, p_k) \in s_{-i} \text{ and } p_k > y \}$. Note that if $y < P$ the reserve price takes effect, the bid $(0, Q, P)$ implies $q_0 = Q$ is an addend in z and consequently $Q_i(y, s_{-i}) = 0$.

Conversely, obtaining at least a quantity z of resource from the auction requires a bid with marginal price

$$P_i(z, s_{-i}) = \inf \{ y \geq 0 : Q_i(y, s_{-i}) \geq z \}.$$

Practically speaking, the progressive second-price auction aims for an ϵ -Nash equilibrium in which an individual buyer's utility can not be increased by more than ϵ . At such equilibria successful bid prices will be close but not in general the same. Note, however, that if prices are quantized—for example in dollars and cents—then ties become more likely, especially when ϵ is small. For the simulations in this paper we do not quantize prices and take $\epsilon = 5$ throughout. Even so, a tie breaking condition appears useful. Following Jia and Caines [38] we allocate bids whose prices are tied in proportion to the quantities requested.

Thus, the allocation to buyer i is

$$a_i(s) = (q_i/z) \min \{ q_i, Q_i(p_i, s_{-i}) \} \quad (7.2)$$

where $z = 1$ if $q_i = 0$ and otherwise

$$z = \sum \{ q_k : (k, q_k, p_k) \in s \text{ and } p_k = p_i \}. \quad (7.3)$$

Since the cost $c_i(s)$ to buyer i for participating in a second-price auction is the loss incurred by the other buyers due to that participation, changes in the allocations can be used to compute costs. In particular,

$$c_i(s) = \sum_{k \in \mathcal{I} \setminus \{i\}} p_k (a_k(s_{-i}) - a_k(s)).$$

Remark 7.2. *As our second-price auction is progressive the bids s_{-i} reflect the historical influence of buyer i leading up to the present time. Thus, $c_i(s)$ represents the externality obtained by omitting buyer i when determining the allocation and not the true externality that would have resulted if buyer i were never part of the auction. This difference between the true and instantaneous externality is what leads to the zero revenue ϵ -Nash equilibrium in the next section.*

Unlike [45], [12] and [38] we do not consider each buyer to be further constrained by a budget b_i that bounds the cost they are willing to incur for their resource allocation. This is for simplicity and to avoid situations where the second price changes in such a way that a previous bid needs to be updated to stay under budget.

Crucial to the progressive second price auction is the maximum allocation available in a market on the price valuation curve of buyer i . Again following [45] we consider the bid update rule given by

Definition 7.3. *The truthful ϵ -best reply is defined as follows. Let*

$$G_i(s_{-i}) = \{ z \in [0, \bar{q}_i] : z \leq Q_i(\theta'_i(z), s_{-i}) \}. \quad (7.4)$$

The ϵ -best reply to the bids s_{-i} is defined as

$$(v_i, \theta'_i(v_i)) \quad \text{where} \quad v_i = \sup G_i(s_{-i}) - \epsilon/\theta'_i(0).$$

The above bid is truthful since it lies on the graph of the marginal value. Taking v_i slightly less than the supremum ensures $v_i \in G_i(s_{-i})$ and increases the bid price $\theta'_i(v_i)$ due to decreasing marginal value (except in the case of the reserve price). The factor $1/\theta'_i(0)$ ensures the bid $(v_i, \theta'_i(v_i))$ is within ϵ of the best bid. In other words, the truthfull ϵ -best reply of Definition 7.3 provides a truthful bid that cannot be improved by more than ϵ while holding the opposing bids constant.

A proof of the above fact depends on two different ways of viewing the cost—the exclusion-compensation theorem—which we state here as

Theorem 7.4. *In a progressive second price auction under the above assumptions we have*

$$c_i(s) = \int_0^{a_i(s)} P_i(z, s_{-i}) dz. \quad (7.5)$$

Theorem 7.4 appears as equation (8) in [45] along with the observation that repeatedly allowing each buyer to update their ϵ -best replies converges to an ϵ -Nash equilibrium provided ϵ is large enough.

7.3 Zero-Revenue Equilibria

Before studying how communication latency and asynchronous bidding affect the outcomes of the progressive second-price auction, we first address variations in outcomes that result from Remark 7.2 on the computation of cost. This section examines zero-revenue equilibria. It demonstrates how such equilibria emerge under the standard mechanism and explains why a reserve price is required to ensure positive seller revenue.

To construct zero-revenue equilibria, we modify the method of truthful ϵ -best reply given as Algorithm 1 in [45] to alternate between the original bidding strategy and a compromise bid. Simulations indicate this modified algorithm still converges to an ϵ -Nash equilibrium, but in this case one in which each buyer obtains the exact allocation they asked for.

While the ϵ -best reply chooses a near optimal bid without regard for a buyers previous bid, a compromise bid is based on the requested and allocated resource from the previous bid. Namely, we have

Definition 7.5. *Suppose the bid (q_i, p_i) from buyer i receives an allocation of $a_i(s)$. The truthful compromise reply is defined as*

$$(v_i, \theta'_i(v_i)) \quad \text{where} \quad v_i = (a_i(s) + q_i)/2.$$

While compromise bids generally do not lead to an ϵ -Nash equilibrium on their own, interesting behavior happens when bids are alternated between the ϵ -best reply of Definition 7.3 and the compromise reply of Definition 7.5. The compromise bids reduce the externality while the ϵ -best replies advance towards an ϵ -Nash equilibrium.

First recall the method of truthful ϵ -best from [73] as

Algorithm 7.1 ϵ -best Reply

- 1: Each buyer evaluates the ϵ -best reply (Definition 7.3) and updates bids when utility increases by at least ϵ .
- 2: Repeat until no additional ϵ -best replies occur.

Remark 7.6. *At first it seems plausible that only a subset of ϵ -Nash equilibria might be obtained through the method of truthful ϵ -best reply and that the zero revenue case might not be among them. However, if the initial bids start at a particular equilibrium state, then the ϵ -best reply will remain at that equilibrium. Thus, the method of ϵ -best reply can terminate at any equilibrium state simply by starting at that equilibrium.*

Remark 7.7. *The ϵ -best reply for buyer i is constructed so their allocation $a_i(s) = q_i$; however, increasing the utility of buyer i will generally cause some of the opposing buyers to lose their part of their allocations. Thus, after an ϵ -best reply there may be $j \neq i$ such that $a_j(s) < q_j$. This reflects an externality imposed on buyer j . If buyer j makes a compromise bid in reply, this reduces the externality imposed by buyer i on buyer j while at the same time not reducing the allocation of buyer j . In particular, switching between ϵ -best and compromise bids tends to a ϵ -Nash equilibrium in which no other buyer imposes any externality on another buyer.*

Our modified algorithm may now be stated as the alternating ϵ -best with compromise bid method:

Algorithm 7.2 Cooperative ϵ -best Reply

- 1: Each buyer evaluates the ϵ -best reply (Definition 7.3) and updates bids when utility increases by at least ϵ .
- 2: Each buyer submits a compromise bid (Definition 7.5), independent of immediate utility gain.
- 3: Repeat until changes in compromise bids fall below tolerance and no additional ϵ -best replies occur.

Note that compromise bids are made whether or not they increase utility. We do not comment on the rationality of this method of bidding but instead view Algorithm 7.2 as a technique to obtain zero-revenue ϵ -Nash equilibria. In turn such equilibria demonstrate the need for a reserve price even though the allocations are near value maximizing.

To demonstrate the convergence of the alternating ϵ -best with compromise bid method to a zero-revenue ϵ -Nash equilibria, consider a simplified version of the progressive second-price auction with no communication latency or asynchronous bidding. In this case buyers bid round-robin and alternate between bid strategies. The only source of randomness comes from the initial bids. We further illustrate the effects of the reserve price by choosing different values of P ranging from 0 to 16.

Given an equilibrium state s obtained from Algorithm 7.2 we are primarily interested in the aggregate quantities of revenue, total value and total utility. These are given, respectively, by

$$S[c] = \sum_{i \in \mathcal{I}_0} c_i(s), \quad S[v] = \sum_{i \in \mathcal{I}_0} \theta_i(a_i(s)) \quad \text{and} \quad S[u] = S[v] - S[c].$$

To further understand the effects of the reserve price, we also compute an average bid price as

$$E[p] = \frac{1}{S[a]} \sum_{i \in \mathcal{I}_0} a_i(s)p_i \quad \text{where} \quad S[a] = \sum_{i \in \mathcal{I}_0} a_i(s).$$

We study how the above aggregate quantities depend on the initial bids, by averaging them over the ϵ -Nash equilibria corresponding to an ensemble of 100 randomly

chosen initial bids. Specifically, consider a random set of initial bids of the form $\{(i, q_i, \theta'_i(q_i)) : i \in \mathcal{I}\}$ where the q_i are independent and uniformly distributed over $[0, \bar{q}_i]$ for $i \neq 0$ and $q_0 = Q$. Let Ω be the underlying uniform sample space. Given $\omega \in \Omega$ define s_w to be the ϵ -Nash equilibrium obtained from Algorithm 7.2 starting at the initial bid corresponding to ω .

Now let $\mathcal{E} \subseteq \Omega$ be an ensemble of 100 independent realizations of the initial bids. Suppose X is either revenue, total value, total utility or bid price. The ensemble average and sample variance of X is given by

$$\langle X \rangle = \frac{1}{|\mathcal{E}|} \sum_{\omega \in \mathcal{E}} X_w \quad \text{and} \quad V[\langle X \rangle] = \frac{1}{|\mathcal{E}| - 1} \sum_{\omega \in \mathcal{E}} (X_w - \langle X \rangle)^2.$$

Here X_w indicates X measured at the ϵ -Nash equilibrium given by s_w . For example, if $X = S[c]$ then the ensemble-averaged revenue would be

$$\langle S[c] \rangle = \frac{1}{|\mathcal{E}|} \sum_{\omega \in \mathcal{E}} \sum_{i \in \mathcal{I}_0} c_i(s_w).$$

Table 7.1: The effects of reserve price on the bid price, total value, utility and revenue in the ϵ -Nash equilibria obtained from Algorithm 7.2 averaged over 100 different random initial bids. Except for the revenue corresponding to a zero reserve price, the standard deviations—not shown—were less than 1 percent of the averages.

Reserve Price	0	6	12	14	16
Bid Price	13.4024	13.4022	13.3745	14.124	16.1289
Total Value	15544.5	15544.6	15536.8	12623.1	5784.98
Total Utility	15544.5	9544.56	3536.76	1602.27	441.175
Total Revenue	10^{-13}	6000	12000	11020.8	5343.8

Table 7.1 reports the averages obtained through simulation of Algorithm 7.2 over an ensemble of 100 different initial bids. The standard deviations $V[\langle X \rangle]^{1/2}$ were consistently less than 1 percent of the average except for the zero-revenue case corresponding to the reserve price of $P = 0$ where both the average and deviation were numerically equal to zero. Note that for $P \leq 12$ the total revenue is equal to

QP but when $P \geq 14$ the total revenue decreases. This is because as that point the seller starts buying back the resource which does not contribute to revenue.

Remark 7.8. *Since $\theta'_i(z)$ is decreasing, it is invertible. Define*

$$f_i(y) = \begin{cases} (\theta'_i)^{-1}(y) & \text{for } y \in [0, \bar{p}_i] \\ 0 & \text{otherwise.} \end{cases}$$

If the reserve price P is such that $\sum_{i \in \mathcal{I}_0} f_i(P) < Q$ then seller buyback is guaranteed. A fundamental premise of the progressive second-price auction is that the seller does not know the full valuation of each buyer; therefore, we do not explore this condition further in this paper.

Given $w \in \Omega$ observe that the revenue-minimizing equilibrium s_w is contained in a neighborhood of truthful bids which are nearly revenue minimizing and also ϵ -Nash equilibria. It follows that if the initial bids are chosen randomly, then there is a positive probability that those bids are already at a ϵ -Nash equilibrium that is nearly revenue minimizing. We set $P = 12$ near the largest value such that all resource is allocated to reduce the variability in revenue between equilibrium states. We also fix the initial bids made by the buyers. This allows us to focus on the variability caused by asynchronous bidding and communication latency in the sections which follow.

7.4 Latency and Asynchronicity

In this section we characterize the effects of latency and asynchronicity on the outcomes of the ϵ -Nash equilibrium states reached by the distributed progressive second price auction using the method of ϵ -best reply.

We use a renewal processes to model both communication latency and asynchronicity in bid updates. Namely, we employ a sequence of independent Weibull-distributed random variables with probability density

$$\text{pdf}(x) = \frac{\beta}{\lambda} \left(\frac{x}{\lambda} \right)^{\beta-1} e^{-(x/\lambda)^\beta}$$

where β is the shape and λ the scale. The parameter β characterizes the delay regime: $\beta < 1$ corresponds to a heavy-tailed, light-traffic latency distribution (bursty

communication), whereas $\beta > 1$ represents a controlled, scheduled update process with more predictable timing. Inspired by Arfeen, Pawlikowski, D. McNickle and A. Willig [5], see also Arshadi and Jahangir [6], we model the interarrival times of bid messages using $\lambda_c = 1.0$ and $\beta_c = 0.75$ with a translation of $\delta_c = 0.1$ seconds. The decreasing hazard rate represents a bid sent under conditions of light traffic where most messages arrive with minimum latency but when lost experience exponential backoff of retransmission times.

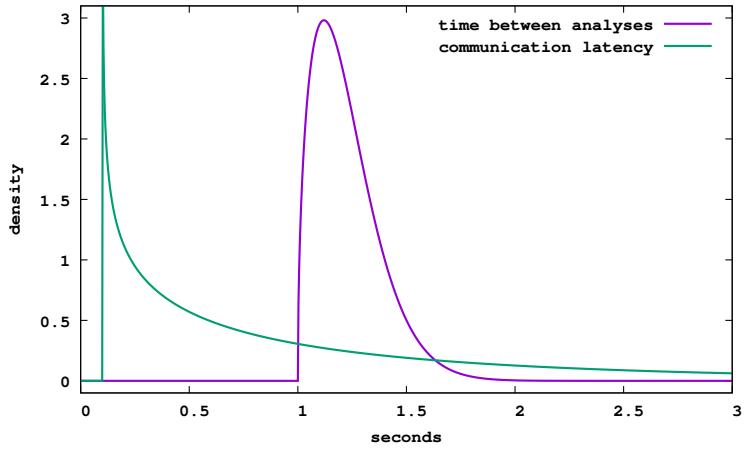


Figure 7.1: Comparison of the probability density functions governing the time between evaluation of bids and the communication latency to transmit a bid to the auction.

As the auction progresses the ϵ -best reply is evaluated by each buyer over intervals characterized by a Weibull distribution with $\lambda_e = 0.25$ and $\beta_e = 1.5$ translated by $\delta_e = 1.0$ seconds. The increasing hazard rate represents a controlled scheduling of market analysis that, in part, results from the assumption that each bid incurs a cost of ϵ that needs to be amortized before making the next bid.

Remark 7.9. *For each fixed buyer the intervals between market analyses follow a Weibull distribution; however, actually sending a message to update a bid is only performed when the utility can be increased by at least ϵ . Moreover, as the state of the auction gets closer to an equilibrium, the rate at which ϵ -best replies lead to a bid update appears to slow down.*

Let Δ_i^c represent the communication delay (time from bid placement to activation) and Δ_i^e be the evaluation interarrival delay (time between attempted rebids). These random variables are independent and distributed according to

$$\Delta_i^c \sim \delta_c + \text{Weibull}(\lambda_c, \beta_c) \quad \text{and} \quad \Delta_i^e \sim \delta_e + \text{Weibull}(\lambda_e, \beta_e).$$

Thus, a bid (q_i, p_i) computed at time t is observed at time $t + \Delta_i^c$. Similarly, a buyer who analyzes the current state of the market at time t will again compute their ϵ -best reply at time $t + \Delta_i^e$. The simulation operates as a discrete-time event system. Events are scheduled and processed in a priority queue, advancing the simulation clock t to the next event.

Figure 7.1 depicts the distributions of the time between bids and the communication latency. Even though the expected time between bid updates is much greater than the expected latency, the heavy tail corresponding to the shape parameter $\beta_c = 0.75$ implies there is a chance—due to network lag—that a new bid update might be considered before the previous bid has been received by the auction.

To characterize the effects of latency and asynchronicity on the outcomes of the progressive second-price auction we consider an ensemble \mathcal{E} of 100 realizations for the random processes given by Δ_i^c and Δ_i^e . Note that the ensemble used in Section 7.3 was over random initial bids. In this section we hold the initial bids fixed. Now, for $\omega \in \mathcal{E}$ let s_ω denote the equilibrium state obtained from the Algorithm 7.1 subject to the communication delays and bid updates specified by ω .

Figure 7.2 on the left depicts the ensemble averages of the bid prices and total utility as a function of the scale of the communications latency with λ_c ranging from the default value of 1.0 shown in Figure 7.1 up to 20. The shadows illustrate the standard deviations of the ensembles given by $V[\langle E[p] \rangle]^{1/2}$ and $V[\langle S[u] \rangle]^{1/2}$. The deviations are small while the bid prices and total utility are horizontal to within the errors represented by those deviations.

Note that the variance of the ensemble average $V[\langle E[p] \rangle]$ is different than ensem-

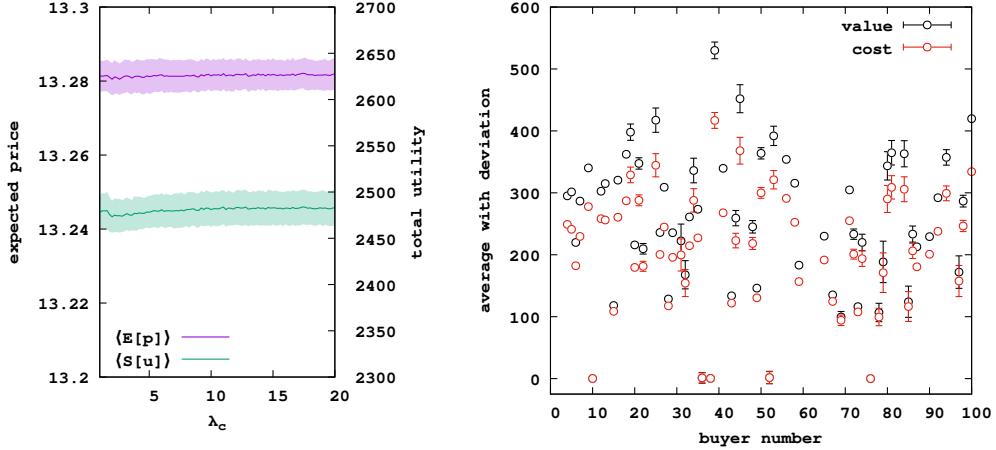


Figure 7.2: Left shows that changing the scale λ_c of the communication latency has minimal effect on the ensemble-averaged price and total utility received by all buyers. Right shows the average value and cost for each individual buyer in the case $\lambda_c = 1$.

ble average of the variance, which instead is given by

$$\langle V[p] \rangle = \frac{1}{|\mathcal{E}|} \sum_{\omega \in \mathcal{E}} \left(\frac{1}{S[a]} \sum_{i \in \mathcal{I}_0} a_i(s)(p_i - E)^2 \right) \Big|_{s=s_w}.$$

Figure 7.2 on the right depicts the predictability of individual outcomes for the 62 buyers who received allocations in the market. The remaining 28 buyers consistently did not receive allocations and are not shown. No resource was purchased by the seller. We study the individual values and costs given by

$$\langle v_i \rangle = \frac{1}{|\mathcal{E}|} \sum_{\omega \in \mathcal{E}} v_i(s_w) \quad \text{and} \quad \langle c_i \rangle = \frac{1}{|\mathcal{E}|} \sum_{\omega \in \mathcal{E}} c_i(s_w)$$

and the standard deviations in these averages. Some buyers experienced much higher deviations in outcomes compared to others. This pattern was repeatable. Similar results but affecting different individuals were obtained for other populations of buyers.

Also computed but not shown, the deviation in individual utility

$$V[\langle u_i \rangle]^{1/2} = \left(\frac{1}{|\mathcal{E}| - 1} \sum_{\omega \in \mathcal{E}} (u_i(s_w) - \langle u_i \rangle)^2 \right)^{1/2}$$

was uniformly small for all buyers. Since demand is perfectly elastic, then arguably

the buyer utility given by $u_i = v_i - c_i$ is more important than either the value or cost on their own.

7.5 Different Latencies

This section studies whether buyers who compute their ϵ -best reply more frequently and experience less latency in their bid messages have any advantage over buyers who analyze the market less frequently and whose bid messages suffer greater latencies. Consider a population of 100 buyers with valuation curves such that

$$\theta_{i+50} = \theta_i \quad \text{for } i = 1, 2, \dots, 50.$$

Thus, the first 50 buyers are identical twins of the last 50 buyers with one difference: The last 50 buyers are lazy and evaluate the market 17 times less frequently and experience 17 times more latency in their bid messages.

Specifically, the first 50 buyers keep the same bid evaluation frequency and latency as in Table 7.1 while the delay and scale parameters for the time between bids and communication latency for the last 50 buyers are multiplied by a factor of 17. The last 50 buyers evaluate the market 17 times more slowly and experience 17 times more latency in their bid messages.

Figure 7.3 depicts the outcomes of the first 50 buyers compared their identical but lazy twins—the last 50 buyers. The graph on the left shows the transient part of the utility received over time for representative pair of buyer twins. After making an ϵ -best reply, either twin obtains the same utility from the market as the other. However, the difference in the time scales allows the industrious twin to maintain non-zero utility for a larger percentage of the time. As the equilibrium state is reached—not shown—the rate of bid updates slows down so much that the fact that one twin is 17 times slower than the other no longer matters.

In the end, both twins obtain essentially the same utility at the equilibrium state. The right graph illustrated the statistics for the same twin pairs taken over an ensemble \mathcal{E} of size 100. The outcomes in terms of individually received utility

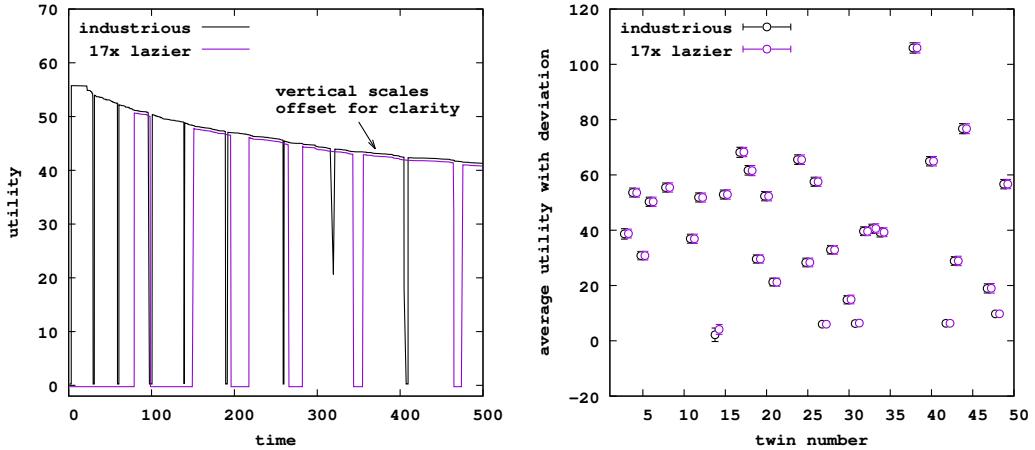


Figure 7.3: The outcomes for lazy buyers who evaluate the market 17 times less frequently and experience 17 times the latency in their bid messages compared to an equal number of industrious buyers with identical valuations.

depicted on the right are indistinguishable between the industrious and lazy twin. Note also there is very little deviation in outcomes due to the asynchronous nature of the individual market evaluations and the communication latencies in the distributed second-price auction.

7.6 Conclusions

We have introduced an algorithm that constructs ϵ -Nash equilibria consisting of truthful bids in which each buyer imposes zero externality on the other buyers and demonstrated through simulation that this algorithm converges. Although it is known that the method of truthful ϵ -best replies stated as Algorithm 1 in [45] may fail to converge unless ϵ is large enough, it is possible the introduction of compromise bids removes the condition on ϵ . A theoretical proof of this fact currently seems out of reach, however, if one assumes the ϵ -best replies converge to an ϵ -Nash equilibrium without the compromise bids, then it may be possible to show our new algorithm continues to converge when alternating these bids with compromise bids.

It is remarkable that the individual utility and bid prices at the ϵ -Nash equi-

libria have so little dependency on the random elements in the market mechanism of the progressive second-price auction. Seller revenue was further stabilized by the introduction a reserve price.

Chapter 8

An Analysis of a Progressive Second-Price Multi-Auction Market with Perfect Substitute and Perfect Information

Abstract

We present a dynamic extension of the Progressive Second-Price (PSP) auction for multi-auction markets in which buyers participate across multiple sellers through a common marginal price. Within this structure, we extend the exclusion-compensation principle to a global setting, showing that aggregate externalities can be represented as a price-ordered composition of local staircases. This formulation provides a unified view of valuation, allocation, and cost in distributed markets and establishes a rigorous foundation for future work: the establishment of a ϵ -best-response for multi-auction markets with perfect information and perfect substitute.

8.1 Introduction

The progressive second-price (PSP) auction, introduced by Lazar and Semret [45], is a decentralized mechanism designed for resource allocation. PSP satisfies truthfulness, individual rationality, and social welfare maximization while operating without centralized knowledge of every buyers' individual valuation. Unlike centralized auctions, PSP unfolds as a distributed process in which buyers and sellers interact locally

through repeated bidding rounds. At its core, each bid is a message consisting of a quantity and the buyer’s unit-price marginal valuation at that quantity. Messages that communicate a linear approximation of the valuation curve rather than the full valuation bounds the cost of communication and ensures tractability in cases where the full curve is unknown even to the buyer or changes over time due to external influences.

Classical Vickrey–Clarke–Groves (VCG) auction theory [83, 24, 32] established that charging each participant the externality they impose on others results in truthful bidding. This is known as the exclusion–compensation principle: a participant pays for the cost their presence creates by displacing others from their allocations. Truthful bidding is intuitively a rational strategy because payments align with marginal effects on others’ welfare. The classical theory assumes every buyer can report their entire valuation function to a central planner. This full revelation enables efficient allocation and exclusion–compensation pricing, but relies on centralized information.

By contrast, PSP restricts communication to a two-dimensional message space. Buyers do not reveal their full valuation functions in a single step. Instead, they gradually reveal information across rounds as the auction progresses through an ordered pair (q, p) . This iterative process is what makes the mechanism progressive: truthful valuations are approached incrementally through repeated bids, each reflecting a more accurate signal of willingness-to-pay. The progressive structure is not merely a design choice, but a necessity imposed by the restricted message space, the costs of distributed communication, and the progressive consumption of an allocated resource.

Extending PSP to multiple concurrent auctions raises new challenges. Each market has its own reserve prices and allocation dynamics, and buyers may face a variety of participation costs such as membership fees, per-round bid fees, or cumulative transaction costs. Rational buyers may choose to avoid participation in some markets when the expected utility does not exceed these fees. We formalize this behavior through an opt-out function, which filters a buyer’s desired allocation according to

cost and benefit. This defines a bounded participation rule: buyers update their bids only when the changing externality causes the previous bid to violate budget constraints or when the improvement in utility exceeds participation costs and tolerance thresholds.

The presence of multiple auctions also complicates the structure of the message space. While each bid is an ordered pair, a buyer active in several markets must communicate multiple bids. This raises questions about complexity, latency, and overhead: does simultaneously bidding in n auctions incur a fee of $n\epsilon$ where ϵ is the per-auction bid fee? Does each auction have a different bid fee? How do different bid communication latencies and costs affect which auctions a buyer participates in? In fragmented markets, message passing may incur delays or overheads that further shape equilibrium. For simplicity, our model begins with minimal assumptions: we treat the multi-auction system as a platform for allocating resources in which there is no latency when communicating bids—abstracting away from the network topology. This allows us to focus on the equilibrium and incentive properties of the auction and opt-out mechanism. A study of the effects of latency in the communication of bids appears in [11].

We remark that changing market conditions naturally affect the allocation of a consumable resource over time, in particular, buyers may enter or exit the market or their needs change. We assume that the timescales over which the market changes are large compared to the speed at which the PSP auction converges. We therefore assume the allocation of resources obtained by the auction represents a quasi-static equilibrium that slowly changes over time.

Auction theory provides mechanisms for allocating goods and resources among agents with private valuations. Traditional second-price auctions, introduced by Vickrey [83] and extended to social welfare optimization by Clarke [24] and Groves [32], incentivize truthful bidding by charging each participant the externality their presence imposes on others. This exclusion-compensation principle ensures that truthful bidding is rational and aligns individual incentives with social welfare. However, these

mechanisms assume centralized information and allow bidders to submit complete valuation functions in a single step.

To address decentralized environments, Lazar and Semret [45] introduced the progressive second-price auction as an iterative, message-passing algorithm. Each bid is restricted to a two-dimensional message (price and quantity). This restriction limits communication overhead but also means buyers cannot reveal their full valuation functions in one step. Instead, valuations must be progressively revealed over successive rounds. The progressive structure is therefore intrinsic to the mechanism: it emerges from the restricted message space and the costs of distributed communication. Our work builds on this foundation and asks how the message space and communication costs interact when multiple auctions run simultaneously.

In Lazar and Semret’s original analysis, the existence of an ϵ -Nash equilibrium relies on a single global ϵ parameter, interpreted as a market-wide tolerance for profitable deviation. In multi-auction PSP markets with participation fees and potential communication overhead, however, ϵ may take multiple forms: a flat entry cost, a per-round cost, a cumulative cost across rounds, or even additional terms reflecting message-passing complexity. Each of these modifies the decision rule by which buyers accept or reject new bids. Consequently, equilibrium must be reinterpreted as an ϵ -Nash equilibrium where ϵ depends on heterogeneous costs across markets.

A main contribution of this paper is to show that the exclusion-compensation principle, which underpins truthfulness, continues to hold in multi-auction PSP markets. This ensures that the externality created by each buyer’s presence is priced correctly even when participation costs are heterogeneous.

The remainder of this paper is structured as follows. Section 8.2 presents the rules of the PSP auction discussed in this work, including modifications given in [12, 10, 70]. Section 8.3 demonstrates the adherence of the multi-auction market to the second-price rule, while Section 8.4 explores the possibility of an absorbing ϵ -Nash equilibrium region and discusses the design of a coordinated best-response. Section 8.5 is a simulation of a set of coupled auctions, showing convergence even with buyers

operating independently. Finally, Section 6.7 summarizes and outlines directions for future work in theoretical and practical applications. Namely, we determine the proof of existence and uniqueness for the ϵ -Nash equilibrium region to be incomplete, and outline this direction for future work.

8.2 PSP Auction Rules

The PSP auction is a decentralized mechanism in which buyers iteratively submit bids to sellers, and sellers update reserve prices and allocations based on received bids. Each auction operates locally, and coordination emerges through repeated interactions across the market graph. In what follows, we define the bid structure, auction dynamics, pricing rules, allocation strategies, and participation behavior that govern the PSP mechanism.

8.2.1 Bid Structure and Strategy Space

Let $\mathcal{I} = \mathcal{B} \cup \mathcal{L}$ denote the set of all agents, partitioned into buyers and sellers. Each seller $j \in \mathcal{L}$ manages a local auction for a divisible resource, and each buyer $i \in \mathcal{B}$ may submit bids to a subset of sellers. The bid profile of auction j is given by the column vector s^j with entries s_i^j , where $(i, j) \in \mathcal{B} \times \mathcal{L}$. A bid

$$s_i^j = (q_i^j, p_i^j) \in S_i^j = [0, Q^j] \times [0, \infty)$$

represents a single interaction between buyer i and seller j , where q_i^j is the quantity requested by the buyer and p_i^j is the unit price offered.

We call the array of active bids \mathbf{s} . This array generates a directed graph, drawing edges between vertex pairs engaged in active bidding. We reserve *vertex* for graph theory and use *node* only when referring generically to buyers or sellers as economic agents. If there is no interaction between two players i and j , then we set $s_i^j = (0, 0)$, and zero the associated entry of the biadjacency matrix. Each nonzero entry defines an edge in the buyer-seller bipartite graph. Denote by $s_i = [s_i^j]_{j \in \mathcal{L}}$ the i -th row of the array \mathbf{s} . Note that s_i represents the bids from buyer i in all the auctions hosted by

the sellers. Let s_{-i} denote the array \mathbf{s} with the i -th row removed. This array consists of entries s_k^j where $k \in \mathcal{B} \setminus \{i\}$ and $j \in \mathcal{L}$. Thus, s_{-i} represents the bid profiles of buyer i 's potential opponents.

The strategy space for auction j is the set of all bid vectors $S^j = \Pi_{i \in \mathcal{B}} S_i^j$, and represents the space of all possible ways in which a player can play a game (also known as the strategy set). Similarly, the (full) strategy space for buyer i is defined as all possible bids at all auctions,

$$\left\{ S_i = \Pi_{j \in \mathcal{L}} S_i^j, \quad S_{-i} = \Pi_{j \in \mathcal{L}} (\Pi_{k \neq i \in \mathcal{B}} S_k^j) \right\}.$$

where the buyer valuation θ of quantity z is a second-order parabolic function,

Table 8.1: Basic sets and notation for a bundle of J independent PSP auctions

Object	Single auction	Multi-Auction
Quantity	q_i	$\sum_j q_i^j$
Price	$p_i = \theta'_i(q_i)$	$p_i = \theta - I'(\sum_j q_i^j)$
Bids	$s = \{(i, q_i, p_i) : i \in \mathcal{I}\}$	$s = \{(i, j, q_i^j, p_i) : i, j \in \mathcal{A} \subset I \times J\}$
Opposing	$\{(k, q_k, p_k) : k \neq i\}$	$\{(k, j, q_k^j, p_k^j) : k \neq i\}$

$$\theta_i(z) = \frac{\theta_i}{2}(z \wedge q_i)^2 + \theta_i q_i(z \wedge q_i),$$

where \wedge represents the minimum, and \mathcal{A} is an index set representing the bipartite connections of the buyers and sellers, where $I = |\mathcal{B}|$, $J = |\mathcal{L}|$.

Each buyer will know the available quantity and the seller's updated reserve price for each market in which they bid. Buyers act strategically by selecting sellers, adjusting bid quantities, and choosing whether to participate based on their expected ability to satisfy demand. In the PSP framework buyers cannot reveal their entire valuation functions in a single step; instead they must request allocations iteratively. To regulate this behavior we introduce a bounded participation rule, as in [12], which endogenously limits the set of sellers a buyer engages with.

8.2.2 Residual Quantity and Allocation

As a market with perfect but incomplete information, sellers can only gain information about demand by observing buyer behavior in their *local* auctions. Sellers are not associated with an opt-out function and update only on information revealed by buyers who have not opted out. In each iteration, every seller completes one update of its local auction.

For each seller j , the reserve price $p^j(t)$ is the price at which seller j is indifferent between selling her final unit of resource and keeping it. Equivalently, the seller may be viewed as submitting an internal bid $(Q^j, p^j(t))$ on her own auction. At the end of each round t , the reserve price is updated with information from the set of active bids, where $\mathcal{B}^j(t)$ is the set of buyers who win strictly positive allocations at seller j in round t , and $\epsilon > 0$.

We define the clearing price at seller j to be the smallest price at which aggregate awarded quantity meets available quantity:

$$\chi^j(t) = \min \left\{ y : \sum_{k: p_k^j(t) > y} q_k^j(a(t)) \geq Q^j(t) \right\}. \quad (8.1)$$

Any residual supply must therefore be allocated among bids that tie at prices just above $\chi^j(t)$, after higher-priced bids are filled. Let

$$\underline{p}^j(t) := \min \{ p_i^j(t) : i \in \mathcal{B}^j(t) \}, \quad \bar{p}^j(t) := \max \{ p_i^j(t) : i \notin \mathcal{B}^j(t) \}, \quad (8.2)$$

be the lowest winning and highest losing bid prices at seller j , and where buyers *not* in $\mathcal{B}^j(t)$ receive zero allocation at seller j . The clearing price satisfies

$$\bar{p}^j(t) < \bar{p}^j(t) + \epsilon \leq \chi^j(t) \leq \underline{p}^j(t) - \epsilon < \underline{p}^j(t)$$

whenever there is at least one winning and one losing bidder at seller j . In particular, $\chi^j(t)$ lies in the open interval between the highest losing and lowest winning bid. At equilibrium, the reserve price $p^j(t)$ coincides with the clearing price at seller j , i.e., the clearing price implied by the PSP allocation rule.

The allocation is determined by the availability implied by competition at the current bid price. The maximum available quantity of data in auction j at unit price y given s_{-i}^j is

$$q^j(y, t) = \left[Q^j(t) - \sum_{k: p_k^j(t) > y} a_k^j(t) \right]^+, \quad (8.3)$$

be the residual quantity remaining after all buyers who bid strictly above y have been served. We refer to the tie-splitting rule originated in the analysis of quantized PSP auctions by Qu, Jia, and Caines [70],

$$a_i^j(t) = \min \left\{ a_i^j(a), \frac{a_i^j(a)}{\sum_{\ell: p_{\ell}^j = y} a_{\ell}^j(a)} q^j(y, t) \right\}. \quad (8.4)$$

For each buyer–seller pair (i, j) at time t , $a_i^j(t)$ is the amount that seller j allocates to buyer i .

We remark that the reserve price $p^j(t)$ that lies in the margin interval determined by the bids

$$\bar{p}^j(t) < p^j(t) < \underline{p}^j(t), \quad (8.5)$$

whenever both $\bar{p}^j(t)$ and $\underline{p}^j(t)$ are defined, and we deliberately leave the precise rule for selecting $p^j(t)$ within the interval (8.5) unspecified. In particular, admissible choices include

$$p^j(t) = \chi^j(t), \quad p^j(t) = \bar{p}^j(t) + \epsilon, \quad p^j(t) = \underline{p}^j(t) - \epsilon,$$

provided that reserve price updates lie within ϵ and the resulting sequence $\{p^j(t)\}_t$ is nondecreasing. Could there be an optimal coordinated reserve price, chosen using buyer feedback, that upholds the best-response dynamics? We admit there may be a *coordinated reserve profile* that improves potential revenue for participating sellers. We leave this conjecture for future work.

8.3 Exclusion–Compensation

The cost to the buyer adheres to the second-price rule for each local auction, this is the “social opportunity cost” of the PSP pricing rule. For a fixed seller j , buyer i ’s

exclusion-compensation (EC) payment equals the loss imposed on other buyers at that seller. For a fixed auction j we write

$$c_i^j(s) = \sum_{k \neq i} p_k^j \left[a_k^j(0; s_{-i}^j) - a_k^j(s_i^j; s_{-i}^j) \right], \quad (8.6)$$

where s_{-i}^j collects the opponents' bids at seller j and for each buyer $k \neq i$ the term $a_k^j(0; s_{-i}^j) - a_k^j(s_i^j; s_{-i}^j)$ denotes the reduction in allocation to k due to buyer i 's participation at j . Equivalently, the integral form uses the opposing buyers' piecewise-constant marginal price function $P^j(\cdot, s_{-i}^j)$ built from s_{-i}^j ,

$$c_i^j(s) = \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz, \quad (8.7)$$

which holds true locally at each auction, where the opposing bids are calculated against the allocated resource to buyer i . Each buyer in the opposing bid vector contributes a price from their awarded allocation that enters into the externality of another.

8.3.1 (Local) Externality Cost.

We examine the continuity of pricing given local residual quantities. For a fixed seller j , the opponents of buyer i are ordered by decreasing prices $p_{(1)}^j > p_{(2)}^j > \dots > p_{(M^j)}^j \geq 0$ where $p_{(n)}^j$ denotes the n -th distinct price level among all buyers $k \neq i$ at seller j . At each price tiers $p_{(n)}^j$, there may be one or more buyers posting that same price. The corresponding quantities requested by those buyers are q_k^j for all k such that $p_k^j = p_{(n)}^j$. Fix buyer i and a strategy profile s ; let $s_{-i} \in R^{(I-1) \times J}$ denote the matrix of other buyers' bids, with s_{-i}^j its j th column. For each seller j , the local inverse price function $P^j(\cdot, s_{-i}^j) : [0, \infty) \rightarrow [0, \infty)$ maps an allocation level z for buyer i at seller j to the smallest price level y such that the residual quantity available to i at price y (after excluding opponents with strictly higher price and resolving ties at y according to the mechanism's tie-breaking rule) is at least z . Equivalently, if

$$Q_i^j(y, s_{-i}^j) = \text{the quantity available to } i \text{ at seller } j \text{ at price } y,$$

then we use the generalized inverse

$$P^j(z, s_{-i}^j) = \inf\{y \geq 0 : Q_i^j(y, s_{-i}^j) \geq z\},$$

which is a right-continuous, nonincreasing step function of z . The corresponding residual quantities are

$$\zeta_0^j = Q^j, \quad \zeta_n^j = \left[Q^j - \sum_{m=1}^n \sum_{k \neq i: p_k^j = p_{(m)}^j} q_k^j \right]^+, \quad (8.8)$$

for $n \geq 1$. The amount of resource available at price $p_{(n)}^j$ is then $\zeta_{n-1}^j - \zeta_n^j \geq 0$. The local inverse price function is

$$P^j(z, s_{-i}^j) = p_{(n)}^j \quad \text{for } z \in (\zeta_n^j, \zeta_{n-1}^j].$$

Buyer i 's per-seller exclusion-compensation cost is the area under the function

$$c_i^j(s) = \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz = \sum_{m=n+1}^{M^j} p_{(m)}^j (\zeta_{m-1}^j - \zeta_m^j) + p_{(n)}^j (a_i^j(s) - \zeta_n^j), \quad (8.9)$$

for $a_i^j(s) \in (\zeta_n^j, \zeta_{n-1}^j]$. The summation adds to the cumulative cost of all the bids at strictly higher prices from buyers receiving their full allocation, whereas the final term $p_{(n)}^j (a_i^j(s) - \zeta_n^j)$ corresponds to the partial allocation at the threshold price, representing the portion of resource assigned to buyer i .

At the threshold price for seller j , (8.7) collapses to $c_i^j(s) = p_*^j a_i^j(s)$, since each unit awarded to buyer i at that price displaces another buyer who would have otherwise received it at the same threshold. It follows that $P^j(\cdot, s_{-i}^j)$ is constant over the awarded interval. In general, the exclusion-compensation rule charges every awarded unit as if it displaces other bids valued at the threshold price $p_*^j(t)$.

This local view refers to the effect of buyer i 's participation at a single seller j , where EC is computed using the local inverse price function $P^j(z, s_{-i}^j)$. The aggregate (or global) view refers to the buyer's total EC cost over all sellers in which it participates, computed by concatenating and price-ordering the local steps to form $P_i(z, s_{-i})$. In the following analysis we will assume a coordinated clock across all auctions.

8.3.2 (Global) Externality Cost

We formalize how the aggregate inverse price function is obtained by price-ordering and concatenating the local ones, and (iii) prove that the buyer's total exclusion-compensation (payment) can be computed either as a sum of local integrals or as a single global integral. This sets up the multi-auction utility rule used later.

For each ordered price y , we have that $P_i(z, s_{-i})$ is defined for the range of z corresponding to the total resource available from all sellers at that price, i.e.,

$$z \in \left(\sum_{p_{(m)}^j > y} (\zeta_{m-1}^j - \zeta_m^j), \sum_{p_{(m)}^j \geq y} (\zeta_{m-1}^j - \zeta_m^j) \right]. \quad (8.10)$$

Define the aggregate available quantity at the price y as

$$Q_i(y, s_{-i}) = \sum_{j \in \mathcal{L}_i} Q_i^j(y, s_{-i}^j) \quad (8.11)$$

and the price at which we obtain a total of at least z units of resource among all available auctions

$$P_i(z, s_{-i}) = \inf \{ y \geq 0 : Q_i(y, s_{-i}) \geq z \}, \quad (8.12)$$

where because $Q_i(y, s_{-i})$ is a right-continuous, nondecreasing step function with finitely many jumps at $\{p_{(m)}^j\}$, the infimum is attained. To build the aggregate inverse price function $P_i(z, s_{-i})$, we collect all local price-quantity pairs with cumulative quantity breakpoints $0 = \zeta_{M^j}^j < \zeta_{M^j-1}^j < \dots < \zeta_0^j$ (where $P^j(z) = p_{(m)}^j$ on $(\zeta_m^j, \zeta_{m-1}^j]$), then the local price-quantity pairs generalize to

$$\left(p_{(m)}^j, (\zeta_{m-1}^j - \zeta_m^j) \right), \quad m = 1, \dots, M^j, \quad (8.13)$$

and $P_i(\cdot, s_{-i})$ is the staircase obtained by taking the union (multiset) of all local price-quantity pairs from the sellers and ordering the pairs by nonincreasing price.

We define the global inverse price function,

$$P_i(z; s_{-i}) = \inf \left\{ y \geq 0 : \sum_j \sum_{m: p_{(m)}^j \geq y} (\zeta_{m-1}^j - \zeta_m^j) \geq z \right\}, \quad (8.14)$$

representing the total resource available at prices greater than or equal to y .

Lemma 8.1 (Aggregate exclusion-compensation). *For every strategy profile s ,*

$$\sum_{j \in \mathcal{L}_i} \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz = \int_0^{\sum_{j \in \mathcal{L}_i} a_i^j(s)} P_i(z, s_{-i}) dz. \quad (8.15)$$

and therefore the aggregate payment decomposes additively,

$$c_i(s) = \sum_{j \in \mathcal{L}_i} c_i^j(s). \quad (8.16)$$

Proof. For each seller j , the local inverse $P^j(\cdot, s_{-i}^j)$ is a nonincreasing step function with steps $\{(p_{(m)}^j, (\zeta_{m-1}^j - \zeta_m^j))\}_{m=1}^{M^j}$ as above. Let $a_i^j(s) \in [0, Q^j]$ be buyer i 's awarded allocation from seller j . There exists a unique index $n(j) \in \{1, \dots, M^j\}$ with

$$a_i^j(s) \in (\zeta_{n(j)}^j, \zeta_{n(j)-1}^j].$$

By direct evaluation of the integral of a piecewise-constant function,

$$\int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz = \sum_{m=n(j)+1}^{M^j} p_{(m)}^j (\zeta_{m-1}^j - \zeta_m^j) + p_{(n(j))}^j (a_i^j(s) - \zeta_{n(j)}^j). \quad (8.17)$$

The summation term represents the cumulative cost of all bids at strictly higher prices that receive their full available quantity, whereas the final term $p_{(n(j))}^j (a_i^j(s) - \zeta_{n(j)}^j)$ corresponds to the partial allocation at the threshold price level, i.e. a_i^j .

The local payment $\int_0^{a_i^j} P^j(\cdot) dz$ equals the sum of the areas of the topmost local blocks truncated at total width $a_i^j(s)$. Summing (8.17) over $j \in \mathcal{L}_i$ yields the total area of the multiset of all such blocks, truncated seller-wise at widths $\{a_i^j(s)\}_j$:

$$\sum_{j \in \mathcal{L}_i} \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz = \sum_{j \in \mathcal{L}_i} \sum_{m=1}^{M^j} p_{(m)}^j \Delta z_m^j,$$

where for each pair (j, m) the nonnegative width Δz_m^j is either $\zeta_{m-1}^j - \zeta_m^j$ (if the entire tier is used) or the appropriate truncated width at the threshold tier (if only part of that tier is used). Thus the total payment is the finite sum of prices $p_{(m)}^j$ by allocations Δz_m^j .

Intuitively, the local inverse curves P^j are descending staircases in allocation-price coordinates. Their union, when resorted by price and concatenated, produces the aggregate staircase $P_i(\cdot, s_{-i})$ by taking the same multiset of rectangles $\{(p_{(m)}^j, \Delta z_m^j)\}_{j,m}$

and placing them consecutively along the z -axis in nonincreasing order of price. By construction, $P_i(z, s_{-i})$ is again a nonincreasing step function and its integral up to the total allocation

$$\sum_{j \in \mathcal{L}_i} a_i^j(s) = \sum_{j \in \mathcal{L}_i} \sum_{m=1}^{M^j} \Delta z_m^j$$

is equivalent to integrating the single aggregate staircase up to the global allocation:

$$\int_0^{\sum_{j \in \mathcal{L}_i} a_i^j(s)} P_i(z, s_{-i}) dz = \sum_{j \in \mathcal{L}_i} \sum_{m=1}^{M^j} p_{(m)}^j \Delta z_m^j.$$

This is a direct instance of a measure-preserving reordering, proving (8.15). Finally, (8.16) follows immediately from the definitions

$$c_i^j(s) = \int_0^{a_i^j(s)} P^j(z, s_{-i}^j) dz, \quad c_i(s) = \int_0^{\sum_{j \in \mathcal{L}_i} a_i^j(s)} P_i(z, s_{-i}) dz.$$

□

8.3.3 Valuation and Utility

A rational buyer sets a uniform bid price $p_i(t)$ across all active sellers $j \in \mathcal{L}_i(t)$. The central object for each buyer i is a single buyer-specific valuation curve θ_i , evaluated on the buyer's total awarded quantity across all active auctions. The total possible value is

$$q_i(q_i \circ a) = \theta_i \left(\sum_{j \in \mathcal{L}_i(t)} q_i^j(a) \right) = \int_0^{\sum_{j \in \mathcal{L}_i(t)} q_i^j(a)} \theta_i'(z) dz \quad (8.18)$$

where $\theta_i : [0, Q_i] \rightarrow [0, \infty)$ is the buyer's elastic valuation function with strictly decreasing derivative θ_i' . A buyer's marginal value depends only on the aggregate quantity it expects to receive across auctions, not on the identity of the seller.

Given a strategy profile s , the utility of buyer i for potential allocation a is dependent on the cost, $c_i(s)$, where the cost to buyer i as a function of the entire strategy profile s . In a multi-auction setting this profile evolves with iteration t , where $c_i(s)$ may represent total participation costs, including membership fees, per-round overhead, and per-auction message costs.

We note that buyers that engage with more sellers may face higher aggregate participation cost. We therefore consider the possibility of a time-dependent participation cost for buyer i , $\epsilon_i(t) = |\mathcal{L}_i(t)|\epsilon$, where $|\mathcal{L}_i(t)|$ is the number of active auctions and represents the incremental participation overhead incurred by buyer i when updating all bids simultaneously. The proportional form extends the original single-auction bound by scaling the threshold with the number of active markets. We acknowledge that this is a reasonable pricing structure. However, in the context of this paper, we make the simple assumption that membership fees are a-priori, and that ϵ is a market-wide, fixed fee, and leave a more complex pricing structure as future work.

Given strategy s , the utility of buyer i is $u_i(s) = q_i - c_i(s)$. Now, in terms of the opposing bid vector s_{-i} , a buyer's realized utility at time t depends on the current state of play. Information propagation across the market affects how the vector of opposing bids s_{-i} is formed, and thus how externalities are computed. Given $\epsilon > 0$, a state s is an ϵ -Nash equilibrium if and only if

$$u_i(s'_i; s_{-i}) - u_i(s_i; s_{-i}) \leq \epsilon, \quad \forall i, \forall s'_i \in S_i.$$

Each buyer's utility u_i is computed from its valuation $\theta_i(\cdot)$ and the externality cost derived from (P^j, P_i) . The equilibrium represents the absorbing region of the dynamics—once reached, no buyer can improve its utility by more than the cost of participation.

In the multi-auction setting, buyer i posts a vector of bids that share a common marginal price w_i across all connected sellers. The utility comparison therefore becomes an aggregate test given by Lemma 8.1. For convenience, define $q_i = \sum_{j \in \mathcal{L}_i} q_i^j(a)$, and $q'_i = \sum_{j \in \mathcal{L}_i} q_i'^j(a)$. We have

$$\Delta u_i = \theta_i(q'_i) - \theta_i(q_i) - \int_{q_i}^{q'_i} P_i(z, s_{-i}) dz. \quad (8.19)$$

8.4 Buyer Best Response

We adapt the ϵ -best reply of Lazar and Semret [45] to a market with multiple simultaneous auctions selling a quantity of an identical resource, creating a market of

perfect substitutes. Unlike the serial network setting considered in [45], we assume all sellers offer perfect substitutes, so buyer i is characterized by a single strictly concave valuation function θ_i defined over the total quantity obtained across all active auctions.

Given opponents' bids s_{-i} , define the total quantity of resource available to buyer i is given in (8.11) as $Q_i(y, s_{-i})$. This leads to a natural generalization of the single-auction framework as follows. Namely, we replace the amount available in one auction at the price y by the total resource available in all auctions. Specifically, let the corresponding feasible set be

$$G_i(s_{-i}) = \left\{ z \in [0, \bar{q}_i] : z \leq Q_i(\theta'_i(z), s_{-i}) \right\}, \quad (8.20)$$

where \bar{q}_i is an upper bound on buyer i 's total allocation. The set $G_i(s_{-i})$ collects all total quantities that buyer i can obtain truthfully at marginal value $\theta'_i(z)$ given the total remaining quantity across all auctions given all opposing bids s_{-i} .

Following the method of ϵ -best reply for a single auction let $\epsilon > 0$ and define the target bid

$$w_i = \theta'_i(v_i), \quad \text{where} \quad v_i = \sup G_i(s_{-i}) - \frac{\epsilon}{\theta'_i(0)}.$$

By construction $0 \leq v_i \leq \bar{q}_i$ and $v_i \leq Q_i(p_i, s_{-i})$. Thus (v_i, w_i) represents an ϵ -relaxed truthful bid with a w_i marginal price valuation on v_i amount of resource. The difficulty remains to proportion how much of the total resource v_i to request from each of the auctions $j \in \mathcal{L}_i$ in which buyer i is allowed to bid. In particular, we look for a multi-auction bid of the form

$$s_i = \{(i, j, v_i^j, w_i) : j \in \mathcal{L}_i\} \quad \text{where} \quad \sum_{j \in \mathcal{L}_i} v_i^j = v_i \quad (8.21)$$

and the quantities v_i^j further satisfy the resource availability constraints in each auction

$$v_i^j \leq Q_i^j(w_i, s_{-i}^j) \quad \text{for all} \quad j \in \mathcal{L}_i. \quad (8.22)$$

Among all such decompositions, buyer i chooses s_i to minimize the total exclusion-compensation payment. In fact, we wish to choose the quantities v_i^j in a way that

satisfies (8.22) and at the same time minimizes the exclusion-compensation payment in each of the respective auctions. Because exclusion-compensation is additive across ordered bids, (Lemma 8.1), this minimization reduces to displacing the lowest-priced units first.

Note that $\sup G_i(s_{-i})$ represents the maximum total quantity buyer i can obtain truthfully from all auctions in \mathcal{L}_i given s_{-i} . Define

$$q_i = \sup G_i(s_{-i}), \quad v_i = q_i - \frac{\epsilon}{\theta'_i(0)}, \quad w_i = \theta'_i(v_i).$$

Thus (v_i, w_i) is an ϵ -relaxed truthful target bid, with $v_i \leq q_i$ and marginal valuation w_i .

The remaining task is to decompose v_i across auctions. We seek a bid

$$s_i = \{(i, j, v_i^j, w_i) : j \in \mathcal{L}_i\}, \quad \text{with} \quad \sum_{j \in \mathcal{L}_i} v_i^j = v_i,$$

subject to the feasibility constraints

$$v_i^j \leq Q_i^j(w_i, s_{-i}^j) \quad \text{for all} \quad j \in \mathcal{L}_i.$$

Among all feasible decompositions, buyer i chooses $(v_i^j)_{j \in \mathcal{L}_i}$ to minimize total exclusion-compensation payment.

To evaluate costs, consider the set of all bids competing with buyer i across the auctions in which i participates:

$$s_{-i} = \{(k, j, q_k^j, p_k^j) : k \neq i, j \in \mathcal{L}_i\} \in S_{-i}.$$

We first impose a global price ordering on these bids; we then allocate the resource among opposing buyers as if buyer i were not present in order to compute the second-prices needed for calculating the utility u_i of buyer i .

Using Lemma 8.1, these bids are ordered globally by price, producing a single aggregate staircase $P_i(\cdot, s_{-i})$. The sorting implies

$$s_{-i} = \{(k(n), j(n), q_{k(n)}^{j(n)}, p_{k(n)}^{j(n)}) : n = 1, \dots, N\},$$

where $k(n) \neq i$ and $j(n) \in \mathcal{L}_i$, and the ordering is decreasing such that

$$p_{k(n)}^{j(n)} \geq p_{k(n+1)}^{j(n+1)}$$

with ties broken deterministically. The bids are ordered such that

$$\begin{cases} p_i^j(n) < p_i^j(n+1) \text{ or} \\ p_i^j(n) = p_i^j(n+1) \text{ and } \gamma_{i(n+1)} < \gamma_{i(n)} \end{cases} \quad (8.23)$$

where γ is a key used to impose a deterministic sort. To simplify the notation write

$$q(n) = q_{k(n)}^{j(n)} \quad \text{and} \quad p(n) = p_{k(n)}^{j(n)} \quad \text{for all } n = 1, \dots, N.$$

For each seller j let $Q^j(n)$ be the remaining resource available in auction j after the bids m with $m < n$ have been allocated $a(m)$. Thus, $Q^j(1) = Q^j$ and for $n = 1, \dots, N$ we obtain

$$Q^\ell(n+1) = \begin{cases} Q^{j(n)}(n) - a(n) & \text{for } \ell = j(n) \\ Q^\ell(n) & \text{otherwise,} \end{cases}$$

where

$$a(n) = \min(q_{k(n)}^{j(n)}, Q^{j(n)}(n)) \quad (8.24)$$

is the allocation awarded to buyer $k(n)$ at seller $j(n)$.

This ordering is across auctions and reflects allocation received for each bid based on the marginal values of the opposing buyers in the absence of buyer i .

For each index n , define

$$H_n = (\theta'_i)^{-1}(p(n)), \quad (8.25)$$

the total quantity at which buyer i 's marginal valuation equals the n th ordered opposing price. Recall $a(n)$ denotes the allocation associated with the n th ordered opposing bid. Buyer i selects the smallest index n such that

$$\sum_{m=n}^N a(m) \leq H_n. \quad (8.26)$$

Because the externality of each displaced unit equals its price, this rule minimizes the total exclusion-compensation cost. Here n indexes the valuation cutoff determined by H_n , while n' indexes the truncation required to meet the ϵ -relaxed quantity v_i . The index n is determined solely by the valuation curve, independent of ϵ ; the index n' is determined by the ϵ -relaxed quantity constraint.

Since θ_i represents strictly decreasing returns then θ'_i is a strictly decreasing function, and so θ'_i is invertible. For convenience we extend the inverse function as

$$(\theta'_i)^{-1}(p) = \begin{cases} q & \text{for } p < \theta'_i(0), \quad \text{where } \theta'_i(q) = p, \\ 0 & \text{for } p \geq \theta'_i(0). \end{cases}$$

Lemma 8.1 allows costs incurred across distinct sellers to be evaluated using a single inverse demand curve. In fact, without the aggregate exclusion-compensation identity, a single global price ordering across auctions would not define a coherent best response.

Now we consider the opposing bids from lowest to highest and whether the second price set by that bid is less than our own marginal valuation for obtaining the corresponding allocation. Thus H_n represents the total quantity on our valuation curve that we would want if we outbid all the bids from n to N .

We then consider the ordered opposing bids from index $n = N$ to $n = 1$, i.e., from the lowest to the highest price, and determine whether displacing each successive allocation is profitable relative to buyer i 's valuation. Because exclusion-compensation is linear in displaced units, any deviation from this ordering strictly increases total cost for a fixed total quantity. Each allocation that we obtain contributes to our cost through the second-price mechanism as the externality from the opposing bid. Since we obtain allocations in reverse order this minimizes the second-cost from as many auctions we can bid in.

There are two cases, either

$$n = 1 \quad \text{or} \quad \sum_{m=n}^N a(m) \leq H_n, \quad \text{but} \quad \sum_{m=n-1}^N a(m) > H_{n-1}.$$

In the first case we are willing to outbid all the bids to get all the allocations.

Thus, the resources from all available markets is less than the amount our valuation says we would want by just outbidding the highest bid, and so we will get our requested quantity. There is no additional cost in bidding more because we do not incur any additional externality, we have already outbid all opposing bids and

$$\sum_{j \in \mathcal{L}_i} Q^j = \sum_{m=1}^N a(m) \leq H_1. \quad (8.27)$$

Taking $q_i^j = Q^j$ it follows that $\sum_{j \in \mathcal{L}_i} q_i^j = \sup G_i(s_{-i})$ where G_i is given by (8.20).

What's left is to set the quantities $v_i^j \leq q_i^j$ in a way that minimizes the cost.

Thus, we want

$$v_i = \sum_{j \in \mathcal{L}_i} v_i^j = \sum_{j \in \mathcal{L}_i} Q^j - \frac{\epsilon}{\theta_i'(0)} \quad (8.28)$$

Since the second cost increases with the bid prices of the opposing bids, we first choose n' such that

$$\sum_{m=n'}^N a(m) \leq v_i < \sum_{m=n'-1}^N a(m). \quad (8.29)$$

Take $\alpha \geq 0$ to be defined as $\alpha = v_i - \sum_{m=n'}^N a(m) < a(n' - 1)$ and define

$$v_i^\ell = \begin{cases} \sum_{m=n'}^N a(m), & \text{for } \ell \neq j(n' - 1) \\ \sum_{m=n'}^{j(m)=\ell} a(m) + \alpha & \text{for } \ell = j(n' - 1). \end{cases} \quad (8.30)$$

Now, $\sum_{\ell \in \mathcal{L}_i} v_i^\ell$ satisfies (8.28). Next we consider the case where $\sum_{m=n-1}^N a(m) > H_{n-1}$, and the case where $\sum_{m=n}^N a(m) \leq H_n$. Recall that $H_n = (\theta_i')^{-1}(p(n))$ represents the total quantity on our valuation curve that we would want if we outbid all the bids from n to N . We consider two subcases:

$$\sum_{m=n}^N a(m) \leq H_{n-1} \quad \text{or} \quad \sum_{m=n}^N a(m) > H_{n-1},$$

where the second inequality represents the case where outbidding the next highest bid would result in a lower desired allocation at a higher price, and the mechanism

discards the bid. In the first case, $\sum_{m=n}^N a(m) \leq H_{n-1}$, we have an increase in demand and $\sum_{m=n-1}^N a(m) \leq H_{n-1}$. Therefore

$$\sum_{m=n}^N a(m) \leq H_{n-1} < \sum_{m=n-1}^N a(m) \implies 0 \leq H_{n-1} - \sum_{m=n}^N a(m) < a(n' - 1).$$

Let $\beta = H_{n-1} - \sum_{m=n}^N a(m)$. Now,

$$q_i^\ell = \begin{cases} \sum_{\substack{m=n \\ j(m)=\ell}}^N a(m), & \text{for } \ell \neq j(n-1) \\ \sum_{\substack{m=n \\ j(m)=\ell}}^N a(m) + \beta & \text{for } \ell = j(n-1), \end{cases} \quad (8.31)$$

and $\sum_{\ell \in \mathcal{L}_i} q_i^\ell = H_{n-1}$, remarking that

$$\sum_{\ell \in \mathcal{L}_i} \sum_{\substack{m=n \\ j(m)=\ell}}^N a(m) = \sum_{m=n}^N a(m).$$

We reduce the quantity q_i^ℓ by $\epsilon/\theta'_i(0)$:

$$\sum_{\ell \in \mathcal{L}_i} v_i^\ell = H_{n-1} - \frac{\epsilon}{\theta'_i(0)} \quad (8.32)$$

so that

$$\sum_{m=n'}^N a(m) \leq H_{n-1} - \frac{\epsilon}{\theta'_i(0)} < \sum_{m=n'-1}^N a(m) \leq H_{n-1},$$

where (8.27) implies that $n' > 1$. Now, $\alpha = H_{n-1} - \epsilon/\theta'_i(0) - \sum_{m=n'}^N a(m) < a(n' - 1)$, and (8.30) satisfies (8.32). Note that the partial allocation at index $n' - 1$ is awarded to buyer $k(n' - 1)$ at seller $j(n' - 1)$ by construction of the ordering.

The preceding construction determines how a buyer should decompose a target total quantity across multiple auctions so as to minimize exclusion-compensation cost. What remains is to show that this construction indeed corresponds to a best response in the induced multi-auction game.

To this end, we separate the argument into two steps. First, we characterize the exact best response for buyer i when the buyer is allowed to request the full utility-maximizing quantity dictated by its valuation curve. This step is independent of

the ϵ -relaxation and establishes the fundamental structure of optimal demand across auctions. Second, we show how this exact best response can be truncated to yield an ϵ -best response that respects the relaxed feasibility constraint while preserving near-optimal utility.

Fix opponents' bids s_{-i} and let the aggregate price function for buyer i , $P_i(\cdot, s_{-i})$ be defined as in Lemma 8.1, and consider the ordered opponent bids, where $p(n)$ is the bid price and $a(n)$ is the quantity allocated according to the n -th bid by (8.24) at that price, and where again the bids are ordered according to (8.23).

Define the cumulative quantity obtained by outbidding all bids between $n + 1$ and N as

$$A_n = \sum_{m=n+1}^N a(m). \quad (8.33)$$

Lemma 8.2. *Suppose $A_n < H_n$ and $a(n) > 0$. Then outbidding the next bid increases total utility.*

Proof. Consider just outbidding bid n . Since bid n has price $p(n)$ our bid price is $p_i = p(n)^+$. By outbidding bid n we may obtain any amount up to $a(n)$. Define

$$\delta_n = \min(H_n - A_n, a(n))$$

which is the truthful amount we want subject to that availability. Therefore, the associated cost and value increments are

$$c_i(n-1) - c_i(n) = \delta_n p_i \quad \text{and} \quad v_i(n-1) - v_i(n) = \theta_i(A_n + \delta_n) - \theta_i(A_n).$$

The corresponding utility increment is

$$u_i(n-1) - u_i(n) = \theta_i(A_n + \delta_n) - \theta_i(A_n) - \delta_n p_i.$$

By the Mean Value Theorem, there exists $\eta \in (0, \delta_n)$ such that

$$\theta_i(A_n + \delta_n) - \theta_i(A_n) = \theta'_i(A_n + \eta) \delta_n.$$

Hence,

$$u_i(n-1) - u_i(n) = \delta_n (\theta'_i(A_n + \eta) - p_i).$$

Since θ'_i is decreasing, then $H_n \geq A_n + \delta_n > A_n + \eta$ implies

$$p_i = \theta'_i(H_n) \leq \theta'_i(A_n + \delta_n) < \theta'_i(A_n + \eta)$$

and consequently $u_i(n-1) > u_i(n)$. \square

Note, if $a(n) = 0$ then outbidding the next bid does not change the externality by taking away any allocation from another buyer and consequently $u_i(n-1) = u_i(n)$ which leaves the utility unchanged.

Our first result is stated as a proposition, establishing an optimality property on a set of ordered bids. The second result then follows as a theorem, quantifying the welfare loss incurred by the ϵ -relaxation.

Proposition 8.3 (Buyer best response). *Let t_i denote the truthful bid from Lemma 8.2 and choose n as in (8.26). Then, t_i satisfies*

$$u_i(t_i, s_{-i}) = \sup_{s_i} u_i(s_i, s_{-i}),$$

consuming the full allocation from each bid $m > n$, as well as the partial amount $\delta_n = H_n - A_n$ from bid n , and zero from all bids $m < n$.

Proof. Provided the allocation a_i^j corresponding to a bid (i, j, q_i^j, p_i^j) satisfies $a_i^j = q_i^j$, then any bids of the form (i, j, q_i^j, p) with $p \geq p_i^j$ receives the same allocation and results in the same cost, value and utility. Consider the multi-auction bid of buyer i given by

$$s_i = \{(i, j, q_i^j, p_i^j) : j \in \mathcal{L}_i\} \tag{8.34}$$

It follows that the bid

$$\{(i, j, a_i^j, p_i) : j \in \mathcal{L}_i\} \quad \text{where} \quad p_i = \max\{p_i^j : j \in \mathcal{L}_i\}$$

receives exactly the same allocation as the other bid and generates the same utility. We claim the truthful bid t_i from Lemma 8.2 satisfies

$$u_i(t_i, s_{-i}) = \sup_{s_i} u_i(s_i, s_{-i}).$$

Following the remark above, we may assume that the supremum is obtained for a bid s_i of the form given in (8.34) where $a_i^j = q_i^j$ for all $j \in \mathcal{L}_i$ and $p_i^j = p_i$ does not depend on j . For each n , define

$$u_i(n) = \theta_i(A_n) - c_i(n),$$

where $c_i(n)$ is the exclusion-compensation cost of the first $n - 1$ bids.

Thus utility increases strictly up to the unique point H_n , defined in (8.25) and decreases thereafter. Consuming bids sequentially from lowest price to highest (equivalently, in decreasing index order) and stopping exactly at this point therefore maximizes u_i . Since the resulting bid t_i is of the same form as s_i then they both provide the same utility. \square

We observe that, at the stopping point of the construction, buyer i 's marginal valuation equals the price of the displaced bid. That is, if n is the index at which the allocation transitions from fully displaced to not displaced, then

$$\theta'_i(A_n + \delta_n) = p(n),$$

with A_n and δ_n defined as above. Consequently, displacing any bid with index $m < n$ would require paying a price strictly exceeding buyer i 's marginal valuation, and therefore yields no positive utility gain. Hence bids with index $m < n$ generate zero externality in the best-response reply.

Theorem 8.4 (ϵ -best response). *Let t_i be the exact best response characterized in Proposition 8.3, and define*

$$q_i = \sup G_i(s_{-i}), \quad v_i = q_i - \frac{\epsilon}{\theta'_i(0)}, \quad w_i = \theta'_i(v_i).$$

Construct a feasible multi-auction bid

$$t_i^{(\epsilon)} = \{(i, j, v_i^j, w_i) : j \in \mathcal{L}_i\} \quad \text{with} \quad \sum_{j \in \mathcal{L}_i} v_i^j = v_i,$$

satisfying the per-auction constraints (8.22), with the quantities v_i^j chosen according to the ordered-bid construction (8.30). Then

$$u_i(t_i^{(\epsilon)}, s_{-i}) \geq u_i(t_i, s_{-i}) - \epsilon.$$

where $u_i(t_i, s_{-i}) = \sup_{s_i} u_i(s_i, s_{-i})$.

Proof. By Proposition 8.3,

$$\sup_{s_i} u_i(s_i, s_{-i}) = u_i(t_i, s_{-i}) = \theta_i(H_n) - c_i(H_n),$$

where $c_i(\cdot)$ is the exclusion-compensation cost induced by the ordered bids.

The bid $t_i^{(\epsilon)}$ requests total quantity $v_i = q_i - \Delta$ and is chosen to minimize exclusion-compensation across the set of possible bids, so its payment is no larger than the payment incurred by requesting $q_i - \Delta$ under any other bid. Let n be the valuation cutoff index defined by (8.26), independent of ϵ . By Proposition 8.3, the exact best response requests a total quantity equal to H_n , and therefore $q_i = \sup G_i(s_{-i}) = H_n$.

(i) If $\epsilon = 0$, then $v_i = q_i$ and any bid of the form

$$t_i^{(0)} = \{(i, j, v_i^j, w_i) : j \in \mathcal{L}_i\}, \quad \sum_{j \in \mathcal{L}_i} v_i^j = v_i,$$

achieves the exact optimum by Proposition 8.3.

(ii) Assume $\epsilon > 0$. Define the truncation amount

$$\Delta = \frac{\epsilon}{\theta_i'(0)}$$

so that $v_i = q_i - \Delta$.

We compare t_i at total quantity $q_i = H_n$ and $t_i^{(\epsilon)}$ at total quantity $q_i - \Delta$. Since $t_i^{(\epsilon)}$ requests a total quantity reduced by Δ relative to t_i , we may write

$$u_i(t_i, s_{-i}) - u_i(t_i^{(\epsilon)}, s_{-i}) = \theta_i(q_i) - \theta_i(q_i - \Delta) - (c_i(q_i) - c_i(q_i - \Delta)).$$

Since exclusion-compensation cost is nondecreasing in the total displaced quantity, we have $c_i(q_i) - c_i(q_i - \Delta) \geq 0$, and therefore

$$0 \leq u_i(t_i, s_{-i}) - u_i(t_i^{(\epsilon)}, s_{-i}) \leq \theta_i(q_i) - \theta_i(q_i - \Delta).$$

Since θ_i' is decreasing and $\theta_i'(H_n) = p(n)$,

$$\theta_i(q_i) - \theta_i(q_i - \Delta) = \int_{q_i - \Delta}^{q_i} \theta_i'(z) dz.$$

Now, since $\theta'_i(z) \leq \theta'_i(0)$ for all z ,

$$0 \leq u_i(t_i, s_{-i}) - u_i(t_i^{(\epsilon)}, s_{-i}) \leq \theta'_i(0) \Delta \leq \epsilon.$$

□

Taken together, Proposition 8.3 and Theorem 8.4 provide a complete characterization of buyer best responses in the multi-auction setting. The proposition shows that, when feasibility constraints are ignored, a buyer's optimal strategy is to displace opposing bids in increasing price order until its marginal valuation equals the prevailing externality price. The theorem then demonstrates that truncating this strategy by at most $\epsilon/\theta'_i(0)$ units yields an ϵ -best response.

Crucially, Lemma 8.1 induces a global ordering of opposing bids across auctions, which allows exclusion-compensation costs to depend only on the total displaced quantity. Without such an ordering, a well-defined best response would not exist.

This structure suggests that repeated best-response updates may exhibit stable behavior. Once bids are arranged so that no buyer can profitably adjust its total requested quantity or its allocation across auctions by more than an ϵ margin, subsequent best-response calculations will reproduce the same bids. This observation motivates the following conjecture:

Conjecture 8.5 (Absorbing ϵ -NE region.). *Under the proposed ϵ -best-response dynamics, the multi-auction market converges to a state in which no buyer admits a unilateral deviation yielding a utility gain exceeding ϵ . Once such a state is reached, all subsequent best-response updates are rejected, and the dynamics remain within an absorbing ϵ -Nash equilibrium region.*

In practical terms, this conjecture asserts that the market stabilizes under repeated ϵ -best-response updates. After a finite number of iterations, all buyers' bids lie within an ϵ -neighborhood of mutual optimality, and no further updates can produce a material improvement in utility. The results of this section establish the structural properties of individual best responses that make such absorbing behavior plausible; a formal analysis of convergence is left to future work.

8.5 Simulations

In this section we study the effect of buyer connectivity on equilibrium outcomes in a market consisting of multiple simultaneous auctions selling a perfectly substitutable resource. The goal of these simulations is to illustrate how the structure of the buyer–seller incidence pattern influences prices and price variability at an ϵ –Nash equilibrium, complementing the best–response analysis developed in the preceding sections.

Following Semret and Lazar [45], each buyer i is endowed with a strictly concave valuation over the total quantity of resource obtained across all auctions. Valuations are taken to be quadratic of the form

$$\theta_i(z) = \kappa_i(\bar{q}_i - z/2)z, \quad z \in [0, \bar{q}_i],$$

where \bar{q}_i represents the maximum desired quantity and $\kappa_i = \bar{p}_i/\bar{q}_i$ scales marginal value so that \bar{p}_i is the maximum marginal valuation at zero consumption. This specification ensures decreasing marginal utility while remaining simple enough to allow large-scale simulation.

Furthermore, at each equilibrium or stopping point, the following quantities are collected,

$$E_j = \frac{1}{A_j} \sum_{i \in \mathcal{B}^j} a_i^j p_i \quad \text{and} \quad V_j = \frac{1}{A_j} \sum_{i \in \mathcal{B}^j} a_i^j (p_i - E_j)^2 \quad \text{where} \quad A_j = \sum_{i \in \mathcal{B}^j} a_i^j$$

across realizations. In addition, buyer classification (winners, zero allocations) and network statistics (fraction of shared buyers, seeds) are collected.

Figure 8.1 shows the average marginal price E_j for the corresponding ϵ –Nash equilibrium in auction j as a function of the number of buyers allowed to bid in both auctions at a corresponding ϵ –Nash equilibrium. Here the equilibrium is obtained by using the ϵ –best reply algorithm with $\epsilon = 5.0$, see Algorithm 1 in [45]. Each buyer is allowed to bid in turn round robin. The shadow beneath each curve represents the region $E_j \pm \sqrt{V_j}$ where V_j is the variance in marginal price. Note that as connectivity

increases, both sellers exhibit convergence in average marginal price and clearing price in their respective auctions.

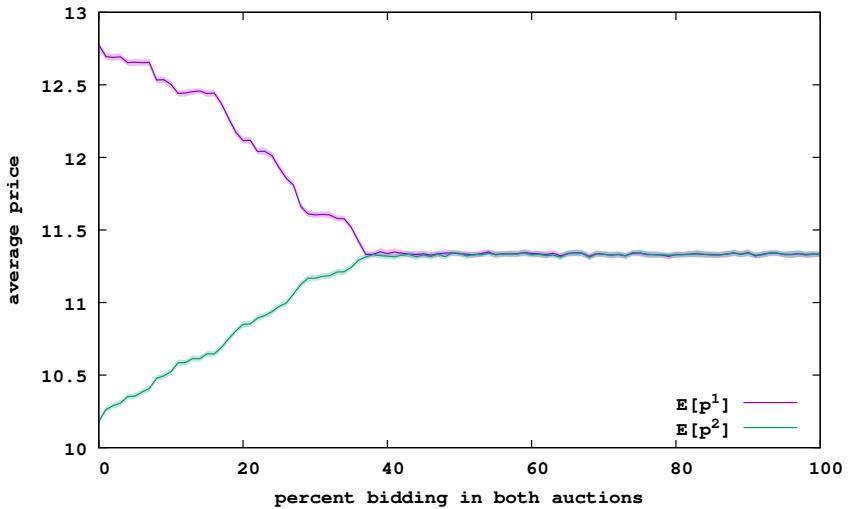


Figure 8.1: Average bid price versus the percentage of buyers participating in multiple auctions. The resource in each auction is given by $Q^1 = 1000$ and $Q^2 = 2000$. The reserve price is $P^j = 6$ in both actions.

Qualitatively, the Y-shape of Figure 8.1 is expected. When all buyers participate in only one auction, the auctions are independent and the average bid price is higher in the auction with less available resource. Seller-seller influence through buyers bidding in both auctions cause the market to function like one auction when a sufficient number of buyers are able to buy from both sellers. Quantitatively, when about 35 percent of the buyers or more are able to bid in both auctions, the average bid price in each auction is about the same. Repeated trials with different buyers sampled according to the same distribution yield similar Y-shaped graphs. We remark, however, that the point at which the bid prices in both auctions align varied to some degree.

Next, we consider a market consisting of four auctions arranged in a linear topology, as illustrated in Figure 8.2. Buyers are initially associated with a primary auction, but may also be allowed to bid in neighboring auctions. The parameter p controls buyer connectivity: it represents the percentage of buyers permitted to participate in adjacent auctions.

When p is small, buyers are largely confined to a single auction and competition

is local. As p increases, participation sets overlap more substantially, and buyers increasingly compete across multiple auctions. In the limit, when p is large, nearly all buyers are able to bid in all auctions and the market approaches a fully connected setting.

Figure 8.2 illustrates the buyer–seller graph used in the simulations for a representative value of p .

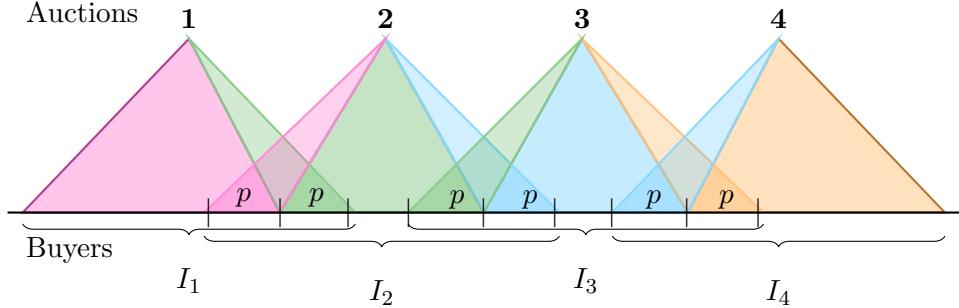


Figure 8.2: A buyer–seller graph for a market with four auctions, 400 buyers and a 6-edge distance between the first and last auctions. Here p is the percentage of buyers allowed to bid in the neighboring auctions. Shown is $p = 20$ where $|\mathcal{I}_1| = 120$, $|\mathcal{I}_2| = 140$, $|\mathcal{I}_3| = 140$ and $|\mathcal{I}_4| = 120$. Note that when $p = 300$ all bidders bid in all auctions.

For each value of the connectivity parameter p , we simulate asynchronous ϵ –best–response dynamics starting from an initial bid profile. Buyers update their bids sequentially, following the best–response construction described earlier, until an ϵ –Nash equilibrium is reached, meaning that no buyer can improve its utility by more than ϵ through a unilateral deviation.

Figure 8.3 reports the resulting average bid prices as a function of buyer connectivity for a market with heterogeneous auction capacities.

To account for heterogeneity, each experiment is repeated over multiple random realizations of buyer valuations. At equilibrium, we record the resulting bid prices in each auction and compute averages and standard deviations across runs.

Figures 8.3 and 8.4 show the average equilibrium bid price in each auction as a function of buyer connectivity p . The shaded regions represent one standard deviation

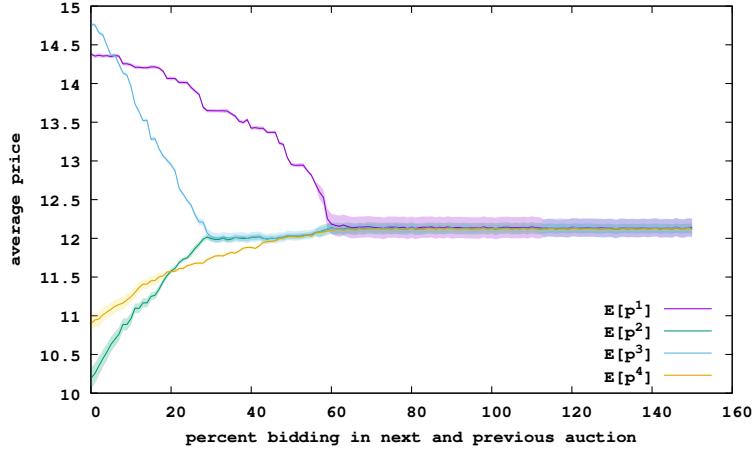


Figure 8.3: Average bid price at an ϵ -Nash equilibrium for a market consisting of four auctions versus the percentage of buyers who can bid in the neighboring auction—see Figure 8.2. The shadow illustrates the standard deviation. The resource in each auction is given by $Q^1 = 500$, $Q^2 = 2000$, $Q^3 = 500$ and $Q^4 = 2000$. The reserve price is $P^j = 10$ in all actions.

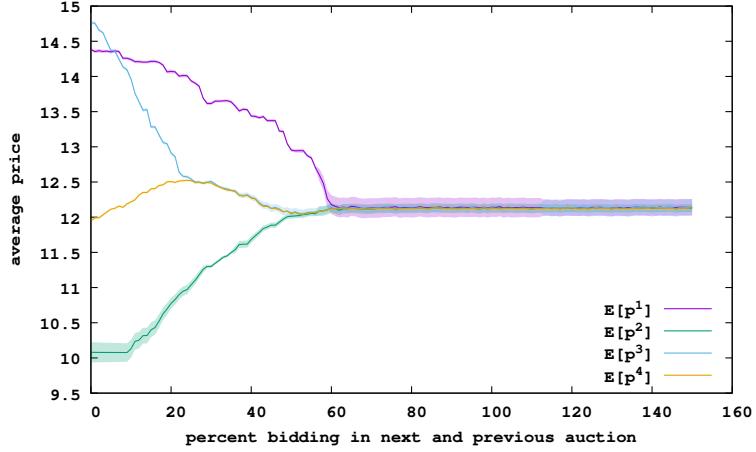


Figure 8.4: The same buyers as in Figure 8.3 except with $Q^1 = 500$, $Q^2 = 2500$, $Q^3 = 500$ and $Q^4 = 1500$. The reserve price is $P^j = 10$ in all actions.

across simulation runs. In both cases, the reserve price is fixed at $P^j = 10$ for all auctions, while auction capacities differ across the two scenarios.

When connectivity is low, auctions behave almost independently and equilibrium prices reflect local supply and demand conditions. As connectivity increases, buyers

are able to shift demand across neighboring auctions, leading to stronger coupling between prices. Price dispersion decreases, and equilibrium prices become increasingly aligned across auctions, despite differences in available capacity.

These observations are consistent with the interval-based best-response characterization established earlier. Higher connectivity enlarges the ordered set of competing bids relevant to each buyer's decision problem, thereby expanding the range of bids that can influence equilibrium outcomes. As a result, local changes in supply or demand propagate more broadly through the market, reducing price differentials while increasing sensitivity to aggregate conditions.

Overall, the simulations highlight how buyer connectivity governs the transition from localized competition to a more integrated market, and illustrate the role of network structure in shaping equilibrium behavior in multi-auction PSP markets.

Statement on Supplementary Material The code for the experiments presented in this paper can be found at:

- <https://github.com/jkblazek/dissertation>

8.6 Conclusion and Future Work

The analysis above shows that buyer best responses in multi-auction PSP markets admit a precise interval structure determined by the interaction between reserve prices and marginal valuations. In particular, for fixed opponents' bids, the globally ordered set of competing bids induces a partition of the feasible quantity space into intervals on which the buyer's utility is monotone. Utility increases as quantity is accumulated through lower-priced bids and decreases once higher-priced bids are reached.

As a result, each buyer's optimal response depends only on a finite collection of price intervals associated with this ordered set of bids. Bids priced outside these intervals cannot affect the buyer's utility or allocation decision, while bids within the intervals are ordered and locally comparable. This induces a strong locality property

in the best-response calculation: only those competitors whose bids appear in the relevant portion of the global ordering can influence the buyer's choice.

Taken together, these results show that the multi-auction best-response problem can be reduced to a one-dimensional comparison over a finite ordered set of bids, despite the presence of multiple sellers and heterogeneous competition. The ordered structure of opponent bids fully determines both the direction and magnitude of profitable deviations, and thereby identifies the subset of bids that are relevant for each buyer's decision problem.

Future Work: An absorbing ϵ -Nash equilibrium region A formal proof of Conjecture 8.5, which asserts the existence of an absorbing ϵ -Nash equilibrium region under the proposed best-response dynamics, is beyond the scope of the present work. Nevertheless, the results above provide strong structural support for this claim. In particular, the characterization the ϵ -best response the autonomy of the buyer's process, and as such, we expect that once bids lie within an ϵ -neighborhood of mutual optimality, no deviation can yield a utility gain exceeding ϵ .

Future Work: Dynamic participation costs Buyer i maintains an active seller set

$$\mathcal{L}_i(t) = j \in \mathcal{L} : a_i^j(s, t) > 0.$$

Suppose, instead of a static ϵ , the buyer faces a participation cost $\epsilon_i(t) = |\mathcal{L}_i(t)|, \epsilon$ that scales with the size of its active set.

Under this rule, a buyer enters a new auction only if the predicted utility gain exceeds the higher threshold associated with the enlarged active set, and exits an auction whenever the utility loss is smaller than the threshold for the current set. This asymmetry creates a band of width ϵ between entry and exit conditions, introducing a simple form of hysteresis.

Entry (add a seller $j \notin \mathcal{L}_i(t)$): add j and compute the predicted improvement

$\Delta u_i(\text{add } j \leftarrow \text{status quo})$. Accept the entry if

$$\Delta u_i(\text{add } j \leftarrow \text{status quo}) > (|\mathcal{L}_i(t)| + 1)\epsilon.$$

Exit (drop a seller $j \in \mathcal{L}_i(t)$): remove j and compute the predicted loss Δu_i ($\text{status quo} \leftarrow \text{drop } j$) (this quantity is nonnegative when dropping j hurts). Accept the exit if

$$\Delta u_i(\text{status quo} \leftarrow \text{drop } j) < |\mathcal{L}_i(t)|\epsilon.$$

Thus, entry is tested against the threshold evaluated with the *larger* active set, while exit is tested against the threshold for the *current* set. The gap between $(|\mathcal{L}_i(t)| + 1)\epsilon$ and $|\mathcal{L}_i(t)|\epsilon$ creates a strict band of width ϵ .

Future Work: Expanding on the bipartite structure Let $A \in \{0, 1\}^{|\mathcal{I}| \times |\mathcal{L}|}$ denote the buyer–seller biadjacency matrix, where $A_{ij} = 1$ if buyer i participates in auction j . For each buyer i and seller j with $A_{ij} = 1$, consider the deviation of buyer i ’s bid price from the seller’s reserve price,

$$z_i^j(t) = |p_i^j(t) - p^j(t)|.$$

Collecting these coordinates yields a buyer state vector $z_i(t)$ defined on the support of A .

Influence between buyers arises only through shared sellers. Accordingly, a natural metric on buyer states is the L_∞ norm restricted to common neighbors,

$$d_\infty(i, k) = \max_{j \in \mathcal{L}} A_{ij} A_{kj} |z_i^j(t) - z_k^j(t)|.$$

Under this metric, influence neighborhoods are axis–aligned regions determined by the largest marginal deviation at any shared seller, reflecting the fact that PSP pricing and displacement are driven by the single closest competitor at the margin.

The reserve–price margin interval

$$\bar{p}^j(t) < p^j(t) < \underline{p}^j(t),$$

naturally defines the boundary of influence at seller j : buyers with bids outside this interval cannot affect marginal allocations or second-cost terms, while buyers within the interval determine displacement at the margin. From this perspective, the choice of $p^j(t)$ selects a boundary within an admissible influence region, and admissible update rules correspond to monotone movements of this boundary.

Future work will investigate coordinated reserve–price selection rules, using feedback from marginal buyers, that remain within the margin interval and preserve ϵ –best–response dynamics. Such rules can be interpreted as SIG boundary updates and may substantially reduce the computational burden of second–cost calculations by restricting attention to buyers within the corresponding influence sets.

Appendix: Proving Price and Cost Equality

Proof. Following [45], we note that “it is readily apparent that”

$$a_i(s) = q_i \wedge \underline{Q}_i(p_i; s_{-i}),$$

where

$$c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})] \quad \text{and} \quad c_i(s) = \int_0^{a_i(s)} P_i(z, s_{-i}), dz. \quad (8.35)$$

Dropping s_{-i} for simplicity, we wish to show that

$$z \leq Q_i(y) \Rightarrow y \geq P_i(z) \quad \text{and} \quad y > P_i(z) \Rightarrow z \leq Q_i(y).$$

We have

$$P_i(z) = \inf y \geq 0 : Q_i(y) \geq z, \quad Q_i(y) = [Q - \sum_{p_k > y, k \neq i} q_k]^+.$$

Assume no ties and $q_{i_n} > 0$. If there are ties, they can be consolidated into a single opposing bid $(q_i + q_k, p)$. In the case of opposing bids with ties, then we consolidate to one opposing bid $(q_i + q_k, p)$.

For example, let

$$p_1 = (1, 3) > p_2 = (2, 2) > p_3 = (1, 1), \quad Q = 5.$$

Then

$$Q_i([3, \infty)) = 5, \quad Q_i([2, 3)) = 4, \quad Q_i([1, 2)) = 2, \quad Q_i([0, 1)) = 1,$$

and

$$P_i((4, 5]) = 3, \quad P_i((2, 4]) = 2, \quad P_i((1, 2]) = 1, \quad P_i([0, 1]) = 0.$$

In general, assuming surplus,

$$\begin{aligned} Q_i([p_{i_1}, \infty)) &= Q, \\ Q_i([p_{i_2}, p_{i_1})) &= Q - q_{i_1}, \\ &\vdots \\ Q_i([p_{i_{n+1}}, p_{i_n})) &= Q - \sum_{m=1}^n q_{i_m}, \\ &\vdots \\ Q_i([0, p_{i_N})) &= Q - \sum_{m=1}^N q_{i_m}, \\ P_i((Q - q_i, Q]) &= p_i, \quad P_i((Q - \sum_{m=1}^2 q_{i_m}, Q - q_{i_1}]) = p_{i_2}, \\ &\vdots \\ P_i((Q - \sum_{m=1}^n q_{i_m}, Q - \sum_{m=1}^{n-1} q_{i_m})) &= p_{i_n}, \end{aligned}$$

runs:

$$P_i([0, Q - \sum_{m=1}^N q_{i_m}]) = 0.$$

Let

$$p_{i_1} = Q - q_i, \quad p_0 = Q, \quad p_{i_n} = Q - \sum_{m=1}^n q_{i_m}, \quad p_{i_{n-1}} = Q - \sum_{m=1}^{n-1} q_{i_m}, \quad p_{i_{N+1}} = 0.$$

Define cumulative remaining capacities

$$\xi_0 := Q, \quad \xi_n := \left[Q - \sum_{m=1}^n q_{i_m} \right]^+, \quad n \geq 1.$$

(If there is no surplus, then for some $M \leq N$, $\xi_M = 0$ and $\xi_n = 0$ for all $n \geq M$.)

On this grid,

$$Q_i(y) = \begin{cases} \xi_0, & y \in [p_{i_1}, \infty), \\ \xi_1, & y \in [p_{i_2}, p_{i_1}), \\ \vdots \xi_n, & y \in [p_{i_{n+1}}, p_{i_n}), \\ \vdots \end{cases} \quad P_i(z) = \begin{cases} p_{i_n}, & z \in (\xi_n, \xi_{n-1}], \quad 1 \leq n \leq M, \\ 0, & z \in [0, \xi_M]. \end{cases}$$

This is the 90° rotation correspondence between Q_i and P_i described in [45]. Fix $a = a_i(s) \in (0, Q]$ and let n satisfy $a \in (\xi_n, \xi_{n-1}]$. Since P_i is constant on each $(\xi_m, \xi_{m-1}]$ with value p_{i_m} ,

$$\begin{aligned} \text{runs} \int_0^a P(z), dz &= \left(\int_{\xi_M}^{\xi_{M-1}} + \cdots + \int_{\xi_{n+1}}^{\xi_n} + \int_{\xi_n}^a \right) P(z), dz \\ &= p_{i_M}(\xi_{M-1} - \xi_M) + p_{i_{M-1}}(\xi_{M-2} - \xi_{M-1}) + \cdots + \\ &\quad p_{i_{n+1}}(\xi_n - \xi_{n+1}) + p_{i_n}(a - \xi_n) \\ &= p_{i_n}(a - \xi_n) + \sum_{m=n+1}^M p_{i_m}(\xi_{m-1} - \xi_m). \end{aligned}$$

We claim that

$$\text{runsc}_i(s) = \sum_{m=1}^M p_{i_m} [a_j(0; s_{-i}) - a_j(s_i; s_{-i})] = \int_0^{a_i(s)} P(z), dz,$$

where

$$Q_i(p_{i_{n+1}}) = \xi_n = \left[Q - \sum_{m=1}^n q_{i_m} \right]^+ = \left[Q - \sum_{p_k > y, k \neq i} q_k \right]^+.$$

By the PSP rule, when i is absent ($s_i = 0$), opponents fill capacity from the lowest price upward; with i present, the first $a = a_i(s)$ units that i acquires displace lower opponents. For each opponent $j = i_m$,

$$a_j(0; s_{-i}) - a_j(s_i; s_{-i}) = \begin{cases} \xi_{m-1} - \xi_m, & m > n \quad (\text{fully displaced block}), \\ a - \xi_n, & m = n \quad (\text{partially displaced block}), \\ 0, & m < n. \end{cases}$$

Multiplying by p_{i_m} and summing $m = 1, \dots, M$ gives

$$c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})],$$

which equals the integral

$$c_i(s) = \int_0^{a_i(s)} P_i(z; s_{-i}), dz.$$

The sum indexes opponents j , while the integral indexes marginal quantity z ; both compute the same total displacement cost. \square

Chapter 9

Conclusions and Future Work

9.1 Conclusions

The research presented here forms a collaborative exchange of ideas from auction theory, network science, and game theory. Our understanding of decentralized design, and through an immersive study of the mathematics supporting the Progressive Second Price Auction and its extensions gives form to a model view where each participant—buyer or seller—as a local deterministic factor who in the iterative update process is able to compute a dynamic solution that we would have never been able to come by otherwise. Through these repeated interactions, we see global collaboration, emergent despite any central authority. Our approach shows that despite solutions that seem to evolve or disappear, the rigor of the mechanism’s foundational work and robustness to perturbed starting conditions calls for a further understanding of this system of dynamic inquiries and responses. Our intent was to demonstrate an understanding of decentralized systems where individualistic control may achieve convergence, efficiency, and incentive compatibility even when information is systemically withheld, incomplete, or distributed unevenly.

The first chapters establish the theoretical and historical foundations, connecting our work with the Vickrey-Clarke-Groves lineage and to distributed versions proposed by Lazar, Semret, and others. Later chapters develop the mathematical formalism of the mechanism, define influence and opt-out behavior, and link the dynamics of bidding to graph structures that represent real networks. We begin our study with

an overview given in Chapter 3.

We remark that the contributions detailed in this work are extensions of the PSP mechanism beyond the work of the authors who originated and validated its design. Through our consideration, verification and simulation of these previous works, mentioned already throughout this dissertation, we mean to describe how agents affect one another both directly and indirectly, and demonstrate how local dependencies create global outcomes.

In Chapters 4 and 5, we describe how the PSP mechanism is expressed mathematically. It bridges market behavior and network representation by formalizing buyer–seller dynamics as systems of equations and potential games. The analysis highlights how second-price and VCG principles can be extended to networked environments, turning theoretical PSP rules into tractable, model-based representations. It also links user behavior and algorithmic response, serving as a transitional point from mechanism design to network analysis.

Chapter 6 extends the analysis to incorporate both graph-based and game-theoretic methods. It presents bipartite and influence-propagation models for buyer–seller networks and defines saturation as a limit of influence diffusion. The model of multi-auction markets, where buyers and sellers are connected through overlapping influence sets, hosts the projection-based influence framework, defining primary and secondary influence sets and demonstrating how local interactions produce market-wide coordination and saturation effects. This work connects the PSP framework to quantized and asynchronous convergence results from Qu, Jia, and Caines [69, 70] and establishes formal stability criteria through graph connectivity and a partial-ordering of relations built on these connections via the monotonicity of PSP price updates. The treatment of monotonicity and influence shells provides a rigorous framework for studying equilibrium formation, as we move from representation to theoretical proof.

A modification to the mechanism that enforces collaborative behavior in Chapter 7 demonstrates the existence of a unique equilibrium state where the auction produces zero revenue. The “zero cost” solution is a natural consequence of the itera-

tive bidding process and the second-price rule, where the price paid is the externality imposed on others. Chapter 7 additionally incorporates a study on latency, using a Weibull distribution to represent network delay, addressing convergence under realistic network conditions, and proposing solutions to encourage convergence speed in an applied analysis.

Each chapter builds on the last to expand and refine our knowledge of this intricate system of optimization designed for the allocation of consumable, networked resources. The final chapter challenges the theoretical framework in a multi-auction setting. We intend to validate the mechanism for resource allocation in a generalized, dynamic framework, positioning PSP as both a theoretical and computational model for studying decentralized coordination, and provide a coherent theory of decentralized market dynamics that connects equilibrium reasoning, network structure, and temporal adaptation.

Our contributions demonstrate how decentralized systems can be described by rules that are both simple and powerful. By connecting the local logic of individual agents with the global behavior of the network, the PSP framework bridges economics, computation, and ecology. This work shows that order can emerge without central control and that systems built on local feedback can adapt, learn, and reach equilibrium through iteration and interaction. We close with the conviction that the principles outlined here—adaptability, decentralization, and resilience—will remain essential in understanding and designing the markets and networks of the future.

9.2 Future Work

There were many things that we were unable to complete, and that have been left for future work. Future research will address theoretical completeness, robustness under uncertainty, and broader real-world implementation of PSP mechanisms.

Theoretical Proofs

Future work will continue to challenge the theory behind the model, extending and validating its proofs through technical analysis and simulation. We will explore the constraints that drive the mechanism's dynamics, and in so doing further explore the solution space. We will elaborate and expand on the formal foundation of the model.

Several elements of Chapter 8 remain incomplete, particularly the proof of the existence of an ϵ -best reply for agents participating in multiple concurrent auctions. Establishing this result is essential to confirming that PSP remains stable under overlapping participation. Additional work should generalize equilibrium convergence proofs for distributed markets with shared influence sets and bounded participation domains. Completing these proofs will strengthen the theoretical foundation of the dynamic multi-auction PSP model.

Noise and Stochastic Dynamics

Future research will continue to explore imperfect models, stochastic environments, and noise-driven adaptation in decentralized markets. Inspired by the work of Nakagaki et al. [61], which showed that the slime mold *Physarum polycephalum* finds efficient paths through noisy environments, PSP could incorporate controlled randomness to avoid suboptimal equilibria. Stochastic modeling traditions established by Uhlenbeck and Ornstein [82] and expanded by Qu, Jia, and Caines [70] provide a foundation for introducing mean-reverting and diffusion processes that link continuous-time uncertainty to discrete iterative dynamics.

Digital Data and Network Reliability

Reliability remains critical for decentralized systems. Insights from Bracha [18] on asynchronous consensus and Shapiro et al. [76] on eventual consistency should inform PSP designs that tolerate missing or delayed bid updates. Xu and Hajek's Supermarket Game [89] offers a valuable analogy for understanding queueing and congestion, mapping communication latency to auction dynamics.

Network Optimization under Imperfection

Building on Barrett et al. [8] and Michael and Newman [56], future research should examine how graph metrics—such as resistance distance and dynamic proximity—characterize evolving connectivity in imperfect markets. This can help model how influence sets deform under link failures or network perturbations, revealing how topology governs equilibrium robustness.

Applications and Computational Extensions

The PSP framework has broad applicability across emerging decentralized technologies. Buterin’s Ethereum [20] explores potential implementations of PSP as automated smart contracts for resource allocation. Brandt et al. [19] point to computational limits in multi-agent equilibrium models; future work should therefore emphasize scalable, approximate computation for PSP networks. Incorporating evolutionary game dynamics from Cressman [25], Chastain [21], and Wang [85] can also yield adaptive, learning-based bidding mechanisms for real-time decentralized markets.

In summary, future work should strengthen theoretical completeness, improve robustness under stochastic conditions, and expand PSP applications to digital, networked, and permissionless systems. Together, these directions will extend the mechanism from a theoretical framework to a practical foundation for understanding and designing decentralized coordination in complex adaptive systems.

References

- [1] Omar Abdul-Ghafoor, Mahamod Ismail, Rosdiadee Nordin, and Musbah Shaat. Resource allocation in multiuser multi-carrier cognitive radio network via game and supermarket game theory: survey, tutorial, and open research directions. *KSII Transactions on Internet and Information Systems*, 8:3674–3710, November 2014. DOI: 10.3837/tiis.2014.11.003.
- [2] P.A. Absil and Jerome Malick. Projection-like Retractions on Matrix Manifolds. *SIAM Journal on Optimization*, 22:135–158, 2012.
- [3] Marcos K. Aguilera and Sam Toueg. The correctness proof of ben-or’s randomized consensus algorithm. *Distributed Computing*, 25:371–381, 2012. DOI: 10.1007/s00446-012-0162-z.
- [4] Osvaldo Anacleto and Catriona Queen. Dynamic Chain Graph Models for Time Series Network Data. *Bayesian Analysis*, 12(2):491 –509, 2017. DOI: 10.1214/16-BA1010. URL: <https://doi.org/10.1214/16-BA1010>.
- [5] Muhammad Asad Arfeen, K. Pawlikowski, D. McNickle, and A. Willig. The role of the weibull distribution in modeling internet access and backbone core networks. *Journal of Network and Computer Applications*, 141:1–22, 2019.
- [6] Laleh Arshadi and Amir H. Jahangir. An empirical study on tcp flow interarrival time distribution for normal and anomalous traffic. *International Journal of Communication Systems*, 30:1–19, 2017.
- [7] AT&T. AT&T mobile share flex plans, \NoCaseChange{<https://www.att.com/plans/wireless/mobile-share-flex.html>} (last accessed Oct 14, 2019), 2019. URL: <https://www.att.com/plans/wireless/mobile-share-flex.html>.
- [8] Wayne Barrett, Emily J. Evans, and Amanda E. Francis. Resistance distance in straight linear 2-trees, 2017. DOI: 10.48550/ARXIV.1712.05883. URL: <https://arxiv.org/abs/1712.05883>.
- [9] Melanie Baur. Combinatorial concepts and algorithms for drawing planar graphs, January 2012.
- [10] Jordana Blazek, Eric Olson, and Frederick C. Harris. Bipartiteness in progressive second-price multi-auction networks with perfect substitute. *Preprint*, 2025. URL: <https://arxiv.org/abs/2511.19225>.

- [11] Jordana Blazek, Eric Olson, and Frederick C. Harris. The effects of latency in a progressive second price auction. *Preprint*, 2025. URL: <https://arxiv.org/abs/2511.17424>.
- [12] Jordan Blocher and Frederick C. Harris. An equilibrium analysis of a secondary mobile data-share market. *Information*, 12(11), 2021. ISSN: 2078-2489. DOI: 10.3390/info12110434. URL: <https://www.mdpi.com/2078-2489/12/11/434>.
- [13] Jordan Blocher and Frederick C. Harris. An equilibrium analysis of a secondary mobile data-share market. *Information*, 12(11), 2021. ISSN: 2078-2489. DOI: 10.3390/info12110434. URL: <https://www.mdpi.com/2078-2489/12/11/434>.
- [14] Jordan Blocher and Frederick C. Harris. An optimization algorithm for the sale of overage data in hong kong's mobile data exchange market. In Shahram Latifi, editor, *17th International Conference on Information Technology–New Generations (ITNG 2020)*, pages 553–561, Cham. Springer International Publishing, 2020. ISBN: 978-3-030-43020-7. DOI: 10.1007/978-3-030-43020-7_73. URL: https://link.springer.com/chapter/10.1007/978-3-030-43020-7_73.
- [15] Immanuel M. Bomze. Lotka-volterra equation and replicator dynamics: a two-dimensional classification. *Biological Cybernetics*, 48:201–211, 2004. DOI: 10.1007/BF00318088.
- [16] George Boole. *The mathematical analysis of logic: being an essay towards a calculus of deductive reasoning*. Macmillan, Barclay, & Macmillan, Cambridge, 1847.
- [17] Stephen P. Borgatti and Martin G. Everett. A graph-theoretic perspective on centrality. *Social Networks*, 28(4):466–484, 2006. DOI: 10.1016/j.socnet.2005.11.005.
- [18] Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *J. ACM*, 32(4):824–840, October 1985. ISSN: 0004-5411. DOI: 10.1145/4221.214134. URL: <https://doi.org/10.1145/4221.214134>.
- [19] Felix Brandt and Tuomas Sandholm. On the existence of unconditionally privacy-preserving auction protocols. *ACM Trans. Inf. Syst. Secur.*, 11(2), May 2008. ISSN: 1094-9224. DOI: 10.1145/1330332.1330338. URL: <https://doi.org/10.1145/1330332.1330338>.
- [20] Vitalik Buterin. Ethereum: a next-generation smart contract and decentralized application platform, 2014. Available: <https://ethereum.org>.
- [21] Erick Chastain, Adi Livnat, Christos Papadimitriou, and Umesh Vazirani. Algorithms, games, and evolution. *Proceedings of the National Academy of Sciences*, 111(29):10620–10623, 2014. DOI: 10.1073/pnas.1406556111. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1406556111>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1406556111>.

- [22] Suh-Wen Chiou. Optimization of robust area traffic control with equilibrium flow under demand uncertainty. *Computers & Operations Research*, 41:399–411, 2014. ISSN: 0305-0548. DOI: <https://doi.org/10.1016/j.cor.2013.06.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0305054813001652>.
- [23] Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 1997. ISBN: 9780821803158.
- [24] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971. DOI: 10.1007/BF01726210.
- [25] Ross Cressman and Yi Tao. The replicator equation and other game dynamics. *Proceedings of the National Academy of Sciences*, 111(supplement_3):10810–10817, 2014. DOI: 10.1073/pnas.1400823111. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1400823111>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1400823111>.
- [26] A. Delenda, Patrick Maillé, and Bruno Tuffin. Reserve price in progressive second price auctions. In 755–759 Vol.2, January 2004. ISBN: 0-7803-8623-X. DOI: 10.1109/ISCC.2004.1358631.
- [27] Matt DeVos, Bojan Mohar, and Robert Šámal. Unexpected behaviour of crossing sequences. *Journal of Combinatorial Theory, Series B*, 101(6):448–463, 2011. ISSN: 0095-8956. DOI: <https://doi.org/10.1016/j.jctb.2010.12.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0095895611000025>.
- [28] Enrico Diecidue and Peter P. Wakker. On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–298, November 2001. ISSN: 1573-0476. DOI: 10.1023/A:1011877808366. URL: <https://doi.org/10.1023/A:1011877808366>.
- [29] Boualem Djehiche, Alain Tcheukam, and Hamidou Tembine. Mean-field-type games in engineering. *AIMS Electronics and Electrical Engineering*, 1(1):18–73, 2017. ISSN: 2578-1588. DOI: 10.3934/ElectrEng.2017.1.18. URL: <https://www.aimspress.com/article/doi/10.3934/ElectrEng.2017.1.18>.
- [30] Ferenc Domes and Arnold Neumaier. Linear and parabolic relaxations for quadratic constraints. *Journal of Global Optimization*, 65, July 2016. DOI: 10.1007/s10898-015-0381-5.
- [31] Linton C. Freeman. Centrality in social networks conceptual clarification. *Social Networks*, 1(3):215–239, 1978. DOI: 10.1016/0378-8733(78)90021-7.
- [32] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973. DOI: 10.2307/1914085.

- [33] E. R. Hansen. Global optimization using interval analysis: the one-dimensional case. *Journal of Optimization Theory and Applications*, 29(3):331–344, November 1979. ISSN: 1573-2878. DOI: 10.1007/BF00933139. URL: <https://doi.org/10.1007/BF00933139>.
- [34] Dirk Helbing. Boltzmann-like and boltzmann-fokker-planck equations as a foundation of behavioral models. *Physica A: Statistical Mechanics and its Applications*, 196(4):546–573, 1993. ISSN: 0378-4371. DOI: [https://doi.org/10.1016/0378-4371\(93\)90034-2](https://doi.org/10.1016/0378-4371(93)90034-2). URL: <https://www.sciencedirect.com/science/article/pii/0378437193900342>.
- [35] Leon J. Helsloot, Gamze Tillem, and Zekeriya Erkin. Badass: preserving privacy in behavioural advertising with applied secret sharing. In Joonsang Baek, Willy Susilo, and Jongkil Kim, editors, *Provable Security*, pages 397–405, Cham. Springer International Publishing, 2018. ISBN: 978-3-030-01446-9.
- [36] Jean-Manuel Izaret and Just Schürmann. Why progressive pricing is becoming a competitive necessity, 2019. URL: <https://www.bcg.com/publications/2019/why-progressive-pricing-becoming-competitive-necessity>.
- [37] Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996. DOI: 10.1006/jeth.1996.0108.
- [38] Peng Jia and Peter Caines. Analysis of a class of decentralized decision processes: quantized progressive second price auctions. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1–6, January 2008.
- [39] Jon Kleinberg. Cascading behavior in networks: algorithmic and economic issues. *Algorithmic Game Theory*, January 2007. DOI: 10.1017/CBO9780511800481.026.
- [40] Leonard Kleinrock. *Queueing Systems, Volume 1: Theory*. Wiley-Interscience, New York, NY, USA, 1975. ISBN: 978-0471491101.
- [41] Yannis A. Korilis, Aurel A. Lazar, and Ariel Orda. Capacity allocation under noncooperative routing. *IEEE Transactions on Automatic Control*, 42:309–325, 1997.
- [42] Yannis A. Korilis, Aurel A. Lazar, Ariel Orda, and Aurel A. Lazar Ariel Orda. The designer’s perspective to noncooperative networks. In *Infocom*, pages 562–570, 1995.
- [43] Ba Lam To, Thi Mai Trang Nguyen, and Guy Pujolle. Flow-based power control in cognitive radio ad-hoc network. In *2012 8th International Wireless Communications and Mobile Computing Conference (IWCMC)*, pages 130–135, 2012. DOI: 10.1109/IWCMC.2012.6314190.
- [44] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978. ISSN: 0001-0782. DOI: 10.1145/359545.359563. URL: <https://doi.org/10.1145/359545.359563>.

- [45] Aurel Lazar and Nemo Semret. Design, Analysis and Simulation of the Progressive Second Price Auction for Network Bandwidth Sharing. Game Theory and Information, University Library of Munich, Germany, 1998. URL: <https://EconPapers.repec.org/RePEc:wpa:wupgpa:9809001>.
- [46] Yongfu Li, Xiao Jiang, Hao Zhu, Xiaozheng He, Srinivas Peeta, Taixiong Zheng, and Yinguo Li. Multiple measures-based chaotic time series for traffic flow prediction based on bayesian theory. *Nonlinear Dynamics*, 85(1):179–194, July 2016. ISSN: 1573-269X. DOI: 10.1007/s11071-016-2677-5. URL: <https://doi.org/10.1007/s11071-016-2677-5>.
- [47] Zonghua Liu. Chaotic time series analysis. *Mathematical Problems in Engineering*, 2010:720190, April 2010. ISSN: 1024-123X. DOI: 10.1155/2010/720190. URL: <https://doi.org/10.1155/2010/720190>.
- [48] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
- [49] Sabita Maharjan, Yan Zhang, and Stein Gjessing. Economic approaches for cognitive radio networks: a survey. *Wireless Personal Communications*, 57:33–51, 2011. DOI: 10.1007/s11277-010-0005-9.
- [50] P. Maillé and B. Tuffin. Multibid auctions for bandwidth allocation in communication networks. In *IEEE INFOCOM 2004*, doi: 10.1109/INFCOM.2004.1354481, volume 1, page 65, March 2004. DOI: 10.1109/INFCOM.2004.1354481.
- [51] Patrick Maillé. Market clearing price and equilibria of the progressive second price mechanism. en. *RAIRO - Operations Research - Recherche Opérationnelle*, 41(4):465–478, 2007. DOI: 10.1051/ro:2007030. URL: <http://www.numdam.org/articles/10.1051/ro:2007030/>.
- [52] Mohammad Hossein Manshaei, Quanyan Zhu, Tansu Alpcan, Tamer Başar, and Jean-Pierre Hubaux. Game theory meets network security and privacy. *ACM Comput. Surv.*, 45(3), July 2013. ISSN: 0360-0300. DOI: 10.1145/2480741.2480742. URL: <https://doi.org/10.1145/2480741.2480742>.
- [53] Robert M. May and George F. Oster. Bifurcations and dynamic complexity in simple ecological models. *The American Naturalist*, 110(974):573–599, 1976. ISSN: 00030147, 15375323. URL: <http://www.jstor.org/stable/2459579> (visited on 10/22/2022).
- [54] Nakagaki T. Meyer B Ansorge C. The role of noise in self-organized decision making by the true slime mold *physarum polycephalum*. *PLoS One.*, March 2017. DOI: 10.1371/journal.pone.0172933,.
- [55] T.S. Michael and T. Quint. Sphere of influence graphs: edge density and clique size. *Mathematical and Computer Modelling*, 20(7):19–24, 1994. DOI: 10.1016/0895-7177(94)90067-1. URL: [https://doi.org/10.1016/0895-7177\(94\)90067-1](https://doi.org/10.1016/0895-7177(94)90067-1).

- [56] T.S Michael and T Quint. Sphere of influence graphs in general metric spaces. *Math. Comput. Model.*, 29(7):45–53, April 1999. ISSN: 0895-7177.
- [57] Bojan Mohar. The laplacian spectrum of graphs. *Graph Theory, Combinatorics, and Applications*, 2:871–898, 1991.
- [58] Tiago Montanher, Arnold Neumaier, and Ferenc Domes. A computational study of global optimization solvers on two trust region subproblems. *Journal of Global Optimization*, 71(4):915–934, August 2018. ISSN: 1573-2916. DOI: 10.1007/s10898-018-0649-7. URL: <https://doi.org/10.1007/s10898-018-0649-7>.
- [59] Joan Morris. *A simulation-based approach to dynamic pricing*. Master’s thesis, MIT, May 2001.
- [60] Danielle Movsowitz-Davidow, Nir Lavi, and Orna Agmon Ben-Yehuda. Bill estimation in simplified memory progressive second price auctions. In Karim Djemame, Jörn Altmann, José Ángel Bañares, Orna Agmon Ben-Yehuda, and Maurizio Naldi, editors, *Economics of Grids, Clouds, Systems, and Services*, pages 54–62, Cham. Springer International Publishing, 2019. ISBN: 978-3-030-36027-6.
- [61] Toshiyuki Nakagaki, Hiroyasu Yamada, and Ágota Tóth. Intelligence: maze-solving by an amoeboid organism. *Nature*, 407(6803):470, 2000. DOI: 10.1038/35035159.
- [62] John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950. URL: <https://EconPapers.repec.org/RePEc:ecm:emetrp:v:18:y:1950:i:2:p:155-162>.
- [63] Dusit Niyato and Ekram Hossain. Competitive pricing for spectrum sharing in cognitive radio networks: dynamic game, inefficiency of nash equilibrium, and collusion. *IEEE Journal on Selected Areas in Communications*, 26(1):192–202, 2008. DOI: 10.1109/JSAC.2008.080117.
- [64] Marcel Nutz, Jaime San Martin, and Xiaowei Tan. Convergence to the mean field game limit: A case study. *The Annals of Applied Probability*, 30(1):259–286, 2020. DOI: 10.1214/19-AAP1501. URL: <https://doi.org/10.1214/19-AAP1501>.
- [65] Olukayode Oki, Thomas Olwal, Pragassen Mudali, and Matthew Adigun. Biologically-inspired foraging decision making in distributed cognitive radio networks. In pages 28–40, January 2018. ISBN: 978-3-319-68384-3. DOI: 10.1007/978-3-319-68385-0_3.
- [66] Lisa Ordóñez and Lehman Benson. Decisions under time pressure: how time constraint affects risky decision making. *Organizational Behavior and Human Decision Processes*, 71:121–140, February 1997. DOI: 10.1006/obhd.1997.2717.

- [67] Lace M. Padilla, Sarah H. Creem-Regehr, Mary Hegarty, and Jeanine K. Stefanucci. Decision making with visualizations: a cognitive framework across disciplines. *Cognitive Research: Principles and Implications*, 3(1):29, July 2018. ISSN: 2365-7464. DOI: 10.1186/s41235-018-0120-9. URL: <https://doi.org/10.1186/s41235-018-0120-9>.
- [68] Billy Pinheiro, Rafael Chaves, Eduardo Cerqueira, and Antônio Jorge Abelém. *CIM-SDN: A Common Information Model extension for Software-Defined Networking*. December 2013, pages 836–841. ISBN: 978-1-4799-2851-4.
- [69] Zheng Qu, Yue Jia, and Peter E. Caines. Convergence and stability of networked progressive second price auctions. *Automatica*, 45(9):2107–2114, 2009.
- [70] Zheng Qu, Yue Jia, and Peter E. Caines. Quantized progressive second price mechanisms for decentralized resource allocation. In *Proceedings of the 46th IEEE Conference on Decision and Control (CDC)*, pages 1475–1482, 2007.
- [71] Thomas Quint. Necessary and sufficient conditions for balancedness in partitioning games. *Mathematical Social Sciences*, 22(1):87–91, 1991. ISSN: 0165-4896. DOI: [https://doi.org/10.1016/0165-4896\(91\)90068-3](https://doi.org/10.1016/0165-4896(91)90068-3). URL: <https://www.sciencedirect.com/science/article/pii/0165489691900683>.
- [72] Thomas Quint and Martin Shubik. On local and network games, June 2003. URL: <https://ssrn.com/abstract=396081>.
- [73] Nemo Semret. *Market Mechanisms for Network Resource Sharing*. PhD thesis, Columbia University, New York, NY, USA, 1999. ISBN: 0-599-30774-9. AAI9930793.
- [74] Nemo Semret, R. R. F. Liao, A. T. Campbell, and A. A. Lazar. Pricing, provisioning and peering: dynamic markets for differentiated internet services and implications for network interconnections. *IEEE Journal on Selected Areas in Communications*, <http://dx.doi.org/10.1109/49.898733>, 18(12):2499–2513, December 2000. DOI: <http://dx.doi.org/10.1109/49.898733>.
- [75] Devavrat Shah. Message-passing in stochastic processing networks. *Laboratory for Information and Decision Systems, Department of EECS, Massachusetts Institute of Technology*, 2011. URL: http://www.mit.edu/~devavrat/papers/message_passing_survey.pdf.
- [76] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. A comprehensive study of Convergent and Commutative Replicated Data Types. Research Report RR-7506, Inria – Centre Paris-Rocquencourt ; INRIA, January 2011, page 50. URL: <https://hal.inria.fr/inria-00555588>.
- [77] Sergey Shary. Krawczyk operator revised. In *International Conference on Computational Mathematics ICCM-2004*, June 2004.
- [78] Michelle Shen. Modeling population dynamics in changing environments, 2017. URL: <https://math.mit.edu/documents/rsi/2017Shen.pdf>.

- [79] Steve Smale. Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73(6):747–817, 1967. DOI: 10.1090/S0002-9904-1967-11798-1.
- [80] Attila Tasnádi. Price vs. quantity in oligopoly games. *International Journal of Industrial Organization*, 24(3):541–554, 2006. ISSN: 0167-7187. DOI: <https://doi.org/10.1016/j.ijindorg.2005.07.009>. URL: <https://www.sciencedirect.com/science/article/pii/S0167718705001074>.
- [81] Godfried T. Toussaint. The sphere of influence graph: theory and applications. *International Journal of Information Technology & Computer Science*, 14(2), May 2014. ISSN: 2091-1610. URL: <https://studylib.net/doc/18368340/the-sphere-of-influence-graph--theory-and-applications>.
- [82] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Physical Review*, 36(5):823–841, 1930. DOI: 10.1103/PhysRev.36.823.
- [83] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961. DOI: 10.2307/2977633.
- [84] Aida Vosoughi, Joseph R. Cavallaro, and Alan Marshall. Trust-aware consensus-inspired distributed cooperative spectrum sensing for cognitive radio ad hoc networks. *IEEE Transactions on Cognitive Communications and Networking*, 2(1):24–37, 2016. DOI: 10.1109/TCCN.2016.2584080.
- [85] Beibei Wang, Yongle Wu, and K.J. Ray Liu. Game theory for cognitive radio networks: an overview. *Computer Networks*, 54(14):2537–2561, 2010. ISSN: 1389-1286. DOI: <https://doi.org/10.1016/j.comnet.2010.04.004>. URL: <https://www.sciencedirect.com/science/article/pii/S1389128610001064>.
- [86] Jingyi Wang, Xinyue Zhang, Qixun Zhang, Ming Li, Yuanxiong Guo, Zhiyong Feng, and Miao Pan. Data-driven spectrum trading with secondary users' differential privacy preservation. *IEEE Transactions on Dependable and Secure Computing*, 18(1):438–447, 2021. DOI: 10.1109/TDSC.2019.2892447.
- [87] John Glen Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952. DOI: 10.1680/ipeds.1952.11259.
- [88] Marcus Weber. Clustering by using a simplex structure. eng. Technical report ZIB-Report 04-03, ZIB, Takustr. 7, 14195 Berlin, 2003. URL: <https://opus4.kobv.de/opus4-zib/frontdoor/deliver/index/docId/778/file/ZR-04-03.pdf>. urn:nbn:de:0297-zib-7782.
- [89] Jiaming Xu and Bruce Hajek. The Supermarket Game. *Stochastic Systems*, 3(2):405–441, 2013. URL: <https://arxiv.org/abs/1202.2089>.
- [90] Jonathan S. Yedidia. Message-passing algorithms for inference and optimization: “belief propagation” and “divide and concur”. *Unpublished Manuscript, Princeton University*, 2022. Available at Disney Research, Boston.

- [91] Renzhi Yu and Zhongjing Ma. A distributed charging coordination of plug-in electric vehicles based on potential game considering feeder overload constraint. In *2015 18th International Conference on Electrical Machines and Systems (ICEMS)*, pages 1114–1118, 2015. DOI: 10.1109/ICEMS.2015.7385205.
- [92] X. Zhang, L. Guo, M. Li, and Y. Fang. Social-enabled data offloading via mobile participation - a game-theoretical approach. In *2016 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, December 2016. DOI: 10.1109/GLOCOM.2016.7842274.
- [93] Xiaonan Zhang, Linke Guo, Ming Li, and Yuguang Fang. Motivating human-enabled mobile participation for data offloading. *IEEE Transactions on Mobile Computing*, 17(7):1624–1637, July 2018. ISSN: 1558-0660. DOI: 10.1109/TMC.2017.2773087.
- [94] Yang Zhang, Qingyu Yang, Wei Yu, Dou An, Donghe Li, and Wei Hua Zhao. An online continuous progressive second price auction for electric vehicle charging. *IEEE Internet of Things Journal*, 6:2907–2921, 2019.
- [95] Liang Zheng, Carlee Joe-Wong, Chee Wei Tan, Sangtae Ha, and Mung Chiang. Secondary markets for mobile data: feasibility and benefits of traded data plans. In *2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 1580–1588, 2015. DOI: 10.1109/INFOCOM.2015.7218537. URL: <http://dx.doi.org/10.1109/infocom.2015.7218537>.
- [96] Yongquan Zhou, Yuxiang Zhou, Qifang Luo, and Mohamed Abdel-Basset. A simplex method-based social spider optimization algorithm for clustering analysis. *Engineering Applications of Artificial Intelligence*, 64:67–82, 2017. ISSN: 0952-1976. DOI: <https://doi.org/10.1016/j.engappai.2017.06.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0952197617301239>.

Appendix A

List of Publications by Jordana Katherine Blazek

- [1] Jordan Blocher and Frederick C. Harris. An equilibrium analysis of a secondary mobile data-share market. *Information*, 12(11), 2021. ISSN: 2078-2489. DOI: 10.3390/info12110434. URL: <https://www.mdpi.com/2078-2489/12/11/434>.
- [2] Jordan Blocher and Frederick C. Harris. An optimization algorithm for the sale of overage data in hong kong's mobile data exchange market. In Shahram Latifi, editor, *17th International Conference on Information Technology–New Generations (ITNG 2020)*, pages 553–561, Cham. Springer International Publishing, 2020. ISBN: 978-3-030-43020-7. DOI: 10.1007/978-3-030-43020-7_73. URL: https://link.springer.com/chapter/10.1007/978-3-030-43020-7_73.
- [3] Jordan Blocher, Vincent R. Martinez, and Eric Olson. Data assimilation using noisy time-averaged measurements. *Physica D: Nonlinear Phenomena*, 376–377:49–59, 2018. ISSN: 0167-2789. DOI: <https://doi.org/10.1016/j.physd.2017.12.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0167278917305262>. Special Issue: Nonlinear Partial Differential Equations in Mathematical Fluid Dynamics.
- [4] Jordan F. Blocher. *Dynamical Systems: Chaotic Attractors and Synchronization using Time-Averaged Partial Observations of the Phase Space*. Master's thesis, University of Nevada, Department of Mathematics and Statistics, May 2016. URL: <https://scholarwolf.unr.edu/items/e480c738-5a10-4e6a-bc3a-c959f49e3c13>. Advisor: Dr. Eric Olson.
- [5] Jordan F. Blocher, Samantha M. Hampton, and Christopher E. Linden. Extremal functions on cayley digraphs of finite cyclic groups. In *2012 12th International Symposium on Pervasive Systems, Algorithms and Networks*, pages 53–57, 2012. DOI: 10.1109/I-SPAN.2012.14.