Algorithms for the Sharing Economy

An Economic Modelling Perspective on Distributed Exchange

Rixt Hellinga

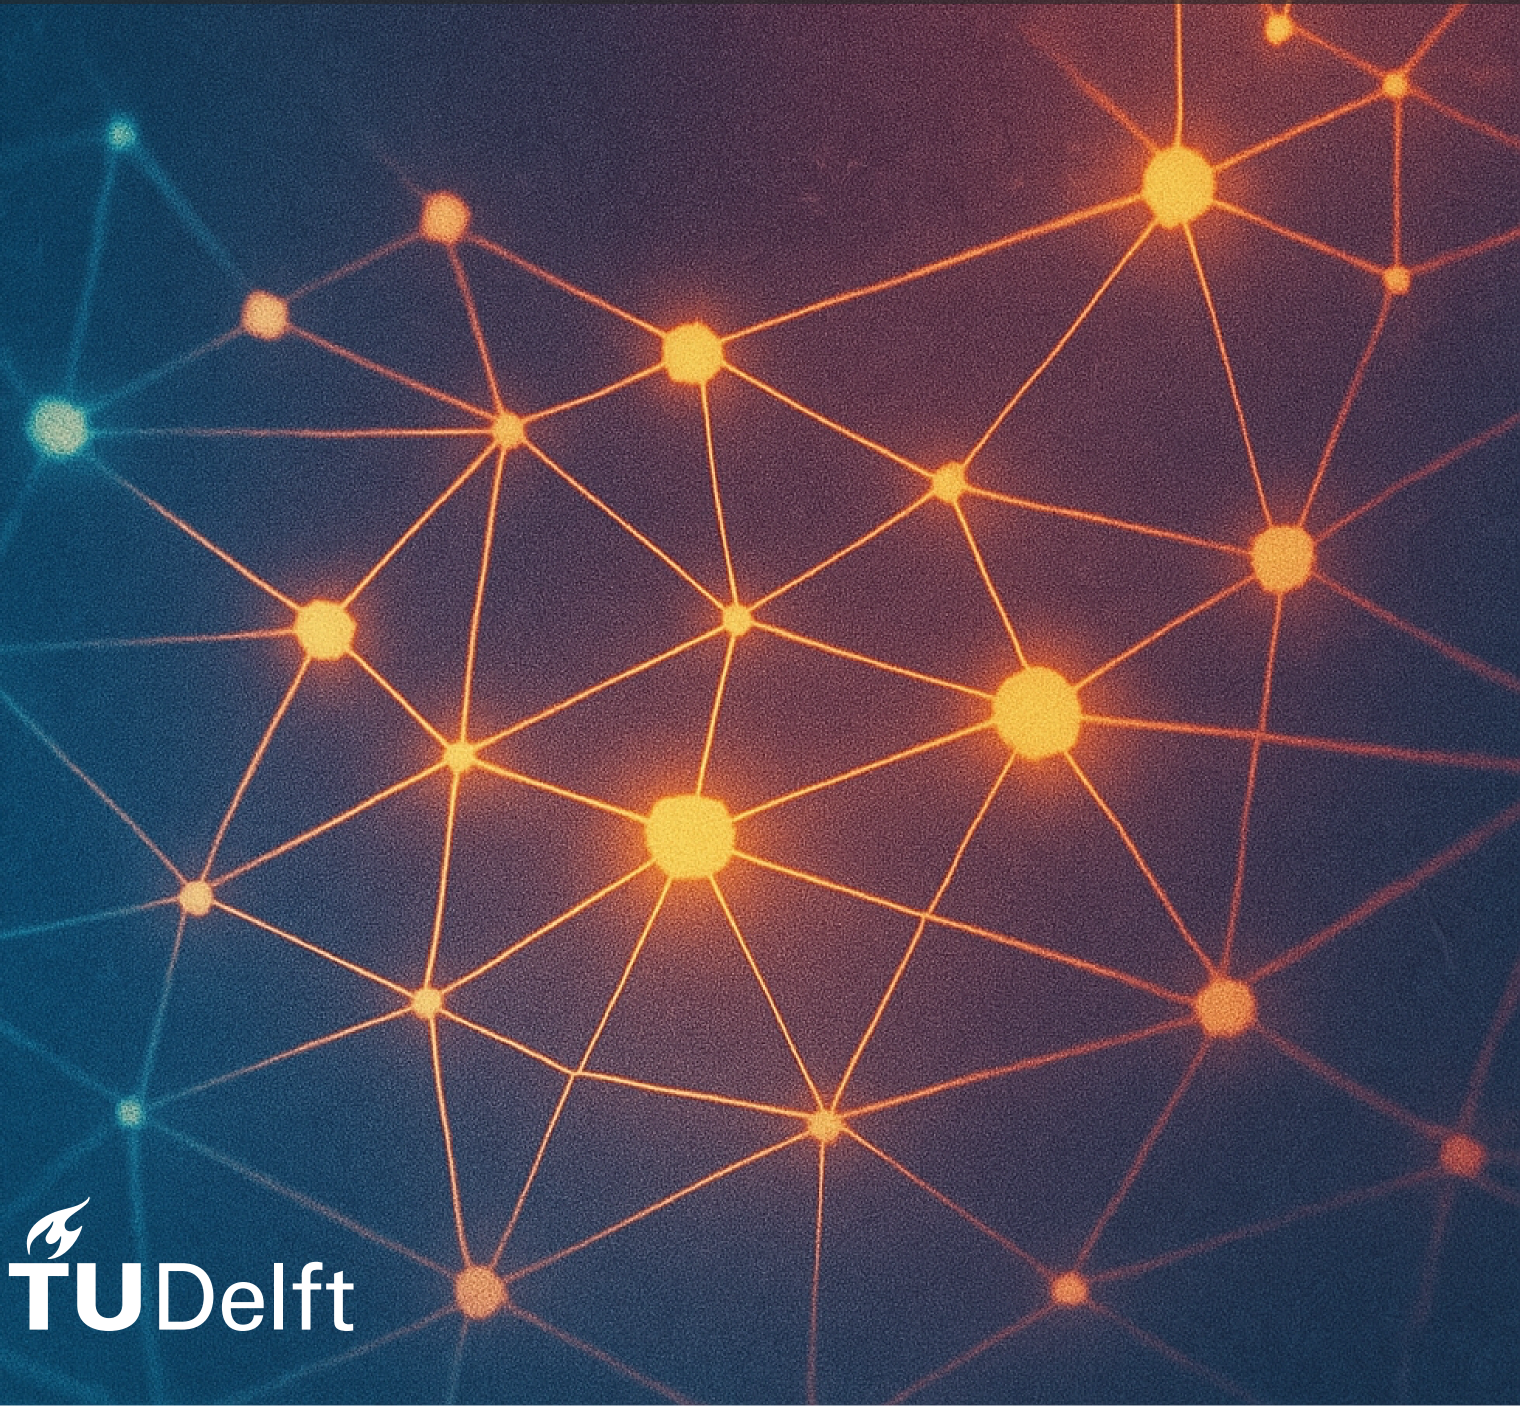

Delft University of Technology

TU Delft

Algorithms for the Sharing Economy

An Economic Modelling Perspective on Distributed Exchange

by

Rixt Hellinga

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday July 17, 2025.

Student number:	4699203	
Project duration:	November 2024 – July, 2025	
Thesis committee:	Dr. N. Yorke-Smith,	TU Delft, supervisor
	Dr. G. Iosifidis,	TU Delft, supervisor
	Dr. ir. E. Aivaloglou,	TU Delft

TU Delft

Acknowledgements

A project like this is never completed entirely alone, and I would like to take this opportunity to thank the people without whom this thesis would not have been possible.

First and foremost, I thank my supervisors, Neil and Georgios, for their guidance and support throughout this process. Your open-mindedness towards what may have been a somewhat unconventional thesis allowed me to approach the work with an explorative mindset. As my final project as a student, I particularly valued the freedom to investigate broadly and creatively.
It was a pleasure to work with supervisors who not only encouraged exploration but also knew when to help me refocus and critically assess what mattered and why. This balance was invaluable — especially given my own tendency to explore perhaps a little too widely.

I would also like to thank my fellow students for the lively and enriching weekly meetings. It was always inspiring to gain insight into the many different topics that everyone was working on.

Finally, a heartfelt thank you to those closest to me. In particular, I want to thank Diederik — for reading, re-reading, and reading my thesis yet again, and for always supporting me along the way. Thank you!

Rixt Hellinga
Delft, July 2025

Abstract

This thesis investigates the gap between the theoretical ideals and practical realities of distributed exchange protocols in peer-to-peer sharing economies. While classical economic models and stylised trading mechanisms have been extensively studied, their assumptions often overlook the complexities and behavioural nuances of real-world agent-based markets. Conversely, more empirically grounded approaches tend to be highly case-specific, limiting their generalisability and disconnecting them from established economic concepts.

The central aim of this work is to develop and analyse an algorithmic economic model that more faithfully captures the dynamics of decentralised resource sharing — balancing fairness, efficiency, and resistance to manipulation, while retaining a useful level of abstraction.

In light of this purpose, the thesis performs a comparative theoretical analysis of centralised and distributed exchange mechanisms, and introduces a novel exchange market model that incorporates heterogeneous strategies, bounded rationality, and asynchronous timing. A series of numerical experiments evaluates the performance and robustness of different protocols under these features of sharing economies.

The results show that distributed, mixed-strategy protocols can achieve stable and desirable outcomes, however their success is sensitive to population diversity, limited information, and strategic behaviour. These findings highlight the importance of integrating behavioural aspects into protocol design, and provide insights towards building more robust models for the sharing economy.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
PRD	Proportional Response Dynamics
AD Market	Arrow-Debreu Market
EG Program	Eisenberg-Gale Program
CE	Competitive Equilibrium
SE	Sharing Equilibrium
KKT Conditions	Karush–Kuhn–Tucker Conditions
IR	Incentive Ratio
P2P	peer-to-peer
SANP	Subgradient Algorithm with Nonlinear Projection
het.	heterogeneous
hom.	homogeneous

Table 1: Abbreviations used throughout this thesis.

Symbols

Symbol	Description	Example / Notes
Sets & Indices		
$\mathcal{N}$	Set of all agents	
n	Number of agents	$n = \lvert\mathcal{N}\rvert$
i, j	Indices for agents	
$\mathcal{E}$	Set of connections between agents	
$\mathcal{G}$	Network structure	$\mathcal{G} = (\mathcal{N}, \mathcal{E})$
$\mathcal{N}_i$	Neighbourhood of agent i	$\{j \mid (i, j) \in \mathcal{E}\}$
Resources & Allocations		
t	Timestep index	$t = 1, \ldots, T$
ε_i	Long-term resource average endowment of agent i	
D_i	Distributable resources by agent i	$\mathbb{E}[D_i] = \varepsilon_i$
x_{ij}	Allocation from agent i to agent j	
$\mathbf{x}_i$	Allocation vector of agent i	$[x_{ij}]_{j \in \mathcal{N}_i}$
$\mathbf{X}$	Allocation matrix	$[x_{ij}]_{i,j \in \mathcal{N}}$

Symbol	Description	Example / Notes
$\mathbf{r}_i$	Resources received by agent i	$[x_{ji}]_{j\in\mathcal{N}_i}$
R_i	Total resources received by agent i	$\sum_{j\in\mathcal{N}_i} x_{ji}$
Utilities		
v_i	Value of agent i's resource	
U_i	Total utility obtained by agent i	e.g., $\sum_{j\in\mathcal{N}_i} u_{ij}$
$\mathbf{U}$	Vector of all received utilities	$[U_i]_{i\in\mathcal{N}}$
$u_i(\cdot)$	Utility function for agent i	$u_i(\mathbf{r}_i) = u_i$
Other Parameters		
ρ_i	Sharing ratio of agent i	
w_i	Strategy weight of agent i	
γ_i	Memory decay of agent i	
ϵ	Convergence error tolerance	e.g., $O(1/\epsilon)$
ζ^M	Incentive Ratio of market M	
Operators		
$\overline{(\cdot)}$	Time-average operator	$\frac{1}{T}\sum_{t=1}^{T}\cdot$
$(\cdot)^*$	Equilibrium operator	e.g., $\mathbf{x}^*$, $\mathbf{u}_i^*$
$\hat{(\cdot)}$	Predicted value of variable	e.g., $\hat{U}_i$
$(\cdot)^f$	Variable under strategy f	e.g., $\mathbf{x}_i^{\pi}$
$(\cdot)(t)$	Variable at time t	e.g., $D_i(t)$
$\log a$	natural logarithm of a	$\log a = \ln a$
$a \Rightarrow b$	Logical implication	
$a \rightarrow b$	Convergence of a toward b	
Algorithms		
π	Greedy resource allocation strategy	
ϕ	Proportional resource allocation strategy	
ψ	Mixed resources allocation strategy	

Table 2: Notation used throughout this thesis.

1

Introduction

1.1. Motivation and Context

The sharing economy has emerged as a rapidly growing and popular alternative to traditional ownership-based economic systems [1]. Driven by technological developments in trading platforms and dissatisfaction with conventional economic outcomes (e.g., recessions, inflation [2]), adoption of online, agent-supported sharing economies (e.g., file-sharing through BitTorrent[1]) has steadily increased [3].

With no universally agreed-upon definition, the term *sharing economy* has been subject to varied interpretations [4, 5]. Generally, it is characterised by consumer-to-consumer interaction, utilisation of underused resources, and temporary access [6, 7, 8]. In this thesis, we adopt a peer-to-peer (P2P) trading perspective — distinct from on-demand models such as Uber or DoorDash, which generate new service jobs to meet immediate demand.

Enabled by the ubiquitous connectivity of modern digital platforms, sharing economies facilitate direct, decentralised connections between users. Each user acts as a *prosumer*, simultaneously consuming and providing resources, and independently deciding how much unused capacity to share with the expectation of future reciprocity. These exchanges aim to promote inclusivity, reduce idle capacity, and strengthen local trade [9, 10, 11]. Several practical examples of such systems include:

- **Resource-sharing via user-provided networks.** Participants share assets such as bandwidth, processing power, or storage. For instance, the Swarm[2] storage network, and Wi-Fi sharing platforms like Guifi.net[3] and FON[4], allow users to form ad hoc mesh networks, dynamically routing traffic across peers [12].
- **Energy sharing in micro-grids.** Decentralised energy systems, such as the Brooklyn Microgrid [13], connect local electricity producers and consumers [14], in a P2P setting [15, 16], trading to improve the network's overall flexibility [17, 18].
- **(AI-driven) agentic markets.** As autonomous software agents become more capable, machine-to-machine exchanges increasingly involve agents acting on behalf of users. An expanding number of online resources are now managed algorithmically [19]. For example, idle computational power can be allocated via AI agents that evaluate supply and demand in real time [20, 21].

[1] https://en.wikipedia.org/wiki/BitTorrent
[2] https://web3edge.io/research/what-is-swarm/
[3] https://guifi.net/
[4] https://fon.com/

While such real-world sharing economies vary in form, their shared features can be represented through an abstract, networked, exchange market, as visualised in Figure 1.1. In such a trading market, users trade (1) in a high-frequency, distributed manner, (2) without using money, (3) using diverse and stochastically available resources, and (4) based on the utility they derive from those resources. Generally, in such a network, some sort of stable and fair state — an equilibrium — is a desired outcome. Such equilibria characterise the steady-state behaviour of markets and allow the assessment of the impact of various economic regulations on market outcomes.

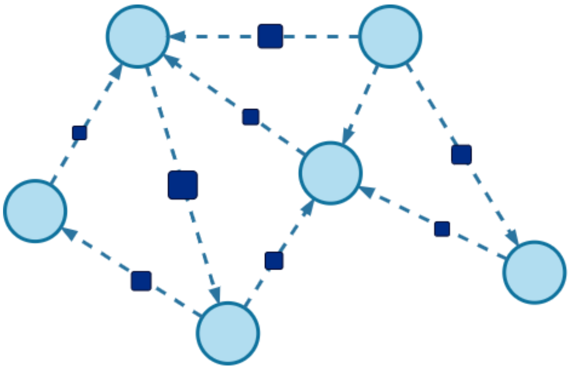

Figure 1.1: Abstract representation of exchange in a sharing economy.

Sharing economies built on agentic P2P interactions often rely on distributed, standardised trading protocols, which introduce several challenges [22]: Issues such as free-riding and fairness require further investigation to enable effective regulation and robust market design [23].
A key issue with such sharing economies is the discrepancy between the theoretical performance of trading protocols and their behaviour in real-world deployments. Many existing studies rely on overly stylised or unrealistic models of exchange [24].
This resulting discrepancy has often been exemplified in real-life trading protocols. For example, the proportional trading protocol *PropShare* [25], was originally considered strategy-proof and fair. Although theoretically capable of achieving equilibrium [26, 27], empirical findings showed that agents often failed to estimate neighbours' contributions accurately due to limited information, thereby reducing market efficiency [27].

One promising direction is to model these trading protocols in more realistic market environments, which allows for the exploration of the protocol's limitations and refine them in accordance with the system's specific priorities and vulnerabilities [28].
Algorithmic approaches offer a powerful means to examine market behaviours [19] and have long been a foundational tool in computer science. Such computational economics enable the simulation of market dynamics and the exploration of interventions, trading structures, and allocation mechanisms, and have traditionally been employed to reason about exchange and pricing systems [29, 30], and to guide the design of economic policies and reforms [31, 32]. More specifically, in the context of sharing economies, algorithmic modelling can inform the design of principled trading protocols aimed at reducing friction in agentic economies [22]. This thesis argues that rather than discarding classical economic theory, we can reinterpret its foundations to expose dynamics that are well-suited to P2P settings.

One might argue that data-driven machine learning approaches can more accurately replicate market behaviour and yield context-specific insights, thereby addressing key research gaps. However, such techniques suffer from limited interpretability, often referred to as the "black box" critique [33]. Consequently, machine learning models are primarily used for predictive tasks on case-specific data, rather than for inferring causal relationships [34].
In contrast, algorithmic economic models (rule-based mechanisms grounded in economic theory) emphasise transparency and interpretability, enabling clearer insights into how specific mar-

ket features affect outcomes. This economic perspective supports the inference of general and causal relationships [35], which is crucial for designing fair and robust market systems. As an increasing number of resources are governed by algorithmic, agentic markets [19, 36], the need for interpretable and transparent allocation mechanisms becomes more pressing [37]. Understanding both what allocations occur, and also why they occur, is essential for anticipating outcomes and fostering trust in sharing economy platforms.

Beyond their interpretability, algorithmic market models gain strength from being grounded in foundational economic principles, such as equilibrium definitions and optimality metrics like Pareto efficiency. These concepts allow us to characterise and evaluate equilibria in stylised market settings. Fairness and stability are core attributes of equilibria and are often central concerns in the design of exchange markets. As such, they provide valuable benchmarks for analysing market performance.
However, the challenge lies in the fact that these foundational economic concepts are not tailored to the kinds of exchange markets considered here. This misalignment complicates efforts to use such concepts not only to evaluate the performance or efficiency of P2P markets, but also to compare them with more conventional markets.

1.2. Research Gap

Despite significant advances, important gaps remain in applying algorithmic economic research to sharing economies. These gaps stem from fundamental differences in how sharing economies operate compared to traditional markets, and from challenges in realistically capturing the decentralised, dynamic behaviour of agents within these systems. Specifically, we focus on two main issues:

- Much of the existing economic literature focuses on classical monetary economies, where equilibrium concepts are well-defined under assumptions such as price-based exchange and fully rational agents. However, these assumptions often do not hold in sharing economies, which rely more on resource-based reciprocity and non-monetary exchanges [19, 38, 39]. As a result, the theoretical tools and equilibrium models developed for traditional markets fail to capture the unique dynamics of sharing economies. This leaves a significant gap on how sharing economies can be modelled using existing knowledge on the modelling of classical economies.
- Algorithmic economics is interdisciplinary, drawing from computer science, economics, and psychology to model market behaviour [40, 41]. Building a model of connected agents demands the design of coordinated protocols that account for trust, technical feasibility, and economic constraints [22]; at the intersection of software engineering and economics, such protocols become highly complex — and, if poorly implemented, can lead to significant financial risks and inefficiencies [42]. However, striking a balance between abstracting markets to enable efficient modelling through computational tractability and incorporating realistic interdisciplinary factors such as trust remains a persistent challenge [43].
  While significant progress has been made in incorporating realistic elements into abstract economic models — such as bounded rationality, diverse trading strategies, and asynchronous trading — many of these insightful results are studied in isolation and in traditional market settings. As previously noted, these findings do not easily translate to the context of the sharing economy. Current sharing economy research remains largely lacking in terms of combining both economic, social, and equilibrium concepts to fully explain how realistic market factors operate in sharing-based systems [44].

A complicating factor in this field is the vast body of relevant literature, dating back to Walras' theorems in 1874 [45]. The lack of standardised terminology across studies makes model definitions ambiguous and hard to distinguish [46, 47]. Assembling this literature into a coherent analysis is therefore a challenge in its own right.

1.3. Research Questions

Having established the research gap concerning the lack of applicability of existing economic theory to sharing economies — both terms of in their general structure and in capturing real-world aspects — we formulate the following research question:

How can models of behaviourally-driven distributed exchange capture the dynamics of sharing economies?

To approach this, we break down the research question into more manageable sub-questions:

1. How can we derive a theoretical model of the sharing economy from existing programs and distributed algorithms?
2. How can we model real-world factors, such as strategy diversity and manipulation incentives, to understand their impact on the trading dynamics of distributed market algorithms?

Through addressing these questions we develop and analyse a market model as a more realistic testing ground for trading protocols. By demonstrating how theoretically optimal algorithms can fail under certain conditions, the model helps identify where market designs or protocols may break down and how they might be improved. This approach advances both theoretical insight and practical guidance for designing robust, efficient sharing economy systems.

1.4. Notation and Model

To expand on the model from Figure 1.1, we formulate a more detailed market model of a competitive, decentralised sharing economy that captures the practical cases described in Section 1.1 to some level of abstraction. The exchange market under consideration is time-slotted market, where agents trade sequentially for a horizon of T timeslots, constrained to some network $\mathcal{G}$. $\mathcal{G}$ is defined on nodes (agents) $\mathcal{N}$ and edges (connections) $\mathcal{E}$: $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. This network represents the trading options of the agents, where each agent i may only trade with direct neighbours that are in its neighbourhood $\mathcal{N}_i = \{j \mid (i,j) \in \mathcal{E}\}$.

There is a set $\mathcal{N}$ of n prosumer agents. Each of these have the intention to trade and obtain some value from this trading. Each agent owns one specific resource, unique to that agent, i.e. agent i owns resource i and agent j owns resource j, and so forth. Agent i's resource has some positive (static) per-unit value v_i. The agents trade these resources without the explicit assumption of a monetary device. This is exactly in accordance with Strathern [48]'s definition of a barter market as a market where diverse objects are exchanged freely without the use of an abstract value measure.

Each timeslot, the agent produces some amount of its resource. Because of the stochastic character of the market, the agent does not produce the exact same amount each timeslot. Instead, the agent produces some amount of distributable resource $D_i(t)$ at timestep t, sampled stochastically with a normal distribution. Consequently, the long-term average of $D_i(t)$ approaches the endowment value ε_i. So $D_i(t) \sim N(\varepsilon_i, \sigma^2)$, $E[D_i(t)] = \varepsilon_i$, and at $t \to \infty$ we have $\overline{D_i}(t) = \varepsilon_i$.

With its distributable resources $D_i(t)$, agent i decides on an allocation, expressed in the allocation-vector $\mathbf{x}_i$, by allocating portions x_{ij} of its distributable resource D_i to other agents j. The allocations are summarised in the allocation matrix $\mathbf{X}$. Each agent allocates no more than its distributable resources each timeslot, i.e. $D_i \geq \sum_{j \in \mathcal{N}_i} x_{ij}$. As a visual aid, Figure 1.2 presents an example of trading.
From the trades in timeslot t, agent i receives resources $\mathbf{r}_i = [x_{ji}]_{j \in \mathcal{N}_i}$. These resources provide the agent with utility, defined by the utility function $u_i(\mathbf{r}_i)$. This utility function is taken to be linear, meaning agent i receives total utility $u_i(\mathbf{r}_i) = \sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji}$.

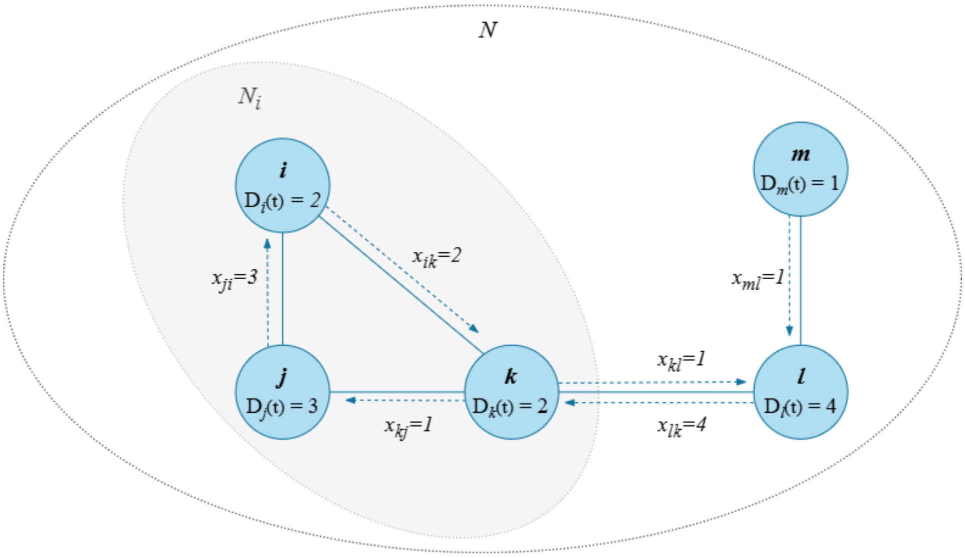

Figure 1.2: An example of how trading works in the defined exchange market: Agent i's neighbourhood consists of agents j and k. At this timeslot t, agent i allocates all of its 2 resources to neighbour k, while receiving 3 resources from neighbour j.

Through these trades, the market may (or may not) approach a desired state known as the equilibrium. This process we call convergence towards the equilibrium. We define the equilibrium as the conditions that the market clears (i.e., each agent allocates all its resources) and all agents maximise their utility. More formally:

Definition 1.4.1 (**Equilibrium**)**.** An allocation matrix $\mathbf{X}$ is an equilibrium allocation matrix $\mathbf{X}^*$ if the following holds:

- Market Clearance: $\forall i \in \mathcal{N}, \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i$
- Utility Maximisation: $\forall i \in \mathcal{N}, \mathbf{X} \in \arg\max_{\mathbf{X}} u_i(\mathbf{r}_i)$

With this definition of an equilibrium, we refine the objective of finding programs that model the exchange market: We investigate programs, both central and distributed, that model the trading decisions $\mathbf{x}_i$ of prosumers in this exchange market. These programs compute an allocation matrix $\mathbf{X}$ that satisfies the conditions of Market Clearance and Utility Maximisation of definition 1.4.1.

1.5. Contributions

Building on this concept of an exchange market, this work makes several important contributions to advancing the theoretical understanding and computational modelling of sharing economies:

1. Using an in-depth literature review, we clarify key conceptual gaps, especially between (price-based) classical models and the features of trading behaviour in real-world sharing economies. It becomes clear that an exchange market framework which is (1) aligned with the characteristics of sharing economies, (2) abstracted sufficiently to express various practical instances, and (3) grounded in classical economic theory, offers a natural way to think about the sharing economy.

2. We prove that the Eisenberg-Gale program [49], Shmyrev's program [50], and Jain's program [51], can be adapted to compute the equilibrium in the exchange market. The main modification being the conversion from explicit prices to inherent resource values. This re-

sult highlights that modelling sharing economies does not require reinventing the wheel; rather, it calls for a principled reinterpretation and reformulation of existing theory.

3. We demonstrate that two distinct distributed trading strategies (one based on proportional response and one on a greedy behaviour) both approach the equilibrium in our exchange market model, despite relying on different behavioural assumptions. We found that the proportional strategy aligns with mirror descent in this context — a connection thusfar unproven in this specific market setting — and elucidate the dependence of distributed algorithms on the consensus on resource values.
   Additionally, we prove that any linear combination of optimal distributed algorithms remains optimal. This reveals the possibility to design more behaviourally diverse protocols without sacrificing fairness and stability.
   We use convergence guarantees established in idealised settings as a basis for analysing and informing more realistic market models, noting that while the proposed strategies are not universally optimal, their simplicity aligns well with explainable mechanism design.

4. Finally, we find that extending distributed algorithms with realistic behavioural features (namely, heterogeneous strategies, asynchronous updates, memory decay, and manipulative behaviour) significantly alters trading dynamics and outcome fairness.
   We find that incomplete information undermines market integrity and increases susceptibility to manipulation; intriguingly, however, the agents who manipulate the market are not always the ones who benefit the most. Effective protocol design must incorporate memory mechanisms and address these vulnerabilities to maintain trustworthiness.
   Additionally, consensus emerges as an influential factor on market robustness and fairness, making it especially important for market regulation and design to encourage alignment in environments characterised by diverse behaviours and principles.
   Overall, we find that effective decentralised market modelling requires managing strategic diversity through coordination and incorporating factors like timing and memory to ensure robustness and fairness in practical applicability beyond idealised theory.

1.6. Roadmap

The thesis has a two-part structure: it presents both a theoretical analysis and numerical experiments, which are closely connected as shown in Figure 1.3. This approach allows us to clearly distinguish between the two main research gaps addressed. Before these core sections, we conduct an in-depth literature review (Chapter 2) to map the extensive work in this field and clarify where current models of sharing economies fall short, motivating Contribution 1.

Part I: Theoretical Analysis

Part I demonstrates how existing economic literature can be leveraged to accurately model the exchange market. Using the insights gathered in the literature review we analyse existing programs and propose adapted distributed models that fit the exchange market (Chapter 3) to form Contribution 2. Then, we analyse distributed algorithms that model the trading decisions of agents (Chapter 4) and compare them to the previously discussed programs. The discussion on these theoretical models, supported by several mathematical proofs, finish up Contribution 3.

Part II: Numerical Experiments

The analysis of Part I provides the support for the proposal of a novel exchange market model (Chapter 6). Additionally, the model is extended with aspects that tackle the limitations of the current state-of-the-art. We conduct numerical experiments to explore how the model's components—strategy heterogeneity, asynchronous dynamics, memory limits, and manipulative strategies—shape market phenomena (Chapter 7). These explorative experiments aim to examine trade-offs and risks in exchange markets and to assess whether the proposed model captures aspects of real-world market behaviour often overlooked by other approaches. Together with the discussion of these results, the experiments conclude Contribution 4, as the last contribution of this research.

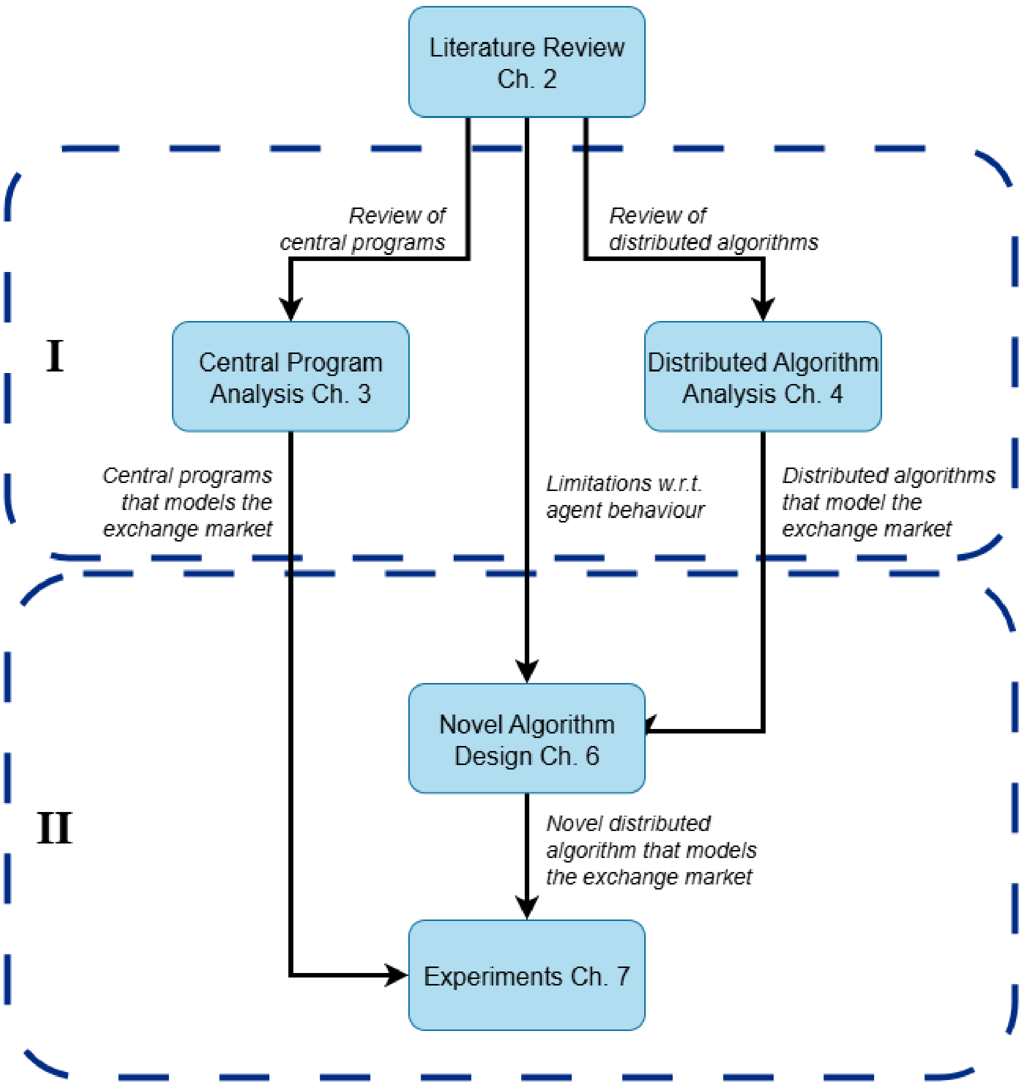

Figure 1.3: Roadmap of this Thesis

2

Literature Review

This literature review covers key models and methods in market equilibrium theory. The selection of previous works, based on scoping [52] and citation chaining (otherwise known as snowballing[1]), traces developments from classical Fisher and Arrow-Debreu models to decentralised trading mechanisms. The goal is to identify core concepts and computational techniques, while critically examining limitations in the recent state-of-the-art, especially regarding agent behaviour, equilibrium convergence, and applicability to real-world sharing economies.

2.1. Preliminaries

2.1.1. Equilibria

Studying equilibria and how to compute them has long been a key focus in economics [53]. From a system-wide perspective, the most desirable market state is such an equilibrium: a market state that is both fair and stable. An often used metric for fairness is the proportionality, where agents receive returns proportional to their contribution (e.g., [54, 55, 56]). From the individual agent's perspective, however, the most desirable state of a market is a state in which it maximises their utility function. In a state in which that agent maximises its utility, the agent has no incentive to deviate from the current state, making that market state inherently stable.
Although many types of equilibria exist, the most common for competitive markets (e.g., Fisher market, AD market) is the Competitive Equilibrium (CE), which combines both system-level interests and agent-level interests. This state is Pareto-optimal[2], meaning that no resources can be redistributed to benefit one agent, without diminishing the utility received by another agent. Given some monetary budget m_i for each agent i, we can define the CE as follows:

Definition 2.1.1 (**Competitive Equilibrium**). An allocation matrix $\mathbf{X}$ and set of prices $\mathbf{p}$ form a Competitive Equilibrium if the following holds:

- Market Clearance: $\forall i \in \mathcal{N}, \sum_{j \in \mathcal{N}} x_{ij} = \varepsilon_i$;
- Budget feasibility: $\forall i \in \mathcal{N}, \sum_{j \in \mathcal{N}} x_{ij} \cdot p_j \leq m_i$;
- Utility Maximisation: $\forall i \in \mathcal{N}, \mathbf{X} \in \arg\max_{\mathbf{X}} u_i(\mathbf{r}_i)$.

Proving these three conditions to be true means that the market has reached equilibrium. Of the three conditions, market clearance, is quite straightforward: Each agent allocates or trades all their resources. The utility maximisation condition is often the most complex to prove. However, budget feasibility is particularly interesting in exchange markets, which lack a monetary device. In

[1] https://pure.tudelft.nl/ws/portalfiles/portal/9993402/how_to_write_a_literature_review_paper_version_after_review_no_track_changes.pdf
[2] https://en.wikipedia.org/wiki/Pareto_efficiency

markets with explicit prices, the budget m_i and endowment ε_i are separate possessions of agent i; in contrast, in exchange markets, agents trade directly using their endowments, so budget and endowment represent the same resources. As a result, the budget feasibility constraint and the market clearance constraint can be merged. This aligns with a well-established result in literature: Equilibrium computation is effectively utility optimisation restricted by packing constraints.
In the special case in which all agents have equal budgets $\forall i \in \mathcal{N}, \varepsilon_i = 1$, the CE is called the Competitive Equilibrium from Equal Incomes (CEEI) [57]. Both the CE and the CEEI have desirable properties, including the maximum Nash Social Welfare and Pareto efficiency, as well as envy-freeness and proportional fairness [58].

A closely related concept is the Sharing Equilibrium (SE) [59], which is specifically applicable to exchange markets. This equilibrium is defined on sharing ratios ρ: an agent's ratio of value received to value allocated.
Definition 2.1.2 (**Sharing Equilibrium**). A vector of sharing ratios qualifies as equilibrium sharing ratios $\rho^* = [\rho_i^*]_{i \in \mathcal{N}}$ in a Sharing Equilibrium if the following holds:

- Utility maximisation: Agent i allocates resources to agent j such that i maximises its expected received resource R_i;
- Feasibility: The allocations $\mathbf{X}$ satisfy the constraint that no agent i receives more resources than it is entitled to by the equilibrium sharing ratio ρ_i^* and endowment ε_i: $R_i \leq \varepsilon_i \cdot \rho_i^*$;
- Rational behaviour: Each agent i expects to receive resource x_{ij}/ρ_j^* from agent j after allocating resources x_{ij}.

Although this definition of an equilibrium is aligned with the CE as described earlier, it is specifically tuned to a market in which there is no explicit monetary tool.

2.1.2. Markets

A well-known market, similar to the exchange market, is the Arrow-Debreu (AD) market[3] [60]. In this market, agents have an initial endowment of resources and utility functions defined on those resources. Agents sell their initial resources and use the proceeds to acquire others' resources. This market was originally formulated by Walras [45], who questioned whether such a market is guaranteed to have an equilibrium. Arrow and Debreu [60] confirmed that indeed this equilibrium exists, provided the agents' utility functions are concave. Hence the name Arrow-Debreu market.

Fisher markets are closely related to AD markets, the key difference being how agents derive their budget: Fisher markets assign budgets externally, while AD markets derive budgets from resource sales. The endogenous value of money in an AD market induce a double effect of prices, affecting both demand and budgets. As such, the AD market is often seen as a more complex market, while the Fisher market is considered to be a simplified case of the AD market [61]. The difference in complexity is exemplified by the lack of a mathematical program with a constructive proof for the AD market, while there is such a program for the Fisher market [51].

Exchange markets differ from Fisher and AD markets in several key ways. While Fisher and AD markets are typically modelled as centralised systems with complete information — where agents act as buyers with fixed budgets (Fisher) or prosumers with initial endowments (AD) — exchange markets feature decentralised, P2P interactions where agents assume trade resources multilaterally. Unlike the monetary mediation of trades in Fisher and AD markets, exchange markets often operate without money, relying on local, platform-facilitated exchanges. Additionally, Fisher and AD models emphasise efficiency, budget feasibility, and equilibrium under idealised assumptions, whereas exchange markets often prioritise finding an equilibrium under trust, reciprocity, and adaptability within networked platforms. Table 2.1 highlights these and other distinctions, showing how traditional models diverge from the fundamentals of exchange markets.

[3] The term *exchange market* is often used synonymously with AD markets, where agents trade money and goods. However, this thesis uses *exchange market* to refer to an abstract model of a sharing economy without explicit money exchange.

	Fisher Markets	AD Markets	Exchange Markets
Agent Roles	Agents are buyers with fixed budgets or sellers with fixed resources; buyers spend budgets to maximise utility.	Agents are prosumers with fixed resources.	Agents are prosumers, flexibly trading in peer-based roles.
Exchange Mechanism	Money and prices mediate trades; agents come to the market with a fixed budget.	Money and prices mediate trades; agents earn a budget by selling initial resources.	Exchange is often moneyless, with decentralised, local platform-matched trades based on reciprocity.
Market Design & Principles	Often modelled as centralised market with complete information. Emphasises efficiency and budget feasibility.	Often modelled as centralised market with complete information. Focus on equilibrium and optimal allocation.	Decentralised systems shaped by platforms and underlying networks. Emphasises trust and reciprocity.

Table 2.1: Fundamental differences between three types of markets.

2.1.3. Computational Efficiency

The computational complexity of methods used for computing equilibria has been a topic of interest for decades [62]. The computation of a CE for the AD market in general has been established to be PPAD-hard (Polynomial Parity argument on a Directed Graph) [63]. The PPAD class of problems refers to a subclass of TFNP (Total Function Nondeterministic Polynomial) that can be reduced, in polynomial time, to the directed graph problem called *End-of-a-Line* [64]. The relation of these complexity classes to well-known classes P and NP can be seen in Fig. 2.1. The PPAD class of problems is guaranteed to have a solution as a result of mathematical proof considering its specific combinatorial structure.

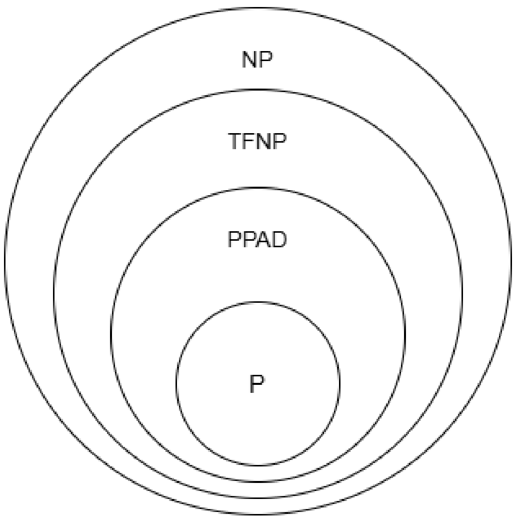

Figure 2.1: The relation of complexity-classes PPAD and TFNP to P and NP.

Although imposing restrictions on the market can decrease the complexity of computing the equilibrium, guaranteeing polynomial-time solutions remains challenging. The computational complexity of finding equilibrium solutions is a significant concern: Jain [51] argued that the market of rational agents can be interpreted as a Turing machine trying to compute the next move. If a more complex Turing machine (a modern computer) cannot compute the equilibrium within a reasonable time, then it is unreasonable to expect a simple Turing machine (the market) to be able to compute the equilibrium. One could thus argue that it is then more valuable to have an uncomplex approximation of the equilibrium, rather than a computationally complex exact equilibrium.

2.2. Seminal Works

This section surveys key results underpinning modern equilibrium computation. While not limited to money-free exchange markets, these models provide essential context for understanding how sharing economies can be modelled within broader market equilibrium theory.

2.2.1. Tâtonnement

Walras [45] introduced the tâtonnement process as an early approach to finding market equilibria. Tâtonnement is a sequential process which adapts prices due to excess demand. Each step, agents experiencing excess demand, adjust their prices upwards, and vice versa for a lack of demand. Cheung et al. [65] later showed this process to converge to the CE for Fisher markets by drawing a parallel between tâtonnement and gradient descent over a convex function expressing the negative excess demand.

2.2.2. Eisenberg-Gale program

Eisenberg and Gale [49] proposed a relatively simple program, based on the geometric mean of agents' utility functions, that can be used to compute the CE in a Fisher market:

$$\max_{\mathbf{X}} (\Pi_{i \in \mathcal{N}} u_i(\mathbf{r}_i)^{m_i})^{1/\sum_{i \in \mathcal{N}} m_i}$$

Constrained to the endowments and the budgets, and taking the log of the above function results in the following definition of the Eisenberg-Gale (EG) program:

$$\begin{aligned} \max_{\mathbf{X}} \quad & \sum_{i \in \mathcal{N}} m_i \log u_i(\mathbf{r}_i) \\ \text{s.t.} \quad & \forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}} x_{ij} \leq 1, \\ & \forall i, j \in \mathcal{N} : x_{ij} \geq 0 \end{aligned} \tag{P1}$$

.

Closely related to the EG program is Shmyrev's program [50], which is actually the dual of the dual of the EG program, up to a change of variables [66]:

$$\begin{aligned} \max_{\mathbf{b}} \quad & \sum_{i,j \in \mathcal{N}} b_{ij} \cdot \log v_{ij} - \sum_{j} p_j \cdot \log p_j \\ \text{s.t.} \quad & \forall j \in \mathcal{N} : \sum_{i \in \mathcal{N}} b_{ij} = p_j, \\ & \forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}} b_{ij} = m_i \\ & \forall i, j \in \mathcal{N} : b_{ij} \geq 0 \end{aligned} \tag{P2}$$

where the agents place bids b_{ij} on resources, and allocations are determined by the ratio of bids and prices ($x_{ij} = b_{ji}/p_i$). v_{ij} denotes how much agent i deems j's resource to be worth. This notation means agents value a resource differently from one another, as opposed to v_j where agent j's resources has the same value to each other agent.

2.2.3. Proportional Response Dynamics

Proportional Response Dynamics (PRD) are a simple, distributed dynamics, which provably converges to the CE [26]. Originally set in the context of bandwidth allocation, much of the research concerning distributed equilibrium-finding has been based on these dynamics, especially in Fisher markets.

In the most common interpretation of PRD, agents allocate their resources sequentially based on how much utility they gained in the previous timeslot: In the first step, each agent i places a bid b_{ij} (within its budget m_i) on the resource of agent j. Next, each agent divides its own resource among its trading partners, in proportion to how much those partners bid on the resource: If agent j offers a higher bid to agent i, then agent i allocates more of its resource to agent j. Once all allocations are made, each agent i receives a bundle of resources x_{ij} from others and gains utility from them. Finally, agents update their bids, proportionally to the utility the agent received from that partner's resource in the previous timeslot. An example of this mechanism is shown in Figure 2.2.

In the case of an exchange market, the bidding step is skipped entirely. Agents instead allocate their resources directly in proportion to the utility they received last timeslot — more closely resembling the original PRD [26]. There have been multiple adaptations of this mechanism, including, for example, asynchronous PRD [67], damped PRD [68], and lazy PRD [69].

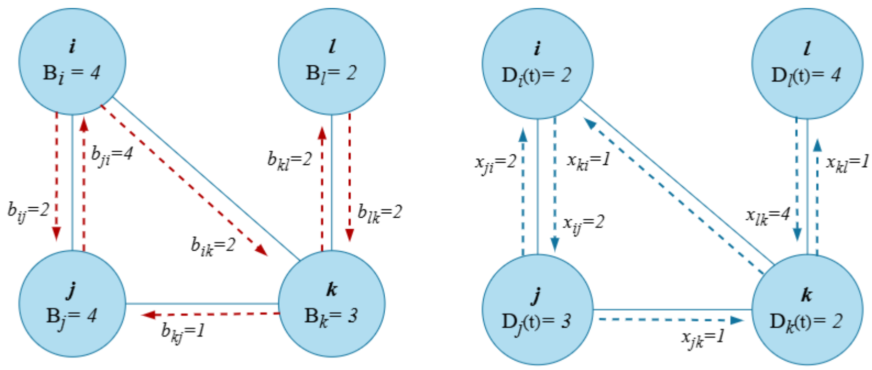

Figure 2.2: Example of a PRD bidding and allocation timeslot

The convergence of PRD in Fisher markets can be better understood by connecting it to mirror descent. The PRD update rule resembles how mirror descent would solve Shmyrev's program [70]. This fundamental insight provides both convergence guarantees and an intuitive interpretation of PRD: PRD is not just a heuristic, but can be viewed as performing a principled optimisation over the utility landscape defined by the market.

2.2.4. AD Market models

Early algorithms to compute the CE in an AD market were approximations through fixed-point or Newton-based methods. Additionally, the running time of these algorithms was exponential, making them computationally expensive [61]. Jain [51] demonstrated how polynomial time solutions could be obtained by defining the equilibrium as a solution to a central convex program:

$$
\begin{aligned}
&\forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}} x_{ij} = 1 \\
&\forall i, j \in \mathcal{N} : x_{ij} \geq 0 \\
&\forall i, j \in \mathcal{N} : \log p_i - \log p_j \leq \left(\frac{\sum_{k \in \mathcal{N}} v_{ik} \cdot x_{ki}}{v_{ij}} \right)
\end{aligned}
\tag{P3}
$$

This program computes any and all equilibrium solutions of the AD market[4]. To complement this convex program, the authors also proposed the first polynomial exact algorithm for computing the CE in the AD market, using the ellipsoid method.
The similarities between the convex program proposed by Jain and the KKT-conditions of the EG program are evident, and there exists a variation on Program P3 that Jain interprets as a generalisation of the EG program. Jain's convex program for the linear AD market is however quite different from the EG program and Shmyrev's reformulation, in the sense that it does not confirm existence of an equilibrium in an AD market, but only computes it if such an equilibrium exists. Additionally, the proof is non-constructive, unlike the proof of the EG program.

Jain's work inspired several new findings; Ye [72] presented an interior point algorithm to solve the program in $O(n^4 \log(1/\epsilon))$ time. Later, Duan and Mehlhorn [61] found the first combinatorial polynomial algorithm for the AD market that performs in $O(n^6 \log(n \cdot U^{max}))$, where U^{max} is the maximum integer utility. It is based on balancing network flow, and at a high-level its price-adaptation mechanism seems to share similarities with the basic ideas of tâtonnement. This algorithm was inspired by Devanur et al.'s work who worked with flows to find algorithms for the linear AD market [74]. Eventually, Garg and Végh [75] developed the first strongly polynomial algorithm for AD markets. Unlike earlier approaches, its convergence does not depend on the desired precision parameter ϵ.
Together, these developments show how ideas from convex programming and flow-based algorithms have gradually improved our ability to compute equilibria in AD markets. While challenges remain, this work has laid useful groundwork for tackling similar problems in more general exchange settings.

2.3. Recent Trends

Having reviewed the seminal works that form the foundation of equilibrium computation, we now turn our attention to recent trends in the field. These state-of-the-art works build directly on the seminal works, but aim to address some of their key limitations. By examining how recent research has narrowed the gaps left by earlier seminal works, we can more clearly identify which challenges have been addressed and which remain open. This, in turn, helps establish a foundation for modelling dynamic distributed exchange in P2P markets based on the fundamental economic concepts discussed earlier in this chapter. We frame our critique around four core assumptions that limit their ability to capture the distinctive features of exchange markets:

1. **Stochasticity of markets**: the modelling of static environments of markets.
2. **Distributed & asynchronous exchange**: the centralised manner of equilibrium computation and assumption of synchronous concurrent trading actions.
3. **Heterogeneous allocation strategies**: the assumption of homogeneity in the type of decision rules agents use.
4. **Bounded rationality**: the assumption of agents to be entirely rational, with full information availability.

While there is undeniably a large body of research on P2P networks that addresses some of these critiques, much of it consists of narrowly scoped, case-specific studies[5]. Though these approaches yield valuable insights, they often fall short of offering generalizable design guidance. This leaves a gap between conceptual understanding and practical application in designing viable P2P systems. Accordingly, our focus in recent trends shifts more toward abstracted models grounded in the seminal works discussed earlier.
The discussed papers are summarised in Table 2.2, to provide an overview of how the current state of equilibrium research compares to the objectives of this research.

[4]Similar convex programs were actually found earlier by Nenakov and Primak [71] in 1983.
[5]A particularly large share of this focusses exclusively on energy markets.

	Market	Contribution
[76]	Fisher market with central allocator and price updates.	Simple price updates under belief heterogeneity and bounded rationality.
[77]	Fisher market with central allocator and bidding buyers.	Dynamic allocation with sequential agents and private stochastic types.
[78, 79, 56]	Fisher market with central allocator and online arrival.	Online allocation mechanism with unknown stochastic item arrivals.
[80]	Allocation market with central allocator to multiple agents.	Online allocation with unknown demand of sequentially arriving agents.
[81]	Fisher market with central seller and buyer bidding.	Online pricing with fluctuating values and private buyer information.
[67]	Fisher market with asynchronous proportional bidding.	Asynchronous bidding converges via market-based game formulation.
[59]	Exchange market with decentralised peer trading.	Greedy, myopic prosumers in markets with stochastic resource availability.
[82]	Arrow-Debreu market with decentralised peer exchange.	Trading under heterogeneous beliefs and learning in non-rational equilibria.

Table 2.2: Overview of state-of-the-art research, highlighting market types, mechanisms, and agent assumptions relevant to current research gaps.

2.3.1. Stochasticity

A long-standing shortcoming in economic theory is the focus on static markets [83] — in which variables such as the budgets and endowments remain fixed over time (e.g., [84, 85]) — when real exchange markets are generally subject to change [81, 86].
Researchers have often opted for simplified static models for the reason that incorporating market dynamics can introduce significant complexity. Modelling a market as a dynamic system leads to a dynamic equilibrium [86], implying that the quality and approximation of the equilibrium changes over time, resulting in a higher complexity of analysing equilibria in dynamic markets.

A number of recent works addressed these issues by including market stochasticity, specifically modelled as dynamically available resources; Yang et al. [81] considered this stochasticity in a Fisher market, with an online version of PRD which allocates stochastically-valued resources to buyers. This online PRD achieved logarithmic convergence rates in markets with strongly convex market objectives.
Another approach, by Gao et al. [78], was to model sequentially arriving resources and leverage dual averaging to drive a Fisher market to equilibrium. In this model, termed PACE, a central allocator must make pricing decisions, by ensuring buyers' time-averaged utilities converge to the equilibrium. So although PACE includes considerations for the dynamic character of a market, it does rest on the assumption of some central allocator. PACE was later extended upon by Liao et al. [79] and Yang et al. [56], who achieved logarithmic convergence bounds when modelling inter-arrival times of the resources as non-stationary and parameter-free, respectively.
Turning to exchange markets, Georgiadis et al. [59] presented a greedy distributed algorithm converging to equilibrium in settings with stochastic resource availability. Beyond a few efforts, research on stochasticity in exchange markets remains limited compared to Fisher markets. Future work should thus build on these to develop dynamic models that handle stochastic resource variations in exchange markets.

2.3.2. Distributed & Synchronous computation

Understanding how agents arrive at an equilibrium is just as important as identifying the equilibrium itself [87]; markets are rarely governed by a single resource-allocating authority, and agents typically interact in a decentralised, uncoordinated fashion [88]. Even if we can compute equilibrium efficiently, it does not follow that independent agents will arrive at that market state. Thus, it is non-trivial to study whether and how equilibrium can emerge from these local interactions.

Nevertheless, foundational models prioritise centralised computation over decentralised process. We refer to such approaches, which ignore the trading process and compute the equilibrium from a centralised perspective, as central programs. In much of the literature, an equilibrium is treated as an output to be computed by an omniscient entity, rather than as an emergent property of agent interaction. A recurrent justification for this choice in literature is the objective to compute the equilibrium state efficiently, rather than simulating the market process. The gradual decentralised processes observed in real exchange markets are not guaranteed to converge to an equilibrium [89] and a distributed, agent-oriented complicates the convergence to equilibria [90]. As such, decentralised processes are often not seen as the right tool for this goal. Therefore, although the process of getting to the equilibrium is important, many works deem it a better choice to compute the equilibrium state with a central program without constraints on the process leading there.

Some seminal works seemingly address the concern of the decision-making process towards the equilibrium: PRD and tâtonnement. However, the latter makes the assumption that no trading actually takes place until the equilibrium prices have been established, and thus tâtonnement cannot really be categorised as a process-oriented model. PRD on the other hand, clearly states the sequential updates of allocations each timeslot, and is thus a process-oriented model. However, PRD assumes that all agents trade proportionally and synchronously: assumptions that are unrealistic in large, decentralised markets. Enforcing synchronous trading can introduce artificial coordination that misrepresents the inherently asynchronous nature of real-world exchange markets, and impose unnecessary overhead that grows with market scale [91, 92]. Therefore, although seemingly there were process-oriented models among the seminal works, these by far do not realistically model the process of trading.

More recently, the restriction of synchronous trading has been addressed to some degree. For instance, Kolumbus et al. [67] showed that a Fisher market, where an adversary asynchronously selects batches of agents to actively trade, can still reach equilibrium. Earlier, Cheung and Cole [93] used a similar strategy to show the convergence of asynchronous tâtonnement. The authors of both papers use the connection of these algorithms to some type of descent algorithm to show that, each asynchronous step improves some potential function that models the market, thus guaranteeing convergence. In a related direction, Sinclair et al. [80] and Jalota and Ye [77] found that the allocation of resources in a Fisher market where agents arrive sequentially is only guaranteed to converge when utility functions are known and uses prediction based on past observations to overcome this limitation.

These approaches suggest that partial information and timing uncertainty do not necessarily prevent emergence of an equilibrium: Asynchronous mechanisms can sometimes cause greater stability in some equilibrium-computing processes [94, 95], while at the same time it can make determining the convergence rate a more difficult task [67]. The effects of asynchronous dynamics depend on the market structure as well as the agent selection strategy and their response to market conditions.

While most of these results are limited to Fisher markets, extending them to exchange settings remains a promising direction. Overall, models of market interaction should move away from unrealistic assumptions of synchronous allocation and centralised computation, and instead reflect the asynchronous, decentralised nature of real-world dynamics.

2.3.3. Heterogeneous Allocation Strategies

Another recurring limitation of exchange market models is the lack of heterogeneity in the employed allocation strategies among the population of agents [96]. Although the presence of different levels of wealth often indicates heterogeneous strategies [97], surveys have shown that many models fall short in modelling this aspect [98]. The vast majority of the heterogeneity that we do see in market-modelling research is rooted in the different preferences and resources agents have (e.g., [99, 56]), leaving strategy heterogeneity relatively unexplored. The justification for the choice of homogeneous strategies across the population is often superficial and relies on broad assumptions without in-depth empirical or theoretical support.

Some variations on well-known trading rules have started to shed light on a heterogeneous strategy population; for instance, several studies have examined variations on PRD, such as lazy PRD [69]. Le and Ramazi [100] studied markets with mixed-strategy populations and found that, although mixed-strategy populations often reach equilibrium, the equilibrium becomes unstable. This suggests that strategy heterogeneity could have adverse effects on the convergence of markets. The authors leave further insights on the convergence of mixed-strategy markets as future work.

However, the most advancement in terms of heterogeneous strategies can be found in terms of manipulative strategies. Multiple works have studies the impact of having one or more manipulative agents among truthful agents (e.g., [101, 102]).
One way to assess a market's robustness to manipulative trading is through the Incentive Ratio (IR) [103], which quantifies how much an agent can potentially gain by employing a manipulative strategy. The incentive ratio w.r.t some agent i and market M is defined as follows:

$$\zeta_i^M = \max_{s,s'_i} \frac{\max_{\mathbf{X}_{(s_{-i},s'_i)}} u_i(\mathbf{r}_i)}{\min_{\mathbf{X}^*_{(s)}} u_i(\mathbf{r}_i)} \tag{2.1}$$

where s is a configuration of truthful strategies, where s_i denotes agent i's strategy and s_{-i} denotes all strategies except that of agent i. Its counterpart, s', denotes the same concepts for untruthful strategies. $\mathbf{X}_{(s)}$ is the allocation matrix that the strategy configuration s has converged to. In other words, the IR computes in the numerator the maximum possible utility obtained by agent i by employing manipulative strategy s'_i, while all other agents keep a truthful strategy, while in the denominator it computes the minimum obtained utility for agent i is the whole population trades truthfully.

The total IR of the market M for misreporting such weights is then:

$$\zeta^M = \max_{i \in \mathcal{N}} \zeta_i^M \tag{2.2}$$

As of yet, most research on the IR has been in the context of Fisher markets [104, 101, 102], where relatively tight bounds have been found for the IR. For instance, when agents misreport their utility functions in Fisher markets, the IR has been shown to have an upper bound of 2 [103, 101]. However, for AD markets, the situation is more complicated. Even basic forms of manipulation, such as misreporting utilities, can lead to unbounded IR values [105]. More progress has been made in the context of IR in AD markets, mostly focussed on the specific case of Sybil attacks[6]. For example, Cheng et al. [106] found that the IR of Sybil attacks in AD markets is at most 3 for general networks. More specific results include an IR of 2 for tree-structured networks [107] and $\sqrt{2}$ for complete networks [108].

While recent studies have started to address behavioural heterogeneity in manipulative strategies, this area remains underexplored [109]. This gap is especially notable for exchange markets as we

[6] https://en.wikipedia.org/wiki/Sybil_attack

define them, since most IR research focuses on Fisher and AD market models, and no formal IR results currently exist for exchange markets. Thus, alongside heterogeneity in truthful allocation behaviour, there is substantial room for progress in understanding manipulative behaviour within exchange markets.

2.3.4. Bounded Rationality

Frequently, researchers assume that agents' capabilities are sufficient to make logical decisions to satisfy a given strategy [110]. Although unrealistic, such an assumption does simplify the model. One concept that could be used to tackle this trade-off is bounded rationality, which acknowledges that agents may have cognitive or informational limitations while still assuming some form of computational ability. Studies have shown that assuming bounded rationality yields more realistic models and can meaningfully affect equilibrium outcomes [111, 112]. Bounded rationality offers a promising middle ground between unrealistic rational models and overly simplistic equilibrium approaches [113].

Although seminal works do not consider any type of bounded rationality, some recent work incorporates bounded rationality explicitly: Dvijotham et al. [76] considered that agents' responses to trading events should be simple, such that the algorithm would be applicable to computationally bounded agents, and thus modelled belief-forming with finite horizons in Fisher markets. Their findings suggest that convergence to the equilibrium is still possible under simple heterogeneous beliefs. Similarly, Choo et al. [82] considered equilibrium-finding with information asymmetries in an AD market and shows that market-wide beliefs can emerge even when some agents are rationally bounded by incorrect information. Angeletos and Sastry [114] proposed a more intricate model, where agents' attentions can vary in an AD market, and provide different conditions under which heterogeneous attention spans may or may not be disruptive to the system.

Still, in otherwise advanced models bounded rationality has often been overlooked. For instance, the works on the PACE algorithm ([78, 79, 56]) can be considered state-of-the-art, however it assumed a fully rational central allocator with perfect memory, conflicting with their explicitly stated aim to model more realistic markets. We stipulate that these algorithms might show different convergence if the allocators were subject to cognitive constraints. Similarly, Kolumbus et al. [67] considered asynchronous dynamics, but left the effects of imperfect information availability as an open question. This leads to the question of whether limitations in information processing could disrupt algorithms that usually converge well in exchange market models. To answer this question and better reflect real behaviour, market models should explicitly incorporate bounded rationality and information-processing constraints.

In the examples above one specific category of bounded rationality becomes apparent: the availability of information. Many equilibrium models assume unrealistically complete knowledge on the part of agents [115]. One way to address this is by incorporating limited memory or memory decay, where agents gradually forget past information. Such simple cognitive limitations offer a pragmatic way to introduce bounded rationality without over-specifying agent psychology, thereby preserving the appeal of equilibrium theory: the freedom of choice [116]. As such, a simple concept of bounded rationality can really add to the realism of the market, without restricting the agents too much, striking balance between over- and under-specification of the cognitive capabilities of agents.

2.4. Conclusion

The discussion on recent trends in this field of research has revealed several limitations concerning the state of current literature: Market models tend to rely on centralised mechanisms and make unrealistic assumptions about agent behaviour, information, and market dynamics. While foundational models offer valuable analytical tools w.r.t. optimality, they typically assume static markets, full rationality, homogeneous strategies, and synchronous decision-making — assumptions that fail to capture the complexity of real-world exchange markets.

In fact, the most progressive state-of-the-art works in terms of modelling dynamic and varied agent behaviour have been focussed on the Fisher market, which is fundamentally different from the exchange market. The limited attention paid to decentralised exchange markets highlights a critical shortcoming in literature.
The identified limitations define a core research gap on the lack of algorithmic models capturing realistic trading behaviour and structural constraints in decentralised sharing markets. This directly informs Contribution 1 in understanding how current models fail to represent realistic trading behaviour in sharing economies.

Part I

Theoretical Analysis

3

Centralised Approaches

Our starting point in modelling the market is to identify centralised programs that compute equilibrium allocations. Although these do not directly capture the distributed character of the market, they play a vital role in assessing and explaining the convergence behaviour of the distributed allocation strategies introduced later.
In light of this, we examine how existing central programs can be adapted to the exchange market. Multiple programs that are significant to the field and connect meaningfully to distributed market models are considered. Each of these programs can model the exchange market described in Section 1.4 and are able to identify its equilibrium state.
With such central programs identified, we discuss the suitability of these programs for computing the equilibrium solution that will serve as a benchmark. This benchmark allows us to evaluate and compare the results of distributed allocation strategies in later chapters.

3.1. Adaptation of the EG Program

The original EG Program (P1) can compute a market equilibrium for the Fisher market. However, the Fisher market is a significantly different market from the one described in Section 1.4, and so the EG Program must be adapted in order to model the exchange market.

First of all, the utility function has to be defined as a linear utility function such that $u_i(\mathbf{r}_i) = \sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji}$. Additionally, the weight m_i represents the value of the agents' resources, rather than a direct monetary budget. Depending on budget m_i, the geometric mean of the agents' utilities will shift, meaning that the computed equilibrium shifts along with it.
Rather than using monetary budgets, we define their weights based on the expected value of their own resources. Because of the stochastic availability of resources, the value of an agent's resources varies over time. However, for any central program considered in this thesis, we assume the long-term expected value of the distributable resource: $\mathbb{E}[D_i] = \varepsilon_i$. This assumption is similar to those made in other works (e.g., [77, 56]. Thus, the weight of agent i becomes the value of its endowment: $v_i \cdot \varepsilon_i$. This gives us the following adapted version of the Program P1:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{i \in \mathcal{N}} \left((v_i \cdot \varepsilon_i) \log \left(\sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji} \right) \right) \\ \text{s.t.} \quad & \forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} x_{ij} \leq \varepsilon_i, \\ & \forall i, j \in \mathcal{N} : x_{ij} \geq 0 \end{aligned} \tag{P4}$$

This program is constrained to the underlying network $\mathcal{G}$, by including the notation of neighbourhoods $\mathcal{N}_i$. The values v_i can, without loss of generality, be scaled to 1 by inversely scaling ε_i. However, retaining v_i offers two advantages: It allows for more straightforward interpretation and representation in terms of distinguishing between value and amount, and it aligns with related literature, where values v and amounts ε are typically treated separately, making comparisons more straightforward. Program P4 offers a clear model of the exchange market, taking on the same intuitive concepts of optimising the Nash Social Welfare, as in Program P1.
The optimality of this program for the market from Chapter 1.4 is proven through the KKT-conditions, as seen below.

Theorem 3.1.1 (**Optimality of Program** P4)**.** *Program P4 computes the equilibrium allocations* $\mathbf{X}^*$ *according to equilibrium definition 1.4.1.*

Proof. The Lagrangian of the objective function in Program P4 is as follows:

$$L(\mathbf{X}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_{i \in \mathcal{N}} \left((v_i \cdot \varepsilon_i) \cdot \log\left(u_i(\mathbf{r}_i)\right) - \lambda_i \cdot (\sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i) \right) + \sum_{i,j \in \mathcal{N}} \mu_{ij} \cdot x_{ij}$$

The derivatives of this function with respect to x_{ji}, λ_i, and μ_{ij} are

$$\frac{\partial L}{\partial x_{ji}} = v_j \cdot \frac{v_i \cdot \varepsilon_i}{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}} - \lambda_j + \mu_{ij}$$

$$\frac{\partial L}{\partial \lambda_i} = \varepsilon_i - \sum_{j \in \mathcal{N}_i} x_{ij}$$

$$\frac{\partial L}{\partial \mu_{ij}} = x_{ij}$$

Often, the dual variable λ_i is interpreted as the relative price of resource i at equilibrium. Given the exchange market, the concept of explicit prices is not a suitable interpretation. Rather, it is better to think of λ as the relative value of resources to be allocated to others (or 'paid') in order to receive an amount of returned value, i.e. the relative cost. Given these derivatives we can construct the KKT-conditions.

- primal feasibility: $\forall j \in \mathcal{N} : \sum_{i \in \mathcal{N}_j} x_{ij} \leq \varepsilon_i$
  $\forall i, j \in \mathcal{N} : x_{ij} \geq 0$
- dual feasibility: $\forall i \in \mathcal{N} : \lambda_i \geq 0$
  $\forall i, j \in \mathcal{N} : \mu_{ij} \geq 0$
- complementary slackness: $\forall i \in \mathcal{N} : \lambda_i \cdot (\sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i) = 0$
  $\forall i, j \in \mathcal{N} : \mu_{ij} \cdot x_{ij} = 0$
- stationarity: $\forall i \in \mathcal{N}, \forall j \in \mathcal{N}_i : v_j \cdot \frac{v_i \cdot \varepsilon_i}{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}} - \lambda_j + \mu_{ij} = 0$

From these conditions we can derive the following statements:

1. $\forall i \in \mathcal{N} : \lambda_i > 0 \Rightarrow \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i$

   (due to the complementary slackness of λ_i)

2. $\forall i \in \mathcal{N}, j \in \mathcal{N}_i : \frac{v_j}{\lambda_j} \leq \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{v_i \cdot \varepsilon_i}$

   (due to the stationarity condition & dual feasibility of μ_{ij})

3. $\forall i \in \mathcal{N}, j \in \mathcal{N}_i : x_{ji} > 0 \Rightarrow \mu_{ij} = 0 \Rightarrow v_j/\lambda_j = \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{v_i \cdot \varepsilon_i}$

   (due to the stationarity condition & complementary slackness of μ_{ij})

Given the assumption that for each resource i, $v_i > 0$, we can derive that $\forall i \in \mathcal{N}, \lambda_i > 0$: If agents value a resource i positively and it were free (meaning $\lambda_i = 0$), then each agent would demand an unbounded amount of that resource. This would automatically drive up λ_i. Thus, $\forall i \in \mathcal{N} : \lambda_i > 0$, and therefore statement 1 tells us that $\forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i$. This proves the market clearance condition of the equilibrium 1.4.1.

To prove the other condition of the equilibrium 1.4.1, utility maximisation, we use notation β_i to denote how much value agent i has had to allocate in order to receive back value (often referred to as the inverse best bang-per-buck). This means that $\beta_i = (v_i \cdot \varepsilon_i)/(\sum_{j \in \mathcal{N}_i} x_{ji} \cdot v_j)$. Given condition 2, we know that $v_j/\lambda_j \leq 1/\beta_i$ so $\beta_i \leq \lambda_j/v_j$. Given condition 3 we know if $x_{ji} > 0$ then $\lambda_j/v_j = \beta_i \leq \lambda_k/v_k$. This means that, if agent i receives resources from agent j, then there must have been no other agent k that could give agent i more value for a lower relative cost. This means that agent i is achieving the best possible value-received to value-allocated ratio, thus proving the utility maximisation condition of the equilibrium.
Given that these conditions of the equilibrium as defined in 1.4.1 are satisfied, we can conclude that Program P4 solves the market described in Section 1.4. □

3.2. Adaptation of Shmyrev's program

Another approach to computing the equilibrium in a Fisher market is Program P2, which is closely related to the EG Program P1. Given that we can adapt the EG program to suit the exchange market, we ask if it is possible to do the same with the Shmyrev program. The program below gives an affirmative answer:

$$\begin{aligned} \min_{\mathbf{X}} \quad & \sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \cdot \log \left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \right) - \sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \log v_j \cdot \varepsilon_j \right) \\ \text{s.t.} \quad & \forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} x_{ij} \leq \varepsilon_i, \\ & \forall i, j \in \mathcal{N} : x_{ij} \geq 0 \end{aligned} \tag{P5}$$

The changes, compared to Shmyrev's Program P2, are as follows: We introduce the endowments ε and replace the bids b_{ij} by the value of the allocated resources $v_i \cdot x_{ij}$. Following this, prices p_j are replaced with $\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}$, which represents the total value received by agent j. Similarly to Program P4, Program P5 can be used to compute the equilibrium of the exchange market; we show in Appendix A.1 that Program P5 yields a market equilibrium.

Theorem 3.2.1 (**Optimality of Program** P5). *Program P5 computes the equilibrium allocations $\mathbf{X}^*$ according to equilibrium definition 1.4.1.*

While Program P4 and Program P5 are likely duals of each other (due to the similar relationship between the EG program and Shmyrev's program) and therefore share a close structural relationship, there are still notable differences: Program P5 appears more complex and less intuitive at first glance. It has a less clear connection to the Nash social welfare, and its components lack immediate economic interpretation. Nevertheless, Program P5 plays a crucial role in understanding the dynamics of distributed strategies. In particular, Program P5 is more naturally connected to (proportional) decentralised trading behaviour, as will become clear in Chapter 4. Thus, while its connection to classical economic equilibrium concepts is less direct, its relevance its highlighted in the analysis of distributed processes.

3.3. Adaptation of Jain's program

As established, the AD market is closely related to the exchange market under study here, which allows us to derive central equilibrium programs from existing formulations. In particular, we adapt the convex feasibility Program P3 proposed by Jain [51].

As the defined market has no prices p, as used in Program P3, we recall the sharing ratio ρ from the SE (Definition 2.1.2) as the ratio of an agent's total value received to total value allocated:

$$\rho_i = \frac{\sum_{j \in \mathcal{N}_i} v_j \cdot \overline{x_{ji}}}{v_i \cdot \varepsilon_i} \tag{3.1}$$

According to [12], these sharing ratios ρ and the prices p coincide. We use this insight to factor out the prices from the program. If we introduce neighbourhoods $\mathcal{N}_i$, increase all endowments from 1 to ε_i, factor out the log-terms, and replace values v_{ij} with v_j[1], we obtain the following program:

$$\begin{aligned}
&\forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i \\
&\forall i, j \in \mathcal{N} : x_{ij} \geq 0 \\
&\forall i \in \mathcal{N}, j \in \mathcal{N}_i : \frac{v_j}{\rho_j} \leq \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{\rho_i \cdot \varepsilon_i} \\
&\forall i \in \mathcal{N} : \rho_i > 0
\end{aligned} \tag{P6}$$

Several elements of Program P6 can be recognised from the KKT conditions of Program P4. Specifically, the first rule corresponds to the capacity constraint, while the second and fourth rules reflect primal and dual feasibility, respectively. The third rule captures the stationarity condition from the KKT conditions of Program P4, which, as shown in the proof of Theorem 3.1, supports the utility maximisation argument. We can therefore interpret this third rule as the utility maximisation clause.

Theorem 3.3.1 (**Optimality of Program** P6). *Program P6 computes the equilibrium allocations* $\mathbf{X}^*$ *according to equilibrium definition 1.4.1.*

A formal proof of the optimality of Program P6 for the exchange market under consideration is provided in Appendix A.2. Despite its strong connection to Program P4, Program P6 lacks one important feature: While the Program P4 and Program P5 come with a constructive proof guaranteeing a solution, Program P6 does not. As a result, Program P6 can verify an equilibrium when one exists, however it does not ensure a solution in all cases — unlike Program P4 and Program P5, which always return a result.

3.4. Benchmark Choice

Each of the discussed convex programs can be used to compute a market equilibrium. However, we adopt Program P4 as our method of choice for several reasons. First, Program P4 is intuitive, offering a clear economic interpretation as the maximisation of Nash Social Welfare. Program P5 on the other hand, lacks a straightforward economic interpretation. More importantly, Program

[1] The reasoning behind this reduction will become clear in Chapter 4.

P4 guarantees to always yield a single solution. In contrast, Program P6 may admit multiple solutions, requiring additional justification for how to select among them. By relying on a optimisation program instead of a feasibility program, we avoid the need to impose such arbitrary selection criteria and ensure interpretability and consistency in our equilibrium computations.

The equilibrium solution computed by Program P4 can serve as a reference to gauge the fairness and stability of other (possibly sub-optimal) distributed allocation strategies. This development marks the completion of Contribution 2, by demonstrating how existing equilibrium programs can be adapted to capture key features of exchange markets.

4

Distributed Approaches

Building on the contents of Chapter 3, we continue to address the gap of modelling exchange markets using existing market models; while the central programs found in Chapter 3 provide an exact solution for the equilibrium, they do not explain how such an equilibrium might be reached through distributed exchange. To address this, we now turn our attention to distributed algorithms that can approximate this equilibrium.
Specifically, we explore two commonly used distributed trading strategies: greedy and proportional. Drawing inspiration from established algorithms in related exchange markets, we adapt these ideas to the market model introduced in Section 1.4 and compare the found market states to the equilibrium found by Program P4.

Along the way, we explain how each strategy converges toward the equilibrium, providing novel insights into the connections between distributed strategies and the central programs discussed earlier. These connections clarify why either strategy converges to equilibrium and also address an open question posed by Birnbaum et al. [70].

This chapter thus completes Contribution 3 by providing a comprehensive model of the distributed trading process in the exchange market and sets the stage for Contribution 4 by establishing a foundation for a new mixed-strategy approach.

4.1. Greedy Strategy

A greedy strategy is an allocation strategy for which each agent i maximises the expected utility $\hat{U}_i$ of next timeslot:

$$\mathbf{x}_i(t) = \max_{\mathbf{x}_i}(\hat{U}_i(t+1))$$

Given the allocations of a greedy strategy, the agents can compute a sharing ratio ρ, as defined in equation 3.1. However, considering that this is a distributed strategy that computes allocations over timeslots t, we introduce instead $\rho(t)$.

$$\rho_i(t) = \frac{\sum_{j\in\mathcal{N}_i} v_j \cdot \overline{x_{ji}}(t)}{v_i \cdot \varepsilon_i}$$

In greedy strategy, which we refer to as π, each agent i allocates its resources $D_i(t)$ to the neighbour j with the lowest value $\rho_j(t)$. The reasoning behind this is as follows: Agent i assumes ρ_j to indicate neighbour j's sharing behaviour. As such, when allocating value $v_i \cdot x_{ij}(t)$ to neighbour j,it expects to receive back $(v_i \cdot x_{ij}(t)/\rho_j(t)$ — similar to the rationality clause in the SE. As such, allocating all resources neighbour $j \mid j \in \arg\min_{k\in\mathcal{N}_i} \rho_k$ is expected to return the highest reward. Based on this strategy π, we propose the following algorithm:

Algorithm 1: Algorithm implementing strategy π

```
for t = 1, 2, ... do
    foreach i ∈ 𝒩 do
        Agent i announces to its neighbours the weighted sharing ratio:
        ρ_i(t) = (Σ_{j∈𝒩_i} v_j · \overline{x_{ji}}(t)) / (v_i · ε_i) ;
    end
    foreach i ∈ 𝒩 do
        Agent i distributes its resources D_i(t) evenly over its neighbour j ∈ 𝒩_i having the
          smallest weighted sharing ratio ρ_j(t) in set N_i;
    end
end
```

Strategy π is primarily inspired by the greedy approach proposed by Georgiadis et al. [59], which was shown to converge to equilibrium in the limit , within a stochastic exchange market closely resembling the one described in Section 1.4, differing in its use of valuation functions v and defines equilibrium based on the SE framework. An interesting aspect of allocation strategy π is that it can be derived from the third constraint of Program P6. Intuitively, we argue that, as π is a greedy strategy, it inherently performs utility optimisation. As such, π logically optimises the third constraint of Program P6: The utility maximisation clause. For the interested reader, we refer to Appendix A.3 for a more detailed explanation.

4.2. Proportional Strategy

The proportional allocation strategy we propose considers, for each agent i, the proportional contribution to its utility that it received from each neighbour j:

$$x_{ij}(t+1) = D_i(t+1) \cdot \frac{v_j \cdot x_{ji}(t)}{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}(t)} \tag{4.1}$$

Each agent i returns in the next timeslot $t+1$ the number of resources proportional to how much it has benefited from each neighbour's last allocation, i.e. a proportional response. We will refer to this strategy as ϕ.

Allocation strategy ϕ is based on the Proportional Response Dynamics (PRD) by Wu and Zhang [26]. For the Fisher market and the AD the allocations of PRD are interleaved with bidding rounds. Wu and Zhang's original formulation is particularly suited to our setting, as it does not involve explicit monetary updates. Although ϕ differs from PRD in terms of notation and interpretation, the underlying mechanism remains the same.
Wu and Zhang established a convergence rate of $O(\log(n/\epsilon)/(1-c))$ for PRD, given strictly concave utility functions, in a deterministic market where each resource is valued the same. Wu and Zhang left convergence for linear utilities in exchange markets open. For both the linear Fisher and the AD market a convergence rate of $O(1/t)$ has been proven [70, 69]. However, as our market is neither a Fisher market nor and AD market and does not have strictly concave utility functions, we may not assume either convergence rate.

4.3. Convergence Analysis

We have now identified two strategies that model a way to exchange resources in the exchange market, both of which are derived from strategies that were proven to converge to the equilibrium in closely related markets to the exchange market of this thesis. However, we have not yet shown why π and ϕ may intuitively be expected to converge to equilibrium in this exchange market.

The reasoning behind the convergence of ϕ can be found in a mirror-descent-type algorithm, specifically, a mirror descent over the objective function of Program P5.

Theorem 4.3.1 (Mirror-Descent of ϕ). *ϕ is equivalent to a mirror descent algorithm over Program P5.*

The proof of this theorem is presented in Appendix A.4. Consequently we state that, due to the convergent, optimal nature of the mirror descent algorithm, ϕ converges to the market equilibrium. Additionally, this insight extends the relationship between proportional allocations strategies and Shmyrev's program, identified by Birnbaum et al. [70], to the exchange market, and provides an explanation for why ϕ may be expected to converge to the equilibrium. Moreover, this proof resolves an open problem posed by Birnbaum et al. regarding the existence of convex programs linked to proportional trading via mirror descent in exchange markets. Although left out in this thesis, Theorem 4.3.1 likely also implies a convergence rate similar to that of $O(1/t)$ determined by Birnbaum et al.

The reasoning behind π converging to the equilibrium is somewhat less formal. We visualise the convergence of π next to that of ϕ, to provide more context. Figure 4.1 below demonstrates the convergence process along a generic convex program (representing, for example, Program P5) solving the market at its minimum. The space, which is usually of a high dimension, is heavily simplified for the purpose of clarity.
In Figure 4.1 (b), the mirror descent process of ϕ is shown. In Figure 4.1 (a), the process of update steps of π is shown. Clearly the update steps of π do not approach the exact minimum (equilibrium) in the way ϕ does, instead, it oscillates around the equilibrium. This is a consequence of π only allowing integral allocations, where agents allocate all resources to the same neighbour. In the limit, at $t \to \infty$, the average of these allocations converges to the point in the middle, the minimum. This averaging over time is what makes it possible for the greedy allocations to converge to the optimal solution. Although not a formal proof, this interpretation sheds light on why π might be expected to converge.

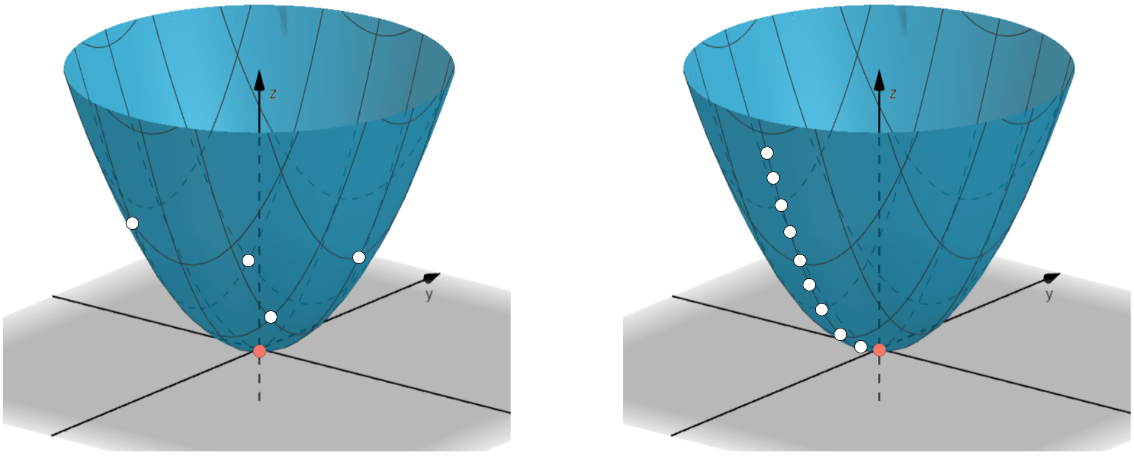

Figure 4.1: Possible allocations in a heavily simplified convex solution space. (a) Allocations chosen by π. (b) Allocations chosen by ϕ over some generic convex program that solves the exchange market.

4.4. Uniqueness & Measuring Optimality

Provided that both π and ϕ converge to the equilibrium, an important difference with Program P4, is that these strategies only approach the equilibrium as t progresses, while Program P4 computes the equilibrium directly. Thus, before some final timeslot T, the exchanges brought on by strategy π or ϕ have not yet brought the market to equilibrium. At any time t, the market's distance to the equilibrium can be measured. We will show how the equilibrium computed by Program P4 can be used to measure that distance, allowing us to characterise the quality of the market state that a distributed strategy converges to.

For a distance measure based on Program P4 to be meaningful, we have to know that, in the limit $t \to \infty$, both π and ϕ converge to the specific equilibrium computed by Program P4, and not just any equilibrium. Program P4 is not unique in equilibrium allocations $\mathbf{X}^*$ as previously established. Therefore, it is not guaranteed that the solutions found by the three different components are equal.
However, although equilibrium allocations $\mathbf{X}^*$ might not be unique, there is another variable in which the exchange market under study *is* unique: equilibrium utilities $\mathbf{U}^*$. Since the market in question uses linear utility function $u_i(\mathbf{r}_i)$ (qualifying as both submodular and supermodular functions), the gross substitutes property holds [117]. Therefore, an increase in the cost for one resource, results in an increase the demand for other resources. Gross substitutability guarantees that the equilibrium utilities $\mathbf{U}^*$ that the agents receive are unique [118]. The insight that linear markets lead to unique utility vectors at equilibrium was already established by Gale [119], who showed that in a linear economy, all agents are indifferent to the different equilibria, indicating that the utility they receive is the same across all equilibria.
Thus, if π or ϕ converge to the market equilibrium, then they also compute the same (unique) vector of utilities $\mathbf{U}$ as Program P4 computes. This means that a distance measure comparing the vectors of received utilities $\mathbf{U}(t)$ to the equilibrium vector $\mathbf{U}^*$ would be a suitable measure to denote how far ϕ and π are removed from equilibrium at time t.

We have concluded that we can use Program P4 to measure the performance of these two strategies. Important to note, however, is that the manner in which we have discussed π and ϕ thusfar, both strategies converge to the equilibrium differently; while π converges when taking the average over all trading iterations $\overline{\mathbf{U}}(T)$, ϕ converges in the last iterate $\mathbf{U}(T)$. Now we will show that ϕ converges in $\overline{\mathbf{U}}(T)$ as well: The Cesaro Mean Theorem[1] states that if some variable a converges to value l in the limit, then the average of that value will converge to the same value. More formally:

$$\lim_{T\to\infty} a_T = l \Rightarrow \lim_{T\to\infty} \frac{a_1 + ... + a_T}{T} = l$$

Thus proving, if in the last iterate $\mathbf{U}(T)$ computed by ϕ converges to $\mathbf{U}^*$, $\overline{\mathbf{U}}(T)$ also converges to $\mathbf{U}^*$. Having ϕ and π converge in the same manner will be helpful later-on when analysing the convergence of a combined strategy. This finding, that ϕ converges on average as well, is supported by the convergence in the base experiments in Section 7.2.

4.5. Common Values

Individual valuations v_{ij} provide more flexibility than common values v_i, and are more suitable to the context of an exchange market, which have been shown to exhibit different preference profiles between agents [120]. However, individual valuations v_{ij} pose a problem for the distributed allocation strategies considered in this thesis.
The challenge in finding such a distributed strategy is as follows: Without a common reference, agents have no common ground to assess the worth of the resources they are exchanging. Assume agent i values a resource differently from agent j. If agent i allocates resources to agent j, agent j might perceive its value differently. When agent j then allocates resources in return, the value of these returns may differ from (1) what agent i finds proportional and (2) what agent i expected to receive in return.
This lack of consensus prevents agents from reliably predicting what they will receive in return and from trading accordingly. Since the trading strategies discussed generally rely on some form of expectation — such as predicted returns or proportional fairness — to guide decisions and support stability, the inability to make reliable predictions becomes a significant issue.

[1] https://proofwiki.org/wiki/Cesaro_Mean

The question then becomes: How does introducing common values v_i resolve this issue? Given common values v_i, we can determine a universal metric of value. All agents agree on the value of resources, meaning that we can express the value of all resources in terms of one type of resource (i.e., normalise the values). Essentially, this is also what the prices p denote in programs such as Program P3; prices can function as a abstract scale by which to measure the value of resources.

Clearly, we must either use common values v_i, or have some external measure of value (i.e., prices) to deal with individual valuations v_{ij}. Barter trade, like the trade in this exchange market, has no prices by definition. Thus we can conclude that, the reduction from v_{ij} to v_j is necessary for agents to trade in this distributed exchange market.
The same reduction from v_{ij} to v_j can be seen in PRD: The original PRD was proposed in an exchange market without prices [26], and featured only common values v_j. Later versions of this allocation strategy are interleaved a price-based bidding round, which enabled them to extend the values to v_{ij}.

Seemingly, bilateral trade (in which trading only happens in pairs) offers the solution for this problem: Agents i and j can measure the value of of each other's resources in the context of each other's valuations v_{ji} and v_{ij}, allowing individual valuations without pricing. However, bilateral trading does not necessarily converge to market equilibrium [121]. This means that distributed strategies that only consider bilateral trading can be sub-optimal.
In theory, bilateral trading can sometimes converge to a market equilibrium, if and only if the transition rate matrix (an adaptation of the equilibrium allocation matrix $\mathbf{X}^*$) is reversible [122]. This insight is based on the connection between exchange markets and (reversible) Markov Chains. Aperjis and Johari [121] explains, more intuitively, that a bilateral allocation only qualifies as an equilibrium if it produces 'supporting prices': a set of prices $\mathbf{p}$ that can form a CE with the found allocation matrix $\mathbf{X}^*$.
Finding a distributed strategy that computes such an allocation matrix (with guaranteed supporting prices) would be valuable, as it would produce a stable and fair solution based on personal valuations v_{ij}, without the use of a universal metric of value. This would provide the freedom to model barter trading with more specific utility functions. Creating such an strategy, however, is a far from trivial task, and will not be solved in this research.

5

Discussion of Theoretical Findings

5.1. Evaluation of Proposed Models

We have discussed variations of the EG Program (P4), Shmyrev's Program (P5), and Jain's Program (P6), all proven to compute the equilibrium. Program P4 was chosen as the benchmark for evaluating convergence to equilibrium due to its economic interpretability (Nash social welfare) and optimisation formulation, which guarantees a solution. Notably, the other programs have close links to the distributed strategies; mirror descent over Program P5 is equivalent to proportional strategy ϕ — resolving an open question by Birnbaum et al. [70] regarding an analogue to Shmyrev's program for resource allocation markets — while greedy strategy π corresponds to the utility-maximising clause of Program P6.

These previously unestablished connections offer valuable insight into the convergence behaviour of ϕ and π, revealing a key distinction between them: While both are ergodically convergent, ϕ also converges in the last iterate. This allows the distance to equilibrium to be estimated using limited information. In contrast, assessing π always requires knowledge of the entire sequence of iterates, making evaluation of π practically unobservable in real-world, infinitely sequential markets.

The convergence of distributed allocation strategies matters to equilibrium — not solely because equilibrium is always the objective in modelling markets — but because it provides a foundational baseline for evaluating the impact of other market aspects (e.g., network topology, asynchronous behaviour). If equilibrium fails to emerge, the deviation stems from added complexities, not the core strategy. Thus, equilibrium serves as a diagnostic tool, allowing platform regulators and researchers to pinpoint inefficiencies by comparing market behaviour against a clear convergence reference point.

Beyond serving as a diagnostic baseline, theoretical equilibria help identify the set of Pareto optimal allocations, which can guide the design and regulation of real-world environments. By narrowing the design space to these efficient outcomes, market design is better equipped to steer exchanges toward desirable allocations [123].

However, while convergence in stylised settings offers analytical value, it is not sufficient for real-world deployment. Many realistic aspects of exchange markets remain overlooked in traditional economic models, limiting their applicability to exchange markets — this includes π and ϕ. Models of distributed exchange must therefore be tailored to specific system objectives, while still retaining some level of abstraction for generalizability. Convergence in simplified settings may provide a baseline, but it does not guarantee optimality in practice.

The proposed central programs illustrate this limitation by failing to account for the dynamic and distributed nature of our exchange market model. The programs require global information of the market, which is rarely available in decentralised exchange markets. Assuming agents in a distributed network to have global information, such as the topology, has lead to inefficiencies in P2P networks [124]. As a result, these approaches face significant privacy challenges, making them impractical for large or sensitive platforms where confidentiality is a concern. Thus, although exchange market model aligns with realistic expectations by restricting agents to local neighbourhood knowledge, the proposed central programs lack in this aspect.

Although the discussed distributed allocation strategies π and ϕ clearly model distributed exchange, they still have limitations regarding their direct applicability as a trading protocol: π and ϕ do not address any complexities outside pure strategies in idealised market settings. Relying solely on these models may result in stylised behaviours that do not reflect real platform dynamics. For an illustration we refer to the case of PropShare: a theoretically optimal, proportional trading protocol for BitTorrent [25], similar to ϕ. This protocol assumed a single, pure principle of trading for all its agents — just like ϕ and π do. However, the protocol exhibited reduced performance due to the unexpected preference of users to sporadically deviate from proportional trading by exploring new allocations [27]. A market model incorporating more diverse trading preferences would likely have revealed equilibrium deviations, highlighting the link between strategy homogeneity and convergence.
Beyond accounting for the strategic diversity of agents, we must also consider the bounded cognitive capacities of agents, as case of PropShare also exemplified that agents were not able to reliably to compute proportional shares of their neighbours contributions [27].
It is thus clear that the theoretical optimality of strategies like π and ϕ carrying over to real-world applications is unlikely. This illustrates the need for protocols to accommodate strategic flexibility and bounded rationality.

5.2. Refining the Models

To better capture the dynamics of agentic exchange markets, several theoretical extensions are proposed. Local feedback mechanisms can substitute global information, drawing on neighbour interactions and historical behaviour rather than objective truths. Introducing asynchronous or stochastic variants of π and ϕ reflects decentralised varied decision-making.
However, incorporating additional complexities into market models inevitably increases their intricacy. As such, we require a balance between flexible strategies and the computational capabilities of agents. Incorporating bounded rationality addresses this challenge, while still allowing allocations to be guided by multiple simple principles (e.g., reciprocity combined with utility maximisation), supporting greater behavioural diversity. The simplicity of such extensions align with Boland [116]'s emphasis on preserving freedom of choice in economic modelling.
It is important to note that, as agent-based markets grow in complexity, explainable mechanism design remains essential [37]. In this context, ϕ and π continue to provide a valuable foundation for protocol design, particularly due to their interpretability and grounding in intuitive trading principles, and do not need to be entirely discarded, but rather adapted.

In conclusion, the distributed strategies π and ϕ examined do not constitute a one-size-fits-all trading mechanism for online P2P platforms. Rather, they provide a solid foundation for modelling distributed trades within exchange markets. Their optimality facilitates the evaluation of how real-world factors affect the market's capacity to reach equilibrium. By developing models that more accurately emulate real-world market phenomena, we can better understand the qualitative effects of various system features. Addressing these identified gaps is essential for developing mechanisms that are both theoretically sound and practically robust in real-world sharing economies.
Crucially, this analysis highlights that technical solutions alone are insufficient to capture the socio-technical complexities at play in algorithmic economics. In line with Sutherland and Jarrahi [44], future models must increasingly account for behavioural constraints and social dynamics if they are to inform effective and equitable market design.

Part II

Numerical Experiments

6

Proposed Model

This chapter addresses the second gap stated in Section 1.2, concerning the lack of a realistic, distributed trading model for the exchange market. To this end, we propose a model grounded in our earlier analysis of π and ϕ in Chapter 4. We introduce a mixed trading strategy, set in an asynchronous market model for agents with limited memory. This corresponds to the key limitations highlighted in Section 2.3, embodying a more realistic and adaptable approach to modelling trades in a decentralised exchange market.

6.1. Mixed Strategy

Mixed strategies can be interpreted in two distinct ways. The first interpretation considers a population composed of agents adhering to pure strategies, as is assumed by Le and Ramazi [100]: for instance, a population in which 40% of agents behave in a purely greedy manner, while the remaining 60% act according to a purely proportional strategy. The second interpretation assumes that each agent individually adopts a mixture of strategies: for example, agent i is 20% greedy and 80% proportional, while agent j might follow a 40% greedy and 60% proportional mix. Although both perspectives yield a population-level mix of strategies, the former lacks behavioural richness: Each agent is still confined to a single pure strategy. In contrast, the second interpretation allows for agents whose decision-making reflects multiple underlying principles, such as greedy utility maximisation and proportional reciprocity. This richer behavioural model captures agents who deal with potentially conflicting objectives or uncertainty in their preferences. In this thesis, we thus adopt the second interpretation.

The proposed mixed strategy ψ consists of the greedy strategy π and the proportional strategy ϕ. Each agent i has a preference for the use of these strategies, as indicated by a weight w_i. A weight of 1 means that an agent is completely greedy and thus follows π, while a weight of 0 means that the agent will follow ϕ. This structure allows the population of trading agents to exhibit mixed behaviour: A population with identical w values is referred to as a homogeneous (hom.) population, while one with varying w values is referred to as a heterogeneous (het.) population. The allocation in each timeslot is:

$$\mathbf{x}_i^{\psi}(t) = w_i \cdot \mathbf{x}_i^{\pi}(t) + (1 - w_i) \cdot \mathbf{x}_i^{\phi}(t) \tag{6.1}$$

where $\mathbf{x}_i^{\phi}$ are the resources i allocates according to ϕ, and $\mathbf{x}_i^{\pi}$ are the resources i allocates according to π. For example, if agent i with $w_i = 0.3$ has greedy allocation vector $\mathbf{x}_i^{\pi}(t) = [2, 3, 1]$ and proportional allocation vector $\mathbf{x}_i^{\pi}(t) = [4, 2, 0]$. Then, at timeslot t, agent i allocates the following resources to its neighbours $\mathbf{x}_i^{\psi}(t) = [0.3 \cdot 2 + 0.7 \cdot 4, 0.3 \cdot 3 + 0.7 \cdot 2, 0.3 \cdot 1 + 0.7 \cdot 0]^\top = [0.4, 2.3, 0.3]^\top$.

Greedy strategy π assumes that agent i receives back a certain amount of value from agent j as a result of allocation $x_{ij}(t)$. In this mixed setting, however, j's expected behaviour has changed. Thus, we might consider that agent i's expectation of returned resources is no longer accurate. However, in Appendix A.6, we show through simple arithmetic that agent j's mixed strategy does not require a change in agent i's expectation of returned resources.

Given the assumed convergence of π and ϕ (as reasoned about in Section 4.3), combined strategy ψ converges as well, as stated in Theorem 6.1.1.

Theorem 6.1.1 (**Weighted Convergence**). *If ϕ and π converge to equilibrium of a market (in the ergodic sense) then the linearly weighted combination of π and ϕ, such as strategy ψ, also converges to equilibrium, provided that strategy weights w are the same across all agents.*

Theorem 6.1.1 is proven in Appendix A.7. This result allows us to distinguish between two sources of 'mixedness': heterogeneity in the population (i.e., agents having different values of w) and mixing at the strategy level (i.e., all agents having homogeneous weight w). Given a linearly mixed rule for which convergence holds in the homogeneous case, any sub-optimal behaviour observed can more confidently be attributed to population heterogeneity.

6.2. Limited Cognition

Bounded rationality refers to the limitations that individuals face in accessing or fully processing all relevant information [125]. Bounded rationality can be modelled by incorporating knowledge limitations and fading-horizon memory (e.g., [82]). Memory effects, which are commonly observed in markets, can improve alignment between models and empirical observations. More specifically, decaying memory provides a simple mechanism for incorporating recency bias — where recent events are more influential than past events — into the market model [126, 127].

To extend the model to these limitations, we implement extensions where (1) agents are not aware of their exact long-term endowments ε_i and must approximate these by recent observations, and (2) the memory of the agents fades away, introducing inaccuracies when agents approximate variables using time-averaged estimates.

Firstly, removing the requirement that each agent knows their own average resource generation ε_i can be done by replacing ε_i with the time-averaged distributed resources $\overline{D_i}(t)$:

$$\overline{D_i}(t) = \frac{\sum_{\tau=1}^{t} D_i(\tau)}{t}$$

Which gives us

$$\rho_i(t) = \frac{\sum_{j\in\mathcal{N}_i} v_j \cdot \overline{x_{ji}}(t)}{v_i \cdot \overline{D_i}(t)} = \frac{\sum_{j\in\mathcal{N}_i} v_j \cdot \overline{x_{ji}}(t)}{\sum_{j\in\mathcal{N}_i} v_i \cdot \overline{x_{ij}}(t)}$$

Additionally, agents experience memory decay. This decay is typically exponential (e.g., [128, 126, 129]). At timeslot t, an array of size t is generated where the entry at each index τ is $e^{(\tau-t)\cdot\gamma}$ ($\Gamma(t) = [e^{-t\cdot\gamma}, ..., e^{-2\cdot\gamma}, e^{-1\cdot\gamma}, e^{0\cdot\gamma}]$). These weights are then applied to some variable, thereby changing the way an agent computes the time-averaged variable. We denote a γ-discounted variable with a subscript $\cdot_\gamma$. For example, for x_{ji} we have:

Regular computation	**γ-discounted computation**
$\overline{x_{ji}}(t) = \frac{\sum_{\tau=1}^{t} x_{ji}(\tau)}{t}$	$\overline{x_{ji}}_{\gamma}(t) = \frac{\sum_{\tau=1}^{t} x_{ji}(\tau) \cdot e^{(\tau-t)\cdot\gamma}}{\sum_{\tau=1}^{t} e^{(\tau-t)\cdot\gamma}}$

As such, the most recent memories have the most influence on the current actions, depending on the factor γ. The lower γ is, the more uniformly the memories are weighted. An agent with

$\gamma = 0$ thus has a perfect memory. A visualisation of the weights in Γ memory decay is presented in Figure 6.1.

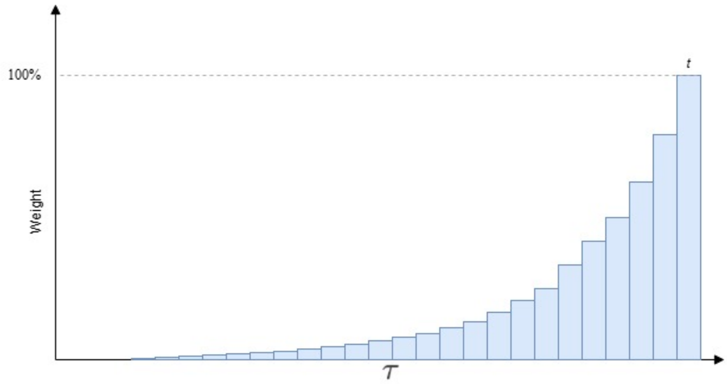

Figure 6.1: Memory weights $^{(\tau-t)\cdot\gamma}$, decaying exponentially

This gives us the following adaptation of the sharing ratio formula:

$$\rho'_{i,\gamma}(t) = \frac{\sum_{j\in\mathcal{N}_i} v_j \cdot \overline{x_{ji}}_\gamma(t)}{\sum_{j\in\mathcal{N}_i} v_i \cdot \overline{x_{ij}}_\gamma(t)}$$

6.3. Asynchronous Dynamics

Following Kolumbus et al. [67], we implement asynchronous update dynamics, where a selected subset of agents update their allocations each timeslot, while the allocations of the remaining agents remain unchanged. This approach reflects markets in which agents may respond less quickly to market changes.
Kolumbus et al. also highlight the open problem of convergence in asynchronous dynamics driven by fully asynchronous information. Similarly, Li et al. [130] emphasise the importance of exploring asynchronous dynamics under partial information. This makes the interaction between asynchronous dynamics and limited information a particularly compelling area for further investigation.

To select which batch of agents is allowed to update their allocations, an adversary chooses the agents that it expects would be the worst-case choice. In Kolumbus et al. [67], the adversary computes this through a potential function that relates closely to the agents' allocation strategy. The authors found that asynchronous PRD converges, regardless of which agents the adversarial chooses, by determining that the potential function increases regardless of batch choice. This potential function is based on the relationship between Shmyrev's program and PRD.
Since no program with similar relationship to the mixed update rule 6.1 has been identified so far, defining a corresponding potential function is less straightforward. Instead a heuristic is used to select agents; for each agent i, a score of potential improvement is computed. This potential Δ_i is measured how close agent i might bring the market to equilibrium: We use Program `P4` to compute a market score based on the market's current state $\mathbf{X}$, then we compute in the same way the market score if only the agent i is allowed to update its allocation $\mathbf{x}_i$. Δ_i is the difference between these two scores, meaning that the higher Δ_i, the closer agent i's allocation has brought the market to the equilibrium. A number of agents with the lowest Δ_i are selected, which are the agents that would improve the market the least. The algorithm is detailed below, where $\mathcal{P}(\cdot)$ denotes Program `P4`.

Algorithm 2: Batch selection process

current market score := $\mathcal{P}(\mathbf{X})$;
foreach $i \in \mathcal{N}$ **do**
    potential market score$_i$:= $\mathcal{P}(\mathbf{X}_{-i}, \mathbf{x}_i)$;
    Δ_i := current market score − potential market score$_i$;
end
selected agent := $\arg\min_{i \in \mathcal{N}} \Delta_i$
batch := selected agents + required agents;
return batch;

This strategy ignores the interaction effects of agents' allocations. However, this is a tractable heuristic, and computing all possible combinations for any batch of size has some tractability concerns. The alternative, namely finding an exact potential function that matches the mixed strategy ψ, is out of scope for this thesis.

6.4. Complete Model

Given mixed strategy ψ for computing the new allocations, the asynchronous mechanism, and a model for memory decay, we propose the following model:

Algorithm 3: Asynchronous market model with mixed strategy ψ and memory decaying agents.

1 **for** $t = 1, 2, \ldots$ **do**
2   **foreach** *agent* $i \in \mathcal{N}$ **do**
3   Update the sharing ratio $\rho'_{i,\gamma}$;
4   **end**
5   Select a batch $\mathcal{B}$ of agents with sub-process 2 ;
6   **foreach** *agent* $i \in \mathcal{B}$ **do**
7   Update $\mathbf{x}_i^{\psi}$;
8   **end**
9   **foreach** *agent* $i \in \mathcal{N}$ **do**
10   Allocate $\mathbf{x}_i^{\psi}$;
11   **end**
12 **end**

6.4.1. Complexity

In each timeslot, the agent updates their sharing ratios (lines 2-4), which takes $O(n)$ time. Next, batch selection is performed (line 5), following Algorithm 2. This step requires computing the planned allocations for each agent to inform decision-making. First, the current market state is computed via matrix operations, which takes constant time. Then, for each agent i, planned allocations are computed by iterating over their neighbours, taking $O(|\mathcal{N}_i|)$ time. In the worst-case scenario (i.e., a complete network) we have $|\mathcal{N}_i| = n - 1$, yielding $O(n)$ time per agent and thus $O(n^2)$ in total. Scoring each planned allocation is a constant-time operation. Sorting the agents by score takes $O(n \log n)$ time, and the subsequent selection and batch creation are again constant time. Therefore, the entire batch selection sub-process takes $O(n^2)$ time. Following this, each agent in the batch $\mathcal{B}$ updates its allocation to the newly updated allocation (lines 6-8). This takes $O(n)$ time in total, with the assumption that the worst-case batch includes all agents in $\mathcal{N}$. Finally, during resource allocation (lines 9-11), each agent loops over their neighbours to distribute resources accordingly. This step also takes $O(n^2)$ time in the worst case. In total, each of the T iterations of Algorithm 3 has a time complexity of $O(n^2)$.

From the agents' perspective, the batch selection process is handled externally. Consequently, their role is limited to computing the sharing ratio and determining subsequent allocations, which reduces their computational burden. This lower complexity supports the interpretation of agents as potentially bounded in terms of computational capacity.

6.4.2. Vulnerability

In the proposed model, agents rely on the information provided by their neighbours (e.g., ρ_j) to compute the new resource allocations. However, making such information public is not always feasible or realistic. In fact, the types of exchange markets described in this thesis often operate under the assumption of private information [131, 132].
Another piece of information missing for the agents in the exchange market is the network topology: The distributed character of the market means that no agent is aware of the entire structure of the market. This lack of information leaves the proposed model vulnerable to two manipulations: misreporting ρ and the Sybil attack.

Although the concept of misreporting ρ is straightforward — agent i reports to its neighbours that it has a sharing ratio of ρ_i' instead of ρ_i — a Sybil attack is somewhat more complicated: Agent i splits into m virtual identities i^m in order to elicit more resources from its neighbours [133]. The utility that agent i obtains is that of all its new virtual identities added. Agent i divides its endowment ε_i over those identities such that $\sum_m \varepsilon_i^m = \varepsilon_i$. Agent i's original neighbours in $\mathcal{N}_i$ see only the Sybil identity i^m and their sharing ratio ρ_i^m, instead of agent i itself. Figure 6.2 illustrates an instance of such a Sybil attack. In this manner, agent i hopes to manipulate its neighbours and receive more resources from them.

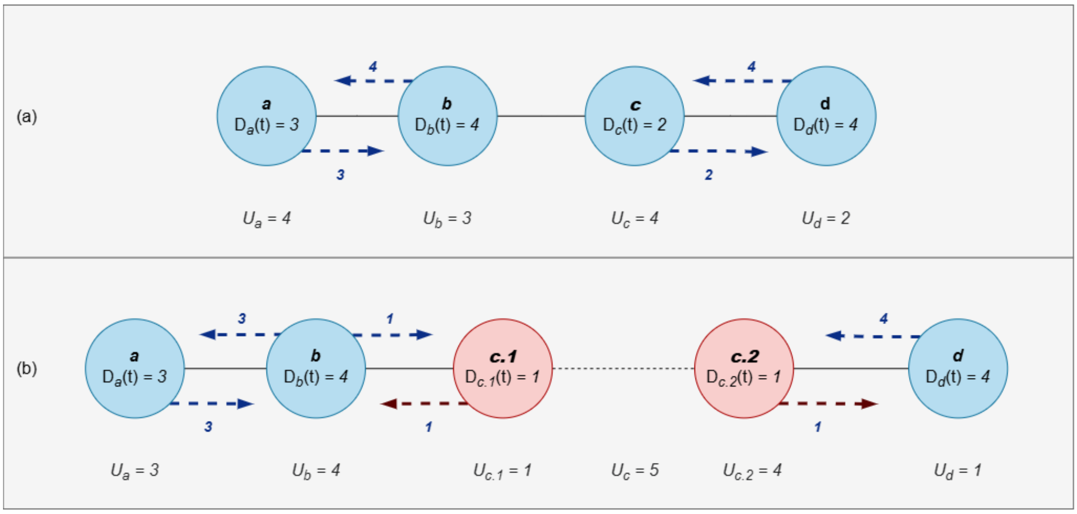

Figure 6.2: Example of a Sybil attack. (a) An initial allocation instance with the utilities of each agent denoted below. (b) An allocation in which agent c employs a Sybil attack and splits itself between two new identities c_1 and c_2, thereby increasing its received utility from 4 to 5.

Both manipulations pose a common threats to exchange markets, largely because they are relatively easy to execute compared to, for example, collusion among multiple agents [134, 108]. They are often considered realistic risks precisely because the distributed nature of exchange markets prevents any single agent from accessing complete global information [135, 136].

7
Experiments

This chapter addresses Research Question 2 through seven numerical experiments on the proposed market model of Chapter 6. The first set of (5) experiments explores how mixed strategies, asynchronous dynamics, and memory decay — individually and in combination — affect convergence properties of the market. We isolate the impact of each aspect and observe their interaction effects as we build toward the full market model 3.

The second set of (2) experiments investigates the model's vulnerability to manipulation, focusing on the allocation rule 6.1 to isolate the effects of non-truthful behaviour from those of asynchrony and memory decay. We examine two key threats: Sybil attacks and misreporting of the sharing ratio. We evaluate whether such manipulation disrupts convergence and to what extent its incentive can be numerically bounded. These experiments deepen our understanding of strategic behaviour and support our broader aim of assessing the performance of trading strategies through realistic market models.

7.1. Experimental Setup

7.1.1. Metrics

Market-level Loss

The market-level loss is defined as the Euclidian distance between the current market state and the equilibrium solution. We express the current market state at timeslot t in the vector of time-averaged received utilities of all agents $\overline{\mathbf{U}}(t)$, as discussed to in Section 4.4. The loss of the market at timeslot t is thus defined as:

$$l^{eq}(t) = \sqrt{\sum_{i=1}^{n} \left(\overline{U_i}(t) - \mathbf{U}_i^*\right)^2}$$

Similar loss functions have been used in previous works (e.g., [81, 77, 69]). To compare two instances of a loss measure, it is important to realise that market instances are randomly generated within some parameter ranges, meaning that the optimal utility vector $\mathbf{U}^*$ might be a different length per instance. As such, we normalise the loss: $l^{eq}_{norm} = l^{eq}/|\mathbf{U}^*|$.

Agent-level Sharing Ratios

In addition to the market-level loss, we take on another metric: the sharing ratios ρ, which can be used to analyse convergence on the agent-level. Although the loss clearly approaches 0 when converging to equilibrium, there is no similar general value for sharing ratios at equilibrium. Depending on the exact market structure, sharing ratios can have different values at equilibrium.

For a complete network, it can be derived from the KKT conditions of Program P4 that $\forall i \in \mathcal{N}$: $\rho_i^* = 1$, provided that $\forall i \in \mathcal{N} : v_i \cdot \varepsilon_i \leq \sum_{j \in \mathcal{N}} v_j \cdot \varepsilon_j$. However, we do not obtain a similar general result for other types of networks.

We refer to plots in which we track each agents sharing ratio as agent-level plots. Each agent-level plot is representative of only a singular network-instance. Shown instances are representative of broader trends[1]. To avoid clutter, we omit the legend in agent-level plots when no statements are made about specific agents.

Incentive Ratio

To measure the effects of manipulative strategies, we measure the IR. We rewrite equation 2.1 to suit the exchange market under consideration. Given unique equilibrium utilities $\mathbf{U}^*$, we know that $u_i(\mathbf{r}_i)$ is the same for all $\mathbf{X}^*$. Provided that equilibrium allocations are computed by Program P4, and thus independent of the agent strategies (s), we write:

$$\zeta_i^M = \max_{s, s_i'} \frac{u_i(\mathbf{r}_{i,(s_{-i}, s_i')})}{U_i^*} \tag{7.1}$$

We collect two IR-based metrics: the maximum observed IR across all network instances, which serves as an estimate of ζ^M, and the average IR across all runs to gauge the expected advantage gained through manipulation. Given the equilibrium sharing ratios $\boldsymbol{\rho}^* = \mathbf{1}$ for complete networks, we can reason about the IR: We know that $\forall i \in \mathcal{N} : \rho_i^* = 1 \rightarrow U_i^* = v_i \cdot \varepsilon_i$. Given that an agent can at most receive all the resources from its neighbours, we can state that in a complete network we have the bound $\zeta_i^M \leq \sum_{j \in \mathcal{N}} v_j \cdot \varepsilon_j / v_i \cdot \varepsilon_i$.

7.1.2. General Parameters

Each experiment iterates over 400 timeslots ($T = 400$), however, some results may show only the first 100 iterations to highlight initial market behaviour.

For each experiment, networks of size n are generated randomly. To account for randomness in the generation process, we generate 30 networks: a rule of thumb for approximating the normal distribution with a limited sample size [137, 138]. To estimate the IR, we increase the number of repetitions to 100: Since we are interested in capturing an extreme value, a high number of repetitions likely improves our chances of coming across the instance that obtains this extreme IR.

Induced by these repetitions is a 95% confidence intervals that illustrates variability of the measured metrics. Since the sharing ratios are specific to a single network instance, instead of an average of multiple instances, no confidence intervals are shown for agent-level graphs.

The choice of underlying network is far from trivial [139]. As such, to capture the diversity of practical instances of exchange markets, we consider different network structures: complete, random, grid, scale-free, or small-world. These network structures reflect patterns observed in real-world exchange markets. For instance, small-world networks with high clustering coefficients are common in internet infrastructure [90], while grid structures resemble energy micro-grids [140]. These are visualised in Figure 7.1 and have the following parameters:

[1] The full set of plots is available at `https://github.com/RixtHellinga/Thesis-Results.git`

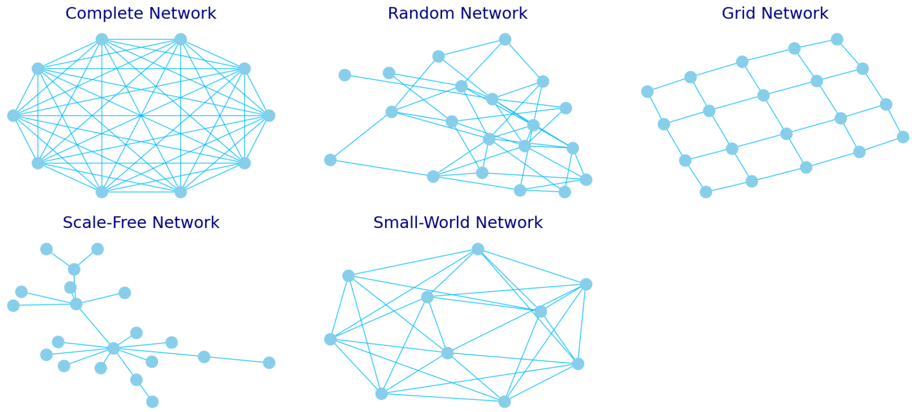

Figure 7.1: The different types of network topologies used in the experiments, visualised in random instances.

	Method	Parameter		Values
complete	-	-	-	-
random network	Erdős–Rényi [141]	probability of edge creation	q	0.2
grid network	-	dimensions	$m \times l$	$\frac{n}{4} \times 4$
scale-free network	Barabasi-Albert [142]	number of edges to attach from a new agent to an existing one	m	1
small-world network	Watts-Strogatz [143]	mean degree	k	$n/2$
		probability to rewire an edge	q	0.05
		clustering coefficient [144]	-	> 1

Table 7.1: Parameter settings for the networks.

A constraint placed on all the networks is that they must be connected, meaning there are no independent sub-networks, and there are no singular loops so agents may not trade with themselves: $i \notin \mathcal{N}_i$. This is because the solution to an allocation problem on a disconnected network is simply the combination of the solutions for each separate component. As such, analysing a disconnected network does not offer any further insight beyond what can be gained from examining its parts individually.

The endowments ε and values v are distributed uniformly over the range of 1 through 5, where the distributable amount of resources each timeslot is sampled with $D_i \sim N(\varepsilon_i, \sigma)$ and $\sigma = 0.1$. Considering that the uniform ranges can be scaled without loss of generality, general conclusions can be drawn even from such a restricted range. These parameters are summarised in Table 7.2.

Parameter		Values
iterations	T	400
endowments	ε	$\{\varepsilon \in \mathbb{Z} \| 1 \leq \varepsilon \leq 5\}$
endowment stochasticity	σ	0.1
resource values	v	$\{v \in \mathbb{Z} \| 1 \leq v \leq 5\}$

Table 7.2: General parameter settings.

7.1.3. Experiment-Specific Parameters

While the general parameters remain the same for all experiments, we vary over other parameters. Specifically, we vary over the heterogeneity of strategies, asynchronous dynamics, level of memory decay, and the presence of a manipulating agent. The parameters are defined as follows:

- **Strategy Heterogeneity**: For parameter w in strategy ψ, multiple settings are used. Pure (non-mixed) strategies π and ϕ fix $w = 1$ and $w = 0$, respectively. For heterogeneous mixed settings (ψ het.), w is sampled uniformly: $w \sim Uniform(0, 1)$. Homogeneous mixed settings (ψ hom.) assign the same non-integer w to all agents.
- **Asynchronous dynamics**: We use three settings: synchronous, asynchronous, and sequential. In the synchronous setting all agents update each timeslot. In the asynchronous setting only a batch selected though process 2 updates each timeslot. In the sequential setting only one agent updates each timeslot.
  The liveliness constraint is set equal to population size n, ensuring feasibility even if in the worst-case where only one agent updates each timeslot.
- **Memory decay**: The default $\gamma = 0$ denotes perfect memory. Values $\gamma = \{0.05, 0.1\}$ model decayed memory. These values are loosely based on Miller's law [2], which states that people can roughly remember the last 7 events. With these values for γ, we guarantee that in the worst case ($\gamma = 0.1$), the half-life of memory is always at least 7.
- **Manipulation - Misreporting** ρ: The manipulation strategy where agents misreport ρ is repeated for increasing values of multiplication factor α_i, where $\alpha_i \cdot \rho_i = \rho'_i$. For the experiments on misreporting ρ we plot the IR against the values of α, and observe if the IR converges as $\alpha \to 0$. We vary of parameter α from 0 to 1 in steps of 0.1: $\alpha \in \{0, 0.1, 0.2, ..., 1\}$.
- **Manipulation - Sybil attack**: The Sybil attack can vary along two dimensions:
  - The first aspect is parameter m: the number of Sybil identities that manipulative agent i creates. It has been shown that the specific case of creating d Sybil identities, where d is the degree of agent i, does not diminish the generality of the obtained IR through the Sybil attack [107]. As such, we employ this setting to represent the Sybil attack, where thus each Sybil identity i^m is connected to only one neighbour j.
  - The second aspect is the distribution of the endowment ε_i into ε_i^m. There seems to be no specific guideline or heuristic to make this division. To tackle this information gap, we employ three different endowment-distribution settings:
    - **Proportional** distribution assigns endowment to Sybil identity i^m proportionally to the value of its neighbour j's endowment: $\mathcal{N}_{i^m} = \{j\} \Rightarrow \varepsilon_i^m = \varepsilon_i \cdot (v_j \cdot \varepsilon_j)/(\sum_{k \in \mathcal{N}_i} v_k \cdot \varepsilon_k)$.
    - **Even** distribution simply divides agent i's endowment equally over all its Sybil identities.
    - **Random** distribution divides the endowment randomly over all the Sybil identities.

    Naturally, for each of the three division types we have $\sum_m \varepsilon_i^m = \varepsilon_i$.

For each experiment, we denote separately which exact configuration is tested and with which purpose.

7.1.4. Supporting Software

For these experiments, most of the implementation was carried out in Python. To generate the networks the package `networkx` was used. To optimise Program `P4`, the Python convex optimisation package `cvxpy`[3] was used, as the Program's convex structure aligns well with the Disciplined

[2] `https://en.wikipedia.org/wiki/The_Magical_Number_Seven%2C_Plus_or_Minus_Two`
[3] `https://www.cvxpy.org/`

Convex Programming (DCP) framework that `cvxpy` follows. Alongside this, the MOSEK solver[4] was employed with the interior-point method. Although the MOSEK solver also allows the use of the simplex method for some problems, which can leverage a warm start, the simplex method may incur exponential runtime in the worst case.

7.2. Experiment 1: Pure Strategies

Focus and Configuration

The first experiment does not relate closely to the model from Chapter 6, rather it is meant to support the claims made in Chapter 4 concerning the convergence of π and ϕ. This experiment examines the convergence behaviour of both ϕ and π in different network types, where the agents operate synchronously, and have their full memory. Since the effect of different network topologies on a greedy strategy similar to π has already been studied [145], we focus instead on comparing ϕ and π within a given network structure, rather than analysing their convergence separately across topologies. The exact configuration of w, the update timing, and γ are shown in Table 7.3.

Network types	n	Reps.
{Complete, Random, Grid, Scale-free, Small-world}	$\{8, 20\}$	30

((a)) Network type, population size, and number of repetitions for Experiment 1.

Allocation Strategy	w	Update Timing	γ	Manipulation
π	1	Sync.	0	None
ϕ	0	Sync.	0	None

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 1.

Table 7.3: Configuration of Experiment 1.

.

Results

Figure 7.2 presents the convergence on market-level given strategies π (labelled Π) and ϕ (labelled Φ). These graphs suggests that both strategies approach the equilibrium, as the loss converges towards 0. In 7.4 the loss for a network of 20 nodes is 0.04376 and 0.00642 after 100 iterations, for π and ϕ, respectively. After 400 iterations the loss is 0.01005 and 0.00161. The loss thus decreases notably over the number of iterations, presumably converging to 0 in the limit $t \to \infty$.

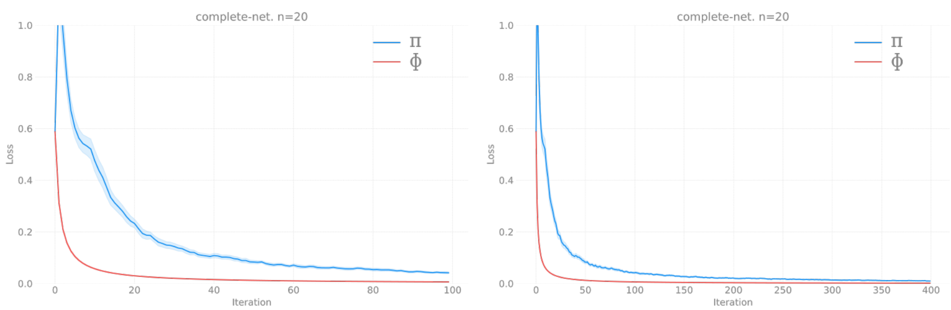

(a) $T = 100$ (b) $T = 400$

Figure 7.2: Market-level converge of loss over time for π (Π) and ϕ (Φ) in a complete network of size $n = 20$ for $T = \{100, 400\}$ iterations.

[4] `https://docs.mosek.com/latest/pythonapi/index.html`

		π	ϕ
$n = 8$	$T = 100$	0.01877	0.00811
	$T = 400$	0.00511	0.00203
$n = 20$	$T = 100$	0.04376	0.00642
	$T = 400$	0.01005	0.00161

Table 7.4: Loss values for π and ϕ in a complete network of size $n = \{8, 20\}$ after $T = \{100, 400\}$ iterations.

Figure 7.2(a) depicts the loss over $T = 100$ and indicates that ϕ converges towards 0 faster and in a smoother manner than π, showing fewer irregularities. This difference in smoothness is visible as well on the agent-level: The results in Figure 7.3 visualise the sharing ratio ρ for the agents in an instance of a complete network of size 8, where the sharing ratios converge in a jagged manner for π and in a smooth manner for ϕ. In both instances, the sharing ratios seemingly converge to a value of 1, indicating a convergence to equilibrium in case of a complete network, as explained in 7.1.1.

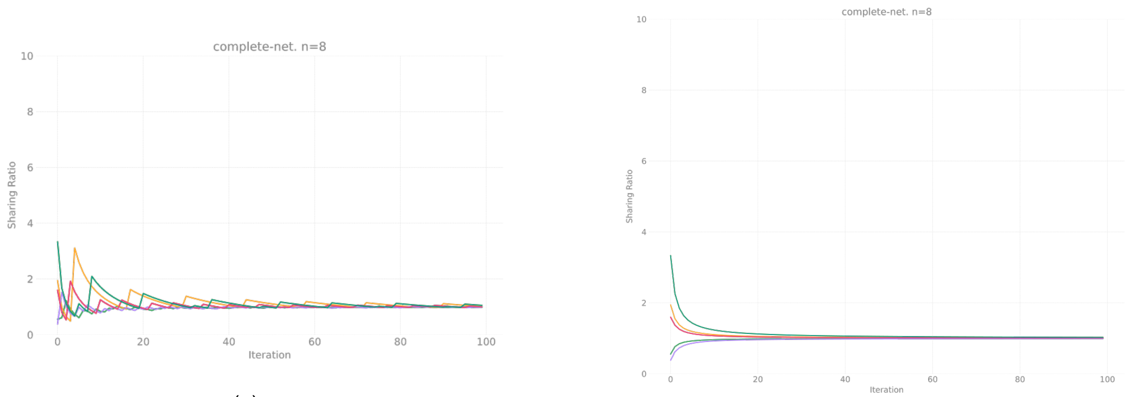

(a) π

(b) ϕ

Figure 7.3: Agent-level convergence of sharing ratios ρ for π and ϕ in a complete network of size $n = 8$ at $T = 100$ iterations. Each agent's sharing ratio is shown as a separate line.

Figure 7.4 shows the convergence of ϕ (labelled Φ) and π (labelled Π) for a population size of $n = 20$ in different network types. In all settings, both ϕ and π approach a loss of 0, suggesting convergence to the equilibrium. This difference in smoothness of convergence and convergence rate diminishes in sparser networks. Especially the results in the scale-free network show that ϕ and π barely differ in terms of convergence speed and irregularities. The results in Appendix B.1 confirm this.[5]

[5]To avoid redundancy, the graphs of Experiment 1 have been merged with the graphs of Experiment 2 in Appendix B.1. For each of the network types we show plots containing all settings of w, which conveys the results relating to both the Experiment 1 as well as the Experiment 2.

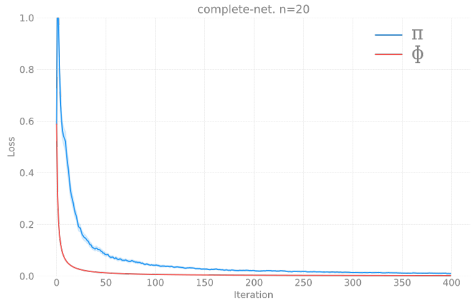

(a) Complete network

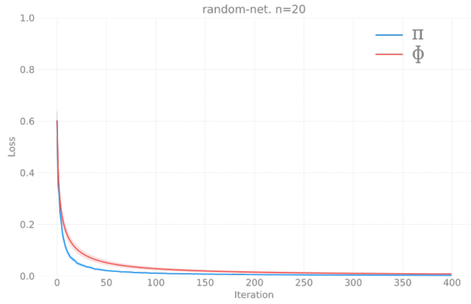

(b) Random network

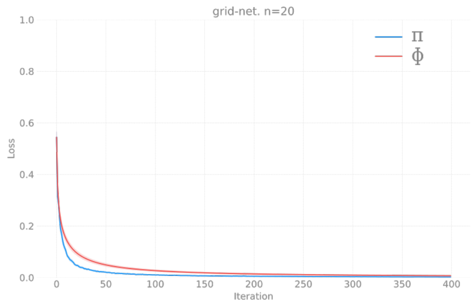

(c) Grid network

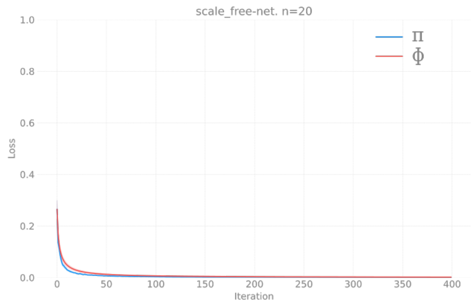

(d) Scale-free network

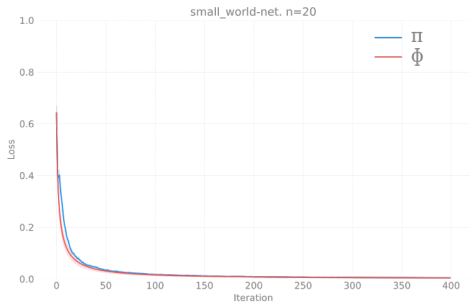

(e) Small-world network

Figure 7.4: Market-level convergence of loss over time for π (Π) and ϕ (Φ) in various network topologies of size $n = 20$ for $T = 400$ iterations.

		Complete	Random	Grid	Scale-free	Small-world
π	$T = 100$	0.04161	0.01098	0.01069	0.00332	0.01725
	$T = 400$	0.00991	0.00282	0.00268	0.00088	0.00456
ϕ	$T = 100$	0.00630	0.02861	0.02716	0.00667	0.01581
	$T = 400$	0.00158	0.00780	0.00725	0.00180	0.00396

Table 7.5: Loss values for π and ϕ in various network topologies of size $n = 20$ after $T = \{100, 400\}$ iterations.

7.3. Experiment 2: Mixed Strategies

Focus and Configuration

The second experiment focuses on the effects of mixed trading strategies, both homogeneous and heterogeneous, where agents have a greedy, proportional, or mixed strategy, determined by weight w. The purpose of this experiment is to investigate how heterogeneity in agent strategies influences both the convergence process and the quality of the resulting market state, specifically how closely the market converges to the equilibrium.

Network types	n	Reps.
{Complete, Random, Grid, Scale-free, Small-world}	$\{8, 20\}$	30

((a)) Network type, population size, and number of repetitions for Experiment 2.

Allocation Strategy	w	Update Timing	γ	Manipulation
π	1	Sync.	0	None
ϕ	0	Sync.	0	None
ψ (hom.)	0.3	Sync.	0	None
ψ (hom.)	0.7	Sync.	0	None
ψ (het.)	$[0, 1]$	Sync.	0	None

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 2.

Table 7.6: Configuration of Experiment 2.

Results

Figure 7.5 depicts how the market converges to equilibrium in the complete network setting, under different settings of w. Homogeneously mixed populations have values $w = 0.3$ and $w = 0.7$ (labelled ψ *hom. w*=0.3 and ψ *hom. w*=0.7, respectively), and the agents in heterogeneously mixed populations may have any $w \in [0, 1]$ (labelled ψ *het.*).
Homogeneous mixed strategies converge at speeds between ϕ and π, with $w = 0.3$ behaving more like ϕ, and $w = 0.7$ resembling π — as expected from the allocation rule's linear structure.

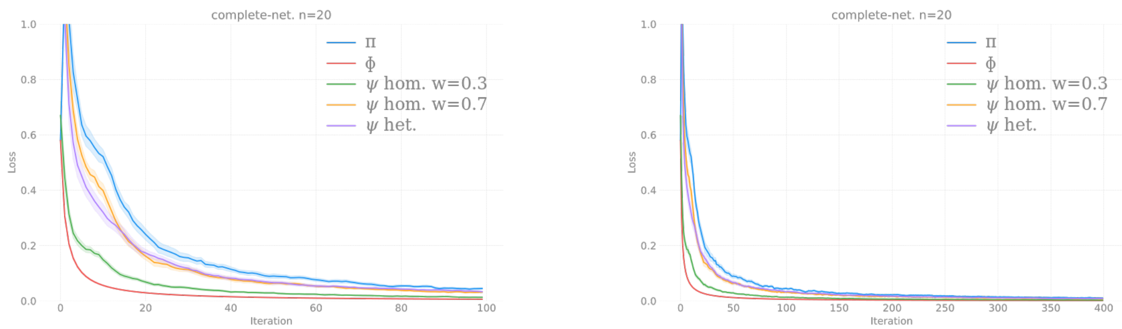

(a) $n = 20$, $T = 100$ (b) $n = 20$, $T = 400$

Figure 7.5: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a complete network of size $n = 20$ for $T = \{100, 400\}$ iterations.

	π	ϕ	ψ hom. $w = 0.3$	ψ hom. $w = 0.7$	ψ het.
$T = 100$	0.04484	0.00618	0.01307	0.03189	0.03327
$T = 400$	0.01040	0.00155	0.00308	0.00732	0.00872

Table 7.7: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a complete network of size $n = 20$ after $T = \{100, 400\}$ iterations.

Although the homogeneously mixed strategies seem to converge to equilibrium across all network types, just as ϕ and π, as can be seen in Appendix B.1, this is not necessarily the case for the heterogeneously mixed strategy. Figure 7.5 shows that, in the complete network, the heterogeneously mixed strategy behaves similarly to the homogeneous variants. However, the results for the heterogeneously mixed strategy in the random network and grid network, as shown in Figures 7.6 and 7.7, show how the market converges to a higher loss, indicating a larger distance to the equilibrium; the heterogeneous mixed strategy ψ seemingly converges just as quickly in initial iterations, but stagnates earlier on in the process. This results in a converged market state further removed from the equilibrium, largely noticeable in the random and grid network. Visually, this might be difficult to notice in the provided graphs, however the respective scores in Tables 7.8 and 7.9 more clearly demonstrate this effect: Between $T = 100$ and $T = 400$ ψ (het.) has improved less than the others, with a mere $8,8\%$ improvement, compared to the average improvement of $74,5\%$ of all other settings in the grid network. The same effect, although less extreme, can be seen for the scale-free and small-world networks in Appendix B.1.

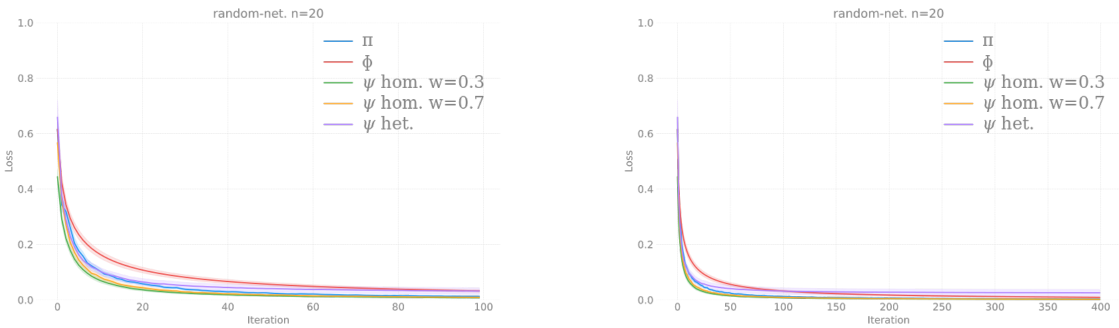

(a) $n = 20, T = 100$ (b) $n = 20, T = 400$

Figure 7.6: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a random network of size $n = 20$ for $T = \{100, 400\}$ iterations.

	π	ϕ	ψ hom. $w = 0.3$	ψ hom. $w = 0.7$	ψ het.
$T = 100$	0.01260	0.03219	0.00749	0.00920	0.03188
$T = 400$	0.00322	0.00918	0.00185	0.00238	0.02554

Table 7.8: Loss values for π, ϕ, and ψ in a random network of size $n = 20$ after $T = \{100, 400\}$ iterations.

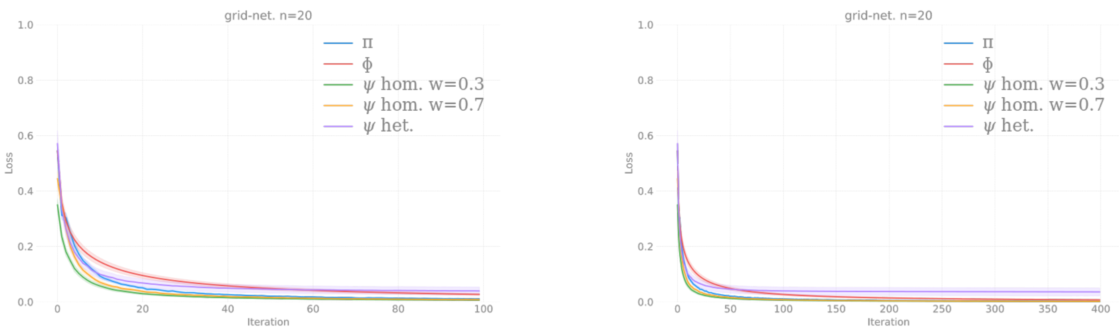

(a) $n = 20, T = 100$ (b) $n = 20, T = 400$

Figure 7.7: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a grid network of size $n = 20$ for $T = \{100, 400\}$ iterations.

	π	ϕ	ψ hom. $w = 0.3$	ψ hom. $w = 0.7$	ψ het.
$T = 100$	0.01094	0.02697	0.00626	0.00819	0.03951
$T = 400$	0.00267	0.00732	0.00153	0.00213	0.03603

Table 7.9: Loss values for π, ϕ, and ψ in a grid network of size $n = 20$ after $T = \{100, 400\}$ iterations.

7.4. Experiment 3: Asynchronous Dynamics

Focus and Configuration

The third experiment centres around the asynchronous dynamics, modelling more realistically distributed systems where agents act at different speeds. This experiment explores how asynchronous dynamics impact market behaviour, including whether convergence occurs, the emergence inequalities, the sensitivity to update patterns, and how convergence speed and efficiency compare to the synchronous case. The exact configuration of w, the update timing, and γ are shown in Table 7.10.

Network types	**n**	**Reps.**
Complete	$\{8, 20\}$	30

((a)) Network type, population size, and number of repetitions for Experiment 3.

Allocation Strategy	**w**	**Update Timing**	γ	**Manipulation**
π	1	Sync.	0	None
π	1	Async.	0	None
π	1	Seq.	0	None
ϕ	0	Sync.	0	None
ϕ	0	Async.	0	None
ϕ	0	Seq.	0	None

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 3.

Table 7.10: Configuration of Experiment 3.

Results

The market-level convergence graphs in Figure 7.8 depict the synchronous and asynchronous versions of π and ϕ (labelled Π, Π *async*, Π *seq*, Φ, Φ *async*, and Φ *seq*) in a complete network of size $n = 20$. Seemingly, at 400 iterations, each of the strategies has converged to a loss close to 0, indicating convergence to equilibrium. Although the asynchronous dynamics do not seem to strongly affect the quality of the converged market state, they do affect the rate of convergence; the asynchronous and sequential versions converge more slowly than their synchronous counterparts. With π, we also see this slowdown also in terms of a more smoothed out market-level response.

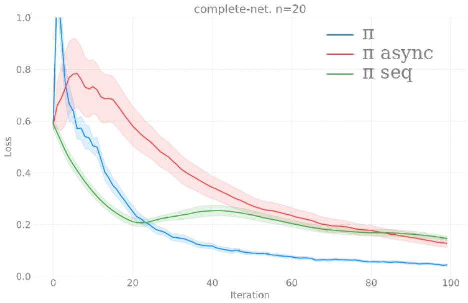

(a) π, $T = 100$

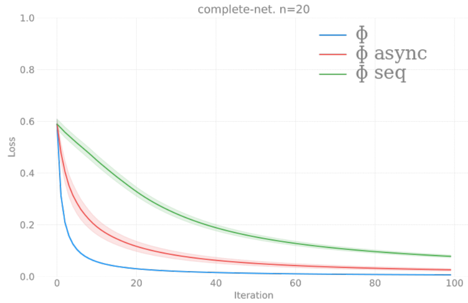

(b) ϕ, $T = 100$

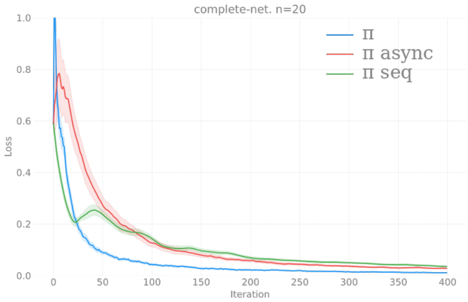

(a) π, $T = 400$

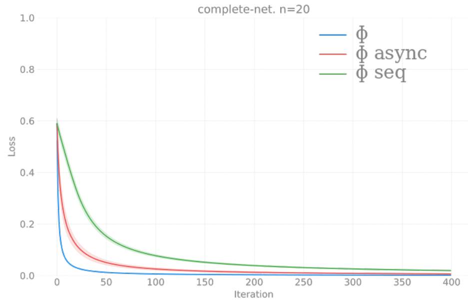

(b) ϕ, $T = 400$

Figure 7.8: Market-level convergence of loss over time for π and ϕ under various update types: synchronous (Π & ϕ), asynchronous (Π *async* & ϕ *async*), and sequential (Π *seq* & ϕ *seq*) in a complete network of size $n = 20$ for $T = \{100, 400\}$ iterations.

	π sync.	π async.	π seq.	ϕ sync.	ϕ async.	ϕ seq.
$T = 100$	0.04383	0.12725	0.14620	0.00628	0.02577	0.07788
$T = 400$	0.01108	0.02798	0.03519	0.00157	0.00645	0.01947

Table 7.11: Loss values for π and ϕ with various update types in a complete network of size $n = 20$ after $T = \{100, 400\}$ iterations.

On the agent-level, asynchronous dynamics show a similarly interesting effect: the convergence of sharing ratios slows down, as seen in Figure 7.9. For π, Figure 7.9 (a) and 7.9 (c) show a smoothing effect, where the typically jagged progression of sharing ratios becomes This suggests that the agents respond more intensely under asynchronous conditions. Such effects do not arise for ϕ, where the convergence of sharing ratios is already completely smooth.

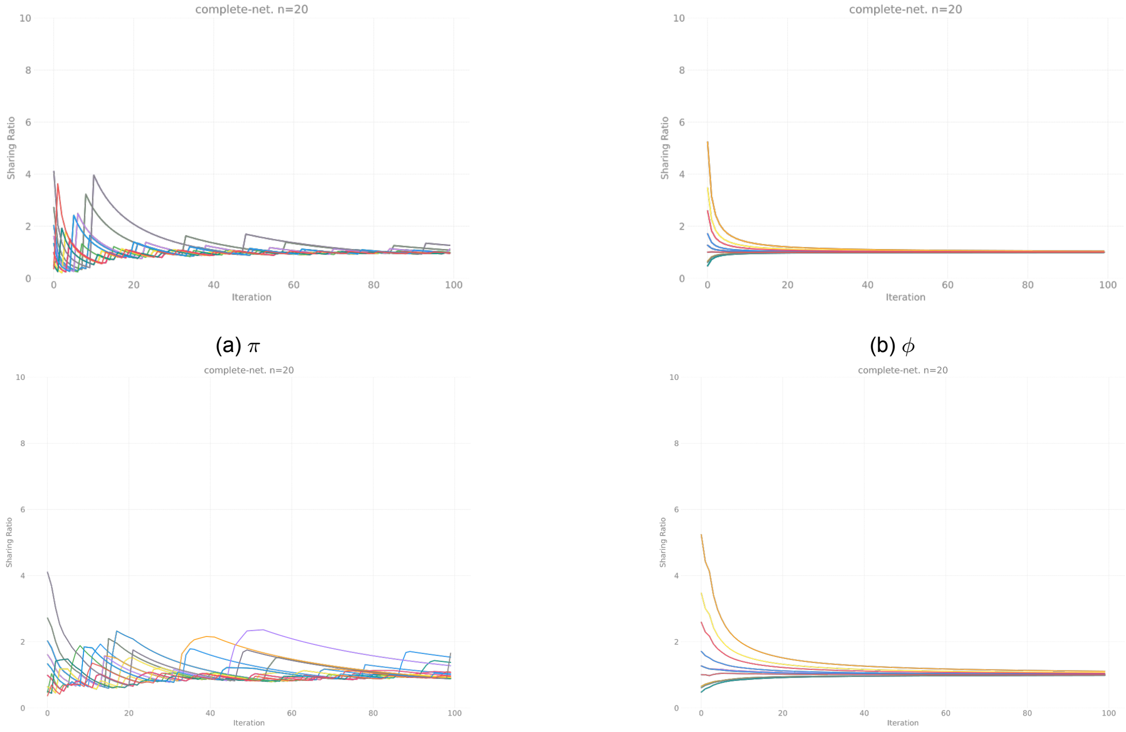

(a) π (b) ϕ

(c) asynchronous π (d) asynchronous ϕ

Figure 7.9: Agent-level convergence of sharing ratios ρ for π and ϕ for synchronous and asynchronous update types in a complete network of size $n = 20$ at $T = 100$ iterations. Each agent's sharing ratio is shown as a separate line.

7.5. Experiment 4: Memory Decay

Focus and Configuration

In this experiment, agents gradually forget experiences, according to decay parameter γ. This experiment addresses how limiting agents' memory of past payoffs affects convergence; whether different trading strategies are equally robust to memory decay; and if memory decay in only a few agents can also affect convergence of the entire market.

Memory decay parameter gamma varies across three values, depicting no memory decay ($\gamma = 0$), some memory decay ($\gamma = 0.05$), strong memory decay ($\gamma = 0.1$). The value for γ is fixed for all iterations. In some variations, only a subset of the population (10%) will be affected by decay, while in others the entire population will be affected. The aim is to explore both global and local effects of memory decay, and to gauge the vulnerability of the market to imperfect information. The exact configuration of w, the update timing, and γ are shown in Table 7.12.

Network types	n	Reps.
Complete	$\{8, 20\}$	30

((a)) Network type, population size, and number of repetitions for Experiment 4.

Allocation Strategy	w	Update Timing	γ	Manipulation
π	1	Sync.	0	None
π	1	Sync.	0.05	None
π	1	Sync.	0.1	None
π	1	Sync.	$0.1, 10\%$	None
ϕ	0	Sync.	0	None
ϕ	0	Sync.	0.05	None
ϕ	0	Sync.	0.1	None
ϕ	0	Sync.	$0.1, 10\%$	None

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 4.

Table 7.12: Configuration of Experiment 4.

Results

As can be seen in Figure 7.10, memory decay is only influential in the cases where the agents exhibit (partial) greediness. This is logical, considering that ϕ already does not need any information on previous trades, with the exception of the trades in the previous timeslot.
Additionally, the higher parameter γ is, the more extreme the effect of the memory decay is on a (partially) greedy strategy. A value of $\gamma = 0.1$ (labelled Π *y*=0.1) converges to a loss of 0.20259, which is further removed from equilibrium than the settings of $\gamma = 0$ (labelled Π) and $\gamma = 0.05$ (labelled Π *y*=0.05), which converge to losses of 0.01117 and 0.12156, respectively.

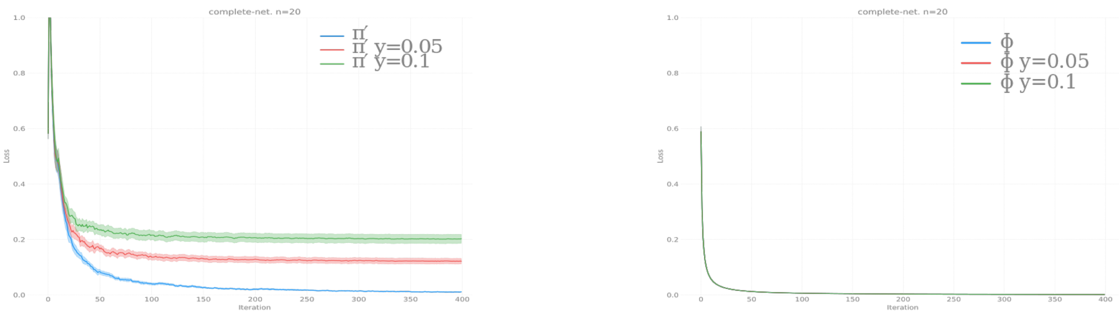

(a) π

(b) ϕ

Figure 7.10: Market-level convergence of loss over time for π and ϕ under various memory decay levels: $\gamma = 0$ (Π & Φ), $\gamma = 0.05$ (Π *y*=0.05 & Φ *y*=0.05), and $\gamma = 0.1$ (Π *y*=0.1 & Φ *y*=0.1) in a complete network of size $n = 20$ for $T = 400$ iterations.

	π $\gamma = 0$	π $\gamma = 0.05$	π $\gamma = 0.1$	ϕ $\gamma = 0$	ϕ $\gamma = 0.05$	ϕ $\gamma = 0.1$
$T = 100$	0.04084	0.13987	0.21393	0.00628	0.00628	0.00628
$T = 400$	0.01117	0.12156	0.20259	0.00157	0.00157	0.00157

Table 7.13: Loss values for π and ϕ with various levels of memory decay in a complete network of size $n = 20$ after $T = \{100, 400\}$ iterations.

On the agent-level, in Figure 7.11, a similar effect can be seen. Sharing ratios tend to oscillate more extremely when agents have shorter memory, especially in greedier strategies. As a result, we see that the convergence of a population suffering from memory decay, in Figure 7.11 (b), is more volatile than the other population, in Figure 7.11 (a), in which agents have perfect memory. Although the oscillation does seem to decrease over time, the sharing ratios converge to many values different from 1. As such, we confirm that, in a complete network, the market converges to a state removed further from the equilibrium due to memory decay.

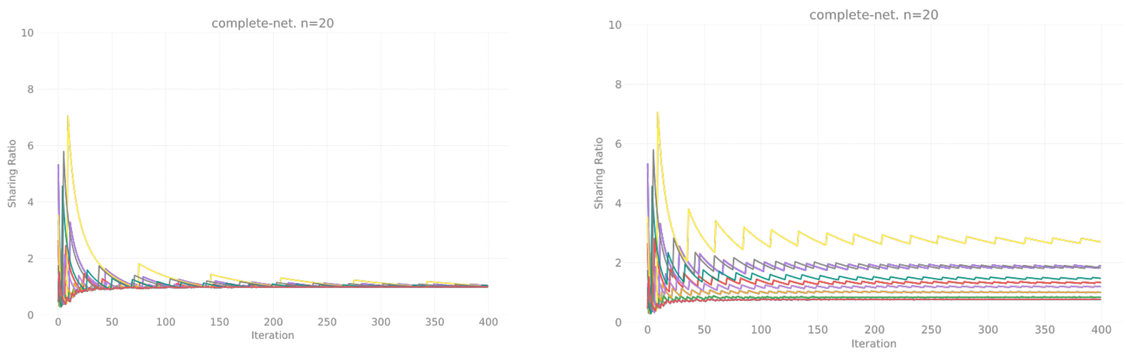

(a) $\gamma = 0$ (b) $\gamma = 0.1$

Figure 7.11: Agent-level convergence of sharing ratios ρ for π under perfect memory ($\gamma = 0$) and memory decay ($\gamma = 0.1$) in a complete network of size $n = 20$ at $T = 400$ iterations. Each agent's sharing ratio is shown as a separate line.

The market-level convergence in which 10% of the population experiences memory decay is illustrated in Figure 7.12. It is evident that, when only a subset of the population is affected by memory decay (labelled as Π *y*=0.1 10%), π converges to a higher loss compared to a greedy population with perfect memory (labelled as Π).

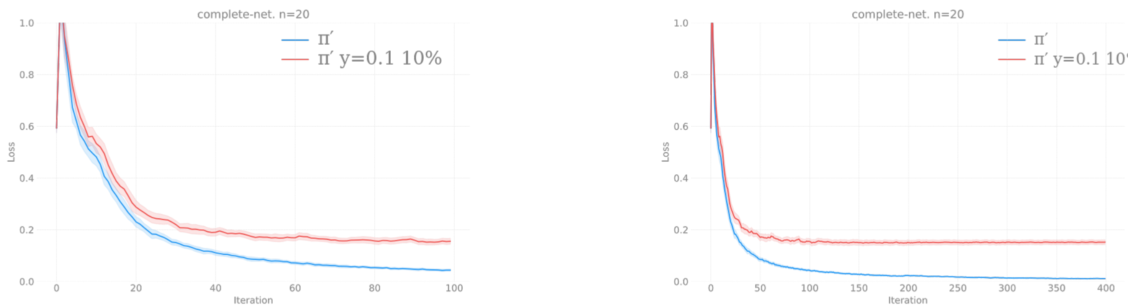

(a) $T = 100$ (b) $T = 400$

Figure 7.12: Market-level convergence of loss over time for π under perfect memory $\gamma = 0$ (Π) and memory decay $\gamma = 0.1$ for 10% of the population (Π *y*=0.1 10%) in a complete network of size $n = 20$ for $T = \{100, 400\}$ iterations.

	$\gamma = 0$	10% $\gamma = 0.1$
$T = 100$	0.04394	0.15538
$T = 400$	0.01122	0.15212

Table 7.14: Loss values for π under perfect memory $\gamma = 0$ and memory decay $\gamma = 0.1$ for 10% of the population in a complete network of size $n = 20$ after $T = \{100, 400\}$ iterations.

At the agent-level, the results in Figure 7.13 (b) depict the scenario in which all agents have the full memory and the scenario in which 10% of agents experience memory decay. Similar to the scenario with complete memory decay, partial decay leads to sharing ratios that do not all converge to the equilibrium sharing ratios. However, in this case, the deviations are primarily limited to the agents affected by memory decay — specifically, the agents that consistently reach high sharing ratios: the dashed lines in Figure 7.13 (b).

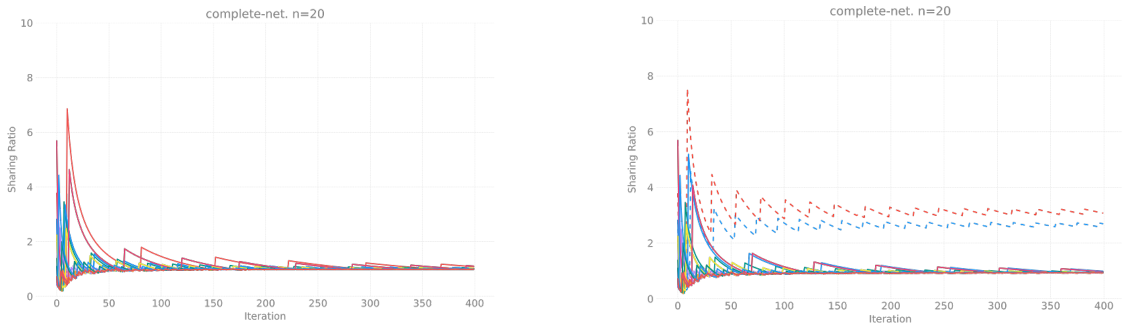

(a) $\gamma = 0$ (b) $\gamma = 0.1$ for 10%

Figure 7.13: Agent-level convergence of sharing ratios ρ for π under perfect memory ($\gamma = 0$) and memory decay ($\gamma = 0.1$) for 10% of the population in a complete network of size $n = 20$ at $T = 400$ iterations. Each agent's sharing ratio is shown as a separate line.

7.6. Experiment 5: Complete Model

Focus and Configuration

This experiment evaluates the complete model as described in Algorithm 3, with heterogeneous strategies, asynchronous dynamics, and memory decay. This experiment explores how asynchronous dynamics, strategy heterogeneity, and bounded memory interact to influence convergence: whether they stabilize, destabilise, or introduce new behaviours, and whether their combination leads to complex market dynamics resembling real-world exchange markets. The exact configuration of w, the update timing, and γ are shown in Table 7.15.

Network types	$\mathbf{n}$	**Reps.**
Complete	$\{8, 20\}$	30

((a)) Network type, population size, and number of repetitions for Experiment 5.

Allocation Strategy	$\mathbf{w}$	**Update Timing**	γ	**Manipulation**
ψ (het.)	$[0, 1]$	Sync.	0	None
ψ (het.)	$[0, 1]$	Async.	0	None
ψ (het.)	$[0, 1]$	Sync.	0.1	None
ψ (het.)	$[0, 1]$	Async.	0.1	None

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 5.

Table 7.15: Configuration of Experiment 5.

Results

Figure 7.14 shows the effects of each of the previously introduced components on the convergence of the market: heterogeneous mixed strategies, asynchronous dynamics, and memory decay. We see the heterogeneous mixed strategy (labelled ψ *het.*), alongside versions with only asynchronous updates (labelled ψ *het. async*), only memory decay (labelled ψ *het. y*=0.1), and the full Model 3 (labelled ψ *het. y=0.1 async*), which includes heterogeneous mixed strategies, asynchronous dynamics, and memory decay. Figure 7.15 puts this full model in direct comparison with the simple and idealised settings for π (labelled Π) and ϕ (labelled Φ).

Figure 7.14 clearly shows that each component affects convergence in distinct ways. Asynchronous dynamics mainly slow down the convergence, which aligns with the findings in Section 7.6. Memory decay, on the other hand, affects the final convergence point — causing the loss to stagnate at a higher value, indicating that the market ends up further from the equilibrium. The combination of effects in ψ *het. y*$= 0.1$ *async* seems to be additive: The market converges more slowly due to the asynchronous dynamics and to a higher loss due to the memory decay.

Figure 7.15 reinforces this: While π and ϕ both converge close to equilibrium, the market full market model clearly stagnates at a much higher loss. The distance to equilibrium seems to be mostly determined by the memory component, while asynchronous updates merely stretch out the timeline and heterogeneous strategies account for only a slight deviation from the equilibrium. Overall, although the full market model 3 still converges in a general sense, it converges to a suboptimal market state.

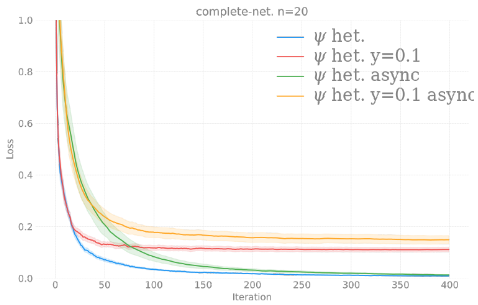

Figure 7.14: Market-level convergence of the heterogeneous mixed strategy (ψ *het.*), a heterogeneous mixed strategy with asynchronous dynamics (ψ *het. async*), a heterogeneous mixed strategy with memory decay (ψ *het. y*$=0.1$), and full model 3 (ψ *het. y*$=0.1$ *async*) in a complete network of size $n = 20$ for $T = 400$ iterations.

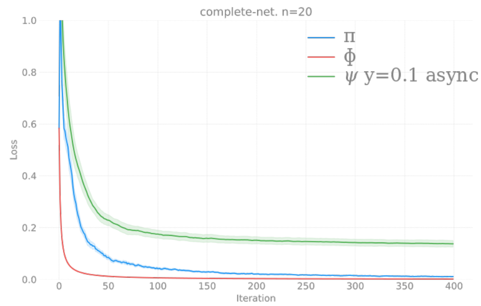

Figure 7.15: Market-level convergence of π (Π), ϕ (Φ), and full model 3 (ψ *y*$= 0.1$ *async*) in a complete network of size $n = 20$ for $T = 400$ iterations.

	π	ϕ	ψ het.	ψ het. y=0.1	ψ het. async.	ψ het. async. y=0.1
$T = 100$	0.04411	0.00622	0.03513	0.11741	0.08429	0.17708
$T = 400$	0.01135	0.00156	0.00928	0.11148	0.01303	0.14337

Table 7.16: Loss values for π, ϕ, and ψ with various combinations of memory decay and asynchronous dynamics in a complete network of size $n = 20$ after $T = \{100, 400\}$ iterations.

On the agent-level, we see the combined effects of asynchronous dynamics and memory decay playing out more dramatically. The result is a rather volatile and chaotic pattern in how agents adjust their sharing ratios over time. In Figure 7.16, we can pick out some familiar patterns from earlier experiments. The asynchronous updates are clearly at work: Agent responses appear rounded off, yet also somewhat more exaggerated in amplitude. At the same time, we see the effects of heterogeneity and memory decay, as the sharing ratios seemingly converge towards suboptimal values. That said, the visual complexity of these agent-level plots makes it difficult to isolate the effects. The interactions between heterogeneity, asynchrony, and memory decay result in noisy dynamics that mask the individual contributions of each component.
Because of this, the clearest takeaways about the roles of memory and timing likely come from the more aggregated market-level results, or from the earlier experiments, which analysed each

component in isolation. The behaviour the agents under model 3 reinforces what we have already seen, however in a significantly more chaotic manner.

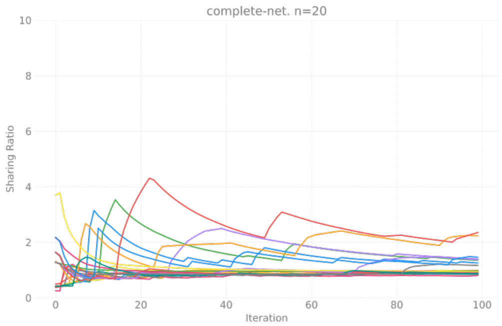

(a) $T = 100$

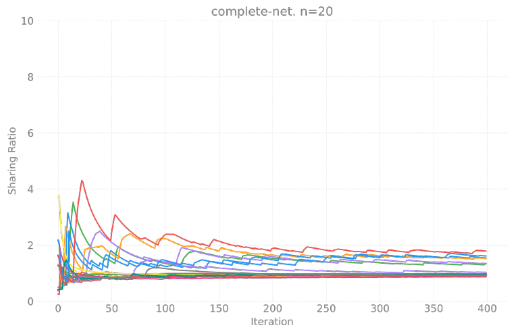

(b) $T = 400$

Figure 7.16: Agent-level convergence of sharing ratios in full model 3 in a complete network of size $n = 20$ for $T = \{100, 400\}$ iterations. Each agent's sharing ratio is shown as a separate line.

7.7. Experiment 6: Misreporting Sharing Ratios

Focus and Configuration

This experiment focusses on identifying a bound for the IR in case that an agent misreports its sharing ratio. We note that agent i misreporting ρ_i can only be beneficial if its neighbourhood has agents with the allocation strategy weight $w > 0$: A purely proportional neighbour does not consider the sharing ratio of agent i, and can thus not be manipulated in such a way. However, the influence in mixed populations is not as clear-cut. This experiment can shed light on this issue.

Additionally, we hypothesize that decreasing the reported sharing ratio ρ'_i as much as possible is expected to result in the highest returned utility, and thus the highest incentive ratio: Given the greedy portion of strategy 6.1, agent j allocates to agent i for which $\rho_i \in \min_{k \in \mathcal{N}_j} \rho_k$. Reporting $\rho'_i = 0$ guarantees that agent i is in $\min_{k \in \mathcal{N}_j} \rho_k$ and thus always receives resources from agent j.

Network types	n	Reps.
Complete	$\{8, 20\}$	100

((a)) Network type, population size, and number of repetitions for Experiment 6.

Allocation Strategy	w	Update Timing	γ	Manipulation
π	1	Sync.	0	Misreporting ρ
ϕ	0	Sync.	0	Misreporting ρ
ψ (hom.)	0.5	Sync.	0	Misreporting ρ
ψ (het.)	$[0, 1]$	Sync.	0	Misreporting ρ

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 6.

Table 7.17: Configuration of Experiment 6.

Results

Figure 7.17 depicts the average and maximum IR for each of the four settings for w, in a complete network setting of size $n = 8$. The IR metrics are plotted over the different values of misreporting factors α, where we clearly see the IR decrease for higher values of α.
The greedy strategy π in Figure 7.17 (a) achieves a significantly higher IR than the mixed strategies ψ in Figure 7.17 (c) and Figure 7.17 (d), both in terms of maximum and average values. Proportional strategy ϕ shows an IR of roughly 1, while the homogeneous and heterogeneous cases yield higher values for IR, similar to one another.

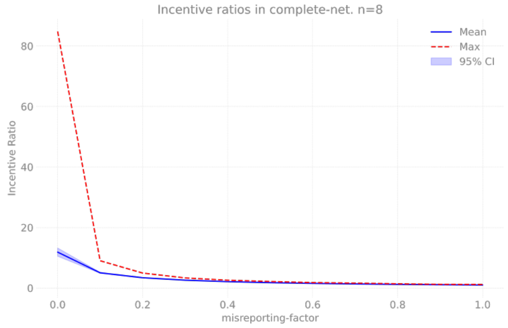

(a) π

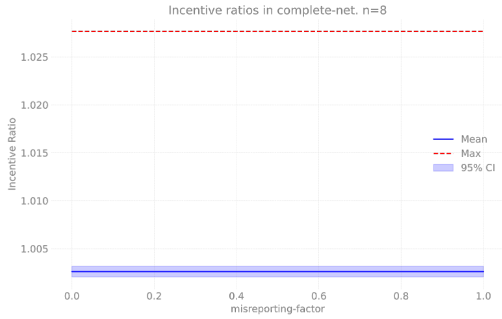

(b) ϕ

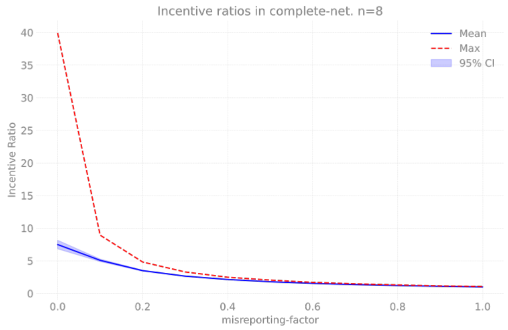

(c) ψ with $w = 0.5$ (hom.)

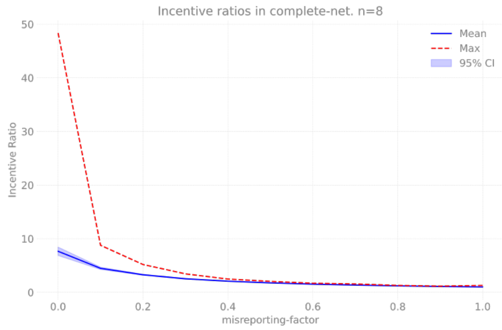

(d) ψ (het.)

Figure 7.17: Maximum and mean IR, including 95% confidence interval, of different market settings for w, in a complete network of size $n = 8$ for $T = 400$ iterations.

We also present two instance-specific results in Figure 7.18, which are representative for all observed network instances, in which convergence of sharing ratios for two specific networks are plotted. In both cases, agent 0 misreports ρ_0 with a factor of $\alpha_0 = 0$, for which agent 0 should obtain the highest utility. This observation, namely that the lowest value for α obtains a the highest IR, and disrupts the market the most, becomes apparent in the results of Appendix B.2, which present the market-level convergence of the loss for different values of α. The legends show the values of the endowments $v_i \cdot \varepsilon_i$ for each agent $i \in \mathcal{N}$, and clearly the Figures in Appendix B.2 consistently present the largest value for loss (and thus the highest disturbances to the market) at $\alpha = 0$.

We observe that for both instances in Figure 7.18 the sharing ratio of the manipulative agent 0 converges to 7.33 for instance 1 (a) and 3.92 for instance 2 (b), both of which equal $\sum_{j \in \mathcal{N}_0} v_j \cdot \varepsilon_j / (v_0 \cdot \varepsilon_0)$. As the results show, ρ_0 converges to $(\sum_{j \in \mathcal{N}_0} v_j \cdot \varepsilon_j)/(v_0 \cdot \varepsilon_0)$. This means that this kind of manipulation strategy seemingly reaches the hypothesised upper bound.

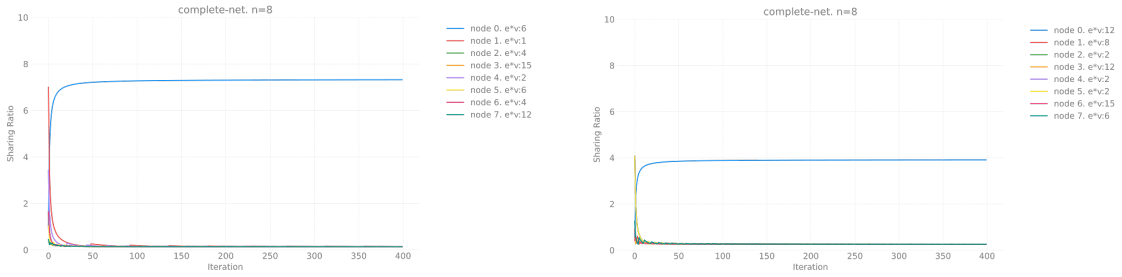

Network instance #1 Network instance #2

Figure 7.18: Sharing ratios in two instances of a complete network of size $n = 8$ for $T = 400$ iterations, where the sharing ratio of the misreporting agent 0 clearly supercedes that of truthful agents.

7.8. Experiment 7: Sybil Attack

Focus and Configuration

This experiment focusses on identifying a bound for the IR of a Sybil attack. Each setting will be repeated for each of the three divisions of endowments (Proportional, Even, and Random), to increase the chances of finding an IR close to the theoretical upper bound. We hypothesize that, as network size decreases, the relative influence of a Sybil identity i^m on some agent j increases, and thus the IR increases.

Network types	$\mathbf{n}$	**Reps.**
{Complete}	$\{4, 8\}$	100

((a)) Network type, population size, and number of repetitions for Experiment 7.

Allocation Strategy	$\mathbf{w}$	**Update Timing**	γ	**Manipulation**
π	1	Sync.	0	Sybil attack
ϕ	0	Sync.	0	Sybil attack
ψ (hom.)	0.5	Sync.	0	Sybil attack
ψ (het.)	$[0, 1]$	Sync.	0	Sybil attack

((b)) Configuration of strategy weight, update type, and memory decay, and manipulation parameters for each line in Experiment 7.

Table 7.18: Configuration of Experiment 7.

Results

The Figures 7.19 through 7.22 below show the results of the experiments of the Sybil attacks. The tables show the maximum found IR for that setting over all 100 randomly generated network instances. The highest out of all three types of endowment divisions for the Sybil attack is marked in bold. Below each table we see the convergence of sharing ratios in the two instances that are in bold for that table, i.e. the highest scoring instances for sizes $n = 4$ and $n = 8$.

For $n = 4$, the random Sybil strategy within a greedy population yields the highest IR, at 2.44536. For $n = 8$, the even Sybil strategy in a greedy population achieves the maximum IR, at 1.48300. These observed IR values for the greedy strategy exceed the previously determined IR bound of $\sqrt{2}$ for Sybil attacks in proportionally trading populations in complete networks. The heterogeneous mixed strategy also exceeds this bound of $\sqrt{2}$, and thus, we can state that both mixed and greedy strategies, demonstrate a lower robustness to Sybil attacks than proportional strategies.

	Proportional	Even	Random
$n = 4$	1.59700	2.30283	**2.44536**
$n = 8$	1.26645	**1.38400**	1.38389

Table 7.19: Maximum IR scores of different Sybil strategies in a greedy (π) population, for a complete network of sizes $n = \{4, 8\}$ for $T = 400$ iterations.

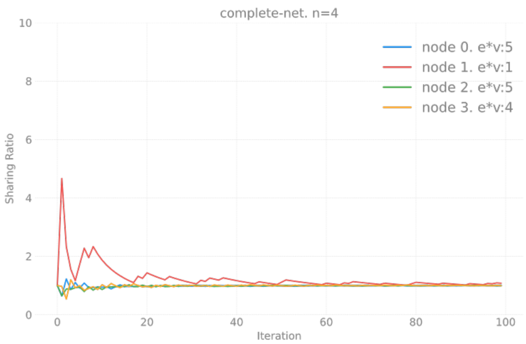

(a) $n = 4$, Sybil attack by agent 1

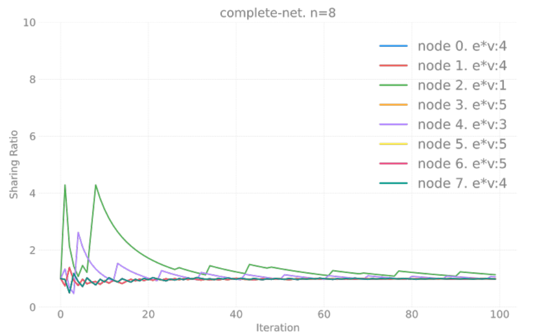

(b) $n = 8$, Sybil attack by agent 2

Figure 7.19: Sharing ratios in the two instances with the highest IR for a greedy (π) population in a complete network of sizes $n = \{4, 8\}$ for $T = 100$ iterations.

	Proportional	Even	Random
$n = 4$	**1.15618**	1.01011	0.97755
$n = 8$	**1.00858**	1.00855	0.98490

Table 7.20: Maximum IR scores of different Sybil strategies in a proportional (ϕ) population for a complete network of sizes $n = \{4, 8\}$ after $T = 400$ iterations.

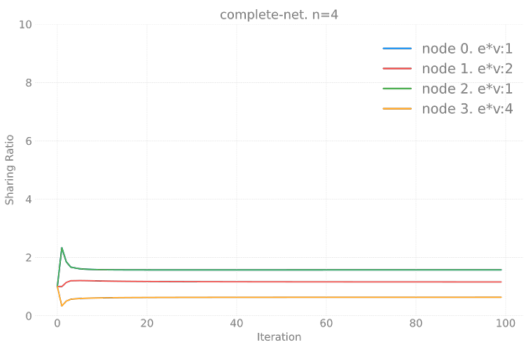

(a) $n = 4$, Sybil attack by agent 1

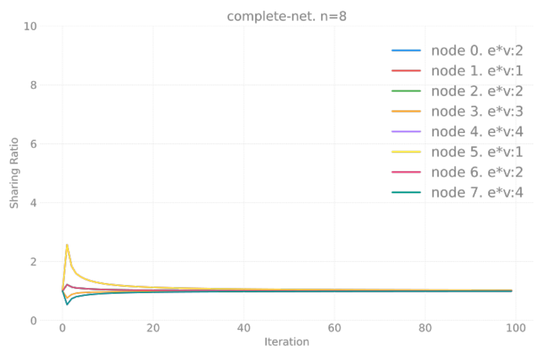

(b) $n = 8$, Sybil attack by agent 3

Figure 7.20: Sharing ratios in the two instances with the highest IR for a proportional (ϕ) population in a complete network of sizes $n = \{4, 8\}$ for $T = 100$ iterations.

	Proportional	**Even**	**Random**
$n = 4$	$\mathbf{1.22133}$	1.06328	1.06722
$n = 8$	$\mathbf{1.02692}$	1.02418	1.02229

Table 7.21: Maximum IR scores of different Sybil strategies in a mixed homogeneous population, for a complete network of sizes $n = \{4, 8\}$ after $T = 400$ iterations.

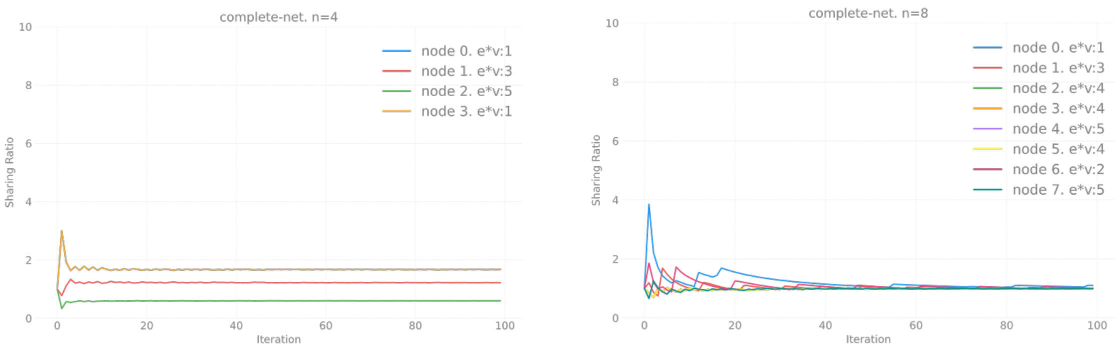

(a) $n = 4$, Sybil attack by agent 1 (b) $n = 8$, Sybil attack by agent 0

Figure 7.21: Sharing ratios in the two instances with the highest IR for a homogeneous mixed population in a complete network $n = \{4, 8\}$ after $T = 100$ iterations.

	Proportional	**Even**	**Random**
$n = 4$	$\mathbf{1.47448}$	1.44915	1.27565
$n = 8$	1.17683	$\mathbf{1.18530}$	1.09794

Table 7.22: Maximum IR scores of different Sybil strategies in a mixed heterogeneous population, for a complete network of sizes $n = \{4, 8\}$ after $T = 400$ iterations.

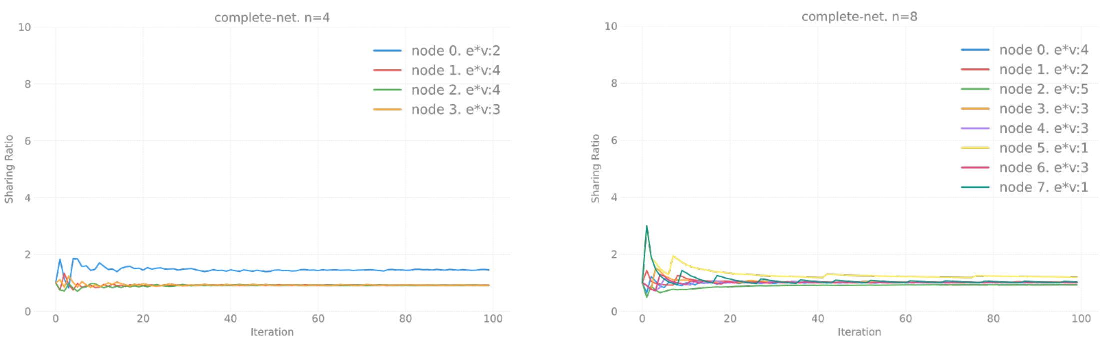

(a) $n = 4$, Sybil attack by agent 0 (b) $n = 8$, Sybil attack by agent 5

Figure 7.22: Sharing ratios in the two instances with the highest IR for a heterogeneous mixed population in a complete network of sizes $n = \{4, 8\}$ for $T = 100$ iterations.

8

Discussion of Numerical Results

8.1. Strategic Diversity

Understanding Observed Behaviour

The results of Experiments 1 and 2 shed light on how differences in agent strategies influence convergence in market settings. In the idealised setting of Experiment 1, both strategies π and ϕ converge toward zero loss, suggesting the market approaches equilibrium. While not a formal proof, these empirical results provide supporting evidence for the equilibrium convergence of both strategies, as claimed in Chapter 4.
Figure 7.2 shows that ϕ follows a convergence pattern roughly consistent with $O(1/t)$, a rate previously observed for PRD in AD and Fisher markets [69, 70], suggesting that a similar bound may hold in exchange markets as well.
Although we lack a formal convergence rate for π, empirical evidence suggests it converges more slowly than ϕ, with significant implications: While a market is out of equilibrium, some agents accumulate disproportionately high average utility ($\exists i \in \mathcal{N} : \overline{U_i}(t) > U_i^*$), creating an incentive to exit the market early.

As discussed in Chapter 5, the idealised conditions of Experiment 1 may not generalise to real-world performance. Experiment 2 introduces mixed strategies, providing a more realistic test of strategic interaction. In homogeneous populations, convergence to the equilibrium occurs reliably across all network structures, with behaviour interpolating between that of π and ϕ, as predicted by Theorem 6.1.1.
In contrast, heterogeneous populations show slightly impaired convergence in some networks. It is still uncertain whether this reflects convergence to a sub-optimal state or slow convergence to an optimal state. The value of obtaining an equilibrium convergence becomes evident here: Due to the difference in convergence of homogeneous and heterogeneous strategies, we can conclude that it is not mixing per se that causes issues in terms of market optimality, but rather the presence of strategic diversity across agents.

Given that Theorem 6.1.1 relies on the homogeneity of an agent population, possible degradation in performance under heterogeneous conditions was anticipated. However, to understand why heterogeneity impedes convergence, we draw a parallel with federated learning — a field long known to be sensitive to heterogeneity [146]. Recall Jain's analogy of the market as a Turing machine: If the market is understood as a distributed machine 'learning' an equilibrium, then strategic heterogeneity introduces a form of noise or non-alignment similar to that which hinders optimisation in federated learning systems.

These findings are consistent with earlier work. Le and Ramazi [100] demonstrated that heterogeneity in strategy design can prevent markets from reaching equilibrium in competitive games.

Additionally, Chien and Sinclair [147] raised concerns about the necessity of explicit coordination mechanisms to ensure convergence in the presence of agent heterogeneity. Together, these findings confirm that the convergence of a market to equilibrium is not an intrinsic property and can be disturbed by the diversity of the agents in it.

Modelling Implications

In terms of market modelling, these insights suggest that maintaining behavioural uniformity is more effective for ensuring convergence to equilibrium than fine-tuning individual strategy parameters. Markets in which agents adopt simple, homogeneous strategies, such as proportional updates, tend to be more robust and predictable. However, overlooking strategic diversity in the name of theoretical optimality can lead models to oversimplify market behaviour, resulting in sub-optimal outcomes and failure to capture real, practical sources of market inefficiencies.

These findings reflect a fundamental trade-off in real-world protocols between allowing strategic freedom and ensuring fairness and stability. While recognising that not each trading protocol is possible in each market (e.g., markets involving indivisible goods require allocation strategies based on integral allocations), we observe that this trade-off can be addressed by incorporating a range of allocation strategies and consensus mechanisms is essential for accurately analysing and predicting efficient market dynamics.

Rather than imposing top-down restrictions on agent behaviour, consensus algorithms allow agents to coordinate through mutual agreement, thereby maintaining convergence to the equilibrium without sacrificing agent autonomy.
Such consensus policies are commonly used in practice to reduce market-level inefficiencies. For example, P2P electricity markets among prosumers often employ consensus-based coordination schemes to achieve stable, system-wide outcomes, even in the presence of participant heterogeneity and network constraints [148].
These results show that even low agent heterogeneity can slow convergence and reduce fairness, prompting protocol design in diverse markets to reassess whether mechanisms truly reflect participants' trading preferences. As shown earlier, a lack of consensus among agents often necessitates standardised measures like money or tokens. Where such standardisation is infeasible, evaluating markets solely through theoretical equilibrium metrics might not be entirely suitable, as our findings suggest these may not be reliably attained in heterogeneous settings.

Overall, we have highlighted a fundamental trade-off in protocol design: Allowing for strategic freedom may enhance flexibility and agent autonomy, but it can also compromise convergence and fairness. Effective market mechanisms must explicitly manage this tension by incorporating both allocation diversity and coordination mechanisms. Without modelling such diversity, there is risk of overestimating market efficiency.

8.2. Temporal Effects

Understanding Observed Behaviour

In Experiment 3, we examined asynchronous dynamics on strategies π and ϕ. Asynchrony slows convergence at the market level for both strategies, with smaller batch sizes causing greater delays (Figure 7.8). At the agent level, updates appear more smooth where sharing ratio changes become more gradual due to staggered updates.
Proportional strategies such as ϕ avoid this issue, as proportional allocations rely only on current local states without correcting past imbalances, yielding inherently smooth convergence.

Given the connection between mirror descent and ϕ, it is natural to view sequential ϕ as a form of coordinate descent. This sheds light on why its sequential and asynchronous variants converge more slowly compared to the synchronous version: In settings where computing the full gradient is no more expensive than computing the partial gradient, gradient descent can converge faster than coordinate descent, in terms of iterations t [149].

There is an important discrepancy between our findings and how asynchronous algorithms usually behave. In computer science, asynchronous algorithms are often seen as faster alternatives to synchronous ones because synchronous algorithms require agents to wait for each other before proceeding. This waiting makes synchronous updates slower in real time, so asynchronous updates can speed up the overall process. However, in our model, time is divided into discrete steps (iterations), and only one agent updates per step in the asynchronous case. This means that, measured by iterations, asynchronous dynamics do not actually speed up the allocation process compared to synchronous dynamics. Instead of multiple agents updating simultaneously (as in synchronous rounds), asynchronous updates happen one at a time, resulting in slower, incremental progress rather than faster, parallel updates. This demonstrates that even the most basic abstractions in market modelling, such as reducing time to iterations, can have significant effects.

The slowing effect of asynchronous dynamics aligns with prior research using similar iteration-based asynchronous models. Earlier studies have consistently found that delays in agent updates tend to slow convergence in market environments [150, 151]. Our findings reinforce this broader understanding of how asynchrony impacts market efficiency.
Moreover, our results build upon the work of Kolumbus et al. [67], who demonstrated convergence of asynchronous PRD in Fisher markets. However, their work does not directly compare asynchronous market behaviour to synchronous baselines. In contrast, we contextualise asynchronous ϕ by demonstrating it converges more slowly than its synchronous equivalent.
Kolumbus et al. also compare asynchronous PRD to an asynchronous best-response strategy, which converges faster than asynchronous PRD. At first glance, this appears to conflict with our observation that π performs worse than ϕ under asynchrony. Although this best-response strategy and π are both forms of greedy algorithms, the key difference is that the best-response strategy optimises the overall market objective, while π focuses on individual gains. This underlines the importance of understanding agents' objectives, and consequently, their strategies: Strategies aligned with global objectives tend to converge more rapidly to equilibrium than those prioritising individual incentives. Since individuals often prioritise personal gains over market objectives (even preferring free-riding above active participation) [152], effective incentives are necessary to encourage actions beneficial to the market.

Modelling Implications

From a system design perspective, these results highlight a fundamental trade-off: Asynchronous updates induce stability by rounding off abrupt sharing ratio spikes and distributing market shocks over time. However, this comes at the cost of slower convergence to equilibrium.
In time-sensitive markets, such as rapid auctions, delayed convergence can lead to inefficiencies. In such cases, synchrony-promoting mechanisms may provide a solution. For example, Feng et al. [153] propose a mechanism that prioritises specific trades, such as those with significant trading imbalances, which accelerates adjustment towards equilibria. This technique represents a best-case asynchronous scenario, in contrast to our worst-case adversarial asynchronous dynamics, explaining the accelerated convergence compared to our observed slowdown.
Conversely, in markets where stability is vital, asynchrony may be a desirable feature. By allowing the market to adjust incrementally, it naturally buffers against overreactions, mitigating the risk of destabilising cascades that synchronous updates might trigger. In this way, asynchronous dynamics act as a built-in stabiliser, trading speed for resilience.

We can thus conclude that protocols designed for real-world environments must be tested under asynchronous conditions, as synchronous assumptions — though simplifying analysis — are rarely realistic. Experiments show that deviations from synchrony can significantly impact protocol behaviour. Asynchronous updates, common in real markets, can both stabilise markets by dampening sharp reactions while also slowing the convergence to equilibrium. Therefore, market models should explicitly incorporate update timing and frequency to capture this trade-off between resilience and efficiency, informing design decisions around synchronisation and update protocols.

8.3. Unreliable Information Challenges

Understanding Observed Behaviour

In Experiment 4, increasing memory decay parameter γ causes agents to discount past performance, leading to oscillations in sharing ratios. These fluctuations reduce market efficiency.
In Experiment 5, combining strategic heterogeneity and asynchronous updates, memory decay remains the dominant influence. Though the system converges, it settles at suboptimal outcomes increasingly far from equilibrium.

Greedy strategies like π, which rely on historical averages, are particularly impacted by memory decay. Even if merely 10% of the population suffers from memory decay, we observe suboptimal market dynamics. Both the extent and intensity of decay substantially degrade global performance. Agents with decayed memory gain disproportionately, demonstrating that misreporting, or lazily reporting information is a simple yet effective strategy to obtain disproportionate gains.

Figure 7.13 shows that when only few agents are affected by memory decay, it is precisely those agents that benefit most. This can be explained as follows: After not receiving resources for several timeslots, the sharing ratio of a memory-decaying agent i drops sharply. The neighbours of agent i, using greedy allocation, then favour i, concentrating inflows towards i. Agent i's sharing ratio spikes due to the inflow. However, this elevated sharing ratio incurs no greater penalty than a moderately high one. From the perspective of a greedy neighbour, there is no strategic difference between i having a somewhat high or very high sharing ratio: both result in zero allocation. Thus, memory-decaying agents, whose ratios frequently drop to the minimum in their neighbourhood, can outperform full-memory agents, whose ratios stay stable and rarely reach the minimum. This indicates an asymmetry inherent to strategy π: Underreporting the sharing ratio (via memory decay) is rewarded, while overreporting is barely penalised.

At the market level, the observed higher loss and suboptimal state corroborate previous work by Temzelides and Yu [154], who argued that lack of public memory introduces frictions requiring money for coordination. Although our model does not include an explicit monetary device, it captures coordination through consensus on values. Nonetheless, memory-decaying populations converge to non-equilibrium states, confirming that a global value measurement alone may be necessary but not sufficient; market equilibrium convergence also depends on factors like the trading protocols used.

Although our results might appear at odds with studies such as Cavalli and Naimzada [155], who find weighted memory stabilises tâtonnement price dynamics, the key difference lies in the nature of the strategies. Tâtonnement is generally not dependent on memory, and the introduction of a weighted memory just meant an increase in information. On the other hand π depends heavily on memory and loses information with decay, causing non-equilibrium behaviour. This raises questions about whether weights other than exponential affect π differently, as identifying memory weights that enable a greedy strategy like π to match or outperform tâtonnement in convergence is challenging due to its reliance on historical data.

Modelling Implications

Agents benefitting more from volatile behaviour than from consistent behaviour creates opportunities for manipulation and encourages lazy participation where agents underreport received utilities. Experiment 6 exemplifies these effects: An agent misreporting its sharing ratio achieves a high IR, indicating significant deviation and increased loss. The greater the misreporting (higher α), the larger the agent's disproportional gains, as can be seen in Appendix B.2. The results of experiment 6 also show that agents with higher endowments relative to their neighbours have reduced incentives to misreport, whereas those with lower endowments benefit more from manipulation, consistent with Chen et al. [156], who find a linear increase in misreporting incentives as relative endowment decreases. Therefore, markets exhibiting significant disparities in endowment value should be mindful of manipulative low-endowment agents, who stand to gain most from manipulation.

As such, markets with more equal endowment distributions are less vulnerable to manipulation. Given the IR bound for some agent i — $(\sum_{j \in \mathcal{N}_i} v_j \cdot \varepsilon_j)/(v_i \cdot \varepsilon_i)$ — equal-resource markets yield an IR of exactly $n-1$, representing the most equitable outcome since any change in initial endowment raises the IR for some agent. Mechanisms aiming to equalise resource provision would thus enhance robustness.

In adversarial environments, protocols that accept agent inputs without verification are vulnerable to manipulation, undermining both fairness and stability. This challenge is especially pronounced in decentralised systems, where enforcing memory or global coordination is often infeasible due to privacy, storage, or design constraints.
Without embedded memory or accountability mechanisms, reliance on greedy or lazy agents leads to volatile and unfair outcomes, even in the absence of malicious intent. To address this, systems should support verifiability, for example through blockchain-based tools, which offer transparency while preserving user privacy.
These findings carry clear implications for real-world sharing economies: When local information, limited memory, and self-interest dominate, decentralised coordination risks becoming fundamentally unstable and unjust.

8.4. Risks of Decentralisation

The underlying network topology further shapes the dynamics discussed. Both pure strategies π and ϕ appear to converge to equilibrium across networks, but the nature of convergence varies by network type. In complete networks, convergence occurs with differing speed and smoothness. In contrast, sparse networks limit agents' allocation options, reducing strategy differences and constraining market dynamics.
This suggests a trade-off in terms of network connectivity: Scale-free networks provide robustness and hierarchy but limit strategic freedom, whereas complete networks encourage sensitivity, cohesion, and potential efficiency. Both allow equilibrium convergence, but fully connected networks enable a more flexible path. These findings align with Celata et al. [139], who highlight community strength and cohesion as key to participation and effective decentralised exchange.

Because the network is distributed, agents lack full topology knowledge. Information propagates node-by-node, making consensus more challenging than in centralised coordination and hindering straightforward policy implementation. Such communication challenges, induced by the topology of a network, have often lead to inefficiencies. For example, users of the *Gnutella* P2P network frequently experienced overload, as its flooding-based communication protocol was poorly aligned with the underlying network topology [157, 158]. Greater awareness of structural effects during the modelling stage could have helped mitigate this issue.

Another risk from this limited knowledge is shown in Sybil attacks (Experiment 7), where IR exceeds 1, confirming agents can benefit from Sybil attacks. The system's vulnerability arises because agents do not know the full network and cannot detect Sybil manipulations by neighbours. Our results identify that proportional endowment division often yields the most effective Sybil attack. However, this division requires knowledge of neighbours' endowments, rarely available in real distributed settings, limiting its practical feasibility.
Though Sybil attacks yield smaller IR than misreporting manipulations, they remain a serious threat; in game-theoretic terms, if Sybil creation is known to all agents to be feasible, creating many Sybils becomes a dominant strategy, triggering an arms race that escalates systemic risk. Many models proactively test for vulnerability to the Sybil attack (e.g., [159, 160, 25]), reflecting the significant concern for these attacks.

An interesting pattern is the inverse relationship between IR and population size: Smaller networks exhibit higher IR under Sybil attack. We hypothesise that as markets grow, self-regulation — the "invisible hand" — curbs manipulation, consistent with Chen et al. [103]. However, if large markets comprise disconnected clusters rather than a complete network, Sybil attacks remain effective locally, replicating small-network dynamics.

Unexpectedly, the manipulating agent is not always the main beneficiary of a Sybil attack. Figures 7.19 to 7.22 show sharing ratios converging away from equilibrium. While the manipulator often improves its ratio, agents with the lowest endowments frequently benefit most.
This suggests potential implicit collusion: Low-value agents might encourage high-value peers to launch Sybil attacks while staying passive. Regulatory focus should thus prioritise low-endowment agents, who disproportionately gain from manipulation. This presents a risk, not only for obtaining fair trading situations: Although findings up until now have pointed towards the value of consensus mechanisms in distributed markets, a Sybil attack could efficiently corrupt such consensus by multiplying the weight of its vote through multiple identities [161]. As such, Sybil attacks are risks in multiple phases of distributed trading.

Overall, these findings reveal a trade-off between privacy and robustness. Protocols relying on unverifiable identities or truthful reporting without validation are vulnerable. Without safeguards, performance and fairness guarantees fail. Our results highlight the essential role of modelling adversaries when evaluating decentralised exchange protocols.

8.5. System-Level Interactions

Experiment 5 combines previously studied imperfections — strategic heterogeneity, memory decay, and asynchronous updates. These factors do not exhibit strong interactive effects; rather, each independently contributes to the delayed convergence to a sub-optimal market state. Here is exemplified that in a complex market model, as Sanders et al. [162] emphasise, the existence of an equilibrium does not imply convergence to it
This additive relationship is particularly relevant to open questions raised by Kolumbus et al. [67] and Li et al. [130], who ask whether informational frictions and asynchronous timing jointly affect equilibrium. While we interpret informational frictions slightly differently, our findings suggest that such imperfections degrade performance primarily through independent, additive effects. Although both reduce convergence quality, we observe no visible interaction between them.

At the agent level, sharing ratios exhibit chaotic oscillations before stabilising (Figure 7.16). These dynamics obscure the individual effects of the market features. Nonetheless, the combined effect is clear: The combination of memory decay, asynchrony, and heterogeneity leads to volatile and unpredictable allocation trajectories.
Such chaotic dynamics are well documented in large-scale multi-agent systems. Interestingly, while much of the literature attributes such chaos to complex learning behaviour, our results demonstrate that even simple trading heuristics can generate similar dynamics — emulating complex agent behaviour with simple concepts [36]. This supports the view that simple behavioural rules may suffice to capture key aspects of market behaviour, as argued by Boland [116].

Analogous patterns can be observed in real-world decentralised systems. For instance, energy micro-grids frequently exhibit early chaotic instability before reaching stable configurations [16]. These systems must balance adaptability and efficiency, often accepting short-term volatility in exchange for decentralised control. Our model replicates this pattern: initial chaos followed by stable yet imperfect outcomes, suggesting that the interplay of heterogeneity, memory, and asynchrony is a structural feature rather than a modelling artefact. Since real equilibria are rarely static, the complexity and unpredictability observed in our model indicate that its dynamic behaviour reflects relevant features of real-world systems and are indeed plausible.

Our modelling abstraction is diagnostic rather than predictive: Its purpose is not to replicate every domain detail, but rather to identify which realistic factors could undermine the equilibrium, thus directing future research and development efforts toward the most impactful shortcomings in current market models.
By examining these targeted relaxations, we reveal the fragility of equilibrium guarantees in idealised models. This suggests a validation process that begins with theoretical proofs of convergence, moves on to simulations that include realistic variations, and finally advances to real-system prototypes combining multiple elements to better capture the complexity and unpredictabil-

ity of actual markets. Future research should build on our P2P-adapted framework, ensuring that assessments of protocol optimality carefully consider these three key dimensions of realism before deployment.

8.6. Limitations

While the results offer valuable insights into decentralised market dynamics, several limitations of the experimental setup should be acknowledged. First, the behavioural model was deliberately simplified to maintain analytical clarity and make simulations feasible. While this enabled straightforward comparisons across strategies and conditions, it necessarily omits many real-world decision-making complexities. Such abstraction is inherent to all modelling — any model is, at best, an interpretation shaped by subjective choices. For example, even our specific interpretation of timing impacted outcomes in asynchronous settings. Recognising this interpretive nature is crucial when applying or extending the model.

Additionally, our experiments were designed to prioritise qualitative understanding over comprehensive quantification, focusing on a limited subset of the parameter space, such as variations in memory, asynchrony, and strategic heterogeneity, within relatively small agent populations of 8 to 20 agents. This constrained setup allowed us to investigate structural relationships in a controlled environment, but it also restricted the generalizability of the results and the depth of statistical analysis. Although outcomes were averaged over 30 repetitions and reported with 95% confidence intervals, the analysis remained largely descriptive. More formal statistical methods, such as hypothesis testing or variance decomposition, were not employed. These choices reflect an intentional focus on identifying broad, interpretable patterns in decentralised allocation rather than producing precise or empirically grounded predictions. Applying the model to real-world contexts would require a wider exploration of the parameter space, more robust statistical techniques, and empirical validation in future, case-specific studies.

Finally, in the last two experiments, we restricted our analysis to complete networks. While this choice simplifies the analysis and aligns well with P2P markets, it is not without loss of generality. Prior work has shown that network topology significantly impacts the achievable IR, particularly for manipulative strategies.
For the misreporting strategy, we expect the complete network to yield relatively high IR values, as every agent can directly influence all others. This makes a misreporting agent effectively act as a resource sink, amplifying its impact. This aligns with earlier findings in Chapter 7 showing that sparse networks were more robust to losses due to their structural constraints.
However, in the case of Sybil attacks, we can make no similar statement, as previous work has shown lower IR bounds for complete networks. Future work should therefore revisit how different network structures affect vulnerability and strategic behaviour.

Together, these limitations suggest improvements to be made in future research on market modelling: Expanding the parameter space and incorporating empirical data would strengthen the applicability to specific cases, and thus the practical relevance of the insights presented in this study.

9 Conclusion

In recent years, sharing economies have gained substantial traction, marked by the increasing prevalence of agentic networks in which autonomous agents engage in distributed exchange according to specific protocols. While such systems have immense potential to streamline resource allocation and decentralise market power, they simultaneously pose significant challenges to conventional algorithmic economic theory. While they can serve as useful reference points, traditional frameworks often rely on assumptions of price-based exchange, full rationality, and centralised coordination, making them ill-equipped to model the dynamic and distributed nature of these emerging systems.

This disconnect between these idealised models and practical realities has led to failures in real-world protocol deployments where mismatches between assumptions and actual agent behaviour resulted in unexpected suboptimal market performance. In light of these shortcomings, we argue for a fundamental reconceptualisation of algorithmic market design — one that is both rooted in fundamental economic theory and tailored to the characteristics of sharing economies.

9.1. Research Aim and Contributions

The central aim of this thesis has been to model the distributed and behaviourally rich nature of exchange in sharing economies and bridge the gap to idealised, classical economic theory. Grounding the model in classical economic theory allowed us to preserve principles such as equilibrium and Pareto optimality while extending applicability to agentic, P2P environments. This approach facilitates a deeper understanding of how structural features of sharing economies and behavioural features of agents interact. More specifically, we sought a model that is both realistic in its assumptions and abstract enough to generalise across various sharing economy settings. To this end, we constructed a principled model of distributed exchange, incorporating asynchronous trading, bounded rationality, and heterogeneous strategies while replacing global information with local feedback mechanisms.

A key technical contribution lies in our analysis of distributed trading strategies. Specifically, we adapted programs well-known in conventional economic theory to align with the exchange market. We then used these to demonstrate, for the first time in this context, why proportional and greedy trading strategies can lead to equilibrium — contributing a novel insight into the dynamics of distributed resource allocation.

Beyond this, we showed how our model produces a range of market phenomena with clear real-world relevance. For instance, we observed a fundamental trade-off between cognitive bounds and robustness: Markets that prioritise autonomy often sacrifice efficiency due to reduced robustness against manipulation. Similarly, markets with discrete allocations showed more erratic be-

haviour, as the limited decision space prevented agents from reaching ideal outcomes, reflecting real-world constraints on feasible exchanges. These findings underline the sensitivity of decentralised systems to timing, cognition, and structure.
Building on these insights, our results further illustrate that maintaining behavioural uniformity is crucial for market robustness and reliable convergence, with globally aligned strategies supporting faster equilibrium. Even minor modelling choices—such as the timing of agent updates—can notably affect system dynamics, underscoring the importance of explicitly accounting for update frequency when balancing resilience and efficiency. Most critically, we identified that markets lacking trustworthy information guarantees are especially vulnerable to manipulation, as passive agents can exploit system asymmetries, leading to volatility and undermining trust.
Overall, our results suggest that sharing economies can self-organise toward fairness and efficiency, however, only under specific behavioural and structural conditions. By simulating these conditions explicitly, we gain insight into how certain market features, such as valuation diversity or memory constraints, can either support or undermine desirable outcomes.

9.2. Research Outcome and Implications

In light of our main research question — *How can models of behaviourally-driven distributed exchange capture the dynamics of sharing economies?* — we conclude that incorporating behavioural and structural realism is crucial to understanding how sharing economies operate; by constructing and analysing an agent-based model extended with bounded rationality, strategic heterogeneity, and asynchronous interaction, we demonstrated that behaviourally-informed models can meaningfully simulate the complex dynamics of decentralised exchange. These models capture how agents adapt to uncertainty, how local interactions produce global effects, and how structural factors like timing, information asymmetry, and allocation granularity shape overall outcomes.

The practical implication of our work is that protocol design must explicitly account for the induced trade-offs between autonomy, information asymmetry, and market robustness. For instance, when agent data is kept private, market efficiency may suffer due to vulnerability to local manipulation. It is thus important to decide where to place a market along these and other trade-offs. Additionally, a system that tolerates a single manipulative strategy may unintentionally incentivise others to follow suit. Effective regulation must target both individual behaviour, as well as the second-order incentives created by interaction rules.
At a more high level, our framework offers guidance to evaluate interactions and trade-offs within markets, helping to anticipate vulnerabilities and design protocols that foster stability, efficiency, and fairness.

Taken together, our findings indicate that the proposed model effectively captures the behaviour of sharing economies, while also highlighting the importance of principled design in doing so. Such models allow us to infer interpretable, causal relationships between protocol design and emergent market behaviour. In turn, the proposed market model can help designers and regulators anticipate market failures and develop appropriate regulations and incentives proactively, rather than retroactively.

9.3. Final Reflections

This thesis has highlighted the value of interdisciplinary perspectives in addressing complex sociotechnical system modelling. As technological progress continues to drive specialisation and fragmentation of markets, the need for principled exchange intensifies. Agentic, distributed sharing economies offer a promising avenue for reducing friction and increasing efficiency, provided that their underlying mechanisms are designed with care.
These developments compel us to consider what kind of inclusive, decentralised, and transparent sharing economy we wish to create. The design choices we make in modelling the markets will shape the effectiveness of our regulation in economic infrastructures.

9.4. Future work

Due to the technical contributions made — such as explaining the convergence of distributed trading strategies along new, central programs of the exchange market — several new open problems have emerged for future research. One key area concerns formal proofs of convergence rates. While the experimental results in Chapter 7 demonstrated a market equilibrium in many settings, these findings remain empirical. A formal proof showing if and how heterogeneous mixed strategies converge to market equilibrium would be a significant advancement for the field. Similarly, the exact convergence rates of proportional strategies and asynchronous greedy strategies remain open questions. Future work looking into more formal convergence proofs of heterogeneous mixed strategies would open doors for more diverse, less restrictive trading protocols being used.

More broadly, further research could explore the dependence of non-monetary distributed trading on consensus. As discussed in Section 4.5, a promising direction is to develop distributed algorithms for exchange markets that converge even without consensus on resource valuations. This would help model fair exchange without relying on explicit pricing mechanisms, making the role of consensus a particularly interesting focus.

Another open direction is to expand the spectrum of agent heterogeneity. This thesis considered linear combinations of two algorithms, however, many other strategy types exist (e.g., satisficing or petty behaviour). Incorporating these would significantly broaden the range of possible strategies within a population.

Additional open questions include a more network-oriented approach; market forces might constrain the neighbours with whom agents might trade [22], and as such, the underlying network representing these restrictions remain an interesting road to explore. More specifically, identifying specific network structures where proportional strategies outperform greedy ones. These directions could yield deeper insights into how different restrictions influence market outcomes. Given the ongoing evolution of real-world economies, this area of research remains relevant for further exploration.

References

[1] X. Zhu and K. Liu. “A systematic review and future directions of the sharing economy: business models, operational insights and environment-based utilities”. In: *Journal of Cleaner Production* 290 (Mar. 2021), p. 125209. ISSN: 0959-6526. DOI: 10.1016/j.jclepro.2020.125209.

[2] J. Taskinsoy. “The Reincarnation of Barter Trade Barter Economy”. In: *SSRN Electronic Journal* (2023). ISSN: 1556-5068. DOI: 10.2139/ssrn.4456717.

[3] K. Rong and Y. Luo. “Toward born sharing: The sharing economy evolution enabled by the digital ecosystems”. In: *Technological Forecasting and Social Change* 196 (Nov. 2023), p. 122776. ISSN: 0040-1625. DOI: 10.1016/j.techfore.2023.122776.

[4] S. Dolnicar. *Peer-to-Peer Accommodation Networks*. Goodfellow Publishers, Dec. 2017. ISBN: 9781911396512. DOI: 10.23912/9781911396512-3454.

[5] S. Netter et al. “Sharing economy revisited: Towards a new framework for understanding sharing models”. In: *Journal of Cleaner Production* 221 (June 2019), pp. 224–233. ISSN: 0959-6526. DOI: 10.1016/j.jclepro.2019.02.225.

[6] K. Frenken and J. Schor. “Putting the sharing economy into perspective”. In: *Environmental Innovation and Societal Transitions* 23 (June 2017), pp. 3–10. ISSN: 2210-4224. DOI: 10.1016/j.eist.2017.01.003.

[7] M. A. Petruzzi et al. “To Share or to Exchange: An analysis of the sharing economy characteristics of Airbnb and Fairbnb.coop”. In: *International Journal of Hospitality Management* 92 (Jan. 2021), p. 102724. ISSN: 0278-4319. DOI: 10.1016/j.ijhm.2020.102724.

[8] A. Mackelprang et al. “Examining sharing economy operations: A process perspective”. In: *Journal of Operations Management* 69.5 (June 2023), pp. 706–718. ISSN: 1873-1317. DOI: 10.1002/joom.1269.

[9] M. Yıldız and M. Altan. “Literature review of the sharing economy: Socio-cultural perspective”. In: *Journal of Community amp; Applied Social Psychology* 33.5 (July 2023), pp. 1112–1136. ISSN: 1099-1298. DOI: 10.1002/casp.2728.

[10] M. Sadiq et al. “Sharing economy benefits and sustainable development goals: Empirical evidence from the transportation industry of Vietnam”. In: *Journal of Innovation Knowledge* 8.1 (2023), p. 100290. ISSN: 2444-569X. DOI: https://doi.org/10.1016/j.jik.2022.100290.

[11] L. E. Kim. *An Inclusive Sharing Economy*. BSR, 2016. URL: https://www.bsr.org/reports/BSR_An_Inclusive_Sharing_Economy.pdf.

[12] K. P. Tsoukatos. “Network Formation by Reciprocity Versus Sparsity Tradeoffs”. In: *IEEE Transactions on Network Science and Engineering* 7.3 (2020), pp. 1752–1763. DOI: 10.1109/TNSE.2019.2952454.

[13] E. Mengelkamp et al. “Designing microgrid energy markets”. In: *Applied Energy* 210 (Jan. 2018), pp. 870–880. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2017.06.054.

[14] P. Vytelingum et al. “Trading agents for the smart electricity grid”. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*. AAMAS ’10. Toronto, Canada: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 897–904. ISBN: 9780982657119. DOI: 10.1145/1838206.1838326.

[15] Y. Parag and B. K. Sovacool. “Electricity market design for the prosumer era”. In: *Nature Energy* 1.4 (Mar. 2016). ISSN: 2058-7546. DOI: 10.1038/nenergy.2016.32.

[16] T. Sousa et al. “Peer-to-peer and community-based markets: A comprehensive review”. In: *Renewable and Sustainable Energy Reviews* 104 (Apr. 2019), pp. 367–378. ISSN: 1364-0321. DOI: 10.1016/j.rser.2019.01.036.

[17] N. Liu et al. “Energy-Sharing Model With Price-Based Demand Response for Microgrids of Peer-to-Peer Prosumers”. In: *IEEE Transactions on Power Systems* 32.5 (Sept. 2017), pp. 3569–3583. ISSN: 1558-0679. DOI: 10.1109/tpwrs.2017.2649558.

[18] C. Long et al. “Peer-to-peer energy trading in a community microgrid”. In: *2017 IEEE Power Energy Society General Meeting*. 2017, pp. 1–5. DOI: 10.1109/PESGM.2017.8274546.

[19] V. Conitzer and R. Freeman. “Algorithmically Driven Shared Ownership Economies”. In: *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers*. Ed. by J.-F. Laslier et al. Cham: Springer International Publishing, 2019, pp. 275–285. ISBN: 978-3-030-18050-8. DOI: 10.1007/978-3-030-18050-8_38.

[20] Y. Izmirlioglu et al. “A Survey of Multi-Agent Systems for Smartgrids”. In: *Energies* 17.15 (July 2024), p. 3620. ISSN: 1996-1073. DOI: 10.3390/en17153620.

[21] Y. Wang et al. *A Consumer-Centric Market for Database Computation in the Cloud*. 2016. DOI: 10.48550/ARXIV.1609.02104.

[22] D. M. Rothschild et al. *The Agentic Economy*. 2025. DOI: 10.48550/ARXIV.2505.15799. arXiv: 2505.15799.

[23] *Collaboration in Cities: From Sharing to ‘Sharing Economy’*. Tech. rep. REF181217. World Economic Forum, 2017. URL: https://www3.weforum.org/docs/White_Paper_Collaboration_in_Cities_report_2017.pdf.

[24] S. He et al. *Fair Peer-to-Peer Content Delivery via Blockchain*. 2021. DOI: 10.48550/ARXIV.2102.04685.

[25] D. Levin et al. “Bittorrent is an auction: analyzing and improving bittorrent’s incentives”. In: *ACM SIGCOMM Computer Communication Review* 38.4 (Aug. 2008), pp. 243–254. ISSN: 0146-4833. DOI: 10.1145/1402946.1402987.

[26] F. Wu and L. Zhang. “Proportional response dynamics leads to market equilibrium”. In: *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. STOC07. ACM, June 2007, pp. 354–363. DOI: 10.1145/1250790.1250844.

[27] A. Sherman. “Guaranteeing performance through fairness in peer-to-peer file-sharing and streaming systems”. AAI3447946. PhD thesis. USA, 2010. ISBN: 9781124537191.

[28] P. Staudt and B. Richter. “Empirical case study of a digitally enabled energy community with prosumers and P2P trading”. In: *Energy, Sustainability and Society* 15.1 (Jan. 2025). ISSN: 2192-0567. DOI: 10.1186/s13705-025-00511-5.

[29] D. J. Brown and G. M. Heal. “The Optimality of Regulated Pricing: A General Equilibrium Analysis”. In: *Advances in Equilibrium Theory*. Springer Berlin Heidelberg, 1985, pp. 43–54. ISBN: 9783642516023. DOI: 10.1007/978-3-642-51602-3_3.

[30] K. Shomalzadeh and N. Monshizadeh. *Towards a Socially Acceptable Competitive Equilibrium in Energy Markets*. 2024. DOI: 10.48550/ARXIV.2409.04157.

[31] E. Nalbach. “Sharing Economy: Dynamic General Equilibrium Effects”. eng. In: Beiträge zur Jahrestagung des Vereins für Socialpolitik 2018: Digitale Wirtschaft - Session: Growth III C03-V2. Kiel, Hamburg: ZBW - Leibniz-Informationszentrum Wirtschaft, 2018. URL: https://hdl.handle.net/10419/181643.

[32] O. Shobande et al. “General Equilibrium Modelling: The State of Art”. In: (Jan. 2021). URL: https://mpra.ub.uni-muenchen.de/105313/8/MPRA_paper_105313.pdf.

[33] M. Buckmann and A. Joseph. "An Interpretable Machine Learning Workflow with an Application to Economic Forecasting". In: *International Journal of Central Banking* 19.4 (2023), pp. 449–522. URL: https://ideas.repec.org/a/ijc/ijcjou/y2023q4a10.html.

[34] A. Tiffin. "Machine Learning and Causality: The Impact of Financial Crises on Growth". In: *IMF Working Papers* 19.228 (Nov. 2019). ISSN: 1018-5941. DOI: 10.5089/9781513518305.001.

[35] J. J. Heckman and R. Pinto. "The Econometric Model for Causal Policy Analysis". In: *Annual Review of Economics* 14.1 (Aug. 2022), pp. 893–923. ISSN: 1941-1391. DOI: 10.1146/annurev-economics-051520-015456.

[36] M. Bichler et al. *Agentic Markets: Game Dynamics and Equilibrium in Markets with Learning Agents*. 2025. DOI: 10.48550/ARXIV.2506.18571.

[37] S. A. Suryanarayana et al. *Explainability in Mechanism Design: Recent Advances and the Road Ahead*. 2022. DOI: 10.48550/ARXIV.2206.03031.

[38] D. Proserpio et al. "You get what you give: theory and evidence of reciprocity in the sharing economy". In: *Quantitative Marketing and Economics* 16.4 (Aug. 2018), pp. 371–407. ISSN: 1573-711X. DOI: 10.1007/s11129-018-9201-9.

[39] S. G. Saunders and V. D. Truong. "Leveraging an 'economy of kindness' through nonmonetary value exchange to improve community well-being". In: *Marketing Theory* 24.3 (Sept. 2023), pp. 449–464. ISSN: 1741-301X. DOI: 10.1177/14705931231202446.

[40] G. Bai and S. R. Velamuri. "Contextualizing the Sharing Economy". In: *Journal of Management Studies* 58.4 (Nov. 2020), pp. 977–1001. ISSN: 1467-6486. DOI: 10.1111/joms.12652.

[41] L. Mekouar et al. "Peer-to-peer's most wanted: Malicious peers". In: *Computer Networks* 50.4 (Mar. 2006), pp. 545–562. ISSN: 1389-1286. DOI: 10.1016/j.comnet.2005.07.025.

[42] D. Kirste et al. *Automated Market Makers in Cryptoeconomic Systems: A Taxonomy and Archetypes*. 2023. DOI: 10.48550/ARXIV.2309.12818.

[43] S. Chudziak. "Studying economic complexity with agent-based models: advances, challenges and future perspectives". In: *Journal of Economic Interaction and Coordination* 20.2 (Dec. 2024), pp. 413–449. ISSN: 1860-7128. DOI: 10.1007/s11403-024-00428-w.

[44] W. Sutherland and M. H. Jarrahi. "The sharing economy and digital platforms: A review and research agenda". In: *International Journal of Information Management* 43 (Dec. 2018), pp. 328–341. ISSN: 0268-4012. DOI: 10.1016/j.ijinfomgt.2018.07.004.

[45] L. Walras. *Elements of Pure Economics: Or the Theory of Social Wealth*. American Economic Association v. 20. American Economic Association, the Royal Economic Association, and the Royal Economic Society, 1874. ISBN: 9780415313407.

[46] L. L. Price and R. W. Belk. "Consumer Ownership and Sharing: Introduction to the Issue". In: *Journal of the Association for Consumer Research* 1.2 (Apr. 2016), pp. 193–197. ISSN: 2378-1823. DOI: 10.1086/686270.

[47] R. Belk. "Sharing Versus Pseudo-Sharing in Web 2.0". In: *The Anthropologist* 18.1 (July 2014), pp. 7–23. ISSN: 0972-0073. DOI: 10.1080/09720073.2014.11891518.

[48] M. Strathern. "Qualified value: the perspective of gift exchange". In: *Barter, Exchange and Value: An Anthropological Approach*. Ed. by C. Humphrey and S. Hugh-Jones. Cambridge University Press, 1992, pp. 169–191.

[49] E. Eisenberg and D. Gale. "Consensus of Subjective Probabilities: The Pari-Mutuel Method". In: *The Annals of Mathematical Statistics* 30.1 (Mar. 1959), pp. 165–168. ISSN: 0003-4851. DOI: 10.1214/aoms/1177706369.

[50] V. I. Shmyrev. "An algorithm for finding equilibrium in the linear exchange model with fixed budgets". In: *Journal of Applied and Industrial Mathematics* 3.4 (Oct. 2009), pp. 505–518. ISSN: 1990-4797. DOI: 10.1134/s1990478909040097.

[51] K. Jain. “A Polynomial Time Algorithm for Computing the Arrow-Debreu Market Equilibrium for Linear Utilities”. In: *45th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, pp. 286–294. DOI: 10.1109/focs.2004.6.

[52] Z. Munn et al. “What are scoping reviews? Providing a formal definition of scoping reviews as a type of evidence synthesis”. In: *JBI Evidence Synthesis* 20.4 (Mar. 2022), pp. 950–952. ISSN: 2689-8381. DOI: 10.11124/jbies-21-00483.

[53] B. Codenotti et al. “The computation of market equilibria”. In: *ACM SIGACT News* 35.4 (Dec. 2004), pp. 23–37. ISSN: 0163-5700. DOI: 10.1145/1054916.1054927.

[54] D. Bertsimas et al. “The Price of Fairness”. In: *Operations Research* 59.1 (Feb. 2011), pp. 17–31. ISSN: 1526-5463. DOI: 10.1287/opre.1100.0865.

[55] X. Yan and W. Zhu. “Efficiency and Fairness in Resource Exchange”. In: *IEEE Transactions on Cloud Computing* 10.4 (Oct. 2022), pp. 2538–2549. ISSN: 2372-0018. DOI: 10.1109/tcc.2021.3066291.

[56] Z. Yang et al. *Online Fair Allocation with Best-of-Many-Worlds Guarantees*. 2024. DOI: 10.48550/ARXIV.2408.02403.

[57] H. R. Varian. “Equity, envy, and efficiency”. In: *Journal of Economic Theory* 9.1 (Sept. 1974), pp. 63–91. ISSN: 0022-0531. DOI: 10.1016/0022-0531(74)90075-1.

[58] E. Segal-Halevi. “Competitive Equilibrium For Almost All Incomes: Existence and Fairness”. In: (2017). DOI: 10.48550/ARXIV.1705.04212.

[59] L. Georgiadis et al. “On the Efficiency of Sharing Economy Networks”. In: *IEEE Transactions on Network Science and Engineering* 7.3 (July 2020), pp. 1094–1110. ISSN: 2334-329X. DOI: 10.1109/tnse.2019.2904431.

[60] K. J. Arrow and G. Debreu. “Existence of an Equilibrium for a Competitive Economy”. In: *Econometrica* 22.3 (July 1954), p. 265. ISSN: 0012-9682. DOI: 10.2307/1907353.

[61] R. Duan and K. Mehlhorn. “A combinatorial polynomial algorithm for the linear Arrow–Debreu market”. In: *Information and Computation* 243 (Aug. 2015), pp. 112–132. ISSN: 0890-5401. DOI: 10.1016/j.ic.2014.12.009.

[62] N. Nisan et al. *Algorithmic Game Theory*. Cambridge University Press, Sept. 2007. ISBN: 9780511800481. DOI: 10.1017/cbo9780511800481.

[63] X. Chen et al. *Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities*. 2009. DOI: 10.48550/ARXIV.0904.0644.

[64] P. W. Goldberg. *A Survey of PPAD-Completeness for Computing Nash Equilibria*. 2011. DOI: 10.48550/ARXIV.1103.2709.

[65] Y. K. Cheung et al. “Tatonnement beyond gross substitutes? Gradient descent to the rescue”. In: *Games and Economic Behavior* 123 (Sept. 2020), pp. 295–326. ISSN: 0899-8256. DOI: 10.1016/j.geb.2019.03.014.

[66] N. R. Devanur et al. *New Convex Programs for Fisher’s Market Model and its Generalizations*. 2016. DOI: 10.48550/ARXIV.1603.01257.

[67] Y. Kolumbus et al. *Asynchronous Proportional Response Dynamics in Markets with Adversarial Scheduling*. New Orleans, LA, USA, 2023. DOI: 10.48550/ARXIV.2307.04108.

[68] Y. K. Cheung et al. “Dynamics of Distributed Updating in Fisher Markets”. In: *CoRR* abs/1806.04746 (2018). DOI: 10.1145/3219166.3219189.

[69] S. Brânzei et al. *Proportional Dynamics in Exchange Economies*. 2019. DOI: 10.48550/ARXIV.1907.05037.

[70] B. Birnbaum et al. “Distributed algorithms via gradient descent for fisher markets”. In: *Proceedings of the 12th ACM conference on Electronic commerce*. EC ’11. ACM, June 2011. DOI: 10.1145/1993574.1993594.

[71] E. Nenakov and M. Primak. “One algorithm for finding solutions of the Arrow–Debreu model”. In: *Kibernetika* 3 (1983), pp. 127–128.

[72] Y. Ye. “A path to the Arrow–Debreu competitive market equilibrium”. In: *Mathematical Programming* 111.1–2 (Dec. 2006), pp. 315–348. ISSN: 1436-4646. DOI: 10.1007/s10107-006-0065-5.

[73] N. R. Devanur et al. “Market equilibrium via a primal–dual algorithm for a convex program”. In: *Journal of the ACM* 55.5 (Oct. 2008), pp. 1–18. ISSN: 1557-735X. DOI: 10.1145/1411509.1411512.

[74] N. R. Devanur et al. *A Rational Convex Program for Linear Arrow-Debreu Markets*. 2013. DOI: 10.48550/ARXIV.1307.8037.

[75] J. Garg and L. A. Végh. *A Strongly Polynomial Algorithm for Linear Exchange Markets*. 2018. DOI: 10.48550/ARXIV.1809.06266.

[76] K. Dvijotham et al. “Convergence of incentive-driven dynamics in Fisher markets”. In: *Games and Economic Behavior* 134 (July 2022), pp. 361–375. ISSN: 0899-8256. DOI: 10.1016/j.geb.2020.11.005.

[77] D. Jalota and Y. Ye. “Stochastic Online Fisher Markets: Static Pricing Limits and Adaptive Enhancements”. In: *Operations Research* (Nov. 2024). ISSN: 1526-5463. DOI: 10.1287/opre.2023.0636.

[78] Y. Gao et al. “Online market equilibrium with application to fair division”. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. NIPS ’21. Red Hook, NY, USA: Curran Associates Inc., 2021. ISBN: 9781713845393. DOI: 10.48550/ARXIV.2103.12936.

[79] L. Liao et al. “Nonstationary dual averaging and online fair allocation”. In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. NIPS ’22. New Orleans, LA, USA: Curran Associates Inc., 2022. ISBN: 9781713871088. DOI: 10.48550/ARXIV.2202.11614.

[80] S. R. Sinclair et al. *Sequential Fair Allocation of Limited Resources under Stochastic Demands*. 2020. DOI: 10.48550/ARXIV.2011.14382.

[81] Y. Yang et al. *Robustness of Online Proportional Response in Stochastic Online Fisher Markets: a Decentralized Approach*. 2024. arXiv: 2406.00160 [cs.GT].

[82] L. Choo et al. “Information aggregation in Arrow–Debreu markets: an experiment”. In: *Experimental Economics* 22.3 (Sept. 2019), pp. 625–652. ISSN: 1573-6938. DOI: 10.1007/s10683-017-9548-x.

[83] M. Shubik. “The General Equilibrium Model is Incomplete and not Adequate for the Reconciliation of Micro and Macroeconomic Theory”. In: *Kyklos* 28.3 (Aug. 1975), pp. 545–573. ISSN: 1467-6435. DOI: 10.1111/j.1467-6435.1975.tb02283.x.

[84] I. Caragiannis et al. “The Unreasonable Fairness of Maximum Nash Welfare”. In: *ACM Transactions on Economics and Computation* 7.3 (Aug. 2019), pp. 1–32. ISSN: 2167-8383. DOI: 10.1145/3355902.

[85] D. Zeng and A. Psomas. “Fairness-Efficiency Tradeoffs in Dynamic Fair Division”. In: *Proceedings of the 21st ACM Conference on Economics and Computation*. EC ’20. Virtual Event, Hungary: Association for Computing Machinery, 2020, pp. 911–912. ISBN: 9781450379755. DOI: 10.1145/3391403.3399467.

[86] Y. K. Cheung et al. “Tracing Equilibrium in Dynamic Markets via Distributed Adaptation”. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. AAMAS ’19. Montreal QC, Canada: International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1225–1233. ISBN: 9781450363099. DOI: 10.48550/ARXIV.1804.08017.

[87] J. Geanakoplos. “Arrow-Debreu Model of General Equilibrium”. In: *General Equilibrium*. Palgrave Macmillan UK, 1989, pp. 43–61. ISBN: 9781349198023. DOI: 10.1007/978-1-349-19802-3_3.

[88] W. A. Jackson. “Markets as dualistic, semi-decentralized organizations”. In: *Review of Evolutionary Political Economy* 5.1 (Nov. 2023), pp. 153–172. ISSN: 2662-6144. DOI: 10.1007/s43253-023-00111-z.

[89] M. Guzman and J. Stiglitz. *Towards a Dynamic Disequilibrium Theory with Randomness*. June 2020. DOI: 10.3386/w27453.

[90] A. Nongaillard. “An agent-based approach for distributed resource allocations”. Theses. Université des Sciences et Technologie de Lille - Lille I, Dec. 2009. URL: https://theses.hal.science/tel-00831365.

[91] W. Tushar et al. “Peer-to-Peer Trading in Electricity Networks: An Overview”. In: *IEEE Transactions on Smart Grid* 11.4 (July 2020), pp. 3185–3200. ISSN: 1949-3061. DOI: 10.1109/tsg.2020.2969657.

[92] C. Shen. “An Introduction to Non-Synchronous Trading”. 2022. URL: http://math.uchicago.edu/~may/REU2022/REUPapers/Shen,Charene.pdf.

[93] Y. K. Cheung and R. Cole. “Amortized Analysis of Asynchronous Price Dynamics”. In: *26th Annual European Symposium on Algorithms (ESA 2018)*. Ed. by Y. Azar et al. Vol. 112. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018, 18:1–18:15. ISBN: 978-3-95977-081-1. DOI: 10.4230/LIPIcs.ESA.2018.18.

[94] S. E. Page. “On Incentives and Updating in Agent Based Models”. In: *Computational Economics* 10.1 (1997), pp. 67–87. ISSN: 0927-7099. DOI: 10.1023/a:1008625524072.

[95] J. Tong et al. “Asynchronous Algorithms for Computing Equilibrium Prices in a Capital Asset Pricing Model”. In: *Asia-Pacific Journal of Operational Research* 36.05 (Oct. 2019), p. 1950023. ISSN: 1793-7019. DOI: 10.1142/s0217595919500234.

[96] M. Feldman and J. Chuang. “Overcoming free-riding behavior in peer-to-peer systems”. In: *ACM SIGecom Exchanges* 5.4 (July 2005), pp. 41–50. ISSN: 1551-9031. DOI: 10.1145/1120717.1120723.

[97] T. Broer et al. “Heterogenous Information Choice in General Equilibrium”. In: 2018. URL: https://api.semanticscholar.org/CorpusID:159226776.

[98] L. Panait and S. Luke. “Cooperative Multi-Agent Learning: The State of the Art”. In: *Autonomous Agents and Multi-Agent Systems* 11.3 (Nov. 2005), pp. 387–434. ISSN: 1573-7454. DOI: 10.1007/s10458-005-2631-2.

[99] B. Light. “General Equilibrium in a Heterogeneous-Agent Incomplete-Market Economy with Many Consumption Goods and a Risk-Free Bond”. In: *Operations Research* 71.4 (July 2023), pp. 1093–1111. ISSN: 1526-5463. DOI: 10.1287/opre.2023.2442.

[100] H. Le and P. Ramazi. “Heterogeneous Mixed Populations of Best-Responders and Imitators: Equilibrium Convergence and Stability”. In: *IEEE Transactions on Automatic Control* 66.8 (Aug. 2021), pp. 3475–3488. ISSN: 2334-3303. DOI: 10.1109/tac.2020.3016903.

[101] N. Chen et al. “Incentive Ratios of Fisher Markets”. In: *Automata, Languages, and Programming*. Springer Berlin Heidelberg, 2012, pp. 464–475. ISBN: 9783642315855. DOI: 10.1007/978-3-642-31585-5_42.

[102] N. Chen et al. “Incentive ratio: A game theoretical analysis of market equilibria”. In: *Information and Computation* 285 (May 2022), p. 104875. ISSN: 0890-5401. DOI: 10.1016/j.ic.2022.104875.

[103] N. Chen et al. “How Profitable Are Strategic Behaviors in a Market?” In: *Algorithms – ESA 2011*. Springer Berlin Heidelberg, 2011, pp. 106–118. ISBN: 9783642237195. DOI: 10.1007/978-3-642-23719-5_10.

[104] Y. Cheng et al. “Tightening Up the Incentive Ratio for Resource Sharing Over the Rings”. In: *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 2020, pp. 127–136. DOI: 10.1109/IPDPS47924.2020.00023.

[105] I. Polak. “The Incentive Ratio in Exchange Economies”. In: *Combinatorial Optimization and Applications*. Springer International Publishing, 2016, pp. 685–692. ISBN: 9783319487496. DOI: 10.1007/978-3-319-48749-6_49.

[106] Y. Cheng et al. “Limiting the Deviation Incentives in Resource Sharing Networks”. In: *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. AAMAS ’20. Auckland, New Zealand: International Foundation for Autonomous Agents and Multiagent Systems, 2020, pp. 1813–1815. ISBN: 9781450375184.

[107] Z. Chen et al. “Agent incentives of strategic behavior in resource exchange”. In: *Discrete Applied Mathematics* 264 (July 2019), pp. 15–25. ISSN: 0166-218X. DOI: 10.1016/j.dam.2018.07.025.

[108] Z. Chen et al. “Limiting User’s Sybil Attack in Resource Sharing”. In: *Web and Internet Economics*. Springer International Publishing, 2017, pp. 103–119. ISBN: 9783319719245. DOI: 10.1007/978-3-319-71924-5_8.

[109] R. Cole and Y. Tao. *Balancing the Robustness and Convergence of Tatonnement*. 2019. DOI: 10.48550/ARXIV.1908.00844.

[110] L. Tesfatsion. “Chapter 16 Agent-Based Computational Economics: A Constructive Approach to Economic Theory”. In: *Handbook of Computational Economics*. Elsevier, 2006, pp. 831–880. DOI: 10.1016/s1574-0021(05)02016-2.

[111] J. Haltiwanger and M. Waldman. “Rational Expectations and the Limits of Rationality: An Analysis of Heterogeneity”. In: *The American Economic Review* 75.3 (1985), pp. 326–340. ISSN: 00028282. URL: http://www.jstor.org/stable/1814803 (visited on 04/09/2025).

[112] X. Dong et al. “Bounded rational expectation: How it can affect the effectiveness of monetary rules in the open economy”. In: *Journal of International Financial Markets, Institutions and Money* 88 (Oct. 2023), p. 101845. ISSN: 1042-4431. DOI: 10.1016/j.intfin.2023.101845.

[113] E. D. Sadler. “Bounded Rationality and the Network Economist’s Dilemma”. In: *SSRN Electronic Journal* (2014). ISSN: 1556-5068. DOI: 10.2139/ssrn.2501690.

[114] G.-M. Angeletos and K. Sastry. *Inattentive Economies*. Oct. 2019. DOI: 10.3386/w26413.

[115] L. A. Boland. “Building models of learning and the equilibrium process”. In: *Equilibrium Models in Economics: Purposes and Critical Limitations*. Oxford University Press, Apr. 2017. ISBN: 9780190274320. DOI: 10.1093/acprof:oso/9780190274320.003.0016.

[116] L. A. Boland. “Recognizing knowledge and learning in equilibrium models”. In: *Equilibrium Models in Economics: Purposes and Critical Limitations*. Oxford University Press, Apr. 2017. ISBN: 9780190274320. DOI: 10.1093/acprof:oso/9780190274320.003.0007.

[117] R. Amir et al. “On the microeconomic foundations of linear demand for differentiated products”. In: *Journal of Economic Theory* 169 (May 2017), pp. 641–665. ISSN: 0022-0531. DOI: 10.1016/j.jet.2017.03.005.

[118] J. Levin. *General Equilibrium*. 2006. URL: https://web.stanford.edu/~jdlevin/Econ%20202/General%20Equilibrium.pdf.

[119] D. Gale. “The linear exchange model”. In: *Journal of Mathematical Economics* 3.2 (July 1976), pp. 205–209. ISSN: 0304-4068. DOI: 10.1016/0304-4068(76)90029-x.

[120] S. Talari et al. “Mechanism design for decentralized peer-to-peer energy trading considering heterogeneous preferences”. In: *Sustainable Cities and Society* 87 (Dec. 2022), p. 104182. ISSN: 2210-6707. DOI: 10.1016/j.scs.2022.104182.

[121] C. Aperjis and R. Johari. "A peer-to-peer system as an exchange economy". In: *Proceeding from the 2006 Workshop on Game Theory for Communications and Networks*. GameNets '06. Pisa, Italy: Association for Computing Machinery, 2006, 10–es. ISBN: 159593507X. DOI: 10.1145/1190195.1190206.

[122] C. Aperjis et al. "Bilateral and Multilateral Exchanges for Peer-Assisted Content Distribution". In: *IEEE/ACM Transactions on Networking* 19.5 (Oct. 2011), pp. 1290–1303. ISSN: 1558-2566. DOI: 10.1109/tnet.2011.2114898.

[123] A. Saumweber et al. "Designing Pareto optimal electricity retail rates when utility customers are prosumers". In: *Energy Policy* 156 (Sept. 2021), p. 112339. ISSN: 0301-4215. DOI: 10.1016/j.enpol.2021.112339.

[124] M. Ripeanu. "Peer-to-peer architecture case study: Gnutella network". In: *Proceedings First International Conference on Peer-to-Peer Computing*. PTP-01. IEEE Comput. Soc, 2001, pp. 99–100. DOI: 10.1109/p2p.2001.990433.

[125] H. A. Simon. "Bounded Rationality". In: *Utility and Probability*. Palgrave Macmillan UK, 1990, pp. 15–18. ISBN: 9781349205684. DOI: 10.1007/978-1-349-20568-4_5.

[126] U. Malmendier and J. A. Wachter. "Memory of Past Experiences and Economic Decisions". In: *The Oxford Handbook of Human Memory, Two Volume Pack*. Oxford University Press, July 2024, pp. 2228–2266. ISBN: 9780190918019. DOI: 10.1093/oxfordhb/9780190917982.013.78.

[127] N. Neligh. "Rational memory with decay". In: *Journal of Economic Behavior amp; Organization* 223 (July 2024), pp. 120–145. ISSN: 0167-2681. DOI: 10.1016/j.jebo.2024.04.030.

[128] N. Igarashi et al. "A two-phase model of collective memory decay with a dynamical switching point". In: *Scientific Reports* 12.1 (Dec. 2022). ISSN: 2045-2322. DOI: 10.1038/s41598-022-25840-9.

[129] G. Arena et al. "How fast do we forget our past social interactions?: Understanding memory retention with parametric decays in relational event models". In: *Network Science* 11.2 (2023). This work was supported by an ERC Starting Grant (758791)., pp. 267–294. DOI: 10.1017/nws.2023.5.

[130] H. Li et al. *Distributed Generalized Nash Equilibria Seeking Algorithms Involving Synchronous and Asynchronous Schemes*. 2024. DOI: 10.48550/ARXIV.2402.03669.

[131] Y. Wu et al. "P2P energy trading: Blockchain-enabled P2P energy society with multi-scale flexibility services". In: *Energy Reports* 8 (Nov. 2022), pp. 3614–3628. ISSN: 2352-4847. DOI: 10.1016/j.egyr.2022.02.074.

[132] M. Belenkiy et al. "Making p2p accountable without losing privacy". In: *Proceedings of the 2007 ACM workshop on Privacy in electronic society*. CCS07. ACM, Oct. 2007, pp. 31–40. DOI: 10.1145/1314333.1314339.

[133] J. R. Douceur. "The Sybil Attack". In: *Peer-to-Peer Systems*. Ed. by P. Druschel et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 251–260. ISBN: 978-3-540-45748-0.

[134] A. Alhussain et al. "Managing Trust and Detecting Malicious Groups in Peer-to-Peer IoT Networks". In: *Sensors* 21.13 (June 2021), p. 4484. ISSN: 1424-8220. DOI: 10.3390/s21134484.

[135] G. Ciccarelli and R. Lo Cigno. "Collusion in peer-to-peer systems". In: *Computer Networks* 55.15 (Oct. 2011), pp. 3517–3532. ISSN: 1389-1286. DOI: 10.1016/j.comnet.2011.06.028.

[136] J. Aspnes. *Notes on Theory of Distributed Systems*. 2020. DOI: 10.48550/ARXIV.2001.04235.

[137] S. G. Kwak and J. H. Kim. "Central limit theorem: the cornerstone of modern statistics". In: *Korean Journal of Anesthesiology* 70.2 (2017), p. 144. ISSN: 2005-7563. DOI: 10.4097/kjae.2017.70.2.144.

[138] D. C. Matz and E. L. Hause. ““Dealing” with the Central Limit Theorem”. In: *Teaching of Psychology* 35.3 (July 2008), pp. 198–200. ISSN: 1532-8023. DOI: 10.1080/00986280802186201.

[139] F. Celata et al. “The sharing economy as community marketplace? Trust, reciprocity and belonging in peer-to-peer accommodation platforms”. In: *Cambridge Journal of Regions, Economy and Society* 10.2 (Jan. 2017), pp. 349–363. ISSN: 1752-1386. DOI: 10.1093/cjres/rsw044.

[140] C.-K. Chau and K. Elbassioni. “Quantifying Inefficiency of Fair Cost-Sharing Mechanisms for Sharing Economy”. In: (2015). DOI: 10.48550/ARXIV.1511.05270.

[141] P. Erdős and A. Rényi. “On random graphs. I.” In: *Publicationes Mathematicae Debrecen* 6.3–4 (July 2022), pp. 290–297. ISSN: 0033-3883. DOI: 10.5486/pmd.1959.6.3-4.12.

[142] A.-L. Barabasi and R. Albert. “Emergence of scaling in random networks”. In: (1999). DOI: 10.48550/ARXIV.COND-MAT/9910332.

[143] D. J. Watts and S. H. Strogatz. “Collective dynamics of ‘small-world’ networks”. In: *Nature* 393.6684 (June 1998), pp. 440–442. ISSN: 1476-4687. DOI: 10.1038/30918.

[144] M. Humphries et al. “The brainstem reticular formation is a small-world, not scale-free, network”. In: *Proceedings of the Royal Society B: Biological Sciences* 273.1585 (Nov. 2005), pp. 503–511. ISSN: 1471-2954. DOI: 10.1098/rspb.2005.3354.

[145] L. Georgiadis et al. *Exchange of Services in Networks: Competition, Cooperation, and Fairness*. 2015. DOI: 10.48550/ARXIV.1504.02052.

[146] O. Marfoq. “Tackling heterogeneity in federated learning systems”. Theses. Université Côte d’Azur, Dec. 2023. URL: https://theses.hal.science/tel-04498083v1.

[147] S. Chien and A. Sinclair. “Convergence to approximate Nash equilibria in congestion games”. In: *Games and Economic Behavior* 71.2 (Mar. 2011), pp. 315–327. ISSN: 0899-8256. DOI: 10.1016/j.geb.2009.05.004.

[148] E. Sorin et al. *Consensus-based approach to peer-to-peer electricity markets with product differentiation*. 2018. DOI: 10.48550/ARXIV.1804.03521.

[149] M. Schmidt. *Lecture Notes: Numerical Optimization for Machine Learning*. https://www.cs.ubc.ca/~schmidtm/Courses/5XX-S22/S3.pdf. Accessed: 2025-06-26. 2022.

[150] A. Bernales and J. Daoud. “Algorithmic and High Frequency Trading in Dynamic Limit Order Markets”. In: *SSRN Electronic Journal* (2013). ISSN: 1556-5068. DOI: 10.2139/ssrn.2352409.

[151] R. Ao et al. *Asynchronous Gradient Play in Zero-Sum Multi-agent Games*. 2022. DOI: 10.48550/ARXIV.2211.08980.

[152] P. Golle et al. “Incentives for sharing in peer-to-peer networks”. In: *Proceedings of the 3rd ACM conference on Electronic Commerce*. EC01. ACM, Oct. 2001, pp. 264–267. DOI: 10.1145/501158.501193.

[153] C. Feng et al. “Connection-aware P2P trading: Simultaneous trading and peer selection”. In: *Applied Energy* 377 (Jan. 2025), p. 124658. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2024.124658.

[154] T. Temzelides and J. Yu. “Lack-of-recall and centralized monetary trade”. In: *International Economic Review* 45.4 (Oct. 2004), pp. 1221–1227. ISSN: 1468-2354. DOI: 10.1111/j.0020-6598.2004.00303.x.

[155] F. Cavalli and A. Naimzada. “A tâtonnement process with fading memory, stabilization and optimal speed of convergence”. In: *Chaos, Solitons amp; Fractals* 79 (Oct. 2015), pp. 116–129. ISSN: 0960-0779. DOI: 10.1016/j.chaos.2015.05.009.

[156] N. Chen et al. “Incentives for Strategic Behavior in Fisher Market Games”. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 30.1 (Feb. 2016). ISSN: 2159-5399. DOI: 10.1609/aaai.v30i1.10044.

[157] H.-C. Hsiao et al. “Resolving the Topology Mismatch Problem in Unstructured Peer-to-Peer Networks”. In: *IEEE Transactions on Parallel and Distributed Systems* 20.11 (2009), pp. 1668–1681. DOI: 10.1109/TPDS.2009.24.

[158] Q. Lv et al. “Search and replication in unstructured peer-to-peer networks”. In: *ACM SIGMETRICS Performance Evaluation Review* 30.1 (June 2002), pp. 258–259. ISSN: 0163-5999. DOI: 10.1145/511399.511369.

[159] U. Asfia et al. “A Blockchain Construct for Energy Trading against Sybil Attacks”. In: *2019 27th Mediterranean Conference on Control and Automation (MED)*. IEEE, July 2019, pp. 422–427. DOI: 10.1109/med.2019.8798489.

[160] O. Samuel et al. “Towards sustainable smart cities: A secure and scalable trading system for residential homes using blockchain and artificial intelligence”. In: *Sustainable Cities and Society* 76 (Jan. 2022), p. 103371. ISSN: 2210-6707. DOI: 10.1016/j.scs.2021.103371.

[161] M. K. Siam et al. “Securing Decentralized Ecosystems: A Comprehensive Systematic Review of Blockchain Vulnerabilities, Attacks, and Countermeasures and Mitigation Strategies”. In: *Future Internet* 17.4 (Apr. 2025), p. 183. ISSN: 1999-5903. DOI: 10.3390/fi17040183.

[162] J. B. T. Sanders et al. “The prevalence of chaotic dynamics in games with many players”. In: *Scientific Reports* 8.1 (Mar. 2018). ISSN: 2045-2322. DOI: 10.1038/s41598-018-22013-5. URL: http://dx.doi.org/10.1038/s41598-018-22013-5.

[163] A. Beck and M. Teboulle. “Mirror descent and nonlinear projected subgradient methods for convex optimization”. In: *Operations Research Letters* 31.3 (May 2003), pp. 167–175. ISSN: 0167-6377. DOI: 10.1016/s0167-6377(02)00231-6.

[164] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Mar. 2004. ISBN: 9780511804441. DOI: 10.1017/cbo9780511804441.

A

Proofs

A.1. Optimality of Program P5

Proof. To prove the optimality of Program P5 for the exchange market we employ the Lagrangian and the KKT-conditions, to show that each of the conditions of the equilibrium as defined in 1.4.1 holds. The Lagrangian of the objective function of Program P5 is:

$$L(\mathbf{X}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \cdot \log \left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \right) - \sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \log v_j \cdot \varepsilon_j \right)$$
$$+ \sum_{i \in \mathcal{N}} \left(\lambda_i \cdot (\sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i) \right) - \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} (\mu_{ij} \cdot x_{ij})$$

The derivatives are:

$$\frac{\partial L}{\partial x_{ij}} = \left[\sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \log \sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right]' - \left[\sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \log(v_j \cdot \varepsilon_j) \right]'$$
$$+ \left[\sum_{i \in \mathcal{N}} \lambda_i \cdot (\sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i) \right]' - \left[\sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{N}_j} \mu_{ij} \cdot x_{ij} \right]'$$
$$= v_i \cdot (\log \sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} + 1) - v_i \cdot \log(v_j \cdot \varepsilon_j) + \lambda_i - \mu_{ij}$$
$$= v_i \cdot (\log \left(\frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j} \right) + 1) + \lambda_i - \mu_{ij}$$

$$\frac{\partial L}{\partial \lambda_i} = \sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i$$

$$\frac{\partial L}{\partial \mu_{ij}} = -x_{ij}$$

Which makes the KKT-conditions as follows:

- primal feasibility: $\forall j \in \mathcal{N} : \sum_{i \in \mathcal{N}_j} x_{ij} \leq \varepsilon_i$
  $\forall i, j \in \mathcal{N} : x_{ij} \geq 0$
- dual feasibility: $\forall i \in \mathcal{N} : \lambda_i \geq 0$
  $\forall i, j \in \mathcal{N} : \mu_{ij} \geq 0$
- complementary slackness: $\forall i \in \mathcal{N} : \lambda_i \cdot (\sum_{j \in \mathcal{N}_i} x_{ij} - \varepsilon_i) = 0$
  $\forall i, j \in \mathcal{N} : \mu_{ij} \cdot x_{ij} = 0$
- stationarity: $\forall i, j \in \mathcal{N} : v_i(\log\left(\frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j}\right) + 1) + \lambda_i - \mu_{ij} = 0$

So we can derive the following statements:

1. $\forall i \in \mathcal{N} : \lambda_i > 0 \Rightarrow \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i$
   (due to the complementary slackness of λ_i)
2. $\forall i, j \in \mathcal{N} : v_i(\log\left(\frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j}\right) + 1) + \lambda_i - \mu_{ij} = 0$
   $$\Rightarrow (\log\left(\frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j}\right) + 1) - \frac{\mu_{ij}}{v_i} = -\frac{\lambda_i}{v_i}$$
   $$\Rightarrow (\log\left(\frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j}\right) + 1) \geq -\frac{\lambda_i}{v_i}$$
   $$\Rightarrow \frac{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j} \geq e^{-\lambda_i/v_i - 1}$$
   $$\Rightarrow \frac{v_j \cdot \varepsilon_j}{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}} \leq e^{1+\lambda_i/v_i}$$
3. $\forall i, j \in \mathcal{N} : x_{ij} > 0 \Rightarrow \mu_{ij} = 0$
   $$\Rightarrow v_i(\log\left(\frac{\sum_{i \in \mathcal{N}} v_i \cdot x_{ij}}{v_j \cdot \varepsilon_j}\right) + 1) + \lambda_i = 0$$
   $$\Rightarrow \frac{v_j \cdot \varepsilon_j}{\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij}} = e^{1+\lambda_i/v_i}$$
   (due to the complementary slackness of μ_{ij})

Given the assumption that for each resource i, $v_i > 0$, we can derive that $\forall i \in \mathcal{N}, \lambda_i > 0$, because if agents value a resource i positively and it were free (meaning $\lambda_i = 0$), then each agent would want to obtain that resource infinitely much, automatically driving up the value for λ_i. So then, condition 1 above tells us that $\forall i \in \mathcal{N}, \sum_{j \in \mathcal{N}_i} x_{ij} = \varepsilon_i$, thus proving market clearance.

For proving the utility maximisation we denote with β_j the inverse best-bang-per-buck, i.e., how much value agent j has had to distribute in order to receive back value: $\beta_j = (v_j \cdot \varepsilon_j)/(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij})$. Given the derived statements 2 and 3, we know that if $x_{ij} > 0$ then $e^{1+\lambda_i/v_i} = \beta_j \leq e^{1+\lambda_k/v_k}$. We can regard $e^{1+\lambda_i/v_i}$ as a(n) (increasing) function over the relative cost per value λ_i/v_i of agent i, meaning that if $e^{1+\lambda_i/v_i} > e^{1+\lambda_k/v_k}$ then $\lambda_i/v_i > \lambda_k/v_k$. We can see thus that, agent j receives resources from neighbour i ($x_{ij} > 0$) only if neighbour i has the smallest relative cost per value, compared to other neighbours k: $\forall i, k, \in \mathcal{N}_j : \lambda_i/v_i \leq \lambda_k/v_k$. This means that there is no other agent k that could give j a better value per cost than what it is obtaining now. This means that j is achieving the best possible value-received to value-allocated ratio, thus proving the utility maximisation condition of the equilibrium.

With these two conditions proven, we have shown definitively that Program P5 optimises the exchange market according to equilibrium definition 1.4.1.

□

A.2. Optimality of Program P6

This appendix proves optimality of Program P6 for the market described in 1.4.

Proof. Taking the third condition of the program and multiplying both sides by $x_{ji} \cdot \rho_j$ produces:

$$\forall i \in \mathcal{N}, j \in \mathcal{N}_i : v_j \cdot x_{ji} \leq \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{\rho_i \cdot \varepsilon_i} \cdot x_{ji} \cdot \rho_j$$

Summing over j in i's neighbourhood gives:

$$\forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji} \leq \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{\rho_i \cdot \varepsilon_i} \cdot \sum_{j \in \mathcal{N}_i} x_{ji} \cdot \rho_j$$

We know, given the earlier assumption that $\forall i \in \mathcal{N} : v_i > 0$, that the denominator of this will not be zero. As such, we can rewrite:

$$\forall i \in \mathcal{N} : \rho_i \cdot \varepsilon_i \leq \sum_{j \in \mathcal{N}_i} x_{ji} \cdot \rho_j$$

Summing over i gives us:

$$\sum_{i \in \mathcal{N}} \rho_i \cdot \varepsilon_i \leq \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}_i} \rho_j \cdot x_{ji}$$

Given the inherent symmetry of neighbourhoods

$$\forall i, j \in \mathcal{N} : j \in \mathcal{N}_i \Leftrightarrow i \in \mathcal{N}_j$$

we can write

$$\sum_{i \in \mathcal{N}} \rho_i \cdot \varepsilon_i \leq \sum_{j \in \mathcal{N}} \rho_j \cdot \sum_{i \in \mathcal{N}_j} x_{ji} \Rightarrow$$
$$\sum_{i \in \mathcal{N}} \rho_i \cdot \varepsilon_i \leq \sum_{j \in \mathcal{N}} \rho_j \cdot \varepsilon_j$$

Given that this inequality clearly reduces to an equality, all the previous inequalities must have clearly been equalities as well. This shows that whenever x_{ji} is strictly positive, the third constraint in program P6 is an equality as well, which means that at equilibrium, an agent allocates to those neighbours that give the maximum bang-per-buck, proving the utility maximisation constraint of the equilibrium. In combination with the first constraint of the program, which guarantees market clearance, this proves optimality for the exchange market. □

A.3. Deriving π from Program P6

Given the third constraint of Program P6

$$\forall i \in \mathcal{N}, \forall j \in \mathcal{N}_i : \frac{v_j}{\rho_j} \leq \frac{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}}{\rho_i \cdot \varepsilon_i} \tag{A.1}$$

we can derive strategy π. Constraint A.1 can be interpreted as the utility-maximisation clause. To derive π, we assume that we are not at the equilibrium, and thus clause A.1 is violated for some agent i. This means that there is some error ϵ such that $\epsilon = \max\{\epsilon_j\} = \max\{v_j/\rho_j - (\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki})/(\rho_i \cdot \varepsilon_i)\}_{j \in \mathcal{N}_i}$. This market state with $\epsilon > 0$ is 'invalid' according to Program P6. From such an invalid state, we then aim to make a step towards a valid market state (i.e., the equilibrium).

To make such a step, agent i must make an allocation that decreases ϵ. Given the inequality in clause A.1 and agent i's action space consisting of allocation vector $\mathbf{x}_i$, increasing x_{ij} for neighbour j with the largest violation ϵ_j will decrease ϵ the most, and thus make the most progress toward satisfying the clause A.1 and finding the equilibrium. The right-hand side of clause A.1 does not depend on the choice of neighbour j, however, the left-hand side, v_j/ρ_j, does depend on j. Therefore, neighbour j with the highest value of v_j/ρ_j is the one that causes has largest violation ϵ_j.

We note that we can scale all v_i w.l.o.g. by inversely scaling up all ε_i. With this scaling, we can infer that choosing the neighbour j with the highest v_j/ρ_j, reduces to the strategy of choosing the neighbour with the lowest ρ_j, which is exactly the strategy employed in π. So, employing this distributed greedy strategy for the exchange market, implies the direct optimisation of the utility-maximising constraint of Program P6.

A.4. Mirror Descent of ϕ

Proof. In order to prove Theorem 4.3.1, and show a mirror-descent-based connection between Program P5 and ϕ, we first confirm that Program P5 is a convex program. This involves showing that both the solution space and the objective function are convex. The solution space is clearly convex, due to the linear constraints. The objective on the other hand requires a more in-depth analysis.

Lemma A.4.1 (Convexity of Program P5 in $\mathbf{X}$). *Program P5 is convex in the space of allocations $\mathcal{X}$.*

Lemma A.4.2 (Strict Convexity of Program P5 in $\mathbf{U}$). *Program P5 is strictly convex in the space of utilities $\mathcal{U}$.*

In Appendix A.5 we present the proof for both Lemma A.4.1 and Lemma A.4.2, showing both that the objective function is convex in $\mathbf{X}$ and that it is strictly convex in $\mathbf{U}$ and thus computes the unique equilibrium vector $\mathbf{U}^*$. We employ the same technique for both: splitting the objective function into a left-hand side and a right-hand side and showing the curvature for each side separately.

Given that Program P5 is convex, we can construct a type of mirror descent algorithm. For this we use the Subgradient Algorithm with Nonlinear Projection (SANP), which is equivalent to the mirror descent algorithm [163]. Before defining the SANP, we convert the solution space to the unit-simplex, to obtain a representation of the space that aligns with convential SANP spaces. To achieve this, we scale all v to 1, by inversely scaling each ε. Then, we scale all ε such that $\sum_{i \in \mathcal{N}} \varepsilon_i = 1$. This scaling is w.l.o.g. as the relative worth of each agent's endowment remains the same. Each descent step in SANP, over allocation space $\mathcal{X}$, looks as follows:

$$\mathbf{X} = \arg\min_{\mathbf{Y} \in \mathcal{X}} \{\langle \mathbf{Y}, f'(\mathbf{X}) \rangle + B_\varphi(\mathbf{Y}, \mathbf{X})\}$$

where the distance function B_φ is:

$$B_\varphi(\mathbf{a}, \mathbf{b}) = \varphi(\mathbf{a}) - \varphi(\mathbf{b}) - \langle \mathbf{a} - \mathbf{b}, \nabla\varphi(\mathbf{b}) \rangle$$

defined on some strongly convex, continuously differentiable function φ. The choice of φ can change depending on the geometry of the problem at hand. $f(\mathbf{X})$ is defined as the objective function of Program P5, thus our problem is convex and for φ we choose the entropy function:

$$\varphi(\mathbf{X}) = \sum_{i,j \in \mathcal{N}} x_{ij} \log x_{ij}$$

where $\mathbf{X} \notin \mathcal{X} \Rightarrow \varphi(\mathbf{X}) = \infty$ and $\mathbf{X} = \mathbf{0} \Rightarrow \varphi(\mathbf{X}) = 0$, to ensure continuous differentiability. This choice of φ results in Entropic Mirror Descent. The update step becomes:

$$\mathbf{X} = \arg\min_{\mathbf{Y}} \left\{ \sum_{i,j \in \mathcal{N}} y_{ij} \cdot f'_{ij}(\mathbf{X}) + \sum_{i,j \in \mathcal{N}} y_{ij} \log\left(\frac{y_{ij}}{x_{ij}}\right) \right\} \tag{A.2}$$

Assuming that the non-negativity constraint ($\forall i, j \in \mathcal{N} : y_{ij} \geq 0$) and the packing constraint ($\forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} y_{ij} \leq \varepsilon_i$) still hold, we write the Lagrangian to find $\mathbf{Y}$:

$$L(\mathbf{X}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_{i,j \in \mathcal{N}} \left[y_{ij} \cdot f'_{ij}(\mathbf{X}) + y_{ij} \cdot \log\left(\frac{y_{ij}}{x_{ij}}\right) \right] + \sum_{i \in \mathcal{N}} \lambda_i \cdot \left(\sum_{j \in \mathcal{N}_i} y_{ij} - 1 \right) - \sum_{i,j \in \mathcal{N}} (\mu_{ij} \cdot y_{ij})$$

where f'_{ij} is $\partial f / \partial y_{ij}$. The complementary slackness condition and stationarity condition of the KKT-conditions become:

$$\forall i, j \in \mathcal{N} : \mu_{ij} \cdot y_{ij} = 0$$

$$\forall i, j \in \mathcal{N} : f'_{ij}(\mathbf{X}) + \left(\log\left(\frac{y_{ij}}{x_{ij}}\right) + 1 \right) + \lambda_i - \mu_{ij} = 0$$

Provided that we initialise with a point in the interior of the allocation space ($\mathbf{Y}(1) \in \mathcal{X}^{\mathrm{o}}$), we know that $\forall i, j \in \mathcal{N} : y_{ij} > 0$, and thus, due to complementary slackness, $\forall i, j \in \mathcal{N} : \mu_{ij} = 0$. Rewriting the stationarity condition thus becomes:

$$y_{ij} = x_{ij} \cdot e^{-f'_{ij}(\mathbf{X}) - \lambda_i - 1}$$

Defining $Z_i = e^{\lambda_i + 1}$, we write:

$$y_{ij} = \frac{x_{ij} \cdot e^{-f'_{ij}(\mathbf{X})}}{Z_i} \tag{A.3}$$

Given that $\sum_{j \in \mathcal{N}_i} y_{ij} = \varepsilon_i$, we obtain the following:

$$\sum_{j \in \mathcal{N}_i} y_{ij} = \sum_{j \in \mathcal{N}_i} \frac{x_{ij} \cdot e^{-f'_{ij}(\mathbf{X})}}{Z_i} = \frac{\sum_{j \in \mathcal{N}_i} x_{ij} \cdot e^{-f'_{ij}(\mathbf{X})}}{Z_i} = \varepsilon_i$$

meaning that

$$Z_i = \frac{1}{\varepsilon_i} \sum_{j \in \mathcal{N}_i} x_{ij} \cdot e^{-f'_{ij}(\mathbf{X})} \tag{A.4}$$

So with equation A.3 and equation A.4, we can then write:

$$x_{ij} = \varepsilon_i \cdot \frac{x_{ij} \cdot e^{-f'_{ij}(\mathbf{X})}}{\sum_{k \in \mathcal{N}_i} x_{ik} \cdot e^{-f'_{ik}(\mathbf{X})}} \tag{A.5}$$

Now we can start filling in f'_{ij}. For f with all v scaled to 1 we have:

$$f(\mathbf{X}) = \sum_{j\in\mathcal{N}} \sum_{i\in\mathcal{N}_j} x_{ij} \log \sum_{i\in\mathcal{N}_j} x_{ij} - \sum_{i,j\in\mathcal{N}} x_{ij} \log \varepsilon_j$$

and $f'_{ij}(\mathbf{X}) = 1 - \log\Big(\varepsilon_j / \sum_{k\in\mathcal{N}_j} x_{kj}\Big) = 1 - \log(\varepsilon_j / R_j)$. Filling this definition of f'_{ij} into equation A.5:

$$y_{ij} = \varepsilon_i \cdot \frac{x_{ij} \cdot e^{-1} \cdot (\varepsilon_j / \sum_{k\in\mathcal{N}_j} x_{kj})}{\sum_{k\in\mathcal{N}_i} x_{ik} \cdot e^{-1} \cdot (\varepsilon_k / \sum_{l\in\mathcal{N}_k} x_{lk})} = \varepsilon_i \cdot \frac{\varepsilon_j \cdot (x_{ij}/R_j)}{\sum_{k\in\mathcal{N}_i} \varepsilon_k \cdot (x_{ik}/R_k)} \tag{A.6}$$

Knowing that each agent j computes its allocation x_{ij} to agent i proportionally to allocation x_{ij} and its received resources R_j, we can clearly recognise the proportional allocation rule of strategy ψ here. □

A.5. Convexity of Program P5

Convexity w.r.t. x_{ij}

Proof. This proof will show that the objective function of Program P5 is convex w.r.t. variable x_{ij}. Showing this convexity can be done by separating the objective of Program P5 into the right-hand side (RHS) and the left-hand side (LHS):

- **RHS**: $\sum_{j\in\mathcal{N}} \left(\left(\sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}\right) \log v_j \cdot \varepsilon_j\right)$
- **LHS**: $\sum_{j\in\mathcal{N}} \left(\left(\sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}\right) \cdot \log \left(\sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}\right)\right)$

The RHS of the objective function is clearly linear in x_{ij}, so this proof will consider only the LHS, with the knowledge that subtracting a linear function from a convex function still results in a convex function. In this proof, we ease notation by defining a function on x as $s_j(\mathbf{X}) = \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}$.

For the LHS, just proving convexity for one entry $LHS_j = \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij} \log \left(\sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}\right)$ suffices, considering that the addition of multiple convex functions is still convex. There is, however, interaction between several x-variables in this function, so we construct the Hessian matrix to determine convexity. The first derivative of this w.r.t. any x_{ij}:

$$\begin{aligned}
\frac{\partial LHS_j}{\partial x_{ij}} &= \frac{\partial \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij} \log \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}}{\partial x_{ij}} \\
&= \frac{\partial s_j}{\partial x_{ij}} \cdot \log(s_j) + s_j \cdot \frac{\partial \log(s_j)}{\partial x_{ij}}
\end{aligned}$$

With the chain rule again we have:

$$\frac{\partial \log(s_j)}{\partial x_{ij}} = \frac{1}{s_j} \cdot \frac{\partial s_j}{\partial x_{ij}} = \frac{1}{s_j} \cdot v_i$$

So we get the first derivative:

$$\begin{aligned}
\frac{\partial LHS_j}{\partial x_{ij}} &= \frac{\partial s_j}{\partial x_{ij}} \cdot \log(s_j) + s_j \cdot \frac{\partial \log(s_j)}{\partial x_{ij}} \\
&= v_i + \log(s_j) + s_j \cdot \frac{1}{s_j} \cdot v_i \\
&= v_i \cdot \log(s_j) + v_i
\end{aligned}$$

For the second derivative, w.r.t. x_{kl}, we have to consider that it is possible that $k \neq i$ (j is fixed by choice of LHS_j):

$$i = k \Rightarrow \frac{\partial^2 LHS_j}{\partial x_{kj}^2} = \frac{\partial(v_i \cdot \log(s_j) + v_i)}{\partial x_{kj}} = v_i \cdot \frac{1}{s_j} \cdot v_i = \frac{v_i^2}{s_j}$$
$$i \neq k \Rightarrow \frac{\partial^2 LHS_j}{\partial x_{kj}^2} = \frac{\partial(v_i \cdot \log(s_j) + v_i)}{\partial x_{kj}} = v_i \cdot \frac{1}{s_j} \cdot v_k = \frac{v_i \cdot v_k}{s_j}$$

So we essentially have the rule that constructs a Hessian matrix $\mathbf{M}$, with some fixed j, as follows:

$$\mathbf{H}^{(j)} = [\nabla^2 LHS_j]_{(i,j),(k,l)} = \frac{v_i \cdot v_k}{s_j} \tag{A.7}$$

To prove that $\mathbf{H}^{(j)}$ is indeed positive semi-definite, we have to find the eigenvalues. We know that $\mathbf{H}^{(j)} = (1/s_j) \cdot \mathbf{v}\mathbf{v}^\top$, where $\mathbf{v}$ is the vector of values v_i. Any matrix of the form $\mathbf{v}\mathbf{v}^\top$ has a single nonzero eigenvalue $\lambda_1 = ||\mathbf{v}||^2$, and all other eigenvalues $\lambda_2 = ... = \lambda_n = 0$. Multiplying by $1/s_j$ gets us the following set of eigenvalues for matrix $\mathbf{H}^{(j)}$:

$$\{\frac{||\mathbf{v}||^2}{s_j}, 0, ..., 0\}$$

Each of these eigenvalues is non-negative, and thus the Hessian matrix $\mathbf{H}^{(j)}$ proves the convexity of a single entry LHS_j of the LHS.

Given the convexity of this single entry LHS_j, we already know that the entire summation of LHS is also convex. Alternatively, we can also reason about this from the perspective of the full Hessian matrix $\mathbf{H}$. The matrices $\mathbf{H}^{(j)}$ constructed by rule A.7 are actually blocks ordered along the diagonal of $\mathbf{H}$.

The eigenvalues of a block-diagonal matrix are the union of the eigenvalues of its diagonal blocks. So we get the set of eigenvalues:

$$\{\frac{||\mathbf{v}||^2}{s_1}, 0, ..., 0, ..., \frac{||\mathbf{v}||^2}{s_n}, 0, ..., 0\}$$

Which, again, are all non-negative, and thus prove the convexity of the entire LHS w.r.t. x_{ij}, and by extension, of the entire objective function.

□

Strict convexity w.r.t. U

Proof. This proof will show the strict convexity of the objective of Program P5 w.r.t. $\mathbf{U}$. Each i^{th} entry of $\mathbf{U}$ is the result of that agent i's utility function, as follows:

$$U_i = u_i(\mathbf{r}_i) = \sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji}$$

This is exactly the definition of $s_i = \sum_j v_j \cdot x_{ji}$ as used before.

Showing this convexity can be done by separating the objective of Program P5 into the right-hand side (RHS) and the left-hand side (LHS):

- **RHS**: $\sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \log v_j \cdot \varepsilon_j \right)$
- **LHS**: $\sum_{j \in \mathcal{N}} \left(\left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \cdot \log \left(\sum_{i \in \mathcal{N}_j} v_i \cdot x_{ij} \right) \right)$

proving linearity for the former and convexity for the latter. A linear function being subtracted from a convex function results in a convex function. For the RHS we can clearly see the linearity w.r.t. s:

$$\sum_{i,j\in\mathcal{N}} v_i \cdot x_{ij} \log(v_j \cdot \varepsilon_j) = \sum_{j\in\mathcal{N}} \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij} \log(v_j \cdot \varepsilon_j) = \sum_{j\in\mathcal{N}} s_j \log(v_j \cdot \varepsilon_j)$$

For the LHS, we will have to find the derivative w.r.t. $s_j = \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}$. No Hessian matrix $\mathbf{H}$ seems necessary for this convexity proof, as there is no interaction between s_i and s_j $\forall i, j \in \mathcal{N}$, i.e. the LHS is separable w.r.t. all s. As a result, all non-diagonal entries of $\mathbf{H}$ are guaranteed to be 0, creating a diagonal matrix. The eigenvalues of diagonal matrices are simply the entries along the diagonal, meaning that, in order to prove that $\mathbf{H}$ is positive definite, we only have to prove that the entries along the diagonal are positive, i.e. $\forall j \in \mathcal{N}, \frac{\partial^2 LHS_j}{\partial s_j^2} > 0$.

The first derivative of $LHS_j = \sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij} \log\left(\sum_{i\in\mathcal{N}_j} v_i \cdot x_{ij}\right) = s_j \log s_j$ w.r.t. s_j is:

$$\begin{aligned}\frac{\partial LHS_j}{\partial s_j} &= \frac{\partial s_j}{\partial s_j} \cdot \log(s_j) + s_j \cdot \frac{\partial \log(s_j)}{\partial s_j} \\ &= \log(s_j) + s_j \cdot \frac{1}{s_j} \\ &= \log(s_j)\end{aligned}$$

The second derivative then is:

$$\frac{\partial^2 LHS_j}{\partial s_j^2} = \frac{\partial \log(s_j)}{\partial s_j} = \frac{1}{s_j}$$

Again, given $\forall j \in \mathcal{N} : s_j \geq 0$, we know that this is strictly positive $\frac{1}{s_j} > 0$, and thus we have strict convexity of the LHS w.r.t. s_j (and by extension of the entire objective function).

□

A.6. Expected Utility for ψ

Allocation strategy π is based on the assumption that if an agent i allocates the value $v_i \cdot x_{ij}$ to agent j, then the next timeslot $t+1$ agent i expects back value $(v_i \cdot x_{ij}(t))/\rho_j(t)$. This is based on the historical trading behaviour of its neighbours. We will show in this appendix that including a mixed strategy into this expectation does not change its expectation, even in a mixed setting where agents may have $0 < w < 1$.

We denote the utility agent i expects back from agent j with $\hat{U}_{i\leftarrow j}$, based on agent j's mixed strategy, as follows:

$$\hat{U}_{i\leftarrow j}(t+1) = w_j \cdot \hat{U}^{\pi}_{i\leftarrow j}(t+1) + (1-w_j) \cdot \hat{U}^{\phi}_{i\leftarrow j}(t+1)$$

Which, if expanded, with $\hat{U}^{\pi}_{i\leftarrow j} = v_j \cdot x^{\pi}_{ji}$ and $\hat{U}^{\phi}_{i\leftarrow j} = v_j \cdot x^{\phi}_{ji}$ becomes the following:

$$\begin{aligned}\hat{U}_{i\leftarrow j}(t+1) &= w_j \cdot (v_j \cdot x^{\pi}_{ji}(t+1)) + (1-w_j) \cdot v_j \cdot x^{\phi}_{ji}(t+1) \\ &= w_j \cdot \left(\frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}\right) + (1-w_j) \cdot \left(v_j \cdot D_j(t+1) \cdot \frac{v_i \cdot x_{ij}(t)}{\sum_{k\in\mathcal{N}_j} v_k \cdot x_{kj}(t)}\right)\end{aligned}$$

which combines both the expected returns of greedy allocation strategy π on the left-hand side, where agents expect back the value of $x_{ij} \cdot v_i/\rho_j$, and the proportional value returned from each neighbour $v_j \cdot D_j(t+1) \cdot (v_i \cdot x_{ij}/\sum_{k\in\mathcal{N}_j} v_k \cdot x_{kj})$ on the right-hand side.

An issue with the above formula arises for the proportional part of this prediction, where agent i has to predict how agent j responds proportionally in the next timeslot $t+1$, which is based on the allocations of the other agents $\sum_{k\in\mathcal{N}_j} v_k \cdot x_{kj}(t)$ in the current timeslot t, and the amount of j's distributable resources in the next timeslot $D_j(t+1)$. This means that agent i is trying to base its allocation of the current timeslot t on allocations of others in timeslot t, which have not yet taken place, and on j's distributable resource of the next timeslot, which has not been determined yet either. This is clearly not possible. To make the formula feasible again, we replace $\sum_{k\in\mathcal{N}_j} v_k \cdot x_{kj}(t)$ with an approximation based on its historical average $\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t)$, and we do the same for approximating $D_j(t+1)$, which we approximate with j's endowment ε_j. The actual expectation of returned value from agent j to agent i thus becomes:

$$\hat{U}_{i\leftarrow j}(t+1) = w_j \cdot \left(\frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}\right) + (1-w_j) \cdot \left(v_j \cdot \varepsilon_j \cdot \frac{v_i \cdot x_{ij}(t)}{\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t)}\right) \tag{A.8}$$

Now we will show how the right-hand side $v_j \cdot \varepsilon_j \cdot (v_i \cdot x_{ij}(t))/(\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t))$ reduces to $(v_i \cdot x_{ij}(t))/\rho_j(t)$, and thus by extension how formula A.8 reduces to the original expectation of $(v_i \cdot x_{ij}(t))/\rho_j(t)$.

We know that $\rho_j(t) = (\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t))/(v_j \cdot \varepsilon_j)$ by its definition, so $1/\rho_j(t) = (v_j \cdot \varepsilon_j)/(\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t))$. So we can write

$$v_j \cdot \varepsilon_j \cdot \frac{v_i \cdot x_{ij}(t)}{\sum_{k\in\mathcal{N}_j} v_k \cdot \overline{x_{kj}}(t)} = \frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}$$

As such, we can rewrite equation A.8 to be

$$\begin{aligned}\hat{U}_{i\leftarrow j}(t+1) &= w_j \cdot \left(\frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}\right) + (1-w_j) \cdot \left(\frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}\right) \\ &= \frac{v_i \cdot x_{ij}(t)}{\rho_j(t)}\end{aligned}$$

We can thus conclude that the original rational expectation of returned utility, that strategy π is based on, holds even in the mixed case.

A.7. Mixed Strategy Convergence

Proof. Here we will show that if two strategies, such as π and ϕ converge to the equilibrium of a market, then so will a linearly combined strategy ψ, such as equation6.1. To prove Theorem 6.1.1, we first define the following convex function on the utilities $\mathbf{U}(t)$ and equilibrium utilities $\mathbf{U}^*$ in a market:

$$V(t) = (\mathbf{U}(t) - \mathbf{U}^*)^2 = \sum_{i\in\mathcal{N}} (U_i(t) - U_i^*)^2 \tag{A.9}$$

For this function we can state that it is non-negative ($V(t) \geq 0$). Additionally when the function equals 0, then the found utilities are the equilibrium utilities ($V(t) = 0 \Leftrightarrow \mathbf{U}(t) = \mathbf{U}^*$), and thus $V(t) = 0$ implies that an equilibrium is found.

We know by definition that we compute the received utility U_i as such:

$$\begin{aligned}U_i(t) &= u_i(\mathbf{r}_i(t)) \\ &= u_i([x_{ji}(t)]_{j\in\mathcal{N}})\end{aligned}$$

Filling in the mixed allocation rule 6.1 to compute the allocations x, we get

$$U_i(t) = u_i(\Big[w_i \cdot x^{\pi}_{ji}(t) + (1-w_i) \cdot x^{\phi}_{ji}(t)\Big]_{j\in\mathcal{N}})$$

Now we make a non-trivial assumption: We assume that all w have the same value: $\forall i, j \in \mathcal{N}$: $w_i = w_j$. The consequence of this assumption is that from here on out the proof is only valid for population in which each agent has the same value for w: homogeneous populations. This assumption allows us to write the following:

$$\begin{aligned}
U_i(t) &= u_i(\Big[w \cdot x^{\pi}_{ji}(t) + (1-w) \cdot x^{\phi}_{ji}(t)\Big]_{j\in\mathcal{N}}) \\
&= w \cdot u_i(\big[x^{\pi}_{ji}(t)\big]_{j\in\mathcal{N}}) + (1-w) \cdot u_i(\Big[x^{\phi}_{ji}(t)\Big]_{j\in\mathcal{N}}) \\
&= w \cdot u_i(\mathbf{r}^{\pi}_i(t)) + (1-w) \cdot u_i(\mathbf{r}^{\phi}_i(t)))
\end{aligned}$$

Filling this definition of $U_i(t)$ back into equation A.9 gives us

$$V(t) = \sum_{i\in\mathcal{N}} (U_i(t) - U^*_i)^2 = \sum_{i\in\mathcal{N}} (w \cdot u_i(\mathbf{r}^{\pi}_i(t)) + (1-w) \cdot u_i(\mathbf{r}^{\phi}_i(t)) - U^*_i)^2 \tag{A.10}$$

For any convex function $f(\cdot)$ we know that $f(w \cdot a + (1-w) \cdot b) \leq w \cdot f(a) + (1-w) \cdot f(b)$ [164]. Thus, given equation A.10 and the convexity of equation A.9, we write

$$\begin{gathered}
\sum_{i\in\mathcal{N}} \Big(w \cdot u_i(\mathbf{r}^{\pi}_i(t)) + (1-w) \cdot u_i(\mathbf{r}^{\phi}_i(t)) - U^*_i\Big)^2 \\
\leq \\
w \cdot \sum_{i\in\mathcal{N}} (u_i(\mathbf{r}^{\pi}_i(t)) - U^*_i)^2 + (1-w) \cdot \sum_{i\in\mathcal{N}} \Big(u_i(\mathbf{r}^{\phi}_i(t)) - U^*_i\Big)^2 \\
= \\
w \cdot V_\pi(t) + (1-w) \cdot V_\phi(t)
\end{gathered}$$

With these (in)equalities, and equation A.10 we write

$$V(t) \leq w \cdot V_\pi(t) + (1-w) \cdot V_\phi(t)$$

If we then average both sides over T timeslots we obtain:

$$\frac{1}{T}\sum_{t=1}^{T} V(t) \leq w \cdot \frac{1}{T}\sum_{t=1}^{T} V_\pi(t) + (1-w) \cdot \frac{1}{T}\sum_{t=1}^{T} V_\phi(t) \tag{A.11}$$

Given the assumption that both π and ϕ converge to equilibrium in the ergodic sense, as established in Section 4.3, we have have in the limit $T \to \infty$ the following

$$\frac{1}{T}\sum_{t=1}^{T} V_\pi(t) \to 0 \frac{1}{T}\sum_{t=1}^{T} V_\phi(t) \to 0$$

This means that the right-hand side of inequality A.11 converges to 0 as $T \to \infty$. As such, we can state that

$$\begin{aligned}
&\frac{1}{T}\sum_{t=1}^{T} V(t) \to 0 \Rightarrow \\
&\frac{1}{T}\sum_{t=1}^{T}\sum_{i} (U_i(t) - U^*_i)^2 \to 0
\end{aligned} \tag{A.12}$$

We apply Jensen's inequality[1], which states that $f(\mathbb{E}[x]) \leq \mathbb{E}[f(x)]$, and obtain:

$$
\begin{aligned}
&(\mathbb{E}[\mathbf{U}(t)] - \mathbf{U}^*)^2 \leq \mathbb{E}[(\mathbf{U}(t) - \mathbf{U}^*)^2] \Rightarrow \\
&(\overline{\mathbf{U}}(t) - \mathbf{U}^*)^2 \leq \frac{1}{T} \sum_{t=1}^{T} (\mathbf{U}(t) - \mathbf{U}^*)^2 \Rightarrow \\
&\sum_{i \in \mathcal{N}} \left(\overline{U_i}(t) - U_i^*\right)^2 \leq \frac{1}{T} \sum_{t=1}^{T} \sum_{i \in \mathcal{N}} (U_i(t) - U_i^*)^2
\end{aligned}
$$

Given equation A.12, we can thus write

$$
\begin{aligned}
&\frac{1}{T} \sum_{t=1}^{T} \sum_{i \in \mathcal{N}} (U_i(t) - U_i^*)^2 \to 0 \Rightarrow \\
&\sum_{i \in \mathcal{N}} \left(\overline{U_i}(t) - U_i^*\right)^2 \to 0 \\
&\overline{\mathbf{U}}(t) \to \mathbf{U}^*
\end{aligned}
$$

So now we have shown that indeed the utilities $\mathbf{U}(t)$ converge to equilibrium utilities $\mathbf{U}^*$ in the ergodic sense. This indicates that a linearly combined rule, such as 6.1, converges to the equilibrium if the components (such as π and ϕ) also converge to the equilibrium[2]. □

[1] https://www.statlect.com/fundamentals-of-probability/Jensen-inequality.

[2] Provided that w is equal across the population

B

Supplemental Experimental Results

B.1. Mixed Strategies over Various Networks

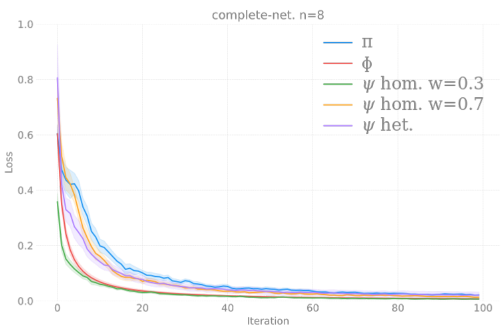

(a) $n = 8$, $T = 100$

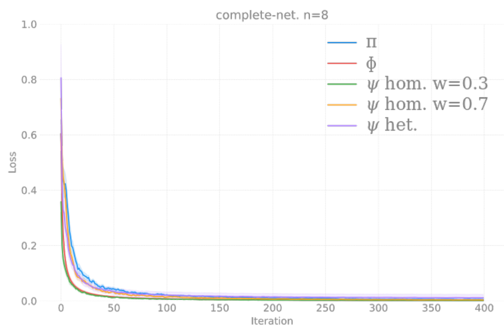

(b) $n = 8$, $T = 400$

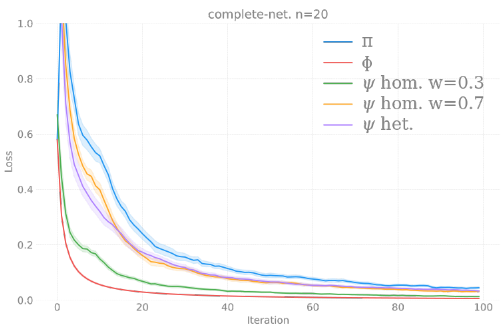

(c) $n = 20$, $T = 100$

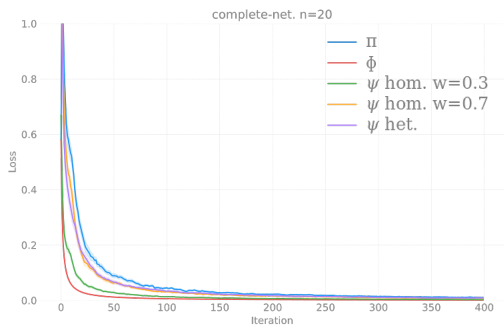

(d) $n = 20$, $T = 400$

Figure B.1: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom.* *w*=0.3 and ψ *hom.* *w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a complete network of size $n = \{8, 20\}$ for $T = \{100, 400\}$ iterations.

		π	ϕ	hom. mixed $w = 0.3$	hom. mixed $w = 0.7$	het. mixed
$n = 8$	$T = 100$	0.02143	0.00757	0.00695	0.01371	0.02084
	$T = 400$	0.00519	0.00190	0.00158	0.00368	0.01151
$n = 20$	$T = 100$	0.04484	0.00618	0.01307	0.03189	0.03327
	$T = 400$	0.01040	0.00155	0.00308	0.00732	0.00872

Table B.1: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a complete network of sizes $n = \{8, 20\}$ after $T = \{100, 400\}$ iterations.

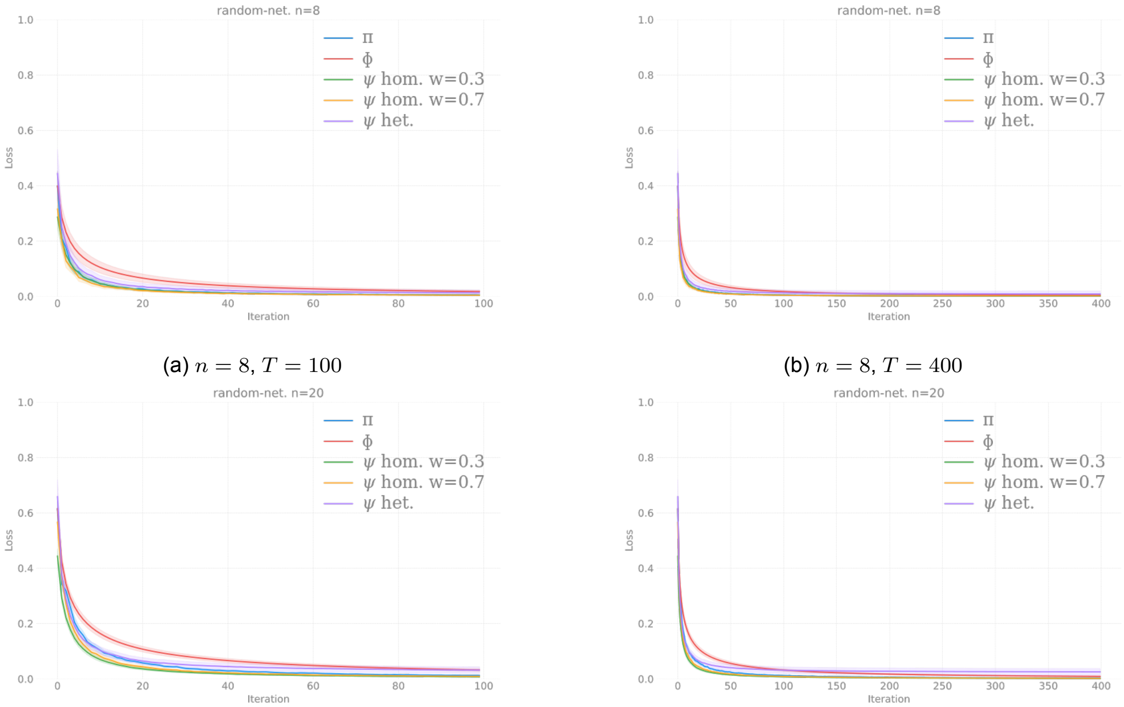

(a) $n = 8$, $T = 100$ (b) $n = 8$, $T = 400$

(a) $n = 20$, $T = 100$ (b) $n = 20$, $T = 400$

Figure B.2: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a random network of size $n = \{8, 20\}$ for $T = \{100, 400\}$ iterations.

		π	ϕ	hom. mixed $w = 0.3$	hom. mixed $w = 0.7$	het. mixed
$n = 8$	$T = 100$	0.00483	0.01756	0.00476	0.00447	0.01296
	$T = 400$	0.00134	0.00472	0.00118	0.00109	0.00881
$n = 20$	$T = 100$	0.01260	0.03219	0.00749	0.00920	0.03188
	$T = 400$	0.00322	0.00918	0.00185	0.00238	0.02554

Table B.2: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a random network of sizes $n = \{8, 20\}$ after $T = \{100, 400\}$ iterations.

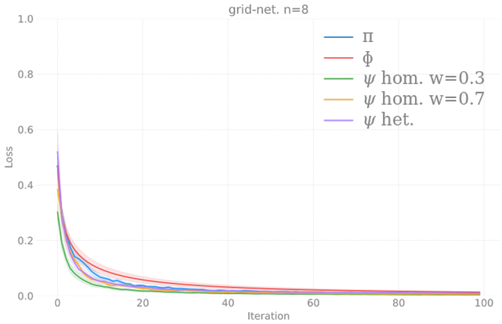

(a) $n = 8$, $T = 100$

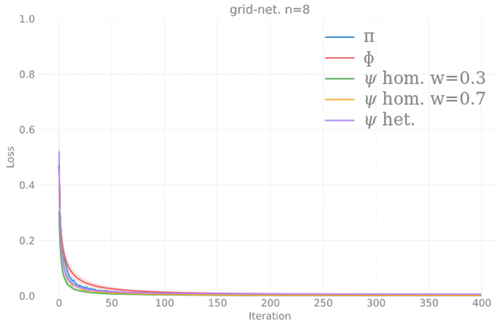

(b) $n = 8$, $T = 400$

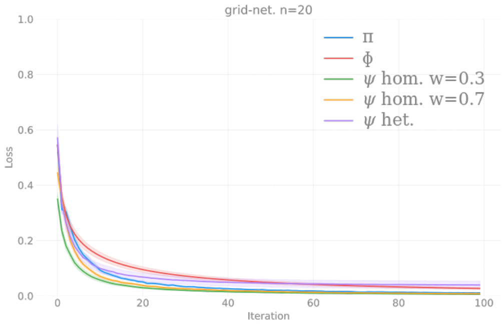

(a) $n = 20$, $T = 100$

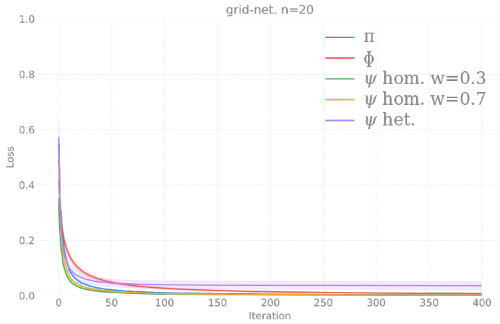

(b) $n = 20$, $T = 400$

Figure B.3: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a grid network of size $n = \{8, 20\}$ for $T = \{100, 400\}$ iterations.

		π	ϕ	hom. mixed $w = 0.3$	hom. mixed $w = 0.7$	het. mixed
$n = 8$	$T = 100$	0.00797	0.01321	0.00381	0.00542	0.00997
	$T = 400$	0.00179	0.00331	0.00093	0.00145	0.00598
$n = 20$	$T = 100$	0.01094	0.02697	0.00626	0.00819	0.03951
	$T = 400$	0.00267	0.00732	0.00153	0.00213	0.03603

Table B.3: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a grid network of sizes $n = \{8, 20\}$ after $T = \{100, 400\}$ iterations.

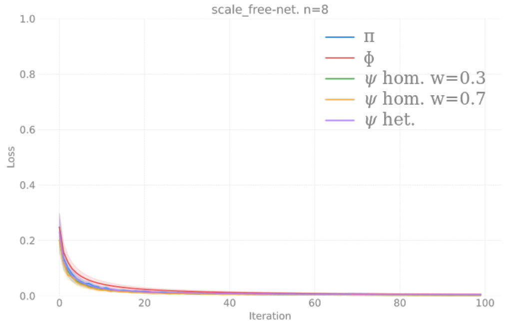

(a) $n = 8$, $T = 100$

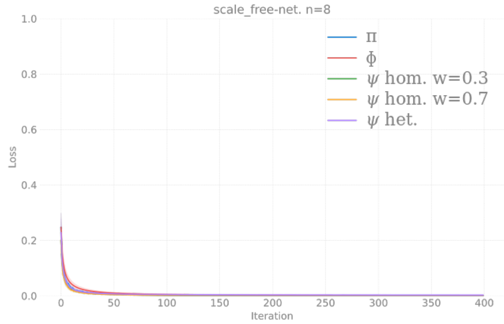

(b) $n = 8$, $T = 400$

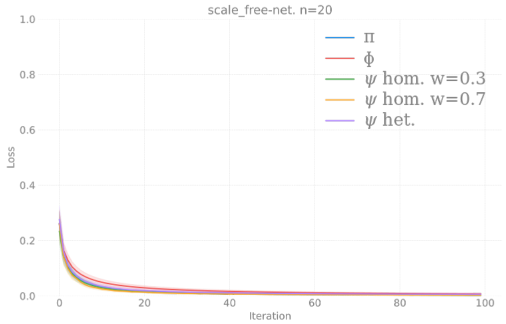

(a) $n = 20$, $T = 100$

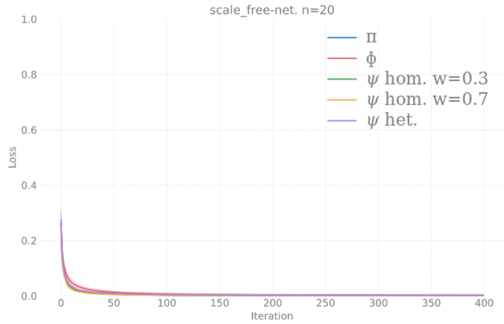

(b) $n = 20$, $T = 400$

Figure B.4: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a scale-free network of size $n = \{8, 20\}$ for $t = \{100, 400\}$ iterations.

		π	ϕ	hom. mixed $w = 0.3$	hom. mixed $w = 0.7$	het. mixed
$n = 8$	$T = 100$	0.00332	0.00520	0.00264	0.00253	0.00463
	$T = 400$	0.00092	0.00130	0.00065	0.00059	0.00266
$n = 20$	$T = 100$	0.00326	0.00727	0.00310	0.00284	0.00579
	$T = 400$	0.00091	0.00192	0.00076	0.00067	0.00321

Table B.4: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a scale-free network of sizes $n = \{8, 20\}$ for $t = \{100, 400\}$ iterations.

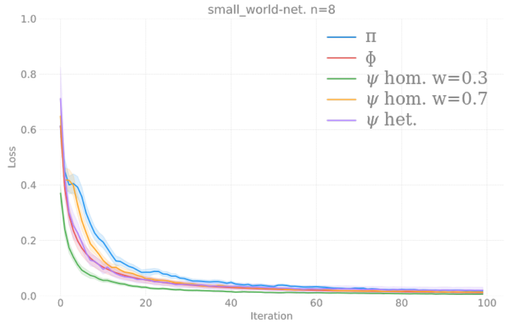

(a) $n = 8$, $T = 100$

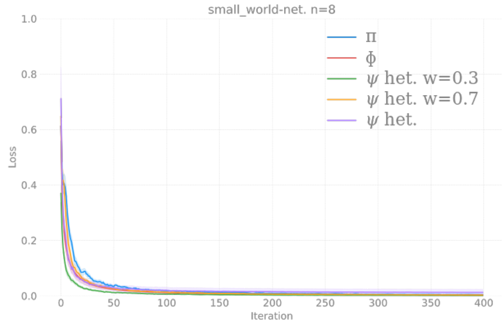

(b) $n = 8$, $T = 400$

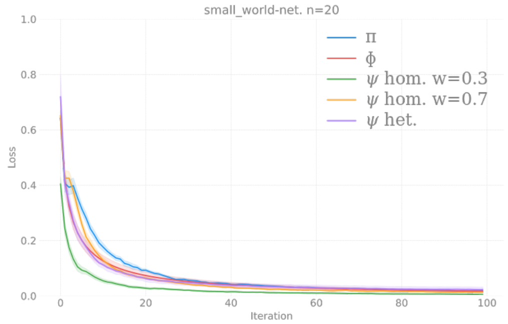

(a) $n = 20$, $T = 100$

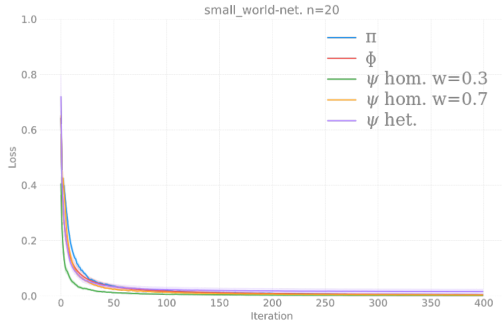

(b) $n = 20$, $T = 400$

Figure B.5: Market-level convergence of loss over time for greedy strategy π (Π), proportional strategy ϕ (Φ), two homogeneous mixed strategies (ψ *hom. w*=0.3 and ψ *hom. w*=0.7), and a heterogeneous mixed strategy (ψ *het.*) in a small-world network of size $n = \{8, 20\}$ for $T = \{100, 400\}$ iterations.

		π	ϕ	hom. mixed $w = 0.3$	hom. mixed $w = 0.7$	het. mixed
$n = 8$	$T = 100$	0.01862	0.01232	0.00666	0.01366	0.02005
	$T = 400$	0.00468	0.00308	0.00140	0.00303	0.01262
$n = 20$	$T = 100$	0.01852	0.01743	0.00584	0.01331	0.02316
	$T = 400$	0.00438	0.00452	0.00155	0.00330	0.01569

Table B.5: Loss values for π, ϕ, ψ with $w = 0.3$ (hom.), ψ with $w = 0.7$ (hom.), and ψ (het.) in a small-world network of sizes $n = \{8, 20\}$ after $T = \{100, 400\}$ iterations.

B.2. Misreporting across all α

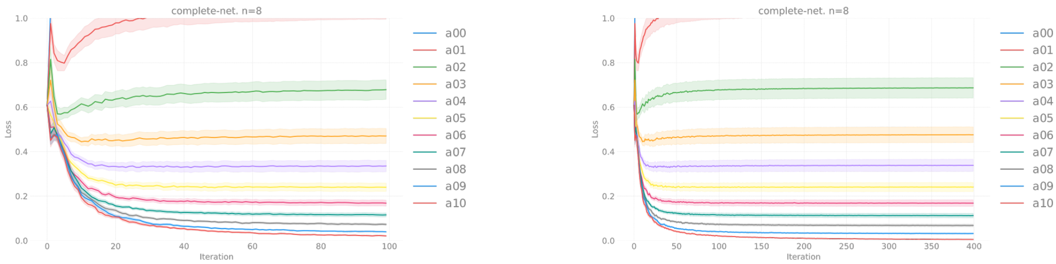

(a) $T = 100$ (b) $T = 400$

Figure B.6: Market-level converge of loss over time for π with various levels for misreporting parameter α, increasing steps of 0.1 from $\alpha = 0$ to $\alpha = 1$ in a complete network of sizes $n = \{8\}$ for $T = \{100, 400\}$ iterations. The loss for $\alpha = 0$, becomes so large that it exceeds the plotted domain.

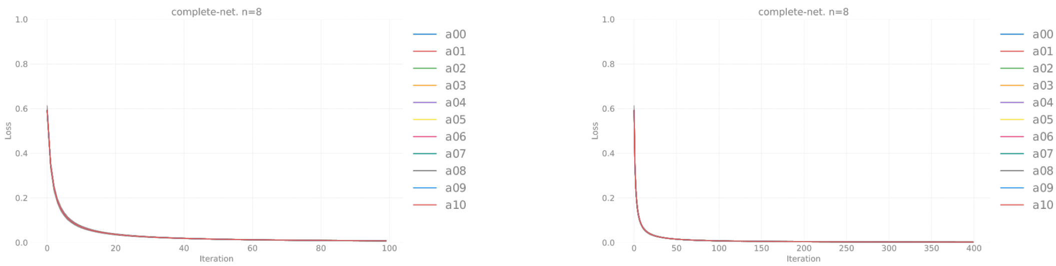

(a) $T = 100$ (b) $T = 400$

Figure B.7: Market-level converge of loss over time for ϕ with various levels for misreporting parameter α, increasing steps of 0.1 from $\alpha = 0$ to $\alpha = 1$ in a complete network of sizes $n = \{8\}$ for $T = \{100, 400\}$ iterations.

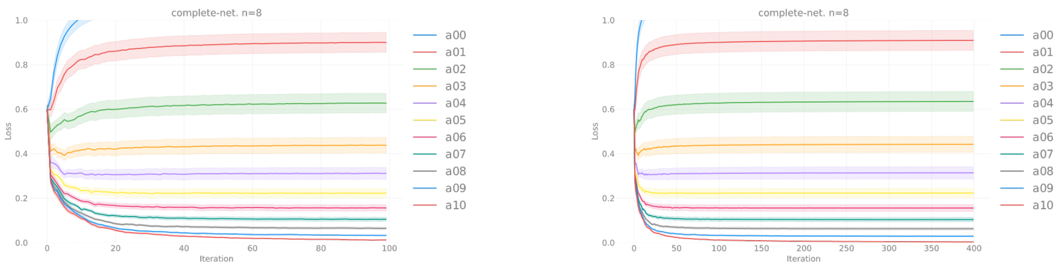

(a) $T = 100$ (b) $T = 400$

Figure B.8: Market-level converge of loss over time for ψ with $w = 0.5$ (hom.) with various levels for misreporting parameter α, increasing steps of 0.1 from $\alpha = 0$ to $\alpha = 1$ in a complete network of sizes $n = \{8\}$ for $t = \{100, 400\}$ iterations. The loss for $\alpha = 0$, becomes so large that it exceeds the plotted domain.

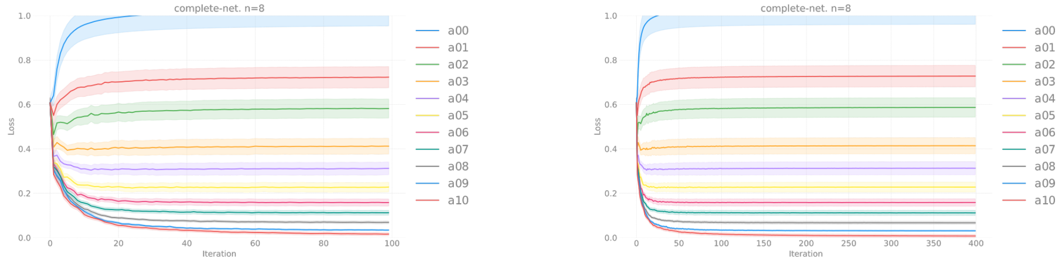

(a) $T = 100$

(b) $T = 400$

Figure B.9: Market-level converge of loss over time for ψ (het.) with various levels for misreporting parameter α, increasing steps of 0.1 from $\alpha = 0$ to $\alpha = 1$ in a complete network of sizes $n = \{8\}$ for $t = \{100, 400\}$ iterations. The loss for $\alpha = 0$, becomes so large that it exceeds the plotted domain.