Towards Strengthening Decentralised Exchange

Extended Abstract

Rixt Hellinga
TU Delft
Delft, Netherlands
r.hellinga@student.tudelft.nl

Georgios Iosifidis
TU Delft
Delft, Netherlands
G.Iosifidis@tudelft.nl

Neil Yorke-Smith
TU Delft
Delft, Netherlands
n.yorke-smith@tudelft.nl

ABSTRACT

This extended abstract considers the gap between the theoretical ideals and practical realities of distributed and money-free resource exchange protocols among self-interested agents. We sketch a convex program that corresponds to the centralised exchange market equilibrium. Secondly we state the convergence of proportional exchange allocation, comprising a distributed solution to the equilibrium. Thirdly, summarising a simulation study, we observe that distributed, mixed-strategy protocols can achieve stable and desirable outcomes, modulated by sensitivity to population diversity, limited information, and strategic agent behaviour.

KEYWORDS

sharing economy; exchange protocols; distributed multi-agent systems; mechanism design; simulation

ACM Reference Format:
Rixt Hellinga, Georgios Iosifidis, and Neil Yorke-Smith. 2026. Towards Strengthening Decentralised Exchange: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/RDIS7876

1 MOTIVATION AND BACKGROUND

More robust exchange protocols for peer-to-peer sharing economies is the motivation for this work. The sharing economy is seen an alternative to traditional ownership-based economic systems, and has gained credence and feasibility in the digital world [10, 14]. We derive a theoretical model of the *money-free* sharing economy paradigm by proposing models (programs) and distributed algorithms (protocols). We also model important real-world factors, such as strategy diversity and manipulation incentives of the interacting agents, to understand their impact on the trading dynamics of distributed market algorithms. By demonstrating how theoretically-optimal algorithms can fail under certain conditions, the model helps identify where such market designs or protocols may break down and how they might be improved.

We consider a time-slotted market, where agents trade sequentially for a horizon of T timeslots, constrained to some network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. This network is defined on nodes $\mathcal{N}$ and bidirectional edges (connections) $\mathcal{E}$; it represents the trading options of the agents, where each agent i may only trade with one-hop neighbours, i.e., $\mathcal{N}_i = \{j \mid (i,j) \in \mathcal{E}\}$. Related classical markets are the Arrow–Debreu (AD) [1] and Fisher market [5], which nevertheless rely on monetary instruments and assume distinct buyer and/or seller agent roles. Our focus instead is on the less-explored direct *exchange markets* as those arising e.g., over social networks [9], where each agent is a prosumer of resources or services.

The market operates as follows. At the start of each timeslot t, every agent $i \in \mathcal{N}$ produces a random amount of $D_i(t) \geq 0$ units of resource and decides its distribution with the allocation vector $\mathbf{x}_i = (x_{ij} \geq 0, j \in \mathcal{N}_i)$ which naturally can be updated every slot. We define $\mathbf{X} = (\mathbf{x}_i, i \in \mathcal{N})$, and assume $\mathbf{E}[D_i(t)] = \varepsilon_i, \forall i$. As a consumer, every agent enjoys utility $u_i(\mathbf{x}_{-i}) = \sum_{j \in \mathcal{N}_i} v_j x_{ji}$, where v_j is the value the resource of j has for all other agents. We are interested to explore the convergence of these interactions. To that end, we define the *exchange equilibrium* as the stationary exchange profile that clears the market and it is individually optimal, namely:

*Definition 1.1 (**Equilibrium**).* An allocation matrix $\mathbf{X}$ is an equilibrium allocation matrix $\mathbf{X}^*$ if the following conditions hold:

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

We note that a closely related equilibrium concept, expressed in terms of exchange ratios, was introduced in Georgiadis et al. [7].

2 MARKET MODEL AND ALGORITHMS

Our first goal is to devise a centralized program capturing the ideal operation of this exchange market. Similar models exist for markets with monetary instruments, but to the best of our knowledge this remains an open question for exchange markets, see also the illuminating discussion of Tsoukatos [11]. In particular, following the seminal work of Eisenberg and Gale [6] which devised a program for the Fisher market [2], we propose the following model:

$$\begin{aligned} \max_{\mathbf{X}} \quad & \sum_{i \in \mathcal{N}} \left((v_i \cdot \varepsilon_i) \log \left(\sum_{j \in \mathcal{N}_i} v_j \cdot x_{ji} \right) \right) \\ \text{s.t.} \quad & \forall i \in \mathcal{N} : \sum_{j \in \mathcal{N}_i} x_{ij} \leq \varepsilon_i, \\ & \forall i, j \in \mathcal{N} : x_{ij} \geq 0. \end{aligned} \tag{P1}$$

In analogy to the EG model, the objective in P1 uses the product $v_i \varepsilon_i$ to measure the contribution of agent i to the market, and weights its utility accordingly, while using a logarithmic transformation to balance the exchanges (means of fairness). The constraints are straight-forward, ensuring simply that $\mathbf{X}$ is feasible. From a technical point of view, Program P1 is useful due to the following result which is proved by Hellinga [8].

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/RDIS7876

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

Table 1: Loss values (lower is better) for greedy (π), proportional (ϕ) and mixed (ψ) strategies with combinations of memory decay (factor γ) and asynchronous dynamics in a complete network of size $n = 20$ after T iterations.

Theorem 2.1 (**Optimality of Program P1**). *Program* P1 *computes the equilibrium allocations* $\mathbf{X}^*$ *according to Definition 1.1.*

Naturally, distributed solution algorithms to P1 can be thought of as trading strategies for the autonomous agents. There is a twofold challenge here: first, to identify an algorithm that admits a physical interpretation, i.e., aligns with the rationale of real-world agents; and secondly, prove this algorithm solves the market program and hence attains the desirable equilibrium. Not surprisingly, this approach has been followed for other types of markets. It was shown, for instance, that proportional bidding can identify the equilibrium [13], and corresponds to a gradient descent algorithm solving the Fisher market (only buyers) [2]; a similar bidding strategy converges in markets with buyers and sellers [3]; Tsoukatos [11] extended this approach to settings where the agents can create links.

For the money-free market, Wu and Zhang [12] studied a proportional bidding; Georgiadis et al. [7] proposed a greedy strategy where each agent serves its best-reciprocating neighbour at each slot. Here, we propose a proportional allocation (without bidding), namely:

$$\phi: \quad x_{ij}(t+1) = D_i(t+1) \cdot \frac{v_j \cdot x_{ji}(t)}{\sum_{k \in \mathcal{N}_i} v_k \cdot x_{ki}(t)}, \forall (i,j) \in \mathcal{E}, \qquad (1)$$

and prove it solves P1, as it is essentially a mirror-descent iteration. We also consider *mixed-response* strategies, where each agent follows a mixture of greedy (π) and proportional reciprocation (ϕ), and evaluate its convergence to market equilibria; see Hellinga [8] for details.

3 EMPIRICAL STUDY

Moving then from theory to empirics, we study the market dynamics and trading strategies through simulation. A first set of experiments explored how mixed strategies, asynchronous dynamics, and memory decay affect convergence to equilibrium. A second set of experiments investigated vulnerability of the market to manipulation, focusing on non-truthful behaviour: Sybil attacks and misreporting of the sharing ratio ρ_i. The latter is defined as an agent's ratio of value received to value allocated.

While the 'component' strategies both converge close to equilibrium, the full market model with mixed-strategy agents stagnates at a much higher loss (Table 1). 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 model still converges in a general sense, it converges to a suboptimal market state. This aligns with the expectation that the real-world features modelled here – strategy heterogeneity, asynchrony, memory decay – can hinder the market's ability to reach perfect equilibrium.

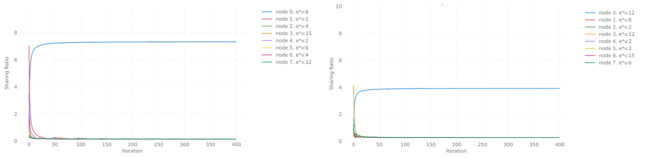

Network instance #1 **Network instance #2**

Figure 1: Sharing ratios in two instances of a complete network of size $n = 8$ for $T = 400$ iterations. The sharing ratio of the misreporting agent 0 exceeds that of truthful agents.

Regarding untruthful agents, we present two instance-specific results in Figure 1, which are representative for all network instances; 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. We observe that for both instances in Figure 1 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.

Regarding Sybil attacks, in our simulations the observed incentive ratio (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 [4]. 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.

4 CONCLUSION

The increasing prevalence of agentic networks, in which autonomous agents engage in distributed exchange according to specific protocols, poses significant challenges to conventional algorithmic economic theory. This paper aligned with a fundamental reconceptualisation of algorithmic market design. We modelled the distributed and behaviourally-rich nature of exchange in P2P sharing economies, constructing a principled model of distributed exchange, incorporating asynchronous trading, bounded rationality, and heterogeneous strategies while replacing global (monetary) information with local feedback mechanisms.

The central technical contribution was an analysis of distributed trading strategies. We adapted programs well-known in conventional economic theory to align with the exchange market, and stated the optimality of the modified EG program. We then explained, for the first time in this context, why proportional trading strategies can lead to equilibrium. Simulation results showed how our model produces a range of market phenomena with clear real-world relevance, and the empirical convergence to a market equilibrium in many settings. Future work includes studying the convergence rates of the various agent strategies.

ACKNOWLEDGMENTS

Thanks to the anonymous reviewers at AAMAS. This work was partially funded by the European Commission through grant number 101139270 (ORIGAMI).

REFERENCES

[1] Kenneth J. Arrow and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265. https://doi.org/10.2307/1907353

[2] Benjamin Birnbaum, Nikhil R. Devanur, and Lin Xiao. 2011. Distributed algorithms via gradient descent for Fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC '11)*. 127–136. https://doi.org/10.1145/1993574.1993594

[3] Simina Brânzei, Nikhil R. Devanur, and Yuval Rabani. 2021. Proportional Dynamics in Exchange Economies. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC'21)*, Péter Biró, Shuchi Chawla, and Federico Echenique (Eds.). 180–201. https://doi.org/10.1145/3465456.3467644

[4] Zhou Chen, Yukun Cheng, Xiaotie Deng, Qi Qi, and Xiang Yan. 2017. Limiting User's Sybil Attack in Resource Sharing. In *Proceedings of the 13th International Conference on Web and Internet Economics (WINE'17)*. Springer, 103–119. https://doi.org/10.1007/978-3-319-71924-5_8

[5] Ran Duan and Kurt Mehlhorn. 2015. A combinatorial polynomial algorithm for the linear Arrow–Debreu market. *Information and Computation* 243 (2015), 112–132. https://doi.org/10.1016/j.ic.2014.12.009

[6] Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168. https://doi.org/10.1214/aoms/1177706369

[7] Leonidas Georgiadis, George Iosifidis, and Leandros Tassiulas. 2020. On the Efficiency of Sharing Economy Networks. *IEEE Transactions on Network Science and Engineering* 7, 3 (2020), 1094—1110. https://doi.org/10.1109/tnse.2019.2904431

[8] Rixt Hellinga. 2025. *Algorithms for the Sharing Economy: An Economic Modelling Perspective on Distributed Exchange.* Master's thesis. Delft University of Technology. https://resolver.tudelft.nl/uuid:1f7448d6-b8b7-4607-979d-c1cd872c2be0

[9] Hirokazu Shirado, George Iosifidis, Leandros Tassiulas, and Nicholas A. Christakis. 2019. Resource sharing in technologically defined social networks. *Nature Communications* 10, 1 (2019), 1079. https://doi.org/10.1038/s41467-019-08935-2

[10] John Taskinsoy. 2023. The Reincarnation of Barter Trade & Barter Economy. *SSRN Electronic Journal* 4456717 (2023), 53. https://doi.org/10.2139/ssrn.4456717

[11] Konstantinos P. Tsoukatos. 2020. Network Formation by Reciprocity Versus Sparsity Tradeoffs. *IEEE Transactions on Network Science and Engineering* 7, 3 (2020), 1752–1763. https://doi.org/10.1109/TNSE.2019.2952454

[12] Fang Wu and Li Zhang. 2007. Proportional response dynamics leads to market equilibrium. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC'07)*. 354–363. https://doi.org/10.1145/1250790.1250844

[13] Li Zhang. 2011. Proportional response dynamics in the Fisher market. *Theoretical Computer Science* 412, 24 (2011), 2691–2698. https://doi.org/10.1016/j.tcs.2010.06.021 Selected Papers from 36th International Colloquium on Automata, Languages and Programming (ICALP 2009).

[14] Xiaoxi Zhu and Kai Liu. 2021. A systematic review and future directions of the sharing economy: Business models, operational insights and environment-based utilities. *Journal of Cleaner Production* 290 (2021), 125209. https://doi.org/10.1016/j.jclepro.2020.125209