Human Interaction and Emerging Technologies (IHIET 2023), Vol. 111, 2023, 200–207
https://doi.org/10.54941/ahfe1004025

Simulating and Quantifying Inequality in Strategic Agent Networks

Mayank Kejriwal

Information Sciences Institute, University of Southern California, Los Angeles CA 90292, USA

ABSTRACT

Transactions are an important aspect of human social life, and represent dynamic flow of information, intangible values, such as trust, as well as monetary and social capital. Although much research has been conducted on the nature of transactions in fields ranging from the social sciences to game theory, the systemic effects of different types of strategic agents transacting in real-world social networks (often following a scale-free distribution) are not fully understood. An influential economic measure that has not received adequate attention in the complex networks and game theory communities, is the Gini Coefficient, which is widely used to quantify and understand wealth inequality. In this paper, we define a network model called a *strategic agent network (SAN)* and present a methodological framework based on game theory for investigating questions of inequality using SANs. We briefly comment on results obtained from a preliminary experimental investigation using a real-world dataset based on Bitcoin.

Keywords: Strategic agent networks, Prisoner's dilemma, Game theory, Gini coefficient, Inequality, Modelling dynamic transactions

INTRODUCTION

With conceptual and methodological advances in both network science (Barabasi, 2013), and computational social science (Alvarez, 2016), it has become possible to study complex research questions by modelling and simulating the evolution of dynamic social systems (Troitzsch, 2012; Dabbaghian and Mago, 2014). Research on social networks within the computational sciences alone now spans over two decades of research, with recent focus on higher-order and 'multiplex' networks (Kanawati, 2015). However, although networks have a long history in modelling and simulating human social systems, such as work in network growth models that aim to model the dynamic aspects of such systems (including how the scale-free structure often observed in such networks comes to be), there is less work on using the network to model a system of *dynamic transactions*. A transaction here does not have to be monetary in nature, since transactions in the real world can involve intangible goods, such as goodwill, social capital, trust, information and other 'goods' on which there has been much exposition in the social sciences (Halpern, 2005; Moore, 1999). We adopt a dictionary definition of the word here, with Merriam-Webster defining the word *transaction* as both (emphases

© 2023. Published by AHFE Open Access. All rights reserved.

ours) "an *exchange* or *transfer* of goods, services, or funds" and as "a communicative action or activity involving two parties or things that *reciprocally affect* or *influence* each other."

Arguably, a significant fraction of interactions in everyday social life is transactional, rather than growth-based where new nodes or actors form a social acquaintance with us. Intuitively, the majority of our interactions tend to be largely limited to people we are already connected to, be they friends, family, neighbours and colleagues. From a utility theory perspective, transactions should only occur between two parties in a rational, non-duress setting if both parties expect to gain from the transaction (mutual benefit). Formal mechanisms, such as contract law, can be used to encode the expectations of both parties; however, in everyday life, transactions are more informal and based on trust. Individuals are also likely to change their behaviour in the presence of external cues and incentives. Finally, once 'betrayed' in a transaction, a betrayed individual is less likely to transact with the betraying individual.

Such behaviour is not fixed, but is also not random. In the last 75 years, game theory has emerged as an influential field of research aiming to model and explain decision making under conditions of uncertainty (Weintraub, 1992; Leonard, 2010). Although game theory historically found its applications in economics and social sciences, much more recently, a growing body of work has explored its utility in modelling agent-based interactions in graphs and networks (Chen et al. 2009; D'souza et al. 2007. Yet, research at the intersection of game theory and networks has tended to focus more on *network growth*, including game-theoretic explanations for models (sometimes, but not always, relying on simulations) such as preferential attachment (Avin et al. 2018). In contrast, there is much less exploration of how game theory can be used to model and simulate dynamic transactions on an existing network, especially in the presence of external incentives and other reasonable constraints guided by sociological observations (e.g., that transactions are less likely to occur in the absence of a pre-existing relationship).

This paper presents an abstract agent-based model that seeks to represent both network structure and the game-theoretic nature of strategic transactions. We focus on strategic transactions as a means to developing a better understanding of *inequality.* Intuitively, if we imagine that the initial condition of a system was one of perfect equality (e.g., everyone has equal wealth or property). Assuming that there was no force or usurpation, inequality can only arise in a system (open or closed) either because of redistribution or if some individuals receive more wealth (in aggregate) over time compared to others. In other words, *microscopic* transactions, broadly defined, allow us to understand the evolution of *macroscopic* properties such as inequality. We note that questions of wealth and income inequality play a major role in economics, with entire books written on the subject over the decades (Atkinson, 2016; Piketty, 2015). This paper is not meant to comment on specific economic policies; but rather, to show that inequality can arise naturally in a system even from the implementation of very local incentives and strategies.

Model: Strategic Agent Network

The model of transacting agents, called a *strategic agent network (SAN)*, is fundamentally defined in this paper as an undirected graph G= (V, E), where V is the set of nodes or vertices, and E is the set of undirected edges. Based on prior discussion, nodes should be thought of as individuals or users ('agents') who are aiming to transact with one another given the Prisoner's Dilemma payoff matrix in Figure 1. The transactions are constrained by the structure of the network i.e., a node cannot transact with another node unless it is directly linked to it via an edge. Sociologically, the edges represent pre-existing relationships that lead to higher likelihood of transactions occurring.

B / A	B stays silent or 'cooperates'	B betrays or 'defects'
A stays silent or 'cooperates'	+1 / +1	+3 / -3
A betrays or 'defects'	-3 / +3	-2 / -2

Figure 1: The Prisoner's Dilemma payoff matrix that can be used as a basis for simulating transactions between any two agents, with a clear incentive and reward structure.

To understand how this framework can be used to simulate transactions, let us consider an example SAN in Figure 2. This SAN looks like an ordinary social network, but each agent is now endowed with a *strategy*. For example, some agents may be naturally inclined to always cooperate, while others may always defect. Many such strategies have been developed and tested in the game theory literature exploring the Prisoner's Dilemma. In the example, all strategies are deterministic and fixed, but in a more complex setting, agents could be adaptive in their adoption of strategies (e.g., in response to increasing or decreasing payoffs), and strategies could be stochastic. A naïve example might be an agent who randomly chooses to cooperate or defect whenever they are given an opportunity to transact.

An actual simulation on this network requires certain decisions to be made that depend on both the experiment and research goal. A robust methodology might be to first determine an 'ordering' on the nodes. Considering the example in Figure 2, suppose the ordering is [A, B, C, D, E] and the system is perfectly equal i.e., all nodes begin with 100 resource-units (e.g., dollars) each. In the simulation, we traverse the nodes in the order above per iteration. For a given node, we randomly sample a neighbouring node and conduct a transaction experiment. Suppose that, for node A, we sample its neighbour

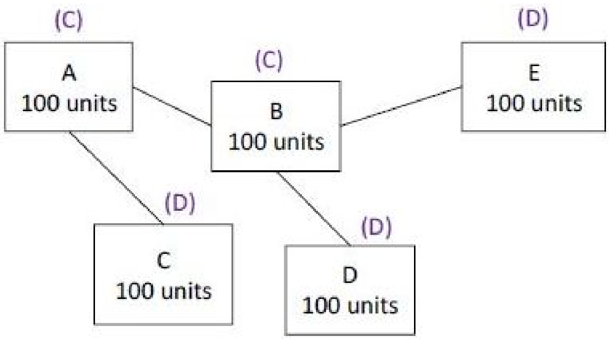

Figure 2: An example of a simple strategic agent network (SAN) where an agent chooses to always cooperate or defect (C or D), as shown above the node. We assume that, initially, each agent has 100 resource-units (e.g., dollars) and there is perfect equality in the network.

node C. Since node A cooperates, and node C defects, node A loses three of its units to node C (see payoff rules in Figure 1). Next, we move on to node B, for whom we randomly sample a neighbour (say, node D). A similar process applies. The details for a complete iteration (ending with node E), are shown at the side of Figure 3, with the resulting network. In the next iteration, we return to the beginning of the list and start the same process again. Because neighbours are randomly sampled, A does not have to necessarily transact with C in the next round. For similar reasons, a node may also end up participating in more than one transaction per iteration.

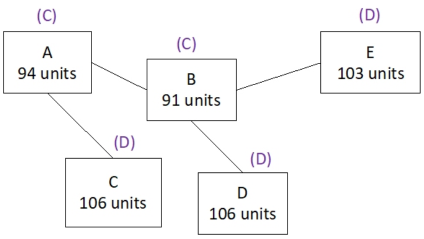

Node traversed	Neighbour sampled	Transaction	Payoff
A	C	A cooperates, C defects	A gives 3 units to C
B	D	B cooperates, D defects	B gives 3 units to D
C	A	C defects, A cooperates	C receives 3 units from A
D	B	D defects, B cooperates	D receives 3 units from B
E	B	E defects, B cooperates	E receives 3 units from B

Figure 3: The resulting network (starting from the network and initial conditions in Figure 2) after a single round of simulation when all nodes have been sampled exactly once, and engaged in exactly one transaction with a randomly sampled neighbouring node (details provided in main text, along with the 'log' of transactions on the right side). We assume that, when both nodes cooperate, or both defect, no resource-units are exchanged between the two.

The sampling and transaction procedures described above are not the only ways to conduct such a simulation using a SAN, but it has the advantage of being 'fair' in terms of the average number of transactions per node. This is especially important in real-world scale-free networks where some nodes may have many neighbours. One might argue that such nodes should have the benefit of conducting 'more' transactions; statistically, this is indeed the case, because each of the (high-degree) node's neighbours have an opportunity to sample the node when it is their turn to transact. Others may argue that each possible transaction should have equal probability of being conducted, and the number of transactions per round should be experimentally fixed, rather than being bounded by the total number of nodes (since each node is traversed once in the simulation algorithm suggested earlier). A different sampling procedure (at the level of edges, rather than nodes) would then apply.

Indeed, when conducting a simulation on a proper dataset, several other details also need to be worked out. If an agent is depleted of resource-units, should that agent be removed from the network? Should the system be allowed to be closed (so the sum of resource-units across all nodes remains constant over time) or should an external entity be modelled that leads to increase or decrease of net overall wealth over time (this is especially relevant when both nodes cooperate, or both defect, as Figure 1 suggests that both should be gaining or losing in wealth as a consequence)? How often should agents be allowed to change their strategy (if at all), and how much are they allowed to observe the dynamics of the overall system (e.g., should perfect information be assumed, as might be the case if the network were modelling a small community where reputations are well known)? These details allow for degrees of freedom, depending both on the research question being investigated and analytical setting assumed or modelled. However, underlying all of these variants, two tenets of the model always hold:

- First, each node in the network is interpreted as an agent making decisions under uncertainty.
- Second, given the decisions of two agents, a game-theoretic payoff matrix is used to determine distribution or redistribution of wealth between the two agents.

Note also that the details and variants mainly concern microscopic properties of the network. Some of these properties may be global (e.g., the payoff matrix, as well as initial conditions) but all are applied and interpreted in the local neighbourhood of the agent. However, the quantity of interest in running a simulation over a large enough number of iterations is typically *macroscopic*, making the model an ideal one for investigating the impacts of strategy and incentive in complex networked systems. As discussed subsequently, in our own preliminary experiments, we focused on how, and to what extent, inequality increases in the system because of this network of transactions being conducted over time.

Returning to our goal of modelling inequality, we can run the per-node traversal simulation for a large number of iterations (say, 1000 or until convergence i.e., resource-units quantities cease to be exchanged between nodes), and at the end of each iteration, compute the *Gini Coefficient*, which is a real value that ranges from 0 (perfect equality, which is the initial state of the network before the first iteration begins) and 1 (perfect inequality, where one individual has all the resource-units). A full formula and explanation may be found in (Lambert and Aronson, 1993).

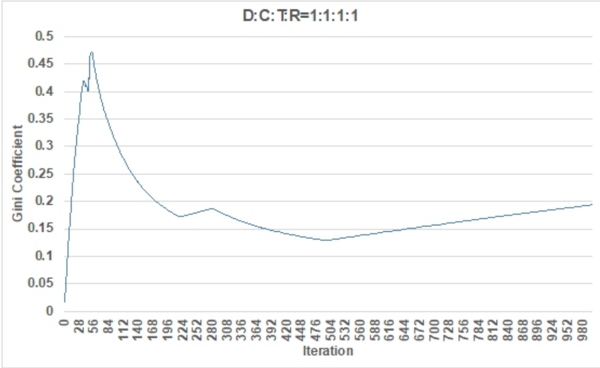

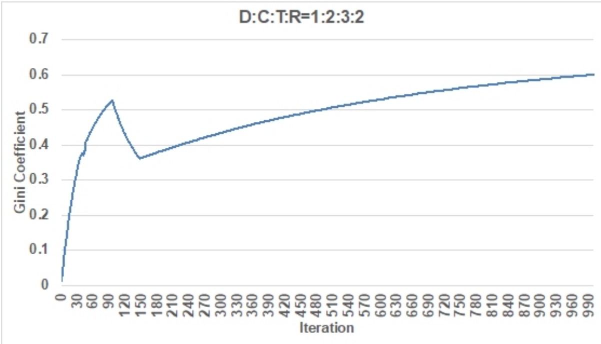

Figure 4: The resulting inequality (measured on the y-axis using the Gini Coefficient) after each iteration (x-axis) on a real-world Bitcoin network under different conditions of strategic-agent mixtures (D:C:T:R = 1:1:1:1 on the top, and D:C:T:R = 1:2:3:2 on the bottom). In both cases, the structure of the network is fixed and static. D, C, T, and R define the *always defect, always cooperate, tit-for-tat, and randomly defect / cooperate* strategies, respectively.

Using a real-world Bitcoin transaction network (Kumar et al. 2016), with more than 5,800 nodes and 35,000 edges, we ran a simulation to demonstrate how the model can give us different descriptions and evolution of inequality. As shown in Figure 4, when nodes are randomly assigned (at the beginning of the simulation) to have fixed strategies of always defect (D), always cooperate (C), tit-for-tat (T) which is a well-known agent detailed in (Rapoport, 2015), and randomly select between cooperate and defect with 50% probability (R), we find that inequality rises initially but then declines before starting a slow ascent again. However, when we change the proportions of the four agents in the proportion D:C:T:R = 1:2:3:2; namely, 12.5% of randomly selected agents (note that this assignment only occurs once at the beginning of the simulation, and is then fixed throughout the simulation) are D, while 25% each are C and R, while the remainder (37.5%) are T, inequality dips only briefly before eventually beginning a steep ascent and nearly crossing 0.6 by the thousandth iteration. The simulation shows that, even after controlling for the network *structure,* the evolution of inequality also depends on the mix of strategies in the network.

CONCLUSION

We proposed and conducted a systematic methodology that uses principles at the intersection of game theory and network science to simulate and quantify inequality in a complex system of inter-connected, transacting agents. Our model specifically relies on nodes in a complex network having 'strategies' resulting in 'local' transactions and redistribution of resource units that ultimately yield interesting 'macroscopic' trends such as rising inequality. An important avenue for future research is to consider the functions that govern the Gini Coefficient distributions shown in Figure 4, and the theoretical derivation of those functions. It may also be valuable to consider applying the Gini Coefficient in dynamical versions of network growth models, where the network's nodes and edges are themselves not static, but allow for incoming nodes and edges with each iteration. In essence, this would require us to model two kinds of dynamic behaviour: the transactional behaviour that we explored in this article, and the edge-formation behaviour that is often used to explain the scale free degree distributions of the kinds of networks employed in this paper. The latter tends to draw on more psychological theories of behaviour, such as a preference for new nodes to 'attach' to nodes that (already) have relatively high degree. Some work in game theory has attempted to explain edge formation using transactions, but we hypothesize that both mechanisms (preferential attachment and game-theoretic modelling of cooperative-competitive transactions) can together be more fruitful in producing a richer, more accurate and more theoretically satisfying model of interactive human systems.

REFERENCES

Alvarez, R. M. (Ed.). (2016). *Computational social science*. Cambridge University Press.

Atkinson, A. B. (2016). Inequality: What can be done. *Practice*, *40*(2), 289–292.

Avin, C., Cohen, A., Fraigniaud, P., Lotker, Z., & Peleg, D. (2018, April). Preferential attachment as a unique equilibrium. In *Proceedings of the 2018 World Wide Web Conference* (pp. 559-568).

Barabási, A. L. (2013). Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *371*(1987), 20120375.

Chen, W., Teng, S. H., Wang, Y., & Zhou, Y. (2009). On the α-sensitivity of Nash equilibria in PageRank-based network reputation games. In *Frontiers in Algorithmics: Third International Workshop, FAW 2009, Hefei, China, June 20-23, 2009. Proceedings 3* (pp. 63–73). Springer Berlin Heidelberg.

Dabbaghian, V., & Mago, V. K. (Eds.). (2014). *Theories and simulations of complex social systems*. Springer Berlin Heidelberg.

D'souza, R. M., Borgs, C., Chayes, J. T., Berger, N., & Kleinberg, R. D. (2007). Emergence of tempered preferential attachment from optimization. *Proceedings of the National Academy of Sciences*, *104*(15), 6112–6117.

Halpern, D. (2005). *Social capital*. Polity.

Kanawati, R. (2015). Multiplex Network Mining: A Brief Survey. *IEEE Intell. Informatics Bull.*, *16*(1), 24–27.

Kumar, S., Spezzano, F., Subrahmanian, V. S., & Faloutsos, C. (2016, December). Edge weight prediction in weighted signed networks. In *2016 IEEE 16th International Conference on Data Mining (ICDM)* (pp. 221–230). IEEE.

Lambert, P. J., & Aronson, J. R. (1993). Inequality decomposition analysis and the Gini coefficient revisited. *The Economic Journal*, *103*(420), 1221–1227.

Leonard, R. (2010). *Von Neumann, Morgenstern, and the creation of game theory: From chess to social science, 1900–1960*. Cambridge University Press.

Moore, M. (1999). Truth, trust and market transactions: What do we know?. *The Journal of Development Studies*, *36*(1), 74–88.

Piketty, T. (2015). *The economics of inequality*. Harvard University Press.

Rapoport, A., Seale, D. A., & Colman, A. M. (2015). Is tit-for-tat the answer? On the conclusions drawn from Axelrod's tournaments. *PloS one*, *10*(7), e0134128.

Troitzsch, K. G. (2012). Simulating communication and interpretation as a means of interaction in human social systems. *Simulation*, *88*(1), 7–17.

Weintraub, E. R. (Ed.). (1992). *Toward a history of game theory* (Vol. 24). Duke University Press.