SINGAPORE UNIVERSITY OF
TECHNOLOGY AND DESIGN

Efficiency, Regret and Inequality in Decentralised Systems

Submitted by

Barnabé MONNOT

Thesis Advisor

Prof. Georgios PILIOURAS

Engineering Systems & Design

A thesis submitted to the Singapore University of Technology and Design in fulfillment of the requirement for the degree of Doctor of Philosophy

2019

PhD Thesis Examination Committee

TEC Chair:	Prof. Costas Courcoubetis
Main Advisor:	Prof. Georgios Piliouras
Internal TEC member 1:	Prof. Shaowei Lin
Internal TEC member 2:	Prof. Zsombor Méder
Internal TEC member 3:	Prof. Bige Tunçer
External TEC member 1:	Prof. Christos Papadimitriou

Declaration

I hereby confirm the following:

- I hereby confirm that the thesis work is original and has not been submitted to any other University or Institution for higher degree purposes.
- I hereby grant SUTD the permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created in accordance with Policy on Intellectual Property, clause 4.2.2.
- I have fulfilled all requirements as prescribed by the University and provided 1 copy of my thesis in PDF.
- I have attached all publications and award list related to the thesis (e.g. journal, conference report and patent).
- The thesis does not contain patentable or confidential information.
- I certify that the thesis has been checked for plagiarism via turnitin/ithenticate. The score is 24%.

Name and signature:

Date:

Abstract

Engineering Systems & Design

Doctor of Philosophy

Efficiency, Regret and Inequality in Decentralised Systems

by Barnabé MONNOT

We consider systems where agents in a society choose actions from their strategy sets and incur a cost. This cost depends not only on their own choices, but also on those of other agents in the society. Guided by their own interests, agents at equilibrium decide on strategies to minimise their individual cost given the actions of everyone else. But a central planner can do better and select for each agent an action such that the sum of costs in the society is minimised. The gap between this minimum and the cost of the worst equilibrium is known as the Price of Anarchy (PoA).

Since its introduction by Koutsoupias and Papadimitriou (1999), PoA has stood as the gold standard for system efficiency. Taken literally, PoA is a measure of the inefficiency accruing to "anarchic" decentralisation. However, in many systems, mechanisms exist to induce optimal equilibria, where the society incurs costs equal to those it would incur under the governance of the central planner. This is a "designed" decentralisation, where a mechanism has modified the incentives and behaviour of the agents. In that case, PoA is simply 1.

The present thesis investigates the ties between efficiency, regret and inequality in decentralised systems. We exhibit four connected questions to conduct this investigation:

1. When optimal equilibria exist but inefficiency arises from the complexity of reaching these equilibria, can a centralised authority help agents coordinate?
2. When agents follow no-regret learning dynamics and reaching a Nash equilibrium is hard, can we guarantee efficiency bounds of the limit states?
3. In a large, real system such as the Singapore transportation network, what are data-driven proxies for the theoretically-defined notions of regret and PoA? Is the latter as pessimistic in real life as its theoretical bounds predict?
4. When agents are endowed with initial wealth, how do mechanisms inducing optimal equilibria modify the distribution of wealth?

Publications

Gemici, Kurtuluş and Elias Koutsoupias and Barnabé Monnot and Christos H. Papadimitriou and Georgios Piliouras (2019). "Wealth inequality and the price of anarchy". In: *36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)*. Leibniz International Proceedings in Informatics (LIPIcs), pp. 31:1–31:16.

Monnot, Barnabé and Francisco Benita, and Georgios Piliouras (2017). "Routing Games in the Wild: Efficiency, Equilibration and Regret". In: *International Conference on Web and Internet Economics*. Springer, pp. 340–353.

Monnot, Barnabé and Georgios Piliouras (2017). "Limits and limitations of no-regret learning in games". In: *The Knowledge Engineering Review* 32.

Monnot, Barnabé et al. (2016). "Inferring Activities and Optimal Trips: Lessons From Singapore's National Science Experiment". In: *Complex Systems Design & Management Asia*. Ed. by Michael-Alexandre Cardin et al. Springer, pp. 247–264.

Monnot, Barnabé and Justin Ruths (2016). "Sensitivity of Network Controllability to Weight-Based Edge Thresholding". In: *Complex Networks VII*. Springer International Publishing, pp. 45–61.

Singapore is, surprisingly, a prime location to study anarchy. Basking in its tropical warmth for 5 years allowed me to formalise many thoughts on the way systems are organised, or organise themselves—the former being much more frequent here in the Little Red Dot. My first acknowledgement is to the city itself, for providing me with a home away from home and funding throughout, via the President's Graduate Fellowship. But places are not the whole story, so the remainder of these acknowledgements will be about the people.

Few of these thoughts would have made any sense without my advisor, Georgios Piliouras. I learned as much from you about the ideas presented here as I did on what makes a great researcher, an honest, driven spirit and a dedicated mentor. Thank you for bringing me along to the places I needed to be to become what I hope is an ounce of that.

Conversations with Zsombor Méder, Marco Scarsini, Shaowei Lin, Elias Koutsoupias and Christos Papadimitriou, the guidance of and research with Bige Tunçer, Costas Courcoubetis, Francisco Benita, Justin Ruths, Katalin Tamási and Selin Ahipasaoglu helped me achieve more than I could have done on my own. A special place is reserved to Sin Chee, and more generally the ESD people, for being a shining beacon of light in the grim landscape of bureaucracy.

Five years is a long time, but it speeds past when you are among friends. Every moment with Priji feels illuminated by a warm glow seeping through trees, with the scent of tea in the air. Zack, Harry, Giorgio, Manos, Eliza, Oindrila, John, Yoon, Ozgün, Ezgi, Zubin, Ying Yi, Francis, Yukari, Giacomo, Ozge, Sandra, Bharath, Josh, David, Hung, Jia Yi, Gulten, Zhangsheng, Preethi, Yeshwant, Viaan, Zaneena, Nat, Gangotri, Ambika, Prince, Jon, Vlad, Aditya, Megan, Sacha, Cheng Liang, Yahsin made the Ph.D. exciting and worth it. I am grateful to have met all of them.

Leaving home, even when it was always the plan, feels bittersweet, but the feeling is lesser when home sometimes come to you. I thank my family for making the trip often—though I suspect I was not the only reason for it—and for their total support throughout. Taking you around my adoptive home registers among the very best memories of the whole experience. I was lucky enough to have one of my two brothers move in the same city and the other tirelessly colour my thoughts (and walls in undisclosable locations) with his spray cans and artworks.

Visiting my home city to reconnect with my family and friends never made me feel like I had left entirely. They made it easy to find anew the places we used to share, despite my brief, random, apparitions. Louis, Alexandre, Rukmini, Salomé, Côme, Alexis, Maïa, Caroline, Adélie, Nina, Morgane, Quentin, Maxence, Lucia, Guilherme and each of your special someones I had the pleasure to meet throughout, Charlotte, Fred *et* Céline, Daddy *et* Lise, Grand-Père *et* Sissi, merci.

Contents

"There is no central authority that designs, engineers and runs the Internet. But what if there were such master puppeteer, a benevolent dictator who, for example, micromanaged its operation, allocating bandwidth to flows so as to maximize total satisfaction? How much better would the Internet run? What is the price of anarchy?"

"Algorithms, Games and the Internet", Christos Papadimitriou (2001)

Chapter 1

Efficiency, regret and inequality in decentralised systems

Complex societies are built on the interactions of agents striving to achieve their own goals. Cities are a prime example: they work as a large meeting place for this multitude of intersecting agents and they grow from resulting socio-economic products (Jacobs, 1961; West, 2017). In this context, it would be a Sisyphean task to plan for all these interactions and optimise such that the society as a whole is in the *best* possible state it can be, its **social optimum** (SO)—for some definition of "best". The reason may only be computational: finding out such a state potentially requires resources far greater than will ever be available.

It is much easier to simply assume that each agent follows a plan of her own, guided by her own motivations and hope that the aggregation of everyone's individual decisions will lead the whole system to a good state. Hayek (1945) thus explains the superiority of decentralised decision-making, and exhumes "knowledge" as the key object that underpins such systems. The planner may indeed never be able to possess "all the knowledge which ought to be used but which is initially dispersed among many different individuals".[1]

What is the consequence of letting agents in the system do their own bidding? People do find their way on the complex routing network of large cities, and packets reliably arrive at destination on the Internet, with no central planner showing the way for everybody. But these two situations come with an important side effect: **congestion**. Resources become slower to access as more agents employ them and negative externalities decrease the global welfare of the system. In other words, decentralisation leads to inefficiency.

This is where, in Hayek's words, the problem may reduce to "designing an efficient economic system". Smart tolling can guide the road network away from suboptimal states (Fleischer, Jain, and Mahdian, 2004). Protocols such as TCP make the Internet

[1] As Mirowski and Nik-Khah (2018) note, since Hayek's work, the word "knowledge" has become problematic and later weaved into the concept of "information", leading to some confusion. We revisit some of these questions later in the thesis, especially in Chapter 3.

robust to overly congested links (Cerf and Kahn, 1974). Both make use of "knowledge": the first in a fully decentralised manner, by setting the price of the toll to the marginal cost of an additional driver employing the road, the second with clients observing the number of packets they lose as they steadily increase their rate of emission.

How can we quantify how well these mechanisms perform? We first need a benchmark to evaluate them against. The **Nash equilibrium** (NE) (Nash, 1950), defined as any state of the system where no agent could profitably change her decision if she was the only one to do so, is a natural candidate. If we believe agents act out of their self-interest, it appears reasonable to conclude that the logical endpoint of a system will be a NE.

Understanding how much improvement is achieved with these mechanisms—or, conversely, how much worse the system is without them—reduces to measuring the gap between the worst of the NE in the system and its SO. This idea is formalised by the **price of anarchy** (PoA), given by

$$\texttt{Price of anarchy} = \frac{\texttt{Cost(Worst Nash equilibrium)}}{\texttt{Cost(Social optimum)}}$$

Since its introduction by Koutsoupias and Papadimitriou, 1999,[2] a large literature followed to analyse the PoA in various settings, e.g., congestion games modelling the traffic on a road network or auctions and queues. Positive results were found, bounding the PoA for congestion games with particular cost functions, independently of the underlying network's topology (Roughgarden, 2015). But as with any metric, it is not an end-all.

First, as a predictive notion of system state, NE is a tough sell. It is in general not unique, leading us to wonder how the game will play out and why some equilibrium is reached instead of another. Even in systems where the social optimum is a Nash equilibrium,[3] this multiplicity can durably affect the costs of the agents. One natural idea is to inject some centralisation, to allow coordination within the game. *If agents with limited knowledge are given an extra bit of information from a central messenger, can the system be steered towards more efficient equilibria?*

Second, it is computationally hard to even exhibit one instance of a NE in most settings (Daskalakis, Goldberg, and Papadimitriou, 2006), a further argument that it may not be sensible to expect convergence to such an equilibrium for most real systems. That most learning dynamics do not in general converge to a NE was indeed known (Shapley, 1964) and prompted the definition of weaker notions of equilibrium for which

[2] Under the name of "coordination ratio", until it was retitled "price of anarchy" in Papadimitriou (2001).

[3] In other words, a system where the price of stability, or ratio of the best equilibrium's cost to the social optimum, is 1 (Anshelevich et al., 2004).

the play converges. One such definition is based on **regret** and expects agents to play in order to minimise their own accumulated regret. *If it is not clear that NE is a reasonable outcome of the agents' interactions, can we find better guarantees with a weaker concept of equilibrium?*

Let us assume for an instant that we have overcome the previous obstacles. We must still consider whether PoA is an adequate predictive measure of system inefficiency. Despite its extensive theoretical analysis for a variety of games, so far few attempts have succeeded at properly estimating the value of PoA for a congestion game played out in the real world, in part due to the lack of data granular enough to measure it.

Second, the notion of PoA is only meaningful if we believe our real system to be in a Nash equilibrium. In the case NE is still too strong to observe empirically, it remains possible that some measurement of regret from the data argues that the system has at least reached a stable state.

There is a third difficulty. Low values of PoA may represent two opposite conditions of the system—namely, a system with very light congestion or one with very high congestion (Colini-Baldeschi et al., 2017). A central planner is not of much use if the streets are empty, or, conversely, if they are in such a gridlock that all helpful links are saturated. *How does PoA hold up to empirical analysis of its measurement in a real instance of a routing game?*

A broader criticism of the emphasis on system efficiency is that the distance from social optimum tells us nothing about the actual distribution of costs among agents. It is a well-known fact that Pareto efficiency, an unavoidable concept in the theory of general equilibrium (Arrow and Debreu, 1954), does not discriminate between a position where one agent owns everything and one where all agents own the same share of wealth. Rawls (2009), §12, notes this obvious shortcoming and frames the ensuing problem as one of selecting the just allocation among all efficient ones. Inequalities may arise, but, according to the second principle, are justified as long as the worst-off individual is in a better situation than in an otherwise equal society.

Rawls thus offered a moral theory underpinning social inequalities. In parallel, a larger literature concerned with the adequate measurement of inequality grew from Gini (1921), a seminal text that introduced the well-known Gini index. The axiomatisation of these measurements, undertaken by Sen et al. (1997) and others, framed the topic of inequality in the terms of welfare economics. While this approach proved fruitful to derive properties of various inequality indices in concert with their implications for social welfare, it developed somewhat orthogonally to the more efficiency-driven excursions of price of anarchy work, leaving the effects of efficient mechanisms on inequality not entirely understood. *If agents are endowed with some initial wealth, how does a game modify this distribution?*

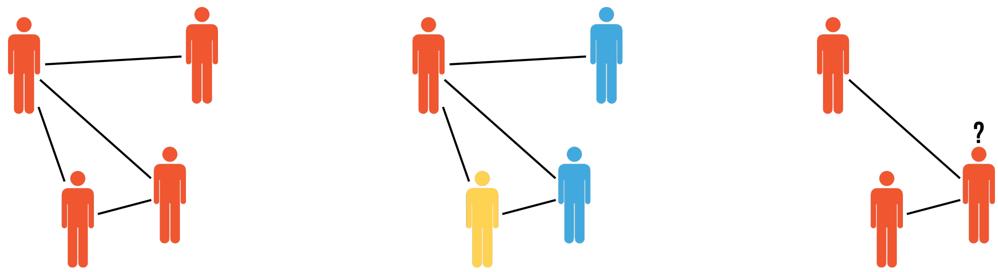

(A) (B) (C)

FIGURE 1.1: The graph colouring game. (a) Players are arranged on a network and start with the same initial colour. (b) A colouring of the network: no two neighbours share the same colour. (c) Players only have access to their local neighbourhood information.

We italicised four questions along the preceding introduction. These connected questions critically assess the price of anarchy as a guarantee for system performance and uncover the relations between efficiency, regret and inequality in decentralised systems. The remainder of this thesis is dedicated to shedding light on each question, employing both experimental methods and theoretical analysis.

1.1 When does sparse information seeding induce efficient equilibrium selection?

The multiplicity of Nash equilibria is the first obstacle towards a predictive notion for game theory. The quality of equilibria may vary drastically and it is not clear that agents left to their own decisions, in a fully decentralised fashion, can reach the best among all. Even when the game possesses additional structure—e.g., if it is a potential game, where each possible improvement by one player is matched with a comparable improvement by society as a whole—, the play can find itself stuck in a local equilibrium. Instead, introducing a restricted intervention by a central authority during the play could help navigate the landscape of equilibria towards more efficient instances.

Our first question employs an abstraction of social situations, where agents have limited information and incur costs that depend on the actions of their immediate neighbours. Players are placed on a network, a structure where each player is connected to a subset of other players (Figure 1.1a). Each player can decide on a colour for herself, among three available options. A player incurs costs proportional to the number of players she is connected to and who share the same colour as herself. In the game, there exists a large number of inefficient equilibria where some pairs of neighbours are still sharing the same colour. Conversely, there is a small number of configurations for which no two neighbours have matching colours, and these configurations are NE (Figure 1.1b).

This setting offers many parallels to Hayek's "circumstances of time and place": agent knowledge is limited to the colour choice of their immediate neighbours. The structure of the network is not known to them, nor the colours of second degree neighbours—neighbours of neighbours (Figure 1.1c). But even when agents have access only to limited information from their local neighbourhood, the game is set up such that "there is hardly anything that happens anywhere in the world that *might* not have an effect on the decision [the player] ought to make" (emphasis in the text, Hayek (1945)). A colour change some distance away from one player may ripple through the network via the actions of connected players reacting to the original change.

Now, in realistic settings, "we cannot expect this problem will be solved by first communicating all this knowledge to a central board which, after integrating *all* knowledge, issues its orders" (Hayek, 1945). For the problem under our consideration, otherwise known as the graph colouring, the complexity of finding such a colouring—where no two neighbours share the same colour—is exponential in the size of the network (Garey, Johnson, and Stockmeyer, 1974). Still, for smaller instances of 30 agents, on which we perform our experiments, one can precompute the colourings in the network. We let the game unfold for some time, after which agents are at or close to equilibrium, and communicate privately to a small fraction of players a suggestion, or seed: their colour in one of the precomputed optimal configuration. The suggestion is non-binding and non-incentivised. Players may disregard it and do not receive any additional payoff from following it.

We find improvements of the play from control rounds with no suggestion to rounds with seeding. The result—up to the complexity of central planning which we largely eschew in our experiment—shows that a centralised authority can help in some instances nudge the game to more efficient equilibria, even when agents have their own incentives. Experiments confirm the importance of the topology of the network and the placement of the receivers for the efficacy of the suggestion. These results constitute the first investigation in this work of the tension between efficiency and decentralisation.

1.2 How efficient are learning agents following no-regret algorithms?

The previous experiment tests the impact of coordination and centralisation on reaching efficient NE. There, the structure of the game—a potential game—made it easy to find NE via an approximate best response dynamics followed by the players. In general games, there is no guarantee that such dynamics converge to a NE or some approximation of it. However, there is a weaker class of equilibria for which dynamics relying only on the information of one's own accumulated costs converge to.

These equilibria are coarse correlated equilibria (CCE) (Young, 2004), which are a superset of NE. A very strong relationship exists between CCE and no-regret learning dynamics: No-regret dynamics are known to converge to the set of CCE, and any particular CCE can be made to be the target of a sequence of profiles that is no-regret for all players. An algorithm is said to have no regret if looking back at a sequence of decisions made by the algorithm and outcomes resulting from these decisions, no fixed move could asymptotically outperform the cost of the obtained outcomes.

The convergence to NE is problematic and seemingly the class of no-regret dynamics offers a way out. It may however be of interest to understand how players can force convergence to NE, even in the presence of malicious agents, while maintaining low regret. The convergence is not natural (since it cannot beat the PPAD[4] hardness of Daskalakis, Goldberg, and Papadimitriou (2006) in any case), and implies a lengthy communication of payoffs as well as a phase of tit-for-tat system of threats and countermeasures to deviations (Section 4.2).

Instead, we could attempt to obtain guarantees on the performance of no-regret learning algorithms. The price of anarchy is concerned with the efficiency gap between the worst NE of a game and its social optimum. Roughgarden (2015), with the definition of $(\lambda, \mu)-$ smooth game, showed a natural extension of the notion to the gap between the worst CCE and the SO. The gap between the best CCE and the SO is much less understood, and gives a reasonable measure of just how much improvement can be obtained by allowing players to follow no-regret strategies instead of focusing on reaching a NE. We introduce respectively the value and the price of learning, measuring the distance between the best (resp., worst) CCE and the best (resp., worst) NE (Section 4.3).

1.3 What does a large scale experiment on Singapore's routing network tell us about the efficiency of the system and the regret of its users?

The price of anarchy finds its first measurement in the famous routing game of Pigou (1920) (Koutsoupias and Papadimitriou, 1999). There, a planner must decide how to route a flow of mass 1 through two links (Figure 1.2a). The first link costs exactly the amount of flow that traverses it. The second link costs always 1 to traverse. What is the minimum attainable cost for the flow? The answer is $\frac{3}{4}$. The planner can split the flow in half, with the half on the constant cost link incurring a total of $\frac{1}{2} \times 1 = \frac{1}{2}$ units, and the half on the variable cost link incurring a cost of $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$ units, for a total of $\frac{3}{4}$ (Figure 1.2b).

[4]Section 2.5 provides precise definitions and results on the complexity of finding a NE

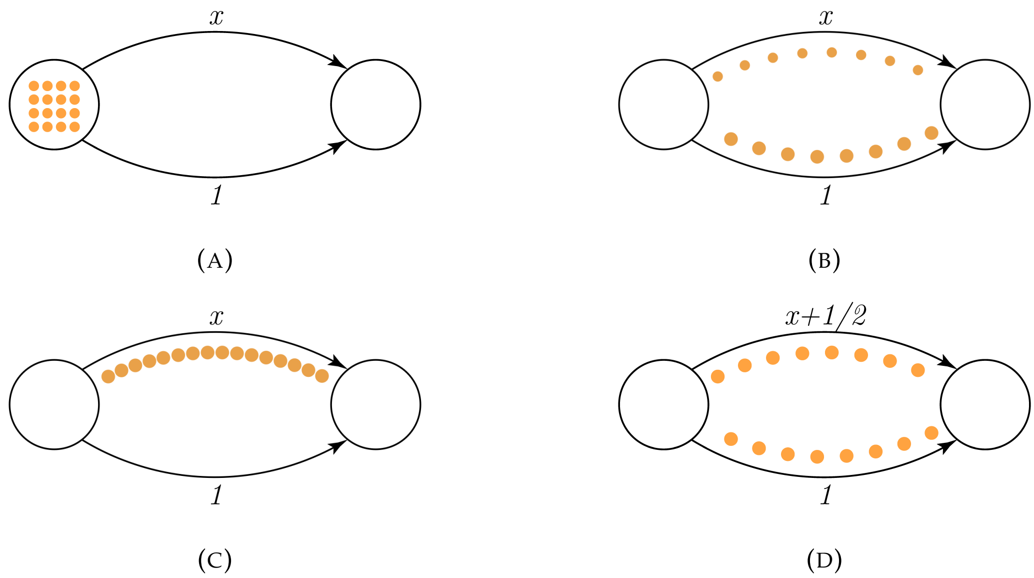

FIGURE 1.2: (a) Pigou (1920) introduced this game as a model for the congestion of "carts" on the roads. A unit demand of flow wishes to go from the left node to the right node. The upper link has latency equal to the fraction of the flow traversing it, while the lower link always has latency 1, irrespective of the flow routed on it. In the following figures, circles represent players, with radius proportional to the cost they incur (latency + price). (b) A central planner wishing to minimise the sum of all costs in the population would route half of the flow in the upper link (with latency is 1/2 and total cost 1/4) and half in the lower link (latency is 1, cost is 1/2). The total cost is then 3/4. (c) At equilibrium however, when nonatomic agents in the flow decide their own route, everyone would use the upper link, for a total cost of 1. The price of anarchy is the ratio of equilibrium cost to optimal cost, equal to 4/3 here. (d) Adding a toll of 1/2 to the upper link, where the cost for an agent is now price + latency, the planner can induce an efficient equilibrium where half of the users now use the upper link and half the lower link. However, we will see in Chapter 6 that this intervention has implications for the distribution of wealth among players (Theorem 6.2.1, the *Inequity Theorem*).

What if each infinitesimal unit of flow is controlled by an agent minimising her own cost? The previous split is untenable: Agents using the constant cost link have the incentive to deviate to the opposite link, where they will experience a cost of $\frac{1}{2}$. The only sustainable configuration (a Nash equilibrium) has all agents from the constant cost link moving to the variable cost link. In that case, the total cost of system (or social cost) is $1 \times 1 = 1$. The gap between this NE and SO is then $\frac{4}{3}$ (Figure 1.2c).

This gap turns out to be an extremely general result: Roughgarden and Tardos (2002) prove that for any network topology with affine link latency (as they were in the previous Pigou example), the most inefficient NE cannot be further than $\frac{4}{3}$ times the cost of SO, with the Pigou network providing the tightness of the bound. For link latencies that are general polynomials, one can derive additional bounds, all finite. The concept can be extended to various settings, such as atomic agents (i.e., a finite number of agents controlling one unit of flow each) (Christodoulou and Koutsoupias, 2005) or risk-averse agents who traverse edges with stochastic latencies (Ordóñez and Stier-Moses, 2010; Nikolova and Stier-Moses, 2015).

By its very nature, the theoretical results pertaining to PoA have all expressed a pessimistic, worst-case view of the system. Intuitively, the notion may be ill-suited to real systems, and refinements of PoA show indeed that average-case efficiency gaps can differ from their worst-case sibling (Panageas and Piliouras, 2016). The meaning of its definition is also questioned by Colini-Baldeschi et al. (2017), who show that regimes of asymptotically light and asymptotically heavy congestion both yield a PoA close to 1 under general conditions, which include typically assumed conditions for road networks—namely, quartic latency functions.

All of these considerations constitute the backdrop to the large experiment conducted in Singapore that will animate our discussion. The data, collected over several months by providing sensors to students which log their location and environmental parameters frequently, is an unprecedented empirical survey of mobility patterns. It is supplemented with algorithmic methods to obtain accurate information from the raw data. As we will show, some of the theoretical ideas we enunciated previously do not map one-to-one but inspire data-driven measurements of the system to uncover its properties: equilibration, efficiency and regret.

1.4 Does increased efficiency naturally lead to increased inequality for agents in congestion games?

Singapore is famous for pioneering innovative congestion control mechanisms. With the introduction of the Area Licensing Scheme (ALS) in 1975, it effectively cordoned off the city centre with an array of tolls collected by officers from drivers entering the area. ALS was replaced by the Electronic Road Pricing System (ERP) in 1998, a system of gantries automatically levying the toll from an onboard unit in the drivers' vehicles. The price of the road varies according to the time period and is updated every three months to meet target speeds under the gantry, with speeds lower than target resulting in an increase of the price and vice versa.

Tolls are complemented by a control of the demand for vehicles with the Certificate of Entitlement (COE) scheme, for which drivers must bid to obtain a 10 year license allowing them to drive the vehicle. As of the time of writing this thesis, the licenses are given to exactly compensate for the number of de-registered vehicles, thus keeping the number of registered vehicles constant.

That tolls and demand control mechanisms are efficient to reduce congestion is clearly established in the literature, from the Pigouvian tax/marginal cost pricing (Figure 1.2d) to nonatomic routing games with type-specific costs (Fleischer, Jain, and Mahdian, 2004). In parallel, a large body of evidence relating to the fairness of congestion pricing has shown the "winners and losers" of such schemes, how the revenue levied from the tolls can be recycled for other projects and the impact on fairness of the procedures (Levinson, 2010). But the intersection between these two branches has not been mined, from the computational results of algorithmic game theory practitioners giving few indications of how the mechanisms they advocate for impact inequality between agents.

Suppose now that heterogeneous users, with different valuations of their time, enter the system. We define the outcome for each agent as its initial wealth minus the cost incurred in the game, weighted by a small constant α representing *the importance of the game for the agent*. We then introduce the *inequity* of the game: as this constant α tends to 0, how does the inequality of wealth vary? In other words, what is the marginal impact of the game on inequality?

We find a general theorem for nonatomic symmetric congestion games showing that the inequity is positive, meaning that the inequality of income increases after the game is played, as measured by the Gini coefficient (Theorem 6.2.1). The theorem can be generalised to any inequality measure that respects four fundamental axioms of inequality indices (Theorem 6.2.2). For the asymmetric case, where agents do not share the same source and destination, we offer some counterexamples building on the inequity theorem of the symmetric case and particular inequality measures that offer decomposition properties (Section 6.3). The inequity, as defined by the marginal effect of the game on some inequality index, is shown to have useful properties when a natural choice for cost functions is made (Section 6.4).

The lesson is that if the quest for efficiency does not necessarily exacerbates inequality among agents, it certainly has a non-neutral impact on inequality in the game. In times of ever greater polarisation of resources and wealth (Piketty, 2014; Milanovic, 2016), it is crucial to understand exactly how the mechanisms that govern our interactions fare with regards other than their efficiency.

Organisation of the thesis

In the next section, we define the game-theoretic concepts that will be needed for the remainder of the thesis. We then dedicate one chapter for each of the four questions.

Chapter 2

Game theory concepts and definitions

In the preceding introduction, we presented some of the concepts that will animate our discussion throughout the four questions. In this chapter, we give a precise definition to these concepts. Later in the text, we supplement the exposition with additional definitions when necessary.

2.1 Definition of a game

The work will be concerned with *games* Γ played by a set of agents $\mathcal{A}$. This set may be continuous or discrete, but let us start with a finite set of agents $\mathcal{A} = \{1, \dots, N\}$ for some integer $N > 0$. We write $i \in N$ to mean that agent i belongs to the set of players $\{1, \dots, N\}$.

The description of the game includes the set of strategies available to each agent i, written S_i unless specified otherwise. This set is finite for all $i \in N$. A *strategy profile* is a vector $s = (s_1, \dots, s_N) \in \mathcal{S} = \prod_{i=1}^N S_i$ specifying the choice of one strategy for each agent. We write s_{-i} to mean the $(N-1)$-dimensional vector of strategies of all agents except for i, i.e., $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N)$.

Agents incur a cost given by their own cost function, which depends on the current strategy profile. Let $c_i : \mathcal{S} \to \mathbb{R}$ be the cost function of agent i, from the set of strategy profiles to the set of real numbers. The objective of all agents is to minimise their own cost. The dependence on the strategy choices of the other players implies that to minimise her own cost, agent i must consider the actions of the other agents.

The rules of the game—its set of agents, strategies and cost functions—are assumed to be common knowledge among the players, unless otherwise stated.[1] This implies that all players know each other player's strategy set and cost function, and all players know that all players know each other player's strategy set and cost function, etc.

[1] This will not be the case in Chapter 4.

2.2 Nash equilibrium

Given a game, how should the agents be expected to play? The search for such an answer brought the *Nash equilibrium* (Nash, 1950). A strategy profile $\bar{s}$ is a NE if no player can profitably deviate unilaterally from $\bar{s}$.

Definition 2.2.1. $\bar{s}$ is a NE if the following condition is verified:

$$\forall i \in N, \forall s_i, c_i(\bar{s}) \le c_i(s_i, \bar{s}_{-i}) \tag{NE}$$

In general, a game Γ may possess several NE. We present such a game, historically known as the *Battle of the Sexes*. Alice and Bob have identical strategy sets composed of two actions: They can either go to the Museum M or to the Park P. They are happy when they coordinate such that both choose the same activity, but each has a preference for a different activity. Indeed, Alice enjoys the Museum better than the Park, while the opposite is true for Bob. This information is summarised by the following cost matrix:

$$\begin{array}{cc} & \begin{array}{cc} M & P \end{array} \\ \begin{array}{c} M \\ P \end{array} & \begin{pmatrix} 0,1 & 2,2 \\ 2,2 & 1,0 \end{pmatrix} \end{array}$$

In this example, Alice is the row player and Bob is the column player. The strategy profile (M, M) yields cost 0 for Alice (who enjoys the Museum) and cost 1 for Bob (who enjoys the Park more, but likes to be with Alice). It is easy to see that this game has two NE, namely (M, M) and (P, P).[2] In either of these profiles, if one player deviates, both end up doing separate activities and thus the deviating player's cost increases.

2.3 The price of anarchy

The notion of NE gives a first solution concept for games of self-interested agents. NE shows that *incentives* are what matters to the agents. Let us analyse the following example, the *Prisoner's Dilemma*, to understand what this means in practice.

Two robbers are arrested and placed in different cells with no communication. The police officer on patrol does not have incriminating evidence, but offers the following bargain.

- If neither of them cooperates with the officer, they are both given a default sentence of one year due to lack of evidence.

[2] And another, distinct, mixed equilibrium (Section 2.4).

- If both of them cooperate, each incriminating the other, there is enough evidence to send both in prison for three years.
- If only one of them cooperates, the "snitch" will be given a special deal from the prosecutor and spend no time behind bars, while his loyal companion will be given four years.

The cost matrix is as such:

$$\begin{array}{cc} & \begin{array}{cc} N & C \end{array} \\ \begin{array}{c} N \\ C \end{array} & \begin{pmatrix} 1,1 & 4,0 \\ 0,4 & 3,3 \end{pmatrix} \end{array}$$

where N is the action of "Not cooperating" while C is that of "Cooperating".

In this classical example, the only NE of the game is (C, C). Although this equilibrium is obviously suboptimal, since both would decrease their cost if they picked (N, N) instead, *neither of the agents has any incentive to not cooperate with the officer*. The definition of *suboptimal* is clear in this example, but let us define precisely how to measure the optimality of a strategy profile.

Definition 2.3.1. The *social cost* of a profile s is given by:

$$\mathtt{SC}(s) = \sum_{i=1}^{N} c_i(s).$$

The *social optimum* of a game Γ is the minimum value of the social cost.

$$\mathtt{SO} = \min_{s \in \mathcal{S}} \mathtt{SC}(s)$$

One can think of the social optimum as being enforced by a centralised planner who has the power to select one strategy for each of the agents. We see in the Prisoner's Dilemma the premise of the question that will agitate most of our work: the discrepancy between the equilibria of self-interested agents and the social optimum as dictated by some central planner. This gap was given the name of *coordination ratio* by Koutsoupias and Papadimitriou (1999), and subsequently renamed **price of anarchy** in Papadimitriou (2001).

Definition 2.3.2. The *price of anarchy* (PoA) of a game Γ is the ratio of the social cost of the worst Nash equilibrium (i.e., the NE with the highest social cost) to the social optimum (SO).

$$\mathtt{PoA}(\Gamma) = \max_{\bar{s} \in \mathtt{NE}(\Gamma)} \frac{\mathtt{SC}(\bar{s})}{\mathtt{SO}} \tag{PoA}$$

2.4 Mixed strategies

We have so far—intentionally—focused on the choice of *pure* strategies by the agents, where players decide on *one* action from their set of strategies. But if we restrict the players to choosing only pure strategies, we are not guaranteed that a NE exists for any finite game Γ. Hence, we must allow players to randomise the choice of their strategies.

Definition 2.4.1. Call $\Delta(S_i)$ the set of mixed strategies of agent i, i.e. the $(|S_i| - 1)$-simplex $\Delta(S_i) = \{(p_j)_{j \in S_i} | \sum_{j \in S_i} p_j = 1;\ p_j \geq 0,\ \forall j\}$.

Theorem 2.4.1. *For any game Γ with a finite number of agents using mixed strategies over a finite set of actions and bounded cost functions, there exists a Nash equilibrium.*

Mixed strategies will be revisited in Chapter 3, with simulations implementing learning dynamics over the space of mixed strategies. In Chapter 4, we devise a protocol for players to check if other players follow truthfully a given mixed strategy, while maintaining no-regret.

2.5 Complexity of finding a Nash equilibrium

A game Γ with costs in $\mathbb{Q}$ can be represented by its strategy sets and the cost function of every player. Assuming we need b bits to encode the costs, and s bits to encode the strategies (with $\max_{i \in N} |S_i| \leq \log_2(s)$), we require up to $(sb)^N$ bits to encode the full cost matrix of all players. This is the input size of the game.

An algorithm is said to be of complexity $O(f(n))$ if its execution is guaranteed to end before some bound $c \cdot f(n)$ for an input of size n. When $f(n)$ is a polynomial function of n, the algorithm can be executed in polynomial time. If no such polynomial exists, the algorithm is executed in exponential time (Garey and Johnson, 2002).

Polynomial time computation is related to a class of problems named P. These problems are yes/no questions, where for an input of size n, one can compute a boolean property of the input in polynomial time. For instance, the problem "Is there a positive integer m such that $N = 8m$?" for some positive integer N can be solved in time polynomial in the size of N, and thus belongs in P.

P is included in a larger class of problems, NP, for which one may not be able to efficiently (i.e., polynomially) find a proof (i.e., an element that satisfies the question, as m in the previous example), but can efficiently verify the proof. For the graph colouring, which is the backdrop of Chapter 3, the problem *Can the graph $G = (V, E)$ be coloured with m colours?* is in NP for $m > 2$. Given a colouring, one can efficiently verify that it is a proper colouring of the graph by iterating over the edges and checking that the endpoints do not share the same colour.

It is not known whether P = NP, although the alternative is preferred by most experts. In this case, there are problems which belong to NP, but not to P. Satisfiability (SAT) is the first problem shown to belong to this class, called NP-complete, by Cook's theorem (Cook, 1971). If a problem L in NP can be reduced to SAT by a polynomial time algorithm, that is, if there is a map from L's input to an instance of SAT, then L is also NP-complete.

Daskalakis, Goldberg, and Papadimitriou (2006) proved that finding a NE of a game Γ belongs to the PPAD class. Indeed, we know by Nash's theorem that every finite game has at least one (possibly mixed) NE. This simplifies the problem somewhat, in contrast with SAT for which we cannot know in advance if there is a circuit satisfying its input predicates. PPAD instances "look like" a total graph search, where the graph is of size 2^n, one for each bit string of length n. This implies that finding a NE is hard, in the sense that provided P $\neq$ NP, PPAD-complete problems are intractable.

2.6 Price of anarchy of routing games

In this section, we define routing games and provide an index of famous results on the price of anarchy in this class of games.

2.6.1 Routing games

Routing games are played on a network $G = (V, E)$, where V is a set of vertices and E a set of (directed) edges connecting vertices. Routing games admit a continuum of players (*nonatomic* routing games) or discrete players (*atomic* routing games). Surprisingly, bounds on PoA in each case are different from each other.[3] The following exposition focuses on nonatomic games, with references to the atomic case when appropriate.

Strategy sets The strategy sets of the players are acyclic paths connecting an origin to a destination. If all players share the same origin and the same destination, the game is called *symmetric* or single commodity. Otherwise, if several origin-destination pairs exist, the game is *asymmetric* or multicommodity. We assume the game is asymmetric and has K such pairs, each with demand $\mu_k > 0$. The strategy set is thus $\mathcal{P} = \cup_{i=1}^{K} \mathcal{P}_k$, where $\mathcal{P}_k$ is the set of all paths connecting the origin and destination of commodity k. A profile (one strategy for each player) determines a flow $f \in \mathcal{P}^{\mathbb{R}}$, returning the demand for path $P \in \mathcal{P}$.

[3]But an atomic routing game with vanishingly small players can be equated with its nonatomic counterpart (Feldman et al., 2016).

Cost functions Edges in the network are associated with a cost function c_e, mapping the congestion on the edge, z, to a non-negative cost. The congestion is exactly equal to the number (or mass) of agents traversing the edge and can be obtained by summing over all paths including the edge. Thus, the cost function is assumed to be non-decreasing in the congestion. The cost of path P under flow f is equal to

$$c_P(f) = \sum_{e \in P} c_e(f_e)$$

where $f_e = \sum_{P \ni e} f_P$ is the congestion on edge e under flow f.

Nash equilibrium A flow $\bar{f}$ is a Nash equilibrium if, for any commodity k and for all paths $P, P' \in \mathcal{P}_k$ such that $\bar{f}_P > 0$,

$$c_P(\bar{f}) \leq c_{P'}(\bar{f})$$

2.6.2 Social cost of a flow

Introduced previously, the price of anarchy is concerned with the efficiency gap between the worst equilibrium of a game and the social optimum. We first need to define the social cost to measure the cost of a flow.

Cost of a flow By summing the costs over the population of players, we obtain the social cost. By equivalence, we can sum over the paths or the edges to obtain the same and thus define

$$\texttt{SC}(f) = \sum_{P \in \mathcal{P}} f_P c_P(f) = \sum_{e \in E} f_e c_e(f_e)$$

Proposition 2.6.1. *If $\bar{f}$ and $\bar{f}'$ are equilibria of a game Γ, then $c_e(\bar{f}) = c_e(\bar{f}')$ for all $e \in E$ and $\texttt{SC}(\bar{f}) = \texttt{SC}(\bar{f}')$.*

Social optimum The socially optimal flow f^* satisfies

$$\texttt{SC}(f^*) = \min_{f \in \mathcal{F}} \texttt{SC}(f)$$

where $\mathcal{F}$ is the set of feasible flows (i.e., flows respecting demand constraints).

TABLE 2.1: Price of anarchy for some classes of cost functions (nonatomic games).

Class	Nonatomic games
Affine costs	4/3
Quadratic costs	1.626
Quartic costs	2.151

2.6.3 Routing games are potential games

A game Γ is called a potential game if there exists a function Φ over game profiles such that, informally, a social cost improvement due to a deviation corresponds to a comparable improvement of the potential. If the improvement of the potential is exactly equal to the improvement in the social cost, the potential is called *exact*.

By defining the following function $\Phi : \mathcal{F} \to \mathbb{R}$, a routing game can be shown to admit Φ as a potential function.

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} c_e(x)dx$$

Additionally, one can prove that equilibrium flows are global minima of the potential function, leading to an optimisation-driven definition of equilibrium.

2.6.4 The price of anarchy of routing games

We define the price of anarchy of a game Γ by

$$\texttt{PoA}(\Gamma) = \frac{\texttt{SC}(\bar{f})}{\texttt{SC}(f^*)},$$

where $\bar{f}$ is an equilibrium flow and f^* is a socially optimal flow.

The price of anarchy can be defined over a class of games, represented by their cost functions. For instance, if all cost functions in Γ belong to some class $\mathcal{G}$, we write $\Gamma \in \Gamma_{\mathcal{G}}$, and the price of anarchy over $\Gamma_{\mathcal{G}}$ is

$$\texttt{PoA}(\Gamma_{\mathcal{G}}) = \max_{\Gamma \in \Gamma_{\mathcal{G}}} \texttt{PoA}(\Gamma).$$

Examples of such classes include linear latencies, quadratic latencies and more generally, any polynomial of the congestion on the edge. We provide in Table 2.1 the bounds obtained for some classes.

2.6.5 (λ, μ)-smooth games

Arguments for the bounds in Table 2.1 were recently collected under the framework of (λ, μ)-smoothness (Roughgarden, 2015).

Definition 2.6.1. A game is (λ, μ)-smooth if for profiles s and s',

$$\sum_{i=1}^{N} c_i(s_i, s'_{-i}) \leq \lambda \cdot \texttt{SC}(s) + \mu \cdot \texttt{SC}(s')$$

With λ and μ obtained to make the inequality as tight as possible, PoA bounds follow easily. These bounds are robust to no-regret learning and hence translate to the larger class of CCE (and, by extension, to mixed NE). However, the additional regret term can be arbitrarily large for a long period of time, as explored in Section 4.4.

2.7 A note on terminology

This short section disambiguates the use of terms important to this thesis. We have introduced equilibria and optimum in Section 2.3. By **efficiency**, we mean the social cost $\texttt{SC}$ of a configuration, i.e., the sum of all costs incurred by all agents in the system. **Efficient equilibria**, or **optimal equilibria**, are equilibria which also mimimise the social cost. When we speak of **decentralisation**, we mean that agents are responsible for their own decisions, and a central authority cannot compel an agent to follow a different choice (but can enforce mechanisms in the game, e.g., tolls).

"Move IN
EVERYONE WINS"

LTA Graciousness campaign for bus and train commuters.

Chapter 3

When does sparse information seeding induce efficient equilibrium selection?

Chapter 1 introduced the broad themes of efficiency, decentralisation and the tension between the two. The selfish decisions of agents can produce outcomes which are suboptimal, e.g., in the Pigou network of Figure 1.2. In some cases, it is possible for a designer to implement mechanisms correcting for this inefficiency, while maintaining decentralisation, for instance using tolls.

But inefficiency, in general, is not only due to selfishness. The difficulty of reaching efficient states is compounded by the possible existence of many Nash equilibria in a game, with varying costs. Agents engaged in a game where some of the equilibria are efficient may simply be unable to reach an optimal state from lack of coordination and get stuck in suboptimal equilibria instead.

In routing games with a positive price of anarchy, the optimal flow is not stable with respect to decentralisation. Letting agents choose their own route will invariably lead to a more inefficient configuration.[1] A central authority thus cannot improve the situation unless it directly changes the incentives of the agents in the game. But if efficient equilibria exist, the same central authority could potentially coordinate agents with meaningful suggestions.

This is the setting of our first experiment, which provides evidence that sparse information seeding from a centralised messenger can improve coordination among the agents and lead the play towards better states. In some sense, its assumptions are stronger than any other in the thesis: efficient equilibria exist, and so does a central authority which has access to sufficient computing power to find these equilibria and can communicate privately to some players.

An approach such as public service announcements (PSA) is often used in practice to coordinate a system, but the conditions for its success are not entirely understood

[1] This is true for a broad class of dynamics (Monderer and Shapley, 1996; Shah and Shin, 2010).

when agents act in their own interests as they do in this experiment. Theoretical research on the topic (Balcan, Blum, and Mansour, 2009; Balcan et al., 2014) has only established broad results whose intuition may not carry over to real systems. Our experiment is played over social networks whose properties are shown to have a significant impact on the outcomes.

Indeed, the experiment features groups of subjects embedded in social networks who are individually incentivised to differentiate themselves from their neighbours, with an optimal system-wide solution being a colouring of the graph (Garey, Johnson, and Stockmeyer, 1974). In contrast with previous designs gauging the ability of the entire group to find such a colouring (Kearns, Suri, and Montfort, 2006; Judd, Kearns, and Vorobeychik, 2010), individual performance is scored. This design locates our experiment closer to models where agents are not concerned with the system reaching global optimum but minimise their own costs, a more natural formulation for multi-agent systems that overcomes the limitations of previous literature. However, a central planner who cares for the system-wide costs intervenes to guide the players towards efficiency.

After the game unfolds for a fixed period of time, usually sufficient for the players to approach a NE, a central messenger selects one of several solutions to the problem and communicates privately to 10 or 20 percent of players the action they would implement under the chosen configuration. Adopting the suggestion contained in the seed may not be individually rational for receivers. This setup is compared to one where players do not receive any information from the central messenger. Since the latter may provide suggestions that are detrimental to ones' own choice, the experiment tests the weakest possible form of information seeding: non-binding, unsubsidised and possibly individually irrational.

Contents of this chapter

We start with the description of the experiment in Section 3.1. The scoring rule is discussed as well as the choice of networks. More precisions are given on the manner in which the central messenger chooses the seeds to communicate to receiving nodes.

Our first results are obtained from the analysis of the experimental data in Section 3.2. The suggestion improves the player costs in some networks, but not others. Seeds also impact the stability of the play, with seeded rounds logging a lower update rate than unseeded ones. We also study the behaviour of receivers and find some evidence of costly deviations in the presence of the suggestion.

In Section 3.3, we provide additional results obtained from simulations of the play via both learning dynamics and a behavioural model obtained from experimental data. The position of the receivers is seen to yield more or less improvement depending on its centrality. Network statistics such as their clustering coefficient as well as the average distance further impact the efficiency of the seeds.

3.1 Methodology

3.1.1 Recruitment

240 students, researchers and staff from the Singapore University of Technology and Design (SUTD) were recruited between May 2017 and March 2018 to participate in one of 8 sessions. The first session, thereafter named "pilot session", was reserved for testing the experiment platform as well as refining the experimental parameter of round duration from 180 seconds to 120 seconds. During the second session, a technical issue – invisible to subjects – prevented the data collection during treatment rounds, leaving control rounds unaffected. Minor technical issues were recorded in further sessions, for which a small number of rounds were dropped. The exact numbers are provided in Table A.1.

In each session, 30 participants were recruited and could not sign up for any further session. Participants were given a SGD5 show-up fee, with the possibility of being awarded up to SGD20 more depending on their performance. The sessions were divided in four parts:

- Subjects are placed randomly in a room, arranged such that each subject has limited visibility to the other subjects' computers.
- Once everybody is seated, subjects start an interactive tutorial describing the rules of the game and the modalities of payment. Subjects must complete three questions throughout the tutorial to assess their understanding of the rules of the game and how their final payoff is computed.
- Once all participants have completed the tutorial, the number of rounds is publicly announced and the game moves to the first round. Between each round, subjects are presented with a loading screen and the index of the next round.
- At the end of the final round, the subjects' screens present the summary of their performance in each round with a bold row indicating the round – chosen randomly – that constitutes their evaluation. The payoff of this round is translated into the variable award, between SGD0 and SGD20. A final piece of information is their total winnings, including the SGD5 show-up and the variable award.

Each session, consisting of 30 subjects, played between 15 and 18 rounds of the graph colouring game, for a grand total of 95 rounds played (see Table A.1).

3.1.2 Description of the game

At the start of a round, all 30 players were randomly assigned one node of a network and provided with a budget of 100 points. All nodes were given the same initial colour. During the play, which lasted for 120 seconds, a player was able to change their node colour among 3 options, including the initial colour, at any point, as many times as needed. The player observed at all times the colour of its immediate neighbours, i.e., nodes connected to her own in the network, and their real time colour changes.

We call *a match* an edge connecting two players who share the same colour. When a player has a match, points are deducted from its budget. The costs were scaled by the degree of their node and the duration of the round. For instance, at time t, if a player matched with 2 of her neighbours out of 5, the player incurred a cost of $\frac{2}{5} \times \frac{100}{120}$.

The graph was considered coloured, or, to have reached a colouring, if no two neighbouring nodes shared the same colour, in which case the round was stopped. Thus, individual costs accumulated as long as matches in the network were present. If players did not reach a colouring before 120 seconds, the round was stopped. Finally, players received a financial incentive proportional to their final cost, additional to the show-up fee (see Materials and Methods).

3.1.3 Choice of the networks

We focused on 4 of the most common network structures, specifically Barabási-Albert (*ba*) (Barabási and Albert, 1999), Erdős-Rényi (*er*) (Erdős and Rényi, 1959), stochastic blocks (*sb*) (Holland, Laskey, and Leinhardt, 1983) and Watts-Strogatz (*ws*) (Watts and Strogatz, 1998). For each type, we selected one network of 30 nodes (Figure 3.1), all 3-colourable.

These networks were chosen as adversarial instances for the graph colouring problem. A measure of hardness is defined in the following way. Random runs of best response dynamics (formally defined in Section 3.3.1) are simulated on a network, leading to equilibrium. The number of matches at equilibrium is then averaged over all runs to obtain a hardness measure. 200 networks from each type were generated and the hardest instance was selected for the experiment. As a result of the correlation between the number of colourings and this measure of hardness (Shirado and Christakis, 2017), each of the four networks had a low number of colourings, ranging from 24 for *sb* and *ws* to 48 for *er* and 120 for *ba*, among 3^{30} possible configurations (including permutations).

3.1.4 Choice of the suggestion

Participants were informed that during the play they may receive a static suggestion in the form of a fixed colour appearing in a special section of the game screen, however

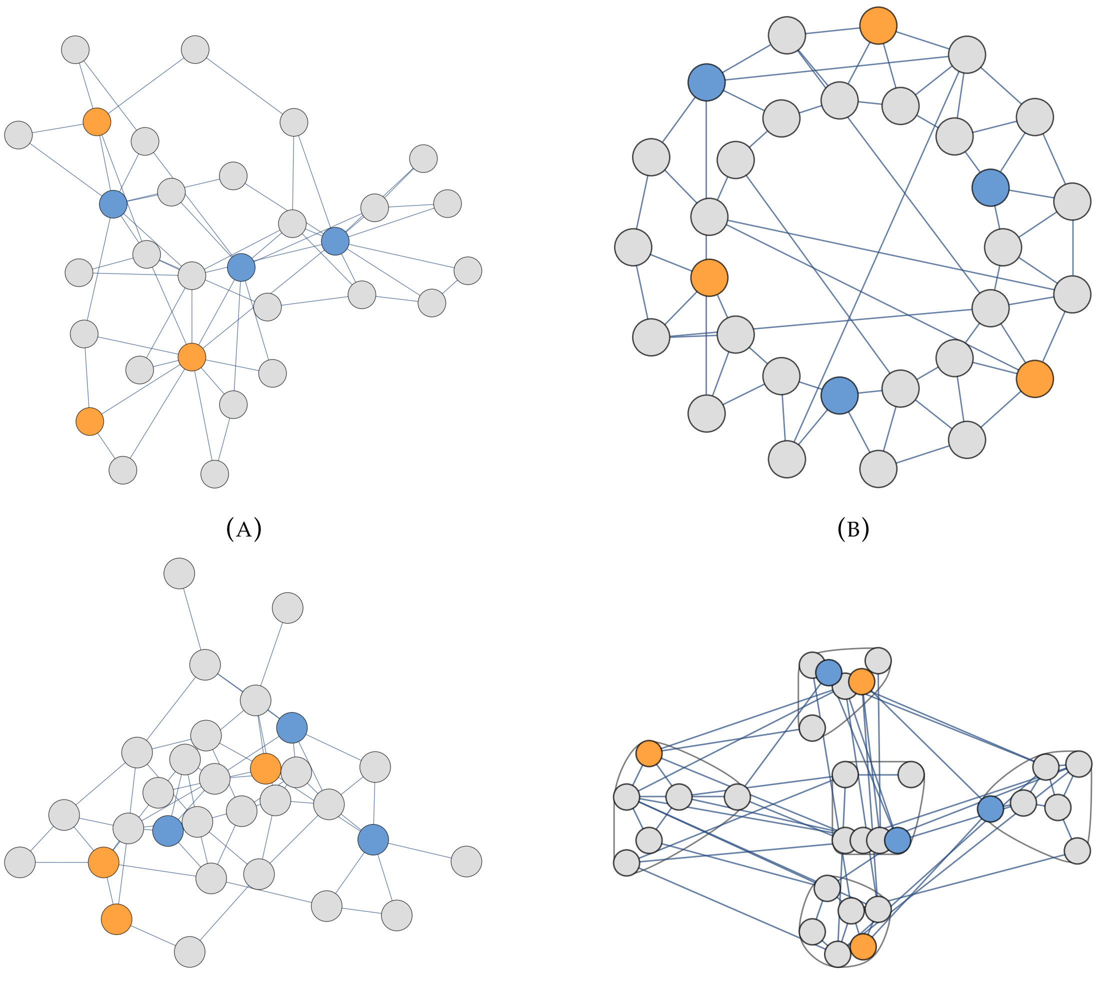

(A) (B) (C) (D)

FIGURE 3.1: The four experimental networks, chosen adversarially. Chosen nodes are coloured in blue or in orange. Receiver nodes in rounds with three seeds are in blue. Additional receiver nodes in rounds with six seeds are in orange. **a.** Barabási-Albert: A few nodes are well-connected (hubs), while others receive fewer neighbours (periphery), connected according to a preferential attachment dynamics. **b.** Watts-Strogatz: Each node is connected to its two nearest neighbours and its two neighbours two steps away, before edges are randomly rewired. **c.** Erdős-Rényi: An edge is drawn between two nodes with some probability p. **d.** Stochastic blocks: 5 families of 6 nodes each have more within-block connections than between-blocks connections.

following this suggestion was not directly incentivised. In the thesis, we use the terms "seed" and "suggestion" interchangeably. The suggestion was given after 20 seconds of play. Among all colourings of the network, the central messenger selected one that minimised some distance from the configuration at $t = 20$. The seed received by the participants was the colour prescribed by the selected colouring for their node.

From the set $\mathcal{C}$ of known colourings for each network, the central messenger selects one to minimise the distance from the configuration c after 20 seconds. The messenger chooses the colouring $\hat{c}$ such that it would yield the minimum number of matches in the network if all receiving nodes were to follow the suggestion.

For instance, we assume the current configuration is $c = (r, r)$, where the colours allowed are $\{r, y\}$ and the graph is a simple edge between two nodes. The first node of the pair is a receiver. Two colourings exist: (r, y) and (y, r), but the latter only would yield zero matches when the receiver follows its recommendation and is thus selected by the centralised messenger.

Unbeknownst to the players, six fixed nodes in each graph were called *chosen nodes*. During control rounds, none of the six chosen nodes received a suggestion. In the two levels of treatment, respectively 3 out of 6 and 6 out of 6 of the chosen nodes (respectively, 10% and 20% of the overall network nodes) received a suggestion, in which case they were called *receiver nodes*. The 3 nodes in the first level of treatment were fixed. Players were blind to the experimental condition of the round, except for receivers who knew they were given a suggestion after 20 seconds of play—but ignored the overall level of suggestion or the network type.

3.2 Experimental results

The suggestions, though sparse and non-binding, decrease player costs under some conditions. Indeed, the cost improvement varies in the four networks, ranging from significantly lower costs to no effect. Second, introducing seeds decreases the update rates of the players and increases the stability of the system. This finding contrasts with previous studies emphasising the importance of noise to move the system to better states (Shirado and Christakis, 2017). Between the networks for which seeds improved equilibration and those where the seeds did not, the most contrasting effect lies in the *quality* of the suggestion, i.e., whether it advertises an individually rational action or not. The quality of the suggestion may explain how reduced exploration entails greater coordination in that case. To explain these variations, we find that the group betweenness centrality (GBC) of the receiver nodes, relative to the average GBC of any set of nodes with equal size, correlates with the efficiency of the suggestions in each network. Furthermore, simulations on additional random networks show the impact of

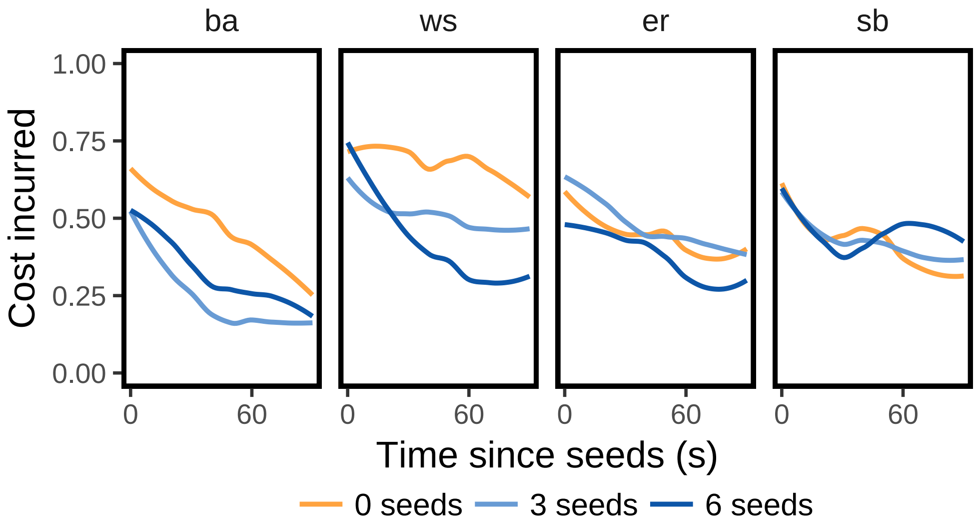

FIGURE 3.2: Throughout the game, players accumulate costs proportional to their number of matches with neighbours. The light blue line (resp. dark blue line) corresponds to the average incurred cost in rounds with 3 seeds (resp. 6 seeds), while the orange line charts the average in the control rounds. The plot gives the average cost after seeding, because all players have the same colour at the start of the game and thus incur a high cost, thereby distorting the scale of the plot. Further, since there is no experimental difference between the first 20 seconds in control and treatments, these are omitted from the plots. Only Barabási-Albert and Watts-Strogatz networks present a significant improvement from 0 to 3 and 6 seeds.

network topology—namely, their clustering coefficient and their average distance—on the cost improvement.

3.2.1 The effect of suggestion on player cost

We consider a measure of partial solution to the problem, embodied by player costs, i.e., the loss of the players from the original 100 points budget given in each round. The series of player costs does not follow a normal distribution after rescaling (Shapiro-Wilk test, $W = 0.86$, $p < 0.01$) and the variances are significantly different in the three groups, $F(2, 2877) = 4.28$, $p < 0.05$.

Figure 3.2 plots the average cost incurred by the players throughout the game. Before receiving seeds from the centralised messenger, average costs are not significantly different between the three conditions and are thus left out of the plot. After reception of the seed, for two networks (*ba* and *ws*), the cost diverges significantly in both treatments. The improvement from 10 to 20 percent of receivers is however not significant for any network type, pointing to a marginally decreasing effect of the treatment.

How are these costs incurred? Players update their colours until a colouring is reached or time runs out, the former option occurring in 9 rounds out of 95—less than 10% of the samples. Thus, some players still incur costs throughout the round. Figure 3.3a reports the high amount of time spent at or near equilibrium, where no player has any strict incentive to deviate from her current colour choice. The stability of inefficient equilibria represents an obstacle to cost improvement and motivates the introduction of seeds to coordinate the players' actions.

3.2.2 Seeds increase system stability

The update rate is computed as the number of updates by a node in a given time period divided by the length of that period (in minutes). Before the seed is provided, the update rate is not significantly different between control and treatment groups, with average rates close to 6.1 updates per minute. However, after broadcasting the seed, the update rates decrease significantly between rounds with suggestions (the average rate is 1.15 for both 3 and 6 seeds) and rounds without (the average rate is 1.49). The contrast is stronger in networks where the suggestion has improved equilibration (Figure 3.3b).

There is no reason for a node to keep updating if it is already achieving its minimum cost, i.e., is in *best response*. Thus, with a greater number of players in best response, we observe longer durations between two moves in the network (Figure S1, SI). As players in *ba* and *ws* are able to reach a more efficient equilibrium in a shorter amount of time (Figure 3.2), the update rate may decrease as a consequence. In Figure 3.3c, the average player cost per second is plotted against the update rate after seeding for each round.

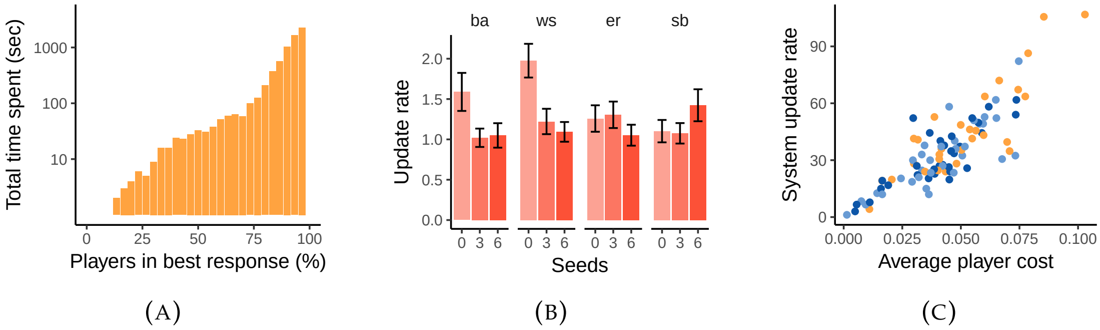

FIGURE 3.3: **a.** A player is in best response if her current colour choice minimises the number of matches with her neighbours. For a given percentage of players x, we plot the total number of seconds over all rounds spent with x percent of players in best response. We observe nodes spend most of the time at or close to equilibrium, where *all* players are in best response. The y-axis is scaled logarithmically. **b.** Average update rates of all nodes after seeding. Only *ba* and *ws* register a significantly lower update rate with respect to the control rounds. **c.** The x-axis represents the average player cost, i.e., the average of the cost incurred during one second by a player. The y-axis gives the system update rate, or the number of colour changes by all nodes divided by the time between seeding and the complete round of 120 seconds. Each round is associated with one point, with colour following the legend of Figure 3.2 (orange = no seed, light blue = 3 seeds, dark blue = 6 seeds). Control rounds appear more frequently towards the top right section of the plot.

The two variables are significantly correlated while a larger share of control rounds is found in the top right quadrant of the chart, indicating higher average costs.

3.2.3 Behaviour of receivers

Twenty seconds into the game, receivers see a suggestion appear on their screen accompanied with a colour choice. The label states: "*This is the colour of your node in a proper colouring of the graph. You are free to follow or disregard this hint. This will not have any direct impact on your score.*" Some receivers are advised to play their current choice, if the selection algorithm deems that it should be their choice in the closest available colouring. Figure 3.4a charts the aggregate behaviour of receivers throughout the game with respect to the seed.

For receivers in the two networks that present a clear improvement over control, *ba* and *ws*, the seed is a best response for a significantly higher portion of the time than in *er* and *sb*. Conversely, the amount of time spent in the suggested colour by receivers in *ba* and *ws* is significantly higher than that of receivers in *er* and *sb* (Figure 3.4b). When more than one colour choice (including the suggestion) is a best response, the seed is played an average of 70% of the time, showing a strict preference for following the suggestion even in the presence of equivalent options.

The seed communicated to receiving nodes may advertise a colour that is not individually rational for the player to follow. Yet, some receivers act against their own

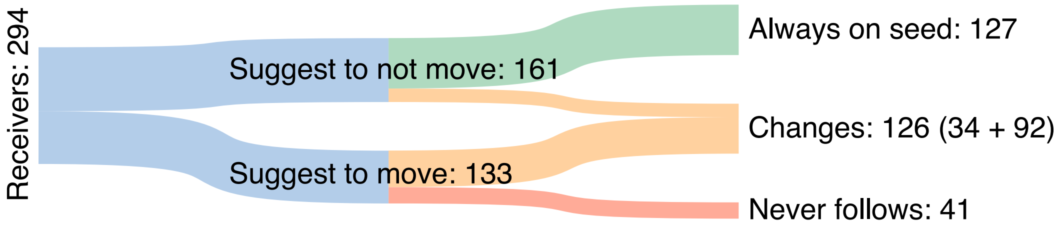

(A)

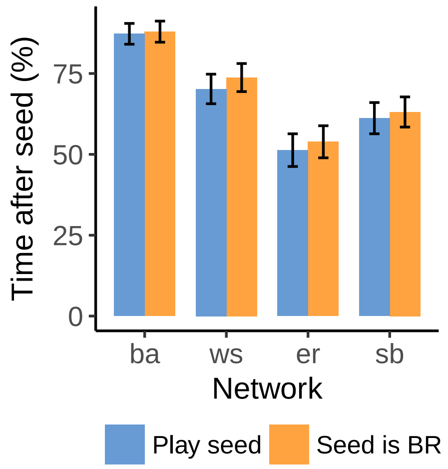

(B)

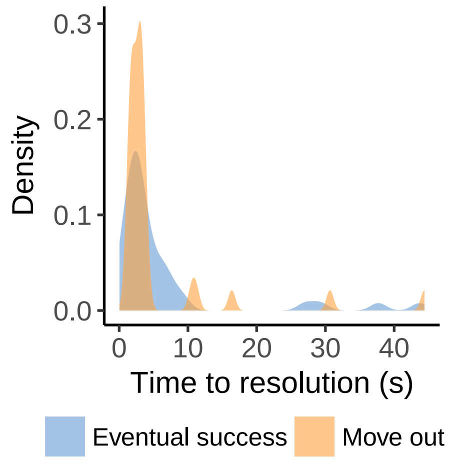

(C)

FIGURE 3.4: **a.** Sankey diagram of receivers. Over half of the seeds suggest the receivers to keep playing their current colour choice. Among those receiving such a seed, over 75% of them keep playing this colour until the end of the game. **b.** For each network, blue columns represent the fraction of time during which the seed is a best response and orange columns represent the fraction of time spent playing the seed. The two values are highly correlated with each other. **c.** For receivers acting against their own self-interest, we plot the density of the time to resolution, where resolution is either brought by the environment changing to bring matches to a level lower than or equal to before the deviation (*eventual success*), or by the receiver flipping out of the seed (*move out*). Most resolutions are quick, under 10 seconds. Receivers seem to test their environment, whether following an *ex ante* unreasonable suggestion will eventually bring the cost down following neighbouring updates. They are quick to revert to a best response (*Med:* 2.9 seconds) otherwise.

immediate self-interest and choose to follow the suggestion even when it is not rational for them to do so. We count 80 instances of such behaviour and determine how and when this deviation is resolved. Three disjoint cases are present: either the receiver's neighbours update their colours such that the number of matches is lower than or equal than the matches before deviating (*eventually successful*, 42 cases); or the receiver moves out of the colour before this happens (*move out*, 33 cases); or the game ends before this happens (5 cases). With the exception of two cases, receivers are in best response before deviating to the non-best response seed, thereby strictly increasing their costs by following the suggestion.

Figure 3.4c shows for the two groups of receivers – eventually successful ones or those who move out – a histogram of the resolution time: how long before they recover from their deviation or before they decide to move out. The plot shows that altruistic receivers may be divided in two groups. Some are willing to wait only for a median time of 2.9 seconds before reverting to a best response if their deviation is not eventually successful. Fewer stay in non-best response for much longer periods, showing a stronger commitment to the suggestion. Meanwhile, most eventually successful deviations are resolved within a 5 to 6 seconds timeframe, indicating that they help the local neighbourhood of the receiver coordinate.

3.3 Impact of network topology on the efficiency

The preceding results obtained from the experimental data hint at a very strong relationship between the behaviour of the receiver and the cost improvements in the game. In this section, we show that the position of the receivers as well as the overall topology of the network both factor into the efficiency of the centralised suggestion. During the experiment, the position of the receivers was kept fixed for each network, and the procedure was only tested on one network (chosen adversarially) of each type. To study the impact of position as well as topology, we are required to simulate the game, which we do in two ways. First, we use two classic learning dynamics from multi-agent systems, best response and multiplicative weight updates. Second, we build a behavioural model of the players and agent-based simulations of the game.

We first introduce the various simulations before turning our attention to the position of the receivers and the topology of the networks.

3.3.1 Simulations of the play

Learning dynamics

The play is simulated with two learning dynamics, best response (BR) and multiplicative weights update (MWU). Each is run following this procedure:

- The game starts from an initial configuration.
- Dynamics are run until an equilibrium is reached.
- Seeds are sent out to players and the play profile is updated.
- Dynamics are run until an equilibrium is reached.
- We compare the number of matches in the first equilibrium with that of the second equilibrium.

Best Response (BR) A BR dynamics selects one-by-one uniformly at random in a player and allows her to change her action given the current choices of her neighbours. The selected action minimises her current cost and in the case where the set of minimising actions includes more than one element, an action is chosen uniformly at random. It is known (Nisan et al., 2007) that a sequence of best responses always converges to a pure Nash equilibrium—if one exists.

The initial configuration is chosen with all nodes sharing the same colour, i.e., a maximal number of matches in the graph, reflecting the experiment setup.

Multiplicative Weights Update (MWU) While BR is a dynamics over the pure strategies of the agent, MWU is a dynamics over their mixed strategies (defined in Section 2.4). Specifically, the probability $p_c^i(t)$ that player i chooses color c at time $t \in \mathbb{N}$ is governed by the following equation:

$$p_c^i(t+1) = p_c^i(t) \frac{1 - \gamma W(c,p)}{1 - \epsilon \sum_{\bar{c} \in \mathcal{S}} W(\bar{c}, p)} \tag{MWU}$$

where

$$W(c,p) = \sum_{\hat{c} \in \mathcal{S}} \sum_{j \in \mathcal{N}(i)} \texttt{cost}(c, \hat{c}) \cdot p_{\hat{c}}^j(t)$$

is the cost of using strategy c against profile p, γ is the learning rate set to $\gamma = 0.01$ and $\texttt{cost}(c, \hat{c})$ is equal to 1 if $c = \hat{c}$—i.e., if there is a match with the neighbour—and 0 otherwise. Strategies performing worse are penalised and their probability in the mixed strategy is reduced. The denominator normalises the distribution such that it sums to 1.

To reflect the experiment, each player's initial mixed strategy is of the form $(2/3 + \epsilon, 1/6 - \epsilon/2, 1/6 - \epsilon/2)$, where ϵ is a uniformly distributed noise in $[-1/500, 1/500]$ to break the symmetry between moves. Dynamics evolve such that at equilibrium, all colours that are a best response have positive weight. As seeds are sent, receivers update their strategy to (w_i, w_{-i}), where $w_i = 2/3 + \epsilon$ where i is the recommended colour and $w_{-i} = (1/6 - \epsilon/2, 1/6 - \epsilon/2)$.

Convergence of dynamics Best response dynamics are run until a Nash equilibrium is found, i.e., when no player can strictly decrease her number of matches by changing her current colour. MWU dynamics converge for each player to a mixed strategy that assigns positive probability to any colour that is a best response to the current environment. For instance, a player may have only one action weighted with probability 1 in the mixed strategy, or a probability p on Red and $1-p$ on Blue if both Red and Blue are colours that would yield the same number of matches under the current profile. The value of p depends on the initialisation of the dynamics, $p(0)$.

Behavioural model

Players in the game make decisions based on their local neighbourhood and information environment (e.g., with seed or without seed). A behavioural model of their decision-making is built from two pieces, whether to switch to a different colour or not, and which colour to choose in a switch. To obtain the data for fitting the model, we get for each second of play and each player in every round a vector containing local information at the start of the round (fraction of neighbours in each colour choice, presence of a seed, colour of seed, current colour choice) and events occurring during that second (whether the player switches or not, and to which colour).

Switching model A logistic regression returning a binary decision variable is constructed from the following factors: fraction of matches, i.e., number of matches divided by node degree; seconds elapsed since previous move; whether the player has received a seed or not; whether the seed is equal to the current colour choice of the player. The model is fit with all of the data. The output variable represents whether the player deviates from her current colour choice or not.

Colour choice model A logistic regression returning a decision variable with three levels is constructed from the following factors: fraction of neighbours in colour 1 (red); fraction of neighbours in colour 2 (yellow); fraction of neighbours in colour 3 (blue); colour of the seed (categorical variable with 4 levels: no seed, red, yellow, blue). The model is fit on vectors where the player makes a colour switch. The output variable represents the colour choice of a switching player.

Agent-based simulation The simulation strictly follows the rules of the experiment. All players are endowed with 100 points, the duration of the game is 120 seconds (called steps of the simulation hereafter) and seeds are given out after 20 steps. At each step, the switching model is queried to decide whether each player should switch. For those who do, the colour choice model is queried to choose their move. If the model

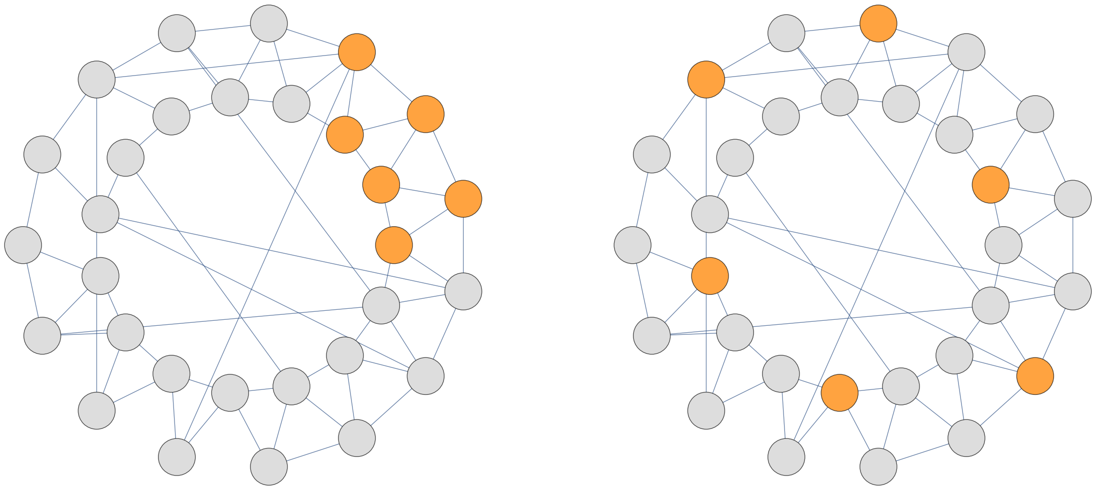

(A) (B)

FIGURE 3.5: Comparison of the group betweenness centrality of two sets of receivers with the sum of their individual betweenness centralities Σ_b. The GBC more accurately measures the centrality of the set of receivers. **a.** GBC = 63.84, $\Sigma_b = 300.57$. **b.** GBC = 240.43, $\Sigma_b = 286.8$.

returns the current colour of the player, it is ignored. Costs are computed using the same scoring rule, with the resulting history of moves in the simulation.

3.3.2 Influence of receiver position on costs

One distinguishing feature of networks where player costs have improved after the suggestion is that the seeds advise to play a best response for a larger fraction of time. However, there is not a significantly higher fraction of receivers for which the seed is their current colour choice in the two successful networks. But by the end of the round, in *ba* and *ws*, a significantly larger fraction of receivers never deviated from the suggestion than in *er* and *sb*. In other words, their environment maintained the seed as a best response.

The preceding observation indicates that network effects are at play, such as the topology of the graphs and the position of the receivers. The group betweenness centrality (Everett and Borgatti, 1999) of a set of nodes is employed here to explain the qualitatively different results obtained in the four networks. Let $\mathcal{C}$ be a set of nodes in the network. Their group betweenness centrality (GBC) is measured by the sum over all pairs of sources and destinations s, t (not included in $\mathcal{C}$) of the fraction of shortest paths connecting s to t running through at least one point in $\mathcal{C}$ among all shortest paths connecting s to t.

We compare in Figure 3.5 the GBC of a set of nodes with the sum of their individual betweenness centralities. The latter does not discriminate between a highly clustered set of receivers, for which the seeds may only have local effects, and a set that covers

TABLE 3.1: Spearman rank correlation coefficient between group betweenness centrality and average number of matches at final equilibrium, for each network and communication level. Significance levels are given next to the values: (*) $p < 0.05$, (**) $p < 0.01$, (***) $p < 0.001$. Best response dynamics are run by sampling one node at a time and choosing at random a best response, until all nodes are at equilibrium. As suggestions are sent, receivers follow the seed unconditionally. Multiplicative weight updates (MWU) keeps track of the players' mixed strategies. Finally, we test the correlation between the GBC of the set of receivers and the average player cost for simulations based on the behavioural model of players. Each dynamics is run for n times on m random receiver sets.

		Networks			
Dynamics	Receivers	ba	ws	er	sb
BR	10%	-0.355 (***)	-0.215 (**)	-0.144 (*)	-0.107
$(n = 200, m = 100)$	20%	-0.449 (***)	-0.137	-0.279 (***)	-0.391 (***)
MWU	10%	-0.220 (*)	-0.282 (**)	0.007	-0.166
$(n = 100, m = 100)$	20%	-0.119	-0.112	-0.099	0.113
Behavioural model	10%	-0.549 (***)	-0.176 (*)	-0.363 (***)	-0.419 (***)
$(n = 200, m = 100)$	20%	-0.592 (***)	-0.236 (***)	-0.352 (***)	-0.553 (***)

more appropriately the network. This is unlike the GBC which gives a starkly different measure for both sets, and is thus a relevant metric for the effect of receiver positions.

We test this hypothesis by sampling at random k receivers on the network. A dynamics is run until an equilibrium is reached, at which point a colouring is selected in the same manner as in the game. The k receivers immediately update to the seed, and the dynamics is run again until a new equilibrium S_f is reached. We repeat this procedure for m times and average the number of matches obtained in S_f. The set of receivers is sampled n times, for $k = \{3, 6\}$ and we compare the series of receivers' GBC with the corresponding average number of matches, using Spearman's rank correlation coefficient. The full table of results is presented in Table 3.1.

With best response (BR) dynamics, for almost all networks and all values of k, the coefficient is significantly smaller than 0, indicating that a higher GBC translates to a lower average number of matches. The dynamics is obtained by sampling uniformly at random one node at a time and allowing this node to play her current best response strategy, or one of them in case of ties.

The results are more contrasted with multiplicative weight update (MWU) dynamics (Littlestone and Warmuth, 1994), perhaps reflecting the noisier approach of the procedure. The dynamics is carried over mixed strategies of the agents, or probability distribution over their colour choices. At each step, every player's mixed strategy is updated by decreasing the probability of a colour that yields a higher cost against the expected choice of one's neighbours.

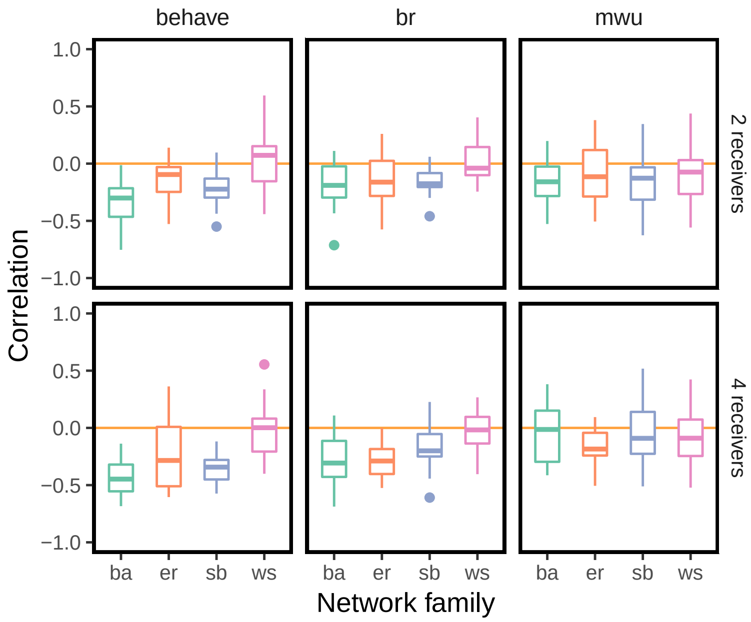

FIGURE 3.6: Spearman correlation coefficient for networks of 20 nodes and three dynamics (BR, MWU, behavioural model). For each network, 25 sets of 2 and 4 receivers are selected randomly. For each set of receivers, each dynamics is simulated 50 times and its resulting statistic is computed, the average number of matches for BR and MWU and the resulting average player cost for the behavioural model. The correlation is obtained for each network between the GBC of the receiver set and the resulting statistic. The boxplot represents the distribution of these coefficients for each network family and number of receivers. For most networks, the correlation coefficient is significantly below 0, indicating that more central receivers (as measured by the GBC) tend to improve the play.

Finally, a behavioural model obtained from the experimental data (detailed in 3.3.1) yields strong indication that the GBC of receivers is inversely correlated with the average of all player costs, computed with the same experimental scoring function. The model is obtained from two logistic regressions encoding respectively the choice to deviate from one's current action and the subsequent colour choice.

3.3.3 Influence of network topology on costs

The three dynamics are repeated on 80 networks, 20 per type, with 20 nodes each. The additional networks were generated from the same rules which yielded the four networks in the study. For each of these networks, 25 sets of receivers were sampled randomly for $k = \{2, 4\}$. For each set of receivers, the three dynamics were sampled from 50 times to obtain the average number of matches resulting from BR and MWU or the average player cost from the behavioural model. Once again, the correlation between these statistics and the centrality of the receivers were computed and shown to be significantly negative for most networks.

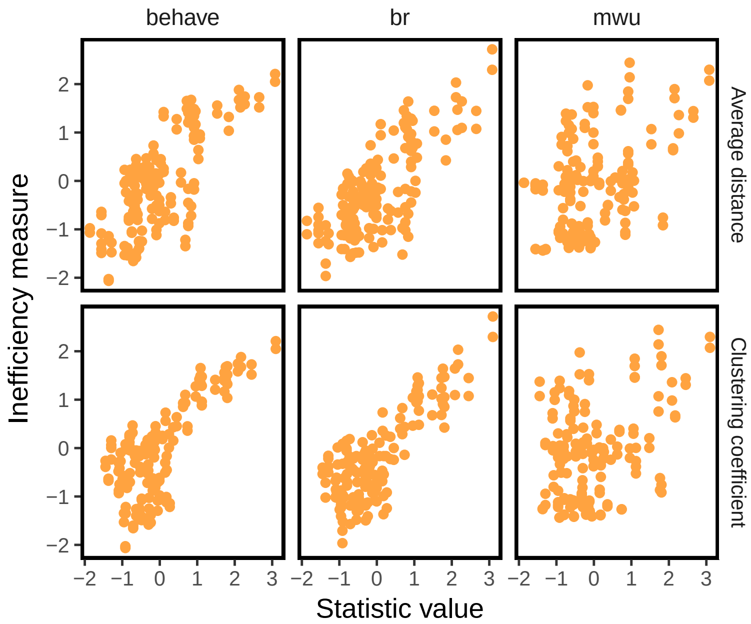

FIGURE 3.7: Scatterplot of the measure of cost for each dynamics (average resulting number of matches for BR and MWU, average player cost for the behavioural model). Simulations with 2 or 4 receivers appear on the same plot. Values of the network statistics and cost are rescaled for plotting (no incidence on the Spearman correlation coefficient). As either the average distance or the clustering coefficient of the network increases, the cost increases.

Figure 3.6 charts the distribution of these correlation coefficients for the two experimental conditions and each of the four types. Again, contrasts exist between the four different types. Preferential attachment networks such as *ba* have the strongest negative correlation between centrality and efficiency of the suggestion. On the other hand, over half of Watts-Strogatz networks have a positive correlation for the behavioural model simulations with 2 receivers and typically do not feature as strong negative correlations as the three other types in remaining treatments.

The differences are explained by the topology of the networks, as measured by network statistics such as the clustering coefficient or the average distance between two vertices. On the same set of 80 networks, a strong negative correlation is found between any of the two previous statistics and cost (Figure 3.7 and Table 3.2). Thus, a more clustered network leads to increased cost, as does a network with larger average distance. Watts-Strogatz networks, among the four considered types, are indeed more clustered and have higher average distance than any other type (Table 3.3).

TABLE 3.2: Spearman rank correlation coefficient between one network statistic (average distance or clustering coefficient) and a measure of system inefficiency (average number of matches for BR and MWU, average player cost for the behavioural model), for each simulated dynamics. Simulations were carried over 80 networks of 20 nodes each, 20 networks per type. 25 receiver sets were sampled for each network at sizes 2 and 4 (resp. 10% and 20% of nodes), with each dynamics run for 50 times on each set. Averages are obtained over all receiver sets. Significance levels are given next to the values: (*) $p < 0.05$, (**) $p < 0.01$, (***) $p < 0.001$.

Dynamics	Statistic	Correlation
BR	Average distance	0.647 (***)
	Clustering coefficient	0.628 (***)
MWU	Average distance	0.362 (***)
	Clustering coefficient	0.207 (**)
Behavioural model	Average distance	0.681 (***)
	Clustering coefficient	0.628 (***)

TABLE 3.3: Average network statistics (average distance and clustering coefficient) for 80 additional generated networks, 20 per types.

Network	Average distance	Clustering coefficient
ba	2.12	0.208
ws	2.42	0.342
er	2.20	0.194
sb	2.23	0.180

3.4 Discussion

The experiment investigates the question of coordination. Given a large landscape of equilibria and transient configurations, can we expect agents guided by their own interests to reach an efficient profile? The answer on that front is largely no, which prompts a natural second question: Can this coordination be helped?

The answers here are more positive, with some nuances. A centralised messenger with perfect information on the current state of the game, its efficient equilibria and able to intervene once is shown to significantly affect payoffs in some cases. The topology of the network matters, as does the position of the receivers, but the suggestion does not appear to hurt the play in any case.

This result is the first complication to the relationship between decentralisation and efficiency. A centralised, but local, suggestion can draw the play away from under-performing equilibria. Can the same be said in more general games? Perhaps not, as Figure 1.2 shows. Advertising a socially optimal configuration in the Pigou network might not lead to any stable profile, with agents deviating back to the variable cost link once they observe their (higher) cost on the constant latency link. Chapter 5 and 6 study respectively the efficiency loss to the selfishness of agents and a measure to reduce the loss via a toll mechanism, with its side-effects on inequality.

"You can plan a pretty picnic
But you can't predict the weather"

Ms. Jackson, André 3000

Chapter 4

How efficient are learning agents following no-regret algorithms?

The graph colouring game shows experimentally that faced with a hard coordination problem, a society of agents can reach more efficient states by making use of a weak centralised suggestion. The suggestion induced the selection of more efficient equilibria, as evidenced by the proportion of the play spent on a best response choice. However, the game has a particular structure, that of a potential game, where any best response sequence of moves leads to a NE, efficient or not. In general games, it is not the case that such a sequence converges to a NE. How can we then guarantee that the play will reach good states?

We focus our attention on no-regret learning dynamics. Multiplicative Weight Updates (MWU) (Young, 2004), introduced in Section 3.3.1 is a general algorithm that yields no-regret dynamics for agents implementing it. MWU is widely applied: The diversity of genetic properties and their evolution through time can be modelled by MWU, as can be the selection of a good model in machine learning. Agents play repeatedly, observe the costs they incur from the play and decrease the probability of selecting relatively high cost actions in the future. The algorithm has low complexity, since it does not keep track of the sequence of play but only the total cost for each strategy. In this sense, a society of agents implementing no-regret algorithms to govern their play presents a high degree of decentralisation.

Contents of this chapter

The convergence to the set of NE cannot be guaranteed in general games with no-regret agents, but the play does converge to a weaker set, that of coarse correlated equilibria (CCE). Precise definition of the set is given in Section 4.1, as well as related concepts. For now, an intuitive definition of a CCE is the following. The profile s is a CCE if a centralised messenger is able to communicate to each agent its strategy s_i under s, and agents must publicly commit to playing their strategy s_i before the play. In other words, if agent i can improve her payoff by publicly *not* committing to s_i, even when

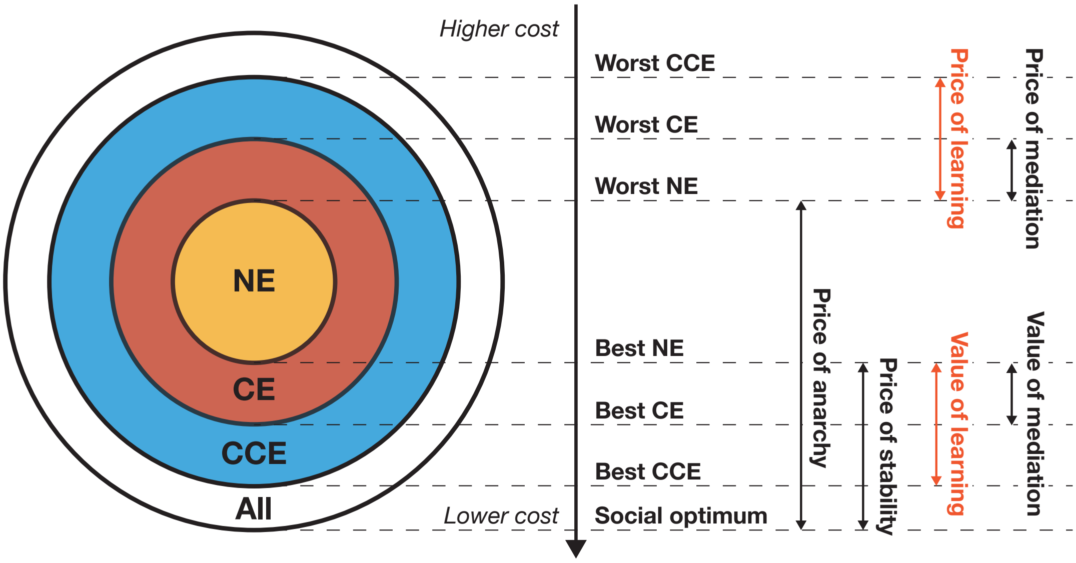

FIGURE 4.1: The hierarchy of equilibrium concepts, with efficiency gaps. We introduce the Price of Learning as well as the Value of Learning, to extend respectively the Price and the Value of Mediation.

the other players can implement a coordinated "punishment" (since they all know i is deviating), then s is not a CCE.

Convergence of the play to the set of CCE is guaranteed by agents following no-regret learning algorithms. At each step, agents play their current strategy s_i^t. They update their strategy following the reception of their costs, and repeat. For large enough t, the sequence $(s^t)_t$ approaches the set of CCE. In Section 4.2, we show a procedure that converges to a NE while maintaining the no-regret property. This convergence is unnatural and involves player elaborating tests to prove the honesty of their opponents.

To quantify the potential improvement of agents converging to general CCE instead of NE, we must find the gap between the best NE and the best CCE. We introduce to that effect in Section 4.3 the **Value of Learning**, or the gap from the best CCE to the best NE. Results from Roughgarden (2015) provide the converse bound from the worst of the CCE to the social optimum, via an extension of PoA-type bounds in (λ, μ)-smooth games. A diagram showing the relationship between efficiency gaps is given in Figure 4.1.

However, the bounds of Roughgarden (2015) are surprisingly lax for no-regret learning. By expanding the regret term in Section 4.4, we find an immense number of time steps are required for the bound to be effective in routing games. This theoretical result becomes more pertinent as we study the regret of players in a real routing game, in Chapter 5.

4.1 Definitions

We first give the definition of a correlated equilibrium (Aumann, 1974).

Definition 4.1.1. A *correlated equilibrium* (CE) is a distribution π over the set of action profiles $S = \prod_i S_i$ such that given any player i and pair of distinct strategies $s_i, s_i' \in S_i, s_i \neq s_i'$,

$$\sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i})\pi(s_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} c_i(s_i', s_{-i})\pi(s_i, s_{-i})$$

Given now is the definition of coarse correlated equilibrium (Young, 2004).

Definition 4.1.2. A *coarse correlated equilibrium* (CCE) is a distribution π over the set of action profiles $S = \prod_i S_i$ such that given any player i and any strategy $s_i \in S_i$,

$$\sum_{s \in S} c_i(s)\pi(s) \leq \sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i})\pi_i(s_{-i})$$

where $\pi_i(s_{-i}) = \sum_{s_i \in S_i} \pi(s_i, s_{-i})$ is the marginal distribution of π with respect to i.

The next definitions outline a framework to what "learning" means for agents involved in a game. They stem from the literature of online algorithms, where an agent with a restricted set of actions repeatedly makes choices that yield a certain payoff.

An online learning algorithm is an online algorithm for choosing a sequence of elements of some fixed set of actions, in response to an observed sequence of cost functions mapping actions to real numbers. The t-th action chosen by the algorithm may depend on the first $t - 1$ observations but not on any later observations; thus the algorithm must choose an action at time t without knowing the payoffs of any actions at that time. More formally,

Definition 4.1.3. An online sequential problem consists of a feasible set $F \in R^m$, and an infinite sequence of cost functions $\{c^1, c^2..., \}$, where $c^t : R^m \to R$.

At each time step t, an online algorithm selects a vector $x^t \in R^m$. After the vector is selected, the algorithm receives f^t, and collects a payoff of $f^t(x^t)$. All decisions must be made online, in the sense that an algorithm does not know f^t before selecting x^t, i.e., at each time t, a (possibly randomised) algorithm can be thought of as a mapping from a history of functions up to time t, $f^1, \ldots, f^{t-1}$, to the set F.

Given an algorithm A and an online sequential problem $(F, \{c^1, c^2, \ldots\})$, if $\{x^1, x^2, \ldots\}$ are the vectors selected by A, then the cost of A until time T is $\sum_{t=1}^{T} c^t(x^t)$. Regret compares the performance of an algorithm with the best static action in hindsight.

Definition 4.1.4. The regret of algorithm A at time T is defined as

$$R(T) = \sum_{t=1}^{T} c^t(x^t) - \min_{x \in F} \sum_{t=1}^{T} c^t(x)$$

An algorithm is said to have no-regret or that it is Hannan consistent (Young, 2004), if for every online sequential problem, its regret at time T is $o(T)$. For the context of game theory, which is our focus here, the following definition of no-regret learning dynamics suffices.

Definition 4.1.5. The regret of agent i at time T is defined as

$$R(T) = \sum_{t=1}^{T} c_i(s^t) - \min_{s'_i \in S_i} \sum_{t=1}^{T} c_i(s'_i, s^t_{-i})$$

We will also make use of Hoeffding's inequality (Hoeffding, 1963) and the Borel-Cantelli lemma (Émile Borel, 1909).

Theorem 4.1.1 (Hoeffding's inequality). *Suppose $(X_k)_{k=1}^n$ are independent random variables taking values in the interval $[0,1]$. Let Y denote the empirical mean $Y = \frac{1}{n}\sum_{k=1}^n X_k$. Then for $t > 0$*

$$\mathbb{P}(|Y - \mathbb{E}[Y]| \geq t) \leq 2\exp\big(-2nt^2\big)$$

Theorem 4.1.2 (Borel-Cantelli). *Let $(A_n)_{n=1}^\infty$ be a sequence of events in a probability space. Suppose $\sum_n \mathbb{P}(A_n) < \infty$, then $\mathbb{P}(\limsup_{n\leftarrow\infty} A_n) = 0$, or, equivalently, with probability 1, only a finite amount of events A_n will happen.*

4.2 Convergence to Nash by no-regret dynamics

A key question in the analysis of extremal behaviour of no-regret dynamics in general games is whether there exists a hidden implicit tension between achieving no-regret guarantees against malicious agent behaviour while at the same time converging to Nash equilibrium in self-play. The following theorem establishes that this is not the case.

Theorem 4.2.1. *In a finite game with N players, for any $\epsilon > 0$, there exist learning dynamics that satisfy simultaneously the following two properties:*

- *Against arbitrary opponents their average regret is at most ϵ,*
- *In self-play they converge pointwise to a ϵ-Nash equilibrium with probability 1.*

Proof. We divide the play in four stages. In the first stage, players explore their strategy space sequentially and learn the costs obtained from every action profile. In the second stage, they communicate their costs, using a procedure akin to cheap talk (Aumann and Hart, 2003). For example, they can use their actions as encoders for the payoffs previously revealed during the exploration stage, and transmit the knowledge they gained then to the other players. In the third stage, they compute the desired ϵ-Nash equilibrium that is to be reached, for $\epsilon > 0$. In the fourth stage, players are expected to use their equilibrium strategies and they monitor other players in case these deviate from equilibrium play. Below the proof, we give a pseudocode version of the algorithm implemented by the players, to summarise the four stages (in Algorithm 1)

The players are expected to follow a communication procedure and implement a no-regret strategy in the case of another player's deviation. Since the first three stages have finite length (though very long: exponential in the size of the cost matrix (Hart and Mansour, 2007)), the no-regret property follows. The restriction on convergence to an ϵ-NE, instead of a mixed NE (so $\epsilon = 0$) arises from the fact that even games with rational costs can possess equilibria that are irrational (Nash, 1951).

Settlement on a particular NE can be decided by a fixed rule before play, such as lexicographically in the players' actions or the NE that has the lowest social cost.

In the fourth stage, players have settled on an equilibrium and will implement it. To fulfil the requirement of pointwise convergence, it is not enough for the players to stick to a deterministic sequence of plays. We want them to pick randomly a move from their equilibrium distribution of actions. During this process, the generated sequence of play of an opponent may not match his equilibrium distribution. In that case, the players need to decide whether the opponent has been truthful but "unlucky" or deliberately malicious.

We achieve this by dividing the fourth stage in blocks of increasing length. Let $n \in \mathbb{N}$ denote the block number, we set block n to have a length of $l(n) = n^2$ turns. On these blocks, the players will make use of tests to verify that all other opponents are truthful, in the sense that they follow the prescribed mixed NE. We want to find a test such that a truthful but possibly unlucky player will fail almost surely a finite number of these tests, while a malicious player will almost surely fail an infinite number of these.

We first look at the case where we have N players with only two strategies, 0 and 1. We can then identify the equilibrium distribution of a player i, to the probability p_i^* that he chooses action 1.

Suppose the play is at the n-th block and player i chooses to implement the mixed strategy p_i. Let $(X_k^i)_{k=1,\dots,l(n)}$ denote the sequence of strategies chosen by player i, such that $X_k^i \sim \mathcal{B}(p_i)$ and all are independent. Let Y_n^i be the empirical frequency of strategy 1 during block n.

$$Y_n^i = \frac{1}{l(n)} \sum_{j=1}^{l(n)} X_j^i$$

If the player is truthful and implements the prescribed NE, then we have $p_i = p_i^*$ and we expect the empirical frequency of strategy 1 Y_n^i to be close to p_i^*. Otherwise, a malicious player will choose $p_i \neq p_i^*$.

Let A_n^i denote the event $A_n^i = \{|Y_n^i - p_i^*| \geq t_n\}$. In other words, we are trying to determine how far the empirical frequency of strategy 1 is from the expected equilibrium distribution. If the event A_n^i is realised, then the test is failed: the empirical distribution of play is too far from the expected NE distribution. The idea is to make block after block the statistical test more discriminating, i.e. get a decreasing sequence $(t_n)_n$ such that a truthful player will only see a finite number of events A_n^i happen, while a malicious one will face an infinite number of failures.

We claim that picking $t_n = n^{-\alpha}$ with $0 < \alpha < 1$ is enough. Indeed by Hoeffding's inequality we have that

$$\mathbb{P}(A_n^i) \leq 2\exp\big(-2n^2t_n^2\big)$$

if the player is truthful (remember that block n has length $l(n) = n^2$).

Extending the proof to the case where a player i has finite strategy set S_i is not hard. Let $(p_s^i)_{s\in S}$ be the distribution that the i-th player decides to implement, while $(p_s^{i,*})_{s\in S}$ is the NE distribution for player i. Let $X_k^{i,s}$ follow a multinomial distribution of parameters $(p_s^i)_{s\in S}$. Then $Y_n^{i,s}$ is the empirical frequency of strategy s during block n for player i. We define events

$$A_n^{i,s} = \{|Y_n^{i,s} - p_s^{i,*}| \geq t_n\}.$$

Then we define our test A_n^i to be $\cup_{s\in S_i} A_n^{i,s}$. Using Hoeffding's inequality again we obtain:

$$\begin{aligned}\mathbb{P}(A_n^i) &= \mathbb{P}(\cup_{s\in S_i} A_n^{i,s})\\ &\leq \sum_{s\in S_i} \mathbb{P}(A_n^{i,s}) \qquad \leq |S_i| \times 2\exp(-2n^2t_n^2)\end{aligned}$$

Thus $\sum \mathbb{P}(A_n^i) < +\infty$ for $0 < \alpha < 1$, so by Borel-Cantelli we know that the A_n^i will only ever happen a finite number of times if the player is truthful, i.e. if $\mathbb{E}[Y_n^{i,s}] = p_s^{i,*}$.

To satisfy the no-regret property, we do the following: if one of the opponents failed the statistical test described earlier, then all players will implement a no-regret strategy for a time $n^{2+\delta}$ to compensate for that. We call this block of size $n^{2+\delta}$ a *compensating block*.

If a finite number of tests fails, then the whole sequence satisfies the ϵ-regret property, since players are arbitrarily close to the ϵ-Nash equilibrium. When one of the tests fails, say, at block n, the maximum regret accumulated is of size n^2. The following compensating block guarantees that overall regret has grown by a value bounded by $n^{1-\delta}$, i.e., sublinearly.

We also guarantee that the expected turn number that ends the last of the truthful player's potential failed block is not infinity. Indeed let B_n be the event that the last failed block is the n-th one. Then, if $A_n = \cup_{i=1}^{N} A_n^i$,

$$\begin{aligned}\mathbb{P}(B_n) &= \mathbb{P}(A_n) \times \mathbb{P}(A_{n+1}^c) \dots \\ &\leq 2\exp(-2n^2t^2) \times 1 \dots \\ &\leq 2\exp(-2n^2t^2)\end{aligned}$$

We use A^c to denote the complement of event A. The first equality holds by independence of the blocks, the second inequality is true from Hoeffding's and the fact that a probability is less or equal to 1. We then define L to be the index of the turn that ends the last compensating block of a truthful player. L is a random variable on the integers. We have

$$\mathbb{E}[L] \leq \sum_n \Big(\sum_{k=1}^{n} (k^2 + k^{2+\delta}) \Big) \times 2\exp(-2n^2t^2) < +\infty$$

We bound $\mathbb{E}[L]$ by assuming a truthful player got every test wrong up to the latest failed one. Then the last turn L occurs at index $\sum_n (n^2 + n^{2+\delta})$. We multiply this by the bound on $\mathbb{P}(B_n)$ and use the property of the exponential to conclude that $\mathbb{E}[L]$ is bounded. $\square$

4.3 The Value of Learning

4.3.1 Social welfare gaps for different equilibrium concepts

We define a measure to compare equilibria obtained under no-regret algorithms to Nash equilibria: *the value of learning*. This measure quantifies by how much the players are able to decrease their costs when relaxing the equilibrium requirements from Nash to CCE.

Definition 4.3.1. Define the value of learning in cost games *VoL* as the ratio of the social cost of the best Nash equilibrium to that of the best coarse correlated equilibrium.

$$\text{VoL}(\Gamma) = \frac{\text{best NE}}{\text{best CCE}}$$

Since the set of NE is included in the set of CCE, then the best NE in terms of social cost will always be greater than the best CCE. Thus we take the ratio so that the value of learning is always greater than or equal to 1, a convention also found in other papers

Data: N players
Data: $\epsilon > 0$
Data: $0 < \alpha < 1$
Data: $\delta > 0$
Step 1: Exploration
begin
    while *One profile has not been played* **do**
        Play new profile
    end
end
Step 2: Communication
begin
    Players communicate their costs by encoding them using their actions
end
Step 3: Computation
begin
    The ϵ-Nash Equilibrium to be played is computed from the costs
end
Step 4: Implementation
Data: ϵ-NE p^*
begin
    $n \leftarrow 1$ `// n is the block number`
    while $n > 0$ **do**
        $l \leftarrow n^2$
        $t \leftarrow n^{-\delta}$
        $C_{i,s} \leftarrow 0, \forall i, s$ `// C counts use of strategy s by player i`
        for $j \leftarrow 1$ **to** l **do**
            Players move according to $S = (s_1, \ldots, s_N)$
            for $i \leftarrow 1$ **to** N **do**
                $C_{i,S(i)} \leftarrow C_{i,S(i)} + 1$
            end
        end
        if $\exists i, s$ *such that* $|\frac{C_{i,s}}{l} - p_s^{i,*}| \geq t$ **then**
            All players play a no-regret procedure for $n^{2+\delta}$ rounds
        end
        $n \leftarrow n + 1$
    end
end

Algorithm 1: Proof of Theorem 4.2.1. Players converge to an ϵ-NE in self-play while maintaining no-regret.

related to the price of anarchy (Ashlagi, Monderer, and Tennenholtz, 2008; Bradonjic et al., 2009).

Conversely, we define *the price of learning* as the ratio of the worst CCE to the worst NE.

Definition 4.3.2. Define the price of learning *PoL* in a cost game Γ as the ratio of the social cost of the worst coarse correlated equilibrium to that of the worst Nash equilibrium.

$$\mathrm{PoL}(\Gamma) = \frac{\text{worst CCE}}{\text{worst NE}}$$

The ratio of the worst CE to the worst NE was previously defined as the price of mediation (PoM) (Bradonjic et al., 2009). With the help of Proposition 4.3.1, we can extend this result to learning algorithms that possess the no-regret property.

4.3.2 CE = CCE for N agents 2 strategy games

We present a result allowing us to collapse two equilibrium classes in a specific case: any number N of players having 2 strategies each.

Proposition 4.3.1. *For games where all players have only two strategies, the set of coarse correlated equilibria is the same as the set of correlated equilibria.*

Proof. Let i be one of the players, suppose his two strategies are A and D, where we pick D to be the deviating one. Then the requirement for correlated equilibrium states that

$$\sum_{s_{-i}\in S_{-i}} u_i(s_{-i}, D)\pi(s_{-i}, A) \geq \sum_{s_{-i}\in S_{-i}} u_i(s_{-i}, A)\pi(s_{-i}, A)$$

while the corresponding one for coarse correlated equilibrium is

$$\sum_{s_{-i}\in S_{-i}} u_i(s_{-i}, D)(\pi(s_{-i}, A) + \pi(s_{-i}, D)) \geq$$
$$\sum_{s_{-i}\in S_{-i}} (u_i(s_{-i}, D)\pi(s_{-i}, D) + u_i(s_{-i}, A)\pi(s_{-i}, A))$$

which is equivalent after removing the $\sum_{s_{-i}\in S_{-i}} u_i(s_{-i}, D)\pi(s_{-i}, D)$ term on both sides. □

4.3.3 The VoL in 2x2 games

Denote by $\Gamma_{2\times2}$ the class of 2×2 games. We are interested in the best-case scenario: how high the ratio of the value of learning can get for all 2×2 games.

Definition 4.3.3. Denote by $\mathrm{VoL}(\Gamma_{2\times2}) = \sup_{\Gamma\in\Gamma_{2\times2}} \mathrm{VoL}(\Gamma)$ the value of learning for the class of 2×2 games.

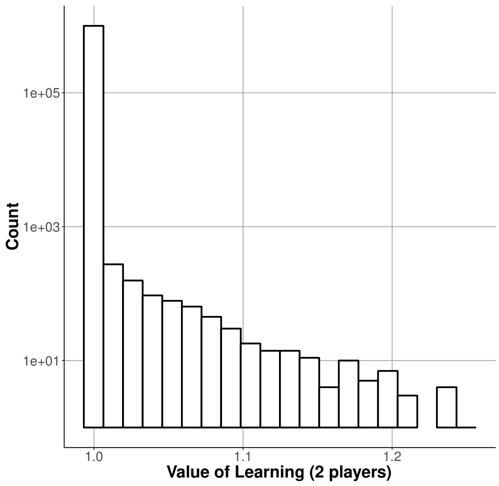

FIGURE 4.2: Histogram of values of learning obtained over 10^7 simulations for 2×2 games. A $\log_{10}$ scale is used for the y-axis.

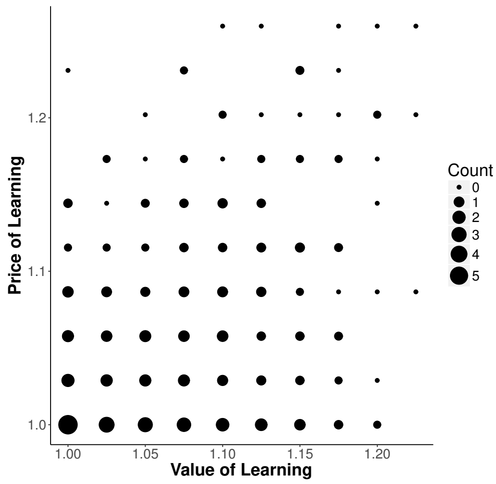

FIGURE 4.3: 2D histogram of VoL and PoL over 10^6 simulations for 2×2 games. The count legend is to be interpreted as a power of ten (where a count of 5 is 10^5 observations)

Proposition 4.3.2. *VoL*$(\Gamma_{2\times 2}) \geq \frac{3}{2}$

Proof. Consider the following cost game for $x > 1$

$$\begin{array}{ccc} & L & R \\ \begin{array}{c} T \\ B \end{array} & \multicolumn{2}{c}{\begin{pmatrix} 0, x-1 & x, x \\ 1, 1 & x-1, 0 \end{pmatrix}} \end{array}$$

The game admits three NE: (T, L), (B, R) and $((0.5, 0.5), (0.5, 0.5))$. The first two have social cost equal to $x - 1$ while the mixed equilibrium's is x. The minimum social cost is thus obtained for the pure equilibria, at $x - 1$.

The correlated equilibrium that minimises social cost assigns probability $1/3$ to every action profile except for (T, R). Its social cost is $2x/3$. Hence, in this game, VoL $= \frac{3(x-1)}{2x}$. Taking $x \longrightarrow +\infty$, we derive VoL$(\Gamma_{2\times 2}) \geq \frac{3}{2}$. □

We conjecture that this $\frac{3}{2}$ bound is tight, i.e., there is no 2×2 game Γ such that VoL$(\Gamma) > 3/2$. To support this claim, we run numerical simulations on games generated from a random uniform distribution. A notable result is the predominance of games for which the ratios are 1, i.e. mediation does not better the social welfare/cost. We then observe higher ratios at a lower rate, hence our histograms look like those of a power law (Figure 4.2). The obtained ratios come close to the 3/2 threshold, without going further (only a few ratios approaching 1.4 were observed over 10^7 simulations).

Proposition 4.3.3. *PoL*$(\Gamma_{2\times 2}) = 2$

Proof. By Proposition 4.3.1, the social cost of the worst CE is equal to the social cost of the worst CCE, since the set of CE is the same as the set of CCE. Then by (Bradonjic et al., 2009), we have that $PoL(\Gamma_{2\times 2}) = 2$. $\square$

In Figure 4.3 we present a 2D histogram of the joint distribution of the VoL and PoL. 10^6 games were generated and for each we compute both values. The size of the dot is representative of how many games possess particular values for the VoL and the PoL.

4.3.4 The VoL in larger games

Next, we examine larger games, i.e., games with more than 2 players and/or more than 2 strategies per player. Let Γ_{m_1,m_2} denote a 2 player game with respectively m_1 and m_2 strategies for each player.

Proposition 4.3.4. *For sets of games* Γ_{m_1,m_2}*,* $\max(m_1, m_2) > 2$*, we have VoL*$(\Gamma_{m_1,m_2}) = +\infty$.

Proof. Consider for $\epsilon < \frac{1}{2}$ the game

$$\begin{array}{cc} & \begin{array}{ccc} L & C & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} 1-\epsilon, 1-\epsilon & 2\epsilon, \frac{3\epsilon}{2} & 2\epsilon, \frac{1}{2} \\ \frac{1}{2}, 2\epsilon & \epsilon, 1-\epsilon & 1, 2\epsilon \end{pmatrix} \end{array}$$

The game admits three NE: (L, B), $((0, 1), (2/3, 0, 1/3))$ and $(2/3, 1/3), (0, 1-\epsilon, \epsilon)$. Of the three, the latter has the lowest social cost, equal to $1/3 + o(\epsilon)$, where $o(\epsilon) \longrightarrow_{\epsilon\to 0} 0$.

We can define the following correlated equilibrium π:

$$\begin{array}{cc} & \begin{array}{ccc} L & C & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} 0 & 1-\frac{5\epsilon}{2} & \epsilon \\ \epsilon & 0 & \epsilon/2 \end{pmatrix} \end{array}$$

The best social cost in a correlated equilibrium will be lower than that of π, which is $o(\epsilon)$. We also have that the best social cost in a CCE will be lower than that of a CE. Thus taking $\epsilon \to 0$, we obtain an unbounded VoL. $\square$

The set of CE being included in the set of CCE, we can again extend some results from previous papers to coarse correlated equilibria.

Proposition 4.3.5. *For games* Γ_{m_1,m_2}*,* $\max(m_1, m_2) > 2$*, we have PoL*$(\Gamma_{m_1,m_2}) = +\infty$.

Proof. Since CE $\subseteq$ CCE, the social cost of the worst CCE is higher than that of the worst CE. By (Bradonjic et al., 2009) we have that PoM $= +\infty$, hence PoL $= +\infty$. $\square$

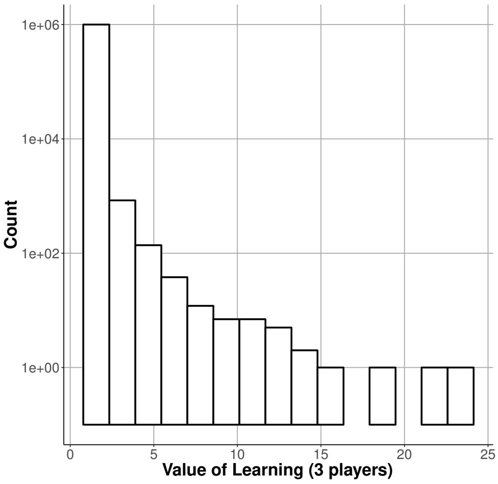

FIGURE 4.4: Histogram of ratios best NE/best CCE (VoL) obtained over 10^6 simulations for 3×3 games.

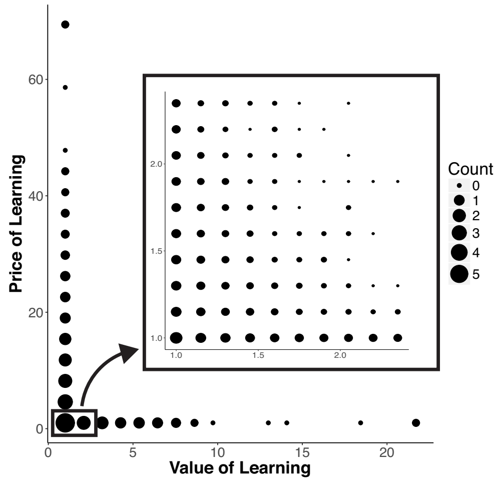

FIGURE 4.5: 2D histogram of VoL and PoL over 10^6 simulations for 3×3 games. The count legend is to be interpreted as a power of ten (where a count of 5 is 10^5 observations). We zoomed in the portion $[1, 2.5]^2$ to show finer results.

We run a number of simulations to see how VoL is distributed for random games (Figure 4.4). We have also included a 2D histogram (Figure 4.5) showing (VoL, PoL) for a number of generated games. Some sampled games have high VoL and some high PoL but not both, indicating a competitive relationship between the two quantities.

4.4 Limits of smoothness bounds for congestion games

We look at the (λ, μ)-bound in Roughgarden (2015) where agents implement no-regret learning procedures. In this case, it is shown that

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[C(s^t)] \leq \frac{\lambda}{1-\mu}\mathbb{E}[C(s^*)] + R(T)$$

with $R(T) \rightarrow_{T\rightarrow\infty} 0$. In other words, after a large enough T time steps, the bound for the cost under no-regret learning naturally converges to the same bound for Nash equilibria. This implies that the PoA over the set of coarse correlated equilibria is the same as the PoA over the set of NE.

But what is in this bound? When agents implement a no-regret strategy and the cost in each time step is $c_t : S \rightarrow \mathbb{R}$, with $c_t(s) \in [0, M]$, we can show

$$\mathbb{E}[\sum_{t=1}^{T} c_t(s_t)] < \mathbb{E}[\min_{s\in S}\sum_{t=1}^{T} c_t(s)] + (M+1)\sqrt{T\log(n)}$$

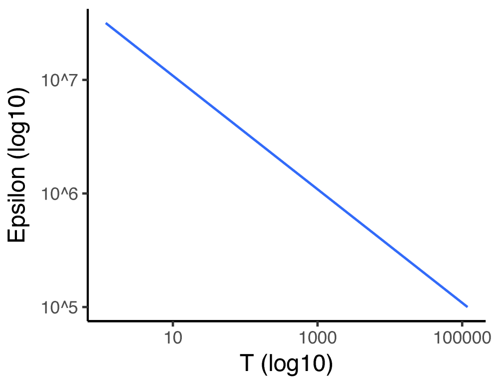

FIGURE 4.6: The tighter the desired bound, the larger T time steps must take place to achieve it.

where n is the number of strategies available to the agents (see also (Cesa-Bianchi and Lugosi, 2006, Sections 2.6, 2.8, Remark 2.2) for a discussion on why the term $O(M\sqrt{T\log(n)})$ cannot be improved upon when no preliminary information is available on the sequence of observed payoffs).

So

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[C(s^t)] \leq \frac{\lambda}{1-\mu}\mathbb{E}[C(s^*)] + \frac{1}{T}\sum_{t=1}^{T}\frac{\Delta(s^t)}{1-\mu}$$
$$\leq \frac{\lambda}{1-\mu}\mathbb{E}[C(s^*)] + \frac{N(M+1)\sqrt{\log(n)}}{(1-\mu)\sqrt{T}}$$

where

$$\frac{1}{T}\sum_{t=1}^{T}\frac{\Delta(s^t)}{1-\mu} = \frac{1}{T(1-\mu)}\sum_{t=1}^{T}\sum_{i=1}^{N}\delta_i(s^t)$$
$$= \frac{1}{T(1-\mu)}\sum_{t=1}^{T}\sum_{i=1}^{N}(C_i(s^t) - C_i(s_i^*, s_{-i}^t))$$
$$\leq \frac{N(M+1)\sqrt{\log(n)}}{(1-\mu)\sqrt{T}}$$

and N is the number of agents in the system.

Translated to the scale of Singapore, where approximately $N = 2$ million individuals commute each day, the $N\sqrt{\log(n)}$ term is already on the order of 10^6. This is additionally scaled by the upper bound of the cost function, M. If we desire $R(T) \leq \epsilon$, we must have

$$\frac{N(M+1)\sqrt{\log(n)}}{(1-\mu)\sqrt{T}} \leq \epsilon$$

$$\Leftrightarrow T \geq \left(\frac{N(M+1)\sqrt{\log(n)}}{(1-\mu)\epsilon}\right)^2$$

which for ϵ relatively small, of the order of 10^{-3} say, implies a number of time steps greater than 10^{18}! For comparison, the age of the Universe in days is of the order of 10^{12}. This appears unreasonable. We represent in Figure 4.6 the number of time steps T required to achieve a certain approximation ϵ.

4.5 Discussion

The property of no-regret is shared by many natural learning procedures implemented in multi-agent settings. Due to their convergence (to the set of coarse correlated equilibria) they are useful in practice. But if we look closer, it is not clear where this convergence leads the play. We have first shown that we can steer it using a somewhat unnatural algorithm to any NE of the one-shot game, while maintaining the no-regret property. In the next sections, we have understood better how the class of CCE relates to no-regret dynamics, and to the smaller class of CE.

This led us to define more general measures of the price of anarchy: if it is hard to predict where the play following no-regret dynamics will go, we are at least able to give some price of anarchy-type bounds on the resulting payoffs. We have concluded with experimental results that show a concentration of small ratios, indicating a closeness to NE payoffs. Proving our conjecture about the Value of Learning for 2×2 games remains an open question, with the lower bound of $3/2$ derived in 4.3.3, which we believe to be tight.

"Meanwhile along Orchard Road, bit of a traffic snarl building up with the evening rush hour... Otherwise it's been yet another ***beautiful day*** *here on our little island paradise!"*

The Art of Charlie Chan Hock Chye, Sonny Liew

Chapter 5

What does a large scale experiment on Singapore's routing network tell us about the efficiency of the system and the regret of its users?

Chapter 3 presented a controlled experiment where a game with large price of anarchy does not settle in the most inefficient states *and* can be nudged towards even better performing equilibria. On the other hand, Section 4.4 showed that no-regret bounds are very weak for large-scale congestion games. This chapter continues our exploration with an analysis of a "snapshot" of a large, real system, supported by a unique dataset on routing in Singapore, to answer what efficiency is in practice.

Let us first frame the discussion in this chapter. Since its inception, the literature following price of anarchy sought to establish theoretical performance guarantees on wider and wider classes of games. It obtained critical results showing that for congestion games (Rosenthal, 1973), which represent for instance the game played on a network by commuters—formally defined in Section 2.6—, the cost of an equilibrium is bounded away from optimal by a constant (Koutsoupias and Papadimitriou, 1999; Roughgarden and Tardos, 2002). This remains true even when agents follow no-regret learning strategies, as presented in Chapter 4 (Roughgarden, 2015).

These findings have strong implications for the design of routing networks, limiting the improvements that a central planner can achieve. But as they stand, they are incomplete as long as the predictive power of PoA has not been established in real networks these models represent. What does 2.151, the PoA upper bound for routing games with cost functions classically employed to represent congestion on real routes, mean in practice?

This work more generally leverages a granular data source, the National Science Experiment, to investigate questions related to efficiency in routing games. The data informs new metrics derived from concepts familiar in routing games, such as regret.

As such, this work is inscribed in the current effort to perform econometric measurements and experiments related to the larger algorithmic game theory literature (Bajari, Hong, and Nekipelov, 2013; Syrgkanis, 2015; Nekipelov, Syrgkanis, and Tardos, 2015).

Contents of this chapter

Our goal is to perform the first large scale, multimodal and granular experiment on routing in a city. The data is collected by distributing sensors to students throughout Singapore and incentivising correct use. Noisy by nature, the data is cleaned thoroughly with a set of validated algorithms until a collection of over 34,000 morning trips are obtained. References on Singapore's transportation landscape, the experiment and data cleaning processes are presented in Section 5.1.

Three sections are devoted to building increasingly sophisticated measures of the network to check for its equilibrium properties (Section 5.2), the individual optimality of commuters (Section 5.3) and finally the optimality of the system itself (Section 5.4). Before proceeding, we provide an overview of each section.

Equilibration of the system The price of anarchy is concerned with the cost of a game at a Nash equilibrium.[1] However, we first ask for a weaker definition of equilibrium, understood as stasis. Are agents currently updating their routing decisions? If not, this behaviour is consistent with a stationary point reached, for instance, by best response strategies.

Individual optimality Equilibration is the first step towards framing the data in game-theoretic terms. We now ask if the system is such that agents are close to optimality in the following sense: if two agents leave from the same origin around the same time, with the same mode of transportation, going to the same destination, are their travel times comparable? We must first observe that this optimality criterion is weaker than one which would ask of agents to choose the *best* route possible, given the congestion on all the links. The data however may not provide us with such information if we sample a limited set of agents on the network. But as sample sizes grow and groups of comparable subjects get larger, we must be more and more confident that the fastest student approaches strong optimality.

It turns out that the weak optimality criterion has close ties with regret, interpreted here as hindsight with respect to the best action of another comparable subject instead of any best route. If being at a Nash equilibrium implies that all agents have zero regret, the converse does not hold in general. Once again, we do not show agents are at the game-theoretic Nash equilibrium, but a low regret, given the extensions of PoA in such

[1] For a routing game Γ, all equilibria have equal cost (Section 2.6)

situations (Roughgarden, 2015), guarantees that our estimate of the numerator of PoA, concerned with the cost at equilibrium, cannot stray too far from its real value.

System optimality And yet, we shall not directly attempt to estimate the PoA. Once again, the denominator, which is the cost at optimum, asks more from our data than is available. We must be confident about the demands of all source-destination pairs, not only the ones in our dataset, to compute an appropriate optimal assignment for the subjects. Once obtained, the estimate of PoA yields the inefficiency loss due to selfish routing.

For a real network however, we argue that PoA does not adequately measure the loss of efficiency due to tragedy of the commons effects. The demands on the road resources are such that if it is individually efficient to join the network, it may not be collectively optimal to do so. As more agents join, the resources become stressed while latencies increase. In contrast, PoA is optimised in cases when edges are either empty or completely saturated (Colini-Baldeschi et al., 2017), the latter being hardly what one would describe an efficient network. For this reason, we introduce the Stress of Catastrophe (SoC), an upper bound to the PoA which measures system stress. We estimate low SoC, lower indeed than the theoretical bounds on PoA, casting further doubt on its applicability for real networks.

Result Snippets

- We show that most subjects use the same means of transportation across trips and that a large number of them consistently selects the same route. For example, when controlling for those who use consistently the same means of transportation across different days, the percentage of subjects selecting the same route is very high, in the order of 94%. (see Section 5.2).

- The empirical regret distribution has a median value of 5 minutes 15 seconds and mean approaching 7 minutes for an average travel time of around 27 minutes (see Section 5.3).

- We define and estimate the Stress of Catastrophe for subjects in private transportation across various scenarios. Even at its most pessimistic, the SoC is much lower than the corresponding PoA bound on real road networks. We find a marked contrast when discriminating by mode of transportation (see Section 5.4).

5.1 The Singapore National Science Experiment

5.1.1 A short guide to transportation in Singapore

We briefly discuss here some salient points on the topography of Singapore and its transportation network, to provide some context for the coming empirical results. All data is collected from data.gov.sg, for which we give the latest available measurement.

Singapore is a city state home to 5.6 million inhabitants (2017), of which 4 million are part of the resident population. The Land Transportation Authority (or LTA) oversees the questions related to transportation in the city. According to its classification, rail, bus and taxi modes are all considered public transportation, however, in the following, taxis will be understood as being private vehicles. Five mass rapid transit (MRT) lines and three light rail transit (LRT) lines operate in Singapore, for a total of 230 kilometres of rail (2017). Bus services are present throughout Singapore with a wide coverage (260 lines) and high frequency (under 15 minutes at peak hours, with half arriving under 10 minutes).

Private transportation in Singapore is controlled by two major mechanisms. The Certificate of Entitlement (COE) is a 10-year license auctioned at regular intervals to compensate exactly for the de-registration of motor vehicles (motorbikes, individual cars and trucks). The Electronic Road Pricing (ERP) is a dynamic tolling mechanism active during peak hours, with gantries located around the city centre and major expressways. Roads cover approximately 12% of land, with 550,000 private cars, 50,000 rental or for-hire cars and 25,000 taxis registered in 2016. Related to the number of households in Singapore, this entails that there is on average one car in 45% of households, circulating for 17,500 kilometres per year (2017). Public transportation is widely employed, with over 60% of trips completed with public transportation.

5.1.2 Data collection

The National Science Experiment (NSE) is a large-scale experiment realised in SUTD in collaboration with the National Research Foundation of Singapore and several industry experts. A custom-built sensor, SENSg, was designed to be carried throughout the day by students from primary, secondary and junior colleges, for one week each. Over 90,000 students took part in the experiment.

The sensor records at high frequency (up to every 13 seconds) a measurement made up of its location (determined by scanning surrounding MAC addresses) as well as several environmental factors, such as relative temperature and humidity, noise levels and illumination. Although this study focuses primarily on the geographical data, the additional factors are notably used in the machine learning-assisted mode identification algorithm described below in Section 5.1.5. Overall, over 130 million measurements were recorded.

5.1.3 Limitations of active methods of collection

The goal of the experiment is to overcome previous limitations of self-reported (or active) transportation data (e.g., via Household Interview Transportation Surveys). Active reporting tasks subjects with providing the researchers a detailed account of their trips, including the nature of the trip. This provides verifiable information with the limitation that subjects are likely to under-report their trips, due to the difficulty of logging their activities consistently (Du and Aultman-Hall, 2007).

Passive trip reporting offers a promising alternative that also allows for scale. Raw data only is collected periodically, such as the agent's latitude and longitude. It is then the task of the researchers to do additional transformations of the data to obtain meaningful statistics. The burden of verifying subject logs to constitute a dataset is avoided, which makes the procedure more suitable for large data sets.

Previous studies employed passive trip reporting. For instance, Axhausen et al. (2003) track private vehicles of their subjects. A trip is defined as a sequence starting from the powering on of the car and ending when the contact is off—with additional criteria to filter out smaller trips or include stops during which the car is left on. Schüssler and Axhausen (2008) track the subject's location instead of their vehicle, which compounds processing difficulties due to the continuous nature of human mobility. The authors apply a smoothing method on the geographical data points to overcome the deficiencies of GPS-based data collection, also noted by Jun, Guensler, and Ogle (2006).

One of the most vexing problems facing researchers studying large-scale mobility patterns is the unwillingness of participants to spend their smart-phone battery energy for the collection of location data from power-hungry GPS services. Several research groups (including this one) have proposed down-sampling location estimates as a solution (Jariyasunant, Sengupta, and Walker, 2012; Kumar et al., 2013). An alternative approach was taken in the design of the NSE sensor, described in detail in Section 5.1.2. MAC addresses of surrounding Wi-Fi hotspots are scanned and recorded to locate the sensor's position. This method is shown to be accurate to within 20 to 30 metres and tends to be more effective at lower speeds (Tsui et al., 2010).

Passive methods collect raw data from the subjects, in the background. Given the size of the dataset and the often noisy nature of its collection, the data must then be interpreted by the analysts using algorithms. We describe in the next two subsections two algorithms yielding semantic data out of the collected raw data. Finally, we address concerns about the representativeness of the sample (students in Singapore) and limitations of the dataset for our research questions in Section 5.1.6.

5.1.4 Trip identification

The collected data is a list of raw geographical positions (latitude and longitude). First, smoothing algorithms eliminate the noise from the stream of locations. Second, a dwell time-based algorithm decides on a list of Point of Interests (POIs) (latitude and longitude). These POIs are locations where the subject spent a significant amount of time at low velocity, i.e., a stationary activity in a smaller area.

To identify home and school locations in the list of POIs, additional information such as the timestamp is checked. Late night measurements typically hint at the home location, while the sensor spends the clearest part of daytime hours near or at school.

5.1.5 Mode identification

The mode identification algorithm uses a mix of the measurements operated by the sensor and geographical information, such as the position of train lines and roads or bus and train stations. The information is taken as input by a machine learning procedure based on random forests, able to discriminate for our purposes between walking, travelling by train, by bus or by car. Accuracy reaches 85% on a validation set, indeed superior to similar recent studies (Sankaran et al., 2014; Shin et al., 2015; Zhu et al., 2016). Figure 5.1 presents a trip as a collection of segments.

5.1.6 The clean trip dataset: Opportunities and limitations

The NSE 2016 dataset contains measurements from 49,526 students who participated in the experiment and wore the sensor. The experiment was designed to analyse homogeneous users, i.e., primary, secondary or junior college students, reducing the complexity of understanding mobility patterns. This work focuses on morning travels of students who get to their schools from their homes. Two main reasons were considered for this choice.

First, in the following analyses, the latency, or duration of the trip, is considered as the primary "cost" of the subjects, discounting any other monetary cost. Morning trips typically feature subjects optimising to minimise their latency. Evening trips are more sparse since the battery of sensor is expected to be charged at night while the subject is home. By the end of the day, if the battery has run out, the evening trip is not recorded. We have however in the dataset 21,065 samples for which both morning and evening trips are recorded. For these pairs, the average duration of the morning trip is

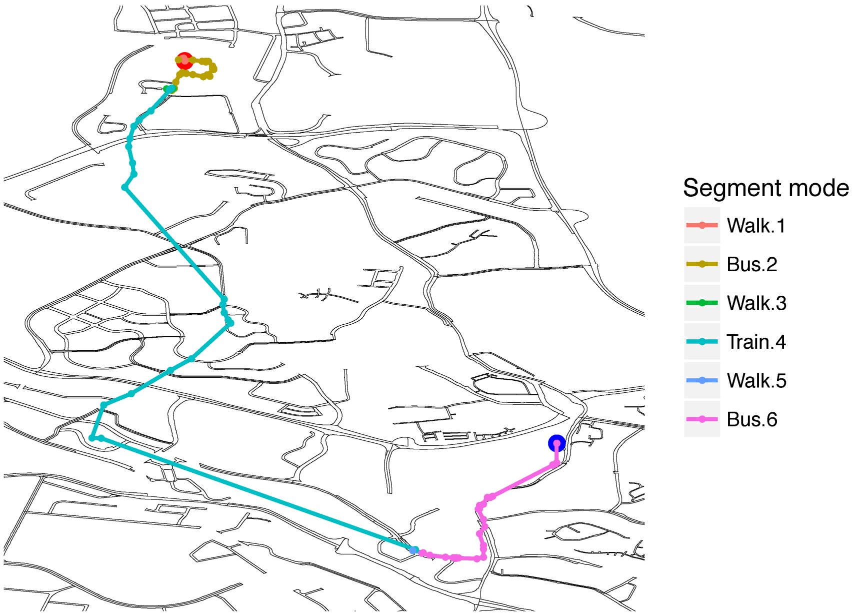

FIGURE 5.1: The subject walks from home (in red) to a bus stop to catch a train, after which a bus is taken to reach destination (in blue). Circles along the way represent one data point each.

29 minutes and 6 seconds, while it is 33 minutes and 33 seconds for the evening trip, with a greater average number of stops (possibly extraneous activities).

Second, the data source—students of Singapore—is not an exact representation of the Singapore population. However, their exposure to traffic during the morning hours—which are effectively the most congested conditions—allows us to infer properties of the system. The geographical distribution of their homes broadly correspond to the population density of Singapore, and thus provides additional confidence that the traffic and public transport conditions experienced by the subjects in the dataset are similar to other commuters in Singapore.

To ensure the quality of our empirical results, we perform a strict data cleaning process. For instance, trips where too few points (e.g., at a frequency greater than 1 per minute) were recorded are filtered out, as are trips where many students have simultaneous behaviour, indicating school buses. A total of 34,121 clean trips are considered, with 16,563 unique students and 89 schools. The number of students by school type is approximately equally distributed, hence capturing the routing behaviour of students over a large space in Singapore.

5.2 Equilibration and empirical consistency of routing decisions

The first key question we address is the equilibrium property of the Singapore road network. We ask whether the system has reached stasis, in the sense of consistency

TABLE 5.1: Number of students using consistently the same mode of transportation

Type	Number of students
Public transport only	3,417
Private transport only	2,174
Walk only	296
Multimode	491
Total $\geq$ 2 trips	6,378
Total $\geq$ 1 trip	11,439

of routing decisions by the subjects. If the system is at equilibrium, we should expect that the subjects' routing decisions do not vary greatly between successive days of study. We investigate the issue from two positions. First, we compare the modes of transportation selected by each subject over the days of the experiment. Second, we improve the previous result by considering whether routing choices in terms of the selected path are identical over experiment days.

Each subject carries the sensor for up to 4 days in a week, allowing us to compare the morning trips taken by the same subject between different days. Since the presence of noise in the sensor data and trip detection algorithm does not guarantee us that the whole week of experiment will be available, we filter out subjects for which only one morning trip is available. In our clean dataset, we have 16,563 individual subjects, out of which 11,439 have two or more trips logged in.

We first compare the transportation modes selected by subjects during the morning trips. In Table 5.1, we differentiate the 6,378 subjects who have consistently used the same mode of transportation in a week. These subjects represent 56% of the group that has two or more trips reported.

In a second and more granular analysis, we compare the routing decisions of the subjects at road level. Our aim is to determine whether each subject selects the same route consistently to reach school in the morning. To achieve this task, we need a distance $d_{a,b}$ measuring the similarity between two sequences $a = (a_i)_{i=1}^{n}$ and $b = (b_j)_{j=1}^{m}$ of coordinates. If the sequences are close, we can conclude that the same route is selected. For subjects employing the same mode of transportation across all days of experiment, the percentage of subjects selecting the same route is very high, in the order of 94%. We detail in Appendix B.1 how the distance between two routes in computed.

5.3 Individual optimality and empirical imitation-regret

Analysing the consistency of routing decisions across several trips gives a coarse picture of how well-founded these decisions are. It is possible to perform an additional check on the individual optimality of the subjects' decisions from the data itself. One can observe how different is the behaviour of comparable subjects, i.e., subjects who share the same origin, destination, time of travel and mode of transportation. As a necessary condition for equilibration of the system, we should expect to find these differences to be small.

5.3.1 Definition of the clusters

To find the sets of comparable subjects, we group the subjects by clusters, indexed by four variables:

- **Geographical location** l**:** Students living in the same neighbourhood are grouped together.
- **Time of departure** t**:** It is not accurate to compare a student departing from home at 6 am with one starting at 8 am. For this reason, subjects travelling *on the same day* and *within the same time frame* are grouped together, using a window size of 20 minutes.
- **Destination** s**:** Students going to the same school are grouped together. In the case that two or more schools share the same location (e.g. a Primary and a Secondary school), students attending either one of them are added to the same cluster.
- **Mode of transportation** m**:** The analysis is carried out over two modes of transportation, either private transportation (car, taxi) or public transportation (bus, train).

Two spatial clustering methods are implemented to group by origin and decide on the index l. In the first version, we find the smallest bounding box that contains all the home locations of the students. We divide this bounding box in cells of equal edge size r, e.g. $r = 400$ metres, and assign to the same geographical clusters students with home locations inside of the same cell. This is a grid-based method that partitions the space into a finite number of cells from a grid structure. Its main advantage is its fast processing time.

The second version of the spatial clustering approach is based on a distance rule where all home locations of the students in the same cluster should be within r metres of each other. This is a hierarchical clustering method using decision trees based in the

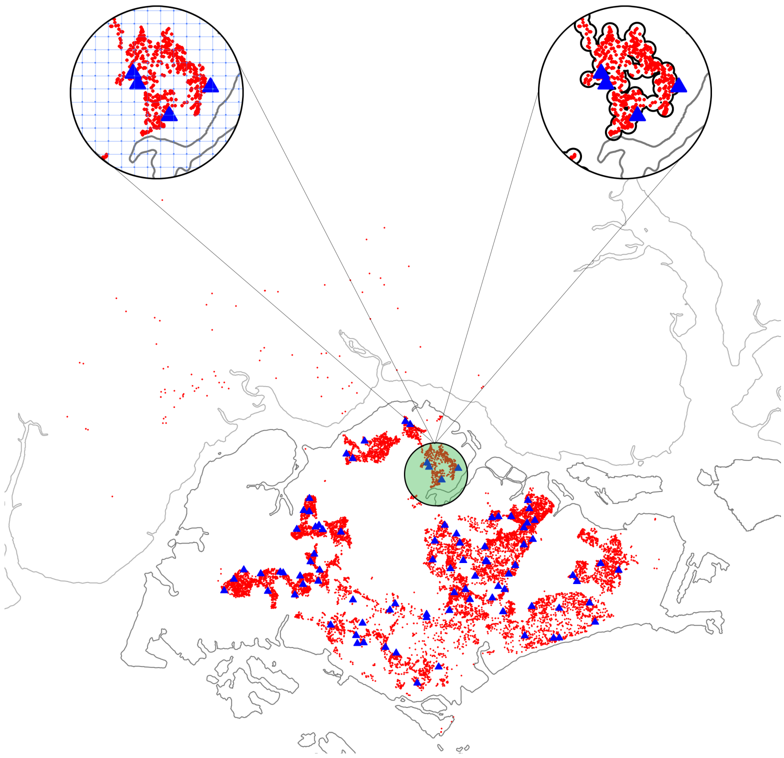

FIGURE 5.2: Home locations (red dots), school locations (blue triangles) and spatial clustering methods. Grid and circle clustering. (*Figure produced by F. Benita, with data from the author*)

geodesic distance matrix of all trips. This technique, although computationally more expensive, ensures that the distance rule holds for all the trips.

Figure 5.2 shows the visual comparison of the two different spatial clustering methods for $r = 400$ metres. The red dots mark the home locations of the subjects. The blue triangles mark the school locations of the 89 schools.[2] The grid-based method (top left) is a simple but efficient strategy, grouping the points that fall in specific cells of the mesh. On the other hand, the distance rule approach (top right) can be visualised by circles of diameter equal to 400 metres. Inside each circle, the maximum distance between any two home locations is 400 metres. The algorithm optimises a criterion function and the centroid of each spatial cluster can be easily identified, making it a powerful method to build the clusters. Recall that Figure 5.2 is presented only for visualisation purposes since inside each spatial cluster (cells/circles), students might be mixed among different transportation modes and different destination schools.

[2]It is interesting to note that some students have home location in Malaysia, and commute from Malaysia to Singapore daily for study.

5.3.2 Definition of the empirical imitation-regret

We obtain a set of clusters $\{C_{l,t,s,m}\}_{l,t,s,m}$ where each $C_{l,t,s,m}$ contains the trip durations $t^i_{l,t,s,m}$ of students in the cluster. If several students belong to the same cluster, we find the student whose trip has the minimum duration among all trips in the cluster. We call this trip the *baseline* $t^b_{l,t,s,m}$, against which the remaining trips will be compared.

We next define the imitation-regret for student i in cluster $C_{l,t,s,m}$ by

$$R^i_{l,t,s,m} = t^i_{l,t,s,m} - t^b_{l,t,s,m}$$

The imitation-regret for the baseline student is zero, and nonnegative for everyone else. We are interested in seeing how large the deviations from the baseline can be, as a necessary condition for the system to be at equilibrium is that these deviations must be close to zero. Sensitivity analysis is performed to account for the parameters used in the clustering method.

In this analysis, only students using one mode of transportation (public or private) are considered. Multimodal trips present the additional difficulty of comparing students that may have a different mix of transportation modes, with a low number of clusters with at least two comparable students. They are however discussed in the following findings.

Our notion of empirical imitation-regret shares its name with the traditional regret measure, commonly found in the learning and multi-agent systems literature. The subjects are faced with multiple strategies that they can choose from: all the routes that go from their neighbourhood to the destination. They may not know about current traffic conditions or which route will take the least amount of time but nevertheless have to make a decision. A posteriori, this decision can be compared with the best action they could have implemented on that day, and the difference is the imitation-regret. The appearance of the word "imitation" is due to the fact that we compare the decision solely with other players' choices of routes. A better route that is not used by any of the students in the cluster will therefore not be considered here. This restriction is shared with many natural learning dynamics (e.g., follow-the-leader dynamics) and thus can be interpreted as a reasonable assumption on subjects' decisions.

The measure of empirical imitation-regret depends naturally on the geographical area covered by the neighbourhood. As the area increases, so does the accumulated imitation-regret, since the minimum is taken over a larger set of students. However, geographical clusters lose in precision as their size gets larger, since two different subjects in the same cluster may have widely diverging trip durations. The results in this section use a geographical cluster size of about 400 metres, found to balance the two issues well. Additionally, we perform sensitivity analysis to show the robustness of our findings, presented at the end of this Section.

The value of 400 metres is picked for the following reasons. First, assuming a uniformly random distribution in the cell their expected distance would be a little over 200 metres. In practice this distribution is concentrated on blocks of flats and two students could easily be living in the same block. Let's assume that the students are at distance of 200 metres. There are two cases, either both the students drive/are driven or they take bus/metro. If they drive, this distance is noise. If they use metro/bus then since they go to the same school, they typically would use the same bus/metro and board it at the same regular time. The only differentiating factors are the difference in the distance they cover on foot (in a geometric world via triangle inequality less than 200 metres) and the difference between the amounts of buffer time (arrive a little earlier) at the stop. The rest of the route is identical.

How is the notion of empirical imitation-regret relevant to understand the decision-making of subjects and system properties? On the one hand, low empirical imitation-regret is a necessary condition for equilibrium. Indeed, at equilibrium, all comparable subjects should perform their trip in roughly the same amount of time. If an individual's imitation-regret is large enough, say, 10 minutes, she may be better off switching to a different route, e.g., the one used by the fastest individual in the cluster.

On the other hand, a high empirical imitation-regret warns us that some users are unable to find the fastest route to reach their destination. We see two possible directions to explore following such a conclusion. If we assume that individuals are solely interested in minimising their trip duration—a fair assumption for the morning trip, constrained by the hard deadline of the class start, also supported by the significantly longer trip durations in the evening commute—, then the network may benefit from the injection of information on how to traverse it. Otherwise, a high empirical imitation-regret reveals that other factors enter into consideration when the student is selecting the route, such as finding the least expensive one, the cooler one (in terms of temperature) or one that is shared with other students. The additional data collected by the sensor (e.g. temperature, proximity to other sensors) could indeed be articulated to uncover the nature of these factors.

5.3.3 Estimates of the empirical imitation-regret

In Figure 5.3, left, we plot the complementary cumulative distribution of the empirical imitation-regret. A point on the curve indicates the fraction of individuals (read on the y-axis) who have an empirical imitation-regret greater or equal than x (read on the x-axis). We also give the mean (solid red line) and median (dashed red line) experienced

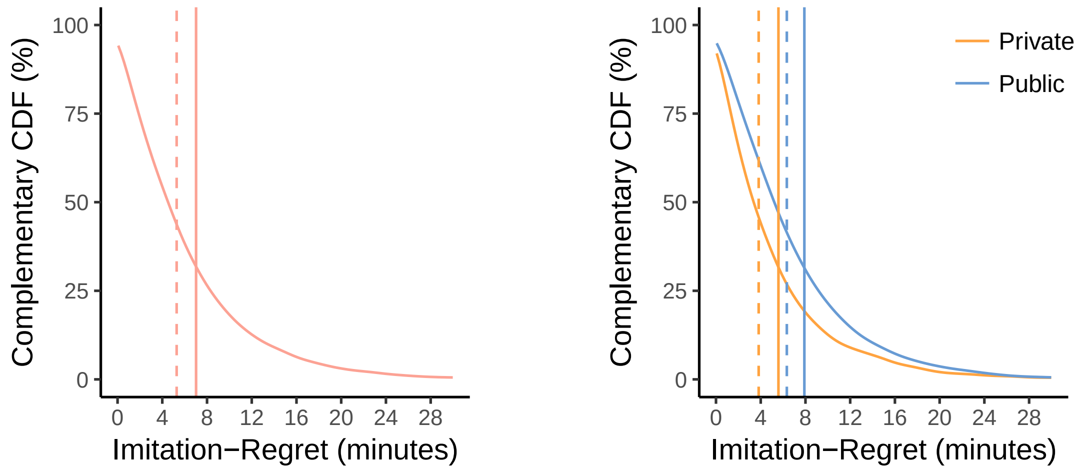

FIGURE 5.3: *Left:* Complementary cumulative distribution function of the imitation-regret. We aggregate all days of the experiment in a single figure and remove students with zero imitation-regret (in other words, the baseline students). The mean imitation-regret signalled by the red line is equal to 7 minutes, while the median imitation-regret plotted with the dashed blue line is equal to 5 minutes and 15 seconds. *Right:* Comparison of complementary CDF of imitation-regret per mode of transportation.

empirical imitation-regret. It should be noted that the empirical imitation-regret distribution and its moments do not include the students for which the imitation-regret is zero, i.e., the best in the cluster.

In the dataset of subject who appear in clusters with multiple trips, the average length of a trip is 25 minutes and 12 seconds. This locates the average regret of the trips at about 28% of the trip duration. This result motivates the introduction of a solution parametrised by two values, ϵ and δ. The reported measurements constitute an (ϵ, δ)-equilibrium if we find that a fraction $1 - \delta$ of users experience at most a quantity ϵ of imitation-regret. The experiment yields values $\epsilon = 15$ minutes and 20 seconds, and $\delta = 0.1$.

Second, one can investigate the imitation-regret by mode of transportation. In Figure 5.3, right, we plot the same complementary cumulative distribution function for both subjects who use only private transportation and subjects who use only public transportation. While the regret for private transportation is lesser than that of public transportation, the regret as a percentage of the trip duration inverts this relation. Indeed, the average duration of a trip made with a private vehicle is 13 minutes and 50 seconds, while the same average for public transportation is 31 minutes and 40 seconds. However, for private vehicle trips, the average is at 5 minutes, while it is 6 minutes and 40 seconds for trips taken in public transport. This locates the regret at respectively 37% and 21% of the trip duration. The greater variability of road conditions during car trips may explain why the regret as a fraction of the trip duration is greater.

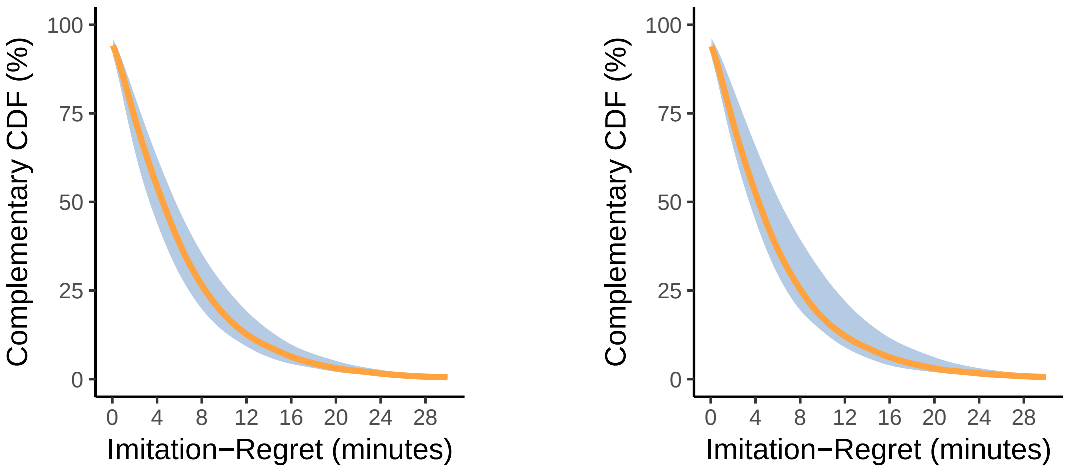

FIGURE 5.4: *Left:* We use as baseline the curve in orange representing the results for the grid clustering with cell size 400 metres. The blue area behind the line represents the variations when using ball clusters instead, of sizes 200, 400, 600, 800, 1,000 metres. *Right:* The baseline (in orange) is now the curve for the ball clustering with 400 metres diameter, with the blue area showing variations with the grid clustering.

For students using a mix of both modes of transportation, one can expect the complementary CDF to locate between the two curves, and thus the regret to be bounded by that of private transportation trips and public transportation trips. It is not represented on the Figure 5.3, right.

Third, we study the imitation-regret between modes, i.e., taking the regret with respect to the fastest individual in the cluster, irrelevant of transportation mode. We focus our analysis on mixed clusters, where at least one individual using public transportation *only* and one individual using private transportation *only* appear. This analysis is carried over all clustering procedures, i.e., both ball and grid clustering with the five different sizes.

From the smallest cluster cells to the largest ones, we have between 439 and 1,953 such mixed clusters.[3] However, for all cluster sizes, the fraction of trips completed faster in public transportation than in a private vehicle is located around 5%.

We show in Figure 5.4 that our measure of imitation-regret is stable with respect to the clustering procedure. Although different cluster sizes as well as the two cluster methodologies (ball or grid) differ, the qualitative results hold for all procedures, while descriptive statistics are tightly bounded.

[3]As the size of the cluster cell decreases, it is indeed less likely to contain both subjects using public transportation and subjects using private transportation.

5.4 The Stress of Catastrophe

To measure the impact of congestion on the efficiency of the network, we compare the travel times in our dataset with estimated travel times in optimistic conditions. The ratio between the two is defined as the stress of catastrophe (Section 5.4.1). We present our methodology to estimate the free-flow travel times, using a graph representation of Singapore's road network for private transportation trips and queries to an online service for public transportation (Section 5.4.2). Finally, the estimations of the stress of catastrophe are given for all modes and discussed (Section 5.4.3).

5.4.1 Definition of the stress of catastrophe

The stress of catastrophe (SoC) is introduced to give a measure of the weight of externalities in the system. As more agents join the road network, congestion increases on the links. Classically, the PoA has been employed to quantify how bad the selfish decision-making of these agents affects the efficiency of the system, compared to the social optimum implemented by a central planner. However, PoA does not fully capture the effects of a tragedy of the commons that congestion presents. In such a scenario, it is not costly for one additional individual to enter the system, but since all agents do so, the global welfare is very much diminished. Similarly, congestion can reach levels after which the action of a central planner has little effect, yielding a low PoA that does not reflect just how congested the system is (Colini-Baldeschi et al., 2017).

On top of this, estimating the equilibria and optima of a routing system are data-intensive tasks. First, demands need to be known or estimated from samples for every origin-destination pair of the agents. Second, latency functions for every edge of the network need to be estimated. Third, the global optimum and equilibrium flow need to be computed, so as to compare their respective cost and estimate the PoA. Recent works have followed this approach, finding that PoA is low, or even equal to 1 (no inefficiency) (Zhang et al., 2018; Wu, Möhring, and Xu, 2018).

The method we offer in comparison does not require to extrapolate from a sample or record the behaviour of the whole system, and instead provides an estimate of an upper bound of PoA via empirical data. We are only concerned with obtaining a dataset for which our samples experience congestion as any individual in the general population would, a point we discussed in Section 5.1.6 with respect to our dataset.

The SoC eschews these pitfalls by providing an optimistic lower bound to the socially optimal trip durations. It stems from the simple fact that a crude lower bound to the optimal trip duration is one in which no one else is present on the road. With the road street network of the city at our disposition and a shortest path algorithm, free-flow trip durations are obtained. They give us a "blue sky"—i.e., ideal scenario—lower bound. Comparing the actual recorded trip duration length to this lower bound in turn yields a ratio of how much faster the trip could have been in a no-externality scenario. Formally, we define the SoC from our data as such:

$$\texttt{SoC} = \frac{\texttt{Cost}\,\texttt{(Recorded trip duration)}}{\texttt{Cost}\,\texttt{(Free-flow trip durations)}}.$$

Since the denominator is a lower bound to the socially optimal cost, we also have the following corollary:

$$\texttt{PoA} = \frac{\texttt{Cost}\,\texttt{(Recorded trip durations)}}{\texttt{Cost}\,\texttt{(Optimal trip durations)}} \leq \texttt{SoC}.$$

5.4.2 Estimation of the free-flow trip duration

From a road map of Singapore is computed a graph representation, where each vertex is located at an intersection or a bend in the road. An edge connecting two vertices indicates the presence of a segment of road going from one vertex to the other. Edges also possess additional metadata: their physical length (in meters) as well as the road type—such as expressway, local street, arterial road, etc.

Given this information, we create five scenarios, corresponding to five different speed profiles on each edge type, from a very fast profile to a very slow one, defined from existing speed limits in Singapore.[4] The time to traverse one edge is computed as the length of the segment divided by the speed on the segment. Finally, for each private transportation trip for all students in the dataset, we associate its origin and destination with the closest vertex in the graph before running a shortest path algorithm to estimate the free-flow travel time of the trip.

For public transportation trips, we query an online oracle, the Google Maps API. The API is not time-dependent, i.e., will not return results that depend on the congestion. To obtain a best case trip duration, we remove potential waiting time at train or bus stations. To minimise the number of requests to the API, we employ the grid clustering method described in the previous section and query for the best route between a non-empty cluster's centroid and a school. Some of the requests do not return satisfactory results, either due to a "Not Found" error or when the algorithm repositions the starting point of the trip too far from the student's home. These unsuccessful requests are dropped in the analysis.

[4] Exact speed profiles are given in Appendix B.2.

TABLE 5.2: Estimated free-flow mean trip duration and stress of catastrophe.

Scenario	**Mean trip duration (min)**	**Stress of catastrophe**
Very fast	8′05	1.83
Fast	9′06	1.63
Medium	10′22	1.43
Slow	12′12	1.22
Very slow	15′15	0.97
Public transport	32′03	1.05

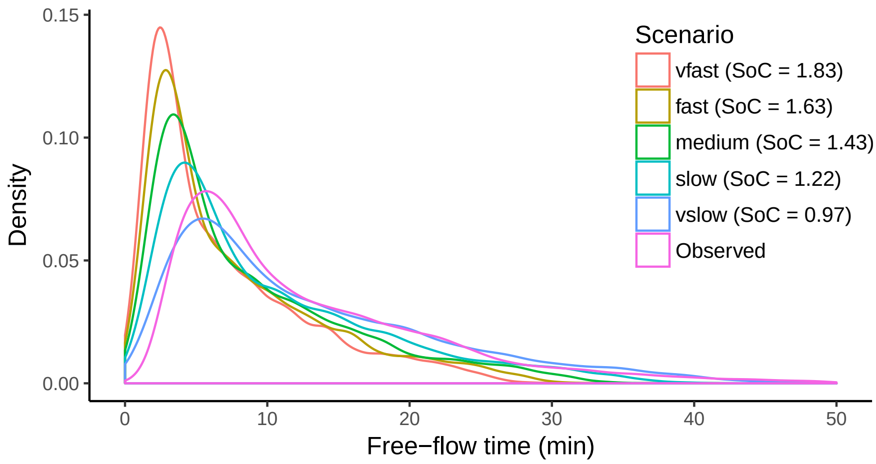

FIGURE 5.5: Private stress of catastrophe, density distribution.

5.4.3 Estimation of the stress of catastrophe

For each trip, as described above, we obtain five estimates of free-flow duration, one per scenario. Overall, all scenarios effectively give a lower bound to the recorded private vehicle trip durations, with the exception of the "Very slow" speed profile for which the SoC is below 1. To obtain the SoC under one particular scenario, we sum up all recorded durations from the dataset and divide by the sum of free-flow estimates. The obtained measures are recorded in Table 5.2.

Figure 5.5 displays the densities of free-flow estimates of trip duration. Under the optimistic scenario, where all speeds are taken to approach the enforced speed limit on real Singapore roads, the stress of catastrophe equals 1.83 whereas the most pessimistic profile yields a stress of catastrophe of 0.97.

The five scenarios offer more or less stringent lower bounds on the real travel time of experiment subjects. The result strikingly reveals that even under the most optimistic scenario, the SoC is still well under the theoretically-derived PoA for real road networks. For nonatomic congestion games with quartic cost functions—classically employed to model real road networks—the PoA is 2.151. Thus, the experiment data reveals that the worst case estimate given by the price of anarchy may be largely pessimistic.

For subjects travelling on the public system of Singapore, the estimate of the SoC is 1.05, indicating a much lower impact of congestion on the travel time. In a city where 65% of the population daily travels in public transportation, the low level of SoC for this mode has important policy implications.

To put it all together, games featuring links with a congestion element (e.g., the monomial $\alpha \cdot x^4$ for quartic cost functions) represent a private transportation network. Adding constant cost links is akin to adding public transportation options to the game. Whether we estimate the SoC for cars only, or mixing public and private transportation, we obtain an upper bound of PoA below that given to us by the worst-case example—and more so when both modes are mixed. The next section will provide arguments to explore why this may be the case in real transportation networks.

5.5 Why is PoA small in real transportation networks?

In this section, we add a strong assumption on the strategy sets of commuters. We find that this assumption yields lower PoA bounds, perhaps explaining the empirical results of Section 5.4. We also provide experimental justification for the assumption.

We note that such an assumption was developed independently in Bilò and Vinci (2018). While their work computes precise values of the PoA for linear congestion games, we provide more general bounds for parallel links networks using a different strategy. Future work may provide a more complete picture of PoA behaviour with this assumption.

5.5.1 Definitions and model

We use the framework introduced in Section 2.6 to define a routing game Γ. In this section, cost functions are specified as Bureau of Public Roads-type, i.e. affine monomials (Public Roads, 1964), expressed as

$$c_e(f_e) = \texttt{FF} + \texttt{FF} \cdot \alpha \Big(\frac{f_e}{\texttt{CAP}} \Big)^\beta$$

where α and β are conventionally taken to be 0.15 and 4, respectively, while $\texttt{CAP}$ is the capacity of the edge and $\texttt{FF}$ is the free-flow time, or the time to traverse the edge when congestion is nil.

The *path cost* of path P under flow f is hence given by

$$c_P(f) = t_P + \sum_e \alpha_e f_e^d, \text{ where } f_e = \sum_{P \ni e} f_P$$

We introduce the concept of δ-free-flow routing games in the next definition.

Definition 5.5.1. A routing game is a δ-free-flow routing game if for all $P \in \mathcal{P}$, we have $t_P \leq (1+\delta)t_M$, where $t_M = \min_P t_P$.

In a δ-free-flow routing game, the paths available to agents have free-flow time contained within a fraction δ of each other. We provide two preliminary justifications for this model, with empirical evidence in the following Section 5.5.4.

What is the value of δ in the Pigou network of Figure 1.2? Since the variable cost link has latency $c(x) = x$ and the constant cost link $c(x) = 1$, δ ought to be set to ∞. However, in real road networks, it is physically impossible to "jump" from a vertex to the next, so δ must be strictly bounded. Could this be the reason why Pigou, a worst-case network for many PoA bounds, behaves so differently than real transportation networks, where the PoA is lower (Section 5.4)?

Second, we must ask why δ, if bounded, is low in general. Faced with the exponentially large number of paths connecting its origin to its destination, it is possible that agents use a simple heuristic to prune the space of strategies down to a few possible options. It does not appear realistic for an agent to consider *all* possible paths connecting its origin to its destination.[5] Letting the agent decide between paths for which the free-flow time (or intuitively, the *directness*) is low seems a priori reasonable.

Note too that the concept of δ-free-flow routing must be taken over the available paths of the agents. In Section 2.6, we had defined the path cost as the sum of edge costs $c_P(f) = \sum_e c_e(f)$. Here, we have explicitly separated the free-flow duration of the path from the cost incurred over each edge, as $c_P(f) = t_P + \sum_e \alpha_e f_e^d$.

5.5.2 δ-free-flow routing for $\delta = 0$

We start the discussion with the case of $\delta = 0$. This is equivalent to setting all free-flow costs of the available paths to the same value t_M, and thus is strategically similar to routing games with monomial latency functions without an affine constant. For $\delta = 0$, price of anarchy is 1, a fact noted in Roughgarden and Tardos, 2002 for linear cost functions, without the fixed free-flow time latencies in our case. We provide a proof using a non-linear programming approach to the optimal flow f^*.

Proposition 5.5.1. *For parallel links networks, if* $\delta = 0$, $\text{PoA}(\delta) = \text{PoA}(0) = 1$.

[5] To take a contrived example, one would not drive through Kuala Lumpur to reach the East of Singapore from the West.

Proof.

$$\begin{aligned} C(f) &= \sum_e f_e(t_M + \alpha_e f_e^p) \\ &= \mu t_M + \sum_e \alpha_e f_e^{p+1} \end{aligned}$$

At equilibrium, $\alpha_e f_e^p = K$ for all e for some $K > 0$.

The optimum flow f^* is attained by solving:

$$\begin{aligned} \min \sum_e \alpha_e f_e^{p+1} \\ \text{s.t.} \sum_e f_e = \mu \end{aligned} \tag{OPT-PL}$$

We can set up the Lagrangian $L(\lambda) = \sum_e \alpha_e f_e^{p+1} + \lambda(\mu - \sum_e f_e)$. We get

$$\frac{\partial L}{\partial f_e} = 0 \Leftrightarrow \frac{\lambda}{p+1} = \alpha_e f_e^p,\ \forall e$$

which implies that all edge latencies are equal at optimum and thus f^* has the same cost as the equilibrium. □

A modification of the proof yields the result for general networks, with several commodities $k = 1, \ldots, K$, demands (μ_k), $\sum_{k=1}^K \mu_k = \mu$, paths $P_k \in \mathcal{P}_k$ and free-flow time for commodity k, t_k.

Proposition 5.5.2. *For general networks with several commodities, if $\delta = 0$, $\texttt{PoA}(\delta) = \texttt{PoA}(0) = 1$.*

Proof.

$$\begin{aligned} C(f) &= \sum_{k=1}^K \sum_{P \in \mathcal{P}_k} f_P(t_k + \sum_{e \in P} \alpha_e f_e^p) \\ &= \sum_{k=1}^K \mu_k t_k + \sum_{k=1}^K \sum_{P \in \mathcal{P}_k} f_P \sum_{e \in P} \alpha_e f_e^p \\ &= \sum_{k=1}^K \mu_k t_k + \sum_e \alpha_e f_e^{d+1} \end{aligned}$$

At equilibrium, $\sum_{e \in P} \alpha_e f_e^p = X$ for all P for some $X > 0$.

The optimum flow f^* is attained by solving:

$$\begin{aligned} \min \sum_e \alpha_e f_e^{p+1} \\ \text{s.t. } \forall k, \sum_{P\in\mathcal{P}_k} f_P = \mu_k \\ \forall e, \sum_{P\ni e} f_p = f_e \end{aligned} \tag{OPT-GN}$$

We can set up the Lagrangian

$$L((\lambda_e)_e, (\gamma_k)_k) = \sum_e \alpha_e f_e^{p+1} + \sum_k \gamma_k(\mu_k - \sum_{P\in\mathcal{P}_k} f_P) + \sum_e \lambda_e(f_e - \sum_{P\ni e} f_P)\,.$$

We get

$$\frac{\partial L}{\partial f_e} = 0 \Leftrightarrow -\lambda_e = \alpha_e(p+1)f_e^p,\, \forall e$$

and for $P \in \mathcal{P}_k$

$$\begin{aligned} \frac{\partial L}{\partial f_P} = -\sum_{e\in P}\lambda_e - \gamma_k = 0 &\Leftrightarrow \gamma_k = -\sum_{e\in P}\lambda_e \\ &\Leftrightarrow \frac{\gamma_k}{p+1} = \sum_{e\in P}\alpha_e f_e^p \end{aligned}$$

which implies that all path latencies are equal at optimum and thus f^* is an equilibrium. Since all equilibria have same cost in symmetric nonatomic games, $\texttt{PoA}(0) = 1$. □

5.5.3 δ-free-flow PoA for general networks

We have now a game Γ played on a network, with social cost function C. We let again C_T represent the social cost of the game Γ_T played on the same network with free-flow path costs all set to t', with $t' \leq (1+\delta)t_P$ for any path P.

Proposition 5.5.3. *Let f^* be the optimal flow for Γ, and f_T^* the optimal flow for Γ_T. $C(f^*) \geq \frac{C_T(f_T^*)}{1+\delta}$.*

Proof. Let f be a feasible flow.

$$\begin{aligned} C_T(f) &= \sum_P f_P(t' + \sum_{e\in P}\alpha_e f_e^d) \\ &\leq \sum_P f_P((1+\delta)t_P + (1+\delta)\sum_{e\in P}\alpha_e f_e^d) \\ &= (1+\delta)\sum_P f_P(t_P + \sum_{e\in P}\alpha_e f_e^d) \end{aligned}$$

$$= (1+\delta)C(f)$$

Thus $C(f^*) \geq \frac{C_T(f^*)}{(1+\delta)} \geq \frac{C_T(f_T^*)}{1+\delta}$. □

For parallel links network, we can prove that if $\bar{f}_T$ is an equilibrium flow of Γ_T, then $C(\bar{f}) \leq C_T(\bar{f}_T)$.

Lemma 5.5.4. $C(\bar{f}) \leq C^T(\bar{f}_T)$

Proof. Suppose $C(\bar{f}) > C^T(\bar{f}_T)$. Let W be the equilibrium latency on a link of the original game Γ and W_T that of Γ_T. $C(\bar{f}) > C^T(\bar{f}_T)$ is equivalent to $W > W_T$.

Note first that at equilibrium all links $\bar{f}_T$ are used and so $\bar{f}_{T,e} > 0$ for all e.

Two cases can take place:

1. $\exists e^*$ such that $\bar{f}_{e^*} = 0$. This implies that $t_{e^*} \leq W_T$, so W is not the latency of an equilibrium, contradiction.

2. $\bar{f}_e > 0$ for all $e > 0$. We then have

$$\begin{aligned} & t_e + \alpha_e \bar{f}_e^d > (1+\delta)t_M + \alpha_e \bar{f}_{T,e}^d, \ \forall e \\ \Rightarrow & \alpha_e \bar{f}_e^d - \alpha_e \bar{f}_{T,e}^d > (1+\delta)t_M - t_e \geq 0, \ \forall e \\ \Rightarrow & \alpha_e \bar{f}_e^d > \alpha_e \bar{f}_{T,e}^d, \ \forall e \\ \Rightarrow & \bar{f}_e > \bar{f}_{T,e}, \ \forall e \end{aligned}$$

   where the second implication comes from the definition of $(1+\delta)t_M$, weakly greater than all other free-flow times. This is a contradiction since both flows sum up to the demand.

□

We also have $C_T(f_T^*) = C_T(\bar{f}_T)$ and thus the following bound holds.

$$\texttt{PoA}(\delta) = \frac{C(\bar{f})}{C(f^*)} \leq (1+\delta)\frac{C_T(\bar{f}_T)}{C_T(f_T^*)} = 1+\delta$$

5.5.4 Empirical justification of δ-free-flow routing

The graph introduced in Section 5.4 to estimate free-flow travel time is employed here. We look for two measurements:

- One, we estimate the *best free-flow time* for a morning trip in the dataset, as we have done in Section 5.4.

- Two, we also estimate the *data free-flow time* of the route selected by the subject for its morning trip, presented in the next paragraph.

Comparing the two values allows us to estimate the δ for our population.

Trip reconstruction

A private transportation trip segment measured by the sensor consists of a stream of geographical locations. For each datapoint, we associate the closest edge in the graph. The size of the graph (61151 vertices and 65596 edges) implies a lengthy lookup phase to associate the point to its closest edge. For this reason, we consider a smaller dataset of 449 car segments out of the 17,897 segments in the larger dataset. These selected segments are well distributed across Singapore.

The direction in which the subject traversed the edge is assigned by a heuristic based on the distance of each endpoint to the endpoints of edges preceding and following the edge under consideration. In other words, the heuristic attempts to minimize the amount of back and forth, selecting the direction that least creates deviations.

Information on the origin and destination of the trip as well as the list of directed edges traversed by the subject does not suffice. Where the sensor does not record a datapoint,[6] we must provide a best guess on which edges were crossed during the trip.

- For gaps of small length between two directed edges e_1 and e_2 (in that order), we compute the average speed between the two edges and drive a straight line between the target of e_1 and the source of e_2. The duration to cross this gap is obtained as the geographical distance divided by the average speed.

- For gaps of larger length, we run a shortest path algorithm between the target of e_1 and the source of e_2.

The data free-flow time is finally obtained as the sum of durations of redirected edges, small gaps and large gaps.

The final free-flow duration of the selected route is obtained as the sum of durations to traverse the redirected edges, small gaps and large gaps.

Estimation of δ

For each trip segment, two estimates are obtained: the best free-flow time and the data free-flow time. We call *deviation* the ratio between these two estimates. The deviation is strongly related to the parameter δ we introduce in Section 5.5.1. It measures the free-flow time difference between the best route the subject could have chosen and the route actually selected, both in a situation of no congestion. The distribution of the deviation among subjects provides a clue to estimating δ for the routing game of Singapore. A

[6] Geographical location is obtained by scanning surrounding WiFi access points. The method does not always yield accurate enough measurements, but the issue can be mitigated with proper data processing (Monnot et al., 2016).

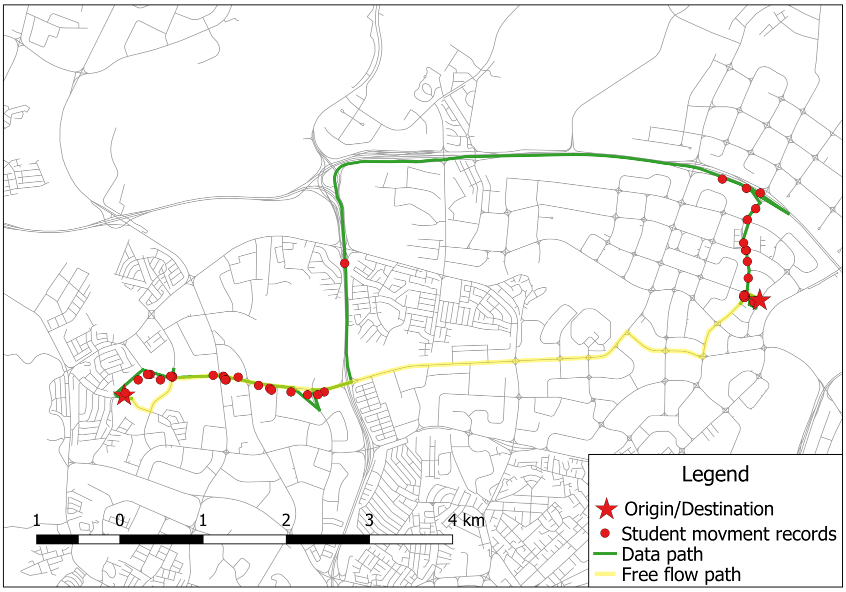

FIGURE 5.6: For each trip segment, we find the best free-flow time and the data free-flow time. The reconstruction of the selected route uses datapoints logged along the trip. In yellow, the fastest route in free-flow condition is highlighted. The reconstructed route is in green, along which we find the data free-flow time. (*Figure produced by F. Benita, with data from the author.*)

Quartile	0%	25%	50%	75%	100%
Deviation	-0.68	0.17	0.45	0.88	3.53

TABLE 5.3: Quartiles of deviation.

small δ yields support to the hypothesis that agents only consider routes which connect origin and destination in a straightforward manner (under no congestion) as part of their strategy set (Figure 5.7 and Table 5.3).

5.6 Discussion

The price of anarchy of real routing systems is a question that has long eluded experimental researchers. To our knowledge, this experiment was among the first to investigate from such granular data the efficiency of a mixed transportation system. Another effort, led by Zhang et al. (2018), has employed a different data source, namely the rate of passage of cars on a reduced section of the Boston transportation network. By estimating the demands and cost functions on the road, the authors were able to estimate an optimal assignment and compare it with the recorded flows. A large dataset obtained from cell phone location data was employed by Çolak, Lima, and González (2016) to measure the impact of selfish routing on congestion, in five different cities.

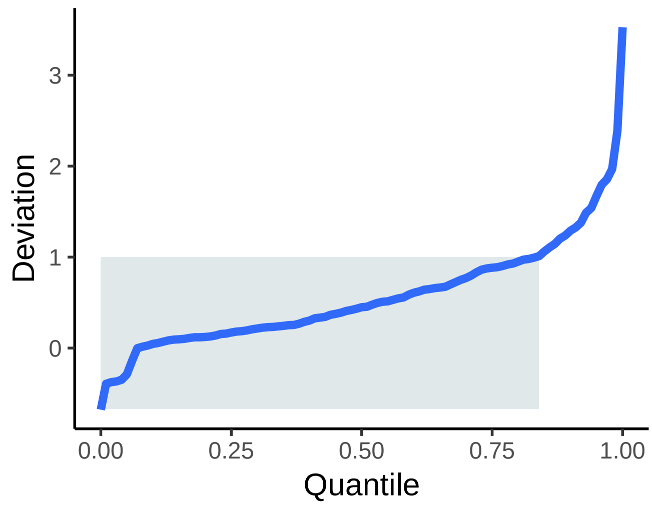

FIGURE 5.7: The deviation is measured by the ratio of the selected route free-flow time to the minimum free-flow time among all routes between the origin and the destination. Close to 80% of values are below 1, implying that the free-flow time of the selected route is rarely twice as long as the best free-flow time.

Demands were estimated from the data while road information was obtained from OpenStreetMaps, with optimal assignments given by an online provider.

All studies, including ours, rely on so-called "Big Data" methods. This is not a coincidence: The biggest obstacle to investigating the efficiency of routing system is the collection of data, of which large amounts are needed for accurate estimations. On the one hand, it is possible to have a relatively coarse understanding of the system's performance, from Household Interview Travel Surveys (HITS), but these surveys do not reliably contain granular routing decisions. Although they represent a sample of the population, they remain useful to derive origin and destination demand matrices.

It is however clear that with the propagation of low cost sensors and data policies for smart cities (e.g., requiring detailed travel data from taxis and/or ride-sharing operators), access to this type of information will be facilitated. We expect further studies to not only experimentally assess the current efficiency of commuting systems but to predict the behaviour of the system under stress or variations in its design. Areas of game theory such as behavioural studies can be profitably invoked to understand how traffic changes in response to the introduction of tolls or other pricing mechanisms. In this drive for efficiency, it is important to consider which side effects may arise, a question we pose with respect to inequality in the next Chapter 6.

Chapter 6

How does the drive for efficiency affect wealth inequality?

Our estimates of the stress of catastrophe in Chapter 5 showed a clear difference between subjects in private transportation, for whom the SoC was larger than 1, and subjects in public transportation, with an SoC (and by extension, a PoA) much closer to 1. The latter thus reduce the average inefficiency due to selfish routing, but incur the cost of higher latencies due to a lower transportation speed.

Is this a general phenomenon? Price mechanisms for transportation, be they "embodied" in the cost of buying a private vehicle, calling a car from a ride-sharing operator or a traditional taxi company, or explicitly levied with tolls, have long been hailed as an effective way to deal with congestion. Resources are rare and one private mode of transportation imposes more externalities than any other, in the form of increased pollution, noise or land use waste, and thus adequately pricing these externalities is critical.

But inefficiency, as PoA reveals, also appears distinctly as a result of self-interested routing decisions. Even assuming all previous externalities are inexistent, selfish routing introduces more inefficiency from routing decisions alone, which in general are bounded away from a socially optimal state (Figure 1.2b). There again, pricing with tolls can induce efficient equilibria.

In its most basic formulation—marginal cost pricing where agents trade time with money exactly one-to-one—the celebrated Pigouvian tax internalises the externalities of selfish routing and updates the cost functions such that an equilibrium flow in this new game is also an optimal flow of the original game. If the "one-to-one" assumption on elasticity is too strong,[1] one can introduce agent types who trade time with money according to different rates, and obtain similar results.

As we move closer to representing the true value a user has for outcomes of a game, the question of distributional justice becomes inevitable. That a user would trade, e.g., time against money says little if we do not consider as input to the game some initial

[1] In the Pigou model, cost is obtained by a simple latency + price addition, see Figure 1.2d.

distribution of wealth. And once we do, it is natural to ask how this distribution is affected by the mechanism we choose to design.[2]

The question of inequality is all the more timely as a surprisingly wide range of topical issues finds its roots in the increasing polarisation of income and wealth, from climate change (Boyce, 1994; Torras and Boyce, 1998; Cardenas, 2007; Baek and Gweisah, 2013) to the degradation of democracies (Stiglitz, 2012; Acemoglu et al., 2015; Giridharadas, 2018). Inefficiency from selfish routing, as we have seen, can be mitigated with tolls, yielding a lower overall congestion. But how to do so while maintaining equity—a desiderata expressed in the Paris agreement of 2015 (UNFCCC, 2015)—remains a contentious debate.

The effects of congestion pricing, for this reason, have been analysed from the two viewpoints exposed above: efficiency and equity. On the one hand, its positive effects on overall congestion were noted both theoretically (Pigou, 1920; Fleischer, Jain, and Mahdian, 2004) and experimentally (Olszewski and Xie, 2005). On the other hand, the efficiency-driven models of routing games have not addressed the question of equity, discussed more thoroughly in empirical surveys of transportation policies (see Levinson (2010) for a review). This gap is all the more troubling as, we argue in this chapter, it is not an insurmountable one—as long as we correctly model it. Once inequality can be reasoned about in the framework of algorithmic game theory, its tension (or convergence) with efficiency is better understood and opens up broader avenues of research.

The model we propose explicitly integrates distributional effects into its formulation. When agents, with varying wealth and values of time to money, participate in the game, their choices are informed by their resources. As such, the model is well-suited to represent the effects of any form of pricing, in particular tolls, on the distribution of wealth. Once general results are established for broad classes of games, we propose extensions and examples of applications to designers who integrate equity as part of their objective function.

Contents of this chapter

The legwork presenting routing games formalism was done previously in Section 2.6. Our discussion of inequality in routing games will reuse key concepts while introducing wealth-endowed agents into the model in Section 6.1.[3] We also provide additional

[2]Note that even if we do not care about each user's elasticity for time, the question of inequality lies just below the surface. The equilibrium flow in the original Pigou example is perfectly fair: all users incur the same latency. But what can be said about the optimal flow, where one half of the users must make an individually suboptimal choice, a half usually left unspecified? The question may be settled by appealing to a lottery, where each user has a 50% chance to be assigned to the slow link, or a repeated game where on even days, the first half gets to travel on the fast link while on odd days, the second half goes.

[3]There are important differences between inequality of wealth and inequality of income, which we do not tackle in this chapter. The word "wealth" is taken in the sense usually given in microeconomic models of "some initial endowment".

definitions related to the measurement of inequality.

Section 6.2 establishes a broad result on inequality in symmetric routing games, showing that any assignment of tolls must increase inequality in the population. Economic equality can be measured in myriad different ways, but we show that any index satisfying four well-known axioms will yield the same result.

For the asymmetric case, treated in Section 6.3, we cannot obtain such a clear-cut result. We show via a decomposable inequality index, the Theil L index (Theil, 1967), and a counterexample that inequality can go in either direction. The symmetric case helps us argue why this may be the case.

Computationally, one may care not only about the sign of the deviation, but also about its magnitude. We introduce the inequity index in Section 6.4, a measure of the marginal impact of a game on inequality. The inequity index in our model has natural properties such as scale invariance or robustness to no-regret learning agents. Following this, a Pigou network with wealth is presented in Section 6.5, exemplifying the inequality increase in symmetric games and the tension with PoA.

We explore in Section 6.6 how a designer could modify the game to recycle revenues raised from the tolls, or effect a tradeoff between efficiency and inequality. This thread, explored in transportation policy literature and favoured by policy makers, can naturally be pursued in our model.

Finally, in Section 6.7, we come back to the NSE data and uncover experimental validation of our model. A methodology is proposed to estimate the inequity index, in spite of limitations due to the absence of precise wealth data.

6.1 Congestion games with wealth

Section 2.6 presented routing games with homogeneous agents who care solely about their latency. By modifying the edge cost functions to include an edge price discounted by wealth, we integrate monetary costs into the game. This translates to the *ex post* wealth distribution, where the cost of the game is subtracted from the initial wealth of each agent.

Wealth We have a continuum of *types* which lie in $[0, 1]$. Type x has wealth $q(x)$, where q is the *quantile function* of the wealth of a population of agents — that is, $|z : q(z) \leq q(x)| = x$, where $| \cdot |$ is the Lebesgue measure. We shall further assume that $q(0) > 0$ and q is measurable and nondecreasing. Typically, we will assume a continuum of types and a strictly increasing, continuous q. In this case, if we treat wealth as random variable, then q expresses the inverse of its cumulative distribution function.

Flow A *flow* $F : [0,1] \mapsto \mathcal{P}$ is a mapping from types to paths. We shall only need to consider *canonical* flows, that is, flows F which divide $[0,1]$ into finitely many intervals, and map the interiors of those intervals to one path in $\mathcal{P}$; that is, F is specified by a finite number of reals $a_1 = 0 \leq a_2 \leq \cdots \leq a_{N+1} = 1$ such that $F(b) = F(c)$ for all i and $b, c \in (a_i, a_{i+1})$. Section 6.2 expands on this issue.

Let F be a flow. The *congestion* of this flow, c^F, is a function mapping E to the nonnegative reals, where $c^F(e) = |\{x : e \in F(x)\}|$, where $|\cdot|$ denotes the Lebesgue measure.

Agent cost The *value of travel time* (VTT) relates the price of the chosen path to the latency incurred by the player. Thus, the monetary cost of a player x on path $F(x)$ is defined to be

$$\text{cost}^F(x) = \sum_{e \in F(x)} \ell_e(c^F(e)) \cdot \texttt{VTT} + \tau_e \,.$$

There is an extensive discussion in the transportation literature of the true cost of transportation to the traveler and the value of time, see (Abrantes and Wardman, 2011; Börjesson, Fosgerau, and Algers, 2012) for some of the most recent papers. This field has established and studied the *income elasticity* of the value of (travel) time (informally, the precise nature of the term $\texttt{VTT}$ above) and validated and measured it through extensive surveys and other studies over three decades. The upshot is that the cross-sectional elasticity (that is, the elasticity with regressive corrections across causal parameters such as having children and living in the capital) is constant across long periods of time, and that the precise relationship seems to be $\texttt{VTT} = q^\beta$ where $\beta \leq 1$ is conventionally taken to be one, even though certain countries, such as the UK, use value 0.8. The model follows $\beta = 1$, meaning that if player x traverses an edge with price m and delay d, the perceived cost is, in monetary terms, $d \cdot q(x) + m$.

Put together, the cost of player x under flow F is

$$\text{cost}^F(x) = \sum_{e \in F(x)} \ell_e(c^F(e)) \cdot q(x) + \tau_e \tag{CAN}$$

We call this agent cost function *canonical* and show further that it is a natural choice with good properties (Section 6.4).

Nash equilibrium We say that a flow F is a *Nash equilibrium* if for all types x and for all paths $P \in \mathcal{P}$

$$\text{cost}^F(x) \leq \text{cost}^{F,P}(x) = \sum_{e \in P} \ell_e(c^F(e)) \cdot q(x) + \tau_e \tag{NE}$$

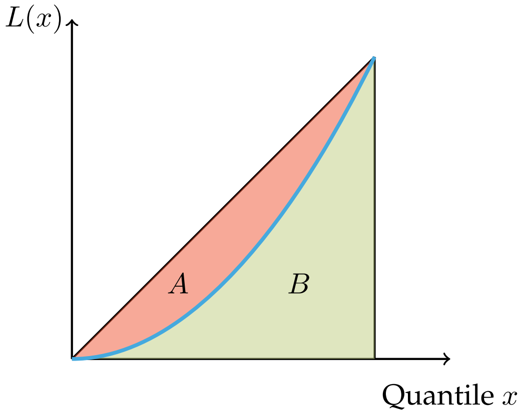

FIGURE 6.1: The Lorenz curve is plotted in blue. The green area is $B = \int_0^1 L(t)dt$. The Gini coefficient is then $G = 1 - 2B = 2A$.

that is, if no type x would be better off by deviating to another path $P \in \mathcal{P}$. We let $\text{cost}^{F,P}(x)$ denote the cost for agent x in flow F if it was using path P instead of $F(x)$.

Gini coefficient The Gini coefficient (Gini, 1921) is a central measure of inequality.

Definition 6.1.1. The Gini coefficient of wealth distribution q is given by

$$G(q) = 1 - 2\int_0^1 L(t)dt$$

where $L(t)$ is the Lorenz curve, or the fraction of total wealth held by individuals under and at quantile x.

$$L(t) = \frac{1}{\mu}\int_0^t q(x)dx = \frac{1}{\mu}Q(t) \tag{LC}$$

for $Q(t) = \int_0^t q(x)dx$, the cumulative wealth up to quantile t. We show in Figure 6.1 the relationship between the Lorenz curve and the Gini coefficient.

A Gini coefficient equal to zero corresponds to perfect equality (everyone has the same wealth), whereas a Gini coefficient of one corresponds to maximal inequality (the emperor owns all the wealth). The Gini coefficient satisfies four axioms that we will encounter again to justify a more general result in Section 6.2.

Scale invariance The Gini coefficient does not change after rescaling wealth (e.g., change of units/currency). For a distribution q and $\lambda > 0$, $G(q) = G(\lambda q)$.

Population invariance The Gini coefficient does not depend on the size of the population in the following sense. Supposing all agents of distribution q are cloned and inserted back into the original population to obtain q', we have $G(q) = G(q')$.

Anonymity The Gini coefficient does not use any other attribute of the population apart from its wealth distribution. For instance, it cannot weigh an agent more due to its ranking in the population.

Transfer principle If wealth (less than the difference[4]), is transferred from an agent with higher wealth to an agent with lower wealth, the resulting distribution is more equal (i.e., the Gini decreases).

Our motivating problem We consider how Nash equilibrium flow F affects the wealth of the population. In particular, we assume that the wealth of type x changes from $q(x)$ to $q(x) - \alpha \cdot \text{cost}^F(x)$ for some (intuitively small) $\alpha > 0$. We call the resulting wealth distribution $\hat{q}(x)$. Notice that, in general, $\hat{q}(x)$ may be different from $q(x) - \alpha \cdot \text{cost}^F(x)$, since the cost of F may rearrange the order of types (recall that distributions such as $q(x)$ are assumed to be nondecreasing). As we shall see in the inequity theorem proof of Section 6.2, this turns out to never be the case and moreover the inequality increases as a result.

In the following, we define q to be the wealth distribution of agents before playing the game. We let $q_0 = q - \alpha \cdot \text{cost}_0^F$ be the wealth distribution after playing the game *without tolls*, where $\text{cost}_0^F(x) = \sum_{e \in F(x)} \ell_e(c^F(e))$. The move from q to q_0 is defined as **the impact of travel**, the variation that is due only to the fact that players are engaged in a game. When tolls are levied, we have a second move, from q_0 to $\hat{q}$, defined as **the impact of tolls**.

6.2 The inequity theorem

Tolls can be used in congestion games so as to induce socially optimal flows (from the perspective of total cost) as Nash equilibrium (Cole, Dodis, and Roughgarden, 2003; Fleischer, Jain, and Mahdian, 2004). We next prove a general theorem showing that tolls always exacerbate societal inequality in symmetric routing games. So, in a sense to achieve optimality from the perspective of social welfare we have to pay a hidden cost in terms of fairness.

6.2.1 Formulation and proof of the theorem

Theorem 6.2.1 (Inequity theorem)**.** *In any Nash equilibrium of any symmetric congestion game with type-specific costs, any set of positive edge tolls τ_e increases the inequality of the population. More specifically,*

[4]If the wealth transfer is less than the difference of their wealth, the relative ordering of the wealth of the two agents does not change.

- *The impact of travel is zero: the Gini coefficient of the ex ante wealth distribution q is equal to the Gini coefficient of the toll-free wealth distribution q_0, $G(q_0) = G(q)$.*
- *The impact of tolls is nonnegative: the Gini coefficient of the ex ante wealth distribution is lower than (or equal to) the Gini coefficient of the ex post wealth distribution $\hat{q} = q - \alpha \cdot cost^F$, or $G(\hat{q}) \geq G(q) = G(q_0)$.*

Additionally, if the quantile distribution of wealth is increasing, the Gini coefficient increases strictly.

Canonical Nash equilibrium flows are piecewise constant

We show in this section that there exists $0 = a_1 \leq \cdots \leq a_{N+1} = 1$ real numbers such that if $\bar{F}$ is an equilibrium flow, then $\bar{F}(x) = \bar{F}(y)$ for $x, y \in (a_i, a_{i+1})$, $\forall i \in 1, \ldots, N$.

Cole, Dodis, and Roughgarden (2003) show that for routing games with type-specific cost functions, a Nash equilibrium exists that satisfies the interval definition above. The authors call such an equilibrium *canonical*, when in if $x \leq y$, then

$$\sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) \geq \sum_{e \in \bar{F}(y)} \ell_e(c^{\bar{F}}(e)) \text{ and } \sum_{e \in \bar{F}(x)} \tau_e \leq \sum_{e \in \bar{F}(y)} \tau_e \,.$$

Higher wealth agents thus incur a lower latency but higher tolls than lower wealth agents.

Canonical equilibrium flows exist even if the quantile distribution is discontinuous or not everywhere increasing.

Proof of the theorem

The proof of the impact of tolls is done in three steps. First, we show that if two wealth distributions with equal means cross at one point, one has a higher Gini coefficient than the other. This is equivalent to the transfer principle, or Pigou-Dalton principle of wealth inequality measures. Second, we show that when a distribution is obtained by decreasing proportionally less higher wealth than lower wealth—in other words, a regressive tax—then the resulting distribution has a higher Gini coefficient than the original one, i.e., is more unequal. Third, we show that under equilibrium in the game, players with higher wealth have a relatively lower path cost than players with lower wealth. Finally, Theorem 6.2.1 is obtained as a corollary of the three lemmas.

We present in Figure 6.2 a somewhat visual representation of our lemmas. By fixing the wealth of the median player to 1, we observe the two quantile distributions (before and after the game) cross at this exact point, leading to *ex post* polarised wealth by Lemma 6.2.1. This in turn implies that the Lorenz curve before the game dominates the Lorenz curve after the game.

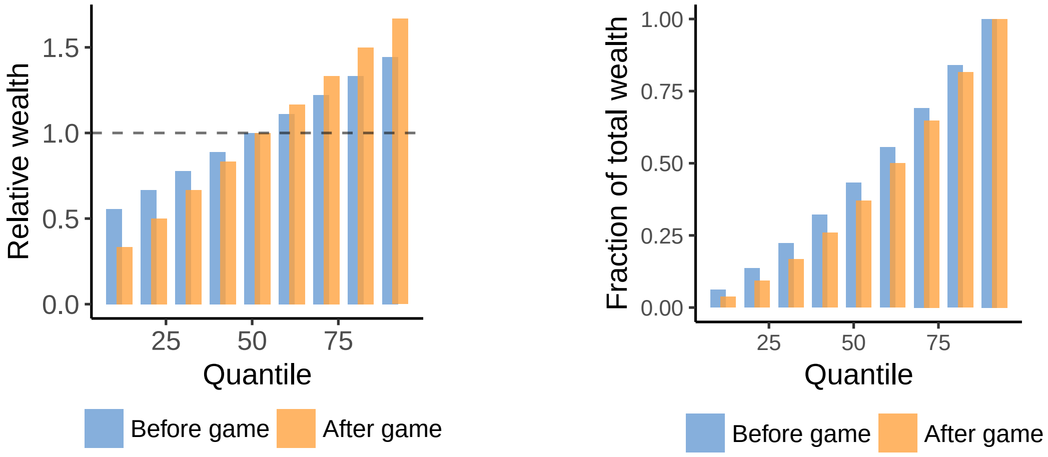

FIGURE 6.2: As the game is played, lower wealth agents lose relatively more of their wealth than high wealth players. **Left:** Quantile distribution of wealth before and after the game. **Right:** Lorenz distribution of wealth before and after the game, obtained at t by computing the total fraction of wealth owned by agents up to quantile t. The distribution after the game is Lorenz dominated by that before the game.

Lemma 6.2.1. *Suppose q and $\hat{q}$ are two wealth distributions (represented by their quantile functions) of equal means, i.e., $\mu = \int_0^1 q(x)dx = \int_0^1 \hat{q}(x)dx = \hat{\mu}$. If there exists x^* such that $\hat{q}(x) \leq q(x), \forall x \leq x^*$, and $\hat{q}(x) \geq q(x)$ otherwise, then $G(q) \leq G(\hat{q})$.*

Proof. We will show that $L(x) \geq \hat{L}(x)$, $\forall x$ where $L(x)$ is the Lorenz curve of q defined in (LC) and thus $G(q) \leq G(\hat{q})$. First note that for $x \leq x^*$,

$$\hat{L}(x) \leq L(x)$$

since $q(x) \geq \hat{q}(x)$.

For $x \geq x^*$, we have

$$\begin{aligned}
\hat{L}(x) \leq L(x) \iff & \int_0^{x^*} \hat{q}(y)dy + \int_{x^*}^{x} \hat{q}(y)dy \leq \int_0^{x^*} q(y)dy + \int_{x^*}^{x} q(y)dy \\
\iff & \int_0^{x^*} [q(y) - \hat{q}(y)]dy \geq \int_{x^*}^{x} [\hat{q}(y) - q(y)]dy
\end{aligned}$$

The last inequality is true since by $\mu = \hat{\mu}$ we get

$$\int_0^{x^*} [q(y) - \hat{q}(y)]dy = \int_{x^*}^{1} [\hat{q}(y) - q(y)]dy$$

$\square$

Lemma 6.2.2. *Suppose two wealth distributions (represented by their quantile functions) q and $\hat{q}$ are such that $\hat{q}(x) = \beta(x) \cdot q(x)$ and $1 \geq \beta(y) \geq \beta(z) > 0$ for $y \geq z$.*[5] *Then $G(q) \leq G(\hat{q})$.*

Proof. If we rescale distribution $\hat{q}$ to $q_1(x) = \frac{\mu}{\hat{\mu}}\hat{q}(x)$ the Gini remains invariant. It suffices to show that $G(q) \leq G(q_1)$. Since $\mu = \mu_1$, we can compare the two distributions using Lemma 6.2.1.

Introduce $\beta_1(x) = \frac{\mu}{\hat{\mu}}\beta(x)$, the transformation of income from q to q_1. $\beta_1(x)$ is a nondecreasing function of x. If $\beta_1(x) < 1$ for all x, then obviously we cannot have $\mu = \mu_1$ and the same holds if $\beta_1(x) > 1$ for all x. Hence there exists some x^* such that $\beta_1(x^*) = 1$. Moreover, by the monotonicity properties of β_1, the income order is preserved and x^* satisfies all the properties of Lemma 6.2.1. Thus, $G(q) \leq G(q_1) = G(\hat{q})$. □

Lemma 6.2.3. *Let $0 \leq x \leq y \leq 1$ and $\bar{F}$ be an equilibrium flow. Then $\frac{cost^{\bar{F}}(x)}{q(x)} \geq \frac{cost^{\bar{F}}(y)}{q(y)}$.*

Proof. Suppose $\bar{F}$ is an equilibrium flow. At equilibrium, the following holds:

$$\text{cost}^{\bar{F}}(x) \leq \sum_{e \in P} \ell_e(c^{\bar{F}}(e)) \cdot q(x) + \tau_e \,,\ \forall P \in \mathcal{P}\,.$$

Thus for $x \leq y$,

$$\begin{aligned}
\frac{\text{cost}^{\bar{F}}(y)}{q(y)} &\leq \sum_{e \in P} \ell_e(c^{\bar{F}}(e)) + \frac{\tau_e}{q(y)} \,,\ \forall P \in \mathcal{P} \\
&\leq \sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) + \frac{\tau_e}{q(y)} \\
&\leq \sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) + \frac{\tau_e}{q(x)} \\
&= \frac{\text{cost}^{\bar{F}}(x)}{q(x)}\,.
\end{aligned}$$

□

We can now prove the inequity theorem.

Inequity Theorem. The *ex post* wealth of x under flow F is given by

$$\begin{aligned}
\hat{q}(x) &= q(x) - \alpha \cdot \text{cost}^F(x) \\
&= q(x)\Big(1 - \alpha \cdot \frac{\text{cost}^F(x)}{q(x)}\Big)
\end{aligned}$$

[5] I.e., $\hat{q}$ is obtained from q by a transformation that reduces lower wealth relatively more than higher wealth. Income *order is preserved* and $\hat{\mu} \leq \mu$.

Thus, at equilibrium under flow $\bar{F}$, wealth is scaled by $\beta(x) = 1 - \alpha \cdot \frac{\text{cost}^F}{q(x)}$, where β is nondecreasing by Lemma 6.2.3. In Lemma 6.2.2, if q is our initial wealth distribution, *before the game is played*, then $\hat{q}$ is our wealth distribution *after the game is played*, and the β we have just defined satisfies the assumptions of the Lemma. The inequity theorem follows.

To show that the impact of travel is null, we consider the Nash equilibrium flow $\bar{F}$ of the game with costs

$$\text{cost}^{\bar{F}}(x) = \sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) \cdot q(x) = q(x) \sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) \,.$$

Akin to the Wardrop equilibrium, we can show that all agents incur the same latency $L = \sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e))$. Suppose $x \leq y$ and $\sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) < \sum_{e \in \bar{F}(y)} \ell_e(c^{\bar{F}}(e))$. Then by switching to $\bar{F}(x)$, y can decrease its cost $q(y) \sum_{e \in \bar{F}(y)} \ell_e(c^{\bar{F}}(e))$. The same argument holds if $\sum_{e \in \bar{F}(x)} \ell_e(c^{\bar{F}}(e)) > \sum_{e \in \bar{F}(y)} \ell_e(c^{\bar{F}}(e))$.

Since all players incur equal latency L, we have

$$q_0(x) = q(x) - \alpha \cdot q(x) \cdot L = q(x)(1 - \alpha \cdot L) \,.$$

By the invariance of Gini coefficient to scaling of wealth, $G(q_0) = G(q)$ and the impact of travel is null. $\square$

6.2.2 Strictness of the inequity theorem

The theorem also holds for the discontinuous case, which is typical to statistical analyses of inequality that group individuals under some average wealth. When the quantile function is not strictly increasing, it is possible that the inequality is left unchanged by the price mechanism.

In the Pigou example, one can take the following continuous wealth distribution $q(x) = 0.5$ for $x \leq 0.5$, $q(x) = -1 + 3x$ for $0.5 \leq x \leq 2/3$ and $q(x) = 1$ for $x \geq 2/3$, without any price to the variable cost link. All agents incur a cost equal to their wealth at the Nash equilibrium. We now set the price of the variable latency link to 0.75. Agents from quantile 0 to quantile $2/3$ occupy the constant latency link, and the $1/3$ mass of agents with wealth 1 distributes itself between the two links, with 75% of them using the variable latency link (since $q(x) \cdot \ell(z) + \tau = 1/4 + 3/4 = q(x)$). Again all agents incur a cost equal to their wealth, and thus the inequality is the same as before the price mechanism.

A nonincreasing but not everywhere decreasing edge cost function is unnatural: it would assign the same cost to two agents with different wealth, violating the concept of a decreasing tradeoff from money to time with wealth. The specific edge cost function under study, $\frac{\tau_e}{q} + \ell_e(z)$, is certainly decreasing in wealth. We then show that the second assumption to obtain a strict increase of inequality, a strictly increasing quantile function, is sufficient. In this case, agent costs are a strictly decreasing function of wealth, i.e., $\text{cost}^F(x) > \text{cost}^F(y)$ for $0 \leq x < y \leq 1$ in Lemma 6.2.3. The function β in Lemma 6.2.2 is thus strictly increasing, implying that all inequalities are strict in Lemma 6.2.1.

6.2.3 Extension to general inequality measures

Lemma 6.2.1 expresses that in symmetric games, the *ex ante* wealth distribution Lorenz dominates the *ex post* distribution, in the sense that for any quantile x, we have $L(x) \geq \hat{L}(x)$. Lorenz dominance, denoted by $\geq_L$ is a consistent ranking with all indices of inequality I that satisfy the four axioms given in Section 6.1 (population principle, anonymity, scale invariance and transfer principle) (Sen et al., 1997). Thus, if $q \geq_L \hat{q}$, then $I(q) \leq I(\hat{q})$.

Lorenz domination is only a partial order, e.g., if two Lorenz curves intersect, none Lorenz-dominates the other. However, in our setting, this turns out to never be the case. Thus a large family that includes the Gini, Theil (1967) and Atkinson (1970) indices can be used to measure the increase in inequality of Theorem 6.2.1.

Theorem 6.2.2. *For any wealth inequality measure satisfying the axioms of invariance to population scaling, anonymity, invariance to multiplicative scaling and the transfer principle, the inequity theorem holds and inequality increases as tolls are levied on the players.*

Remark 1 (Necessity and sufficiency of the four axioms). The proof given by Sen et al. (1997) makes clear that each of the four axioms is necessary to show that Lorenz dominance implies the consistent ordering on inequality measures. First, one wishes to compare two populations of *same size* and *same mean*. Lorenz ordering is then equivalent to the consistent ordering of social evaluation functions (SEF), a class of functions aggregating the utility of all agents in the population, under some concavity assumptions (Dasgupta, Sen, and Starrett, 1973).

- To compare populations q_1, q_2 of different sizes, say m and n respectively, one needs to have the population axiom allowing us to replicating the population while keeping the SEF constants. We can then compare the SEF at q_1^n and q_2^m, each population replicated respectively n and m times, which are both of the same size.

- If the two populations have different means, the invariance to rescaling is necessary to compare say q_1 with $\frac{\mu_1}{\mu_2} q_2$. The measure of inequality is unchanged from q_2 to $\frac{\mu_1}{\mu_2} q_2$ and both q_1 and $\frac{\mu_1}{\mu_2} q_2$ now have equal means.

- The anonymity axiom ensures that we can build the Lorenz curve by ordering members of the population according to their quantile. If not, different orderings of the population would yield different measures of inequality.

- Finally, the transfer axiom ensures that regressive transfers worsen the inequality. As seen in the proof of Lemma 6.2.2, the distribution after such transfers is Lorenz dominated by the original distribution. Failing that, we would not obtain a consistent ordering from Lorenz domination for other inequality measures. Additionally, all inequality measures which assign 0 to perfect equality (a property called *normalisation*) would give negative values after a regressive transfer, while the original equal distribution obviously Lorenz dominates the resulting distribution.

Although we do not use the reverse implication, it is also true that for any inequality measure satisfying all four axioms, there is a corresponding consistent ordering on Lorenz curves, showing sufficiency of the axioms.

6.3 The asymmetric case

6.3.1 Decomposition for asymmetric games

Does the inequity theorem hold for the asymmetric nonatomic routing games? In general, the answer is no, with a few caveats. First, asymmetric routing games, by their very nature, introduce populations that are not comparable: they have different sources and destinations as well as different available paths to reach one from the other. As shown in the remainder of this section, we cannot find a strict equivalent to the inequity theorem in the asymmetric case.

However, the inequity theorem in symmetric games carries over in the following way: *among each subpopulation, the inequality worsens.* As tolls are introduced, users in each commodity polarise such that higher incomes use tolled roads, following the construction presented in Section 6.2.

There is a clear parallel between asymmetric routing games and the measurement of inequality in other contexts, including in the rural/urban divide or the global inequality of wealth (Milanovic, 2016). In both cases, one would like to understand how inequality *within* each subpopulation (the rural and the urban populations, in the former; each individual nation, in the latter) is changed due to some economic activity. This alone would not paint the complete picture: one would also have to measure how the inequality *between* subpopulations has changed too.

Among the economic indices of inequality introduced previously, the Theil indices deal well with such a decomposition. The Theil T index is defined by

$$T_T(q) = \int_0^1 \frac{q(x)}{\mu} \ln \frac{q(x)}{\mu} dx$$

and the Theil L index by

$$T_L(q) = \int_0^1 \ln \frac{\mu}{q(x)} dx.$$

In the discrete case of a partition of N agents into M different subgroups, it is known that

$$T_T = \sum_{i=1}^{M} s_i T_i + \sum_{i=1}^{M} s_i \ln \frac{\overline{x_i}}{\mu}$$

where s_i is the share of total income possessed by group i, T_i is the Theil index T_T of group i, $\overline{x_i}$ is the average income of an agent in i and μ is the average income of all agents. The first sum gives us the measure of within-groups inequality, while the second is that of between-groups inequality.

The Gini coefficient does not admit such an additive decomposition, introducing instead an overlap term that does not allow the same intuitive reasoning on variations of inequality in the nonatomic asymmetric routing games. Indeed, by the inequity theorem in the symmetric case, we know for all relative income inequality measures, including the Theil indices, that within-groups inequality increases. Does it mean that the term $\sum_{i=1}^{M} s_i T_i$ always increases after the game? It is not the case for the Theil T index.

Indeed, though we know that $\hat{T}_i \geq T_i$ for all i, it is entirely possible that the shares of total income of each group also change in a way that $\sum_{i=1}^{M} \hat{s}_i \hat{T}_i \leq \sum_{i=1}^{M} s_i T_i$. The T index thus does not exactly reflect the property given by the inequity theorem.

Fortunately, the L index decomposes neatly into a weighted average of the subgroup's L indices with weights depending only on the population shares:

$$T_L = \sum_{i=1}^{M} \frac{N_i}{N} L_i + \sum_{i=1}^{M} \frac{N_i}{N} \ln \frac{N_i/N}{s_i}.$$

This implies that for the Theil L index, by the inequity theorem, the within-groups inequality is always non-decreasing. The inequality thus only decreases if the between-groups component is negative and of a greater magnitude than the within-groups component.

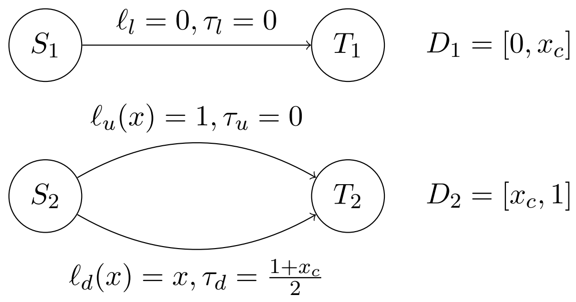

FIGURE 6.3: A multicommodity example: the interval D_1 of agents with income between 0 and x_c wishes to go from S_1 to T_1, while the interval D_2 goes from S_2 to T_2. Tolls and latencies are given along the edges.

TABLE 6.1: Within-groups, between-groups and total inequality differentials. The total inequality differential is equal to the sum of the between- and within-groups differentials, up to rounding produced for this table. We take $\alpha = 0.01$.

	Theil T index		Theil L index	
x_c	$\frac{1}{6}$	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{2}$
Within-groups differential ($\times 10^{-3}$)	3.67	0.88	4.05	0.31
Between-groups differential ($\times 10^{-3}$)	-0.86	-2.66	-2.26	-3.23
Total inequality differential ($\times 10^{-3}$)	2.81	-1.79	1.80	-2.91

6.3.2 Example of decomposition

In Figure 6.3, two populations are travelling along different networks. In a sense, this example is artificial as the two populations never share paths. It illustrates however that inequality can evolve in two different directions, depending on the within- and between-groups components. In this example, we use x_c as a parameter to control the size of the unaffected population on the upper network, and thus indirectly the importance of the between-groups inequality component: as the income of agents on the lower network decreases due to the cost of latency and tolls, their income gets closer to that of the agents on the upper network, and thus between-groups inequality is reduced.

The calculations for the lower network are similar to those done in Section 6.5 and are not repeated here. The Theil T indices are obtained for both populations' ex ante and ex post income distribution, denoted by T_i and $\hat{T}_i$ respectively for population i.

When $x_c = \frac{1}{6}$, a small portion of the population is on the upper network. The within-groups inequality overcomes the between-groups inequality and thus $\hat{T} > T$. On the other hand, when $x_c = \frac{1}{2}$, we have instead $\hat{T} < T$: the inequality increase in the lower network is compensated by the reduction in between-groups inequality. We give in Table 6.1 the computed numbers for the example.

6.4 The inequity index

The inequity theorem shows that under general conditions of the cost functions, the income inequality between agents increases after tolls are levied. In this section, we quantify this deterioration of equality by introducing a new metric. We have captured the importance of the game costs to the agents' income by a parameter $\alpha > 0$, intuitively small. The inequity (index) is defined as the derivative of the Gini coefficient as α goes to zero.

Definition 6.4.1. Let Γ be a nonatomic symmetric congestion game. Agents have an initial *ex ante* distribution $(q(x))_{x\in[0,1]}$ and incur a cost $\text{cost}^F(x)$ under flow F. Let $q_\alpha(x) = q(x) - \alpha \cdot \text{cost}^F(x)$ be the *ex post* income distribution for some $\alpha > 0$. The *inequity* of Γ is defined as

$$I(\Gamma) = \lim_{\alpha\to 0^+} \frac{G(q_\alpha) - G(q)}{\alpha}.$$

Note that this notion is well-defined. The Gini coefficient for distribution q_α is given by

$$G(q_\alpha) = 1 - 2\frac{\int_0^1 \int_0^x (q(t) - \alpha \cdot \text{cost}^F(t))dtdx}{\int_0^1 (q(x) - \alpha \cdot \text{cost}^F(x))dx} = 1 - 2\frac{\int_0^1 Q(x)dx - \alpha \int_0^1 \int_0^x \text{cost}^F(t)dtdx}{\mu - \alpha \cdot SC}$$

where μ is the total income of distribution q and SC is the social cost. This function is indeed differentiable with respect to α, provided the obvious requirement of $\mu > 0$ is satisfied.

6.4.1 Scale invariance of the inequity index

The inequity theorem implies that the inequity is always nonnegative. For the rest of the paper we will focus on the canonical cost functions (CAN). As a reminder, the cost of agent x in edge e is

$$q(x) \cdot \ell_e(c^F(e)) + \tau_e\,.$$

The canonical cost functions, besides having strong experimental justification (Abrantes and Wardman, 2011; Börjesson, Fosgerau, and Algers, 2012) provide also significant advantages in the theoretical study of inequity. Specifically, the inequity index is invariant under multiplicative scaling of the population wealth.

Theorem 6.4.1 (Robustness under scaling of income)**.** *Assume agent cost functions are in canonical form* (CAN) *in a game* Γ*. Then the inequity is scale invariant: If the wealth of each agent is scaled by a constant* $\lambda > 0$ *and optimal tolls are used in the resulting game* Γ_λ*, then* $I(\Gamma) = I(\Gamma_\lambda)$.

Proof. Optimal tolls τ_e^* in the original (unscaled) game Γ minimise the social cost as defined by the social planner, i.e., the sum of all player latencies. In the scaled version of the game, Γ_λ, the new optimal tolls $\hat{\tau}_e^*$ should be such that the resulting flow is identical to the minimising flow in the original game. In that case, the social optimum is realised for Γ_λ.

It is possible to show that this result holds for $\hat{\tau}_e^* = \lambda\tau_e^*$. Indeed, since

$$(\lambda q(x)) \cdot \ell_e(c^F(e)) + \lambda\tau_e = \lambda(q(x) \cdot \ell_e(c^F(e)) + \tau_e)$$

all costs are scaled by the same constant, and thus the strategic content of the game is unchanged (i.e., all players act as if the costs were the same as in Γ).

The Gini coefficient is scale invariant in the sense that if the wealth distribution q is multiplied by $\lambda > 0$, then $G(q) = G(\lambda q)$. In the canonical form (CAN), the *ex post* distribution is

$$\hat{q}(x) = q(x) - \alpha \cdot \sum_{e \in F(x)} \Big(q(x) \cdot \ell_e(c^F(e)) + \tau_e\Big)\,, \forall x\,.$$

If wealth is scaled by $\lambda > 0$ and optimal tolls are selected by the social planner, the new distribution is

$$\lambda q(x) - \alpha \cdot \sum_{e \in F(x)} \Big(\lambda q(x) \cdot \ell_e(c^F(e)) + \lambda\tau_e\Big) = \lambda\hat{q}(x)\,, \forall x\,.$$

for which the Gini coefficient is equal to $G(\hat{q})$. This further implies that the inequity of game Γ_λ is equal to that of Γ. $\square$

6.4.2 Robustness to no-regret learning

So far we have looked at the inequity index in the context of agents playing the Nash equilibrium of the routing game. However, it is possible to relax this assumption and let agents implement a no-regret strategy of their own.

Let $F_1, F_2, \ldots$ be a sequence of flows obtained from agents repeatedly playing the game. Agent x is implementing a no-regret algorithm if it has vanishing regret, i.e.,

$$R(T) = \frac{1}{T}\sum_{i=1}^{T} \mathrm{cost}^{F_i}(x) - \min_{P \in \mathcal{P}} \frac{1}{T}\sum_{i=1}^{T}\sum_{e \in P} \Big(q(x) \cdot \ell_e(c^{F_i}(e)) + \tau_e\Big) \to 0 \text{ as } T \to \infty\,.$$

We also call an ϵ-approximate Nash equilibrium a flow F_ϵ such that

$$\int_0^1 \mathrm{cost}^{F_\epsilon}(x)dx - \int_0^1 \min_{p \in \mathcal{P}} \sum_{e \in P} \Big(q(x) \cdot \ell_e(c^{F_\epsilon}(e)) + \tau_e\Big)dx \le \epsilon\,.$$

Following the results from Blum, Even-Dar, and Ligett (2006), we can show that under regret-minimising agents, the flow converges to that of an approximate equilibrium under the assumption of a finite number of wealth levels $w_1, \ldots, w_K$. This assumption is rather realistic since in practice there can only be a finite number of wealth levels. Additionally, any continuous distribution over wealth can be approximated to arbitrarily high accuracy by a distribution of finite but large enough support.

Theorem 6.4.2 (Robustness under no-regret learning)**.** *Given a finite number of income levels, the inequity index is uniquely defined under the assumption of no-regret learning agents. Specifically, if all agents follow a no-regret algorithm, we have*

$$\lim_{\alpha \to 0; \alpha > 0} \lim_{T \to \infty} \frac{\frac{1}{T} \sum_{t=1}^{T} G(\hat{q}^t) - G(q)}{\alpha} = I(\Gamma)$$

where $\hat{q}^t$ is the ex post wealth distribution of the t-th instance of the game.

Before the proof, we need the two following technical lemmas.

Lemma 6.4.1. *$\bar{F}$ is an equilibrium flow for $\Gamma = (G, q, cost)$ if and only if $\bar{F}$ is an equilibrium flow for $\Gamma_t = (G, q, cost_t)$, where*

$$cost_t^F(x) = \frac{cost^F(x)}{q(x)} .$$

$cost_t^F(x)$ is the perceived cost in terms of time to agent x under flow F.

Proof. Follows from (NE).

$$\begin{aligned} \text{cost}^F(x) &\leq \sum_{e \in P} \ell_e(c^F(e)) \cdot q(x) + \tau_e \,, \forall P \in \mathcal{P} \\ \Leftrightarrow \frac{\text{cost}^F(x)}{q(x)} &\leq \sum_{e \in P} \ell_e(c^F(e)) + \frac{\tau_e}{q(x)} \,, \forall P \in \mathcal{P} \end{aligned}$$

□

Lemma 6.4.2. *If F is a ϵ-NE of Γ_t, then F is a ϵ'-NE of Γ, where $\epsilon' = (1 - \sqrt{\epsilon})\sqrt{\epsilon} \cdot q_M + o(\epsilon)$ and $q_M = \sup_{x \in [0,1]} q(x) = q(1)$.*

Proof.

$$\begin{aligned} F \text{ is } \epsilon\text{-NE of } \Gamma_t \Rightarrow \int_0^1 \text{cost}_t^F(x) dx - \int_0^1 \min_{p \in \mathcal{P}} \sum_{e \in P} \left(\ell_e(c^F(e)) + \frac{\tau_e}{q(x)} \right) dx \leq \epsilon \\ \Rightarrow \text{cost}_t^F(x) - \min_{p \in \mathcal{P}} \sum_{e \in P} \left(\ell_e(c^F(e)) + \frac{\tau_e}{q(x)} \right) \leq \sqrt{\epsilon} \\ \text{for more than } (1 - \sqrt{\epsilon}) \text{ agents} \end{aligned}$$

$$\Rightarrow \text{cost}^F(x) - \min_{p\in\mathcal{P}} \sum_{e\in P} \Big(q(x)\cdot \ell_e(c^F(e)) + \tau_e\Big) \leq q(x)\sqrt{\epsilon} \leq q_M\sqrt{\epsilon}$$

$$\text{for more than } (1-\sqrt{\epsilon}) \text{ agents}$$

$$\Rightarrow \int_0^1 \text{cost}^F(x)dx - \int_0^1 \min_{p\in\mathcal{P}} \sum_{e\in P} \Big(q(x)\cdot \ell_e(c^F(e)) + \tau_e\Big)dx$$

$$\leq (1-\sqrt{\epsilon})\sqrt{\epsilon}\cdot q_M + o(\epsilon)$$

□

We can now prove the result for the stability of the inequity index to no-regret learning agents.

Proof. The proof consists of two steps. In the first step we will show that our symmetric game of type-specific costs $\Gamma = (G, q, \text{cost})$ reduces to an asymmetric congestion game $\hat{\Gamma}$. In the second step, we will apply results about the behaviour of no-regret dynamics in asymmetric congestion games from (Blum, Even-Dar, and Ligett, 2006) to prove the robustness of the inequity index.

Let $\Gamma_t = (G, q, \text{cost}_t)$ the game with cost equal to the perceived latency. For every edge e in Γ_t, construct the parallel edges $(\hat{e}_i)_{i=1}^K$ linking e to its original endpoint. The price of edge $\hat{e}_i$ is constant and equal to $\frac{\tau_e}{w_i}$. Now for each path $P \in \mathcal{P}$, the player of type i has an associated path $\hat{P} \in \hat{\mathcal{P}}$ that uses all the edges in P as well as $\hat{e}_i$. Since Γ_t and $\hat{\Gamma}$ are payoff-equivalent games, under the assumption that our latency functions have bounded slope we can use the results from Blum, Even-Dar, and Ligett (2006) to show that the inequity is stable under no-regret learning.

As long as latency functions are of bounded slope, $\text{cost}_t^F(x)$ is of bounded slope, since the income of players and tolls are bounded. No-regret algorithms will therefore approach an approximate ϵ-Nash equilibrium of Γ_t for some $\epsilon > 0$, and thus a ϵ'-Nash equilibrium of Γ where ϵ' is defined as in 6.4.2, sometimes takeb here to be a function of ϵ. In the following, requiring the regret to be under ϵ yields a $f(\epsilon)$-NE equilibrium of Γ_t in most time steps, which is thus a $g(\epsilon)$-NE equilibrium of Γ in most time steps, for $g = \epsilon' \circ f$.

Let α and some ϵ be fixed. There exists a time T_ϵ such that after T_ϵ, $R(T) \leq \epsilon$. By Blum, Even-Dar, and Ligett (2006), at most a fraction $K\epsilon^{\frac{1}{4}}$ of the first T_ϵ games is not a $g(\epsilon)$-NE, where g is a function that goes to zero when ϵ also goes to zero and K is a constant.[6] T_ϵ is a function in $O(\frac{1}{\epsilon^2})$ so as ϵ goes to zero we have that T_ϵ goes to ∞. Call A_ϵ the set of time periods where $\hat{q}^t$ is obtained from a $g(\epsilon)$-approximate NE and B_ϵ

[6] Precisely, there is a fraction $ms^{\frac{1}{4}}\epsilon^{\frac{1}{4}}$ of time periods where the flow is not an $\epsilon'(\epsilon+2\sqrt{s\epsilon n}+2m^{\frac{3}{4}}s^{\frac{1}{4}}\epsilon^{\frac{1}{4}})$-NE, where m is the number of edges, s is a bound on the slope of edge cost functions and n is the largest number of edges in a path.

where this is not the case. We have

$$\frac{1}{T_\epsilon}\sum_{t=1}^{T_\epsilon} G(\hat{q}^t) = \frac{1}{T_\epsilon}\sum_{t\in A_\epsilon} G(\hat{q}^t) + \frac{1}{T_\epsilon}\sum_{t\in B_\epsilon} G(\hat{q}^t) \tag{6.1}$$

$$\text{and } \frac{1}{T_\epsilon}\sum_{t\in B_\epsilon} G(\hat{q}^t) \leq \frac{1}{T_\epsilon} K\epsilon^{\frac{1}{4}} \cdot T_\epsilon = K\epsilon^{\frac{1}{4}} \tag{6.2}$$

where the last inequality holds due to the Gini coefficient being bounded by 0 and 1. Thus, the maximum distance between the Gini coefficient of $\hat{q}^t$ and that of the NE is 1, for the at most ϵT_ϵ time steps where we do not have a $g(\epsilon)$-NE.

It remains to prove that for the other time steps in A_ϵ, we are approaching the NE costs that give rise to $\hat{q}$ as ϵ goes to zero. Indeed, let F_ϵ be the flow corresponding to an ϵ-approximate NE. In congestion games, the vector of path costs $(C_p)_{p\in\mathcal{P}}$ realised at a NE is unique. Take a decreasing sequence of $(\epsilon_k)_k \to 0$ as $k \to \infty$ and flows F_{ϵ_k} that are ϵ_k-NE. Their associated cost vectors are $(C_p^{\epsilon_k})_p$. Suppose that as $\lim_{k\to\infty}(C_p^{\epsilon_k})_p \neq (C_p)_p$. By compactness, up to a subsequence, F_{ϵ_k} converges to some flow F that is a NE. But for this flow F the path costs are different from $(C_p)_p$. This is a contradiction.

On the other hand, the Gini coefficient is a continuous function of the agents' costs, so any sequence of ϵ-NE approximating a NE of the game will approach its Gini coefficient. Coming back to Equations (6.1) and (6.2), we know that

$$\frac{1}{T_\epsilon}\sum_{t=1}^{T_\epsilon} G(\hat{q}^t) \sim_{\epsilon\to 0} (1-\epsilon)G(\hat{q}) + O(\epsilon^{\frac{1}{4}})$$

and thus $\frac{1}{T}\sum_{t=1}^{T} G(\hat{q}^t) \to G(\hat{q})$ as $T \to \infty$. □

6.5 Inequality in the Pigou network

We illustrate the model with the Pigou network, where a mass 1 of agents travels between two nodes along two edges. The upper link has variable latency $\ell(z) = z^d$ and toll τ, while the lower link has constant latency equal to 1. The quantile function of wealth in the population is $q(x) = x$.

One can prove that for a fixed toll τ, at equilibrium, there exists a quantile x^* such that individuals above quantile x^* use the upper link, while individuals below quantile x^* use the lower link. Since agent x^* must be indifferent between the upper link and the lower link, for $d = 1$, we solve for x^*

$$1 = (1 - x^*) + \frac{\tau}{x^*} \Leftrightarrow x^* = \sqrt{t}.$$

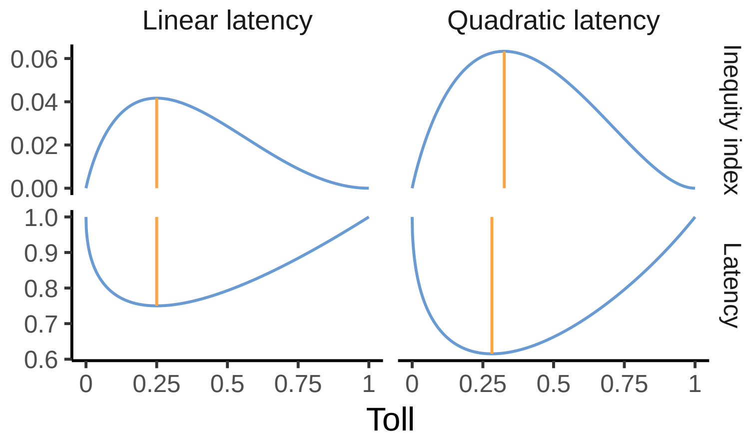

FIGURE 6.4: When the variable latency is set to $\ell(z) = z$, the optimal toll $\tau_1 = 1/4$ induces the highest inequity. This is not the case for $\ell(z) = z^2$, where the optimal toll is $\tau_2 = 0.281766$ but the maximiser of inequity is $\tau = 0.325487$. The inequity remains nonnegative for any toll and positive for flows routing a positive mass on the upper link.

The social planner who wishes to induce the optimal latency flow must set $\tau = 1/4$ so that $x^* = 1/2$. We let τ_d represent the optimal toll, i.e., the toll which induces the optimal flow. For $d = 1$, τ_d is also the maximiser of inequity: the marginal impact of the game on wealth and inequality is at its highest when latency is minimised!

However, as Figure 6.4 shows, this is no longer true for τ_2, when $\ell(z) = z^2$. There, the optimal toll does not induce the highest inequity. By virtue of Theorem 6.2.1, the inequity remains nonnegative for any toll and positive for flows routing a positive mass on the upper link.

6.6 Designing for equity

Inequality may increase after tolls are levied, due to the regressive nature of the tax. In this section, we investigate three mechanisms via which a system designer can either redistribute or optimise for equity. The simplest mechanism operates a direct transfer from the levied tolls to the population, for which we provide justification in Section 6.6.1. Recent proposals instead suggest to reinvest the tolls into funding for public transportation. Given the relatively inelastic travel times to congestion seen in Section 5.4 for this class of commuters, we adapt the Pigou example by equating the constant cost link as "public transportation" and improving the travel time on the link proportionally with the collected tolls (Section 6.6.2). Finally, we provide the model and result for a social planner who wishes to operate a tradeoff between efficiency and inequality by modifying its social cost function (Section 6.6.3).

6.6.1 Tax redistribution

We turn our attention to a redistributive model of tolls. In the simplest case, the collected tolls are given back uniformly to all agents. This progressive transfer reduces the Gini coefficient from that of $\hat{q}$, the *ex post* distribution, as we shall see.

Proposition 6.6.1. *Let $\bar{F}$ be an equilibrium flow on a routing game Γ with tolls τ and wealth q. Players incur a cost equal to cost$^{\bar{F}}(x) = \sum_{e \in \bar{F}(x)} q(x) \cdot \ell_e(c^{\bar{F}}(e)) + \tau_e$. The total collected tolls under a flow F are $T = \sum_{e \in E} F_e \tau_e$. Let cost$_r^F(x) =$ cost$^F(x) - T$ be updated costs for the agents, and Γ_r the game played with cost$_r$ cost functions. Then $\bar{F}$ is an equilibrium flow of Γ_r.*

Proof. By (NE), a flow $\bar{F}$ is at equilibrium if and only if for all agents $x \in [0,1]$, and for all paths $P \in \mathcal{P}$,

$$\text{cost}^{\bar{F}}(x) \leq \sum_{e \in P} q(x) \cdot \ell_e(c^{\bar{F}}(e)) + \tau_e$$

Since

$$\begin{aligned}
\text{cost}_r^{\bar{F}}(x) &= \sum_{e \in \bar{F}(x)} q(x) \cdot \ell_e(c^{\bar{F}}(e)) + \tau_e - \sum_{e \in E} \bar{F}_e \tau_e \\
&= \text{cost}^{\bar{F}}(x) - \sum_{e \in E} \bar{F}_e \tau_e \\
&\leq \sum_{e \in P} q(x) \cdot \ell_e(c^{\bar{F}}(e)) + \tau_e - \sum_{e \in E} \bar{F}_e \tau_e \,, \forall P \in \mathcal{P} \\
&= \text{cost}_r^{\bar{F},P}(x) \,, \forall P \in \mathcal{P}
\end{aligned}$$

$\bar{F}$ is an equilibrium flow of Γ_r. □

6.6.2 A subsidised Pigou network

Take the Pigou example, unit demand, flow $f \in [0,1]$ is routed on the variable cost link of x, flow $1-f$ is routed on the constant cost link of 1. Agent wealth follows the distribution $q(x) = x$.

Our data has shown that the Stress of Catastrophe of users in public transportation is very close to 1. In turn, this implies that their travel is close to being independent of any level of congestion. This is especially true as the use of buses in our dataset is much lower than that of train services. We can therefore equate the constant latency link with public transportation, normalised to 1.

Following recent proposals to channel the tax levied through tolls into improved public transport infrastructure, we modify the Pigou example. If a fraction f of users incurs a toll τ, then the total collected tax is τf. We can assume that a fraction $\gamma > 0$ of that toll is employed to reduce the constant cost on the lower link—perhaps via the

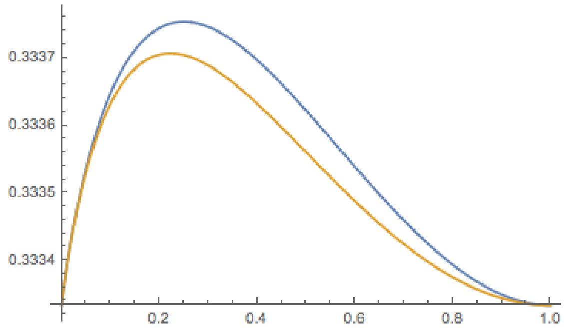

FIGURE 6.5: We plot the Gini coefficient as a function of the tolls in the Pigou example. The blue line is the Gini coefficient of the *ex post* wealth distribution in the non-subsidised Pigou, while the orange line is the Gini coefficient of the *ex post* wealth distribution in the subsidised game.

construction of an additional train line, bus service or cycling lane. The new cost of the lower link is then $1 - \gamma\tau f$.

The agent of quantile x^* must be indifferent between choosing to use the upper link or the lower link, and thus faces the equation

$$\begin{aligned} x(1-x)+\tau = x(1-\gamma\tau(1-x)) \Leftrightarrow \quad & (1+\gamma\tau)x^2 - \gamma\tau x - \tau = 0 \\ \Leftrightarrow \quad & x^* = \frac{\gamma\tau + \sqrt{(\gamma\tau)^2 + 4\tau(1+\gamma\tau)}}{2(1+\gamma\tau)} \end{aligned}$$

As γ or τ increases, the quantile of the "switching" agent also increases, such that agents are diverted from the variable cost link towards the constant cost link. The Gini coefficient of *ex post* wealth distribution of the modified Pigou is lower than that of the *ex post* wealth distribution of the classical Pigou. We show this result in Figure 6.5.

6.6.3 Balancing efficiency and inequity

We briefly present results obtained in Gemici et al. (2019). In a parallel links network, a planner can efficiently compute a precise tradeoff between efficiency and equity, as measured by the Gini coefficient. The model assumes agents equally distributed over a fixed number of wealth levels. Cost functions on the links are piecewise continuous and constant with increments at multiples of the fraction of agents on one wealth level. The planner seeks to minimise

$$\sum_e \ell_e(c^F(e)) + \lambda G(q_\alpha)$$

where q_α is the wealth distribution after the game and $\lambda > 0$ measures the importance of equity to the planner. A dynamic programming algorithm can be found to set tolls

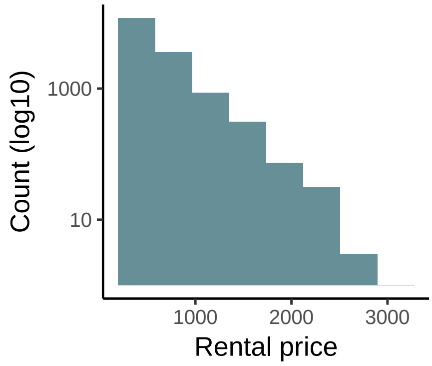

FIGURE 6.6: Distribution of rental prices

such that the objective is minimised. More details are presented in Gemici et al. (2019).

6.7 Wealth and mobility in Singapore

We use detailed transportation data gathered through Singapore's National Science Experiment (NSE) (presented in Section 5.1) to test how income inequality affects the distribution of transportation delays in a representative sample of students (Monnot et al., 2016; Monnot, Benita, and Piliouras, 2017; Wilhelm et al., 2016). Although Singapore is the third most densely populated country in the world, the modern infrastructure, cost of private cars, and significant tolls in Singapore minimise congestion on the roads. We examine whether this gain in efficiency incurs costs in terms of income inequality, as predicted by the theoretical results in this paper. The NSE dataset enables us to accurately split student trips in the morning—the time of the day when tolls are most onerous—by the transportation mode (bus, car, walk, and train) (Wilhelm et al., 2017). We then combine the travel data with a dataset on property prices to assess the relationship between income and the average duration and average distance of trips by transportation mode.

Rental prices are shown to correlate with income and thus provide here a reasonable proxy to the study of wealth distribution. Our rental prices dataset associates a mean rental price to each populated subzone of Singapore (219 in total). The distribution of rental prices among subjects in our dataset appears to follow a power law, with numerous students in areas with lower prices and few students in high prices areas (Figure 6.6).

Two trends are observed in the data. For subjects outside of the 0.1% top rental prices, the relationship between rental prices and average travel time is significantly negative: higher rental prices correlate strongly with lower travel time. On the contrary, for subjects in the top 0.1% rental prices, travel time is significantly higher than any

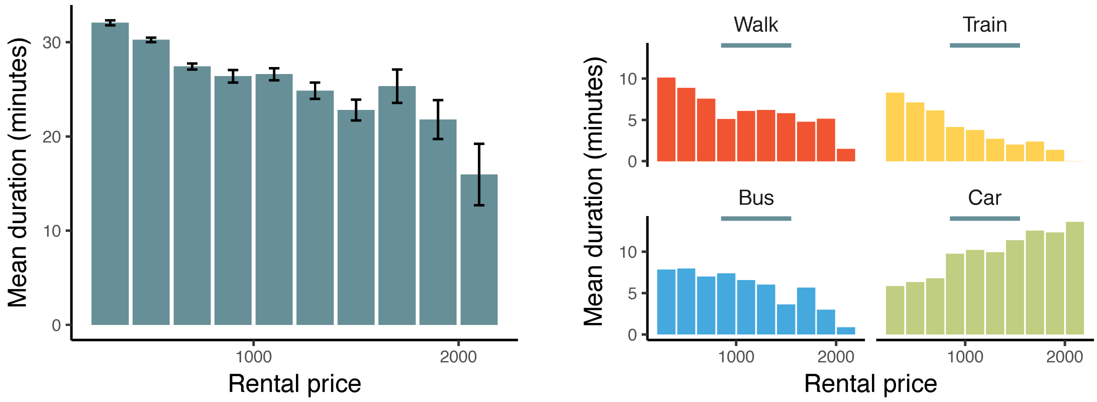

FIGURE 6.7: Average travel time of individual subjects ($n = 16563$), by rental prices (bin size = SGD200). Outliers of the top 0.1% subjects are omitted. Standard error is represented on the left histogram. The average commute duration per rental price group decreases with rental price. On the right, we present the duration is broken down by mode of transportation.

other group. The low number of samples for the latter however does not allow to conclude definitely that this is significant.

As can be seen in Figure 6.8, when one compares low-rental and high-rental groups, there is a notable increase in car usage and decrease in the use of walking and public transportation. Because cars are much faster than using bus and walking, the use of cars is associated with a sizeable difference in the average duration that subjects spend in traveling to school (Figure 6.7). Subjects in the lowest two rental groups spend on average 7 to 5 minutes more on their commute compared to middle-income groups. Thus, the Singaporean case—which is an ideal setting to examine the relationship between inequality and transportation delays—offers positive evidence on the inequity theorem.

6.8 Discussion

The question of efficiency and equality has a long history in economics. At the turn of the 20th century, Vilfredo Pareto and Arthur C. Pigou elucidated the concept of social welfare, with the first providing an efficiency-driven optimality criterion (the eponymous Pareto-efficiency) and the second a broad inquiry into the nature of "welfare" (Pigou, 1920). In both works, inequality and distributional fairness feature prominently alongside discussion of efficiency, embodied by the "national dividend" in Pigou.

The definition of the "social welfare function" (SWF), was not fully formalised until Samuelson (1948), after interpersonal comparisons of utility were deemed too problematic (Baujard, 2013). SWFs attempt to provide orderings on states of the society, which, hopefully, respect the individual orderings of its members. However, Kenneth J. Arrow

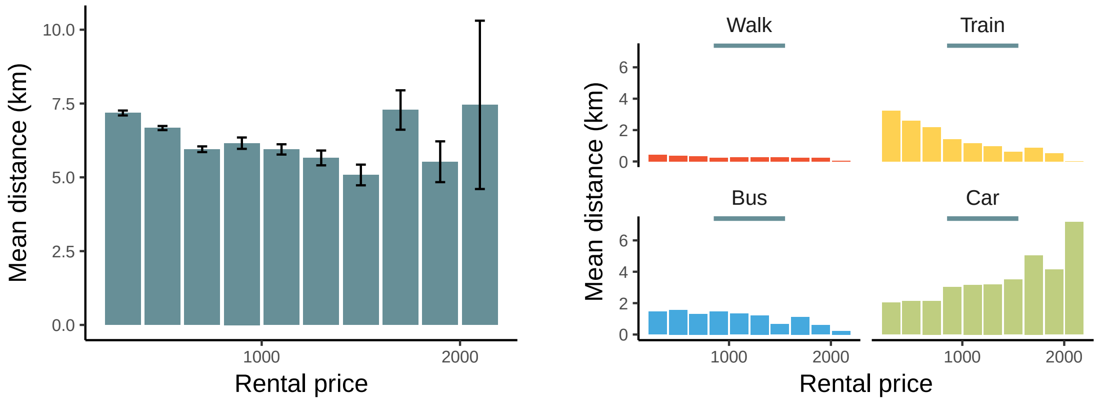

FIGURE 6.8: Average travel distance of individual subjects ($n = 16563$), by rental prices (bin size = SGD200). Outliers of the top 0.1% subjects are omitted. Standard error is represented on the left histogram. The average commute distance does not follow a regular trend with rental price, but car usage distance does increase while train distance decreases.

showed that under "reasonable" axioms,[7] in general such a consistent SWF cannot be found (Arrow, 1951). Yet, the separation between efficiency and inequality appears to have been sealed by the second theorem of welfare economics which justifies that any Pareto-efficient distribution can be reached by tuning the initial allocation of goods in an economy (Arrow and Debreu, 1954). The issue of distributional fairness is taken by the next generation of market designers as the responsibility of the policy-maker, with the economist merely supplying a set of designs with varying values of efficiency, fairness or complexity (Li, 2017), again in the spirit of the welfare theorems.

Appending some measure of inequality to the social welfare function—here, a simple version of SWF, where heterogeneity appears as a result of differing values for time among agents—, as we have done in this Chapter, may appear unnatural, but makes the question of distributional fairness hard to look away from. The tension between efficiency and equity is not clear-cut, as revealed by the examples presented throughout. Yet, the inequity index offers a quantitative approach to inequality, perhaps of interest in games other than routing.

[7]The independence of irrelevant alternatives being perhaps the most controversial (Sen, 2018).

Chapter 7

Coda

Decentralisation, in this thesis, is understood in contrast with the “master puppeteer”, a social planner anxious to optimise some function of the costs in the population. Its meaning thus implies agency of the players, whose rules of interactions are described by the incentives of the game and their own self interest. A key tension exists between efficiency, a state on first observation accessible only to the social planner, and decentralisation, as exemplified by the price of anarchy. And yet, mechanisms presented in this thesis, whether involving a central authority as in Chapter 3, lengthy communication and sophisticated punishments as in Chapter 4 or tolls as in Chapter 6, can coordinate agents to achieve if not the social optimum, a state strictly better than the worst equilibrium, in a decentralised fashion.

Does decentralisation always imply anarchy? These examples certainly argue that it does unless anarchy is kept in check by some mechanisms. We turn our attention to the rapidly growing field of economic systems governed by blockchains for a recent instance. “Decentralisation” in blockchains is often understood as the censorship- and fault-resistant properties of data replication supported by distributed consensus. In plain English, it is easy enough for one agent to operate a server receiving and broadcasting transactions, if a user trusts the operator to check for the correctness of the data (e.g., two transactions cannot contradict each other) or to maintain access to it. A malicious operator however might decide to corrupt (censor) incoming transactions for its own benefit, or simply shut down as a result of random failure (fault).

By replicating data across a multitude of agents operating their own server, one can prevent such faults, while opening another can of worms: who gets to decide which data is to be written into the log? The Nakamoto consensus (Nakamoto, 2008) created for Bitcoin relies on an elegant workaround to avoid confusion. A so-called miner must solve a hard puzzle by brute force, and once it manages to find such a solution, is allowed to disseminate a set of transactions (a block) to the other operators, who can check for the validity of the set against previous transactions. Operators thus compete to be the first to find such a block. The successful miner is rewarded with freshly minted currency native to the protocol, and fees from transactions included in its block. The consensus protocol does not rule out that two miners or more may each provide a distinct solution and transaction sets, but defines rules to cope with such situations, so that one winner emerges in what is called the "canonical" chain.

Decentralisation is thus imperative for the correctness of the data. Failing a large enough number of operators checking for the validity of the chain or competing in the mining race, a single operator can corrupt history or become sole producer of blocks, defeating the security desiderata of plurality. Since the launch of the first clients mining and producing blocks in 2009, focus was given to maintain decentralisation as the value (or at the very least, the amount of data) contained in the chain increased with each new block.

However, the puzzle is set up in a way that new blocks appear on average every 10 minutes, regardless of the intensity of the competition among miners. This led to an "arms race" of computing power, as the probability of successfully mining a new block increases with the rate at which a miner can produce candidate solutions for the puzzle. A consumer CPU may have been enough to engage in the race in the first few years, but with price increases of Bitcoin, it became profitable to invest large resources in first GPUs and later ASICs custom-built for mining. Amateur miners started to invest in mining pools (Liu et al., 2018; Leonardos, Leonardos, and Piliouras, 2019), where resources and rewards are shared, affording a more constant payoff rate, such that by now most of the new block creation is done by one of a handful of these pools.

In a paradoxical way, decentralisation as understood in this thesis has lead to centralisation as understood in the context of blockchains (Arnosti and Weinberg, 2018; Kwon et al., 2019). As agents followed their own incentives to profit from the system, the dynamics naturally led to increased concentration in the hands of a few. A useful reference here is Buterin (2017). The author notes the difference between *architectural* decentralisation, referring to the number of *physical computers* running the protocol, versus *political* decentralisation, or how many *virtual entities* actually control these computers. Hence, "[a]rchitectural centralization often leads to political centralization, though not necessarily". But in the case of the Bitcoin protocol and its miners, "can we really say that the uncoordinated choice model is realistic when 90% of the Bitcoin network's mining power is well-coordinated enough to show up together at the same conference?", in reference to a picture of a panel of seven people.

We could refer to political decentralisation in blockchains as *deconcentration*, to emphasise that the property we look for is high entropy in the block producers rather than agency (or anarchy) for the miners and participants to collude. The master puppeteer would, if it cared for deconcentration, limit the decentralisation of the agents and against their incentives have them mine such that all participate without a clear set of winners taking over. The arms race described previously finds striking parallels with the issue of congestion and tragedy of the commons explored in this work, with a difference.

The master puppeteer also cares to maximise architectural decentralisation, as a higher hash rate (i.e., total computer power devoted to solving the puzzle) in the network provides more security by making it harder to corrupt the *logically centralised* data (everyone has the same view of the transaction log), e.g., by launching a 51% attack.[1] Hence, given a fixed budget of hashing power, we define the *value of deconcentration* as the improvements to the distribution of mining power that a master puppeteer could achieve. If a cost function measures the entropy in the distribution, as the Gini did for instance in Chapter 6, the social optimum for a given hash rate is the entropy of the uniform distribution among miners. A more thorough model is left for future work.

In Papadimitriou (2001), our opening reference, the author asked "Of which game is TCP/IP congestion control the Nash equilibrium?". This question, along with its close relative, price of anarchy, spurred two decades of research into the equilibrium properties of large systems of decentralised agents.

Our previous chapters show that this investigation is far from over. The quantitative analysis of real systems from the viewpoint of PoA (Chapter 5 and 6) yields new research directions towards a better understanding of fairness in such systems. New systems such as blockchains, built on top of the very same Internet that launched the original line of questioning, are direct descendents of the mechanisms that PoA helped scrutinise under the algorithmic lens. Understanding the properties of these mechanisms in practice, how they can be measured, which are their distributional effects and clarifying the tenuous relationships between decentralisation, efficiency and inequality hopefully provides a solid footing for the questions to come.

[1] A 51% attack takes place when a single miner controls a strict majority of the computing power, in which case it can singlehandedly take control of the chain, since other miners cannot produce blocks faster.

Appendix A

Appendix to Chapter 3

A.1 Experiment sessions

We report in the following table the distribution of rounds among our 8 experimental sessions (including pilot).

TABLE A.1: Experiment details

Session date	**Number of rounds**	**Rounds dropped**
May 3rd, 2017	12	Pilot
June 6th, 2017	16	12
June 8th, 2017	16	2
June 15th, 2017	18	1
June 23rd, 2017	16	1
November 15th, 2017	15	1
March 21st, 2018	15	1
March 23rd, 2018	17	0
Total	125	12 (Pilot) + 18 (Dropped)

Appendix B

Appendix to Chapter 5

B.1 Comparison of two routes from sensor measurements

The distance $d_{a,b}$ is obtained with an estimation of the area enclosed by the polygon formed by the concatenation of a and b. More formally, since $a_1 = b_1 =$ Home and $a_n = b_m =$ School, we consider the area enclosed by the polygon $(a_1, a_2, \ldots, a_n, b_{m-1}, b_{m-2}, \ldots, b_1)$. The distance gets more precise as the number of data points logged along the trip increases, although it is still possible to achieve good results for very sparse trips.

We need a criterion to decide whether the previously obtained area is sufficiently small for the two sequences of coordinates to be considered consistent. To this end, we construct the outer contour of each sequence a and b, defined by a polygon containing the sequence of coordinates. Intuitively, we construct a band around the stream of locations, and the area of that band allows us to determine what constitutes an acceptable deviation. We show in Figure B.1b a representation of the outer contour for a trip with 4 points. Given d^o_a and d^o_b, respective areas of the outer contour for a and b, we use the following criterion to classify the routes a and b as consistent:

$$d_{a,b} < \frac{d^o_a + d^o_b}{2} \Rightarrow a \text{ and } b \text{ are consistent.}$$

The average is taken to ensure that both trips are considered equally in our criterion. Indeed, without the average, the algorithm may not be able to recognise a small deviation if only the outer contour of the shortest path appears on the right-hand side.

A parameter w controls the width of the band. If we set w to a value that is too large, we run the risk of incorrectly classifying different trips as consistent. On the other hand, a w that is too small may mark as non-consistent trips that make use of the same route. We have set the value of w by creating negative examples, comparing a trip with translated versions of the same sequence of coordinates. Visual analysis further confirmed the validity of its choice.

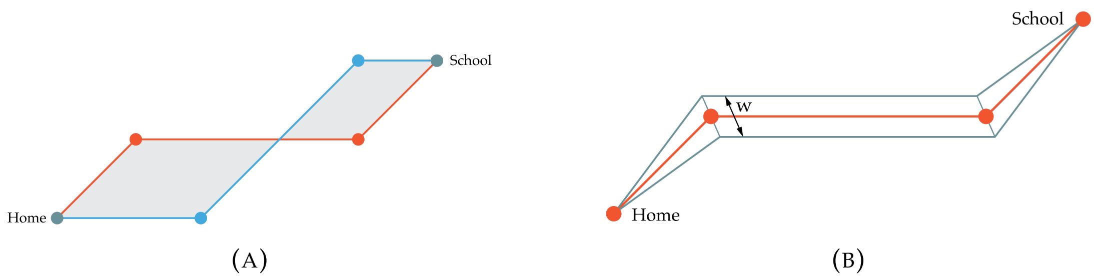

FIGURE B.1: **(A)** Two trips (red and blue) are plotted, with endpoints in dark blue. The distance between the trips is measured by the area of the darkened surface. **(B)** A four-point trip is plotted in red. The outer contour is obtained by fixing a band width w.

B.2 Scenarios and speed profiles on road segments

TABLE B.1: Scenarios configuration

	Scenario (speed km/h)				
Road type	Very Fast	Fast	Medium	Slow	Very Slow
Expressway	70	65	60	55	50
Semi Expressway	50	45	40	35	30
Arterial Road	40	35	30	25	20
Primary Road	30	25	20	15	10
Local Road	25	20	15	10	5
Expressway Slip Road	50	45	40	35	30
Slip Road	25	20	15	10	5

Bibliography

Abrantes, Pedro AL and Mark R Wardman (2011). "Meta-analysis of UK values of travel time: An update". In: *Transportation Research Part A: Policy and Practice* 45.1, pp. 1–17.

Acemoglu, Daron et al. (2015). "Democracy, redistribution, and inequality". In: *Handbook of income distribution*. Vol. 2. Elsevier, pp. 1885–1966.

Anshelevich, E. et al. (2004). "The price of stability for network design with fair cost allocation". In: *Foundations of Computer Science (FOCS)*. IEEE, pp. 295–304.

Arnosti, Nick and S Matthew Weinberg (2018). "Bitcoin: A Natural Oligopoly". In: *arXiv preprint arXiv:1811.08572*.

Arrow, Kenneth J and Gerard Debreu (1954). "Existence of an equilibrium for a competitive economy". In: *Econometrica: Journal of the Econometric Society*, pp. 265–290.

Arrow, Kenneth Joseph et al. (1951). "Social choice and individual values". In:

Ashlagi, Itai, Dov Monderer, and Moshe Tennenholtz (2008). "On the value of correlation". In: *Journal of Artificial Intelligence Research*, pp. 575–613.

Atkinson, Anthony B et al. (1970). "On the measurement of inequality". In: *Journal of economic theory* 2.3, pp. 244–263.

Aumann, Robert J (1974). "Subjectivity and correlation in randomized strategies". In: *Journal of mathematical Economics* 1.1, pp. 67–96.

Aumann, Robert J and Sergiu Hart (2003). "Long cheap talk". In: *Econometrica* 71.6, pp. 1619–1660.

Axhausen, Kay W et al. (2003). "80 weeks of GPS-traces: Approaches to enriching the trip information". In:

Baek, Jungho and Guankerwon Gweisah (2013). "Does income inequality harm the environment?: Empirical evidence from the United States". In: *Energy Policy* 62, pp. 1434–1437.

Bajari, Patrick, Han Hong, and Denis Nekipelov (2013). "Game theory and econometrics: A survey of some recent research". In: *Advances in Economics and Econometrics, 10th World Congress*. Vol. 3, pp. 3–52.

Balcan, Maria-Florina, Avrim Blum, and Yishay Mansour (2009). "Improved equilibria via public service advertising". In: *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, pp. 728–737.

Balcan, Maria-Florina et al. (2014). "Near-optimality in covering games by exposing global information". In: *ACM Transactions on Economics and Computation* 2.4, p. 13.

Barabási, Albert-László and Réka Albert (1999). "Emergence of scaling in random networks". In: *science* 286.5439, pp. 509–512.

Baujard, Antoinette (2013). "Welfare economics". In:

Bilò, Vittorio and Cosimo Vinci (2018). "The Price of Anarchy of Affine Congestion Games with Similar Strategies". In: pp. 48–59.

Blum, Avrim, Eyal Even-Dar, and Katrina Ligett (2006). "Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games". In: *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. ACM, pp. 45–52.

Börjesson, Maria, Mogens Fosgerau, and Staffan Algers (2012). "On the income elasticity of the value of travel time". In: *Transportation Research Part A: Policy and Practice* 46.2, pp. 368–377.

Boyce, James K (1994). "Inequality as a cause of environmental degradation". In: *Ecological Economics* 11.3, pp. 169–178.

Bradonjic, Milan et al. (2009). "On the price of mediation". In: *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, pp. 315–324.

Buterin, Vitalik (2017). *The Meaning of Decentralization*. Ed. by Medium.com. Medium.com [Online; posted 06-February-2017].

Cardenas, Juan-Camilo (2007). "Wealth inequality and overexploitation of the commons: field experiments in Colombia". In: *Inequality, cooperation, and environmental sustainability*, pp. 205–233.

Cerf, Vinton G. and Robert E. Kahn (1974). "A Protocol for Packet Network Intercommunication". In:

Cesa-Bianchi, Nicolo and Gabor Lugosi (2006). *Prediction, learning, and games*. Cambridge university press.

Christodoulou, George and Elias Koutsoupias (2005). "The price of anarchy of finite congestion games". In: *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, pp. 67–73.

Çolak, Serdar, Antonio Lima, and Marta C González (2016). "Understanding congested travel in urban areas". In: *Nature communications* 7, p. 10793.

Cole, Richard, Yevgeniy Dodis, and Tim Roughgarden (2003). "Pricing network edges for heterogeneous selfish users". In: *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*. ACM, pp. 521–530.

Colini-Baldeschi, Riccardo et al. (2017). "The asymptotic behavior of the price of anarchy". In: *International Conference on Web and Internet Economics*. Springer, pp. 133–145.

Cook, Stephen A (1971). "The complexity of theorem-proving procedures". In: *Proceedings of the third annual ACM symposium on Theory of computing*. ACM, pp. 151–158.

Dasgupta, Partha, Amartya Sen, David Starrett, et al. (1973). "Notes on the measurement of inequality". In: *Journal of economic theory* 6.2, pp. 180–187.

Daskalakis, Constantinos, Paul W. Goldberg, and Christos H. Papadimitriou (2006). "The complexity of computing a Nash equilibrium". In: ACM Press, pp. 71–78.

Du, Jianhe and Lisa Aultman-Hall (2007). "Increasing the accuracy of trip rate information from passive multi-day GPS travel datasets: Automatic trip end identification issues". In: *Transportation Research Part A: Policy and Practice* 41.3, pp. 220–232.

Émile Borel, M (1909). "Les probabilités dénombrables et leurs applications arithmétiques". In: *Rendiconti del Circolo Matematico di Palermo (1884-1940)* 27.1, pp. 247–271.

Erdős, P and A Rényi (1959). "On random graphs". In: *Publicationes Mathematicae Debrecen* 6, pp. 290–297.

Everett, M. G. and S. P. Borgatti (1999). "The centrality of groups and classes". en. In: *The Journal of Mathematical Sociology* 23.3, pp. 181–201. ISSN: 0022-250X, 1545-5874. DOI: 10.1080/0022250X.1999.9990219. URL: http://www.tandfonline.com/doi/abs/10.1080/0022250X.1999.9990219 (visited on 06/18/2018).

Feldman, Michal et al. (2016). "The price of anarchy in large games". In: *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, pp. 963–976.

Fleischer, Lisa, Kamal Jain, and Mohammad Mahdian (2004). "Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games". In: *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*. IEEE, pp. 277–285.

Garey, Michael R and David S Johnson (2002). *Computers and intractability*. Vol. 29. wh freeman New York.

Garey, Michael R, David S Johnson, and Larry Stockmeyer (1974). "Some simplified NP-complete problems". In: *Proceedings of the sixth annual ACM symposium on Theory of computing*. ACM, pp. 47–63.

Gemici, Kurtulus et al. (2019). "Wealth Inequality and the Price of Anarchy". In: *36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Gini, Corrado (1921). "Measurement of inequality of incomes". In: *The Economic Journal* 31.121, pp. 124–126.

Giridharadas, Anand (2018). *Winners take all: The elite charade of changing the world*. Knopf.

Hart, Sergiu and Yishay Mansour (2007). “The communication complexity of uncoupled nash equilibrium procedures”. In: *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, pp. 345–353.

Hayek, Friedrich August (1945). “The use of knowledge in society”. In: *The American economic review* 35.4, pp. 519–530.

Hoeffding, Wassily (1963). “Probability inequalities for sums of bounded random variables”. In: *Journal of the American statistical association* 58.301, pp. 13–30.

Holland, Paul W, Kathryn Blackmond Laskey, and Samuel Leinhardt (1983). “Stochastic blockmodels: First steps”. In: *Social networks* 5.2, pp. 109–137.

Jacobs, Jane (1961). “The Death and Life of Great American Cities”. In:

Jariyasunant, Jerald, Raja Sengupta, and Joan Walker (2012). “Overcoming battery life problems of smartphones when creating automated travel diaries”. In: *proceedings of the 13th International Conference on Travel Behavior Research*.

Judd, Stephen, Michael Kearns, and Yevgeniy Vorobeychik (2010). “Behavioral dynamics and influence in networked coloring and consensus”. In: *Proceedings of the National Academy of Sciences* 107.34, pp. 14978–14982.

Jun, Jungwook, Randall Guensler, and Jennifer Ogle (2006). “Smoothing methods to minimize impact of Global Positioning System random error on travel distance, speed, and acceleration profile estimates”. In: *Transportation Research Record: Journal of the Transportation Research Board* 1972, pp. 141–150.

Kearns, Michael, Siddharth Suri, and Nick Montfort (2006). “An experimental study of the coloring problem on human subject networks”. In: *Science* 313.5788, pp. 824–827.

Koutsoupias, Elias and Christos H. Papadimitriou (1999). “Worst-case Equilibria”. In: *STACS*, pp. 404–413.

Kumar, Sumeet et al. (2013). “Integrating on-board diagnostics speed data with sparse GPS measurements for vehicle trajectory estimation”. In: *SICE Annual Conference (SICE), 2013 Proceedings of*. IEEE, pp. 2302–2308.

Kwon, Yujin et al. (2019). “Impossibility of Full Decentralization in Permissionless Blockchains”. In: *arXiv e-prints*, arXiv:1905.05158, arXiv:1905.05158. arXiv: 1905.05158 [cs.CR].

Leonardos, Nikos, Stefanos Leonardos, and Georgios Piliouras (2019). “Oceanic Games: Centralization Risks and Incentives in Blockchain Mining”. In: *arXiv preprint arXiv:1904.02368*.

Levinson, David (2010). “Equity effects of road pricing: A review”. In: *Transport Reviews* 30.1, pp. 33–57.

Li, Shengwu (2017). “Ethics and market design”. In: *Oxford Review of Economic Policy* 33.4, pp. 705–720.

Littlestone, Nick and Manfred K Warmuth (1994). “The weighted majority algorithm”. In: *Information and computation* 108.2, pp. 212–261.

Liu, Xiaojun et al. (2018). "Evolutionary game for mining pool selection in blockchain networks". In: *IEEE Wireless Communications Letters* 7.5, pp. 760–763.

Milanovic, Branko (2016). *Global inequality: A new approach for the age of globalization*. Harvard University Press.

Mirowski, Philip and Edward Nik-Khah (2018). *The Knowledge We Have Lost in Information: The History of Information in Modern Economics*.

Monderer, D. and L. S. Shapley (1996). "Potential Games". In: *Games and Economic Behavior*, pp. 124–143.

Monnot, Barnabé, Francisco Benita, and Georgios Piliouras (2017). "Routing Games in the Wild: Efficiency, Equilibration and Regret". In: *International Conference on Web and Internet Economics*. Springer, pp. 340–353.

Monnot, Barnabé et al. (2016). "Inferring Activities and Optimal Trips: Lessons From Singapore's National Science Experiment". In: *Complex Systems Design & Management Asia*. Ed. by Michael-Alexandre Cardin et al. Springer, pp. 247–264.

Nakamoto, Satoshi et al. (2008). "Bitcoin: A peer-to-peer electronic cash system". In:

Nash, J. (1950). "Equilibrium points in n-person games". In: *Proceedings of the National Academy of Sciences*, pp. 48–49.

Nash, John (1951). "Non-cooperative games". In: *Annals of mathematics*, pp. 286–295.

Nekipelov, Denis, Vasilis Syrgkanis, and Eva Tardos (2015). "Econometrics for Learning Agents". In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. Portland, Oregon, USA: ACM, pp. 1–18. ISBN: 978-1-4503-3410-5. DOI: 10.1145/2764468.2764522. URL: http://doi.acm.org/10.1145/2764468.2764522.

Nikolova, Evdokia and Nicolas E Stier-Moses (2015). "The burden of risk aversion in mean-risk selfish routing". In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, pp. 489–506.

Nisan, Noam et al. (2007). *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press. ISBN: 0521872820.

Olszewski, Piotr and Litian Xie (2005). "Modelling the effects of road pricing on traffic in Singapore". In: *Transportation Research Part A: Policy and Practice* 39.7-9, pp. 755–772.

Ordóñez, Fernando and Nicolás E Stier-Moses (2010). "Wardrop equilibria with risk-averse users". In: *Transportation Science* 44.1, pp. 63–86.

Panageas, Ioannis and Georgios Piliouras (2016). "Average case performance of replicator dynamics in potential games via computing regions of attraction". In: *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, pp. 703–720.

Papadimitriou, Christos H (2001). "Algorithms, games, and the internet". In: *International Colloquium on Automata, Languages, and Programming*. Springer, pp. 1–3.

Pigou, Arthur Cecil (1920). *The economics of welfare*. Palgrave Macmillan.

Piketty, Thomas (2014). “Capital in the 21st Century”. In:

Public Roads, Bureau of (1964). “Traffic assignment manual”. In: *US Department of Commerce*.

Rawls, John (2009). *A theory of justice*. Harvard university press.

Rosenthal, R.W. (1973). “A class of games possessing pure-strategy Nash equilibria”. In: *International Journal of Game Theory* 2.1, pp. 65–67.

Roughgarden, Tim (2015). “Intrinsic robustness of the price of anarchy”. In: vol. 62. 5. ACM, p. 32.

Roughgarden, Tim and Éva Tardos (2002). “How bad is selfish routing?” In: *Journal of the ACM (JACM)* 49.2, pp. 236–259.

Samuelson, Paul Anthony (1948). “Foundations of economic analysis”. In:

Sankaran, Kartik et al. (2014). “Using mobile phone barometer for low-power transportation context detection”. In: *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*. ACM, pp. 191–205.

Schüssler, Nadine and Kay W Axhausen (2008). “Identifying trips and activities and their characteristics from GPS raw data without further information”. In:

Sen, Amartya (2018). *Collective choice and social welfare*. Harvard University Press.

Sen, Amartya et al. (1997). *On economic inequality*. Oxford University Press.

Shah, Devavrat and Jinwoo Shin (2010). “Dynamics in congestion games”. In: *ACM SIGMETRICS Performance Evaluation Review* 38.1, pp. 107–118.

Shapley, Lloyd (1964). “Some topics in two-person games”. In: *Advances in game theory* 52, pp. 1–29.

Shin, Dongyoun et al. (2015). “Urban sensing: Using smartphones for transportation mode classification”. In: *Computers, Environment and Urban Systems* 53, pp. 76–86.

Shirado, Hirokazu and Nicholas A Christakis (2017). “Locally noisy autonomous agents improve global human coordination in network experiments”. In: *Nature* 545.7654, p. 370.

Stiglitz, Joseph E (2012). *The price of inequality: How today’s divided society endangers our future*. WW Norton & Company.

Syrgkanis, Vasilis (2015). “Algorithmic game theory and econometrics”. In: *ACM SIGecom Exchanges* 14.1, pp. 105–108.

Theil, H (1967). *Economics and information theory.* English. OCLC: 710718565. Amsterdam: North-Holland.

Torras, Mariano and James K Boyce (1998). “Income, inequality, and pollution: a reassessment of the environmental Kuznets curve”. In: *Ecological economics* 25.2, pp. 147–160.

Tsui, Arvin Wen et al. (2010). “Accuracy performance analysis between war driving and war walking in metropolitan Wi-Fi localization”. In: *Mobile Computing, IEEE Transactions on* 9.11, pp. 1551–1562.

UNFCCC (2015). *The Paris Agreement*. Conference of the Parties Twenty-first session Paris, 30 November to 11 December 2015.

Watts, Duncan J and Steven H Strogatz (1998). "Collective dynamics of 'small-world' networks". In: *nature* 393.6684, p. 440.

West, Geoffrey (2017). *Scale: the universal laws of growth, innovation, sustainability, and the pace of life in organisms, cities, economies, and companies*. Penguin.

Wilhelm, Erik et al. (2016). "Sensg: Large-scale deployment of wearable sensors for trip and transport mode logging". In: *Transportation Research Board 95th Annual Meeting*. 16-4353.

Wilhelm, Erik et al. (2017). "Evaluation of Transport Mode Using Wearable Sensor Data from Thousands of Students". In: *Proceedings of the Transportation Research Board 96th Annual Meeting*. Transportation Research Board, pp. 1–18.

Wu, Zijun, Rolf H Möhring, and Dachuan Xu (2018). "Selfishness need not be bad: a general proof". In: *arXiv preprint arXiv:1805.07762*.

Young, H.P. (2004). *Strategic Learning and Its Limits*. Arne Ryde memorial lectures. Oxford University Press. URL: https://books.google.fr/books?id=3oUBoQEACAAJ.

Zhang, Jing et al. (2018). "The price of anarchy in transportation networks: Data-driven evaluation and reduction strategies". In: *Proceedings of the IEEE* 106.4, pp. 538–553.

Zhu, Qiuhui et al. (2016). "Identifying Transportation Modes from Raw GPS Data". In: *Social Computing: Second International Conference of Young Computer Scientists, Engineers and Educators, ICYCSEE 2016*. Ed. by Wanxiang Che et al. Singapore: Springer, pp. 395–409.