

Wealth Inequality and the Price of Anarchy

Kurtuluş Gemici

Department of Sociology, National University of Singapore, Singapore
kgemici@nus.edu.sg

Elias Koutsoupias

Department of Computer Science, University of Oxford, United Kingdom
elias@cs.ox.ac.uk

Barnabé Monnot

Engineering Systems & Design, Singapore University of Technology and Design, Singapore
monnot_barnabe@mymail.sutd.edu.sg

Christos H. Papadimitriou

Department of Computer Science, Columbia University, United States of America
christos@cs.columbia.edu

Georgios Piliouras

Engineering Systems & Design, Singapore University of Technology and Design, Singapore
georgios@sutd.edu.sg

Abstract

The price of anarchy quantifies the degradation of social welfare in games due to the lack of a centralized authority that can enforce the optimal outcome. It is known that, in certain games, such effects can be ameliorated via tolls or taxes. This leads to a natural, but largely unexplored, question: what is the effect of such transfers on social inequality? We study this question in nonatomic congestion games, arguably one of the most thoroughly studied settings from the perspective of the price of anarchy. We introduce a new model that incorporates the income distribution of the population and captures the income elasticity of travel time (i.e., how does loss of time translate to lost income). This allows us to argue about the equality of wealth distribution both before and after employing a mechanism. We establish that, under reasonable assumptions, tolls always increase inequality in symmetric congestion games under any reasonable metric of inequality such as the Gini index. We introduce the inequity index, a novel measure for quantifying the magnitude of these forces towards a more unbalanced wealth distribution and show it has good normative properties (robustness to scaling of income, no-regret learning). We analyze inequity both in theoretical settings (Pigou's network under various wealth distributions) as well as experimental ones (based on a large scale field experiment in Singapore). Finally, we provide an algorithm for computing optimal tolls for any point of the trade-off of relative importance of efficiency and equality. We conclude with a discussion of our findings in the context of theories of justice as developed in contemporary social sciences and present several directions for future research.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases congestion games, inequality

Digital Object Identifier 10.4230/LIPIcs.STACS.2019.31

Acknowledgements Kurtuluş Gemici acknowledges NUS Strategic Research Grant (WBS: R-109-000-183-646) awarded to Global Production Networks Centre (GPN@NUS). Elias Koutsoupias acknowledges ERC Advanced Grant 321171 (ALGAME). Barnabé Monnot acknowledges the SUTD Presidential Graduate Fellowship. Christos Papadimitriou acknowledges NSF grant 1408635 “Algorithmic Explorations of Networks, Markets, Evolution, and the Brain”. Georgios Piliouras acknowledges SUTD grant SRG ESD 2015 097, MOE AcRF Tier 2 Grant 2016-T2-1-170, NRF grant NRF2016NCR-NCR002-028 and a NRF fellowship. Barnabé Monnot and Georgios Piliouras would like to thank the other members of the National Science Experiment team at SUTD: Garvit Bansal, Francisco Benita, Sarah Nadiawati, Hugh Tay Keng Liang, Nils Ole Tippenhauer, Bige Tunçer,



© Kurtuluş Gemici, Elias Koutsoupias, Barnabé Monnot, Christos H. Papadimitriou, and Georgios Piliouras;
licensed under Creative Commons License CC-BY

36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019).

Editors: Rolf Niedermeier and Christophe Paul; Article No. 31; pp. 31:1–31:16

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Darshan Virupashka, Erik Wilhelm and Yuren Zhou. The National Science Experiment is supported by the Singapore National Research Foundation (NRF), Grant RGNRF1402.

1 Introduction

Inequality in wealth and income have been rampant worldwide in the past four decades [32,39], considered by many the scourge of modern societies. Economic analysis, on the other hand, traditionally focuses on efficiency, that is to say, Pareto optimality of the allocation. Whether, and to what extent, efficiency and equality are at loggerheads has been debated in economics, and the verdict appears to depend on context and assumptions.

Modern societies also give rise to a plethora of strategic scenarios, in which the behavior of one agent affects the others, and the outcome of which ultimately affects the agents' overall well-being. In game theory, we study the inefficiency of these strategic situations through the so-called *price of anarchy*, the relative efficiency of the game's Nash equilibria over the social optimum [28]. For congestion games in particular, it is known that the price of anarchy can be combatted through the introduction of *tolls* which enforce the optimal outcome as equilibrium, see [15,18] among an extensive literature. However, the effect that tolls may have on the level of inequality in the society does not appear to have been addressed in the literature.

The present paper is a first attempt to articulate and study this issue. We consider games (here only congestion games) in which the agents' utility and behavior depend explicitly on their income or wealth, and study the effect the game's equilibria have on inequality.

Example: Transportation in Singapore, seen as a congestion game with tolls, has a price of anarchy that is close to one [29]. The main arteries are almost never clogged, and public transportation is accessible and runs smoothly. This is the result of bold policy decisions: car ownership in Singapore is significantly taxed, and dynamically adaptive tolls are in place. Interestingly, transportation delays seem to be a decreasing function of income (see Section 7 on data). This is no accident: In this paper we show that there is an inherent tension between efficiency and equality in the context of congestion games.

We are interested in the ways in which optimal (or more generally efficiency-enhancing) mechanisms affect inequality. Inequality is measured in many ways, but perhaps most often through what is known as the *Gini coefficient (or Gini index)*. Informally (see Section 3 for the definition), the Gini coefficient of a distribution of income or wealth is *twice* the area between the 45° line and the normalized convex cumulative wealth/income curve (see Figure 1). That is, we compute the cumulative income/wealth $Q(y)$ of the lowest y fraction of the population for all $0 \leq y \leq 1$, we normalize it so that $Q(1) = 1$, and then we integrate $y - Q(y)$ from 0 to 1. At total equality the Gini index is zero, while at total inequality (i.e., when the emperor owns everything) it is one. In 2015, the income Gini in OECD countries ranged from the .20s (Northern Europe) to the .40s and .50s (USA and East Asia).

Our contributions. We study nonatomic congestion games with tolls, where we introduce a new model that incorporates the wealth distribution of the population and captures the income elasticity of travel time (i.e., how loss of time translates to lost income). This allows us to argue about the equality of wealth distribution both before and after participating in a mechanism (with or without tolls). The basics of our modeling are thus: We consider a

continuum of agents, each agent of a type $x > 0$ standing for their income.¹ We assume that the distribution of types is known. Suppose these agents engage in a game Γ and that, at equilibrium, type x receives a cost c_x . This cost is expressed in the same units as income, dollars, say; after incorporating the losses due to time spent in traffic in dollars as well as any possible costs due to tolls/taxes. As a result, the agent's total wealth becomes $x' = x - \alpha c_x$, where α is a small constant standing for the importance of the game under consideration to an individual's well-being. In Section 4 we establish a broad qualitative result, the Inequity Theorem (Theorem 2), showing that tolls always increase inequality in symmetric congestion games under the most classic inequality measure, the Gini coefficient. In fact, participating in a toll-free symmetric congestion game has no impact on the Gini, whereas optimal tolls on the other hand have a negative impact on the Gini. Theorem 6 broadly expands Theorem 2 to *any* inequality measure that satisfies four fundamental axioms: *invariance to population scaling, anonymity, invariance to income scaling and the transfer principle* (see Section 3). These measures include, besides Gini, some of the most widely employed indices, such as [40] or [3].

At a technical level, the proof of the Inequity Theorem combines game-theoretic properties of congestion games with tolls and the axioms of inequality measures. In order to argue that the Gini of the final income distribution is worse than that of the original, it suffices to argue that the Lorenz curve of the original distribution (see Figure 1) dominates the latter. Lorenz curve domination is established via the combination of Lemmas 3, 4, 5, implying Theorem 2. In fact, Lorenz curve domination suffices to argue something stronger. Any inequality measure satisfying the four axioms is also consistent with the Lorenz domination order, yielding Theorem 6.

In Section 5 we introduce the inequity index, a novel measure for quantifying the magnitude of these forces towards a more unbalanced wealth distribution. Let q be the initial income distribution of the population of agents under consideration, let $G(q)$ be its Gini coefficient, and suppose that \hat{q} is the distribution of the income *after* each x becomes $x - \alpha \cdot c_x$ (that is, after the game has been played). We are interested in the way the game affects the Gini coefficient; we express this, informally, as the coefficient of α in $G(\hat{q}) - G(q)$, ignoring terms that are $o(\alpha)$; in other words, we are interested in the *derivative* of $G(q)$ with respect to α . We call this quantity the *inequity* of the game. We show that from a theoretical perspective it has attractive properties. Specifically it is robust to scaling of income (Theorem 9) and it remains unaffected if instead of immediate equilibration we assume that all agents apply regret-minimizing algorithms (Theorem 10).

We analyze inequity both in theoretical settings (Section 6) as well as experimental ones (Section 7). Specifically, these effects become apparent already in the well-trodden Pigou's network [31]. This network has two parallel links, one with constant delay function 1, and another with delay function x (that is, a delay proportional to the percentage of agents that take this option). Its price of anarchy is $\frac{4}{3}$, and the inequity turns out to be zero. It is well-known that the price of anarchy, *in the case of equal incomes*, can be rendered to one by adding tolls, and it is not hard to see that the same can be done for any income distribution [15] – *but then the inequity becomes substantial*. If tolls decrease, we have a full-fledged trade-off between inequity and price of anarchy. In Theorem 11 we calculate the precise price of anarchy to inequity trade-off of any variant of Pigou's network with income distributions of the form y^β .

¹ Or wealth; we write “income” henceforth in this paper, but “wealth” would also be appropriate everywhere.

In Section 7 we perform data analytics on a dataset capturing the routing behavior of tens of thousands of Singaporean students. This dataset captures the movement of each individual at a high frequency (one new datapoint per individual every 13 seconds) and allows us to distinguish between different modes of transportation (walking, bus, train, car). We can pinpoint each individual's home location which allows us to compute estimates about their wealth. Given the level of data granularity, we can control for different parameters and identify a statistically significant increased commute time for the lower-income students, which corroborates our theoretical analysis. Interestingly, the Singapore case also points out some of the successful policies (e.g., polycentric urban development model) that can be implemented to alleviate the trade-off between efficiency and equality. Finally, we provide an algorithm for computing optimum tolls for any point of the trade-off of relative importance of efficiency and equality for symmetric networks on parallel links. We also present inequity results in asymmetric settings, which prompt several open questions. The full online text [22] provides any missing supplementary material and analysis.

2 Related work

Price of anarchy was introduced in congestion games [28], leading to a long sequence of influential papers in the area [13, 16, 20, 37, 38]. A similarly long line of research on tolls is existing in the AGT community, starting with [15, 18]. Recently, [25] have found efficient algorithms to compute tolls that minimize latency where not all edges can be tolled, putting their work close to the current situation in many cities. [6] have shown that taxes depending on the congestion level of a resource for weighted agents increase the efficiency of congestion games with polynomial latency functions. [5] have proved that without knowing the latency functions and using only tolls and an efficient number of queries to an oracle, target equilibrium flows can be reached. The data analytics, experimental part of our paper shows, perhaps unsurprisingly, that the use of public transportation plays a critical, but not well-understood, role in the functioning of a traffic network. [19] introduced a model of congestion games with buses, and hopefully more research will follow along these lines.

Given the proliferation of the usage of algorithms in all aspects of our lives (from suggesting Airbnb hosts to identifying convicts eligible for early parole), the theoretical computer science community has recently focused on understanding issues of fairness, equality and justice. Surprisingly, the intersection of price of anarchy, i.e., efficiency in games, and fairness has not been explored so far. On a related tangent, the issues of altruism and efficiency have been tackled, e.g. by [9] and [10]. [11, 12] have recently used the Gini coefficient over the probabilities of the agent winning probabilities as an inequality measure of different mechanisms and design mechanisms with such good properties. Although syntactically similar, these works do not model wealth distributions nor do they examine the differential effects of mechanisms to equality, which is our focus.

We continue the line of work of [29] where price of anarchy in congestion games is studied using field experiments with thousands of participants. In this paper, our data analytics corroborate our theoretical insights and give rise to novel questions for future research.

3 Model Description

We describe a game-theoretic model where a continuum of agents participates in a traffic congestion game with tolls. The total disutility for each agent depends both on their traffic-induced latency as well as on the tolls, whose effects are experienced differentially based on each agent's income level.

Congestion game. A symmetric congestion game with type-specific costs consists of a finite set E of edges, and a finite subset of 2^E called the set of paths \mathcal{P} , common to all types. We shall only deal with *network* congestion games, where the set of paths consists of all possible paths between two nodes s and t in a graph with edge set E .

Income. We have a continuum of types which lie in $[0, 1]$. Type x has income $q(x)$, where q is the *quantile function* of the income 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 income as random variable, then q expresses the inverse of its cumulative distribution function.

Flow. A flow $F : [0, 1] \rightarrow \mathcal{P}$ is a mapping from types to paths; we shall only need to consider *finitary* 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 $a_0 = 0 < a_1 < a_2 < \dots < a_k = 1$ such that $F(b) = F(c)$ for all i and $b, c \in (a_i, a_{i+1}]$

Edge cost. Our main result, the Inequity Theorem, holds under general conditions on the edge cost functions. For simplicity of exposition, we look at a specific case that has natural properties and leave to the full online text [22] a discussion on more general results for larger classes of edge cost functions. Each agent x using edge e experiences the edge cost $f_e(q, z, \tau_e)$, where q is the agent's income, z is the level of congestion on edge e and τ_e is the fixed toll paid by the agent.

We are interested in the following edge cost function:

$$f_e(q, z, \tau_e) = \frac{\tau_e}{q} + \ell_e(z)$$

The *path cost* for P is $\sum_{e \in P} f_e(q, z, \tau_e)$.

There is an extensive discussion in the transportation literature of the true cost of transportation to the traveler and the value of time, see [2, 8] for some of the most recent papers, with dozens of references therein. This field has established and studied the *income elasticity* of the value of (travel) time (informally, the precise nature of the formula $\frac{\tau_e}{q}$ 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 $\frac{\tau_e}{q^\beta}$ where $\beta \leq 1$ is conventionally taken to be one, even though certain countries, such as the UK, use value 0.8.

Agent cost. 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. The *agent cost* under flow F to an agent of type x is some function of its income and path cost.

The model allows for some degree of flexibility when designing the overall cost of the agents. We focus our attention to the following agent cost:

$$\text{cost}^F(x) = q(x) \cdot \sum_{e \in F(x)} f_e(q(x), c^F(e), \tau_e) \quad (\text{AC})$$

Edge costs are scaled by the income of the agent and thus the agent cost is given in the units of the toll, i.e. money.

31:6 Wealth Inequality and the Price of Anarchy

For $f_e(q, z, \tau_e) = \frac{\tau_e}{q} + \ell_e(z)$, we have

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


We call this agent cost function *canonical* and show further that it is a natural choice with good properties (Section 5).

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

$$\text{cost}^F(x) \leq q(x) \cdot \sum_{e \in P} f_e(q(x), c^F(e), \tau_e) \quad (\text{NE})$$

that is, if no type x would be better off by deviating to another path $P \in \mathcal{P}$.

In the following, we define q to be the income distribution of agents before playing the game. With our definition of agent costs, one can study q_0 , the income distribution after playing the game *without tolls*, where $f_e(q, z, \tau_e) = \ell_e(z)$. The move from q to q_0 is defined as **the impact of travel**, the variation that is due only to the presence of a game. When tolls are levied, we have a second move, from q_0 to \hat{q} , defined as **the impact of tolls**. We will be mostly concerned with the latter impact.



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

Gini coefficient. The Gini coefficient [23] is a central measure of inequality.

► **Definition 1.** The Gini coefficient of income 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 income held by individuals under and at quantile x .

$$L(t) = \frac{1}{\mu} \int_0^t q(x)dy = \frac{1}{\mu} Q(t) \quad (\text{LC})$$

for $Q(t) = \int_0^t q(x)dx$, the cumulative income up to quantile t . We show in Figure 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 income), whereas a Gini coefficient of one corresponds to maximal inequality (e.g., one person has all the income). The Gini coefficient has several desirable properties such as:

■ **Scale independence.** The Gini coefficient does not change after rescaling incomes (e.g. change of units/currency).

- **Population independence.** It does not depend on the size of the population.
- **Anonymity.** It does not depend on the identity of the rich/poor individuals.
- **Transfer principle.** If income (less than the difference²) is transferred from a rich person to a poor person the resulting distribution is more equal (i.e., the Gini decreases).

Our motivating problem. We consider how Nash equilibrium flow F affects the incomes of the population. In particular, we assume that the income 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 income 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 4, this turns out to never be the case and moreover the inequality increases as a result.

4 The Negative Impact of Tolls on Inequality

4.1 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 [15, 18]. We next prove a general theorem showing that tolls always exacerbate societal inequality. So, in a sense to achieve optimality from the perspective of social welfare we have to pay a hidden cost in terms of fairness.

► **Theorem 2 (The 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,*

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

Additionally, if the quantile distribution of income is increasing and edge cost functions are decreasing in income, the Gini coefficient increases strictly.

Proof Sketch. First, we show that the impact of travel is null. In the toll-free version of the game, the Nash equilibrium is the usual Wardrop equilibrium and all agents incur the same cost C , regardless of their route choice. This implies that all agents lose the same share of income exactly equal to $\alpha \cdot C$. Since the Gini coefficient and all inequality measures we will be concerned with are scale invariant, the inequality is not affected by the impact of travel.

The proof of the impact of tolls is done in three steps. First, we show that if two income distributions with equal means cross at one point, one has a higher Gini coefficient than the other (Lemma 3). This is equivalent to the transfer principle, or Pigou-Dalton principle of income inequality measures. Second, we show that when a distribution is obtained by decreasing proportionally less the higher incomes than the lower incomes – in other words, a regressive tax – then the resulting distribution has a higher Gini coefficient than the original one, i.e., is more unequal (Lemma 4). Third, we show that under equilibrium in the game, players with higher incomes have a lower path cost than players with lower incomes (Lemma 5). Finally, Theorem 2 is obtained as a corollary of the three lemmas.

² If the income transfer is less than the difference of their incomes, the ordering of the wealth of the users does not change.

► **Lemma 3.** *Suppose q and \hat{q} are two income 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 $G(q) \leq G(\hat{q})$.*

► **Lemma 4.** *Suppose two income 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 $y \geq z$ ³ then $G(q) \leq G(\hat{q})$.*

► **Lemma 5.** *Let $0 \leq x \leq y \leq 1$ and F be an equilibrium flow. If agent costs are given by the path cost $\sum_{e \in F(x)} \frac{r_e}{q} + \ell_e(z)$ then $\text{cost}^F(x) \geq \text{cost}^F(y)$.*

The resulting income distribution in the game is given by $\hat{q}(x) = q(x) \cdot (1 - \alpha \cdot \text{cost}^F(x))$. At equilibrium costs decrease with income (Lemma 5). Thus, distribution q Lorenz-dominates distribution \hat{q} (Lemma 3 and 4), i.e., the Lorenz curve of q is always above that of \hat{q} . This implies that the inequality in \hat{q} is greater than in q . ◀

Theorem 2 can actually be generalized. Lorenz domination, a partial order, is respected by *all* inequality coefficients that satisfy the same four axioms as the Gini coefficient, yielding the following more general version of the Inequity Theorem.

► **Theorem 6.** *For any income inequality measure satisfying the axioms of invariance to population scaling, anonymity, invariance to income scaling and the transfer principle, the Inequity Theorem holds and inequality increases as tolls are levied on the players.*

The remaining of the paper focuses on the case of the Gini coefficient.

4.2 Computing the efficiency-equality trade-off

In a *parallel links network* serving a population with a known income distribution, the routing and tolls that optimize any desired trade-off between efficiency and equality can be computed via dynamic programming.

4.2.1 The model

Because of the computational nature of this section, and for the sake of simplicity, we will stick to a simplified, discrete model. Very few of these simplifications are crucial. We assume a population whose income is presented in n *quantiles* q_1, \dots, q_n , where q_1 stands for the average income of the lowest $\frac{1}{n}$ of the population – if n were 100, these would be the income percentiles.

We have K parallel links – we assume that K is fixed. Each link e has a delay function $f_e(x)$ which we assume for simplicity to be piecewise constant with increments at values of x that are multiples of $\frac{1}{n}$ (so that each link accommodates full quantiles), and that the delays have integer values in the set $[D]$, where D is the maximum delay. Evidently, the problem is one of allocating each quantile to a link, and imposing appropriate tolls. It is easy to see that at equilibrium each link will be assigned a *contiguous* set of quantiles.

³ I.e., \hat{q} is obtained from q by a transformation that reduces lower incomes relatively more than higher incomes. Income order is preserved and $\hat{\mu} \leq \mu$.

4.2.2 The objective

We seek to optimize a trade-off between efficiency and equality, that is to say, a weighted sum of total delay and the Gini coefficient resulting from this game, say of the form “minimize total delay + λ times the Gini after the game,” where $\lambda > 0$ is the relative importance of equality over efficiency. It is important to note that the Gini coefficient before the game is in this case captured by (ignoring additive terms and a factor of $\frac{-2}{n}$)

$$\frac{\sum_{i=1}^n (n+1-i)q_i}{\sum_{i=1}^n q_i}.$$

This is on account of the fact that, in the sum that approximates the double integral in Equation (LC), the lowest quantile appears n times, the second lowest $n-1$, etc.

After the game imputes a cost to the i th quantile, the Gini coefficient is captured by

$$\frac{\sum_{i=1}^n (n+1-i)(q_i - \delta_i)}{\sum_{i=1}^n (q_i - \delta_i)},$$

where $\delta_i = q_i d_i + \tau_i$ is the cost of the equilibrium to the i th percentile, and d_i is the delay and τ_i is the toll incurred by the i th quantile. Now, since it is reasonable to assume that $\delta_i \ll q_i$, this quantity can be adequately represented by its numerator divided by the sum of the q_i 's⁴. Thus, omitting constant terms (it is important to recall that the q_i 's are constant), both additive and multiplicative, we conclude that what is minimized is a linear function of the delays d_i and the tolls τ_i . Adding to them the total delay⁵, we conclude that the objective is of the form

$$\min_{\text{allocation of quantiles to links}} \sum_{i=1}^n (\alpha_i d_i + \beta_i \tau_i),$$

for some known positive parameters α_i, β_i .

4.2.3 The algorithm

The algorithm is dynamic programming; namely, we compute the quantity $\text{cost}[S, m, d]$ with $S \subseteq [K]$, $m \leq n$, and $d \leq D$, which is the smallest value of the objective that can be achieved by allocating the *lowest* m percentiles to the set S of links (in the optimum order) with the (largest) delay of the m -th percentile equal to d . The algorithm is presented in Algorithm 1.

By $\tau(d, d', r, m-r)$ we denote the toll required to equalize, for the $m-r$ th quantile, the delay d' with the greater delay d . In conclusion (here $D^* \leq n$ is the number of different values of the delay in the network):

► **Theorem 7.** *The optimum trade-off between total delay and the Gini coefficient can be computed in time $O(nD^*)$*

But of course, the O -notation hides the constant K^{22K} .

⁴ For more accuracy, the computed value of $\sum_i \delta_i$, can be plugged in here and repeat the computation.

⁵ Note that even the *total weighted delay* $\sum_i q_i d_i$ can be similarly accommodated as part of the trade-off.

Algorithm 1: A dynamic programming algorithm to compute the trade-off between efficiency and equality.

Data: Calculate the values $\text{cost}[\{e\}, m, d]$ for all links $e, m \in [n], d \in [D]$

```

begin
  for  $s \leftarrow 2$  to  $K$  do
    for All sets  $S \subseteq L$  with  $|S| = K$  do
      for  $m \leftarrow 1$  to  $n$  do
        for  $d \leftarrow 1$  to  $D$  do
           $\text{cost}[S, m, d] = \min_{e \in S, r < m: \ell_e(r)=d; d' \leq d} \text{cost}[S - \{e\}, m - r, d']$ 
             $+ \sum_{j=m-r+1}^m (\alpha_j d' + \beta_j t(d, d', r, m - r))$ 
        end
      end
    end
  end
end
end
end
end

```

4.3 The asymmetric case

In the case of multiple source-destination pairs the inequality within each set of players in any single commodity is again worsened as a result of tolls. Such a statement is not obtainable for the society as a whole. In the full online text [22], we show how to create, admittedly contrived, counterexamples where despite the fact that within each subpopulation the inequality worsens the population as a whole becomes more equal (e.g. the rich and poor use different subnetworks and only the rich get taxed). We believe that such adversarial counterexamples may be circumvented by imposing more realistic models, and pose this as one of the possible directions for future work.

5 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 8.** 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

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

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

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

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.

5.1 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

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

The canonical cost functions, besides having strong experimental justification [2,8] provide also significant advantages in the theoretical study of inequity. Specifically, the inequity index is invariant under scaling of the population incomes.

► **Theorem 9** (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 all incomes are scaled by a constant $\lambda > 0$ and optimal tolls are used in the resulting game Γ_λ , then $I(\Gamma) = I(\Gamma_\lambda)$.*

Proof Sketch. To give the main idea of the proof, we introduce a scaling parameter $\lambda > 0$. This parameter can be understood as a redenomination of the value of money in the game, for both income and the tolls, where one unit of the “new” currency is effectively as valuable as λ units of the previous currency. As such, this does not affect the strategic content of the game (no change of actions) nor the costs, by the scale invariance property of the Gini coefficient. ◀

5.2 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, \dots 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.

$$R(T) = \frac{1}{T} \sum_{i=1}^T \text{cost}^{F_i}(x) - \min_{p \in \mathcal{P}} \frac{1}{T} \sum_{i=1}^T \sum_{e \in p} f_e(q(x), c^{F_i}(e), \tau_e) \rightarrow 0 \text{ as } T \rightarrow \infty$$

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

$$\int_0^1 \text{cost}^{F_\epsilon}(x) dx - \min_{p \in \mathcal{P}} \sum_{e \in p} f_e(q(x), c^{F_\epsilon}(e), \tau_e) \leq \epsilon.$$


Following the results in [7], we can show that under regret minimizing agents, the flow converges to that of an approximate equilibrium under the assumption of a finite number of wealth/income levels w_1, \dots, w_K . This assumption is rather realistic since in practice there can only be a finite number of income levels. Also, any continuous distributions over incomes can be approximated to arbitrary high accuracy by a distribution of finite but large enough support.

► **Theorem 10** (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*

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

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

31:12 Wealth Inequality and the Price of Anarchy



■ **Figure 2 a.** The Inequity index as a function of the income coefficient β . It is 0 when there is no inequality ($\beta = 0$) because tolls have the same effect on everybody, and rises as the inequality increases. At some point, $\beta \approx 1.688$, the Inequity index starts decreasing, because the toll $(\beta + 1)2^{-(\beta+1)}$ becomes small and has little effect on the inequality index. **b.** Values of the inequality for different β .

Proof Sketch. In the first step of the proof, we show that the symmetric game of type-specific costs Γ reduces to an asymmetric congestion game $\hat{\Gamma}$. In the second step, results on the behavior of no-regret dynamics in asymmetric congestion games [7] imply the robustness of the inequity index. This is due to the Gini coefficient being a bounded and continuous function. ◀

6 Computing the Inequity in Pigou

To illustrate the interplay between wealth or income and congestion games, we consider the well-studied Pigou network, which consists of two parallel links with latency functions $\ell_u(r) = 1$ and $\ell_d(r) = r$. Assume that this transportation network is used by a population of (normalized) size 1 and with wealth or income function $q(x) = x^\beta$, for some nonnegative parameter β .

The perceived cost for quantile x is $\text{cost}(x) = \ell_e(c(e)) \cdot q(x) + \tau_e$, where $e = e(x)$ is the edge used by the quantile, $c(e)$ is the flow through link e , and τ_e the price of link e . It is not hard to argue that at equilibrium the c_u fraction of the population that uses the constant cost link is the poorest c_u part of the population. We will assume further that *the social designer selects tolls to minimize the actual latency* on the network,⁶ and that, without loss of generality, the price of the constant cost edge is set to $\tau_u = 0$, while $\tau = \tau_d$ is the optimal price on the variable cost edge.

By continuity, *at equilibrium the perceived cost of quantile c_u must be the same in both links*, from which we get $\tau = q(c_u)c_u$. We want to investigate the effects of price τ on the Gini coefficient.

Let $\hat{q}(x) = q(x) - \alpha \cdot q(x) \cdot \text{cost}(x)$ be the perceived income when we take into account the effects of perceived latency into the actual income, where α indicates the importance of transportation. We are interested in first order effects, so we will always assume that α is very small and that in fact it tends to 0. Let's define as $G(\alpha, \tau)$ the inequality coefficient when we take into account the effects on the income of the transportation cost, assuming that the social designer selects toll τ .

⁶ There are reasonable alternatives for the social planner, such as minimizing the social cost, that we do not explore in this work.

We can now compute directly the Gini coefficient. For the Pigou network the optimal switching point is $c_u = 1/2$. For income distribution $q(x) = x^\beta$, this optimal switching point corresponds to toll $\tau = 2^{-(\beta+1)}$. Since at equilibrium, $\tau = q(c_u)c_u = c_u^{\beta+1} = 2^{-(\beta+1)}$, we have

$$G(\alpha, (\beta + 1)2^{-(\beta+1)}) = \frac{\beta}{\beta + 2} + \frac{\beta(\beta + 1)}{(\beta + 2)2^{\beta+3}}\alpha + O(\alpha^2).$$

The last expression comes from the Maclaurin expansion of the function, from which we derive the following Theorem.

► **Theorem 11.** *For the Pigou network with two links and latency functions 1 and x , and for a population with income distribution $q(x) = x^\beta$, when tolls are selected to minimize the actual latency, the toll at Nash (Wardrop) equilibrium is $\tau = q(c_u)c_u = 2^{-(\beta+1)}$, and the Inequity index is $I(\Gamma) = \frac{dG(0)}{d\alpha} = \frac{\beta(\beta+1)}{(\beta+2)2^{\beta+3}}$.*


The values of the inequity as a function of the income coefficient β are shown in Figure 2. The maximum occurs when the income coefficient β is close to 2 (actually when $\beta \approx 1.688$), which means that real-life income distributions have almost the maximum Inequity index.

7 Tolls and Inequality: Empirical Findings

We use detailed transportation data gathered through Singapore’s National Science Experiment (NSE) to test how income inequality affects the distribution of transportation delays in a representative sample of students [29, 30, 43]. Although Singapore is the third most densely populated country in the world, the modern infrastructure, cost of private cars, and significant tolls in Singapore minimize 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) [42]. 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.

By relying on the sociological literature pertaining to income inequality and Singapore’s urban development, we divide the students into 9 wealth brackets based on residence. We then conservatively classify these brackets as low-income, middle-income, and high-income groups [1, 14, 17, 24, 26, 36, 41]. The differences in the means of trip distance and duration are statistically significant among these three groups; and, they lend strong support to the predictions of the Inequity Theorem. When one compares low-income and high-income 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 sizable difference in the average duration that students spend in traveling to school (Figure 3).

Although students from high-income groups travel a longer distance compared to middle-income groups, this translates into minor differences in travel duration. The opposite is the case when we compare low-income and middle-income groups. These students experience on average 7 to 5 minutes delay compared to middle-income groups, despite the fact that the distance they travel is roughly comparable to high-income groups (Figure 3). 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. It also provides some lessons on the policies that can be implemented to mitigate the trade-off between efficiency and inequality, as we discuss in the full text [22].



■ **Figure 3** The average trip duration per wealth bracket is presented in the two figures above. *Left:* Average duration in each transport mode. It is notable that brackets 1 and 2 have respectively 7 and 5 minutes more travel time on average than the other brackets. *Right:* The left plot is split along the transport mode, to show the relative durations spent in each mode across wealth brackets. Given that the least-affluent groups spend much more time on the roads compared to middle-income groups, there is a quasi-monotonic increase in the use of car as wealth increases, while we observe that the uses of walking and public transportation decrease as wealth increases.

8 Discussion

The Inequity Theorem raises important questions pertaining to distributive justice and the efficiency of decentralized decision-making mechanisms such as markets [4, 21, 27, 33–35]. Namely, if efficiency can be obtained through purposeful intervention but only at a price of increase in inequality, what are the implications of this trade-off for the organization of markets, industries, and society in general? We offer three potential avenues of research:

- What is the opportunity cost of inequity?
- How does inequity affect cooperation among members of society?
- How does inequity affect the formation of groups and thus cooperation between different groups?

We believe that these questions hold the promise of opening up new lines of research for algorithmic game theory. We hope that future work in this area will shed light on important but largely unexplored issues about the interplay between efficiency and optimality in a wide range of economic scenarios and mechanisms.

References

- 1 Tilak Abeysinghe and Keen Meng Choy. The Aggregate Consumption Puzzle in Singapore. *Journal of Asian Economics*, 15(3):563–578, June 2004. doi:10.1016/j.asieco.2004.05.007.
- 2 Pedro AL Abrantes and Mark R Wardman. Meta-analysis of UK values of travel time: An update. *Transportation Research Part A: Policy and Practice*, 45(1):1–17, 2011.
- 3 Anthony B Atkinson et al. On the measurement of inequality. *Journal of economic theory*, 2(3):244–263, 1970.
- 4 Jens Beckert. *Beyond the Market: The Social Foundations of Economic Efficiency*. Princeton University Press, Princeton, NJ, 2002.

- 5 Umang Bhaskar, Katrina Ligett, Leonard J Schulman, and Chaitanya Swamy. Achieving target equilibria in network routing games without knowing the latency functions. *Games and Economic Behavior*, 2018.
- 6 Vittorio Bilò and Cosimo Vinci. Dynamic taxes for polynomial congestion games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 839–856. ACM, 2016.
- 7 Avrim Blum, Eyal Even-Dar, and Katrina Ligett. 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*, pages 45–52. ACM, 2006.
- 8 Maria Börjesson, Mogens Fosgerau, and Staffan Algers. On the income elasticity of the value of travel time. *Transportation Research Part A: Policy and Practice*, 46(2):368–377, 2012.
- 9 Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, and Evi Papaioannou. The Impact of Altruism on the Efficiency of Atomic Congestion Games. In Martin Wirsing, Martin Hofmann, and Axel Rauschmayer, editors, *Trustworthy Global Computing*, pages 172–188, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- 10 Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. The robust price of anarchy of altruistic games. In *International Workshop on Internet and Network Economics*, pages 383–390. Springer, 2011.
- 11 Zhou Chen, Qi Qi, and Changjun Wang. Balancing Efficiency and Equality in Vehicle Licenses Allocation. In *International Conference on Web and Internet Economics*. Springer, 2017.
- 12 Zhou Chen, Qi Qi, Changjun Wang, and Wenwei Wang. Mechanism Design with Efficiency and Equality Considerations. In *International Conference on Web and Internet Economics*, pages 120–132. Springer, 2017.
- 13 George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73. ACM, 2005.
- 14 Beng-Huat Chua. *Communitarian Ideology and Democracy in Singapore*. Routledge, London, 1995.
- 15 Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 521–530. ACM, 2003.
- 16 José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
- 17 Robert H. Edelstein and Sau Kim Lum. House Prices, Wealth Effects, and the Singapore Macroeconomy. *Journal of Housing Economics*, 13(4):342–367, December 2004. doi:10.1016/j.jhe.2004.09.006.
- 18 Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 277–285. IEEE, 2004.
- 19 Dimitris Fotakis, Laurent Gourvès, and Jérôme Monnot. Selfish transportation games. In *International Conference on Current Trends in Theory and Practice of Informatics*, pages 176–187. Springer, 2017.
- 20 Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2–3):226–239, 2005. Automata, Languages and Programming: Algorithms and Complexity (ICALP-A 2004) Automata, Languages and Programming: Algorithms and Complexity 2004. doi:10.1016/j.tcs.2005.09.024.
- 21 Kurtuluş Gemici. The Neoclassical Origins of Polanyi’s Self-Regulating Market. *Sociological Theory*, 33(2):125–147, June 2015. doi:10.1177/0735275115587389.
- 22 Kurtuluş Gemici, Elias Koutsoupias, Barnabé Monnot, Christos Papadimitriou, and Georgios Piliouras. Wealth inequality and the price of anarchy. arXiv:1802.09269.
- 23 Corrado Gini. Measurement of inequality of incomes. *The Economic Journal*, 31(121):124–126, 1921.

31:16 Wealth Inequality and the Price of Anarchy

- 24 Sun Sheng Han. Polycentric Urban Development and Spatial Clustering of Condominium Property Values: Singapore in the 1990s. *Environment and Planning A: Economy and Space*, 37(3):463–481, March 2005. doi:10.1068/a3746.
- 25 Tobias Harks, Ingo Kleinert, Max Klimm, and Rolf H Möhring. Computing network tolls with support constraints. *Networks*, 65(3):262–285, 2015.
- 26 W. Gregg Huff. The Developmental State, Government, and Singapore’s Economic Development Since 1960. *World Development*, 23(8):1421–1438, 1995.
- 27 Christopher Jencks. Does Inequality Matter? *Daedalus*, 131(1):49–65, 2002.
- 28 Elias Koutsoupias and Christos H. Papadimitriou. Worst-case Equilibria. In *STACS*, pages 404–413, 1999.
- 29 Barnabé Monnot, Francisco Benita, and Georgios Piliouras. Routing Games in the Wild: Efficiency, Equilibration and Regret. In *International Conference on Web and Internet Economics*, pages 340–353. Springer, 2017.
- 30 Barnabé Monnot, Erik Wilhelm, Georgios Piliouras, Yuren Zhou, Daniel Dahlmeier, Hai Yun Lu, and Wang Jin. Inferring Activities and Optimal Trips: Lessons From Singapore’s National Science Experiment. In Michael-Alexandre Cardin, Saik Hay Fong, Pao Chuen Lui, and Yang How Tan, editors, *Complex Systems Design & Management Asia*, pages 247–264. Springer, 2016.
- 31 Arthur Cecil Pigou. *The economics of welfare*. Palgrave Macmillan, 1920.
- 32 Thomas Piketty. *Capital in the 21st Century*. Harvard University Press Cambridge, MA, 2014.
- 33 Karl Polanyi. *The Great Transformation: The Political and Economic Origins of Our Time*. Beacon Press, Boston, MA, 1957.
- 34 John Rawls. *A Theory of Justice*. Belknap Press, Cambridge, MA, rev. edition, 1999.
- 35 John Rawls. *Justice as Fairness: A Restatement*. Harvard University Press, Cambridge, MA, 2001.
- 36 Sean F. Reardon and Kendra Bischoff. Income Inequality and Income Segregation. *American Journal of Sociology*, 116(4):1092–1153, January 2011. doi:10.1086/657114.
- 37 Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. of STOC*, pages 513–522, 2009.
- 38 Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.
- 39 Joseph E Stiglitz. *The price of inequality: How today’s divided society endangers our future*. WW Norton & Company, 2012.
- 40 H Theil. *Economics and information theory*. North-Holland, Amsterdam, 1967. OCLC: 710718565.
- 41 Carl A. Trocki. *Singapore: Wealth, Power and the Culture of Control*. Psychology Press, 2006.
- 42 Erik Wilhelm, Don MacKenzie, Yuren Zhou, Lynette Cheah, and Nils Ole Tippenhauer. Evaluation of Transport Mode Using Wearable Sensor Data from Thousands of Students. In *Proceedings of the Transportation Research Board 96th Annual Meeting*, pages 1–18. Transportation Research Board, 2017.
- 43 Erik Wilhelm, Yuren Zhou, Nan Zhang, Jacksheng Kee, George Loh, and Nils Tippenhauer. Sensg: Large-scale deployment of wearable sensors for trip and transport mode logging. In *Transportation Research Board 95th Annual Meeting*, 2016.