

Routing Games in the Wild: Efficiency, Equilibration and Regret

Large-Scale Field Experiments in Singapore

Barnabé Monnot^(✉), Francisco Benita, and Georgios Piliouras

Engineering Systems and Design, Singapore University of Technology and Design,
Singapore, Singapore

monnot.barnabe@mymail.sutd.edu.sg,
{francisco.benita,georgios}@sutd.edu.sg

Abstract. Routing games are amongst the most well studied domains of game theory. How relevant are these theoretical models and results to capturing the reality of everyday traffic? We focus on a semantically rich dataset that captures detailed information about the daily behavior of thousands of Singaporean commuters and examine the following basic questions:

- Does the traffic equilibrate?
- Is the system behavior consistent with latency minimizing agents?
- Is the resulting system efficient?

The answers to all three questions are shown to be largely positive. Finally, in order to capture the efficiency of the traffic network in a way that agrees with our everyday intuition we introduce a new metric, the *stress of catastrophe*, which reflects the combined inefficiencies of both tragedy of the commons as well as price of anarchy effects.

1 Introduction

Congestion games are amongst the most historic, influential and well-studied classes of games. Proposed in [27] and isomorphic to potential games [19] (in which learning dynamics equilibrate), they have been successfully employed in a myriad of modeling problems. Naturally, one application stands above the rest: modeling traffic. Having strategy sets correspond to the possible paths between source and sink nodes in a network is such a mild and intuitive restriction that routing/congestion games are effectively synonymous to each other and jointly mark a key contribution of the field of game theory.

Routing games have also played a seminal role in the emergence of algorithmic game theory. The central notion of Price of Anarchy (PoA), capturing the inefficiency of worst case equilibria, was famously first introduced and analyzed in routing games [15,30]. Routing games have set the stage for major developments in the area such as the introduction of regret-minimizing agents [5] that eventually led to the consolidation of most known PoA results under the umbrella of (λ, μ) -smoothness arguments [28]. Impressively, this work established

that PoA guarantees are robust for a wide variety of solution concepts such as regret-minimizing agents. Finally, congestion games still drive innovation in the area with results that extend the strength and applicability of PoA bounds for large routing games [9], as well as dynamic populations [17].

With every successive analytical achievement seemingly chipping slowly away at the distance between theoretical models and everyday reality, the PoA constants for routing games, e.g. the $4/3$ for the nonatomic linear case [30] have become something akin to the universal constants of the field. Small, concise, dimensionless, they seem almost by their very nature to project purity and truth. But do they? After all, there are many of them. In the case of quadratic cost functions PoA ≈ 1.626 , whereas for quartic functions, which have been proposed as a reasonable model of road traffic, PoA ≈ 2.151 [29, 32]. What do these “small constants” mean in practice? Quite a lot. An increase of inefficiency from $4/3$ to 2.151 in Singapore would translate to the loss of approximately 730,000 work hours *every single day*. Do any of these “back-of-the-envelope” theoretical calculations have any predictive power *in practice*?

At the antipodes of the aforementioned theoretical work, other, similarly recent theoretical approaches hint that PoA analysis might actually not be reflective of the realized behavior in real networks. One type of work focuses on the instability of worst case equilibria, e.g. [14, 18]. Specifically, [25] show that although bad equilibria may exist, an average case analysis which “weighs” each equilibrium proportionally to its region of attraction typically reveals a picture that is much closer to optimal than PoA analysis. So, PoA analysis may be over-pessimistic. Distressingly, [7, 8] argue something orthogonal, which at first glance appears rather counterintuitive. They argue that networks with low PoA, e.g. PoA = 1, which are typically considered optimal, might actually reflect traffic flows which are deadlocked in severe traffic jams.¹ Finally, PoA calculations can be invalidated if we move into theoretical models that allow for risk averse agents [2, 24, 26]. At this point, as theory alone does not suffice to provide a definitive answer, it makes sense to examine some real world networks at a fine level of detail.

Our goal is to perform the first-to-our-knowledge game theoretic modeling and investigation of a real world traffic network (specifically Singapore’s traffic network) based on repeated large scale field experiments with thousands of participants. Our dataset includes granular information that allows us to inspect minute-by-minute the concurrent decision-making of thousands of commuters, as they respond and adapt to traffic conditions. We focus on arguably the three most basic questions: Is the system at equilibrium? Is this equilibrium consistent with the hypothesis of latency-minimizing agents? Is the resulting system efficient? Before we explore the answers to these questions as provided by the data, let’s try to disambiguate the questions themselves.

¹ Indeed, if we keep increasing the total flow in e.g. Pigou’s example, eventually both in the optimal and equilibrium flows almost all flow will be routed through the slow link.

Is the system at “equilibrium”? Here, we should clearly point out that by equilibrium we mean the formal mathematical notion of equilibrium, i.e. a stationary point. At the first level of inspection, we are not concerned with whether the outcome that the system equilibrates upon is necessarily stable in a game-theoretic sense. We are merely asking “are the agents continuously adapting their behavior from day-to-day” (i.e. the paths they choose, the modes of transportation, and so on)? If significant number of agents choose the same actions from day-to-day this would indicate that the system has indeed reached a fixed point (stasis) and furthermore that at this stable system state there is little entropy/randomness. Such a result is consistent with best response and best response dynamics, with the instability results of mixed Nash equilibria for multiplicative weight update algorithms [14, 18, 25] as well as with some other concurrent dynamics (e.g. imitation dynamics) [1, 10]. On the other hand, it is not a universal consequence of no-regret learning in congestion games [5, 21].

Is the equilibrium “economically stable”? Naturally, from a game theoretic perspective, we wish to understand whether the resulting equilibrium is a Nash equilibrium (or at least if in the case of adapting agents most have low regret when comparing their performance with the best path in hindsight). For a real traffic network, however, it is not practically feasible to compute true “best responses”, since there is an astronomically large number of paths to consider and we do not have data on all paths. We instead estimate inefficiencies at the individual level by quantifying the empirical “imitation” regret for each agent, i.e. how much faster could each agent have reached their destination if they had clairvoyant access to all the routing choices/information from our dataset and chose the best such route with hindsight.

Is the system “efficient”? Traditionally, ever since its inception, the notion of Price of Anarchy has been considered the gold standard for system efficiency with a low PoA considered equivalent to system optimality. The results in [7, 8] in which hopelessly deadlocked traffic jams score perfect PoA scores point out a clear dichotomy between what PoA analysis identifies as an efficient traffic network and what we in our everyday experience identify as a well-functioning network. The reason for this divide lies on the fact that PoA analysis completely disregards any inefficiency that is connected to tragedy of the commons effects.

In order to shed some light on these effects, we define a new inefficiency metric that is defined as the ratio of the social welfare at equilibrium divided by the optimal social welfare when we discount for congestion effects. Namely, although the numerator is as in the PoA, the denominator is computing the average social “blue-sky” optimal welfare as follows: Each agent imagines the scenario where she alone was in the network and computes the best path (minimum length/latency) for herself. This makes sense from an everyday experience perspective, as the typical commuter has an intuitive grasp of how long it would take to cover this distance if the externality costs imposed by the other drivers were removed. We call this ratio, the *Stress of Catastrophe (SoC)*. As this ratio grows the system’s long term persistence is jeopardized. Practically successful networks should have small SoC, which implies small PoA but not the other way around.

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 Sect. 3.1).
- The empirical regret distribution has a median value of 4 min 40 s and mean approaching 6 min for an average travel time of around 29 min (see Sect. 3.2).
- Finally, we define and estimate the Stress of Catastrophe at 1.34, with marked contrast when discriminating by mode of transportation (see Sect. 3.3). These findings are shown to be consistent across different days.

2 Description of the Data

We focus on a semantically rich dataset from Singapore’s National Science Experiment (NSE), a nationwide ongoing educational initiative led by researchers from the Singapore University of Technology and Design (SUTD). This dataset includes precise information about the daily behavior of tens of thousands of Singapore students that carry custom-made sensors for up to 4 consecutive days, resulting in millions of measurements. Indeed, every 13 s, the sensor is able to accurately log its geographical location as well as other environmental factors such as relative temperature and humidity or noise levels.

The students are dispersed throughout the city-state and their daily commutes to school are reasonably long for them to meaningfully interact and experience the daily traffic. For this reason, we focus on the morning trip they undertake to reach their school from their home. The morning trip is also characterized by a lesser number of stops on the way to school or Pre-university, thus it lends itself better to an analysis based solely on travel times. Other types of costs may be included to complement travel time, such as price of the route (based on tolls or public transport fees) or environmental factors. In this study however, our scope is limited to the trip duration, to be extended in future work.

The mode of transportation chosen by the students can be identified using accurate algorithms, e.g. car (driving or being driven to school) versus bus or metro, estimate source and sink destinations (focusing on home-school pairs) as well as their mode-dependent available routes. Some descriptive charts are given in Fig. 1 relating the durations and distances traveled for private and public transportation trips. To guide the reader unfamiliar with Singapore’s road and public transport network, we give a brief optional introduction in our online full version of the paper [20].

Representativeness of the Sample. Students are a restricted class of residents, but we argue that they however provide a tangible idea of Singapore’s mobility. First, as of 2015, the size of the student population up to Pre-University level totals about 460,000 residents. In contrast, the active population’s size, as of 2015, is about 2.2 million.² Our clean dataset includes 32,588 trips taken by

² Statistics were compiled from data.gov.sg.

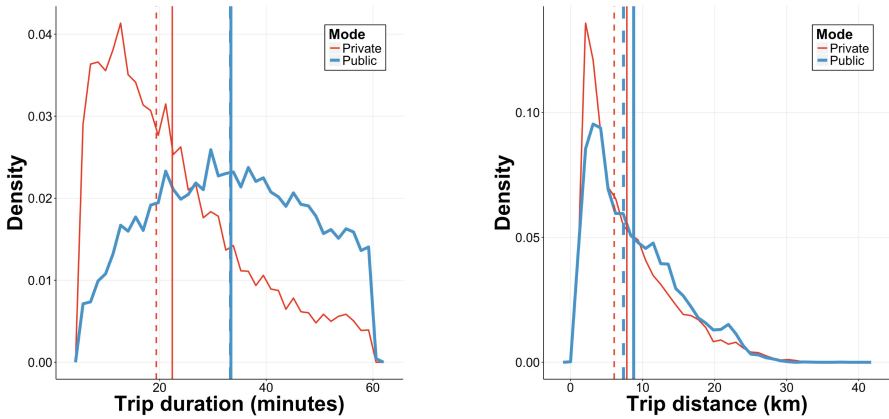


Fig. 1. *Left:* Density plots of trip durations per mode. We note that car trip durations are typically short and more concentrated around a peak value of 15 to 20 min, while public transportation trip durations are scattered between 20 to 50 min. *Right:* Density plots of trip distances per mode. The two densities are close, indicating that distance may not factor in the choice of transportation mode. Median is represented by a dashed line, mean by a solid line.

15,875 unique students, distributed between the three main type of institutions in Singapore (Primary, Secondary and Pre-University).

For the purpose of our study, most of the analysis does not require a complete sample of the population. Students in private transportation experience the same level of congestion as their peers and active individuals, hence estimates over their population translate to estimates over the whole of Singapore’s mobility users. It is even more true for students in public transportation: their trips are possibly the same as those of the active population. Indeed, we find that the ratio of public to private transportation users in our sample closely mirrors that of the population as a whole³, as 57% of students in our dataset use public transportation.

As shown in Fig. 2, the sample of home locations is geographically distributed, so is not focused on a particular area of the city. However, the distribution of schools may not reflect endpoints of trips made by the active population. As an example, it can be observed that few schools are located in the city center, which houses a large number of office buildings. This constitutes one limitation of our dataset, perhaps softened by the fact that active population and students may still share a sizeable part of their route and thus experience the same congestion.

³ Household Interview Travel Survey 2012: Public Transport Mode Share Rises To 63%, LTA News Release.




Fig. 2. *Right:* Home locations (red dots), school locations (blue triangles) and spatial clustering methods, discussed in Sect. 3.2 and the full version online [20]. *Left:* Density map of Singapore. Blue areas are less populated while red areas are denser. (Color figure online)

3 Findings

3.1 Equilibration and Empirical Consistency

If the traffic system is at equilibrium, then we should expect that the subjects' route decisions do not vary substantially between successive days of study. We investigate the issue from three different angles. First, we compare the modes of transportation selected by each individual student over the days of the experiment. Second, we improve the previous result by considering whether the selected routes are identical (e.g. always use the same combination of bus and train, or always use the same road on car). Third, building on our geographical clustering method described in the following Section, we investigate the question of whether the fastest student in the cluster on one day remains the fastest over all days of experiment.

The first analysis shows that more than 60% of subjects have used the same principal mode of transportation in all morning trips available in our dataset. We are here discriminating between trips where the principal mode of transportation is either the train, the bus or the car. We define as principal the mode with which the student has traveled the longest distance. The fraction increases to close to two thirds (65%) of the samples if we simply discriminate between the subjects using public transit from those who use private transportation.

For the second analysis, we have implemented a novel algorithm to determine whether two route choices are identical. We find that for subjects using the same mode of transportation across all days, the percentage of subjects selecting the same route is very high, in the order of 94%. We detail the algorithm used to obtain this result in our Methodology section of full paper available online [20].

Finally, we identify a restricted set of clusters that have the property of being consistent throughout at least two days of experiment, i.e. the members of

the cluster are the same in distinct days of the same week. Members may drop out of their cluster if their starting time or starting point are different from one morning to the next, or if they use another mode of transportation. We find that for these consistent clusters, close to 50% of them have the property that the fastest individual on one day remains the fastest for all days where this cluster appears, showing again a certain degree of consistency in the population.

3.2 Individual Optimality and Empirical Imitation-Regret

To answer the question of individual optimality, we compare the durations of the morning trip for the subjects. A fair comparison is only achieved when looking at students leaving from the same neighborhood on the same day and at roughly the same time, going to the same school and using the same mode of transportation. The notion of neighborhood is expanded upon in our Methodology section, available in the full version of our paper [20], where we describe how the clustering of the data was achieved.

In the cases where the class of comparable subjects has more than two individuals, we collect the **empirical imitation-regret** encountered by every student in the class. To do so, we find the student in the class with minimal trip duration and set her imitation-regret to zero. For other members of the class, the empirical imitation-regret is equal to the (non-negative) difference between their trip duration and the minimal trip duration.

Our notion of empirical imitation-regret shares its name with the traditional regret measure, commonly found in the learning and multi-agent systems literature, for the following reason. The players here are faced with multiple strategies that they can choose from: the routes that go from their neighborhood 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 measured against the best action implemented by a comparable subject on that day, and the difference is the imitation-regret. The introduction 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 subjects in the cluster will therefore not be considered here. This drawback is shared with many natural learning 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 neighborhood. As the area increases, so does the accumulated imitation-regret, since the minimum is taken over a larger set of subjects. However, neighborhoods that are too large lose in precision, as two different subjects in the same cluster may have very different trip lengths. The results in this section use a geographical cluster size of about 400 m, while sensitivity analysis is performed in the Methodology section of our full paper [20] to show the robustness of our findings.

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 one individual encounters a imitation-regret

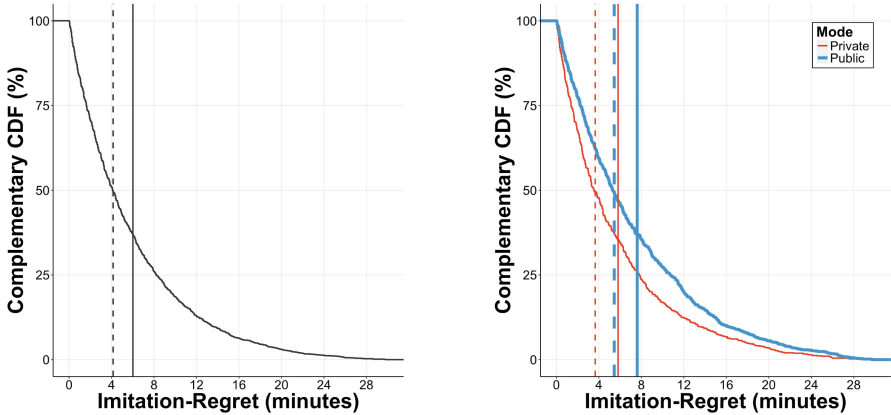


Fig. 3. *Left:* Complementary cumulative distribution function of the imitation-regret (decreasing curve). We aggregate all days of the experiment in a single figure and remove subjects with zero imitation-regret – i.e. the baseline subjects. The mean imitation-regret signalled by the solid vertical line is equal to 6 min (around 27% of the mean travel time), while the median imitation-regret – dashed line – is equal to 4 min and 40 s (around 21% of the median travel time). Sensitivity analysis results are presented in the full paper online [20]. *Right:* Comparison of complementary CDF of imitation-regret per mode of transportation. (Color figure online)

of say, 10 min, she may be better off by 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 after such a conclusion. If we assume that individuals are solely interested in minimizing their trip duration—perhaps a fair assumption for the morning trip, constrained by the hard deadline of the class start—, 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 more climatized one or one that is shared with other students.

In Fig. 3, we plot the complementary cumulative distribution of the empirical imitation-regret. A point on the curve indicates which fraction of individuals (read on the y -axis) have empirical imitation-regret greater or equal than x (read on the x -axis). We also give the mean (solid red line) and median (dashed blue line) experienced empirical imitation-regret. It should be noted that the empirical imitation-regret distribution and its moments do not include the subjects for which the imitation-regret is zero, i.e. the best in the cluster.

Larger geographical cluster sizes give rise to larger average empirical imitation-regrets, but the results are relatively robust. The mean empirical imitation-regret oscillates around 27% of the mean travel time in the dataset

(around 6 min), while the median empirical imitation-regret is at 21% of the median travel time (around 4 min and 40 s). 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 $\epsilon = 22$ min and $\delta = 0.05$.

Finally, 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 and one individual using private transportation appear. We have over 1,400 such clusters, and in close to 80% of them, the fastest individual is a private transportation user. Over these 1,400 clusters, the average imitation-regret incurred by public transport users compared with the fastest private transportation user in their cluster is close to 8 min. For the same population of bus and train users, the average duration of a trip is close to 25 min, indicating that the fastest car user spends roughly two thirds of this time to reach destination. Figure 3 plots the distributions of imitation-regret for the two classes of users.

3.3 Societal Optimality and the Stress of Catastrophe

The Stress of Catastrophe 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 Price of Anarchy 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 that a central planner implements.

But estimating the social optimum of a system from the data is a perilous task. First, exact demands need to be known 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 flow maximizing the social optimum function needs to be computed.

On the other hand, the 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 enter, the global welfare diminishes. 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.

The Stress of Catastrophe 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. Using Google Directions API, free-flow trip durations are obtained and 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 Stress of Catastrophe (SoC) from our data as such:

$$\text{SoC} = \frac{\text{Cost(Recorded trip duration)}}{\text{Cost(Trip durations (free-flow / light traffic))}}$$

To give an idea of the measure in our dataset, we plot in Fig. 4 the histogram of percentages of deviations from the free-flow optimal trip duration. We see that most subjects are relatively close to this minimum bound while as the gap grows, fewer subjects are found.




Fig. 4. *Left:* Histogram of deviations from the free-flow optimal trip durations. *Right:* Stress of Catastrophe computed across the five days with the highest record of unique subjects (sample size > 1,500). The values are between 1.23 and 1.37.

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

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

The question is now how pessimistic is this upper bound? Our results show that $\text{SoC} = 1.34$, when the SoC is computed with both car and transit users. But discriminating between the two yields a much more contrasted picture: the SoC for transit users is found to be 1.18, indicating that students using public transportation have little room to improve their trip duration. Conversely, the SoC varies significantly depending on the traffic conditions for subjects taking private transportation to school. In free-flow conditions, we find the SoC to be equal to 1.86. Details of the SoC for individual days of the experiment can be found in Fig. 4.

It is remarkable that such an pessimistic upper bound is however so close to 1. How does the PoA overestimate the inefficiency of the network then? Consider

PoA results found in the literature, such as the 2.151 ratio of derived in [31] in the case of degree 4 polynomial cost functions. The latest class is often used by network engineers to model the congestion on real roads, following the Bureau of Public Roads standard.

But the average estimated free flow time travel of the sample is 21 min. Assuming the SoC to be as large as the 2.151 bound, on average a commuter would spend $2.151 - 1.34 = 0.811$ times more in transit, i.e. 17 min more per commuter. In other words, pessimistic predictions of the PoA would entail a loss of over 730,000 h per day, if we assume all of the 2,200,000 active individuals and 400,000 students were commuting on that day, a large mismatch with the actual system performance.

4 Connections to Other Work

Algorithmic Game Theory and Econometrics. Recently there has been a surge of interest in combining techniques from algorithmic game theory with the traditional goals of econometrics [3, 33]. These works employ a data-driven approach to analyzing the economic behavior of real world systems and agent interactions. In [22] the authors developed theoretical tools for inferring agent valuations from observed data in the generalized second price auction without relying on the Nash equilibrium assumption, using behavioral models from online learning theory such as regret-minimization. They apply their techniques on auction data to test their effectiveness.

Following this work, [13] studies the behavior of real housing market agents based on data from an online bidding platform. The results inform the design of the auction platform and point towards data-driven policies helping the agents make decisions. The latter idea is made more explicit in a recent article by some of the authors [23]. In a sense, our present work also advocates using data to gauge users interactions but our focus is on routing games, for which it is harder to gather sanitized data. Furthermore, we develop new metrics that are more informative about the state of the system than the price of anarchy.

In [12] the authors provided tools for estimating an empirical PoA of auctions. The PoA is defined as the worst case efficiency loss of any auction that could have produced the data, relative to the optimal. However, auctions and routing games each pose a totally distinct set of challenges. In our setting, the problem of translating data streams to game theoretic concepts adds a rather nontrivial layer of complexity. For example, even identifying the action chosen by each agent, i.e. their routes, is tricky as it requires to robustly map a noisy stream of transportation data into a discrete object, a path in a graph. It should be noted however that our notion of Stress of Catastrophe provides an upper bound on the empirical PoA, as detailed in Sect. 3.3.

Another active strand of research is concerned with fair division of resources. Here too, experimental studies are conducted to determine whether agents exhibit a behaviour close to predictions from the theory and in fact, their behaviour and feedback from using fair division systems pose new theoretical questions [11, 16]. This fruitful cycle can hopefully be replicated in congestion games.

Price of Anarchy for Real World Networks. One earlier paper tangentially connected to estimating the PoA of congestion games is [6]. This is a theoretical paper that provides PoA bounds for perturbed versions of congestion games. As a test of their techniques, they heuristically approximate the PoA on a few benchmark instances of traffic networks available for academic research from the Transportation Network Test Problems [4] by running the Frank-Wolfe algorithm on them. No experiments were performed and no measurements were made. Naturally, this approach cannot be used to test PoA predictions, since it presumes that PoA reflects the worst case possible performance and then merely tests where do these constants lie for non-worst case routing networks.

In effectively parallel independent work [34] focused on quantifying the inefficiencies incurred due to selfish behavior for a sub-transportation network in Eastern Massachusetts, US. They use a dataset containing time average speed on road segments and link capacity in their transportation sub-network. The authors estimate daily user cost functions as well as origin-destination demand by means of inverse optimization techniques using this dataset. From this formulation they compute estimates of the PoA, whose average value is shown to be around 1.5. In contrast to their approach our dataset contains detailed individual user information, which allows for estimates not only of systemic performance but also of individual optimality (e.g. imitative-regret) as well as test to what extent is the system indeed near stasis (i.e. in equilibrium). Also, their approach does not capture how bad the resulting traffic is, i.e. the tension between Price of Anarchy and Tragedy of the Commons, whereas our approach addresses both. Finally, our estimations are derived from explicit online measurements of the system performance and are not reverse engineered by estimating user cost functions which inevitably introduce new errors that cascade through all the calculations.

5 Conclusion

This is hopefully not the end but the beginning of a thorough experimental investigation into the rich game theoretic literature of routing games. Clearly, there are many open questions and challenges to be addressed. Due to space limitations we refer the reader to the online version of our paper for the full discussion of these directions [20].

Acknowledgements. The authors would like to thank the National Science Experiment team at SUTD for their help: Garvit Bansal, Sarah Nadiawati, Hugh Tay Keng Liang, Nils Ole Tippenhauer, Bige Tunçer, Darshan Virupashka, Erik Wilhelm and Yuren Zhou. The National Science Experiment is supported by the Singapore National Research Foundation (NRF), Grant RGNRF1402.

Barnabé Monnot acknowledges the SUTD Presidential Graduate Fellowship. Francisco Benita acknowledges CONACyT CVU 369933 (Mexico). Georgios Piliouras acknowledges SUTD grant SRG ESD 2015 097, MOE AcRF Tier 2 Grant 2016-T2-1-170 and a NRF fellowship. Part of the work was completed while Barnabé Monnot and Georgios Piliouras were visiting scientists at the Simons Institute for the Theory of Computing.

References

1. Ackermann, H., Berenbrink, P., Fischer, S., Hoefer, M.: Concurrent imitation dynamics in congestion games. *Distrib. Comput.* **29**(2), 105–125 (2016)
2. Angelidakis, H., Fotakis, D., Lianenas, T.: Stochastic congestion games with risk-averse players. In: Vöcking, B. (ed.) *SAGT 2013*. LNCS, vol. 8146, pp. 86–97. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-41392-6_8
3. Bajari, P., Hong, H., Nekipelov, D.: Game theory and econometrics: a survey of some recent research. In: *Advances in Economics and Econometrics, 10th World Congress*, vol. 3, pp. 3–52 (2013)
4. Bar-Gera, H.: Transportation network test problems (2011). <https://github.com/bstabler/TransportationNetworks>. Accessed 10 Nov 2017
5. Blum, A., Hajiaghayi, M., Ligett, K., Roth, A.: Regret minimization and the price of total anarchy. In: *STOC*, pp. 373–382 (2008)
6. Buriol, L., Ritt, M., Rodrigues, F., Schäfer, G.: On the smoothed price of anarchy of the traffic assignment problem. In: *ATMOS*, pp. 122–133. *ATMOS* (2011)
7. Colini-Baldeschi, R., Cominetti, R., Mertikopoulos, P., Scarsini, M.: On the asymptotic behavior of the price of anarchy: is selfish routing bad in highly congested networks? *ArXiv e-prints* (2017)
8. Colini-Baldeschi, R., Cominetti, R., Scarsini, M.: On the price of anarchy of highly congested nonatomic network games. In: Gairing, M., Savani, R. (eds.) *SAGT 2016*. LNCS, vol. 9928, pp. 117–128. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53354-3_10
9. Feldman, M., Immorlica, N., Lucier, B., Roughgarden, T., Syrgkanis, V.: The price of anarchy in large games. In: *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 963–976. ACM (2016)
10. Fotakis, D., Kaporis, A.C., Spirakis, P.G.: Atomic congestion games: fast, myopic and concurrent. In: Monien, B., Schroeder, U.-P. (eds.) *SAGT 2008*. LNCS, vol. 4997, pp. 121–132. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-79309-0_12
11. Gal, Y.K., Mash, M., Procaccia, A.D., Zick, Y.: Which is the fairest (rent division) of them all? In: *EC*, pp. 67–84. ACM (2016)
12. Hoy, D., Nekipelov, D., Syrgkanis, V.: Robust data-driven guarantees in auctions. In: *Preliminary version at 1st Workshop on Algorithmic Game Theory and Data Science* (2015)
13. Jalaly, P., Nekipelov, D., Tardos, É.: Learning and trust in auction markets. [arXiv:1703.10672](https://arxiv.org/abs/1703.10672) (2017)
14. Kleinberg, R., Piliouras, G., Tardos, É.: Multiplicative updates outperform generic no-regret learning in congestion games. In: *STOC* (2009)
15. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: *STACS*, pp. 404–413 (1999)
16. Kurokawa, D., Procaccia, A.D., Shah, N.: Leximin allocations in the real world. In: *EC*, pp. 345–362. ACM (2015)
17. Lykouris, T., Syrgkanis, V., Tardos, É.: Learning and efficiency in games with dynamic population. In: *SODA*, pp. 120–129. SIAM (2016)
18. Mehta, R., Panageas, I., Piliouras, G.: Natural selection as an inhibitor of genetic diversity: multiplicative weights updates algorithm and a conjecture of haploid genetics. In: *ITCS* (2015)
19. Monderer, D., Shapley, L.S.: Potential games. *Games Econ. Behav.* **4**(1), 124–143 (1996)

20. Monnot, B., Benita, F., Piliouras, G.: Routing games in the wild: efficiency, equilibration and regret (Large-Scale Field Experiments in Singapore). arXiv preprint [arXiv:1708.04081](https://arxiv.org/abs/1708.04081) (2017)
21. Monnot, B., Piliouras, G.: Limits and limitations of no-regret learning in games. *Knowl. Eng. Rev.* **32** (2017)
22. Nekipelov, D., Syrgkanis, V., Tardos, E.: Econometrics for learning agents. In: EC, pp. 1–18. ACM (2015)
23. Nekipelov, D., Wang, T.: Inference and auction design in online advertising. *Commun. ACM* **60**(7), 70–79 (2017)
24. Nikolova, E., Stier-Moses, N.E.: A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.* **62**(2), 366–382 (2014)
25. Panageas, I., Piliouras, G.: Average case performance of replicator dynamics in potential games via computing regions of attraction. In: EC, pp. 703–720. ACM (2016)
26. Piliouras, G., Nikolova, E., Shamma, J.S.: Risk sensitivity of price of anarchy under uncertainty. *ACM Trans. Econ. Comput.* **5**(1), 5:1–5:27 (2016)
27. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theor.* **2**(1), 65–67 (1973)
28. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: STOC, pp. 513–522. ACM (2009)
29. Roughgarden, T.: *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, Cambridge (2016)
30. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM (JACM)* **49**(2), 236–259 (2002)
31. Roughgarden, T., Tardos, É.: Bounding the inefficiency of equilibria in nonatomic congestion games. *Games Econ. Behav.* **47**(2), 389–403 (2004)
32. Sheffi, Y.: *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, New Jersey, US (1985)
33. Syrgkanis, V.: Algorithmic game theory and econometrics. *ACM SIGecom Exch.* **14**(1), 105–108 (2015)
34. Zhang, J., Pourazarm, S., Cassandras, C.G., Paschalidis, I.C.: Data-driven estimation of origin-destination demand and user cost functions for the optimization of transportation networks. arXiv preprint [arXiv:1610.09580](https://arxiv.org/abs/1610.09580) (2016)