The Price of Fragmentation in Mobility-on-Demand Services

THIBAUT SÉJOURNÉ, Ecole Polytechnique, France
SAMITHA SAMARANAYAKE, Cornell University, USA
SIDDHARTHA BANERJEE, Cornell University, USA

Mobility-on-Demand platforms are a fast growing component of the urban transit ecosystem. Though a growing literature addresses the question of how to make individual MoD platforms more efficient, much less is known about the cost of market fragmentation, i.e., the impact on overall welfare due to splitting demand between multiple independent platforms. Our work aims to quantify how much platform fragmentation degrades the efficiency of the system. In particular, we focus on a setting where demand is exogenously split between multiple platforms, and study the increase in supply rebalancing costs incurred by each platform to meet this demand, vis-à-vis the cost incurred by a centralized platform serving the aggregate demand. We show under a large-market scaling, this *Price-of-Fragmentation* undergoes a phase transition, wherein, depending on the nature of the exogenous demand, the additional cost due to fragmentation either vanishes or grows unbounded. We provide conditions that characterize which regime applies to any given system, and discuss implications of these findings on how such platforms should be regulated.

CCS Concepts: • **Networks** → **Network performance analysis**; **Network performance modeling**;

Additional Key Words and Phrases: Mobility on Demand; ridesharing; platform competition; large-market scaling; phase transition

ACM Reference Format:

Thibault Séjourné, Samitha Samaranayake, and Siddhartha Banerjee. 2018. The Price of Fragmentation in Mobility-on-Demand Services. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 2, Article 30 (June 2018), 26 pages. <https://doi.org/10.1145/3224425>

1 INTRODUCTION

The ability to provide efficient and sustainable urban transportation has been a major challenge for cities throughout the world in recent decades. This challenge is becoming even more difficult with rapid urbanization – it is projected that two thirds of the world population will live in urban areas by 2050 [14]. Traditionally, urban mobility has been satisfied via either personal vehicle ownership (which is convenient, but has high per-passenger congestion/pollution costs), or mass transit systems (which are efficient, but prone to first-and-last-mile problems). Mobility-on-Demand (MoD) systems (e.g. Uber, Lyft, Didi, Ola) provide a new alternative that preserves the convenience of personal vehicles, while reducing its burdens. Their growing popularity is a testament to their success, and they have the potential to help the overall efficiency of the transportation system as a whole.

Authors' addresses: Thibault Séjourné, Ecole Polytechnique, Palaiseau, Île-de-France, 91128, France, thibault.sejourné@polytechnique.edu; Samitha Samaranayake, Cornell University, Ithaca, New York, 14850, USA, samitha@cornell.edu; Siddhartha Banerjee, Cornell University, Ithaca, New York, 14850, USA, sbanerjee@cornell.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

2476-1249/2018/6-ART30 \$15.00

<https://doi.org/10.1145/3224425>

The popularity of MoD systems has given rise to a growing literature on fleet management, dispatch and pricing problems in these systems. However, there is limited work in understanding the effects of having multiple service providers (i.e., platforms). The economic benefits of competition for passengers and drivers (in terms of prices, wages, service quality, etc.) are well-established; sustaining competition may however result in higher overall operational costs. In particular, the concern is that when multiple platforms fragment demand, and myopically optimize their own systems to serve their own customers, they may potentially result in a higher cost to society (in terms of pollution/fuel consumption/congestion, etc) as compared to a monopoly.

Our work aims to understand and quantify the increased operational costs due to competition in MoD ecosystems. There are two challenges to doing so: First, since MoD systems involve complex network dynamics involving many vehicles, any analysis requires careful modeling of demand/supply/rebalancing and stochastic dynamics. Secondly, given the growth of MoD systems, the pertinent question is not so much what the costs are currently, but rather, how they will behave as demand scales in future; such counterfactual analysis is not directly accessible from data. Our techniques address both these questions.

In more detail: we approach this problem by defining a novel notion of the *Price of Fragmentation* – a metric that captures how the efficiency of a MoD ecosystem degrades as demand is randomly split between platforms. We focus on a setting where passengers are exogenously split between multiple platforms (i.e. we ignore endogenous competition effects), and study the increase in system costs incurred by each platform to meet this demand, vis-à-vis the cost incurred by a monopolist serving the aggregate demand. This captures the multiplexing loss (i.e., the net increase in vehicle miles traveled when Lyft customers are matched to the nearest Lyft driver although there is a closer Uber driver, and vice versa) due to having competing MoD providers. By studying the scaling behavior of the PoF with increasing demand, we uncover a surprising phase-transition phenomena (see Figure. 1), wherein the PoF either vanishes or grows unbounded, depending on a structural property of the underlying aggregate demand. We corroborate these results via synthetic experiments, as well as data-driven studies using the NYC taxi-cab data [1].

To summarize, our work is to the best of our knowledge the first study on the loss in efficiency (i.e. operational costs) of MoD systems due to the fragmentation of demand across multiple platforms. While further studies are needed to understand the effects of platform competition, we argue that understanding the systemic inefficiencies uncovered in our work is an important first step for guiding public policy for the MoD ecosystem.

1.1 Summary of our Results

We consider an MoD system with either a monopolist, or two competing firms – in the latter case, we assume the demand gets exogenously split between the two firms. Given this demand split, we focus on the *cost of rebalancing*, i.e. the total cost of aligning vehicle supply with passenger demand. In this context, demand fragmentation leads to a total rebalancing cost in a duopoly that is necessarily greater than or equal to that of the monopolist cost. To quantify this, we consider a stochastic model of demand splitting, where we assume the demand for each origin-destination pair is split between the firms according to some random process; we define the *Price of Fragmentation* (PoF) to be the expected excess rebalancing cost incurred by the duopoly. To incorporate the growth of MoD systems, we study the PoF under a *large-market scaling*: we scale all demands simultaneously by a factor θ , before randomly splitting the demand between the firms (cf. Definition 2.2 for details).

For ease of presentation, we initially focus on the case of homogeneous market shares, wherein the average fraction of demand going to each firm is the same across the network; we extend our results to heterogeneous market shares in Section 4. It is important to note that our setup assumes an independent and sufficient vehicle supply, and exogenous demand splitting. In other words, we

Fig. 1. Simulation demonstrating the scaling behavior of the Price of Fragmentation (the difference between the rebalancing cost of the duopoly and the monopoly) in the fragmentation-resilient (green) and fragmentation-affected (blue) regimes. The demand distributions and costs are based on the NYC taxi data; in particular, the two regimes shown here correspond to two successive hours (noon to 1 p.m. and 1 to 2 p.m.) on May 10. These time slots were chosen as their initial PoF (i.e., before scaling) are approximately equal, while our analysis shows that one is fragmentation-affected while the other is fragmentation-resilient.

assume that consumers have idiosyncratic demand preferences and we do not model competition between the platforms. While we acknowledge that this is an important consideration for real world systems, the insights we gain under our simpler model are by themselves significant and non-obvious, and will likely continue to hold under more complex models as well.

Our results are summarized as follows: for general networks under the large-market scaling, we show that the PoF undergoes a phase transition, depending on the structure of the underlying demand vector. In the first regime the PoF decays exponentially under scaling – we refer to this as the *fragmentation-resilient* regime. In contrast, in the second regime, the PoF grows under scaling, at a rate equal to the square root of the scaling parameter – we refer to this as the *fragmentation-affected* regime. Figure 1 shows an example of these two cases, based on demand distributions estimated using the NYC taxi data over two different one-hour time slots; note that though both the settings appear to have identical PoF at the original demand, the two curves are markedly different as we scale the average demands.

More specifically: the rebalancing cost in an MoD system is characterized by a minimum-cost circulation problem (cf. Equation (1)), and our results aim to understand how the solution of this program changes under random demand splitting. Our main results (Theorems 4.2, 4.3 and 4.4) together show that under homogeneous demand splitting, the system is in the *fragmentation-affected regime* if and only if the underlying demand vector is dual-degenerate with respect to the min-cost circulation program, even when the market is in-homogeneously shared in the city. We also extend the results for non-homogeneous splits (Theorem 4.4), where the behavior may be more complex.

The above results, though providing an explicit characterization of the phase transition, are less interpretable as they are in terms of structural properties of the optima of an LP. We next provide an elegant combinatorial characterization for the fragmentation-affected regime (Theorem 4.5). Our characterization suggests that the troubling case is when the demand induces *local balanced clusters*, i.e., the network can be partitioned in two or more clusters such that there is (roughly) balanced underlying demand going to and fro between clusters, while the cost of traveling between

clusters is high. We conjecture that the structure of urban traffic patterns can indeed lead to such balanced partitions. In Section 5, we provide empirical evidence from the NYC taxi data, where we show that local balanced clusters do often arise in practice. We also perform simulations on synthetic examples, as well as on the NYC taxi data, to verify our PoF characterizations. Finally, in Section 6, we discuss implications of our results on how MoD ecosystems should be regulated, and discuss possible extensions of our work.

1.2 Related Work

As alluded to before, MoD systems have gained widespread recent interest in the research community. A large body of this work studies the design of optimal vehicle rebalancing policies, which we discuss in detail below. Other topics of study of less relevance to our work include dynamic pricing [6] and ride-pooling [3, 19].

For studying rebalancing, a common approach is to formulate the problem as that of controlling a closed queueing network. Assuming that the number of cars and demand are very high, one can then use a fluid approximation to find rebalancing policies, and test them numerically afterwards [16, 20]. Related approaches have been used to optimize weighted throughput [11] and minimize rebalancing costs [24]. All the above works are primarily numerical, and have limited formal guarantees.

More recently, several papers [5, 7, 9, 15] formally characterize the optimal control policy under fluid (or large-market) limits of closed queueing models of ride-sharing systems. In particular, the work of Banerjee et al. [5] provides a unified framework for designing asymptotically optimal rebalancing, dispatch and pricing policies, with rigorous approximation guarantees in finite settings.

Finally we note that there is an extensive literature on competition between platforms [17, 18, 23], including in network settings [2, 4, 8] and power markets [10]. These are of less relevance to our work as we focus on the operational costs of demand fragmentation. However, it may be of interest to develop models synthesizing our work with such strategic competition settings.

2 SETTING AND PRELIMINARIES

2.1 Basic setting

Network Model: We consider a city represented as a complete directed graph $G(V, E)$, comprising of n nodes (or *stations*) V corresponding to neighborhoods, linked together by edges E corresponding to *fastest routes* between any pair of stations. Each edge has an associated *average travel-time* τ_{ij} , which can be interpreted as the cost incurred by a trip from i to j (in terms of driving time, or after appropriate re-weighting, miles driven/fuel consumption/pollution). Travel times need not be symmetric, but by definition must obey the triangle inequality.

Demand Model: Within this city, we study the operations of an MoD ecosystem (i.e., a one or more MoD providers) over a period of time during which the underlying demand for rides remains stationary; this corresponds to a period over which the average demand is roughly constant (for example, the morning rush hour on weekdays, or the afternoon hours on weekends, etc. Our results extend to time-varying arrivals, at the cost of additional notation). Since our focus in this work is on operational costs, we assume that this demand is exogenous (i.e., after pricing), and moreover, that all demand must be serviced with equal priority.

Over the period of time under consideration, we use $\{\Lambda_{ij}\}_{(i,j) \in E}$ to denote the *demand vector*, where for each edge $(i, j) \in E$, Λ_{ij} denotes the (random) number of customers requesting rides starting at node i and terminating at node j . For any node i , we define $\Lambda_i = \sum_{j \in V} (\Lambda_{ji} - \Lambda_{ij})$ to denote the *net inflow* of passengers into node i . When clear from context, we sometimes abuse notation to use Λ to denote the vector $\{\Lambda_i\}$.

The MoD Ecosystem: To model fragmentation, we consider two models of the MoD ecosystem: (i) a *monopolist* setting, where all demand is serviced by a single MoD platform, and (ii) a *duopolist* setting, where we assume the demand is split between two firms. In particular, in the duopoly, for any edge $(i, j) \in E$, we assume the two firms a and b split the demand as $(\Lambda_{ij}^a, \Lambda_{ij}^b)$ such that $\Lambda_{ij}^a + \Lambda_{ij}^b = \Lambda_{ij}$. This split is assumed to be exogenously determined (i.e., unaffected by firms' strategies) – in particular, we assume $\{\Lambda_{ij}^a, \Lambda_{ij}^b\}$ are randomly generated according to some specified distribution.

Supply and Rebalancing: Next we describe the supply side of the model for the monopolistic setting; in the duopolist setting, we assume the same model applies to each firm independently. This corresponds to a setting where each firm has an independent fleet within the time-period of interest; this is the case in some current markets (for example, where drivers are employed by the firms), and moreover, likely to be the case with self-driving fleets. Note this model does not capture *multi-homing drivers*, who work simultaneously for multiple firms; in Section 6, we discuss how our results suggest that a sufficient amount of multi-homing may potentially help reduce overall social costs.

Our chief object of study is the *cost of rebalancing* $RC(\Lambda)$, i.e., the operational cost of rerouting empty supply to meet the demand. This given by the following *min-cost circulation LP*:

$$\begin{aligned} \min \quad & \sum_{(i,j)} \tau_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_j (x_{ji} - x_{ij}) = \Lambda_i \quad \forall i \in V \quad ; \quad x_{ij} \geq 0 \quad \forall (i,j) \in E, \end{aligned} \tag{1}$$

where $\Lambda_i = \sum_j (\Lambda_{ij} - \Lambda_{ji})$ is the net inflow of demand at station i , and x_{ij} is the number of rebalancing trips from station i to j . The rebalancing cost $RC(\Lambda)$ can also be expressed via its dual:

$$\begin{aligned} \max \quad & \sum_i \alpha_i \Lambda_i \\ \text{s.t.} \quad & \alpha_i - \alpha_j \leq \tau_{ij} \quad \forall (i,j) \in E \end{aligned} \tag{2}$$

The formulation of the rebalancing cost in terms of the min-cost circulation (or *optimal transport*) problem arises in many studies of MoD systems, and has different justifications based on different underlying assumptions. The critical point to observe is that $RC(\Lambda)$ is in a sense the best ex-post cost that a platform can obtain while satisfying all demand, and that dynamic rebalancing policies used in practice are known to match this bound very closely in various settings; we discuss these in more detail in Appendix B. Consequently, we henceforth focus on the rebalancing cost $RC(\Lambda)$ as per the LP in (1) (alternately, the Dual LP (2)), and ignore additional losses due to sub-optimality in dynamic scheduling.

Before proceeding, we present a structural characterization of the function $RC(\cdot)$, which we use extensively in what follows.

PROPOSITION 2.1. *$RC(\cdot)$ is convex and homogeneous, i.e., $\forall \theta \geq 0$, $RC(\theta\Lambda) = \theta RC(\Lambda)$.*

PROOF. Consider the characterization of $RC(\cdot)$ in (2): homogeneity follows from factoring out θ from the objective, while convexity is seen by noting that $RC(\Lambda)$ can be written as the maximum of a finite number of linear functions $\alpha^\top \Lambda$, with α being the corner points of the dual polytope. $\square$

Finally, we note that though we focus on rebalancing cost in terms of time, an alternate commonly-used metric is the *earth-movers distance* (EMD), which measures cost in terms of additional distance traveled for rebalancing. The two definitions are however equivalent up to scaling constants. In our work, we use cost/travel-time/distance interchangeably to represent these costs.

2.2 The Price of Fragmentation

Given the above setting, we can now define our measure for capturing the *price of fragmentation* (PoF): the increase in rebalancing costs in a duopoly as compared to a monopoly, for serving the same demand with exogenous demand splitting. Such a notion captures the price paid by society in terms of additional vehicle miles in order to satisfy the same demand.

We first need some additional notation. Given a demand vector $\{\Lambda_{ij}\}$ in the monopolist setting, our aim is to understand the increase in rebalancing costs due to exogenous splitting of the demand into vectors $\{\Lambda_{ij}^a\}$ and $\{\Lambda_{ij}^b\}$ for firms a and b respectively. Moreover, to capture the growth of MoD platforms, we analyze the system under the so-called *large-market scaling*, wherein we scale all the incoming demands by a factor $\theta \in \mathbb{N}$.

Given the above setting, one way to model the price of fragmentation is to consider the worst case increase in rebalancing costs over all feasible demand splits. We discuss this notion in detail in Appendix C; however, under large-market scaling, it is easy to see that such a worst-case PoF must scale linearly in θ (this follows directly from the homogeneity of $RC(\Lambda)$). This suggests that for a meaningful analysis, we need a more refined notion.

We instead consider a *stochastic model of demand splitting* between the firms. At a high level, such a model captures the idea that when users idiosyncratically choose between the two firms, then on a macro-level, the demand split between the firms in each location can be modeled as a random variable centered at the market share of the respective firms. For ease of presentation, we assume that each Λ_{ij} is an integer for every pair (i, j) , and the demand-splits between firms correspond to Binomial random variables with equal mean (i.e., homogeneous market share) across the network. *Our results however hold for general demand and demand-split distributions as well as heterogeneous market share* (cf. Section 4).

Given firms a, b , we define $\rho \in [0, 1/2]$ to be the market share of firm a , and $1 - \rho$ to be the market share of b . Consequently, for any edge $(i, j) \in E$, we assume that the demand Λ_{ij} is split in two parts, with firm a receiving $\Lambda_{ij}^a \sim \text{Binomial}(\Lambda_{ij}, \rho)$, and firm b receiving $\Lambda_{ij}^b = \Lambda_{ij} - \Lambda_{ij}^a$. We abuse notation to define $\Lambda^a \sim \text{Binomial}(\Lambda, \rho)$ to be a vector such that each coordinate is an iid. rv. $\Lambda_{ij}^a \sim \text{Binomial}(\Lambda_{ij}, \rho)$. Now the increase in rebalancing cost is given by:

$$\gamma \triangleq \mathbb{E}[RC(\Lambda^a) + RC(\Lambda^b)] - RC(\Lambda)$$

Where (Λ^a, Λ^b) denotes the random vector of demands for the two firms. Now, under the *large-market scaling* with demand splits $\Lambda^{a,\theta} \sim \text{Binomial}(\theta\Lambda, \rho)$, and $\Lambda^{b,\theta} = \theta\Lambda - \Lambda^{a,\theta}$ (for $\theta \in \mathbb{N}$), we formally define the *Price of Fragmentation* as follows:

Definition 2.2. (Price of Fragmentation) Given monopolist demand Λ and market-share ρ , the Price of Fragmentation is given by:

$$\gamma^\theta \triangleq \mathbb{E}[RC(\Lambda^{a,\theta}) + RC(\Lambda^{b,\theta})] - RC(\theta\Lambda)$$

To get more insight into the PoF, we note first that given the homogeneity and convexity of $RC(\cdot)$, Jensen's inequality implies that $\gamma \geq 0$, i.e., competition necessarily reduces the system efficiency under stochastic splitting. More precisely, noting that $\mathbb{E}[\Lambda^{a,\theta}] = \rho\theta\Lambda$, we get that:

$$\mathbb{E}[RC(\Lambda^{a,\theta}) + RC(\Lambda^{b,\theta})] \geq RC(\rho\theta\Lambda) + RC((1 - \rho)\theta\Lambda) = RC(\theta\Lambda)$$

Note that since $RC(\cdot)$ is continuous and $\Lambda^{a,\theta}$ is binomial (and hence, has sub-Gaussian tails), the law of large numbers guarantees that $\gamma^\theta/\theta \rightarrow 0$. Our results essentially characterize the rate of this convergence, and in the process, show that depending on the underlying demand vector, the stochastic PoF(θ) = γ^θ either decays to 0 or grows unbounded. We refer to the former as the *fragmentation-resilient* regime, and the latter as the *fragmentation-affected* regime.

To conclude this section, we make some observations:

- From a technical perspective, we note that though we restrict to the case of a duopoly, our results generalize to any number of competitors. However, the monopoly to duopoly shift is in a sense the most drastic in terms of costs.
- Similarly, our results also generalize to heterogeneous market shares, i.e., where the demand at different locations may split differently. Here though, in addition to our two regimes, we obtain a third wherein γ^θ scales linearly with θ . We discuss this in more detail later in the paper.
- Finally note that, in a sense, our model of exogenous demand splitting isolates the effects of firm-level competitive decisions (re. pricing, capacity etc.) in a manner similar to price-theoretic vs. microeconomic models for marketplaces [23]. An advantage of this is that our results only depend on the structure of total demand, rather than particular assumptions of competitive dynamics. Consequently any global demand model accounting for competitive effects can be plugged in to extend the range of applications of our results.

3 WARM UP: PRICE OF FRAGMENTATION IN A TWO-STATION NETWORK

To build some intuition into our main result, we first analyze a simple network with only two nodes. One advantage of this setting is that it allows us to solve for the rebalancing cost in closed form, thereby providing an analytic insight into the system behavior in different regimes. These insights prove useful in establishing the result for general networks in the next section.

In the case where our network only has two stations, we have the following system parameters:

- Two stations $\{1, 2\}$ with inter-station travel times $\tau_{12} = 1$ and $\tau_{21} = \tau \geq 1$
- Monopolist demand vectors $\Lambda_{12} = \lambda$ and $\Lambda_{21} = \mu$,
- Optimal dual variables for each station under the monopolist demand vectors given by $\alpha_1 = 0$ and α_2 ,
- Random demand-splits in the duopoly setting given by
 $\Lambda_{12}^a \sim \text{Binomial}(\theta\lambda, \rho)$, $\Lambda_{12}^b = \Lambda_{12} - \Lambda_{12}^a$,
 $\Lambda_{21}^a \sim \text{Binomial}(\theta\mu, \rho)$, $\Lambda_{21}^b = \Lambda_{21} - \Lambda_{21}^a$

The setting is summarized in the following diagram:

Note that our choice of τ_{12}, τ_{21} is without loss of generality (up to multiplicative scaling). Moreover, the dual constraint implies $-1 \leq \alpha_2 \leq \tau$. Consequently, the dual polytope has two corner points $(0, \tau)$ and $(0, -1)$, and:

$$\begin{aligned} RC(\lambda, \mu) &= \max\{\tau(\mu - \lambda), (\lambda - \mu)\} \\ &= \frac{(\tau - 1)(\mu - \lambda) + (\tau + 1)|\mu - \lambda|}{2} \end{aligned}$$

Given this explicit formula, we can now derive the asymptotic behavior of γ^θ . To do so, we approximate the binomial distribution by a Gaussian distribution¹; in particular, we assume the demand splits obey $\Lambda_{12}^a \triangleq X \sim \mathcal{N}(\rho\theta\lambda, \sigma_X^2)$ and $\Lambda_{21}^a \triangleq Y \sim \mathcal{N}(\rho\theta\mu, \sigma_Y^2)$ with $\sigma_X^2 = \theta\lambda\rho(1 - \rho)$ and $\sigma_Y^2 = \theta\mu\rho(1 - \rho)$. Now we have the following result

¹We make this assumption primarily to make the calculations more transparent, thereby giving some intuition for the general case. Note that this may introduce a bias as the demand may be negative; however, this is a second order effect, and is easily handled by our later theorems, which only require tail bounds.

PROPOSITION 3.1. *For sufficiently large θ , $\exists A_1, A_2 \in \mathbb{R}$ such that:*

- $\lambda = \mu \Rightarrow \gamma^\theta = A_1 \theta^{1/2}$
- $\lambda \neq \mu \Rightarrow \gamma^\theta = A_2 \theta^{-1/2} e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)} \cdot \theta} + o\left(\theta^{-1/2} e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)} \cdot \theta}\right)$

The formal proof is presented in Appendix A; however, we can use this setting to point out some features which will help in understanding the general system in the next section:

- Consider a particular realization of the demand splits – now, if all three firms end up having to rebalance in the same direction (say from 1 to 2), then we get 0 PoF for that sample. Note that this corresponds to all three LPs (for rebalancing under the monopolist demand, and individually for each firm) attaining their maximum value at the same corner point of the dual polytope. In contrast, in the unbalanced case, the loss in efficiency comes from an inversion of the direction of rebalancing flow.
- When $\lambda > \mu$, the difference of demand flows increases linearly with θ , while the probability that $\lambda^a < \mu^a$ decays exponentially. This is the reason why we get an exponential decay in γ^θ , and we see that the higher the discrepancy between input demands, the faster the convergence.
- On the contrary, when $\lambda = \mu$, the decay is much slower because the total demand is perfectly balanced. Consequently, γ^θ accounts for the fluctuations of $X - Y$, which has variance on the order of $\sqrt{\theta}$.

To understand what is the loss in efficiency more formally in a general setting, consider the dual LP (cf. (2)), and a realization of the demand-split $\Lambda^a = X$ and $\Lambda^b = \lambda - X$. Suppose the optimal dual corner point under this realization for the monopoly and each firm is given by α^a , α^b and α^* . Then we have:

$$RC(\Lambda^a) + RC(\Lambda^b) - RC(\Lambda) = \Lambda^a \cdot \alpha^a + \Lambda^b \cdot \alpha^b - \Lambda \cdot \alpha^*$$

If $\alpha^a = \alpha^b = \alpha^*$ then the loss is zero; inefficiency arises when the dual solutions are different. This suggests that the overall loss in rebalancing costs is proportional to the probability of having different dual solutions. We formalize this in the next section.

4 PRICE OF FRAGMENTATION: THE GENERAL NETWORK SETTING

The network with two nodes in the previous sections gives us important intuition to understand the behavior of the stochastic PoF in any general network with given demands and edge costs. In particular, Proposition 3.1 showed that the phase transition in the PoF γ^θ from exponential decay to square-root growth was induced by a change in the optimal corner point of the dual polytope. We will now extend this property to general networks.

We start by defining some mathematical objects we use in our further results. First, we define $\mathcal{E}$ to be the set of corner points of the dual polytope (cf. equation (2)); note that the rebalancing cost $RC(\Lambda)$ is the maximum inner product of Λ over all corner points in $\mathcal{E}$. Next, for each $\alpha \in \mathcal{E}$, we define the following set:

$$C_\alpha = \{ \lambda \text{ s.t. } RC(\lambda) = \alpha^\top \lambda \}$$

PROPOSITION 4.1. *$\forall \alpha \in \mathcal{E}$, C_α is a closed cone. Furthermore, the intersection of two cones $(C_\alpha, C_{\alpha'})$ is a hyperplane.*

PROOF. It is immediate that C_α is a cone due to the homogeneity of $RC(\cdot)$. Secondly, if we define $h(\lambda) = RC(\lambda) - \alpha^\top \lambda$, then it is clear that h is continuous and we have $C_\alpha = h^{-1}(\{0\})$ – this verifies that it is closed. Finally, for distinct $(\alpha, \beta) \in \mathcal{E}$ consider $\lambda \in C_\alpha \cap C_\beta$ – we have $RC(\lambda) = \alpha^\top \lambda = \beta^\top \lambda$, thus $(\alpha - \beta)^\top \lambda = 0$, which is the equation of a hyperplane. $\square$

Fig. 2. In Fig. 2(a) we demonstrate the geometry of the *fragmentation-resilient regime* (i.e., where γ^θ undergoes exponential decay), while in Fig. 2(b) we do the same for the *fragmentation-affected regime* (i.e., where $\gamma^\theta = \Omega(\sqrt{\theta})$). In both plots, the straight dashed line represents the expected demand for the first firm as demand scaling θ increases, while the two solid circles each represent the expected demand under two scaling values $\theta_1 < \theta_2$. Note that the expected demand line is within the dual cone in Fig. 2(a), while it is on the cone boundary in Fig. 2(b). The dotted circles around the expected value represent a high-confidence ball for the realized demands; the radii of the balls are based on the variance of the rescaled demand r.v. λ^θ/θ , which decreases as we scale θ . Due to random fluctuations, the probability of having an extra loss is high in Fig. 2(a), inducing a loss of the same order of the fluctuations.

These properties now allow us to derive a nice geometric interpretation of what happens when we scale the demand to infinity; we depict this in Figures. 2(a) and 2(b) respectively for the fragmentation-resilient and fragmentation-affected regimes. The solid circle shows the average demand $\mathbb{E}[\Lambda^{a,\theta}] = \rho\theta\Lambda$, which moves along a line as we scale θ ; the circles depict a high-probability envelope for the (normalized) realized demand $\Lambda^{a,\theta}/\theta$. The radius of the envelope decreases as we rescale by θ .

In the fragmentation-resilient setting, the expected demand remains on a line inside a given dual cone, while in the fragmentation-affected regime, it lies on a hyperplane defined by the intersection of two cones. Since ρ is homogeneous, the expected demands after splitting for both firms in the duopoly also remain on the same line upon scaling. The true demands however have random fluctuations around its expectation, but since the cone becomes wider while scaling, while the demand distributions concentrate (assuming well-behaved tails), the probability of escaping the cone decays fast.

On the other hand, if the optimal dual solution is degenerate, then the expected regime lies in the intersection of 2 or more cones. Now under scaling the expected monopolist demand vector remains on the hyperplane, but despite concentrating about its expected value, the probability of having each duopolist firm's realized demand-split fall on either side of the boundary remains significant.

We now state our main results, which formalize these intuitions. Hereafter, for any set $\mathcal{A}$, $\mathring{\mathcal{A}}$ denotes its interior. First, we characterize the exponential decay regime:

THEOREM 4.2. (Fragmentation-resilient regime) *Given network G , monopolist demand vector Λ and demand-split ρ , suppose that the following conditions hold:*

- $\Lambda \in \mathring{C}_\beta$ for some $\beta \in \mathcal{E}$ (i.e., the monopolist demand vector lies in the interior of a dual cone)
- There is a random vector Z such that $\Lambda^{\theta,a} = \rho\theta\Lambda + \sqrt{\theta}Z$ with $\mathbb{E}[Z_e] = 0 \ \forall e \in E$ and $\mathbb{E}[|Z|_1] < \infty$
- There is a function f such that $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(f(t))$

Then $\gamma^\theta = O\left(\theta f\left(\sqrt{\theta}\right)\right)$

In the particular case of binomial demand splits, note that as long as Z_i has sub-Gaussian tails, $\gamma^\theta = O(\theta e^{-c\theta})$ for some constant c .

PROOF. Define $x = \mathbb{E}[\Lambda^{\theta,a}] = \rho\Lambda$ and $y = \mathbb{E}[\Lambda^{\theta,b}] = (1-\rho)\Lambda$. As we said earlier, the boundaries between dual cones are hyperplanes $\mathcal{P}_{(\alpha,\beta)}$, $\forall (\alpha,\beta) \in \mathcal{E}$. Thus we can define the orthogonal projections of x and y on each of those boundaries. Since there is a finite number of corner points, a strictly positive minimum distance is reached over all planes. We define the minimum distance between the expected demand vectors of both companies and the cone's boundary as:

$$\delta = \min \left\{ \min_{\alpha \in \mathcal{E}} d(x, \mathcal{P}_{(\alpha,\beta)}), \min_{\alpha \in \mathcal{E}} d(y, \mathcal{P}_{(\alpha,\beta)}) \right\}$$

Moreover, since $\Lambda \in \mathring{C}_\alpha$, we have that the Euclidean balls around these points are contained in the cone C_α , i.e., $B(x, \delta) \subset C_\alpha$ and $B(y, \delta) \subset C_\alpha$. Note that for any demand-split, we have $\lambda \in B(x, \delta)$ if and only if $\Lambda - \lambda \in B(y, \delta)$.

When we scale the demand vectors, since we use orthogonal projections and the distance is homogeneous, then δ is also homogeneous with respect to θ . Thus, if our random demand-splits lie within a $\theta\delta$ ball around their expected values, we preserve the same corner point and undergo no loss in efficiency. To find an upper bound on the loss of efficiency when our random variable is outside those spheres, we define $\tilde{\beta} \triangleq \max_{\beta \in \mathcal{E}} \|\beta\|_\infty$.

We also denote α_1, α_2 and β the corner points such that $RC(\Lambda^{\theta,a}) = \alpha_1^\top \Lambda^{\theta,a}$, $RC(\theta\Lambda - \Lambda^{\theta,a}) = \alpha_2^\top (\theta\Lambda - \Lambda^{\theta,a})$, and $RC(\theta\Lambda) = \beta^\top \theta\Lambda$. What is important to notice is that the corner points α_1 and α_2 implicitly depend on Z and are consequently random variables.

Now, we can bound the normalized loss in rebalancing costs under scaling as follows:

$$\begin{aligned} \gamma^\theta &= \mathbb{E} \left[\alpha_1^\top \Lambda^{\theta,a} + \alpha_2^\top (\theta\Lambda - \Lambda^{\theta,a}) - \beta^\top \theta\Lambda \right] \\ &= \theta(\rho\alpha_1 + (1-\rho)\alpha_2 - \beta)^\top \Lambda + \sqrt{\theta} \mathbb{E}[(\alpha_1 - \alpha_2)^\top Z] \\ &\leq 2\theta\tilde{\beta} \|\Lambda\|_1 \mathbb{P} \left[\exists e \in E \text{ s.t. } \sqrt{\theta}|Z_e| > \theta\delta \right] + \sqrt{\theta} \|\alpha_1 - \alpha_2\|_\infty \mathbb{E} \left[\|Z\|_1 \cdot \mathbb{1}_{\{\exists e \in E \text{ s.t. } \sqrt{\theta}|Z_e| > \theta\delta\}} \right] \\ &\leq \left(2\theta\tilde{\beta} \|\Lambda\|_1 + 2\tilde{\beta} \sqrt{\theta} \mathbb{E} \left[\|Z\|_1 \mid \exists e \in E \text{ s.t. } |Z_e| > \sqrt{\theta}\delta \right] \right) \times \mathbb{P} \left[\exists e \in E \text{ s.t. } |Z_e| > \sqrt{\theta}\delta \right] \end{aligned}$$

The first inequality follows from applying Hölder's inequalities, and the requirement that for positive PoF, the demand-split vector must lie in a cone other than $\mathring{C}_\beta$. Now we want to prove that the expectation inside the brackets in the last inequality is finite. For this, recall we assumed that $\mathbb{E}[|Z|_1]$ is finite; furthermore, due to the law of total expectation, and because the expectations are positive, the finiteness of $\mathbb{E}[|Z|_1]$ implies the finiteness of any conditional expectation on $|Z|_1$, and thus the one we want. The only remaining point is to upper bound the probability of having Z out of the ball. Via a union bound, we get:

$$\mathbb{P} \left[\exists e \in E \text{ s.t. } |Z_e| > \delta\sqrt{\theta} \right] \leq N^2 \mathbb{P} \left[|Z_e| > \delta\sqrt{\theta} \right] = O(f(\sqrt{\theta}))$$

Combining the inequalities, we get the desired result:

$$\gamma^\theta = O\left(\theta f(\sqrt{\theta})\right)$$

□

Similarly, we can also characterize the square-root decay regime:

THEOREM 4.3. (Fragmentation-affected regime) Given network G , monopolist demand vector Λ and demand-split ρ , suppose that the following conditions hold:

- $\exists \alpha \neq \beta \in \mathcal{E}$ such that $\Lambda \in C_\alpha \cap C_\beta$ (i.e., the monopolist demand vector lies in the hyperplane defined by the intersection of two or more dual cones)
- There is a r.v. ξ such that $\Lambda^{\theta,a} = \rho\theta\Lambda + \sqrt{\theta}Z$ with $\mathbb{E}[Z_e] = 0 \ \forall e \in E$ and $\mathbb{E}[\|Z\|_1] < \infty$
- $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(f(t))$ for some $f(t) = O(t)$

Then $\gamma^\theta = \Theta(\theta^{1/2})$

PROOF. First we show the lower bound. Note that in this setting, we have $RC(\Lambda) = \alpha^\top \Lambda = \beta^\top \Lambda$ due to the hypothesis of lying on the interface of two cones. Furthermore, $\alpha^\top \Lambda^{\theta,a} \geq \beta^\top \Lambda^{\theta,a}$ if and only if $\beta^\top (\theta\Lambda - \Lambda^{\theta,a}) \geq \alpha^\top (\theta\Lambda - \Lambda^{\theta,a})$. Due to this, it is sufficient to separate two cases depending on which of these two corner points gives the best value. Furthermore, ignoring all other corner points gives a lower bound because the definition of the cones is such that changing to another corner point means that this third point yields a higher score. Thus we have:

$$\begin{aligned} \gamma^\theta &= \mathbb{E} \left[RC(\Lambda^{\theta,a}) + RC(\theta\Lambda - \Lambda^{\theta,a}) - RC(\theta\Lambda) \right] \\ &\geq \mathbb{E} \left[(\alpha^\top \Lambda^{\theta,a} + \beta^\top (\theta\Lambda - \Lambda^{\theta,a}) - \theta\beta^\top \Lambda) \cdot \mathbb{1}_{\alpha^\top Z \geq \beta^\top Z} + (\beta^\top \Lambda^{\theta,a} + \alpha^\top (\theta\Lambda - \Lambda^{\theta,a}) - \theta\alpha^\top \Lambda) \cdot \mathbb{1}_{\beta^\top Z \geq \alpha^\top Z} \right] \\ &= \theta^{1/2} \cdot \mathbb{E} \left[((\alpha - \beta)^\top Z) (\mathbb{1}_{\alpha^\top Z \geq \beta^\top Z} - \mathbb{1}_{\beta^\top Z \geq \alpha^\top Z}) \right] \\ &= \theta^{1/2} \mathbb{E} [|(\alpha - \beta)^\top Z|] = \Omega(\theta^{1/2}) \end{aligned}$$

Thus the first part of the proof is finished.

Next we focus on the upper bound of γ^θ . The proof will reuse ideas of the previous proof. In fact, let's separate α and β of the other corner points. The minimum radius δ such that a ball of radius δ can fit in $C_\alpha \cup C_\beta$ is well defined and is strictly positive. Thus, we can separate the rebalancing loss in two parts such that:

$$\begin{aligned} \gamma^\theta &= \mathbb{E} \left[\left(RC(\Lambda^{\theta,a}) + RC(\Lambda^{\theta,b}) - RC(\theta\Lambda) \right) \left(\mathbb{1}_{\{\forall e \in E, Z_e \leq \sqrt{\theta}\delta\}} + \mathbb{1}_{\{\exists e \in E \text{ s.t. } Z_e > \sqrt{\theta}\delta\}} \right) \right] \\ &\leq \theta^{1/2} \cdot \mathbb{E} \left[|(\alpha - \beta)^\top Z| \mathbb{1}_{\{\forall e \in E, Z_e \leq \sqrt{\theta}\delta\}} \right] + O(f(\sqrt{\theta})) \\ &\leq \theta^{1/2} \cdot \mathbb{E} [|(\alpha - \beta)^\top Z|] + O(f(\sqrt{\theta})) = O(\theta^{1/2}) \end{aligned}$$

If there are more than two corner points the proof is similar but has to be slightly modified. The lower bound is the same by considering only two of those points. For the upper bound, there are still two parts by defining the ball fitting in all cones. One part decays exponentially, and the second one can be upper bounded by summing the absolute value of the random variable $(\alpha - \beta)^\top Z$ over all pairs of corner points which are optimal for Λ . □

We state Theorems 4.2 and 4.3 in greater generality than required, so as to admit different distributions for random demand-splitting in the duopolist setting. For the setting of constant Λ and binomial splitting we defined earlier, we have that Z is sub-Gaussian, i.e., $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(e^{-at^2})$ for some constant a ; thus we get an exponential decay when the

demands are unbalanced. If instead Λ_{ij} comes from a Poisson process with mean $\mathbb{E}[\Lambda_{ij}] = \lambda_{ij}$, and furthermore, under scaling, its expectation is $\theta\lambda_{ij}$. In this case, assuming probabilistic demand splitting, the random demand experienced by each firm follows a Poisson distribution with mean and variance linear in θ ; note that this distribution has sub-Exponential tails, i.e., $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(e^{-at})$ for some constant a . The decay in γ^θ is now slower than the binomial case; in particular, when the demand is unbalanced, we have $\gamma^\theta = O(\theta e^{-a\sqrt{\theta}})$.

Till now we considered the case of homogeneous market share ρ ; however, we can easily extend our results to the case of inhomogeneous market shares percentage ρ . This generalization is interesting as market share can be neighborhood-dependent in practice. For instance, different parts of the city support different demographics, and depending on this, one firm might be preferred over the second one due to its quality of service or cheaper price. To capture this, we consider $\rho = \{\rho_{ij}\}$ to be a vector, and define the demand-split for firm a as:

$$\mathbb{E}[\Lambda^{\theta,a}] = \theta\rho \odot \Lambda \quad \mathbb{E}[\Lambda^{\theta,b}] = \theta(1 - \rho) \odot \Lambda$$

where $\odot$ is the Schur (or element-wise) product of ρ and Λ .

Due to the inhomogeneity in splitting, we can only take advantage of the homogeneity of $RC(\cdot)$ with θ . Geometrically, this means that the expectation of the split demands no longer lie on the same line as the monopolist demand. Now let α , β and η denote the optimal corner points of the dual polytope for demand vectors $\rho \odot \Lambda$, $(1 - \rho) \odot \Lambda$ and Λ respectively – note that unlike before, these corner points can now be different. However, we can still characterize the PoF as before, via the following theorem:

THEOREM 4.4. *Given network G , demand vector Λ and demand-split vector ρ , assume the following hold:*

- $\Lambda \in C_\eta$ for some $\eta \in \mathcal{E}$ (i.e., the monopolist demand vector lies in the interior of a dual cone)
- There is a r.v. Z such that $\Lambda^{\theta,a} = \theta\rho \odot \Lambda + \sqrt{\theta}Z$ with $\mathbb{E}[Z_e] = 0 \forall e \in E$ and $\mathbb{E}[||Z||_1] < \infty$
- $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(f(t))$ with $f(t) = O(t)$

Now suppose we define:

$$L = \alpha^\top \rho \odot \Lambda + \beta^\top (1 - \rho) \odot \Lambda - \eta^\top \Lambda$$

Then we have:

- $\rho \odot \Lambda \in \dot{C}_\alpha$ and $(1 - \rho) \odot \Lambda \in \dot{C}_\beta \Rightarrow \gamma^\theta = L\theta + O(\theta f(\sqrt{\theta}))$
- $\rho \odot \Lambda \in C_\alpha \setminus \dot{C}_\alpha$ or $(1 - \rho) \odot \Lambda \in C_\beta \setminus \dot{C}_\beta \Rightarrow \gamma^\theta = L\theta + \Theta(\theta^{1/2})$

PROOF. The proof is the along the same lines as the homogeneous case, except that each companies' bounds have to be treated separately. The details are given in Appendix A. $\square$

Taken together, the above results thus completely characterize the decay rates of γ^θ . Note that in the general heterogeneous case, we actually get three regimes for γ^θ : if $L = 0$, then we get the exponential decay and square root growth as before; in addition, if $L > 0$, then we now get a *linear growth regime*. Also, note that all the above proofs treats each company separately and then sums to get the required bounds. Thus, the results can easily be adapted for more than 2 firms, leading to similar conclusions. For sake of simplicity we stick to the case of 2 firms, but the general result is presented as Theorem A.4 in Appendix A.

The above results thus indicate that the PoF undergoes a phase transition, going from converging to 0 to diverging to ∞ as the monopolist demand vector Λ approaches a boundary hyperplane between two dual cones. From an operational point of view, it is thus critical to understand conditions for diverging PoF; in particular, we desire conditions which are more intuitive rather than our characterization in terms of dual degeneracy. We now try to obtain such a condition.

Consider a setting with demands $\lambda_i = \sum_j \lambda_{ji} - \lambda_{ij}$, and let (x^*, α^*) denote a pair of optimal solutions for the primal and dual respectively; then by complementary slackness, we have that $x_{ij}^* > 0 \Rightarrow \alpha_i^* - \alpha_j^* = \tau_{ij}$. Now consider subgraphs H and $\hat{H}$ be two subgraphs defined by the set of all nodes V and edges:

$$E_H = \{(i, j) \mid x_{ij} > 0\} \quad E_{\hat{H}} = \{(i, j) \mid \alpha_i^* - \alpha_j^* = \tau_{ij}\}$$

H thus represents the support of positive rebalancing flows for the monopolist setting; from the non-negativity of τ_{ij} , it is easy to observe that E_H forms a directed acyclic graph. $\hat{H}$ is the subgraph defined by tight dual constraints. By complementary slackness, we have $E_H \subset E_{\hat{H}}$. We now have the following theorem that characterizes dual degeneracy in the rebalancing cost LP (1).

THEOREM 4.5. *For the optimal rebalancing solution to be dual degenerate, a necessary condition is that E_H contains at least two connected components. Moreover, this is sufficient with probability 1 under small random perturbations of the edge-weights τ_{ij} .*

PROOF. First, suppose E_H has a single connected component. Now by complementary slackness, the dual variables α are uniquely defined by: $\forall (i, j) \in E_H, \alpha_i - \alpha_j = \tau_{ij}$. This shows the necessity of our condition.

Next, note that under random perturbations of the edge weights τ_{ij} , we have $E_H = E_{\hat{H}}$ with probability 1. Now suppose that E_H (and hence $E_{\hat{H}}$) comprises of multiple connected components. We now show that we can extend it without modifying the LP value or the saturated inequalities.

Let A be one of the connected components and $\delta(A)$ denote the edges in the cut of A ; we define $\epsilon = \inf_{(i,j) \in \delta(A)} |\alpha_i - \alpha_j - \tau_{ij}| > 0$ (by definition of $E_{\hat{H}}$). We now consider a new dual corner point defined by $\beta_i = \alpha_i + \epsilon$ if $i \in A$ and $\beta_i = \alpha_i$ otherwise. This has no incidence with respect to the saturated inequalities, and it will saturate one inactive inequality thanks to its definition. If (i, j) are in A then the argument of the uniform translation applies (see proposition A.1 in Appendix A). Consequently, this new edge is now active and connects A to another connected component. Furthermore, the score for β is the same as for α since we have:

$$\sum_i \alpha_i \lambda_i = \sum_{(i,j)} (\alpha_i - \alpha_j) x_{ij}^* = \sum_{(i,j) \in E_H} (\alpha_i - \alpha_j) x_{ij}^*$$

Thus $\alpha_i - \alpha_j \neq \beta_i - \beta_j \Leftrightarrow x_{ij}^* = 0$, and this modification has no impact on the score. Finally, since we have more constraints than necessary to uniquely characterize a corner point, this means that two corner points are optimal for the dual LP. $\square$

Theorem 4.5 gives some intuition into when dual degeneracy happens; more importantly, it gives an easily testable condition. Note though that for a setting to be fragmentation-affected, we need to display two disconnected components in the rebalancing flows, but which has non-zero demand flowing between the components (otherwise, the underlying demand itself is disconnected!). We next use this to validate our results based on simulation of synthetic and real-world settings.

5 SIMULATIONS AND EXPERIMENTS

We first use synthetic experiments to numerically verify our phase transition theorems; for this, we ran simulations on a simple four nodes network, which allows us to easily create instances which are fragmentation-resilient and fragmentation-affected. Next, we look at the NYC taxi data, and use Theorem 4.5 to classify the fraction of instances falling in each class – surprisingly, we detect the fragmentation affected on a significant fraction of the data. We also perform several robustness checks to validate our findings.

5.1 Synthetic Experiments

First, we experimentally study our characterization of the phase transition in the PoF in terms of dual degeneracy of $RC(\Lambda)$. Although the two-nodes network we used in Section 3 already allowed us to exhibit both regimes, it has a very simple degenerate subset wherein $\Lambda_{ij} = \Lambda_{ji}$. In order to build more complex synthetic examples that exhibit both regimes, we consider a four node network depicted in Fig. 3, comprising of 4 nodes, with inter-node average travel-times as indicated². By varying the monopolist demands between the nodes, we use the network to verify the conditions for the two regimes given in Theorem 4.5.

Recall that Λ_i denotes the net *inflow* of demand into node i . To generate instances corresponding to the two regimes, we consider two different demand vectors, $\Lambda^1 = [2, 3, -4, -1]$ and $\Lambda^2 = [2, 3, -3, -2]$; note that this is sufficient for solving $RC(\Lambda)$, and we do not need to specify Λ_{ij} completely. In both settings, the relevant inter-node travel-times are $\tau_{13} = 1$, $\tau_{14} = 3$, $\tau_{23} = 2$, $\tau_{24} = 5$ (the rest does not matter for solving $RC(\Lambda)$ for the given demands). These parameters, as well as the support of the corresponding optimal rebalancing flows are indicated by bold arrows in Figures 3(a) and 3(b). Note that in the former case, the rebalancing flows comprise a single connected component, and hence the solution is not dual degenerate (via Theorem 4.5); in contrast, for Λ^2 , optimal rebalancing flow comprises of 2 connected components, and hence is dual-degenerate.

Fig. 3. Four node graph with nodes $\{1, 2, 3, 4\}$ used for synthetic experiments. The dashed arrows indicate the direction of net flow of passengers at a node (i.e., entering or exiting; note Λ_i corresponds to total demand flow exiting i). Fig. 3(a) has demand vector $\Lambda^1 = [2, 3, -4, -1]$, while Fig. 3(b) has demand vector $\Lambda^2 = [2, 3, -3, -2]$. The relevant edge costs τ_{ij} (given in the figures) are the same in both settings. The support of the optimal rebalancing flows are indicated by bold arrows in both settings. Note that in the first case, edges with positive rebalancing flow define a single connected component, while in the second, they define two disconnected components $\{1, 3\}$ and $\{2, 4\}$.

In each case, to compute γ^θ / θ for each θ , we simulate 500 instances of our random splitting process, and aggregate the random rebalancing costs.

Fig. 4(a) and 4(b) show the scaling behavior of the PoF (in terms of $\log(\gamma^\theta / \theta)$ vs. $\log \theta$ in the two settings. The green curve plots PoF estimates obtained from simulations, while the blue curve is a reference line of slope -0.5 . Note that in the dual-degenerate demand case, the asymptotic regime indeed shows square-root growth, while in the other regime, the decay is sub-polynomial.

²We only indicate travel-times between nodes $(1, 3)$, $(1, 4)$, $(2, 3)$ and $(2, 4)$ as these are the only ones relevant to the circulation LP for the demands we consider.

Fig. 4. Simulation showing Price of Fragmentation in 4-node network in Fig. 3: The above plots show scaling behavior of the (normalized) increase in rebalancing costs λ^θ / θ in the settings described above. In particular Fig. 4(a) corresponds to the fragmentation-resilient setting (Fig. 3(a)), while Fig. 4(b) corresponds to the fragmentation-affected setting (Fig. 3(b)).

In addition, we also compare PoF under Binomial and Poisson original demands (and random splitting) under the fragmentation-resilient setting in Fig. 5. Here we see that the latter decays slower than the former, as suggested by our theoretical results in Theorem 4.2.

Fig. 5. Comparison of Binomial and Poisson demand splitting in the fragmentation-affected regime (Fig. 4(a)).

5.2 Experiments on NYC Dataset

While we are able to experimentally demonstrate the two regimes in synthetic networks, the question remains as to whether both regimes are only witnessed in carefully engineered examples (in particular, since fragmentation-affected regimes occur under dual-degenerate demands, which may be unlikely in practice). To answer this question, we extend our experimental analysis to a simplified network model of New York City, with demand data from the New York City Taxi Dataset [1] as a proxy for the actual system demand.

Experimental setup: Our experiments use taxi data recorded for the months of May and June 2016. The dataset has about 335,000 trips per day; on an hourly basis, the number of trips ranges from 3,600 to 20,000 with an average number of around 14,000 thousands trips per hour. We model the city as consisting of a finite number of pickup and drop-off *stations* (20-80 in our experiments),

where we assume each trip begins and ends. The stations are determined by taking all of the user demand for the full period of study, clustering the data (pickup and drop-off locations) into the desired number of stations (clusters), and picking the centroid for each cluster. For each time window of interest, the demands are then aggregated to the total demand per origin-destination station pair Λ_{ij} . For convenience, we use the distance d_{ij} as a proxy for the travel-times τ_{ij} to parametrize edge-costs; we compute the Manhattan distance between the stations to get the distances d_{ij} . The pickup-dropoff station locations are shown in Fig. 10 in Appendix D.

Fig. 6. Detailed plots for simulated data shown in Fig. 1. Fig. 6(a) shows PoF scaling for a wider range of θ , allowing us to observe the asymptotic behavior; here the blue line corresponds to the fragmentation-affected regime, and the green to the fragmentation-resilient regime. Fig. 6(b) displays the same data in log-log scale, allowing us to observe precisely the scaling.

(a) Effect of number stations on affected regime probability (b) Effect of time window on affected regime probability

Fig. 7. Probability of observing fragmentation-affected demand in NYC based on the NYC taxi dataset. Fig. 7(a) shows the influence of the number of stations, while Fig. 7(b) demonstrates the influence of the length of the time windows. Statistics plotted are based on aggregation of the data into time windows of length 2 hours.

Experimental Findings: In Figs. 6(b) and 6(a), we show the simulation data from Fig. 1 in more detail; this provides a clear indication that both PoF regimes can be observed using demand distributions estimated from the NYC data. In particular, the log-log plot in Fig. 6(a) clearly demonstrates

the square-root growth and exponential decay. Note also that the two instances are very similar in terms of time (12pm vs. 1pm) and PoF for $\theta = 1$; the difference really manifests under scaling.

Since we have established that the two regimes only depend on the structure of the underlying monopolist demand flows (and is testable via Theorem 4.5), this allows us to quantify how often the fragmented-affected regime occurs in the NYC data. One issue in performing such an analysis is the choice of different estimation parameters, in particular, number of stations and size of aggregation time window. To this end, in Figs. 7(a) and 7(b), we plot the probability of observing the fragmentation-affected regime (along with the 95 percentile error bounds) for different time windows, and varying number of stations, using the data for the months of May and June 2016. We note that we compute the probability of dual-degeneracy conditional on having a connected demand graph (this turns out to be the case in almost all settings).

In Fig. 7(a), we observe that the probability of observing dual-degenerate demands increases significantly as we increase the number of stations. This can be partially explained by the following heuristic argument: in a system with N stations, the demand is scattered over N^2 edges. Consequently, the sparsity of the demand matrix increases quadratically, and when solving the LP, fewer edges are active, leading to more disconnected components (and therefore more fragmented-affected regimes, using Theorem 4.5). Increasing the time window to aggregate more demands reduces the likelihood of this phenomenon, but does not eliminate it. In practice, ridesharing systems currently pickup and drop-off passengers at a spatial granularity corresponding to each street corner, which is approximately 4000 stations for NYC; based on the trend in Fig. 7(a), this may imply a very high probability of observing the fragmented affected regime.

One concern regarding the above discussion is that increasing number of stations also decreases the average inter-station travel cost; hence the overall effect on PoF is unclear. To this end, we also directly computed the PoF scaling behavior of the system as a function of the number of stations, using a fixed demand (again corresponding to 1-2pm on the 10th of May 2016). In Fig. 8, we plot this scaling corresponding to [20, 40, 60, 80] stations. We observe that the setting appears to be fragmentation-resilient when aggregated using 20 and 40 stations, but fragmentation-affected for higher resolutions. We note though that this trend does not hold for all time slots, and other trends (e.g., being fragmentation-resilient with 20 and 60 stations, but not the others) can also be observed in the data; however, the overall trend (in Fig. 7(a)) seems to suggest that greater resolutions lead to greater probability of observing fragmentation-affected regimes.

(We note also that the initial PoF, i.e., for $\theta = 1$, increases as the spatial granularity increases. This is a numerical artifact caused by demands whose pick-up and drop-off locations are in the same cluster decreasing as the spatial granularity is refined – recall that a station based model assumes no loss for demands that start and end at the same station.)

Another variable that impacts the aggregation of data, and thus the likelihood of degeneracy, is the length of the time window, as shown in Fig. 7(b) (with number of stations fixed to 40). One interpretation for this is that it is caused due to a discretization phenomenon – shorter time windows may lead to sparser entries in the rebalancing LP, due to which there may be a higher probability of obtaining a dual-degenerate rebalancing flow. In any case, since the overall trend (higher probability of fragmentation-affected regimes when moving to smaller time slots) suggests that there indeed is a non-trivial probability of observing fragmentation-affectedness in the NYC data, no matter how we slice it.

Finally, we explore the effect of adding more firms to the ecosystem, as well as the effect of having inhomogeneous market shares; these plots are provided in Appendix D. Adding more firms leads to the similar behavior as with two firms (this confirms what we show in Theorem A.4). Inhomogeneous market shares however cause a significant increase in the PoF, and as predicted

Fig. 8. Influence of the number of stations on PoF scaling (based on 1pm-2pm data for May 10).

in Theorem A.3, exhibit a linear scaling (See Fig. 9(b)). Furthermore the linear divergence gives PoF 10 to 100 times bigger compared to homogeneous shares when you scale. This shows that homogenizing demand might be a key aspect for policy makers to try and control, if the firms indeed have spatially varying market shares.

6 INSIGHTS AND EXTENSIONS

We now briefly discuss some insights we obtain from our theoretical and simulation results, and how they may be used to inform and shape public policy for regulating ridesharing ecosystems.

Insights: At a high level, our results suggest that fragmentation in an MoD ecosystem incurs minimal additional cost as long as the underlying demands are not degenerate with respect to the dual of the minimum cost rebalancing LP. When the demand matrix is degenerate, however, we get the fragmentation-affected regime, wherein we have $\text{PoF}(\theta) = \Theta(\theta^{1/2})$; here, even though there is zero normalized loss between the duopoly and the monopoly in the limit, the fluctuations are such that the actual loss diverges when scaling θ . Theorem 4.5 gives us a testable way of checking for this setting – moreover, it provides an intuitive condition for this in terms of the existence of *locally balanced clusters* – these correspond to subsets of the network wherein all nodes are relatively close by, and which have large balanced flows across the cut³.

Numerically, we observe that the degeneracy condition is sensitive to the spatio-temporal discretization of the demand, i.e. the number of stations and time window. In particular, the probability of degeneracy appears to increase with the spatio-temporal granularity of the model. One important insight from this is that *aggregating pickup and drop off locations, and increasing the flexibility of pickup times* can have a large impact on the performance of a fragmented ridesharing system. Practically speaking, cities do have policy/regulatory tools at their disposal that could help in this direction. For example, given that cities control curbside access, they can designate dedicated pickup and drop off locations (stations) that are spaced out in a manner that decreases the probability of degeneracy. Ridesharing companies can also provide incentives to encourage customers to walk to stations that help with the rebalancing problem. On the temporal aggregation side, incentives can be provided to passengers who have flexibility in their pickup times.

Another important insight of our results is regarding the importance of *multi-homing*, i.e., drivers who work for multiple firms simultaneously. To see the benefit of multi-homing, note that the

³For example, consider the flow of traffic between San Francisco and Berkeley across the Bay Bridge – there is a significant cost in traversing between the two, and moreover, a roughly balanced flow of people traveling in both directions.

average fluctuations in our model are on the order of $\sqrt{\theta}$; now if a constant fraction of drivers are multi-homing, there is always sufficient flexible supply available to prevent the need for rebalancing. Multi-homing already occurs with many ridesharing drivers in the US simultaneously being active on both the Lyft and Uber apps. The situation can be further improved by cities mandating that drivers must be allowed to operate across multiple platforms and creating more awareness about the possibility of doing so. Moreover, with the advent of driverless fleets, our work suggests that it may be desirable to require that MoD platforms continue to have a significant number of multi-homing drivers in their fleets.

Another insight from our analysis is that inhomogeneous market shares may lead to much worse PoF (see Fig. 9(b)). Thus, settings where different firms target different neighborhoods can lead to an increased probability of degeneracy. Unfortunately, this is a difficult problem to tackle via a policy/regulatory approach. One thought is that cities can impose a geographic equity requirement on providers that could help with this. An alternate option is to set up a centralized clearinghouse that controls a fraction of the flow (for example, a third-party app run by the city or local transit agency), and can use this to make demand-splits homogeneous over the city. Moreover, the clearinghouse can directly aim to eliminate the degeneracy condition, if it has complete knowledge about all the service providers in the market.

Finally, we note that the conditions we obtain for fragmentation-affectedness are in a sense the very conditions under which *mass public transit is efficient* – people wanting to travel between some two distant locations, in large numbers, and with equal numbers traveling in both directions. Thus, in a sense, our results provide more support for having good public transit.

Extensions: Though our work presents a simple stylistic model for formally thinking about the PoF in ridesharing systems, it suggests many natural extensions. The main open problem is to incorporate strategic dynamics between the platforms. In particular, a natural question is if we can exploit competitive dynamics to affect a change in the PoF; an extreme version of this is the above idea of clearinghouses, which pools some fraction of the demand and then allocates it across companies. Next, though our work focused only on the cost of rebalancing, one can raise the question of how our results change when we consider other controls which increase flexibility (in particular, reservation mechanisms and ride-pooling). Finally, any new policy that reduces rebalancing costs is meaningful only if it is economically viable. In essence, what we want is a system that can coordinate global demand and global fleet, but also sustains a competitive market. This can be modeled as a Stackelberg game framework, and ideas from existing models that incorporate queuing dynamics [12, 13] may prove useful in our setting.

7 CONCLUSIONS

While competition between multiple service providers leads to better outcomes for consumers in general, our work here shows that demand fragmentation across multiple platforms in MoD markets can lead to significant losses in terms of operating costs. To this end, we propose the *Price of Fragmentation* (PoF) to quantify the increase in operational costs in a duopoly vs. a monopoly under stochastic demand splitting, and demonstrate (both theoretically and experimentally) that the PoF undergoes a phase transition based on the underlying demand. More precisely, we show that the system admits two dramatically different behaviors when the demand is scaled to infinity: one with an asymptotic loss of zero, and another with an infinite asymptotic loss. These results remain valid under inhomogeneous market shares, and with more than two firms.

Our work is an initial treatment of a subject that is of considerable interest with the increasing importance of MoD systems and their potential for being part of a sustainable transportation ecosystem for urban areas in the future. Furthermore, the losses that we observe should be even more pronounced in MoD systems with ridepooling, where multiple demands are served by the

same vehicle and demand aggregation is critical for efficiency. Given that merging all service operators to create a monopoly is not a viable solution (for many reasons), there is a need for developing new mechanisms that can mitigate this efficiency loss. This work aims to motivate future research in that direction.

ACKNOWLEDGMENTS

SB acknowledges support from Army Research Office grant W911NF-17-1-0094.

REFERENCES

- [1] 2017. NYC Taxi Data. (2017). http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml
- [2] Melika Abolhassani, Mohammad Hossein Bateni, Mohammad Taghi Hajiaghayi, Hamid Mahini, and Anshul Sawant. 2014. Network cournot competition. In *International Conference on Web and Internet Economics*. Springer, 15–29.
- [3] Javier Alonso-Mora, Samitha Samaranyake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences* (2017), 201611675.
- [4] Elliot Anshelevich and Shreyas Sekar. 2015. Price Competition in Networked Markets: How do monopolies impact social welfare?. In *International Conference on Web and Internet Economics*. Springer, 16–30.
- [5] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. 2017. Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 517–517.
- [6] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. 2015. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 639–639.
- [7] Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. 2018. The Value of State Dependent Control in Ridesharing Systems. (2018).
- [8] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilklic. 2014. Cournot competition in networked markets.. In *EC*. 733.
- [9] Anton Braverman, JG Dai, Xin Liu, and Lei Ying. 2016. Empty-car routing in ridesharing systems. *arXiv preprint arXiv:1609.07219* (2016).
- [10] Desmond Cai, Subhonmesh Bose, and Adam Wierman. 2017. On the role of a market maker in networked cournot competition. *arXiv preprint arXiv:1701.08896* (2017).
- [11] David K George. 2012. *Stochastic Modeling and Decentralized Control Policies for Large-Scale Vehicle Sharing Systems via Closed Queueing Networks*. Ph.D. Dissertation. The Ohio State University.
- [12] Refael Hassin and Moshe Haviv. 2003. *To queue or not to queue: Equilibrium behavior in queueing systems*. Vol. 59. Springer Science & Business Media.
- [13] David Lingenbrink and Krishnamurthy Iyer. 2017. Optimal Signaling Mechanisms in Unobservable Queues with Strategic Customers. (2017).
- [14] United Nations Department of Economic and Social Affairs. 2015. *World Urbanization Prospects: The 2014 Revision*. Technical Report. United Nations.
- [15] Erhun Ozkan and Amy R Ward. 2016. Dynamic Matching for Real-Time Ridesharing. (2016).
- [16] Marco Pavone, Stephen L Smith, Emilio Frazzoli, and Daniela Rus. 2012. Robotic load balancing for mobility-on-demand systems. *The International Journal of Robotics Research* 31, 7 (2012), 839–854.
- [17] Jean-Charles Rochet and Jean Tirole. 2006. Two-sided markets: a progress report. *The RAND Journal of Economics* 37, 3 (2006), 645–667.
- [18] Marc Rysman. 2009. The economics of two-sided markets. *The Journal of Economic Perspectives* (2009), 125–143.
- [19] Paolo Santi, Giovanni Resta, Michael Szell, Stanislav Sobolevsky, Steven H Strogatz, and Carlo Ratti. 2014. Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences* 111, 37 (2014), 13290–13294.
- [20] Kevin Spieser, Samitha Samaranyake, Wolfgang Gruel, and Emilio Frazzoli. 2016. Shared-vehicle mobility-on-demand systems: a fleet operator’s guide to rebalancing empty vehicles. In *Transportation Research Board 95th Annual Meeting*.
- [21] Kyle Treleven, Marco Pavone, and Emilio Frazzoli. 2013. Asymptotically optimal algorithms for one-to-one pickup and delivery problems with applications to transportation systems. *IEEE Trans. Automat. Control* 58, 9 (2013), 2261–2276.
- [22] John N Tsitsiklis, Kuang Xu, et al. 2012. On the power of (even a little) resource pooling. *Stochastic Systems* 2, 1 (2012), 1–66.
- [23] E Glen Weyl. 2010. A price theory of multi-sided platforms. *The American Economic Review* (2010), 1642–1672.
- [24] Rick Zhang and Marco Pavone. 2016. Control of robotic mobility-on-demand systems: a queueing-theoretical perspective. *The International Journal of Robotics Research* 35, 1-3 (2016), 186–203.

A ADDITIONAL PROOFS

In this appendix, we provide the complete proofs of our results. For convenience, we first restate the relevant results before presenting the proof. First, we present an auxiliary lemma, where we note that the dual LP as defined in 2 is degenerate up to translations:

PROPOSITION A.1. *For any Λ , given any β a feasible solution to the dual LP, then $\forall c \in \mathbb{R}$, $\beta + c\mathbf{1} = \{\beta_i + c\}$ is also a feasible point with the same objective value.*

PROOF. For any $i \in V$, we have $(\beta_i + c) - (\beta_j + c) = \beta_i - \beta_j \leq d_{ij}$ since β is feasible – hence, we have $\beta + c\mathbf{1}$ is also feasible. Furthermore, $(\beta + c\mathbf{1})^\top \Lambda = \beta^\top \Lambda + c\mathbf{1}^\top \Lambda = \beta^\top \Lambda$, as we have:

$$\sum_i \Lambda_i = \sum_{(i,j)} \Lambda_{ji} - \sum_{(i,j)} \Lambda_{ij} = 0$$

Thus both have the same objective value. $\square$

As a consequence, we can add non-negativity constraints $\alpha_i \geq 0 \forall i$ to (2) without affecting its value. Now we turn to the results in Section 3.

PROPOSITION A.2 (REF. PROPOSITION 3.1). *For sufficiently large θ , $\exists A_1, A_2 \in \mathbb{R}$ such that:*

- $\lambda = \mu \Rightarrow \gamma^\theta = A_1 \theta^{1/2}$
- $\lambda \neq \mu \Rightarrow \gamma^\theta = A_2 \theta^{-1/2} e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)} \cdot \theta} + o\left(\theta^{-1/2} e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)} \cdot \theta}\right)$

PROOF. First, when $\lambda \neq \mu$, we have:

$$\begin{aligned} 2\gamma^\theta &= \theta(a-1)(\mu-\lambda) + (a+1) \cdot \mathbb{E}[|X-Y| + |(\mu-\lambda) - (X-Y)|] - \theta((a-1)(\mu-\lambda) + (a+1)|\mu-\lambda|) \\ &= (a+1) \cdot \mathbb{E}[|X-Y| + |(\mu-\lambda) - (X-Y)|] - \theta(a+1)|\mu-\lambda| \end{aligned}$$

Here we have $X-Y \sim \mathcal{N}(\rho\theta(\mu-\lambda), \sigma^2)$, with $\sigma^2 = \sigma_X^2 + \sigma_Y^2$; hence, its absolute value is distributed as the folded normal distribution, with expected value ξ_ρ given by:

$$\xi_\rho = \sqrt{\frac{2}{\pi}} \sigma \cdot \exp\left(-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)} \cdot \theta\right) - \rho\theta|\lambda-\mu| \cdot \text{erf}\left(\frac{-\rho\theta|\lambda-\mu|}{\sqrt{2\theta \cdot \rho(1-\rho)(\lambda+\mu)}}\right)$$

Similarly, $(\mu-\lambda) - (X-Y) \sim \mathcal{N}((1-\rho)\theta(\mu-\lambda), \sigma^2)$, which has a similar expression for its expectation $\xi_{1-\rho}$. Next, to simplify the expression for ξ , we substitute the following asymptotic approximation of the error function:

$$\text{erf}(x) = 1 - \frac{e^{-x^2}}{\sqrt{\pi}x} + \frac{e^{-x^2}}{2\sqrt{\pi}x^3} + o\left(\frac{e^{-x^2}}{x^3}\right)$$

Substituting, we get that for any normal distribution with mean m and variance s , the expectation ξ of the corresponding folded normal distribution is given by:

$$\xi = \sqrt{\frac{2}{\pi}} \cdot s \cdot e^{-\frac{m^2}{2s^2}} + |m| \text{erf}\left(\frac{|m|}{\sqrt{2}s}\right) = |m| + \sqrt{\frac{2}{\pi}} \frac{s^3}{m^2} \cdot e^{-\frac{m^2}{2s^2}} + o\left(\frac{s^3}{m^2} \cdot e^{-\frac{m^2}{2s^2}}\right)$$

Substituting in the earlier expressions, and recalling that we assume $\rho \leq 1/2$, we get that $\exists (A, \tilde{A})$ such that:

$$\begin{aligned}
\gamma^\theta &= A.\theta^{-1/2}.e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)}.\theta} + \tilde{A}.\theta^{-1/2}.e^{-\frac{(1-\rho)(\lambda-\mu)^2}{2\rho(\lambda+\mu)}.\theta} + o\left(\theta^{-1/2}.e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)}.\theta}\right) \\
&= A_2.\theta^{-1/2}.e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)}.\theta} + o\left(\theta^{-1/2}.e^{-\frac{\rho(\lambda-\mu)^2}{2(1-\rho)(\lambda+\mu)}.\theta}\right),
\end{aligned}$$

where A_2 is an appropriately chosen constant. This proves the second result.

On the other hand, suppose that $\lambda = \mu$. In this case, we have $\gamma^\theta = (a+1)\mathbb{E}[|X-Y|]$, and we can directly use the expression for the mean of the folded normal distribution with mean 0 and variance σ^2 to get:

$$\gamma^\theta = (a+1)\sqrt{\frac{2\rho(1-\rho)(\lambda+\mu)}{\pi}} \cdot \sqrt{\theta}$$

Thus we have that $\exists A_1$ such that $\gamma^\theta = A_1\theta^{1/2}$. $\square$

THEOREM A.3 (REF. THEOREM 4.4). *Given network G , demand vector Λ and demand-split vector ρ , assume the following hold:*

- $\Lambda \in C_\eta$ for some $\eta \in \mathcal{E}$ (i.e., the monopolist demand vector lies in the interior of a dual cone)
- There is a r.v. Z such that $\Lambda^{\theta,a} = \theta\rho \odot \Lambda + \sqrt{\theta}Z$ with $\mathbb{E}[Z_e] = 0 \forall e \in E$ and $\mathbb{E}[||Z||_1] < \infty$
- $\forall e \in E$, we have $\mathbb{P}[|Z_e| > t] = O(f(t))$ with $f(t) = O(t)$

Now suppose we define:

$$L = \alpha^\top \rho \odot \Lambda + \beta^\top (1 - \rho) \odot \Lambda - \eta^\top \Lambda$$

Then we have:

- $\rho \odot \Lambda \in \dot{C}_\alpha$ and $(1 - \rho) \odot \Lambda \in \dot{C}_\beta \Rightarrow \gamma^\theta = L\theta + O(\theta f(\sqrt{\theta}))$
- $\rho \odot \Lambda \in C_\alpha \setminus \dot{C}_\alpha$ or $(1 - \rho) \odot \Lambda \in C_\beta \setminus \dot{C}_\beta \Rightarrow \gamma^\theta = L\theta + \Theta(\theta^{1/2})$

PROOF. The proof is the same as the homogeneous case, except that each companies' upper and lower bounds have to be treated separately.

Suppose we are in the first case. We can define δ in the first case and use concentration inequalities as before. In particular, we have for the first firm that $RC(\Lambda^{\theta,a}) = \alpha_1^\top \Lambda^{\theta,a}$, and moreover:

$$\begin{aligned}
\mathbb{E}[\alpha_1^\top \Lambda^{\theta,a}] &= \mathbb{E}\left[\alpha_1^\top \Lambda^{\theta,a} \mathbf{1}_{\{\forall e \in E, Z_e \leq \sqrt{\theta}\delta\}}\right] + \mathbb{E}\left[\alpha_1^\top \Lambda^{\theta,a} \mathbf{1}_{\{\exists e \in E \text{ s.t. } Z_e > \sqrt{\theta}\delta\}}\right] \\
&\leq \theta \alpha^\top \rho \odot \Lambda + \left(\theta \tilde{\beta} ||\Lambda||_1 + \tilde{\beta} \sqrt{\theta} \mathbb{E}\left[||Z||_1 \mid \exists e \in E \text{ s.t. } Z_e > \sqrt{\theta}\delta\right]\right) \mathbb{P}\left[\exists e \in E \text{ s.t. } Z_e > \sqrt{\theta}\delta\right] \\
&= \theta \alpha^\top \rho \odot \Lambda + O\left(\theta f(\sqrt{\theta})\right)
\end{aligned}$$

The upper bounds use the exact same argument as before. Thus we get a fast decay for one company, and if we sum each term of γ^θ then we get the desired result for the first case.

For the second case, we only need only one out of two companies to have a square root decay. Let's assume that the company b is under the affected regime. Concerning the lower bound, the cost for firm a can be lower bounded by its limit since it is suboptimal. Thus we have the following bounding box for any company that is under the fragmentation-resilient regime, thanks to the previous point:

$$\theta \alpha^\top \rho \odot \Lambda + O\left(\theta f(\sqrt{\theta})\right) \geq \mathbb{E}[\alpha_1^\top \Lambda^{\theta,a}] \geq \mathbb{E}[\alpha^\top \Lambda^{\theta,a}] = \theta \alpha^\top \rho \odot \Lambda$$

Concerning firm b , the argument from Theorem 4.3 can be repeated. Let's assume there are only 2 active corner points (we can adapt the proof as mentioned in 4.3) denoted β and $\tilde{\beta}$. We have for

the lower bound:

$$\begin{aligned}
\mathbb{E}[\alpha_2^\top \Lambda^{\theta,b}] &\geq \mathbb{E}[(\tilde{\beta}^\top \Lambda^{\theta,b}) \cdot \mathbb{1}_{\tilde{\beta}^\top Z \geq \beta^\top Z} + (\beta^\top \Lambda^{\theta,b}) \cdot \mathbb{1}_{\beta^\top Z \geq \tilde{\beta}^\top Z}] \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + \theta^{1/2} \cdot \mathbb{E} \left[(\tilde{\beta}^\top Z) \cdot \mathbb{1}_{\tilde{\beta}^\top Z \geq \beta^\top Z} + (\beta^\top Z) (1 - \mathbb{1}_{\tilde{\beta}^\top Z \geq \beta^\top Z}) \right] \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + \theta^{1/2} \cdot \mathbb{E} \left[((\tilde{\beta} - \beta)^\top Z) \cdot \mathbb{1}_{\tilde{\beta}^\top Z \geq \beta^\top Z} + (\beta^\top Z) \right] \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + \theta^{1/2} \cdot \mathbb{E} \left[((\tilde{\beta} - \beta)^\top Z) \cdot \mathbb{1}_{\tilde{\beta}^\top Z \geq \beta^\top Z} \right] \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + \Omega(\theta^{1/2})
\end{aligned}$$

Concerning the higher bound of the fragmented-affected company, we consider the radius δ of a ball strictly included in $C_\beta \cup C_{\tilde{\beta}}$ and write:

$$\begin{aligned}
\mathbb{E}[\alpha_2^\top \Lambda^{\theta,b}] &= \mathbb{E} \left[(\alpha_2^\top \Lambda^{\theta,b}) \times \left(\mathbb{1}_{\{\forall e \in E, Z_e \leq \sqrt{\theta} \delta\}} + \mathbb{1}_{\{\exists e \in E \text{ s.t. } Z_e > \sqrt{\theta} \delta\}} \right) \right] \\
&\leq \mathbb{E} \left[(\alpha_2^\top \Lambda^{\theta,b}) \mathbb{1}_{\{\forall e \in E, Z_e \leq \sqrt{\theta} \delta\}} \right] + O(f(\sqrt{\theta})) \\
&\leq \mathbb{E} \left[(\alpha_2^\top \Lambda^{\theta,b}) \right] + O(f(\sqrt{\theta})) \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + \theta^{1/2} \mathbb{E} \left[(\beta + \tilde{\beta})^\top Z \right] + O(f(\sqrt{\theta})) \\
&= \theta \beta^\top (1 - \rho) \odot \Lambda + O(\theta^{1/2})
\end{aligned}$$

Finally, what remains to do is sum over both company, subtract the cost of the monopoly regime and we get the desired conclusion. $\square$

THEOREM A.4. *Given network G , demand vector Λ , n companies and demand-split vector $(\rho_1, \dots, \rho_n)$ such that their sum is the vector $\mathbb{1}_n$, assume the following hold:*

- $\Lambda \in C_\eta$ for some $\eta \in \mathcal{E}$ (i.e., the monopolist demand vector lies in the interior of a dual cone)
- There is a r.v. Z such that, $\forall i$, $\Lambda^{\theta,i} = \theta \rho_i \odot \Lambda + \sqrt{\theta} Z^i$ with $\mathbb{E}[Z_e^i] = 0 \forall e \in E$ and $\mathbb{E}[||Z^i||_1] < \infty$
- $\forall e \in E$, we have $\mathbb{P}[|Z_e^i| > t] = O(f(t))$ with $f(t) = O(t)$
- $\forall i$, $\rho_i \odot \Lambda \in C_{\alpha_i}$

Now suppose we define:

$$L = \left(\sum_{i=1}^n \alpha_i^\top \rho_i \odot \Lambda \right) - \eta^\top \Lambda$$

Then we have:

- $\forall i$, $\rho_i \odot \Lambda \in \overset{\circ}{C}_\alpha \Rightarrow \gamma^\theta = L\theta + O(\theta f(\sqrt{\theta}))$
- $\exists i$, $\rho_i \odot \Lambda \in C_\alpha \setminus \overset{\circ}{C}_\alpha \Rightarrow \gamma^\theta = L\theta + \Theta(\theta^{1/2})$

PROOF. The previous proof gives bounds for each company depending on the underlying regime of their expected demand. We make distinction by denoting α_i the corner point of the expected demand, and ξ_i the corner point of the random variable $\Lambda^{\theta,i}$. We have for the resilient an affected regime respectively:

$$\begin{aligned}
\theta \alpha_i^\top \rho_i \odot \Lambda + O\left(\theta f(\sqrt{\theta})\right) &\geq \mathbb{E}[\xi_i^\top \Lambda^{\theta,i}] \geq \theta \alpha_i^\top \rho_i \odot \Lambda \\
\theta \alpha_i^\top \rho_i \odot \Lambda + O(\theta^{1/2}) &\geq \mathbb{E}[\xi_i^\top \Lambda^{\theta,i}] \geq \theta \alpha_i^\top \rho_i \odot \Lambda + \Omega(\theta^{1/2})
\end{aligned}$$

Thus depending on the regime, we get different bounds that can then be summed to obtain this result. $\square$

B REBALANCING AND STOCHASTIC DYNAMICS

As we mention in Section 2, the formulation of the rebalancing cost in terms of the min-cost circulation problem arises in many models of MoD systems. We now discuss three different justifications given for it in literature,

- (1) Best-case scenarios: For given Λ_{ij} , the total number of trip requests between pairs of nodes (i, j) , $RC(\Lambda)$ represents the *minimum ex-post transportation cost*, i.e., the minimum cost of rebalancing empty vehicles to meet the demand. This is particularly relevant for fragmentation-affected settings, as the true cost in such cases must grow as fast as the LP.
- (2) Long-term planning models: These arise in settings where the platform knows all trip requests in advance, e.g., reservation models like Lyft Shuttle. They are also valid to an extent in settings where passengers have high patience [16, 20, 24].

In more detail: suppose trip requests between pairs of nodes i, j arrive according to a Poisson process with rate Λ_{ij} . Then, for an MoD platform with K vehicles serving n stations, to satisfy all demands, the number of vehicles must satisfy the following *stability condition* (via a simple time-conservation argument):

$$K > \sum_{(i,j) \in E} \Lambda_{ij} \tau_{ij} + RC(\Lambda),$$

When passengers have infinite patience (i.e., are willing to tolerate any finite delay), then it is not hard to show that the condition is also a sufficient condition – in particular, the system is stabilized by natural static probabilistic rebalancing policies (cf. [21]). The argument for settings with reservations is similar, except that trip planning can now be done in advance.

- (3) Closed-queueing networks and passenger loss models: This is the opposite scenario from reservations, wherein passengers are assumed to have low patience, and leave the system if no vehicle is available nearby. In more detail, consider a setting where passengers arrive via a Poisson process, and leave the system immediately if there is no available vehicle. Note that in this setting, the solution to the LP for $RC(\Lambda)$ is still a lower bound on the rebalancing cost. Now given the stochastic arrival process, it is clear that any policy must incur some demand loss. Somewhat surprisingly, Banerjee et al. [5] show that a natural fluid rebalancing policy comes close to matching $RC(\Lambda)$, in that, as long as $\sum_{i,j} \Lambda_{ij} \tau_{ij} = o(K)$, then static rebalancing with probabilities chosen via the LP in (1) incurs a loss which is at most a $(1 + n/K)$ fraction of $RC(\Lambda)$ (and this was further improved to $1 + e^{-O(K)}$ via state-dependent dispatch policies in [7], under some additional balance assumptions on the demands).

To translate this to our setting, note that all our results consider PoF scaling with respect to the demand, and are independent of the number of vehicles. For the fragmentation-affected setting, our results still hold since $RC(\Lambda)$ is a lower bound on the true costs. On the other hand, the results of [5, 7] show that the true cost under probabilistic rebalancing is off from $RC(\Lambda)$ by a small factor as long as the number of vehicles K scales slightly faster than the demand; thus, we still get vanishing PoF in the fragmentation-resilient setting.

C ADVERSARIAL PRICE OF FRAGMENTATION

An alternative to the stochastic PoF under scaling is to consider the *Adversarial Price of Fragmentation* – the increase in rebalancing costs over all feasible demand splits:

$$\begin{aligned} \max \quad & RC(\Lambda^a) + RC(\Lambda - \Lambda^a) \\ \text{s.t.} \quad & 0 \leq \Lambda_{ij}^a \leq \Lambda_{ij} \quad \forall (i, j) \in E \end{aligned} \tag{3}$$

Though our results (in particular, Theorem 4.4) suggest that this can be arbitrarily bad, nevertheless, such a formulation may be appealing for empirical analysis. The above formulation however leads to an intractable problem; in particular, as a consequence of Proposition 2.1, we have that the function $RC(\Lambda^a) + RC(\Lambda - \Lambda^a)$ is convex in Λ^a , and hence maximizing it over a polyhedron is in general computationally hard.

To circumvent this, we use the following heuristic: First, since the polyhedron on which we maximize over is bounded and the cost function is convex, we know that the worst case split is reached at a corner point of the polytope $[0, 1]^{N^2}$, which corresponds to an ILP in $\{0, 1\}^{N^2}$. Furthermore, to obtain bounds, we can utilize a projected subgradient heuristic for the dual problem, taking advantage of the cost function being continuous and convex, and differentiable almost everywhere. In particular, note that we can write the dual as: $\max_{\alpha \in \mathcal{E}} \alpha^\top \Lambda$, where $\mathcal{E}$ denotes the (finite) set of corner points; hence it is differentiable except on degenerate values of Λ .

Adversarial PoF for NYC Data: Even though the above heuristic does not guarantee optimality at convergence, using it on the NYC data, we obtain a fixed point corresponding to a rebalancing cost increase of 567% compared to the monopoly (for the demand data between noon and 1 pm on May 10th). This observation shows that demand fragmentation can lead to large systemic losses. Furthermore, note that when considering the worst case loss for a higher number of companies across which the demand may fragment, we know that the results can only be worse (i.e. the duopoly gives a lower bound).

D ADDITIONAL PLOTS

In this appendix, we present some additional plots. First in Figure. 10(a), 10(b), 10(c) and 10(d), we show the locations of the cluster centers, for different choices of number of stations.

Fig. 9. Multiple firms and heterogeneous market-shares: In Fig. 9(b), we demonstrate the PoF scaling with heterogeneous market shares (chosen uniformly at random on each edge). The plot demonstrates the linear scaling of PoF with θ . Note that the y-axis is in multiples of 10^5 ; this shows that inhomogeneous demand-splits (and therefore adversarial) can have much higher PoF. Similarly, in Fig. 9(a), we demonstrate the change in PoF scaling when going from 2 to 3 firms; note that the situation gets worse with 3 firms, but qualitatively, the scaling behavior is similar. Both plots are based on the same demand distribution data as in 1.

Next, we simulate and plot the PoF for inhomogeneous market shares, using the same data as 1, with demand-split ratios picked uniformly at random for each edge. The splits we picked result in a non-zero scaling coefficient L (defined in Theorem 4.4; note that our theorem predicts a linear

Fig. 10. Scatter plot of the demand locations with the position of clusters for increasing number of stations.

divergence, which is confirmed by the plot in Fig. 9(b). This suggests that spatially inhomogeneous demand-splitting can lead to much higher costs. Finally, in Figure. 9(a), we compare the two fragmentation regimes when the number of firms is either 2 or 3. Note that the phase transition is qualitatively the same (though quantitatively worse) with more firms. The data for both plots is the same data as 1.

Received February 2018; revised March 2018; accepted June 2018