






Graph-Based Equilibrium Metrics for Dynamic Supply–Demand Systems With Applications to Ride-sourcing Platforms

Fan Zhou, Shikai Luo, Xiaohu Qie, Jieping Ye & Hongtu Zhu


To cite this article: Fan Zhou, Shikai Luo, Xiaohu Qie, Jieping Ye & Hongtu Zhu (2021) Graph-Based Equilibrium Metrics for Dynamic Supply–Demand Systems With Applications to Ride-sourcing Platforms, Journal of the American Statistical Association, 116:536, 1688-1699, DOI: [10.1080/01621459.2021.1898409](https://doi.org/10.1080/01621459.2021.1898409)


To link to this article: <https://doi.org/10.1080/01621459.2021.1898409>

 View supplementary material 



 Published online: 23 Apr 2021.

 Submit your article to this journal 

 Article views: 990

 View related articles 

 View Crossmark data 

 Citing articles: 3 View citing articles 



Graph-Based Equilibrium Metrics for Dynamic Supply–Demand Systems With Applications to Ride-sourcing Platforms

Fan Zhou^a, Shikai Luo^b, Xiaohu Qie^b, Jieping Ye^b, and Hongtu Zhu^b

^aSchool of Statistics and Management, Shanghai University of Finance and Economics, Shanghai, China; ^bDidi Chuxing, Beijing, China

ABSTRACT

How to dynamically measure the local-to-global spatio-temporal coherence between demand and supply networks is a fundamental task for ride-sourcing platforms, such as DiDi. Such coherence measurement is critically important for the quantification of the market efficiency and the comparison of different platform policies, such as dispatching. The aim of this paper is to introduce a graph-based equilibrium metric (GEM) to quantify the distance between demand and supply networks based on a weighted graph structure. We formulate GEM as the optimal objective value of an unbalanced optimal transport problem, which can be formulated as an equivalent linear programming and efficiently solved. We examine how the GEM can help solve three operational tasks of ride-sourcing platforms. The first one is that GEM achieves up to 70.6% reduction in root-mean-square error over the second-best distance measurement for the prediction accuracy of order answer rate. The second one is that the use of GEM for designing order dispatching policy increases drivers' revenue for more than 1%, representing a huge improvement in number. The third one is that GEM can serve as an endpoint for comparing different platform policies in AB test. Supplementary materials for this article, including a standardized description of the materials available for reproducing the work, are available as an online supplement.

ARTICLE HISTORY

Received March 2020
Accepted February 2021

KEYWORDS

Graph-based equilibrium metric; Order dispatching; Ride-sourcing platform; Unbalanced optimal transport; Weighted graph

1. Introduction


Large volumes of data collected from multiple spatio-temporal networks are increasingly studied in diverse fields including climate science, social sciences, neuroscience, epidemiology, and transportation. In addition, those spatio-temporal networks may interact with each other across spatial and/or temporal dimension. A typical example is that the dynamic demand and supply networks of a ride-sourcing platform (Wang and Yang 2019) are two sequences of unnormalized masses measured on the same undirected (or directed) graph $G = (\mathbb{V}, \mathbb{E})$, where \mathbb{V} and \mathbb{E} are, respectively, a vertex set and a set of edges connecting vertex pairs. Figure 1 illustrates how the two complicated networks interact with each other and evolve over time. Specifically, a city is divided into hundreds of nonoverlapping grids as the vertex set \mathbb{V} with the edge structure \mathbb{E} determined by road networks and location functionalities. Both demands and supplies are observed across grids at each time window with possibly different total masses and distributions. The ride-sourcing platform uses some order dispatching policy to match customer requests with possible surrounding idle drivers, while after finishing serving assigned orders, drivers return back to the supply pool to prepare for the next feasible matching. The aim of this paper is to address a fundamental question of interest for the demand and supply networks of two-sided markets.


The fundamental question of interest that we consider here is how to quantify the spatial equilibrium of dynamic supply–demand networks for two-sided markets, particularly ride-sourcing platforms (e.g., Uber and DiDi). To solve

this question, we first introduce a weighted graph structure (G, W, C) to characterize the transport network and transport costs of a city. Specifically, we divide each market into N disjoint areas and regard them as vertices, denoted as $\mathbb{V} = \{v_1, \dots, v_N\}$. Let \mathbb{E} be a set of edges between any possible pair of vertices such that $(v_i, v_j) \in \mathbb{E} \subset \mathbb{V} \times \mathbb{V}$ is an edge equipped with a nonnegative weight w_{ij} (e.g., transportation cost). For all $(v_i, v_j) \notin \mathbb{E}$, we set $w_{ij} = \infty$. The weighted graph structure consists of an undirected (or directed) graph $G = (\mathbb{V}, \mathbb{E})$ as well as a weight matrix $W = (w_{ij})$, where w_{ij} 's are nonnegative weights. A graph-based transport cost from v_i to v_j is defined as $c_{ij} = \min_{K \geq 0, (i_k)_{k=0}^K: v_i \rightarrow v_j} \{\sum_k w_{i_k, i_{k+1}} : \forall k \in \llbracket 0, K-1 \rrbracket, (v_{i_k}, v_{i_{k+1}}) \in \mathbb{E}\}$, where $(i_k)_{k=0}^K : v_i \rightarrow v_j$ denotes any path on G through \mathbb{E} starting from $v_{i_0} = v_i$ and ending at $v_{i_K} = v_j$. Thus, c_{ij} is the geodesic distance from v_i to v_j or the minimal cost of transporting one unit of object from v_i to v_j . Thus, we can define a transport cost matrix on (G, W) , denoted as $C = (c_{ij}) \in R^{N \times N}$. The C may be time variant, since it depends on the real-time traffic and weather conditions for ride-sourcing platform. The C is possibly asymmetric since the graph G can be directed.

Second, we need to introduce a distance (or metric) to quantify the difference between demand and supply masses at each time interval and across time on (G, W, C) . At a given time interval, we define $\nu_j = \nu(v_j)$ and $\mu_j = \mu(v_j)$ as the point masses at vertex v_j for the two measures ν and μ , which, respectively, represent the number of customer requests and available drivers inside the vertex v_j of the ride-sourcing platform (Wang and Yang 2019). The supply and demand systems

CONTACT Hongtu Zhu  bowenhongtu@gmail.com  Didi Chuxing, Beijing, China.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

 These materials were reviewed for reproducibility.

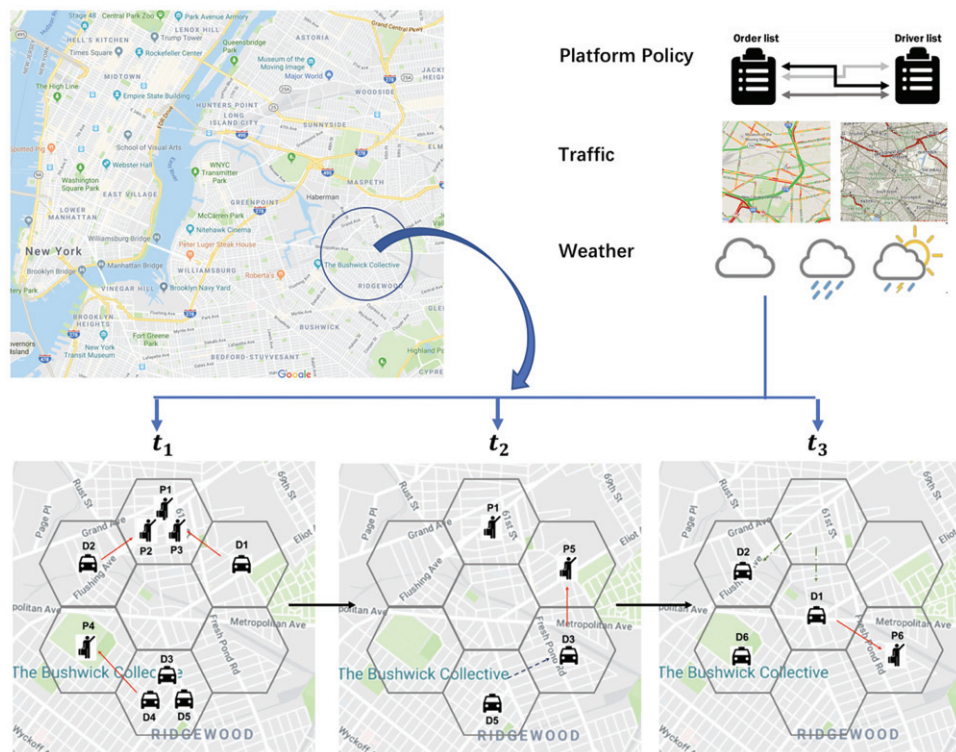


Figure 1. Dynamic supply and demand networks at three time points in a representative ride-sourcing platform. We divide the whole city into multiple hexagon areas.

at each timestamp can be modeled as two discrete Lebesgue measures μ and ν on (G, W, C) with locally finite masses such that $\max(\mu(\mathbb{V}_0), \nu(\mathbb{V}_0))$ is finite for every compact set $\mathbb{V}_0 \subset \mathbb{V}$. We consider a general case that the two measures can be unbalanced, that is, $\mu = \sum_{i=1}^N \mu_i$ and $\nu = \sum_{i=1}^N \nu_i$ may be unequal to each other. Defining a metric between μ and ν falls into the field of optimal transport.

Optimal transport has been widely studied in diverse disciplines, such as statistics, applied mathematics, medical imaging, and computer vision. Wasserstein-based metrics based on the mathematics of optimal mass transport have been proved to be powerful tools for comparing objects in complex spaces. Some successful applications include solving transport partial differential equation (PDE) (Ambrosio and Gangbo 2008), imaging processing (Rabin and Papadakis 2015), statistical inference in machine learning (Solomon et al. 2014), manifold diffeomorphisms (Grenander and Miller 2007), and serving as the cost function for training Generative Adversarial Networks (Arjovsky, Chintala, and Bottou 2017), among many others. However, existing Wasserstein-based metrics are not directly applicable to the comparison of two unbalanced measures defined on (G, W, C) as detailed in Section 2.1.

We introduce a graph-based equilibrium metric (GEM) and formulate it as an unbalanced optimal transport problem. Our main contributions are summarized as follows. First, we propose a novel GEM, which can be regarded as a restricted generalized Wasserstein distance, to quantify the distance between dynamic demand and supply networks on the weighted graph structure. It not only allows the optimal transport guided by asymmetric costs and node connections, but also accounts for unbalanced masses. It also allows one of the two sides (supplies)

to play the transporting role and the other (demands) to be fixed, which satisfies the physical interpretation of ride-sourcing platforms. Second, varying the size of each vertex leads to multilevel GEMs and their corresponding optimal transport functions. At the finest scale, our GEM reduces to solving an unbalanced assignment problem and its corresponding optimal transport function contains many local details. In contrast, at a relatively coarse scale, it gives a coarse representation (or low-frequency patterns) of the optimal transport function. Third, numerically, the calculation of GEM can be reformulated as a standard linear programming (LP) problem. Theoretically, we investigate several theoretical properties of GEM including the convergence of the LP algorithm for computing GEM, the expectation of GEM, and the metric property, additive property and weak convergence of GEM. Fourth, we apply GEM to the “supply-demand diagnostic data set” obtained from the DiDi Chuxing in order to address some important operational tasks, such as the prediction of the efficiency of a given dispatching policy.

The remainder of this article is structured as follows. Section 2 introduces graph-based equilibrium metrics and their computational approach, while discussing their potential applications. Section 3 studies four theoretical properties associated with GEM. Section 4 demonstrates the applications of GEM in the intelligent operations of DiDi Chuxing.

2. Methodologies

2.1. Existing Wasserstein-Type Distances

Many approaches have been proposed to measure the distance between two measures (or distributions) on a metric space. Most

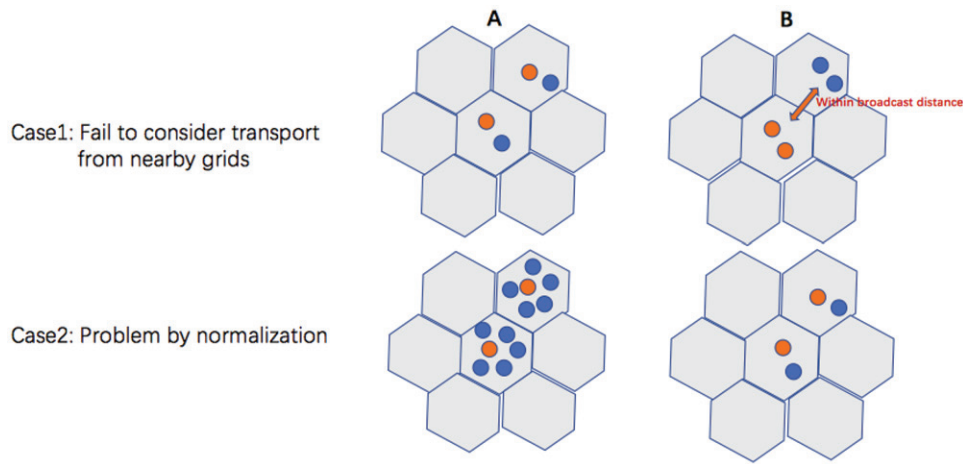


Figure 2. Examples of supply–demand networks for illustrating the limitations of Hellinger distance: blue and orange objects represent orders and idle drivers, respectively. The first row illustrates the consequence of not transporting objects into nearby vertices. Specifically, the left one has a zero Hellinger distance, whereas the right one has a nonzero Hellinger distance without transporting the two drivers into the nearby vertex with two orders. The second row illustrates that the normalization step can remove the imbalance between orders and idle drivers. Specifically, the left one has a large imbalance between orders and idle drivers, whereas after the normalization, the right one has a zero Hellinger distance.

of them fall into two broad categories including the aggregation of pixel-wise differences and the transport cost of moving one measure to match the other. Measurements in the first category include the L_p -distance, the total variation (TV) distance, and the Kullback–Leibler (KL) divergence (Cha 2007), among others. A typical example is the Hellinger Distance reviewed as follows.

Definition 2.1 (Hellinger Distance). Let $M(X)$ and $M_+(X)$ be the vector space of Radon measures and the cone of nonnegative Radon measures on a Hausdorff topological space X . Then, we use μ and ν to denote two probability measures that are absolutely continuous with respect to a third probability measure ν_0 . The square of the Hellinger distance between μ and $\nu \in M_+(X)$ is defined as $D_H^2(\mu, \nu) = 0.5 \int_X (\sqrt{d\mu/d\nu_0} - \sqrt{d\nu/d\nu_0})^2 d\nu_0$, where $d\mu/d\nu_0$ and $d\nu/d\nu_0$ are the Radon-Nikodym derivatives of μ and ν , respectively.

All these metrics suffer from two major issues. Please refer to Figure 2 for details. First, all these metrics not only fail to consider the connections among different locations (vertices in graph), but also ignore the topological (or geometric) structure of X . Second, the use of Hellinger-type distances requires a normalization step to enforce $\mu(X) = \int_X d\mu = \nu(X) = \int_X d\nu$, which can create a false balance issue.

To address these two issues, the second category of metrics, such as the Wasserstein distance (Villani 2008), is proposed by solving an optimal transport problem. In the real world, original supply resources can usually be transported to achieve a better equilibrium between μ and ν . All those distances have deep connections to well-studied assignment problems from the field of combinatorial optimization (Steele 1987).

Definition 2.2 (Wasserstein Distance). Let X and Y be Hausdorff topological spaces and $X \times Y$ be their product space. We introduce a lower semi-continuous function $c : X \times Y \rightarrow R \cup \{\infty\}$, a nonnegative measure (or a transport function) $\gamma \in M_+(X \times Y)$, and an equality constraint $\iota_{\{=\}}(\alpha|\beta)$ which is 0 if $\alpha = \beta$ and

∞ otherwise. Then the optimal transport problem for measures $\mu \in M_+(X)$ and $\nu \in M_+(Y)$ with the same total masses, that is $\mu(X) = \nu(Y)$, can be defined as

$$D_W(\mu, \nu|c) = \inf_{\gamma \in M_+(X \times Y)} \left\{ \int_{X \times Y} c d\gamma + \iota_{\{=\}}(P_\#^X \gamma|\mu) + \iota_{\{=\}}(P_\#^Y \gamma|\nu) \right\}, \tag{1}$$

where $P_\#^X \gamma$ and $P_\#^Y \gamma$ denote the first and second marginals of γ , respectively.

Intuitively, γ denotes a transport plan, measuring how far you have to move the mass of μ to turn it into ν . Standard optimal transport in (1) is only meaningful whenever μ and ν have the same total masses. Whenever $\mu(X) \neq \nu(Y)$, there is no feasible solution γ in Equation (1). For the real-world ride-sourcing platforms, however, it is important to compute some sort of relaxed transport between two arbitrary nonnegative measures. An improved approach is to build an unbalanced optimal transport problem by introducing two divergences over X and Y , denoted as \mathcal{D}_{φ_1} and \mathcal{D}_{φ_2} , respectively (Chizat et al. 2018; Liero, Mielke, and Savaré 2018).

Definition 2.3 (Divergences). Let φ be an entropy function. For $\mu, \nu \in M(T)$, $\frac{d\mu}{d\nu} \nu + \mu^\perp$ is the Lebesgue decomposition of μ with respect to ν . The divergence D_φ is defined by $D_\varphi(\mu|\nu) := \int_T \varphi(\frac{d\mu}{d\nu}) d\nu + \varphi'_\infty \mu^\perp(T)$ if μ and ν are nonnegative and ∞ otherwise.

Now, we can give the formal definition of generalized Wasserstein distance.

Definition 2.4 (Generalized Wasserstein Distance (GWD)). Let $c : X \times Y \rightarrow [0, \infty]$ be a lower semi-continuous function, the unbalanced optimal transport problem is

$$D_{\varphi_1, \varphi_2}(\mu, \nu|c) = \inf_{\gamma \in M_+(X \times Y)} \left\{ \int_{X \times Y} c d\gamma + \mathcal{D}_{\varphi_1}(P_\#^X \gamma|\mu) + \mathcal{D}_{\varphi_2}(P_\#^Y \gamma|\nu) \right\}. \tag{2}$$

Different from standard Wasserstein Distance which normalizes the input measures into probability distributions, GWD quantifies in some way the deviation of the marginals of the transport plan γ from the two unbalanced measures μ and ν by using φ -divergence. Although $D_{\varphi_1, \varphi_2}(\mu, \nu|c)$ enjoys some nice properties, such as metric property (Chizat et al. 2018; Liero, Mielke, and Savaré 2018), the solution to Equation (2), denoted as γ_* , may not have any physical meaning. For the ride-sourcing business, such γ_* is critically important for assigning supplies to demands since it can be regarded as the graph representation of a dispatching policy. Therefore, the use of $D_{\varphi_1, \varphi_2}(\mu, \nu|c)$ still cannot fully cover the “useful” relative size between μ and ν , since it may underestimate unmatched resources by allowing some infeasible transports, that is, the space $M_+(X \times Y)$ is too large to be useful. Three major issues of using $D_{\varphi_1, \varphi_2}(\mu, \nu|c)$ are given as follows. The first issue is that in many applications (e.g., ride-sourcing platform), point masses in only one of the two measures are allowed to be transported and those in the other measure are fixed. In this case, the symmetric property does not hold. The second issue is that neither W nor C can be used to define a standard metric space on $G = (\mathbb{V}, \mathbb{E})$, since transport cost (or weight) matrix may not satisfy the three key assumptions of standard metrics. For instance, the transport cost from v_i to v_j may be unequal to that from v_j to v_i , since transport cost matrix $C \in R^{N \times N}$ can be asymmetric for directed graphs. Moreover, the direct transport cost from v_i to v_j may be larger than or equal to the sum of the transport cost from v_i to v_k and that from v_k to v_j . The third issue is that in some applications, such as supply-demand networks, the transport cost from v_i to v_j may not be a constant and the transport cost from a vertex to itself may not be zero. It is possible that supply units at vertex v_i have their individual transport costs of moving within/outside the vertex v_i . Subsequently, their transport costs from v_i to v_j may follow a distribution instead of being a constant.

2.2. Graph-based Equilibrium Metrics

On (G, W, C) , we formally introduce our GEMs for two discrete measures μ and ν in $M_+(\mathbb{V})$, among which point masses in μ are allowed to be transported and those in ν are fixed. We need to introduce some notations. In this case, we have $X = Y = \mathbb{V}$ and use $P_{\#1}^V \gamma$ and $P_{\#2}^V \gamma$ to represent $P_{\#}^X \gamma$ and $P_{\#}^Y \gamma$, respectively. Let $|\mu| = \sum_{i=1}^N \mu(v_i)$ and $|\mu - \tilde{\mu}| = \sum_{i=1}^N |\mu(v_i) - \tilde{\mu}(v_i)|$. For $i = 1, \dots, N$, we use \mathcal{N}_i to denote the neighboring set of v_i in \mathbb{V} , which contains v_i and its (possibly high-order) neighboring vertices. Moreover, $v_i \in \mathcal{N}_j$ does not ensure $v_j \in \mathcal{N}_i$ since the traffic and road network may prohibit from directly transporting cars from v_j to v_i .

Let $c : \mathbb{V} \times \mathbb{V} \rightarrow R \cup \{\infty\}$ be a function and $\gamma \in M_+(\mathbb{V} \times \mathbb{V})$ be a nonnegative measure. The general form of our GEM on (G, W, C) is written as

$$\begin{aligned} \rho_\lambda(\mu, \nu|G, C) & \quad (3) \\ & = \inf_{\tilde{\mu} \in M_+(\mathbb{V}), \gamma \in M_+(\mathbb{V} \times \mathbb{V})} \left\{ |\nu - \tilde{\mu}| + \lambda \int_{\mathbb{V} \times \mathbb{V}} c d\gamma \right\} \end{aligned}$$

subject to an equality constraint and two sets of transport constraints given by

$$\begin{aligned} |\mu| & = |\tilde{\mu}|, \quad (P_{\#1}^V \gamma)(v_i) = \sum_{v_j \in \mathcal{N}_i} \gamma(v_i, v_j) = \mu_i \quad \text{and} \\ (P_{\#2}^V \gamma)(v_i) & = \sum_{v_j \in \mathcal{N}_j} \gamma(v_j, v_i) = \tilde{\mu}_i, \end{aligned} \quad (4)$$

where λ is a nonnegative hyperparameter. The three sets of constraints in Equation (4) ensure that $\tilde{\mu}$ shares the same total mass with μ and γ transports μ to $\tilde{\mu}$. Thus, the feasible set for Equation (4) is much smaller than that for Equation (2). The integration of $\lambda \int_{\mathbb{V} \times \mathbb{V}} c d\gamma$ and the three sets of constraints in Equation (4) is equivalent to the balanced Wasserstein distance in Equation (1), so GEM is the integration of the balanced Wasserstein distance and the L_1 norm.

In our GEM framework, one of the two measures plays the role of “predator” to move and “catch” the “prey,” which mimics the general supply-demand system of ride-sourcing platforms. Therefore, different from the setting of Piccoli and Rossi (2014), in which both two measures are rescaled, we fix ν but change μ only to make the two sides match each other under the asymmetric distance and transport range constraints. In Figure 3, we consider two simple examples in order to understand the differences between GEM and GWD. Moreover, since we only consider the transport from μ to $\tilde{\mu}$ with ν fixed, $\rho_\lambda(\mu, \nu|G, C)$ is generally asymmetric and can be regarded as a restricted GWD.

Besides GEM, the optimal solution of $(\tilde{\mu}, \gamma)$ to Equation (3), denoted as $(\tilde{\mu}_*, \gamma_*)$, also plays an important role in various two-sided markets, such as ride-sourcing platforms and E-commerce. The $\tilde{\mu}_*$ can be regarded as the supply distribution obtained by optimally transporting μ to match demands ν , whereas γ_* is an optimal transport function associated with $\rho_\lambda(\mu, \nu|G, C)$. If we vary the area of each vertex from the coarsest to the finest scale, then we obtain multilevel GEM and its transport function. At the finest scale, our GEM reduces to solving an unbalanced dispatching problem, so γ_* is able to capture the local structure of the optimal transport function. In contrast, at a relatively coarse scale, we obtain a coarse representation of the optimal transport function, reflecting its global patterns. We will discuss how to apply GEM to ride-sourcing platforms in Section 2.4.

Furthermore, we can simplify γ by defining $\gamma = (\gamma_{ij})$ as an $N \times N$ flow matrix with γ_{ij} being the transport amount from v_i to v_j . Let Γ represent the set consisting of all the feasible solutions γ with all nonnegative elements $\gamma_{ij} \geq 0$. Let $\tilde{\mu} = (\tilde{\mu}_1, \dots, \tilde{\mu}_N)^T \in R^N$ represent the measure μ after transporting γ such that $\tilde{\mu}_i = \sum_{v_j \in \mathcal{N}_i} \gamma_{ji}$ holds for all i . Thus, our GEM is equivalent to solving a discrete optimization problem as follows:

$$\begin{aligned} \rho_\lambda(\mu, \nu|G, C) & = \min_{\gamma \in \Gamma} \{ \|\nu - \tilde{\mu}\|_1 + \lambda \sum_{v_i \in \mathbb{V}} \sum_{v_j \in \mathbb{V}} c_{ij} \gamma_{ij} \} \\ \text{subject to} & \quad \sum_{v_j \in \mathcal{N}_i} \gamma_{ij} = \mu_i, \quad \sum_{v_j \notin \mathcal{N}_i} \gamma_{ij} = 0, \quad \text{and} \\ & \quad \tilde{\mu}_i = \sum_{v_j \in \mathcal{N}_j} \gamma_{ji} \quad \text{for } \forall v_i \in \mathbb{V}, \end{aligned} \quad (5)$$

where $\nu = (\nu_1, \dots, \nu_N)^T \in R^N$ and $\|\cdot\|_1$ corresponds to the L_1 norm. Moreover, $\|\nu - \tilde{\mu}\|_1$ in Equation (5) is equivalent to the first term of the objective function in Equation (3).




Figure 3. Examples illustrating the differences between GEM and GWD. Panel (a): For GEM, the four units can be matched in the left sub-figure, whereas it is infeasible in the right one. There exists a directed edge from A to B, but not from B to A. Panel (b): In the top sub-figure, one “demand” unit at vertex C cannot be matched for GEM since the transport from vertex A to vertex C is not allowed in this case, whereas in the bottom sub-figure, the transport is allowed by GWD. In panel (b), for GEM, it is assumed that the neighboring set \mathcal{N}_i only includes the adjacent vertices of each vertex.

There are two key advantages of using the derived form given in Equation (5) compared to the existing unbalanced optimal transport problem. The first one is that transport is only allowed between a vertex v_i and its neighboring set \mathcal{N}_i based on G .

The second one is that λ can balance the transport cost taken to reallocate point masses and the requirement of assigning μ to satisfy ν . The choice of λ in practice is data-driven. To ensure that the transport only happens among selected vertex pairs under the optimal transport plan, the theoretical upper bound of λ is $2/\max_{v_i \in V, v_j \in \tilde{\mathcal{N}}_i}(c_{ij})$, where $\tilde{\mathcal{N}}_i \subset \mathcal{N}_i$ contains all the neighboring vertices of v_i that receive transport from v_i . In this case, the cost of transporting one unit of supply from vertex v_i to $v_j \in \tilde{\mathcal{N}}_i$, λc_{ij} , is smaller than its contribution to reducing $\|\mathbf{v} - \tilde{\boldsymbol{\mu}}\|_1$, which is 2 (1 for v_i and v_j , respectively), when $v_j - \mu_j \geq 1$ and $\mu_i - v_i \geq 1$. Transport from v_i to v_j keeps decreasing the objective value $\rho_\lambda(\boldsymbol{\mu}, \nu | G, C)$ until either the balance in the destination vertex v_j or that in the origin vertex v_i is achieved. In the real world, we usually let $\lambda \max_{v_i \in V, v_j \in \tilde{\mathcal{N}}_i}(c_{ij})$ fall into the range $[0.4, 0.5]$ with c_{ij} being the geological distance and $\tilde{\mathcal{N}}_i$ containing all the first-order adjacent vertices of v_i in (G, W, C) , which can achieve the best performance in some problems, such as the prediction of order answer rate.

2.3. Computational Approach

Optimal solution γ^* to Equation (5) can be calculated by solving a standard linear programming (LP). We will reformulate (5) as a LP problem and then use a revised simplex method incorporated in a C package GNU Linear Programming Kit (GLPK) to solve it. We have found that GLPK works pretty well for our real data analyses in Section 4.

We need to introduce some notations. Since the transport range constraints in Equation (4) impose $\gamma_{ij} = 0$ for $v_j \notin \mathcal{N}_i$, we only need to assign optimal values to $\tilde{\boldsymbol{\gamma}} = \text{Vec}\{\gamma_{ij}, j \in \mathcal{N}_i\} \in R^{N_0 \times 1}$, where $N_0 = \sum_{i=1}^N n_i$ and $\text{Vec}(\cdot)$ denotes the vectorization of a matrix. With this simplification, the dimension of solvable variables is reduced from $O(N^2)$ to $O(N_0)$, which highly increases the computational efficiency of our algorithm. Let A_1 and A_2 be two $N \times N_0$ matrices. The i th row of A_1 consists of 0's except the $(\sum_{j=1}^{i-1} n_j + 1)$ th to $(\sum_{j=1}^i n_j)$ th elements

being 1. Similarly, all the elements of the i th row of A_2 are zeros except the $(\sum_{p=1}^{j-1} n_p + q)$ th element being 1 when grid v_j is indexed by q in the neighboring set \mathcal{N}_j of vertex v_j . Let $\tilde{C} \in R^{N_0 \times 1}$ be the vector including the unit transport costs for all the corresponding $\gamma_{ij}'s \in \tilde{\boldsymbol{\gamma}}$. Moreover, we define

$$A = [A_1 \quad \mathbf{0} \quad \mathbf{0} \quad \mathbf{0} | A_2 \quad \mathbf{I}_N \quad -\mathbf{I}_N \quad \mathbf{0} | A_2 \quad -\mathbf{I}_N \quad \mathbf{0} \quad \mathbf{I}_N]$$

$$\text{and } \mathbf{b} = [\boldsymbol{\mu}^T, \mathbf{v}^T, \mathbf{v}^T]^T,$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)^T$, $A \in R^{3N \times (N_0 + 3N)}$, $\mathbf{b} \in R^{3N}$, and \mathbf{I}_N is an identity matrix.

The (5) is equivalent to $\min\{\|\mathbf{v} - A_2 \tilde{\boldsymbol{\gamma}}\|_1 + \lambda \tilde{C}^T \tilde{\boldsymbol{\gamma}}\}$ subject to $A_1 \tilde{\boldsymbol{\gamma}} = \boldsymbol{\mu}$ and $\tilde{\boldsymbol{\gamma}} \geq 0$. Let $S \in R^{N \times 1}$, it can be further transferred into a standard linear programming (LP)

$$\begin{aligned} \min\{\mathbf{1}^T S + \lambda \tilde{C}^T \tilde{\boldsymbol{\gamma}}\} \quad \text{subject to} \\ A_1 \tilde{\boldsymbol{\gamma}} = \boldsymbol{\mu}, A_2 \tilde{\boldsymbol{\gamma}} + S \geq \mathbf{v}, A_2 \tilde{\boldsymbol{\gamma}} - S \leq \mathbf{v}, \tilde{\boldsymbol{\gamma}} \geq 0, \text{ and } S \geq 0. \end{aligned} \quad (6)$$

The above LP can be further rewritten as

$$\min_X \{B^T X\} \quad \text{subject to } AX = \mathbf{b}, X \geq 0, \quad (7)$$

where $B = (\lambda \tilde{C}^T, \mathbf{1}^T, \mathbf{0}^T, \mathbf{0}^T)^T$ and $X = (\tilde{\boldsymbol{\gamma}}^T, S^T, \mathbf{w}_1^T, \mathbf{w}_2^T)^T$, in which \mathbf{w}_1 and \mathbf{w}_2 are vectors of slack variables. The dual of Equation (7) is assigned as

$$\max_{\mathbf{y} \in R^{3N}} \{\mathbf{b}^T \mathbf{y}\} \quad \text{subject to } A^T \mathbf{y} \leq B, \quad (8)$$

which further reduces the variable dimension from $N_0 + 3N$ to $3N$.

2.4. Applications of GEM in Ride-sourcing Platforms

To calculate GEM, we need to build a dynamic weighted graph structure over time for each city on the ride-sourcing platform as follows. We first divide a city into $|\mathbb{V}| = N$ nonoverlapping hexagons and regard each hexagon as a vertex in \mathbb{V} . Then, we set $\mathcal{N}_i = \cup_{k=0}^2 \mathcal{N}_i^k$, where \mathcal{N}_i^k includes all the neighboring hexagons within the k th outer layer of v_i for $k > 1$ and \mathcal{N}_i^0 only includes v_i itself. A vertex v_j belongs to the k th outer layer of v_i if k steps are required to walk from v_i to v_j on the hexagonal network.

Thus, we determine $G = (\mathbb{V}, \mathbb{E})$. Second, we set $W_t = (w_{ijt})$, where w_{ijt} is the distance between v_i and v_j in the t th timestamp. Note that w_{ijt} may vary with time due to the real-time locations of drivers and customers. Third, we compute $C_t = (c_{ijt})$ by using W_t in the t th timestamp. Finally, we obtain the dynamic weighted graph structure (G, W_t, C_t) .

We show how to use GEM to address three operational tasks of interest in ride-sourcing platforms. First, we can measure the optimal distance between observed dynamic supply and demand networks across time. We extract the spatio-temporal data $\mathbf{O} = \{(o_{it})\}_t$ and $\mathbf{D} = \{(d_{it})\}_t$ from the dynamic demand and supply systems, where o_{it} and d_{it} represent demands and supplies at vertex v_i in the t th timestamp, respectively. Given \mathbf{O} and \mathbf{D} , we set $\mu_t = (d_{it})_i$ and $\nu_t = (o_{it})_i$ and use the LP algorithm to calculate $\rho(t) = \rho_\lambda(\mu_t, \nu_t | G, C_t)$ and its corresponding solution, denoted as $(\tilde{\mu}_{t*} = (\tilde{d}_{it*})_i, \gamma_{t*})$, in the t th timestamp.

Furthermore, we introduce an optimal supply–demand ratio at each v_i in the t th timestamp defined as the ratio of o_{it} over the ‘optimal’ supplies $\tilde{d}_{it*} + \iota_{\{=\}}(\tilde{d}_{it*} = 0)$, denoted as DSr_{it} , in which we add an extra term $\iota_{\{=\}}(\tilde{d}_{it*} = 0)$ to avoid zero in the denominator. Similarly, we can define an optimal supply–demand difference as $DSd_{it} = o_{it} - \tilde{d}_{it*}$ at each (v_i, t) . It allows us to create the spatiotemporal map of GEM-related measures (DSr_{it}, DSd_{it}) . Furthermore, we extend (DSr_{it}, DSd_{it}) to a wide timespan \mathcal{T}_0 within a large region $\mathbb{V}_0 \in \mathbb{V}$. For instance, we define a weighted average supply–demand ratio over \mathbb{V}_0 in \mathcal{T}_0 and a weighted average absolute supply–demand difference over \mathbb{V}_0 in \mathcal{T}_0 as follows:

$$DSr_{\mathcal{T}_0}(\mathbb{V}_0) = \frac{\int_{t \in \mathcal{T}_0} \sum_{i \in \mathbb{V}_0} w_{it} DSr_{it} dt}{\int_{t \in \mathcal{T}_0} \sum_{i \in \mathbb{V}_0} w_{it} dt} \quad \text{and}$$

$$ADSd_{\mathcal{T}_0}(\mathbb{V}_0) = \frac{\int_{t \in \mathcal{T}_0} \sum_{i \in \mathbb{V}_0} w_{it} |DSd_{it}| dt}{\int_{t \in \mathcal{T}_0} \sum_{i \in \mathbb{V}_0} w_{it} dt}, \quad (9)$$

in which we set w_{it} as either o_{it} or $(o_{it} + \tilde{d}_{it*})/2$ in order to highlight vertices with high demands. A good market equilibrium in ride-sourcing platforms corresponds to small values of $|DSd_{it}|$ and $|DSr_{it} - 1|$ across all (v_i, t) . Please see Section 4.1 for details.

Second, we can use historical supply–demand information contained in $\{(DSr_{it}, DSd_{it}) : (v_i, t) \in \mathbb{V} \times \mathcal{T}_0\}$ to design order dispatching policies for large-scale ride-sourcing platforms. Order dispatch is an essential component of any ride-sourcing platform for assigning idle drivers to nearby passengers. Standard order dispatching approaches focus on immediate customer satisfaction such as serving the order with the nearest drivers (Liao 2003) or the first-come-first-go strategy to serve the order on the top of the waiting list with the first driver becoming available (Zhang and Pavone 2016). Those greedy methods, however, fail to account for the spatial effects of an order and driver (O-D) pair on the other O-D pairs. Thus, they may not be optimal from a global perspective. To improve the users’ experience, some more advanced techniques strive to balance between immediate earnings and future returns. To design better dispatching policy, we will include additional historical supply–demand network information based on GEM to delineate its effects on the average expected gain from serving current order. Please see Section 4.2 for details.

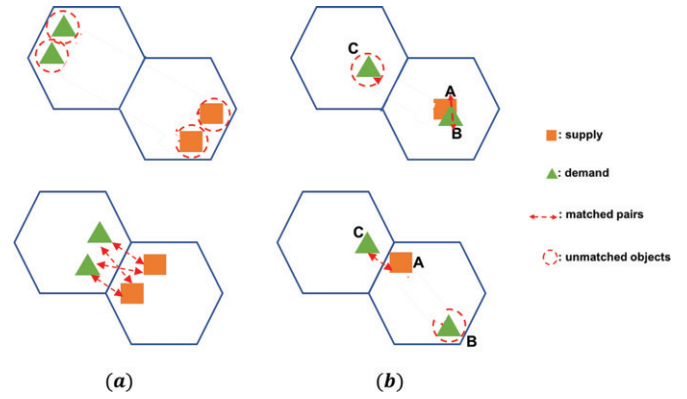


Figure 4. Two examples to illustrate the importance of using random c_{ij} . In Panel (a), two demands and two drivers can only be matched in the lower sub-figure since their corresponding pairs of distance are below a given threshold, whereas it is not the case in the upper sub-figure. In Panel (b), supply A is assigned to demand C in the lower sub-figure when the within-grid cost c_{ij} is nonzero, and to demand B when $c_{ij} = 0$, guided by the optimal transport plan with smaller transport costs.

Third, an important application of $\{\rho(t)\}$ is to use it as a metric to directly compare two (or more) dispatching policies for ride-sourcing platforms. The key idea is to detect whether there exists a significant difference between two sets of GEMs for two competitive policies under the same platform environment. Given the joint distribution of demand and supply in the platform, the smaller GEM is, the better many global operational metrics, such as order answer rate, order finishing rate, and the driver’s working time, are. Compared with those global operational metrics, GEM is a more direct measurement of the operational efficiency for a ride-sourcing platform. Please see Section 4.3 for details.

3. Theoretical Properties

In this section, we study the theoretical properties of our GEM-related methods proposed in Section 2, most of whose proofs can be found in the supplementary document.

First, we establish the convergence property of LP (7) for GEM.

Theorem 3.1. The LP (7) has an optimal basic feasible solution. Furthermore, if X is feasible for the primal problem (7) and \mathbf{y} is feasible for the duality (8), then we have

$$\bar{z} = \mathbf{y}^T \mathbf{b} = \mathbf{y}^T A X \leq B^T X = z \quad (10)$$

If either (7) or (8) has a finite optimal value, then so does the other, the optimal values coincide, and the optimal solutions to both (7) and (8) exist.

An implication of Theorem 3.1 is that Equation (7) is solvable. It demonstrates that there always exist theoretically optimal transport plans (including no transport case) to maximally increase the systematic coherence between the initially unbalanced supplies and demands. However, Theorem 3.1 also indicates that the optimal transport plan may not be unique considering the weighted graph structure and initial supply and demand distributions.

Second, we carry out a probabilistic analysis of our LP (7) for GEM when c_{ij} follows a distribution. Let us start with two

motivating examples of ride-sourcing platforms described in Figure 4. The w_{ij} represents the geological distance or the traffic time, which may vary between each pair of supply at v_j and demand at v_i . For the LP defined in (7), it is assumed that each component of $\tilde{C} = (\tilde{c}_1, \tilde{c}_2, \dots, \tilde{c}_{N_0})^T$ is a nonnegative random variable, whereas all elements in A and \mathbf{b} in Equation (7) are known. Let z^* denote the minimum value of (7). Since z^* is a function of \tilde{C} , it is also a random variable. We provide an upper bound for the expectation of $z^* = z^*(\tilde{C})$ below.

Theorem 3.2 (Expectation Bound). Let $\tilde{c}_1, \dots, \tilde{c}_{N_0}$ be independent nonnegative random variables. Suppose there exist $\alpha_1 \in (0, \infty)$ and $\alpha_2 \in (0, 1]$ such that for $l = 1, 2, \dots, N_0$ and all $h > 0$ with $P(\lambda\tilde{c}_l \geq h) > 0$, we have

$$E(\lambda\tilde{c}_l | \lambda\tilde{c}_l \geq h) \geq \alpha_1 \lambda E(\tilde{c}_l) + \alpha_2 h, \quad (11)$$

where the expectation is taken with respect to \tilde{c}_l . Let $\{\hat{x}_1, \dots, \hat{x}_{3N+N_0}\}$ be any fixed feasible solution to Equation (7). We have

$$E(z^*) \leq \alpha_2^{-1} \left\{ \sum_{l=1}^{N_0} (1 - \alpha_1 \delta_l) E(\lambda\tilde{c}_l) \hat{x}_l + \sum_{l=N_0+1}^{N_0+N} \hat{x}_l \right\}, \quad (12)$$

where $\delta_l \in [0, 1]$ defined in the supplementary document is a pre-defined nonnegative constant for each $l \in \{1, \dots, N_0\}$.

Theorem 3.2 has at least two implications. First, condition (11) holds under some mild conditions. For instance, it can be shown that if \tilde{c}_j is a bounded random variable that takes values in $[c_{j,L}, c_{j,U}]$ such that $P(\tilde{c}_j > c_{j,L}) > 0$ and $\liminf_{h \rightarrow 0} h^{-1} P(\lambda\tilde{c}_j < h + \lambda c_{j,L}) > 0$, then condition (11) holds. Some examples of \tilde{c}_j include uniform, truncated normal, and truncated exponential random variables, among others. For instance, we consider the case that \tilde{c}_j follows Uniform $[c_{j,L}, c_{j,U}]$. It can be shown that $E(\lambda\tilde{c}_j | \lambda\tilde{c}_j \geq h) = 0.5(\lambda c_{j,U} + h)$, yielding $\alpha_1 = c_{j,U} / (c_{j,U} + c_{j,L})$ and $\alpha_2 = 0.5$. Second, Equation (12) gives an upper bound of the expected value of $z^*(\tilde{C})$. If we set $\delta_l = 0$ for all l , then we can obtain a larger upper bound compared with the right-hand side of Equation (12). This result generalizes an existing result of Dyer, Prieze, and Mcdiarmid (1986) for standard linear programs with random costs under a stronger condition corresponding to $\alpha_1 = 1$.

Third, we examine the metric properties of $\rho_\lambda((\cdot), (\cdot)|G, C)$ including non-negativity, identity, symmetry, and the triangle inequality.

Theorem 3.3. The operator $\rho_\lambda((\cdot), (\cdot)|G, C)$ is a semi-metric such that it satisfies nonnegativity, identity, and symmetry, but not necessarily the triangle inequality when (i) $C = (c_{ij}) \in R^{N \times N}$ is symmetric with $c_{ii} = 0$ for all i ; (ii) $j \in \mathcal{N}_i$ if and only if $i \in \mathcal{N}_j$.

Theorem 3.3 indicates that if C is symmetric, then $\rho_\lambda((\cdot), (\cdot)|G, C)$ as a semi-metric satisfies three properties including non-negativity, identity, and symmetry. Although the symmetric assumption of C may be incorrect for all vertexes, it should be valid for most vertexes. Thus, $\rho_\lambda((\cdot), (\cdot)|G, C)$ is approximately a semi-metric.

Fourth, we give the upper and lower bounds of GEM and consider an additivity property in order to better understand how the transport costs and network structures affect GEM.

Theorem 3.4. The following properties hold:

- (i). $\|\mu\| - \|v\| \leq \rho_\lambda(\mu, v|G, C) \leq (\|\mu\| + \|v\|)$
- (ii). Additivity Property. For a nonnegative Δ , when C is symmetric, we have

$$\begin{aligned} |\rho_\lambda(\mu, v + \Delta|G, C) - \rho_\lambda(\mu, v|G, C)| &\leq N\Delta; \\ |\rho_\lambda(\mu + \Delta, v|G, C) - \rho_\lambda(\mu, v|G, C)| &\leq N\Delta. \end{aligned}$$

Property (i) shows that GEM can be bounded from both above and below. Based on the additivity property, the GEM value can either increase or decrease with one-side node-wise augmentation, which depends on the weighted graph structure and the distribution of supply and demand. This indicates that applying proper stimulus at selected vertexes is more efficient than globally increasing supply resources.

Fifth, we examine the weak convergence property of $\rho_\lambda((\cdot), (\cdot)|G, C)$.

Theorem 3.5 (Weak Convergence). Let $\{\mu_n\}$ be a sequence of measures on space \mathbb{V} , and $\mu_n, \mu \in M_+(\mathbb{V})$. If all the transport costs are bounded, that is $c_{ij} \leq R$ holds for $\forall v_i \in \mathbb{V}$ and $v_j \in \mathcal{N}_i$, then $\rho_\lambda(\mu, \mu_n|G, C) \rightarrow 0$ when $\mu_n \rightarrow \mu$ and $\{\mu_n\}$ is tight.

Here is an immediate corollary of **Theorem 3.5**.

Corollary 3.5.1. Let $\{\mu_n\}$ and $\{v_n\}$ be two sequences of measures on space \mathbb{V} , and $\mu_n, v_n, \mu, v \in M_+(\mathbb{V})$. If $c_{ij} \leq R$ holds for $\forall v_i \in \mathbb{V}$ and $v_j \in \mathcal{N}_i$, then we have

$$\begin{aligned} \text{if } \mu_n \text{ (resp. } v_n) \rightarrow \mu \text{ (resp. } v) \text{ and } \{\mu_n\}, \{v_n\} \text{ are tight,} \\ \text{then } \rho_\lambda(\mu_n, v_n|G, C) \rightarrow \rho_\lambda(\mu, v|G, C). \end{aligned}$$

Theorem 3.5 states that the GEM value goes to 0 and no transport is required when the initial distributions of μ and v are getting close to each other.

4. Experiments

In this section, we apply GEM to the supply–demand diagnostic data set in order to address three important operational tasks in ride-sourcing platforms, including answer-rate prediction, the design of order dispatching strategy, and policy assessment. Without special saying, we use the method described in **Section 2.4** to construct the dynamic weighted graph structure across time in all these analyses. We have released the “supply–demand diagnostic dataset” through the DiDi GAIA Open Data Initiative at <https://outreach.didichuxing.com/appEn-vue/dataList> and made the computer codes together with necessary files available at <https://github.com/BIG-S2/GEM>.

4.1. Answer-Rate Prediction

The data set that we use here includes both demand and idle driver information from April 21 to May 20, 2018 in a large city H. We divide the whole city into $N = 800$ nonoverlapping hexagonal sub-regions with side length being 1400 m to form the whole vertex set \mathbb{V} . We let the directed edge weight w_{ij} from v_i to $v_j \in \mathcal{N}_i^1$ be the distance between the centers of the two sub-regions, which is 2400 m if v_j can be directly reached by v_i through traffic without first passing through another vertex.

Otherwise, $w_{ij} = \infty$. We compute the numbers of idle drivers and demands in each vertex per minute and then extract the dynamic supply–demand data set.

The aim of this data analysis is to examine whether the GEM-related measures, such as DSr_{it} , are useful for predicting order answer rate in ride-sourcing platforms. Order answer rate is defined as the number of orders accepted by drivers divided by the total number of orders in a fixed time interval. Specifically, we predict the log-value of order answer rate of the incoming 10 (or 60) minutes by using historical metric values. We computed the Hellinger distance, the L_2 distance, the Wasserstein distance, and GEM for each 10-min interval. The L_2 distance is calculated by using the numbers of orders and available drivers in all vertices across 10 consecutive 1-min timestamps. The Hellinger distance is calculated by normalizing the numbers of orders and available drivers in all vertices and across 10 consecutive 1-min timestamps into probability distributions. For the Wasserstein distance, we first normalize both supplies and demands at each 1-min time interval into two probability distributions and calculate their corresponding Wasserstein distance. Subsequently, we obtain the metric value over each 10-min interval by aggregating the Wasserstein distances computed across the 10 included 1-min timestamps by using their corresponding weights $\sum_{v_j \in V} o_{it_k} / \sum_{t_k \in T} \sum_{v_j \in V} o_{it_k}$. For GEM, we compute the supply–demand ratio map of DSr_{it} per minute and then we calculate $DSr_{T_0}(V)$ for each 10-min interval.

We split the supply–demand dataset into a training dataset consisting of observations from April 26 to May 11, 2018, and a test dataset consisting of observations from May 12 to May 21, 2018. We use linear regression models to predict the log-value of order answer rate of the incoming j th 10 min for $j = 1, \dots, 6$ by using various historical metric variables of the previous $p = 10$ 10-min snapshots and those in the same time windows of the previous 5 days.

We use mean absolute percentage error (MAPE) and root mean squared error (RMSE) as evaluation metrics to examine the prediction accuracy of all the four compared metrics. Table 1 shows their corresponding RMSE and MAPE values based on the test data. Due to the space limitation, we only provide the results corresponding to those at $t + 10$ and $t + 60$ min, which

Table 1. Results from the answer-rate prediction.

			Hellinger	L2-distance	Wasserstein	GEM
t+10	All time	RMSE	0.1362	0.1496	0.1273	0.0552
		MAPE	0.0801	0.0891	0.0718	0.0338
	Peak hour	RMSE	0.2219	0.2187	0.2088	0.0614
		MAPE	0.1494	0.1457	0.1089	0.0422
t+60	All time	RMSE	0.1522	0.1552	0.1413	0.1130
		MAPE	0.0828	0.0868	0.0859	0.0620
	Peak hour	RMSE	0.2395	0.2565	0.2222	0.1530
		MAPE	0.1077	0.1159	0.1317	0.0728

indicate the short-term and long-term prediction capacities of all the four metrics. Moreover, we also include the results during the evening peak hours starting from 6 p.m. to 8 p.m. For both the $t + 10$ and $t + 60$ cases, GEM significantly outperforms all other three metrics, which may not sufficiently capture the dynamic transport and systematic balance of the weighted graph structure.

Figure 5(a) presents the real order answer rates and their predictive values in the last 7 test days (Tuesday to Monday) from May 12 to May 18 based on all the four metrics for the ($t + 10$) case. Compared with all other methods, GEM shows higher consistency between the true and predicted answer rate values, especially for some abnormal extreme cases. Furthermore, Figure 5(b) presents the histograms of RMSEs for the Hellinger distance, the Wasserstein distance, and GEM at each day of the last 7 days, indicating that GEM outperforms the other two metrics consistently in all seven days. Therefore, our GEM is able to capture the short- and long-term variability within the coherence between the two spatial–temporal systems and has strong prediction capacity for future answer rates.

4.2. Order Dispatching Policies

We consider the order dispatching problem of matching N_o orders with N_d available idle drivers, where N_o and N_d denote the total number of orders and that of idle drivers in the current timestamp, respectively. The edge weight $A(k, l)$ in the bipartite graph equals to the expected earnings when pairing driver l



Figure 5. Results from the answer-rate prediction. Panel (a): comparisons of the log-value of real answer rates obtained from May 12 to May 18, 2018 and their predictive values based on the L_2 distance, the Hellinger distance, the Wasserstein distance, and GEM. Panel (b): comparisons of day-wise RMSEs of answer rate prediction obtained from Monday to Sunday within the whole city area.

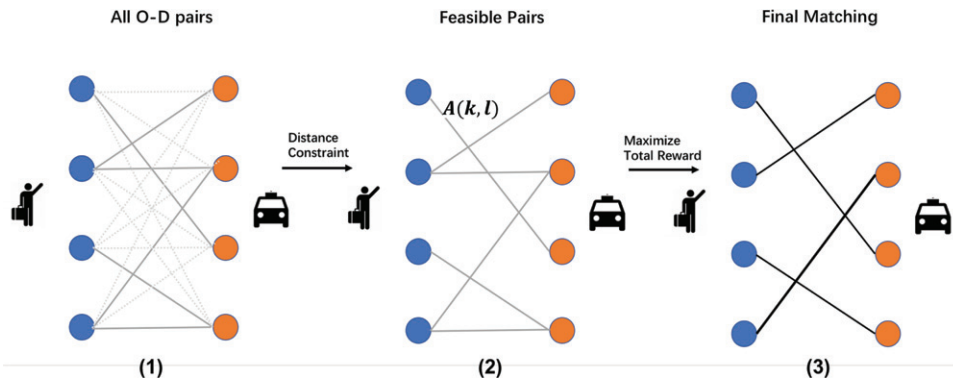


Figure 6. The order dispatch as a bipartite matching problem: (a) available orders and drivers prepared for pairing; (b) quantifying all the potential expected earning $A(k, l)$ for all driver-order pairs (k, l) that satisfy the dispatching constraints; and (c) finding the optimal one-to-one bipartite matching in order to maximize the total revenue.

to order k . Let x_{kl} be 1 if order k is assigned to driver l and 0 otherwise. The global order dispatching algorithm solves a bipartite matching problem as follows:

$$\arg \max_{x_{kl}} \sum_{k=0}^{N_d} \sum_{l=0}^{N_o} A(k, l) x_{kl}, \quad \text{s.t.} \quad \sum_{k=0}^{N_d} x_{kl} \leq 1 \quad \forall l; \\ \sum_{l=0}^{N_o} x_{kl} \leq 1 \quad \forall k; x_{kl} = 0 \text{ if } c_{kl} > \epsilon \quad \forall k, l. \quad (13)$$

See Figure 6 for a graphical illustration of Equation (13). The constraints ensure that each order can be paired to at most one available driver and similarly each driver can be assigned to at most one order. In practice, only drivers within a certain distance could serve the corresponding orders, which means that x_{kl} 's are forced to be 0 when the distance between order k and driver l , denoted as c_{kl} , is beyond the maximal pick-up distance ϵ . The state-of-art algorithm to solve this kind of matching problem is the Kuhn-Munkres (KM) algorithm (Munkres 1957), which will be used to solve the formulated problem here.

In this article, we compare three different dispatching policies based on three different formulations of $A(k, l)$. The first one as a baseline only considers the immediate reward of assigning driver l to order k , which is defined as $A^{(1)}(k, l) = \alpha_1 r_k - \alpha_2 c_{kl}$, where r_k is the driver's earning by serving order k and c_{kl} is the pick-up distance between order k and driver l . Moreover, α_1 and α_2 are tuning parameters such that the two terms are balanced to maximize the drivers' salaries, while reducing customers' waiting time.

The second one is given by $A^{(2)}(k, l) = \alpha_1 r_k - \alpha_2 c_{kl} + \alpha_3 \{\eta^{\Delta t_{lk}} V_1(s'_{lk}) - V_1(s_l)\}$, where η is the discount factor and an additional term $\alpha_3 \{\eta^{\Delta t_{lk}} V_1(s'_{lk}) - V_1(s_l)\}$ is introduced to enhance the long-term effects of current actions on drivers' future income (Xu et al. 2018). Let $V_1(s)$ be the expected earnings from now to the end of the day for a driver located at $s = (v, t)$, where $v \in \mathbb{V}$ and t is the current time. Moreover, $s_l = (v^{(l)}, t)$ and $s'_{lk} = (v^{(k)}, t + \Delta t_{lk})$ here represent the current spatial-temporal state of driver l and his/her estimated finishing state when completing serving order k , where $v^{(l)} \in \mathbb{V}$ is the current region of driver l before order assignment and $v^{(k)} \in \mathbb{V}$ is the destination region of order k and Δt_{lk} denotes the total time required for driver l to finish the whole process of serving

order k . If a driver becomes available to a new order immediately after finishing the ongoing one, then $\eta^{\Delta t_{lk}} V_1(s'_{lk}) - V_1(s_l)$ is the extra future earning for driver l by serving order k other than staying idle.

The third one is given by $A^{(3)}(k, l) = A^{(2)}(k, l) + \alpha_4 \{\eta^{\Delta t_{lk}} V_2(s'_{lk}) - V_2(s_l)\}$, where $\alpha_4 \{\eta^{\Delta t_{lk}} V_2(s'_{lk}) - V_2(s_l)\}$ is introduced to balance the supply-demand coherence. Moreover, $V_2(s) = v_t(v) - \tilde{\mu}_t(v)$ at $s = (v, t)$ is calculated from GEM in Equation (5). Using $V_2(\cdot)$ increases the probability that customers' requests can be quickly answered by nearby drivers, whereas $V_1(\cdot)$ ignores the interaction effects when multiple drivers are heading to the same location. Thus, when the future demand has already been fulfilled by drivers re-allocated by previous completed servings, assigning more drivers might decrease $V_1(\cdot)$ in the target location.

We use a comprehensive and realistic dispatch simulator designed for recovering the real online ride-sourcing system to evaluate the three dispatching policies. The simulator models the transition dynamics of the supply and demand systems to mimic the real on-demand ride-hailing platform. The order demand distribution of the simulator is generated based on historical data. The driver supply distribution is initialized by historical data at the beginning of the day, and then evolves following the simulator's transition dynamics (including drivers getting online/offline, driver movement with passengers and idle driver random movement) as well as the order dispatching policies. The differences between the simulated results and the real-world situation is less than 2% in terms of some important metrics, such as the drivers' revenue, answer rate, and idle driver rate.

To compare the three dispatching policies, we randomly selected a specific city S, which usually has in total 150, 000 to 200, 000 ride demands per day. We still divide the whole city area into $N = 800$ hexagonal vertices and use the geological distance between two nearby grids to be the edge weights. Furthermore, three different days including 2018/05/15 (Tuesday), 2018/05/18 (Friday), and 2018/05/19 (Saturday) were analyzed since the global order answer rates on weekday are usually much lower than those at weekend by looking at the historical data. Both $V_1(\cdot)$ and $V_2(\cdot)$ values were obtained by taking the average of the same weekday or weekend from the previous four weeks since the platform has significant weekly periodicity. The length of time intervals that we used to compute $V_1(\cdot)$ and $V_2(\cdot)$ was

set to be $\mathcal{T}_0 = 10$ min so that all the action windows inside share the same $V_1(\cdot)$ and $V_2(\cdot)$ values. Specifically, $V_2(\cdot)$ is achieved by aggregating the $|\mathcal{T}_0|$ continuous $(v_{kt} - \tilde{\mu}_{kt})$ s. We applied the three dispatching policies with different edge weights to the simulator even based on the same initial input and transition dynamics. We set $\alpha_1 = 1$ and $\alpha_2 = 0.001$ to rescale the order price r_k and the pick-up distance c_{kl} into comparable ranges. The r_k contributes more to the variations of $A(k, l)$ because of the constrained pick-up distance ($c_{kl} \leq \epsilon$). Furthermore, we perform grid search for a wide range of $(\alpha_3; \alpha_4)$ combinations to find its optimal solution, denoted as (α_3^*, α_4^*) , that maximizes an average drivers' revenues for weekdays and weekends in the simulator. Specifically, we fixed $\alpha_4 = 0$ first and use the bisection method to obtain a rough value range of length 0.1 for α_3 with its initial start being $[0, 1]$. Then we apply the grid search method to increase 0.01 amount for α_3 each time within the value range until finding the optimal α_3^* corresponding to the largest averaged drivers' revenue. Subsequently, we fix α_3^* and do the similar grid search to get the optimal α_4^* .

Tables 2 and 3 summarize the collected results corresponding to the baseline policy, $A^{(2)}(k, l)$ with the optimal α_3 , and our

Table 2. Results from order dispatching policies. Comparisons of the three policies with respect to two evaluation metrics including the drivers' revenue and the global answer rate using the simulator for city S on two selected Weekdays.

α_3	α_4	Drivers' revenue (Yuan)	Order answer rate
2018/05/15 (Tuesday)			
0	0	1191316	0.737
0.54	0	1227175(+3.01%)	0.760(+3.12%)
0.54	6	1235037(+3.67%)	0.761(+3.28%)
0.54	7	1236824(+3.82%)	0.763(+3.54%)
0.54	8	1240518(+4.13%)	0.765(+3.82%)
0.54	9	1238850(+3.99%)	0.764(+3.66%)
0.54	10	1231702(+3.39%)	0.761(+3.26%)
2018/05/18 (Friday)			
0	0	13943666	0.539
0.61	0	14701230(+5.43%)	0.557(+3.34%)
0.61	6	14845486(+6.47%)	0.561(+4.08%)
0.61	7	14858400(+6.56%)	0.560(+3.90%)
0.61	8	14865454(+6.61%)	0.560(+3.90%)
0.61	9	14867573(+6.63%)	0.560(+3.90%)
0.61	10	14823121(+6.31%)	0.557(+3.34%)

NOTES: The rows with $(\alpha_3, \alpha_4) = (0, 0)$ correspond to the first (or baseline) policy, those with $\alpha_4 = 0$ and $\alpha_3 \neq 0$ correspond to the second policy, and all other rows correspond to the third policy. The numbers in the parentheses denote the relative improvement of the corresponding policy over the baseline policy for each evaluation metric.

Table 3. Results from order dispatching policies.

α_3	α_4	Drivers' revenue (Yuan)	Order answer rate
2018/05/19 (Saturday)			
0	0	13507568	0.745
0.52	0	13886185(+2.80%)	0.768(+3.09%)
0.52	6	14034453(+3.90%)	0.774(+3.89%)
0.52	7	14008847(+3.71%)	0.772(+3.62%)
0.52	8	14043995(+3.97%)	0.773(+3.76%)
0.52	9	13996021(+3.62%)	0.770(+3.36%)
0.52	10	13934895(+3.16%)	0.768(+3.09%)

NOTE: Comparisons of the three policies with respect to two evaluation metrics including the drivers' revenue and the global answer rate using the simulator for city S on a selected Weekend.

approach with different α_4 values. It reveals that the order dispatching policy based on $A^{(3)}(k, l)$ could achieve higher drivers' revenue and answer rate compared with the other two policies. The optimal α_3 is achieved at 0.54, 0.61, and 0.52 for 2018/05/15, 2018/05/18, and 2018/05/19, respectively. In 2018/05/15 and 2018/05/19, we obtain a smaller optimal α_3 since a higher coherence between supplies and demands is achieved under the baseline policy ($\alpha_3 = \alpha_4 = 0$) than that of 2018/05/18, which indicates that the supply-demand relationship is more related to the policy efficiency than the weekday/weekend status. Moreover, the supply abundance in 2018/05/15 and 2018/05/18 results in a higher order answer rate but a smaller optimal α_4 . Compared to the policy corresponding to $A^{(2)}(k, l)$, adding the GEM-related measurements increases the expected whole-day answer rate and drivers' revenue in more than 1%. It may indicate that the supply-demand difference may affect the expected future gain of a marginal driver.

In practice, we first perform grid search for a wide range of (α_3, α_4) combinations to find its optimal solution (α_3^*, α_4^*) that maximizes the mean drivers' revenue for some representative days in the simulator. Then we fine-tune the parameters via on-line A/B testing, and apply the policy in the real-life dispatching system. The value functions $V_1(\cdot)$ and $V_2(\cdot)$ are updated with the new policy being employed for a period of time, and α_3^* and α_4^* are re-tuned in the real environment to achieve the optimal efficiency.

4.3. Policy Evaluation

We conduct an experiment using another supply-demand dataset of the same city H from December 3rd to December 16th, 2018 in order to compare the effectiveness between two order dispatching policies. We executed them alternatively on successive half-hourly time intervals. Moreover, we start with the baseline policy for the first half hour and change the policy every half hour through the whole day and reverse their order in another day. We include an A/A test, which compares the baseline policy against itself, by using the historical data obtained from November 12 to November 25 as a direct comparison. We calculate GEM within each time window of 30 min as follows. There are in total $M_T = 48$ time intervals per day. To obtain GEM in each time interval \mathcal{T} , we aggregate 30 GEM values, each of which is calculated within the 1-min timestamp, by using normalization weights $o_{it} / \sum_{t \in \mathcal{T}} \sum_{v_i \in \mathbb{V}} o_{it}$.

We first need to introduce some notations. We denote $y_m(t_k)$ as the aggregated GEM value and use $x_m(t_k)$ to denote a 2×1 vector of predictors, which are not strongly influenced by order dispatching policy, including the total number of demands and the total supply time of all drivers in the k th time interval of day m for $k = 1, \dots, M_T$ and $m = 1, \dots, M_D$. Let $a_m(t_k) = 1$ if the new policy is used and $= -1$ otherwise. To examine the marginal effect of policy on GEM, we consider the following regression model:

$$y_m(t_k) = \beta_0(t_k) + \beta_1(t_k)^T \{x_m(t_k) - \bar{x}(t_k)\} + \beta_2(t_k)a_m(t_k) + \eta_m(t_k) + \varepsilon_m(t_k), \tag{14}$$

where $\beta(t_k) = (\beta_0(t_k), \beta_1(t_k)^T, \beta_2(t_k))^T$ is a vector of regression coefficients at t_k , and $\bar{x}(t_k)$ is the sample mean of all

$x_i(t_k)$ s for $k = 1, \dots, M_T$. In addition, we assume that $\eta_m = (\eta_m(t_1), \dots, \eta_m(t_{M_T}))^T$ and $\varepsilon_m = (\varepsilon_m(t_1), \dots, \varepsilon_m(t_{M_T}))^T$ are $M_T \times 1$ vectors of random errors, following mutually independent multivariate Gaussian distributions $N(\mathbf{0}, \Sigma_\eta)$ and $N(\mathbf{0}, \sigma_\varepsilon^2 \cdot \mathbf{I}_{M_T})$, where Σ_η is an $M_T \times M_T$ matrix and σ_ε^2 is a positive scalar. We are interested in testing the following null and alternative hypotheses:

$$H_0 : \int_0^{M_T} \beta_2(t)dt = 0 \text{ v.s. } H_1 : \int_0^{M_T} \beta_2(t)dt \neq 0, \quad (15)$$

where $\int_0^{M_T} \beta_2(t)dt \approx \sum_{k=1}^{M_T} \beta_2(t_k)\Delta t_0$ denotes the average treatment effect per day, in which Δt_0 is the length of each time interval. We propose a joint estimation procedure based on Generalized Estimating Equations (GEE) to iteratively estimate all unknown parameters until a specific convergence criterion being reached (Liang and Zeger 1986). Subsequently, we compute the t -test statistic associated with the average treatment effect per day and its corresponding one-sided (or two-sided) p -value (Mancl and DeRouen 2001).

Furthermore, we consider three global operational metrics including the order answer rates, order finishing rate, and gross merchandise value (GMV) as $y_m(t_k)$ in model (14). We fit the corresponding three regression models in order to study whether the new dispatching policy significantly improves the ride-sourcing platform at the operational level.

Table 4 summarizes all regression analysis results for both the A/A and A/B experimental designs. We can see that in the A/B experimental design, there exists a significant increase in

Table 4. Results from the policy evaluation: relative improvement and two-sided p -value of average treatment effects for the A/A and A/B experiments.

Experiment Design	$y_m(t)$	Relative Improvement(%)	p -value
A/B	Answer Rate	0.76	1.16e-12
	Finish Rate	0.36	4.32e-3
	GMV	0.86	2.91e-6
	GEM	-0.80	4.06e-2
A/A	Answer Rate	0.01	0.96
	Finishing Rate	0.01	0.96
	GMV	-0.08	0.72
	GEM	-0.25	0.43

NOTES: Comparisons of Hellinger, L_2 -distance, Wasserstein, and GEM in predicting answer rate at $t + 10$ and $t + 60$ min. Peak hour denotes the time from 6 p.m. to 8 p.m. MAPE and RMSE denote the mean absolute percentage error and root mean squared error, respectively.

the mean answer rate, finishing rate and gross merchandise value when replacing the old policy by the new one since all the p -values associated with the average treatment effect are smaller than 10^{-3} . The new policy can also significantly reduce the GEM value (p -value smaller than 0.05), which agrees with our assumption that GEM can sufficiently quantify the supply-demand relationship and subsequently affect the examined platform indexes. In contrast, Table 4 shows that in the A/A experimental design, all the four metrics do not show significant treatment effect at the significance level of 5%.

Figure 7(a) presents the GEM value at in total 48 30-min time windows on December 3, 2018 for the A/B experimental design. We observe a significant reduction of GEM value when

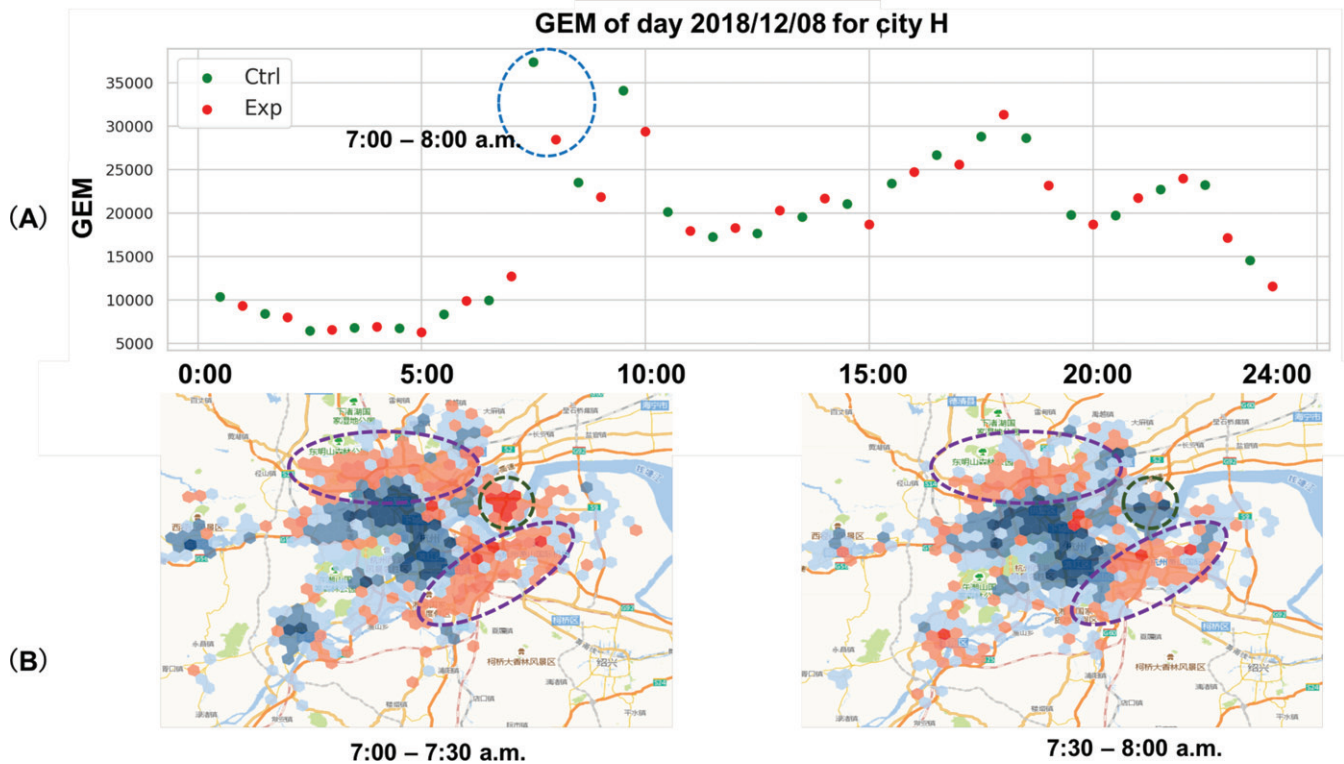


Figure 7. Results from the policy evaluation. (a) The GEM values of a randomly selected day on 2018/12/08 for city H at 30-min scale. Green and red points represent the GEM values generated by the baseline (control) and new (experimental) policies, respectively. In particular, we mark the time period 7:00 a.m. to 8:00 a.m. by a blue circle, which demonstrates a significant reduction of GEM value when changing the policy from the control one to the experimental one. (b) The heatmaps of vertex-wise supply-demand difference DSd_{it} of city H under the control and experimental policies within the 30-min time window from 7:00 to 7:30 a.m. and that from 7:30 to 8:00 a.m., respectively. Hexagons in red and blue colors represent the locations with positive and negative DSd_{it} , respectively, and a deeper color corresponds to a big $|DSd_{it}|$ value.

changing the policy from the control one to the experimental one during the time period from 7:00 to 8:00 a.m. Figure 7(b) presents the heat maps of vertex-wise DSd_{it} within the same time period under the control and experimental policies, respectively. The customer requests in three selected regions marked by green and purple circles were satisfied by the drivers in nearby regions, resulting in the higher supply–demand coherence and thus a smaller GEM value.

Supplementary Material

In the supplementary material, we provide the detailed technical proofs for Theorems 3.1–3.5.

Acknowledgments

This work was done when Dr Hongtu Zhu took the leave of absence from the University of North Carolina at Chapel Hill. The readers are welcome to request reprints from Dr. Hongtu Zhu. Email: bowenhongtu@gmail.com; Phone: 919-966-7272.

Funding

Dr. Fan Zhou's work was partially supported by National Natural Science Foundation of China (12001356), Shanghai Sailing Program (20YF1412300), Fundamental Research Funds for the Central Universities.

References

- Ambrosio, L., and Gangbo, W. (2008), "Hamiltonian Odes in the Wasserstein Space of Probability Measures," *Communications on Pure and Applied Mathematics*, 61, 18–53. [1689]
- Arjovsky, M., Chintala, S., and Bottou, L. (2017), "Wasserstein Generative Adversarial Networks," in *International Conference on Machine Learning*, pp. 214–223. [1689]
- Cha, S.-H. (2007), "Comprehensive Survey on Distance/Similarity Measures Between Probability Density Functions," *City*, 1, 1. [1690]
- Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2018), "Scaling Algorithms for Unbalanced Optimal Transport Problems," *Mathematics of Computation*, 87, 2563–2609. [1690,1691]
- Dyer, M. E., Prieze, A. M., and McDiarmid, C. J. H. (1986), "On Linear Programs With Random Costs," *Mathematical Programming*, 35, 3–16. [1694]
- Grenander, U. and Miller, M. (2007), *Pattern Theory From Representation to Inference*. Oxford: Oxford University Press. [1689]
- Liang, K.-Y., and Zeger, S. L. (1986), "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 73, 13–22. [1698]
- Liao, Z. (2003), "Real-time Taxi Dispatching Using Global Positioning Systems," *Communications of the ACM*, 46, 81–83. [1693]
- Liero, M., Mielke, A., and Savaré, G. (2018), "Optimal Entropy-transport Problems and a New Hellinger–Kantorovich Distance Between Positive Measures," *Inventiones Mathematicae*, 211, 969–1117. [1690,1691]
- Mancl, L. A., and DeRouen, T. A. (2001), "A Covariance Estimator for Gee With Improved Small-sample Properties," *Biometrics*, 57, 126–134. [1698]
- Munkres, J. (1957), "Algorithms for the Assignment and Transportation Problems," *Journal of the Society for Industrial and Applied Mathematics*, 5, 32–38. [1696]
- Piccoli, B., and Rossi, F. (2014), "Generalized Wasserstein Distance and Its Application to Transport Equations With Source," *Archive for Rational Mechanics and Analysis*, 211, 335–358. [1691]
- Rabin, J., and Papadakis, N. (2015), "Convex Color Image Segmentation With Optimal Transport Distances," in *International Conference on Scale Space and Variational Methods in Computer Vision*, ed. Nicolas Papadakis, Berlin: Springer, pp. 256–269. [1689]
- Solomon, J., Rustamov, R., Guibas, L., and Butscher, A. (2014), "Wasserstein Propagation for Semi-supervised Learning," in *International Conference on Machine Learning*, pp. 306–314. PMLR. [1689]
- Steele, J. M. (1987), *Probability Theory and Combinatorial Optimization*. Philadelphia, PA: Society for Industrial and Applied Mathematics. . [1690]
- Villani, C. (2008), *Optimal Transport: Old And New*, Vol. 338. Berlin, Heidelberg: Springer-Verlag. [1690]
- Wang, H., and Yang, H. (2019), "Ridesourcing Systems: A Framework and Review," *Transportation Research Part B: Methodological*, 219, 122–155. [1688]
- Xu, Z., Li, Z., Guan, Q., Zhang, D., Li, Q., Nan, J., Liu, C., Bian, W., and Ye, J. (2018), "Large-Scale Order Dispatch in On-demand Ride-hailing Platforms: A Learning and Planning Approach," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. New York: ACM, pp. 905–913. [1696]
- Zhang, R., and Pavone, M. (2016), "Control of Robotic Mobility-on-demand Systems: A Queuing-Theoretical Perspective," *The International Journal of Robotics Research*, 35, 186–203. [1693]