Spatial Price Integration in Commodity Markets with Capacitated Transportation Networks

John R. Birge
Booth School of Business, University of Chicago, Chicago, Illinois, 60637, john.birge@chicagobooth.edu

Timothy C. Y. Chan
Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 3G8, Canada, tcychan@mie.utoronto.ca

J. Michael Pavlin
Lazaridis School of Business and Economics, Wilfrid Laurier University, Waterloo, Ontario N2L 3C5, Canada, mpavlin@wlu.ca

Ian Yihang Zhu
Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 3G8, Canada, i.zhu@mail.utoronto.ca

Spatial price integration is extensively studied in commodity markets as a means of examining the degree of integration between regions of a geographically diverse market. Many commodity markets that are commonly studied are supported by stable and well-defined transportation networks. In this paper, we analyze the relationship between spatial price integration, i.e., the distribution of prices across geographically distinct locations in the market, and the features of the underlying transportation network. We characterize this relationship and show that price integration is strongly influenced by the characteristics of the network, especially when there are capacity constraints on links in the network. Our results are summarized using a price decomposition which explicitly isolates the influences of market forces (supply and demand), transportation costs and capacity constraints among a set of equilibrium prices. We use these theoretical insights to develop a unique discrete optimization methodology to capture spatiotemporal price variations indicative of underlying network bottlenecks. We apply the methodology to gasoline prices in the southeastern U.S., where the methodology effectively characterizes the price effects of a series of well-documented network disruptions, providing important implications for operations and supply chain management.

Key words: commodity and energy operations; price integration; spatial price equilibrium; supply chain management; network disruptions; congestion; time series analysis; mixed integer programming

1. Introduction

Spatial price integration, defined as the co-movement of prices in a market with geographically separated market participants, is studied extensively in commodity markets. Prices from spatially separated locations that move together are taken as evidence of strong market integration, suggesting that the underlying market is efficient, competitive, and sufficiently well-connected for price

differences to be quickly arbitraged away. On the other hand, prices that are not strongly integrated may suggest the existence of transportation or trade frictions. Studying and characterizing such frictions is of great interest to market participants and policy-makers alike, as it can provide new insights into investment opportunities, trading practices, supply security and consumer welfare. Given the relative ease of acquiring commodity price data over large geographies, measures of price integration are attractive proxies for market integration in large scale studies.

A range of time series econometric methods are typically employed to study price integration in commodity markets; see Dukhanina and Massol (2018) for a review of methods. However, many of these general-purpose empirical methods rely on assumptions that may not be consistent with the equilibrium conditions of a *networked market*, i.e., a market where market participants interact through a transportation network. A notable example illustrating this inconsistency is the notion of a *neutral band* existing between prices at two locations with costly bidirectional transportation (e.g., Goodwin and Piggott 2001). Large price variations can occur within a neutral band defined by the transportation costs without an error-correction (i.e., arbitrage) mechanism, even when the market is efficient. In this pairwise setting, it has been shown that commonly-used concepts such as cointegration are neither necessary nor sufficient for the identification of unexploited arbitrage opportunities or bottlenecks (McNew and Fackler 1997). This example highlights the importance of considering the structural equilibrium conditions when examining spatial price integration.

In this paper, we aim to provide a deeper understanding of spatial price integration in competitive commodity markets with well-defined but capacitated transportation networks. Oil and gas are examples of such markets, where decades of market reforms have led to high levels of competition and where locations typically trade through a stable network of pipelines, railroads and tankers. In these markets, the structure of the network, the transportation costs, and the capacity of the transportation infrastructure may have important implications for price integration. This analysis is increasingly relevant as the growth in supply and demand in these markets consistently outpaces the growth in new infrastructure. In this setting, we establish fundamental connections between spatial price integration and the transportation network. The theoretical results are then used to derive principled methods for market analysis.

We model the market as a network with nodes representing market participants and directed links representing transportation. We study price formation using a spatial price equilibrium model (SPE), which we use to characterize the relationship between prices and market structure. Our results extend the neutral band concept, previously examined only for pairs of nodes with direct connections (i.e., "pairwise" neutral band) to uncapacitated networks where nodes may be connected indirectly through a series of links (i.e., "network" neutral band). We then focus on price

integration in the presence of capacitated links. We study how price shocks generated from bottlenecks are distributed over the market. We then leverage the results of the SPE model to generate a principled and scalable empirical methodology to identify these shocks in the market. Our methodology uses only spatial price data to identify time periods and locations with temporarily inflated prices indicative of capacity constraints in the transportation network. Finally, we demonstrate through a numerical case study using price data alone that our methodology accurately identifies spatiotemporal price variations consistent with well-documented disruptions in the Southeastern U.S. gasoline market. For the remainder of the paper, we use the term *market structure* to refer to the transportation network, which is defined by the network structure (nodes and links), transportation costs on the links, and link capacities.

Our specific contributions and their organization within the paper are as follows:

1. We provide a novel characterization of the relationship between market structure and price integration over general transportation networks. Our results are derived with arbitrary demand and supply functions allowing us to isolate the effect that different components of the market structure have on bounding price differences within the market.
   (a) Uncapacitated and costless transportation (Section 4.1): We show that a structural property of the network, defined as *structural integration*, is a necessary and sufficient condition for the *law of one price* to hold over the nodes.
   (b) Uncapacitated but costly transportation (Section 4.2): We prove existence of a *network neutral band*, which bounds the distribution of prices over the market, and is entirely characterized by the parameters of the network and is independent of the market participants. This result extends the classical *pairwise neutral band* concept to pairs of non-adjacent nodes in the network.
   (c) Capacitated and costly transportation (Section 4.3): We characterize how nodal prices incorporate *congestion surcharges* throughout the network. In particular, we relate these surcharges to the shadow price of capacity constraints in the underlying market allocation problem.
2. Using the previous insights, we identify a novel decomposition of the nodal prices that we use to develop a *surcharge estimation model* (SEM), a unique discrete optimization approach for price-based time series analysis. The SEM is a tractable and interpretable methodology for capturing market characteristics and spatiotemporal variations in price data that are indicative of bottleneck constraints in the underlying market (Section 5).
3. We present a comprehensive case study of the Southeastern U.S. gasoline market, where we show that our methodology can accurately identify spatiotemporal variations in prices that are consistent with well-documented transportation capacity disruptions (Section 6). The results

are used to provide quantitative insights into the cost of the transportation bottlenecks across the market footprint. They are also used to highlight the effects of network topology and operations management on the spatiotemporal distribution of these price shocks.

The main results in the paper are presented under the assumption of a competitive commodity market where the underlying network structure is unchanged over a given time horizon. Our analysis lends itself particularly well to energy markets, where the market structure (e.g., the network of pipelines, railroads or available shipping routes) is generally fixed in the short-term and where the effects of transportation bottlenecks are most pronounced. These markets will be the main focus of our discussion. Nonetheless, we believe our contributions may be relevant to agricultural and mineral markets, where the literature on price integration involves the same theoretical models and empirical methods that we review in Section 2. All proofs are placed in the Appendix.

2. Literature Review

We begin this section by reviewing common network equilibrium models found in the commodity markets literature. We then discuss relevant econometric methods and their applications in these markets.

2.1. Equilibrium models

Our work builds upon spatial price equilibrium models, which were introduced in the seminal papers of Samuelson (1952) and Takayama and Judge (1964). SPE models are derived from a competitive market built over a logistical network where participants are sited at nodes and transportation routes are defined by links between nodes. The equilibrium conditions (i.e., the SPE model) are characterized by the Karush-Kuhn-Tucker (KKT) optimality conditions of the welfare maximizing allocation (optimization) problem.

SPE models have been used extensively in the modeling and analysis of commodity markets (Nagurney et al. 2014, Li et al. 2018, and the references therein). Examples in energy markets include coal (Harker and Friesz 1985), natural gas (Gabriel et al. 2000), crude oil (Bennett and Yuan 2017), and petroleum (Mudrageda and Murphy 2008). The focus of the SPE literature has mainly been on the computation of equilibrium outcomes (Nagurney et al. 2019). Specifically, spatial models of markets are developed and solved using an array of optimization tools to simulate equilibrium prices or flow quantities over different modeling inputs (e.g., transportation costs, supply and demand functions). Two relevant examples that consider transportation constraints are Lochner (2011) and Dieckhöner et al. (2013), which simulate the effects of bottleneck constraints on the spread of consumer prices over different forecasts of demand and supply quantities. SPE models have also been used to analyze infrastructure pricing and investment decisions, for example, to

motivate the pricing structure of the natural gas pipeline capacity (Cremer et al. 2003, Secomandi 2010). Simple network models are considered in this literature, such as two-node or three-node networks. In contrast to both bodies of work, we use an SPE model to provide a characterization of equilibrium prices in relation to general transportation networks (over arbitrary supply and demand functions), and then use this characterization to build empirical models for market analysis *given* equilibrium prices.

An important concept arising from spatial equilibrium theory is that of a neutral band, defining a range of price differences where arbitrage is not possible as a result of transaction costs to trade. In many commodity markets this transaction cost is primarily the result of transportation costs. To the best of our knowledge, the study of the neutral band has been limited to pairwise settings where a pair of locations can trade directly. However, when a market is connected by a more general transportation network, as is common in energy markets such as fossil fuel markets, direct links may not exist between each pair of locations. We introduce the concept of a *network* neutral band in this paper to describe price relationships between nodes over general network topologies.

Imperfect competition has also been studied in networked markets. This describes the setting in which market participants at various levels of the supply chain are assumed to have market power, i.e., the ability to affect prices by adding or withholding demand or supply. This literature has made progress in rich models of the agents, although generally at the expense of the sophistication of the operational models, i.e., by providing analytical models that allow selected participants to exert market power over markets with simple network topologies and operations. For example, Harker (1986), Qiu (1991), Abolhassani et al. (2014) and Bimpikis et al. (2019) provide analytical models of Cournot competition of suppliers over bipartite networks without transportation capacity. Massol and Banal-Estañol (2018) on the other hand considers market power of transportation capacity holders, but only over a single link connecting two perfectly competitive supply nodes. Furthermore, most of this literature (with Massol and Banal-Estañol (2018) being an exception) focuses on policy modeling and deriving analytical solutions, rather than developing empirical methods for analyzing market outcomes. Although we do not incorporate market power in the model, we do consider capacity constraints and allow for arbitrary network topologies. In the following subsection, and in the conclusion, we revisit the topic of imperfect competition from an econometric lens.

2.2. Econometric methods and relevant applications

A wide variety of econometric time series methods are applied to the study of price integration in commodity markets; a review of methods can be found in Dukhanina and Massol (2018). Many of these methods are typically applied pairwise to assess potential frictions between regions in the market footprint. One of the more commonly used methods are cointegration tests, which test

for the stationarity in deviations among prices (Hendry and Juselius 2000). De Vany and Walls (1993), Paul et al. (2001), Brown and Yücel (2008), Holmes et al. (2013), Avalos et al. (2016) are examples of cointegration tests applied to gasoline and natural gas prices across the United States, where lack of cointegration between a pair of regions is often taken as evidence of transportation bottlenecks (Marmer et al. 2007). The literature on price integration is particularly prevalent within North American and European countries, where national energy markets have undergone significant market reforms in recent decades to foster strong market competition (Oliver et al. 2018). These methods are also applied to study energy markets in a global context. For example, Adelman (1984), Weiner (1991), Siliverstovs et al. (2005) and Li et al. (2014) are notable works which examine price integration of crude oil and natural gas across various countries and continents. Globally, limited transportation capacity can also be a major contributor to non-integrated commodity prices (e.g., Siliverstovs et al. 2005). Nonetheless, there may be additional geopolitical factors that can affect price integration globally (Al Kathiri et al. 2017) and these factors should be taken into consideration when drawing conclusions about the effects of transportation capacity.

One major shortcoming of these econometric methods, as pointed out in Dukhanina and Massol (2018), is that the distributional assumptions underlying many of these general-purpose time series methods may not be consistent with spatial equilibrium conditions. This may lead to unreliable conclusions about the market (McNew and Fackler 1997). Some variants of these models have been proposed in light of these concerns. For example, threshold models have been proposed which allow for the existence of a band within which deviations from long-run equilibrium may occur without error-correction, and cointegration is measured only when deviations exceed the threshold (Balke and Fomby 1997). These models have been applied in several commodity markets (e.g., Goodwin and Piggott 2001, Park et al. 2007), limited again to pairwise comparisons. Beyond a pairwise setting, these threshold models offer no definitive or interpretable connection with the logistical network underlying the market.

The methods presented above use only price data which is generally broadly available. Models using additional market data have also been proposed. One example is the regime switching model, which estimates the frequency of being in regimes with unexploited arbitrage opportunities using transportation flow information (e.g., Barrett and Li 2002, Negassa and Myers 2007). Given certain assumptions, such models can also be designed to test for the presence of market power (Massol and Banal-Estañol 2018). While these models may be richer, building accurate models and collecting reliable data is challenging. In practice, proxies of the input variables (such as flow volumes, capacity constraints, and transportation costs) are often required due to operational data being proprietary and dispersed among many market participants. These challenges limit the models to more isolated

environments, such as the pair of regions joined by a single pipeline segment (Micola and Bunn 2007, Oliver et al. 2014, Massol and Banal-Estañol 2018).

In this paper, we study market integration using only spatial time series price data. This has the advantage that this data is readily available, which alongside the tractability of the resulting algorithms, allows for the analysis of larger geographical regions. From an econometric perspective, we relate the congestion within the logistical structure to particular spatiotemporal patterns in market prices and propose a methodology for isolating these patterns from spatial pricing data. This allows us to identify locations and time periods where price changes are consistent with the effects of capacity constraints. However, working within these data constraints limits the ability to differentiate market power from congestion. We discuss the interpretation of our results when market power may be present and potential extensions using additional market data in Section 7.

Finally, the spatial distribution of prices is also extensively studied in electricity markets where the transmission grids are capacitated and subject to frequent congestion (Oren 1997, Hogan 1999, Holmberg and Lazarczyk 2015). However, there are key differences which limit the applicability of the model in this paper to this setting. First, the pricing mechanism is different due to the institutional structure of these markets where a centralized body controls transmission and clears the market. Second, the flow of electricity within a transmission network is governed by physical laws that add additional constraints to the allocation problem and render the equilibrium conditions different from that of SPE models. Kekatos et al. (2014) and Birge et al. (2017) explore the use of pricing data to uncover the topology and constraints underlying the physical grid in deregulated electricity markets. These papers have clear analogies to our paper, but the structural differences in these markets require different approaches.

3. Market Model and Equilibrium Conditions

In this section, we present a model of a competitive market with transportation capacity constraints and review how the optimality conditions of the associated market allocation problem determine market outcomes and equilibrium prices.

We begin by developing a model of a competitive market for a single commodity with spatially separated market participants. Let the market be represented as a network with a set of nodes $\mathcal{N}$ and a set of directed links $\mathcal{E}$. Consumers are located at demand nodes $\mathcal{S}$ and producers are located at supply nodes $\mathcal{K}$, which together form a partition of $\mathcal{N}$. Each demand node may represent many independent, individual consumers in close spatial proximity (e.g., individual car owners purchasing gas within the same city); the same is true for supply nodes. Each demand node $s \in \mathcal{S}$ obtains welfare $W_s(b_s)$ when consuming b_s units of the commodity, representing the aggregate welfare of individual consumers comprising node s. Similarly, each supply node $k \in \mathcal{K}$ bears a

production cost $W_k(b_k)$ for b_k units of the commodity produced. We assume that the welfare function $W_s(\cdot)$ is strictly concave, increasing, and differentiable, while the cost function $W_k(\cdot)$ is convex, increasing and differentiable. The concavity and convexity assumptions are consistent with standard diminishing marginal utility and diminishing return assumptions from the economics literature. Note that the derivative of the welfare and cost functions, i.e., $W_s'(b_s)$ and $W_k'(b_k)$, are the inverse demand and inverse supply functions at a node s and k, respectively; an increase in demand (supply) results in an increase (decrease) in the welfare (cost) functions for a fixed value b_s (b_k). Finally, rather than explicitly modeling storage facilities, we assume they are co-located with demand and supply nodes and on a short-term basis behave similarly to other market participants in that they may influence the aggregate production cost or welfare function at their node.

Nodes are connected by a set of transportation links $\mathcal{E}$. Links will be denoted by either e or (i,j), depending on whether explicit reference to the incident nodes of the link is required. The variable f_{ij} represents the flow of the commodity from node i to j on link $(i,j) \in \mathcal{E}$. We use $I_{(i)} = \{n \in N \mid (n,i) \in \mathcal{E}\}$ to denote the set of incoming nodes to i. Similarly, $O_{(i)} = \{n \in N \mid (i,n) \in \mathcal{E}\}$ is the set of outgoing nodes from i. The flow on each link is non-negative, bounded above by the capacity of the link, u_{ij}, and has a non-negative, per-unit transportation cost of c_{ij}. We use $\mathcal{P}_{(i,j)}$ to denote the set of paths from node i to j, where each element of $\mathcal{P}_{(i,j)}$ represents a sequence of links, and p_{ij}^q to denote the cost of a path $q \in \mathcal{P}_{(i,j)}$, which is the sum of the costs on each link in q. For each pair of nodes i and j, let $\mathcal{P}^*_{(i,j)}$ describe the set of minimum-cost paths between i and j, and p_{ij}^* denote the cost of a minimum-cost path. Finally, for a specific demand node $s \in \mathcal{S}$, we let the set $\mathcal{K}_{(s)} \subseteq \mathcal{K}$ denote the set of supply nodes with a directed path to s.

Using the above notation, the equilibrium of the associated competitive market can be modeled using the following welfare-maximizing market allocation problem:

$$
\begin{aligned}
\underset{\mathbf{f},\mathbf{b}}{\text{maximize}} \quad & \sum_{s\in\mathcal{S}} W_s(b_s) - \sum_{(i,j)\in\mathcal{E}} c_{ij} f_{ij} - \sum_{k\in\mathcal{K}} W_k(b_k) \\
\text{subject to} \quad & -b_s + \sum_{i\in I(s)} f_{si} - \sum_{j\in O(s)} f_{sj} = 0, \quad \forall s \in \mathcal{S}, \\
& b_k + \sum_{i\in I(k)} f_{ik} - \sum_{j\in O(k)} f_{kj} = 0, \quad \forall k \in \mathcal{K}, \\
& 0 \le f_{ij} \le u_{ij}, \quad \forall (i,j) \in \mathcal{E}, \\
& b_s \ge 0, \quad \forall s \in \mathcal{S}, \\
& b_k \ge 0, \quad \forall k \in \mathcal{K}.
\end{aligned} \tag{1}
$$

The equilibrium market allocation in a competitive market maximizes the total social welfare, which, as presented in model (1), is the total consumer welfare minus transportation and production costs (Harker 1986). The first two sets of constraints are the standard flow-balance equations, where

consumers and producers withdraw and inject the commodity into the market, respectively. The third constraint represents capacity constraints on flow. Given that $W_s(\cdot)$ and $W_k(\cdot)$ are strictly concave and convex functions, respectively, formulation (1) is a bounded, convex optimization problem. Equilibrium prices can be deduced from the optimality conditions of (1), which are shown below:

$$\lambda_s = W_s'(b_s) + \alpha_s, \quad \forall s \in \mathcal{S}, \tag{2a}$$

$$\lambda_k = W_k'(b_k) - \alpha_k, \quad \forall k \in \mathcal{K}, \tag{2b}$$

$$\lambda_j - \lambda_i = c_{ij} - w_{ij} + \nu_{ij}, \quad \forall (i,j) \in \mathcal{E}, \tag{2c}$$

$$0 \le w_{ij} \perp f_{ij} \ge 0, \quad \forall (i,j) \in \mathcal{E}, \tag{2d}$$

$$0 \le \nu_{ij} \perp (u_{ij} - f_{ij}) \ge 0, \quad \forall (i,j) \in \mathcal{E}, \tag{2e}$$

$$0 \le \alpha_s \perp b_s \ge 0, \quad \forall s \in S, \tag{2f}$$

$$0 \le \alpha_k \perp b_k \ge 0, \quad \forall k \in \mathcal{K}. \tag{2g}$$

We use $\perp$ to define a complementarity constraint. The non-negative variables λ_s and λ_k are the dual variables corresponding to the two sets of flow balance constraints and represent the marginal cost of obtaining a unit of the commodity at the respective nodes; these variables correspond to *equilibrium prices* at the nodes. The variables α_s and α_k are dual variables of the lower bound constraints of b_s and b_k, respectively. The variables w_{ij} and ν_{ij} are the dual variables corresponding to the lower and upper bound constraints on the flow variables, respectively. Following Cremer et al. (2003), we refer to ν_{ij} as the *shadow price* of the capacity constraint on link (i,j). Equation (2c) establishes a connection between the prices at two nodes connected by a single link. Summing this equation over a path q from node n_1 to n_2 that traverses links in a set $\mathcal{E}_q$ results in

$$\begin{aligned} \lambda_{n_2} - \lambda_{n_1} &= \sum_{(i,j)\in\mathcal{E}_q} c_{ij} - \sum_{(i,j)\in\mathcal{E}_q} w_{ij} + \sum_{(i,j)\in\mathcal{E}_q} \nu_{ij} \\ &= p^q_{n_1 n_2} - \sum_{(i,j)\in\mathcal{E}_q} w_{ij} + \nu^q_{n_1 n_2}, \end{aligned} \tag{3}$$

where $\nu^q_{n_1 n_2} = \sum_{(i,j)\in\mathcal{E}_q} \nu_{ij}$ denotes the sum of shadow prices along path q. Conditions (2d)-(2g) represent the complementary conditions. For example, recall that the variables w_{ij} are non-negative and represent the shadow price of the non-negativity flow constraint. When flow on link (i,j) is positive in an optimal market allocation, the value of w_{ij} must be zero by complementary slackness. Thus, equation (3) can be represented as

$$\lambda_{n_2} - \lambda_{n_1} \le p^q_{n_1 n_2} + \nu^q_{n_1 n_2}, \quad \forall q \in \mathcal{P}(n_1, n_2), \tag{4}$$

$$\lambda_{n_2} - \lambda_{n_1} = p^q_{n_1 n_2} + \nu^q_{n_1 n_2}, \quad \forall q \in \mathcal{P}^+(n_1, n_2), \tag{5}$$

where $\mathcal{P}^+(n_1, n_2)$ is the set of paths from n_1 to n_2 for which there exists positive flow in the optimal market allocation. Equations (4) and (5) are fundamental no-arbitrage results for competitive markets. The pair of equations state that the price at node n_2 must be less than or equal to the price at node n_1 plus the marginal cost of transporting a unit from n_1 to n_2, with equality holding when there is positive flow from n_1 to n_2.

While equations (4) and (5) hold in general, there exist prices that satisfy these conditions that offer no meaningful insight into the relationship between nodal prices in the network. For example, consider a "star network" with a single supply node directly connected to each demand node at zero cost. When demand nodes are not participating (i.e., $b_s = 0$) in the market, their equilibrium prices can be arbitrarily lower than the supply node's equilibrium price. To eliminate such edge cases, we assume (without loss of generality) that all demand nodes participate in the market. Note that our focus going forward is on prices at the demand nodes because the empirical "market price" of a commodity typically refers to the price for end consumers (e.g., price of retail gasoline or price of residential natural gas). Thus, from this point on, we will use the terms demand node prices and consumer prices synonymously. We make an assumption on participation of consumers.

ASSUMPTION 1. *We assume that* $b_s > 0 \ \forall s \in \mathcal{S}$ *in the optimal market allocation.*

Another way to interpret this assumption is that for every consumer, the welfare gained from the first infinitesimally small unit consumed will always exceed the cost of producing and transporting that unit. With this assumption, we can strengthen the equilibrium conditions (4) and (5).

LEMMA 1. *For every* $s \in \mathcal{S}$, $\lambda_s = \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid k \in \mathcal{K}_{(s)}, q \in \mathcal{P}_{(k,s)}\}$.

Lemma 1 states that the equilibrium price at a participating demand node must be equal to the minimum marginal cost of production and transportation (including both explicit transportation costs and the shadow prices along the path) over the set of supply nodes to which the demand node is connected. When there is no congestion in the network, i.e., $f_{ij} < u_{ij} \ \forall (i,j) \in \mathcal{E}$, then $\nu_{ij} = 0 \ \forall (i,j) \in \mathcal{E}$ by equation (2e), and Lemma 1 can be further simplified as shown in Corollary 1.

COROLLARY 1. *For every* $s \in \mathcal{S}$, $\lambda_s = \min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}_{(s)}\}$ *when* $f_{ij} < u_{ij}$, $\forall (i,j) \in \mathcal{E}$.

4. Price Integration in Networks

In this section, we study the relationship between equilibrium prices and the underlying transportation network. To isolate the effects of network topology, link costs, and capacity constraints on the distribution of prices, we consider markets with increasingly general transportation networks. Sections 4.1-4.3 study single market realizations with arbitrary demand and supply functions. In Section 4.4, we consider the implications of these results for the analysis of multiple market realizations when the transportation network is stable.

4.1. Uncapacitated networks without transportation costs

We first define a feature of the market topology that we term *structural integration*. A set of demand nodes is structurally integrated if each node shares the same set of supply nodes. If all demand nodes in the network are structurally integrated, then we refer to the market as being structurally integrated.

DEFINITION 1. A set $\mathcal{S}_I \subseteq \mathcal{S}$ is *structurally integrated* if $\mathcal{K}_{(s)} = \mathcal{K}_{(r)}, \forall s, r \in \mathcal{S}_I$.

The main result in this subsection is that structural integration is a necessary and sufficient condition for the *law of one price* to hold in the absence of transportation frictions (i.e., $c_{ij} = 0$ and $u_{ij} = \infty$). The law of one price refers to a market having a single price for a common commodity irrespective of welfare and cost functions (Parsley and Wei 1996) and represents an extreme level of price integration. It is well known in the literature that in the absence of transportation frictions, the law of one price should theoretically hold for directly connected nodes. We extend this result to more general network topologies.

LEMMA 2. *Consider a market without transportation frictions: $c_{ij} = 0$ and $u_{ij} = \infty$ for all $(i, j) \in \mathcal{E}$. A set of demand nodes $\mathcal{S}_I$ will have common equilibrium prices ($\lambda_s = \lambda_r \ \forall s, r \in \mathcal{S}_I$), for all instantiations of welfare and cost functions if and only if the transportation network is structurally integrated ($\mathcal{K}_{(s)} = \mathcal{K}_{(r)}, \forall s, r \in \mathcal{S}_I$).*

The following example illustrates the difference between markets with and without structural integration.

Example 1 *Consider the network shown in Figure 1a. We assume that transportation costs are zero and there are no capacity constraints on the network. In this network, there exist instances where different producer cost functions can lead to different prices between the demand nodes. For example, suppose both demand nodes have the same welfare function $W_s(b) = b^{1/2}$, while the supply nodes have different linear cost functions: $W_{k_1}(b) = b$ and $W_{k_2}(b) = 2b$. The equilibrium prices under this set of welfare functions are $\lambda_{s_1} = 1, \ \lambda_{s_2} = 2$, since node s_2 can only satisfy its demand from supply node k_2, i.e., the more expensive supplier.*

When we add links that connect k_1 to s_2, either directly (Figure 1b) or indirectly through s_1 (Figure 1c) or k2 (Figure 1d), the market becomes structurally integrated and consumer prices will be equal ($\lambda_{s_1} = \lambda_{s_2} = 1$).

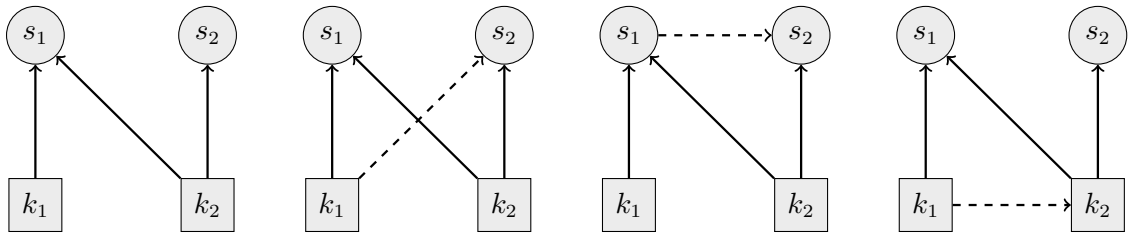

Figure 1 Examples of non-structurally integrated (a) and structurally integrated markets (b), (c) and (d). Dashed lines indicate links which are not present in panel (a).

Structural integration implies that the set of demand nodes are connected to the same set of supply nodes. Thus, in the absence of frictions impeding the movement of goods, the marginal price for all demand nodes will be the same. If the demand nodes are not structurally integrated, it is possible for a supply node that is connected to only a subset of demand nodes to have lower costs, leading to lower prices for this subset of demand nodes.

Structural integration is important for differentiating price differences caused by transportation costs and capacity constraints from price differences due to the topology of the network. In the following sections where we examine richer transportation networks that include transportation costs and capacity constraints, we assume that the market is structurally integrated in order to isolate price effects that result from these network features.

Remark 1. Any connected network where all links are bidirectional is structurally integrated.

Remark 2. A network that is not structurally integrated can be reformulated into one that is structurally integrated if certain assumptions on the welfare and cost functions are met. Specifically, if we assume that the peripheral supply nodes (i.e., the supply that is accessible only by a strict subset of demand nodes) are small, with supply that is insufficient to satisfy any demand node under a subset of "realistic" instantiations of welfare and cost functions, then we can simply account for this supply implicitly in the welfare functions of demand nodes, and the modified network without the peripheral supply nodes is structurally integrated.

4.2. Uncapacitated networks with transportation costs

Next, we consider markets where transportation costs are non-zero but links remain uncapacitated (i.e., $c_{ij} \geq 0$ and $u_{ij} = \infty$). This setting is representative of the majority of commodity market models in the literature. We show that in this setting, structural integration is necessary and sufficient to guarantee a well-defined neutral band, which we refer to as a *network neutral band*. Extending the pairwise neutral band to a network setting enables insight into price integration when demand nodes are not directly adjacent.

THEOREM 1. *Consider a market with an uncapacitated transportation network. A pair of demand nodes $s, r \in \mathcal{S}$ are structurally integrated if and only if*

$$\min\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\} \leq \lambda_s - \lambda_r \leq \max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\} \tag{6}$$

for all instantiations of welfare and cost functions.

When two demand nodes are not structurally integrated, there exist welfare and cost functions that can lead to equilibrium price differences that exceed any given bound. However, when two nodes are structurally integrated, the price difference will always be bounded and the bound is characterized entirely by the network structure and link costs. When the entire market is structurally integrated, the prices at any two demand nodes are still related because they have access to the same set of supply nodes, even though the cost to access the supply may vary. This is reflected in the key part that the differences between shortest path distances to suppliers play in equation (6). Lemma 2, in the previous section, shows a special case of Theorem 1: since all transportation costs are zero, the shortest paths $p^*_{ks} = p^*_{kr}$ are also zero for all pairs of demand nodes so that equation (6) implies common prices. Next, we show that the bound in (6) is tight.

PROPOSITION 1. *Given any value Δ within the neutral band for a pair of structurally integrated demand nodes $s, r \in \mathcal{S}_I$, there exist welfare and cost functions for the market participants that will result in equilibrium prices λ_s, λ_r such that $\lambda_s - \lambda_r = \Delta$.*

Proposition 1 implies that the bound from the network neutral band, described in equation (6), will be at least as tight a bound on $\lambda_s - \lambda_r$ as the bound from the pairwise neutral band. Section EC.1 in the Appendix provides a simple example where the network neutral band is strictly tighter than the pairwise one. For convenience in our analysis and exposition, we will refer to the network neutral band as simply the neutral band. Furthermore, we define the mid-point and half-width of the neutral band between nodes r and s, ρ_{rs} and α_{rs}, as follows:

$$\rho_{rs} = \frac{1}{2}\left(\min\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}\} + \max\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}\}\right),$$
$$\alpha_{rs} = \frac{1}{2}\left(\max\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}\} - \min\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}\}\right).$$

The network neutral band can be used to illustrate the role of the "position" of supply nodes in the network on the degree of price integration, which we explore in the following example.

Example 2 *This example explores the impact of supply proximity to demand nodes, measured by transportation costs, on the neutral band. Figure 2 shows three cases of two demand nodes supplied by two supply nodes in a structurally integrated market.*

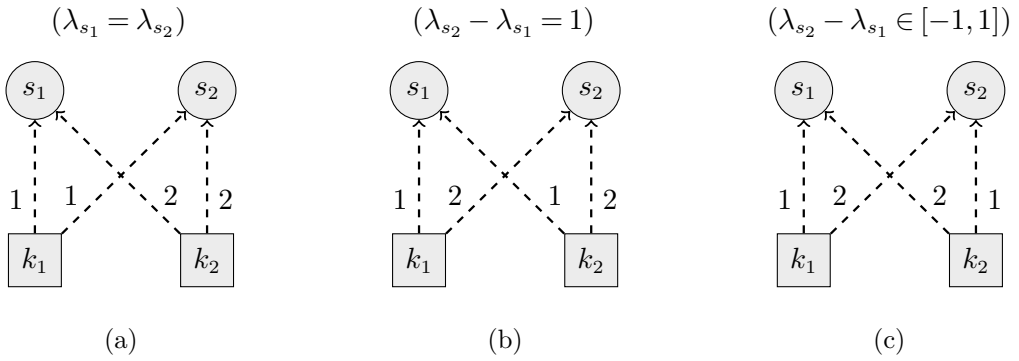

Figure 2 Three instances of an uncapacitated network where the transportation costs on each link are as labeled.

In Figure 2a, both demand nodes face the same transportation costs from each supply node, i.e., $c_{k1s1} = c_{k1s2} = 1$ *and* $c_{k2s1} = c_{k2s2} = 2$*. In this case, the half-width and midpoint of the neutral band is zero and the equilibrium prices for both demand nodes are always equal. In Figure 2b,* s_1 *faces lower transportation costs than* s_2 *over both supply nodes, i.e.,* $c_{k1s1} = c_{k2s1} = 1$ *and* $c_{k1s2} = c_{k2s2} = 2$*. The midpoint is shifted and the equilibrium price at node* s_1 *will always be 1 unit lower than* s_2*. Finally, in Figure 2c, consumers at each demand node can access a subset of the supply nodes with cheaper transportation costs. The neutral band midpoint is at zero but the half-width is 1. This implies that the absolute price difference between* s_1 *and* s_2 *can be up to 1 unit, and the exact difference will vary depending on supply costs.*

This example provides insight into how the distribution of supply and demand over a market footprint can impact the neutral band, and, in turn, market integration. When demand is clustered together, shortest path costs from different supply nodes will be similar for all demand nodes and might result in a situation as in Figure 2a. This results in a small neutral band centred around zero, which leads to a common market price for all demand nodes. When supply is clustered together, transportation costs for a demand node will be similar irrespective of the supplier. This is the case in Figure 2b, which leads to a narrow neutral band, though the midpoint of the neutral band may be far from zero. This results in stable differences in consumer prices. The gasoline market studied in this paper features refining capacity clustered in the Gulf of Mexico region of the U.S. and is an example of this type of market. Finally, the implications for price integration are different in a market where supply nodes are more dispersed with respect to consumers. In this case, certain suppliers will have lower transportation costs for certain consumers as is illustrated in Figure 2c. The resulting heterogeneous consumer preferences for suppliers leads to a wider neutral band within which demand and supply shocks may propagate throughout the market footprint leading to less integrated consumer prices. We refer the reader to Section EC.3 in the Electronic Companion for a more thorough discussion of this example using equilibrium flows.

REMARK 3. We note that although the neutral bands are derived using a model of a competitive market, the bands are, in fact, independent of whether there is market power at the sources of supply. Specifically, they hold even if producers are monopolies or are engaged in oligopolistic competition, which may occur when there are few supply nodes with few independent producers within each node. Equilibrium price differences exceeding the neutral band would by definition indicate the existence of unexploited arbitrage opportunities (since the neutral bands are only a function of transportation costs). This can only occur if transportation is capacitated; otherwise, any market participant would be able to exploit these arbitrage opportunities which would cause price differences to re-equilibrate to be within the neutral band.

4.3. Capacitated networks with transportation costs

We now allow links in the transportation network to be both costly and subject to capacity constraints (i.e., $c_{ij} \geq 0$ and $u_{ij} \leq \infty$). In this setting, positive shadow prices on capacity constraints can lead to a *congestion surcharge* borne by a subset of demand nodes. Without capacity constraints, as described in Corollary 1, each demand node price will be equal to the minimum of the sum of the price at a supply node and the cost of transportation between the supply node and demand node. We define the congestion surcharge as the part of the demand node price above this value:

DEFINITION 2. The *congestion surcharge* w_s for a demand node $s \in \mathcal{S}$ is the amount that the equilibrium price at s exceeds the uncapacitated delivery price to node s:

$$w_s = \max\{\lambda_s - \lambda_k - p^*_{ks} \mid k \in \mathcal{K}\}. \tag{7}$$

We can rearrange equation (7) to obtain

$$\lambda_s = \min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}\} + w_s. \tag{8}$$

Corollary 1 shows that in the absence of capacity constraints, the congestion surcharge is zero. We will study the dynamics of these charges in driving apart equilibrium prices and creating local pricing discrepancies that would not otherwise exist.

Combining the result from Lemma 1 and equation (8), we can write w_s as

$$w_s = \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid k \in \mathcal{K}, q \in \mathcal{P}_{(k,s)}\} - \min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}\}. \tag{9}$$

Equation (9) shows that w_s can be described as the difference between the cost of acquiring a unit when considering shadow prices in the network and and the cost when shadow prices are not considered. Using equation (9), we extend the neutral band described in Theorem 1 to the setting with capacity constraints:

THEOREM 2. *Let $r, s \in \mathcal{S}$. The price difference between r and s is bounded by*

$$\min\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\} + w_s - w_r \leq \lambda_s - \lambda_r \leq \max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\} + w_s - w_r \tag{10}$$

over all welfare and cost functions.

Equation (10) shows that a pair of demand nodes sharing the same congestion surcharge will have the same neutral band as in the setting with no capacity constraints. When the congestion surcharge differs between a pair of demand nodes, the midpoint of the neutral band will be shifted. Notably, the width of the neutral band is not affected by the congestion surcharge. When there are no capacity constraints, $w_s = 0$ for all s (Corollary 1), equation (10) is equivalent to equation (6).

For the subsequent analysis of data, it is useful to assume the existence of a *root node* which is a demand node with a congestion surcharge of zero. Using equation (10) we can derive a simple bound on each consumer price relative to the price of the root node $o \in \mathcal{S}$:

$$\min\{p^*_{ks} - p^*_{ko} \mid k \in \mathcal{K}_{(s)}\} + w_s + \lambda_o \leq \lambda_s \leq \max\{p^*_{ks} - p^*_{ko} \mid k \in \mathcal{K}_{(s)}\} + w_s + \lambda_o. \tag{11}$$

The windows for consumer prices described in the bounds in Equation (11) will be shifted both by the congestion surcharge from a congested link and by the price of the root node. A corollary of equation (9) shows that such a node s will exist if there are no congested links on the path minimizing $\min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}\}$. A sufficient condition for a node s to be a root node is thus that s is not downstream of any congested links. Root nodes are further discussed following Example 3.

4.3.1. Congestion on a single link. To best elucidate the relationship between market structure and the propagation of congestion surcharge throughout a network, we study the case where there is exactly one capacitated link in the network. We first consider price integration between the pair of nodes at either ends of this capacitated link.

PROPOSITION 2. *Consider a market where demand nodes $i, j \in \mathcal{S}$ are joined by the link (i, j). If link (i, j) is the only congested link in the network, then $w_i = 0$ and $w_j = \nu_{ij}$.*

Proposition 2 is intuitive and states that when the flow on link (i, j) in a network reaches its capacity, node j incurs a congestion surcharge equal to the full shadow price of the link. The more interesting case is the impact of the capacitated link (i, j) on prices at nodes $s \in \mathcal{S} \backslash \{i, j\}$ that are not directly adjacent. Below, we illustrate that it is possible for different nodes in the network to incur different amounts of congestion surcharge in the presence of a single congested link.

We first require some additional formalization. Let $e \in \mathcal{E}$ denote the single congested link. Recall that p^*_{ks} is the cost of the minimum-cost path from k to s. Let $p^{*,\neg e}_{ks}$ be the cost of the minimum-cost "replacement" path from k to s which does not include link e and let $\delta_e(k, s) = p^{*,\neg e}_{ks} - p^*_{ks}$.

The value $\delta_e(k,s)$ can be viewed as the maximum cost of continuing commerce between k and s in the absence of link e. If all paths from k to s include link e, then $\delta_e(k,s) := \infty$. Finally, let $\delta_e^{\min}(s) = \min\{\delta_e(k,s) \mid k \in \mathcal{K}\}$ and $\delta_e^{\max}(s) = \max\{\delta_e(k,s) \mid k \in \mathcal{K}\}$.

THEOREM 3. *Suppose there is a single congested link $e \in \mathcal{E}$ in the network with shadow price ν_e. Then, for all $s \in \mathcal{S}$,*

$$w_s \in [\min\{\nu_e, \delta_e^{\min}(s)\},\ \min\{\nu_e, \delta_e^{\max}(s)\}]. \tag{12}$$

Theorem 3 describes, for each node in the network, the effects of replacement paths for a congested link e. For some node $s \in \mathcal{S}$, if the additional cost of rerouting around e from any supplier is high ($\delta_e^{\min}(s) > \nu_e$), then the congestion surcharge at node s reflects the full shadow price ν_e of link e. This describes a situation where link e is an integral component in the transportation of supply to node s. On the other hand, if the opposite were true, i.e., that for node s the additional cost of rerouting is always low ($\delta_e^{\max}(s) < \nu_e$), then the incurred congestion surcharge would be lower than ν_e. The implications of the theorem are consistent with intuition on how network connectivity can mitigate costs of congestion. In a densely connected network, the cost of rerouting around a link (and by proxy $\delta_e^{\max}(s)$) is likely to be low, limiting the set of nodes whose price will reflect the full shadow price of a congested link. On the other hand, in a sparse network the cost of rerouting (and by proxy $\delta_e^{\min}(s)$) may be large, implying that the shadow price of a congested link can be fully reflected in many downstream nodes.

We use the following example to provide a comprehensive illustration of the relationship between network structure and pricing for three cases characterized by Theorem 3: a) the absence of any paths that avoid a congested link e ($\delta_e^{\min}(s) = \infty$), b) when all alternative paths have the same cost ($\delta_e^{\min}(s) = \delta_e^{\max}(s)$), and c) when alternative paths have different cost ($\delta_e^{\min}(s) < \delta_e^{\max}(s)$).

Example 3 *We examine outcomes for three markets illustrated in Figures 3a, 3b, and 3c. Each market features three demand nodes, s_1, s_2 and s_3, and two supply nodes, k_1 and k_2. Each supply node has the cost function $W_k(b_k) = b_k^2$ which possesses increasing marginal costs. The demand node welfare functions are $W_s(b_{s_1}) = 10\sqrt{b_{s_1}}$, $W_s(b_{s_2}) = 20\sqrt{b_{s_2}}$, and $W_s(b_{s_3}) = 20\sqrt{b_{s_3}}$, which possess diminishing marginal utility.*

The markets differ only in the transportation network. Market 3a is connected by the illustrated network where all links have zero transportation costs and only link (s_1, s_2) (highlighted in red) has a capacity of 1 unit. Market 3b differs from market 3a by having the additional link (s_1, s_3) with a transportation cost of 1 unit. Market 3c differs from market 3a by having the additional link (k_1, s_3), also with a transportation cost of 1 unit. Figure 3 shows the equilibrium prices beside each node. Positive flows in the market allocation are shown by solid lines.

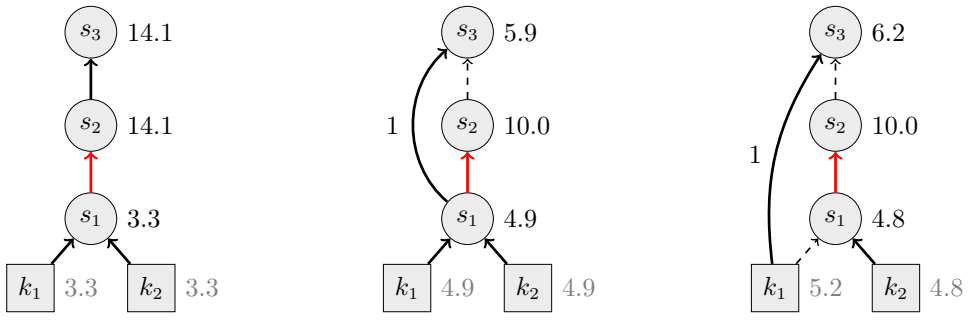

(a) $\delta_e^{\min}(s_3) = \infty$ (b) $\delta_e^{\min}(s_3) = \delta_e^{\max}(s_3) = 1$ (c) $\delta_e^{\min}(s_3) < \delta_e^{\max}(s_3)$

Figure 3 Equilibrium prices with a congested link (in red) in three different markets.

The shadow price for the link (s_1, s_2), denoted by ν_{s_1,s_2}, is equal to 10.8, 5.1, and 5.2 units respectively in markets 3a, 3b and 3c.

In Example 3, if the capacity constraint is removed, the equilibrium prices would be identical across all three markets and equal to 6.1 units, since all three markets are connected by the same subnetwork of zero cost paths. When link (s_1, s_2) has a capacity constraint which is reached, each market has a different set of equilibrium prices. Note that in each market, s_1 is a root node, since the minimum-cost paths from each supply node to s_1 does not include link (s_1, s_2). Since the neutral band is zero for all demand nodes, the price difference $\lambda_{s_2} - \lambda_{s_1}$ and $\lambda_{s_3} - \lambda_{s_1}$ directly reflect the congestion surcharge of the nodes s_2 and s_3, respectively. In all three markets, the equilibrium price at s_2 is equal to the price at s_1 plus the shadow price of the link (s_1, s_2), which can be derived by observing that $\delta_e^{\min}(s_2) = \infty$ in equation (12). Practically, all flow to s_2 must come through s_1.

However, the options available for serving s_3 differs in the three markets. In market 3a, s_3 incurs the full shadow price of 10.8 units at equilibrium since, like node s_2, there do not exist any alternative paths for the commodity to reach s_3 ($\delta_e^{\min}(s_2) = \infty$). In markets 3b and 3c, there are alternative paths to s_3. In market 3b, the equilibrium price at node s_3 is 1 unit higher than at s_1, which can be explained by $\delta_e^{\min}(s_3) = \delta_e^{\max}(s_3) = 1$; any shadow price that exceeds one unit would result in flow being rerouted onto link (s_1, s_3), implying that the price difference between s_3 and s_1 would never exceed 1 unit. In market 3c, s_3 obtains all of the commodity from k_1 directly, with an equilibrium price that is 1.4 units higher than s_1. Since $\delta_e^{\min}(s_3) = 1, \delta_e^{\max}(s_3) = \infty$, equation (7) suggests that the congestion surcharge on node s_3 can be any value between 1 unit and the shadow price of 5.2 units, depending on the supply and demand functions.

Note that in the market 3c, the direct connection from k_1 to s_3 surprisingly results in s_3 incurring a higher price than it did in the market 3b. This outcome results from the fact that in market 3b, node s_3 could access both k_1 and k_2 cheaply, whereas s_3 can only access k_1 cheaply in market 3c. The more concentrated demand on k_1 in market 3c results in a higher production price at k_1

(due to the marginally increasing production cost), leading to a higher equilibrium price at s_3. Market 3c also highlights that examining only the direction of flows in a network may result in the misleading conclusion that s_3 is in a disjoint market from s_1, s_2. On the other hand, the equilibrium prices clearly highlight that both s_2 and s_3 do incur a positive congestion surcharge as a result of the congestion link, albeit different in magnitude.

Finally, note that if link (k_1, s_1) is the capacitated link, then the congestion surcharge cannot be fully observed in consumer prices because the shadow price of the congested link is applied to all consumers (i.e., $w_s > 0 \ \forall s \in \mathcal{S}$, and we do not observe the portion that is cancelled out by the $w_s - w_r$ term in equation (10)). Any market equilibrium will have a root node except in the case where a congested link is upstream of all demand nodes; in such a setting, price differences exceeding the neutral band reflect an underestimate of the total surcharge.

4.4. Observations over multiple market realizations

Up to this point, we focused on the distribution of prices in a single market realization. We now extend our previous results to the case where we have multiple observations over a market. In particular, at each distinct "period", indexed by $t \in \mathcal{T} = \{1, \ldots, T\}$, we observe prices from an independent realization over a market with fixed network structure and link costs, although potentially different welfare functions, cost functions, and capacities. These dynamics are typical of energy markets where the transportation network is capital intensive and can be assumed to be static over the medium term, whereas demand can shift quickly with consumer preferences (e.g., as a result of poor weather) while the network is prone to potential disruptions that can reduce link capacities. Under these conditions, we derive a price decomposition which describes the spatiotemporal dynamics of a set of prices over a time horizon $\mathcal{T}$.

PROPOSITION 3. *The set of equilibrium prices for nodes $s \in \mathcal{S}$ over a market with a fixed network structure and link costs can be expressed as*

$$\lambda_s^t = \eta^t + \rho_s + \epsilon_s^t + w_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{13}$$

where $\epsilon_s^t \in [-\alpha_s, \alpha_s]$ and $w_s^t \geq 0$.

Equation (13) decomposes the distribution of prices over a set of market realizations into a few different components which vary over time, nodes or both (as indexed). More specifically, the set of prices can be decomposed into a node-invariant "market trend" η^t, a set of time invariant terms ρ_s and α_s representing the network neutral band bounding the idiosyncratic movement of ϵ_s^t, and a term w_s^t representing the congestion surcharge. Note that since the ϵ_s^t is bounded by the time

invariant terms, the term w_s^t is the only term that can be unconstrained both spatially (i.e., per node) and temporally (i.e., per $t \in \mathcal{T}$).

The price decomposition provides intuition into the interplay between prices, market forces and the transportation network. When there are no binding capacity constraints in the network over the set of periods $\mathcal{T}$, i.e., $w_s^t = 0,\ \forall s \in \mathcal{S}, t \in \mathcal{T}$, changes in the participant's welfare and cost functions between market realizations may affect the value of η^t or the value of ϵ_s^t within the bound $[-\alpha_s, \alpha_s]$. Price shocks generated from mild local demand and supply shifts are likely to be contained within the neutral band without affecting the overall market, whereas sufficiently large local demand or supply shocks will shift prices throughout the market by changing the value of η^t. In the setting where there are binding capacity constraints, the additional terms w_s^t reflect the congestion surcharge experienced by different nodes in the market. Depending on the network configuration, it is possible that large price shocks generated from significant local demand changes can remain locally contained, i.e., w_s^t will be positive for a small subset of nodes without changing η^t. We refer the reader to Section EC.3 in the Electronic Companion for a more thorough discussion of these ideas.

5. Estimating the Congestion Surcharge from Price Data

In this section, we leverage the previous results and present a framework for analyzing price data and estimating congestion surcharges in a commodity market with limited information transparency, i.e., where precise information of flow directions, costs, volumes and capacities is generally unavailable. Using only spatial price data observed over a given time horizon, our methodology aims to identify spatiotemporal price variations that are indicative and reflective of binding capacity constraints in the underlying network.

5.1. Surcharge Estimation Model

The surcharge estimation model (SEM) is based on the price decomposition shown in equation (13). It takes as input a set of spatial prices $\boldsymbol{\lambda} = \{\lambda_s^t\}_{s\in\mathcal{S}, t\in\mathcal{T}}$ and a set of user-selected parameters and outputs an estimate of the congestion surcharge at each node over the given time horizon. In its most generic form, the SEM model can be presented as follows:

$$\underset{\eta^t, \rho_s, \epsilon_s^t, w_s^t, \alpha_s}{\text{minimize}} \quad \sum_{s\in\mathcal{S}} \alpha_s \tag{14a}$$

$$\text{subject to} \quad \lambda_s^t = \eta^t + \rho_s + \epsilon_s^t + w_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{14b}$$

$$|\epsilon_s^t| \le \alpha_s, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{14c}$$

$$w_s^t \in \mathcal{W},\ \forall s \in \mathcal{S}, t \in \mathcal{T}. \tag{14d}$$

Constraints (14b) and (14c) are derived directly from the price decomposition presented in equation (13). The variables η^t capture a node-invariant underlying trend, while the variables ϵ_s^t capture price variations within time-invariant neutral bands of width α_s centered at ρ_s. All remaining price variation is captured by the variables w_s^t, representing congestion surcharges. Constraints on w_s^t are represented by $\mathcal{W} \subseteq \mathbb{R}^+$.

If price movements are perfectly synchronized across all nodes, i.e., price differences are constant, the optimal objective value will be zero and the prices λ_s^t can be entirely explained by the node-invariant term η^t and time-invariant term ρ_s. The variables ϵ_s^t and w_s^t capture deviations from price integration, which are attributed to variation within the neutral band and transient congestion in the transportation links, respectively. While model (14) is derived directly from the price decomposition, the model without any additional constraints in the form of $\mathcal{W}$ is underdetermined, as shown by the following remark.

REMARK 4. If $\mathcal{W} = \mathbb{R}^+$, the optimal objective value of (14) will always be zero.

When w_s^t is unconstrained, there is a free variable w_s^t for every price λ_s^t, and an optimal solution would simply be to set $w_s^t = \lambda_s^t$ (assuming all λ_s^t are positive) and all other variables to zero. This solution reflects the hypothesis that there is congestion at every time period across all nodes. The other extreme solution is setting $\mathcal{W} = \{0\}$, which represents the hypothesis that nodes do not incur congestion surcharges over the observed period and all price variations can be explained fully through changes in supply and demand. A judicious choice of $\mathcal{W}$ can be used to more finely differentiate between the two extreme explanations for non-integrated prices, and the strategies to do so are discussed in detail in the following subsection.

Finally, we note that it is generally reasonable to assume the existence of periods both with and without congestion events over typical lengths of time studied in the commodity markets literature. For example, in oil and gas markets, transportation congestion events are expected to occur periodically but are on average not present for large proportions of time when considered over time intervals of a few months or years (Oliver et al. 2014, Massol and Banal-Estañol 2018).

5.2. Approach to congestion surcharge identification

Our identification strategy is based on: 1) limiting the proportion of periods that the congestion surcharges may be active; 2) resolving precise values for the congestion surcharges using a conservative strategy; 3) determining the proportion of congested periods in a principled manner. This subsection addresses these points in turn.

5.2.1. Identifying periods of congestion. We use a set of *time-limiting constraints* to force $w_s^t = 0$ for a fraction of total time periods, while allowing the precise set of periods to be selected by the model. Let $\beta \in [0,1]$ define an input parameter representing a fraction of the time horizon for which w_s^t is unconstrained, with $w_s^t = 0$ for all other time periods. The time-limiting constraints can be written as

$$w_s^t \le \psi^t M,\ \sum_{t\in\mathcal{T}} \psi^t \le \lfloor \beta T \rfloor,\ \psi^t \in \{0,1\}, \quad \forall s \in \mathcal{S},\ t \in \mathcal{T}. \tag{15}$$

The binary variables ψ^t determine periods for which the congestion surcharge is free ($\psi^t = 1$) or fixed to zero ($\psi^t = 0$), and M represents a sufficiently large value such that w_s^t will never reach its upper bound when $\psi^t = 1$. When we add equations (15) to model (14), an optimal solution of the model will identify $\lfloor \beta T \rfloor$ periods in $\mathcal{T}$, which, when removed, will minimize the sum of the neutral band widths over each node in the remaining $\lceil (1-\beta)T \rceil$ periods. Put differently, the model seeks to identify the $\lfloor \beta T \rfloor$ periods (by setting $\psi^t = 1$) where price variations across the set of nodes most greatly exceed any set of neutral bands that can be fitted over all other periods (where $\psi^t = 0$). Finally, we note that the periods where $\psi^t = 0$ are used to fit the variables ρ_s and α_s, which are then used to estimate w_s^t over periods where $\psi^t = 1$.

5.2.2. Identifying congestion surcharge values. Next, we introduce a set of *conservative-estimation* constraints that we use to remove one degree of freedom from the variable estimates. First, we note that if (ϵ_s^t, w_s^t) represents a pair of solutions to the SEM where t is a period for which w_s^t is free, it is possible to modify the solution to $(\epsilon_s^t - \delta, w_s^t + \delta)$ without changing the objective value. The range of possible values of w_s^t for which the solution remains optimal is potentially large ($\delta \in [-\alpha_s + \epsilon_s^t, \alpha_s + \epsilon_s^t]$). To handle this ambiguity, we enforce conservative estimates of the surcharges w_s^t, and capture only surcharges resulting in price movements that exceed the neutral band. So, w_s^t will only capture parts of the price that strictly exceed α_s. This is enforced using the following constraints:

$$w_s^t \le \pi_s^t M,\ \epsilon_s^t + (1-\pi_s^t)M \ge \alpha_s,\ \pi_s^t \in \{0,1\}, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}. \tag{16}$$

Identification issues also exist between w_s^t and η^t during congested periods. In particular, without impacting the optimality of a solution, we can make the values of w_s^t arbitrarily larger by shifting value from η^t. That is, $(\epsilon_s^t + w_s^t + \delta, \eta^t - \delta)$ is also a solution for any $\delta \ge 0$, since the negative and positive δ values will cancel each other out. We rectify this problem by adding a constraint to select the maximal value of η^t from the set of optimal solutions, leading to the minimal estimate of w_s^t. This is enforced with the following set of constraints,

$$\sum_s \gamma_s^t \ge \psi^t,\ \epsilon_s^t \le -\alpha_s + (1-\gamma_s^t)M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{17}$$

forcing $\epsilon_s^t = -\alpha_s^t$ for at least one s in each period $t \in \mathcal{T}$ for which $\psi^t = 1$. These constraints ensure that η^t (w_s^t) is the largest (smallest) possible value out of the set of optimal solutions.

5.2.3. Complete mixed-integer linear optimization formulation. We now formulate the SEM as a mixed-integer linear optimization model, where all price data is represented by λ_s^t:

$$
\begin{aligned}
z(\beta) := \text{ minimize } & \sum_{s\in\mathcal{S}} \alpha_s && \text{(18a)}\\
\text{subject to } & \lambda_s^t = \eta^t + \rho_s + \epsilon_s^t + w_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18b)}\\
& \epsilon_s^t \geq -\alpha_s, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18c)}\\
& \epsilon_s^t \leq \alpha_s, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18d)}\\
& w_s^t \leq \psi^t M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18e)}\\
& \sum_{t\in\mathcal{T}} \psi^t \leq \lfloor \beta T \rfloor, && \text{(18f)}\\
& w_s^t \leq \pi_s^t M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18g)}\\
& \epsilon_s^t + (1-\pi_s^t)M \geq \alpha_s, \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18h)}\\
& \sum_{s} \gamma_s^t \geq \psi^t, \quad \forall t \in \mathcal{T}, && \text{(18i)}\\
& \epsilon_s^t \leq -\alpha_s + (1-\gamma_s^t)M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18j)}\\
& \gamma_s^t, \psi^t, \pi_s^t \in \{0,1\}, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(18k)}\\
& w_s^t \geq 0 \quad \forall s \in \mathcal{S}, t \in \mathcal{T}. && \text{(18l)}
\end{aligned}
$$

Constraints (18e) and (18f) define the time-limiting constraints, and constraints (18g)-(18j) define the conservative-estimation constraints. It suffices to set $M = \max\{\lambda_r^t - \lambda_s^t \mid r, s \in \mathcal{S},\ t \in \mathcal{T}\}$, which is the largest absolute price difference between any two nodes across the entire time horizon.

A simpler reformulation. Finally, we show that model (18) can be simplified by introducing a new variable $\bar{w}_s^t := w_s^t + \epsilon_s^t - \alpha_s$. With this new variable, constraints (18b) and (18c) can be rewritten as

$$
\begin{aligned}
& \lambda_s^t = \eta^t + \rho_s + \alpha_s + \bar{w}_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, && \text{(19a)}\\
& \bar{w}_s^t \geq -2\alpha_s, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}. && \text{(19b)}
\end{aligned}
$$

In this new representation, the neutral bands in the absence of congestion are defined as bands of $\bar{w}_s^t \in [-2\alpha_s, 0]$ around mid-points $-\alpha_s$. Since $w_s^t \geq 0$ and $\epsilon_s^t - \alpha_s \leq 0$, then $\bar{w}_s^t > 0$ if and only if $w_s^t > 0$, which implies that constraints (18e) and (18f) can be equivalently defined on $\bar{w}_s^t$.

Constraints (18d), (18g) and (18h) in model (18) ensure that the variable w_s^t can only be positive if $(w_s^t + \epsilon_s^t)$ exceeds the neutral band. These constraints are no longer necessary in the new representation; when $\bar{w}_s^t > 0$, it by definition represents the setting in which the neutral band has been

exceeded. In other words, when both ϵ_s^t and w_s^t appear in constraint (18b), we could add constants with opposite signs to each variable while retaining the objective value, and the constraints were added to avoid this setting. When we combine these two variables in $\bar{w}_s^t$, this issue is resolved. As a result, constraints (18g) and (18h) can be removed entirely and the set of binary variables π_s^t that appear only in these constraints can also be removed.

$$\underset{\eta^t,\rho_s,\alpha_s,\bar{w}_s^t,\psi^t,\gamma_s^t}{\text{minimize}} \quad \sum_{s\in\mathcal{S}} \alpha_s \tag{20a}$$

$$\text{subject to} \quad \lambda_s^t = \eta^t + \rho_s + \alpha_s + \bar{w}_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{20b}$$

$$\bar{w}_s^t \geq -2\alpha_s, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{20c}$$

$$\bar{w}_s^t \leq \psi^t M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{20d}$$

$$\sum_{t\in\mathcal{T}} \psi^t \leq \lfloor \beta T \rfloor, \tag{20e}$$

$$\bar{w}_s^t \leq -2\alpha_s + (1-\gamma_s^t)M, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{20f}$$

$$\sum_{s} \gamma_s^t \geq \psi^t, \quad \forall t \in \mathcal{T}, \tag{20g}$$

$$\gamma_s^t, \psi^t \in \{0,1\}, \quad \forall t \in \mathcal{T}. \tag{20h}$$

Given optimal values of $\bar{w}_s^t$ from model (20), we can calculate values of the original w_s^t and ϵ_s^t variables as follows. If $\bar{w}_s^t \geq 0$, then $w_s^t = \bar{w}_s^t$ and $\epsilon_s^t = \alpha_s$. If $\bar{w}_s^t < 0$, then $w_s^t = 0$ and $\epsilon_s^t = \bar{w}_s^t + \alpha_s$.

Computation over large datasets. A potential complication may arise when solving model (20) over large datasets. Since model (20) does not explicitly account for the sequential nature of time periods and instead treats each $t \in \mathcal{T}$ independently, model (20) includes the problem of choosing the best βT independent time periods out of $\mathcal{T}$. The solution space of this problem can be particularly large when the number of time periods T is large and β is close to 0.5. We propose a set of (optional) constraints to reduce the complexity of solving the SEM over large datasets. For a fixed $t \in \mathcal{T}$ and positive integer m, let $T_{ub}(t,m) = \min\{T, t+m\}$ and $T_{lb}(t,m) = \max\{0, t-m\}$. We can add the following set of constraints to the model:

$$\sum_{t^*=t}^{T_{ub}(t,m)} \psi^{t^*} \geq \nu^t \cdot (T - T_{ub}(t,m) + 1), \quad \forall t \in \mathcal{T}, \tag{21a}$$

$$\psi^t \leq \sum_{t^*=T_{lb}(t,m)}^{t} \nu^{t^*}, \quad \forall t \in \mathcal{T}, \tag{21b}$$

$$\nu^t \in \{0,1\}, \quad \forall t \in \mathcal{T}. \tag{21c}$$

These constraints explicitly link adjacent time periods by enforcing the following condition: a period t can be selected by the SEM, i.e., $\psi^t = 1$, if and only if a *block* of adjacent time periods, including t, of minimum size m is selected. Large datasets by definition are highly granular (e.g.,

daily prices over an extended period of time). Thus, the addition of constraints (21a)-(21c) when solving the SEM with large datasets reflect the observation that empirically, network congestion events are unlikely to appear and resolve instantaneously (e.g., within a single day), but are instead more likely to persist across several adjacent time periods (e.g., across several days).

5.3. Exploring market characteristics with β

The β parameter determines the proportion of time periods where the surcharge terms w_s^t can take positive values. In practice, this parameter determines the tendency of the algorithm to classify pricing deviations as those caused by capacity constraints ($w_s^t > 0$) and those caused by idiosyncratic changes in supply and demand ($w_s^t = 0$). Increasing β corresponds to increasing the algorithm's sensitivity to the effects of capacity constraints. If β is set too low, the algorithm may not be able to fully identify the price movements associated with capacity constraints, and if it is set too high it may misclassify idiosyncratic demand shocks. This challenge is analogous to problems in unsupervised learning such as determining the correct number of clusters in k-means clustering. Our approach is to use complementary methods to ensure that the results, in particular the collection of estimated pricing deviations, are consistent with expected characteristics of capacity bottlenecks.

Manual examination. We examine the congestion surcharge estimates over different β values. Bottlenecks in the underlying network are likely to manifest as price increases over contiguous locations and adjacent time periods. On the other hand, we have low confidence over low magnitude, surcharge estimates that are heterogeneously dispersed geographically and temporally. By increasing the values of β, we can examine the surcharge estimates to identify thresholds for which further increases in β result in more frequent appearance of surcharge estimates exhibiting qualities for which we have low confidence. We elaborate further using empirical data in Section 6.

Changes in model metrics. To improve our understanding of where capacity constraints are active, we consider how the objective value $z(\beta)$ and the total surcharge $\sum_s \sum_t w_s^t$ vary as β increases. The objective value corresponds to an estimation of the total idiosyncratic pricing deviations and the total surcharge corresponds to the fit of the prices with respect to the neutral band. Examining these metrics for "elbows", points where the rate of change of the metric exhibits a discrete sustained change, allows us to pinpoint levels of β where any new detected events have distinctly different characteristics. This method is analogous to the practice in k-means clustering of examining changes in within-cluster variation to identify the correct measure of clusters as k is increased.

Examining the change in objective value $z(\beta)$ as β increases shows the degree to which newly identified capacity events substitute for unexplained pricing variation. As changes in $z(\beta)$ slow down with respect to β, the value of increasing β at explaining pricing variation in the dataset

also slows. Examining the change in the total surcharge as β increases provides an indication of the magnitude and duration of the pricing effects of newly identified capacity events. Prices which exceed the neutral band significantly over a sustained period of time, such that $\sum_s \sum_t w_s^t$ is large, are more likely to be a signal of irregularities in the underlying market. On the other hand, if prices only exceed the neutral band sightly and are short-lived, such that $\sum_s \sum_t w_s^t$ is small, then the confidence we have on w_s^t being congestion surcharge rather than noise is low. Discrete changes in the rate of total surcharge increase indicates transitions between such phenomena.

6. Case study: The Southeastern U.S. Gasoline Market

In this section, we present a case study on the southeastern U.S. gasoline market and demonstrate the effectiveness of our surcharge estimation method in capturing the effect of major network disruptions on gasoline prices. The market, discussed in detail below, relies on highly specialized pipeline infrastructure for commodity transportation, much like crude oil and natural gas markets, as well as an expansive network of tanker vessels, rail cars and trucks. Extensive government deregulation on the supply side, competition and open-access regulations on the transportation side, and price transparency on the demand side have made the gasoline market one of the most competitive commodity markets in the U.S. (Paul et al. 2001, Holmes et al. 2013). The U.S. gasoline market thus provides a natural setting for which to study the effects of transportation bottlenecks on spatial price integration.

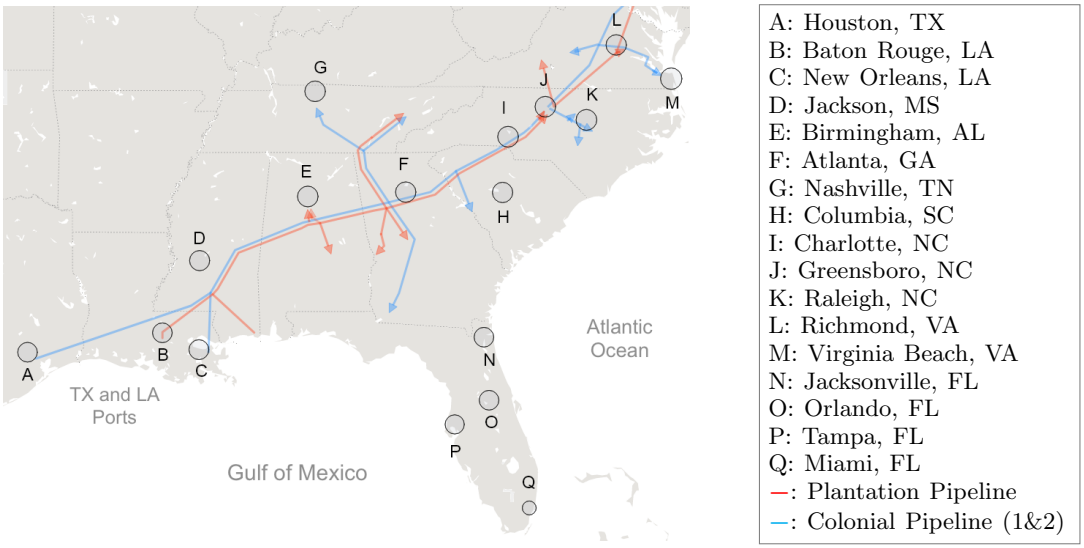

Figure 4 The cities considered in this study, along with the major interstate pipeline networks. The rail, truck and tanker networks are not shown.

6.1. Market and data description

The region that we consider, shown in Figure 4, spans from Texas to Virginia and has substantial refining, transportation and consumption activities. Dozens of major refineries, representing nearly 50% of all the United States' refining capacity, are located along the coasts of Texas and Louisiana (EIA 2012). The majority of gasoline transportation from these refineries to southeastern and eastern states is carried out using two pipelines, namely the Colonial Pipeline (in blue) and the Plantation Pipeline (in red). On the other hand, most transportation to Florida relies on tankers which deliver from ports in Texas and Lousiana to ports in Florida. Finally, there exists an extensive network of rail cars and trucks that further supports the movement of gasoline, particularly to locations that are not in close proximity to pipeline exit points or ports. The fact that the southeastern market is serviced almost exclusively by the refineries, coupled with the close proximity of refineries to each other, suggests that strong price integration should be expected in the absence of transportation capacity constraints.

We obtained daily gasoline prices of the seventeen cities marked in Figure 4 from January 1, 2016 to December 31, 2017. The cities chosen include some of the most populated cities in the region. Daily prices are calculated from the last daily price of regular gasoline at all gasoline stations within the United States Postal Service (USPS) designated boundaries of the city. This data was collected through a combination of fleet card transactions, crowd-sourcing and direct retail pricing, and acquired from a data aggregator. A set of summary statistics for the data is found in Table EC.1 in the Electronic Companion. Finally, since federal and state motor fuel taxes are typically updated on the first day of each year (EIA 2019), we split our data into two sets, one per year. Having two sets of data also offers two instances with which to test the methodology.

6.2. Setup and preliminary results

We begin by analyzing the estimated market characteristics over different β values. To ensure that the SEM can be solved to optimality within an appropriate time window, in this case a few hours, we consider the SEM with the addition of constraints (21a)-(21c) with $m = 7$, corresponding to identifying blocks of adjacent time periods that are a week or more in duration. For each year, the SEM is solved with $\beta = \{0, 0.01, 0.02, ...0.40\}$.

Figure 5 shows a snapshot of the estimated congestion surcharge values over $\beta = \{0.10, 0.20, 0.30, 0.40\}$. Specifically, each subfigure shows the solutions of w_s^t in model (14) for the corresponding β value. While the locations are not labeled in this figure, the figure highlights the the magnitude, trajectory and persistence of estimated surcharge values over increasing β values. It is evident from the figure that there are two major "clusters" of surcharge estimates, the first

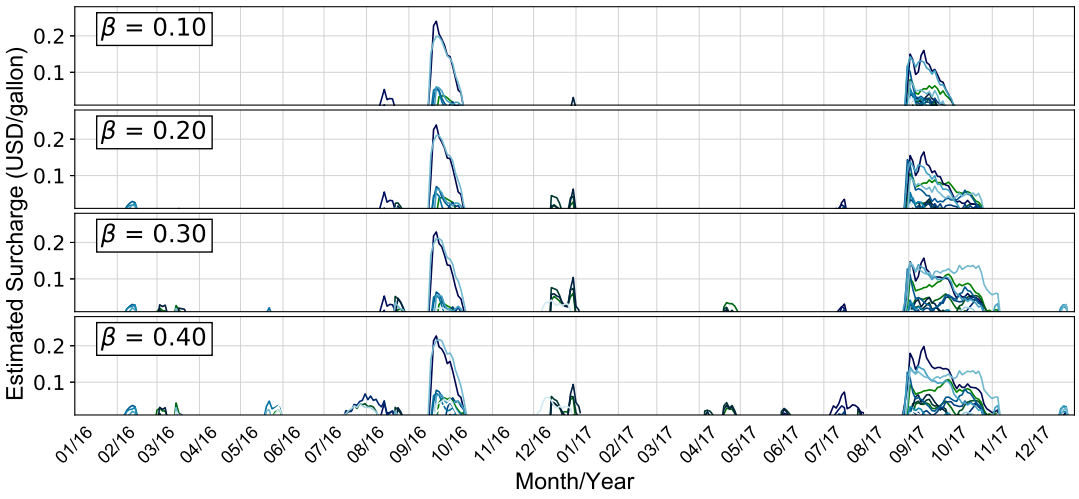

Figure 5 Estimated congestion surcharge in 2016 and 2017 over different β values.

spanning several weeks during the months of September and October of 2016, and the second spanning the months of September to November of 2017. These two clusters are the first to be identified over low β values and persist as β is increases. The surcharge estimates within these two clusters are large in magnitude and are contiguous both temporally and geographically (the latter will be shown in the next section).

We use the metrics proposed in Section 5.3 to further guide the analysis. The objective value and the total surcharge over different β values are shown in Figures 6a and 6b, whereas the change in these values are shown in Figures 6c and 6d. For the year 2016, we observe a distinct change in both the rate of objective value decrease and the rate of total surcharge increase (although to a lesser degree) at $\beta_{(2016)} = 0.08$, as seen in Figures 6c and 6d. Specifically, this value marks a transition point where increasing β no longer generates a large decrease in objective value nor does it capture the same amount of surcharge across the locations. Empirically, this β value marks a point in which the surcharge estimates of the first major cluster become stable while new clusters with lower magnitude and less contiguous in time and space begin to be captured. For the year 2017, we observe a distinct change in the rate of total surcharge increase, this time at $\beta_{2017} = 0.1$ and $\beta_{2017} = 0.22$ as shown in Figure 6d. Empirically, $\beta_{2017} = 0.22$ again marks a point at which the second major cluster becomes stable with increases in β. Beyond this value, further increases in β results in the appearance of new surcharge estimates with lower magnitude and shorter duration.

The two major surcharge clusters that we detect, namely those in September - October 2016 and September - November 2017, coincide with periods of severe weather events and network disruptions. On the other hand, surcharge estimates found over other time periods are relatively lower in magnitude and exhibit features that are consistent with congestion events during normal market

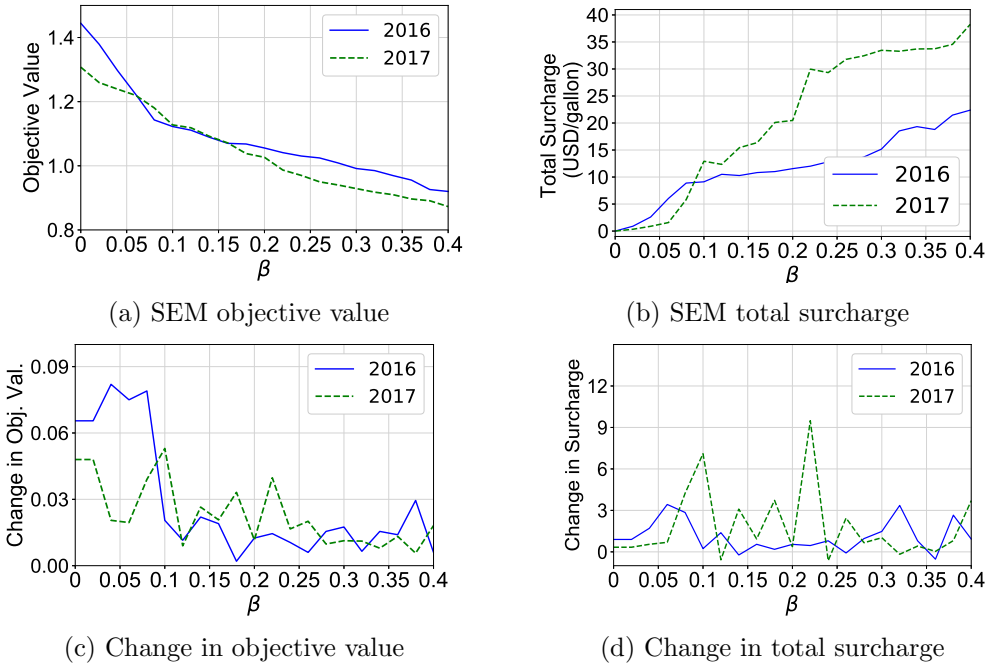

(a) SEM objective value

(b) SEM total surcharge

(c) Change in objective value

(d) Change in total surcharge

Figure 6 The calibration metrics over different β values.

operations. Since we cannot validate these smaller events from secondary sources, for the remainder of this section, we focus on studying gasoline price dispersion over the two time periods with well-documented market disruptions, using the estimated congestion surcharge values (obtained with $\beta_{2016} = 0.08$ and $\beta_{2017} = 0.22$) along with publicly available data of the pipeline system.

Finally, we note that we consider the application of two additional econometric models to our data set in Section EC.4 of the Electronic Companion. Specifically, we apply a pairwise cointegration test and a panel regression model, and compare their results to the surcharge estimates derived by our SEM model. We show that, unlike the SEM model, neither of the two models are capable of isolating the spatiotemporal price effects generated by the capacity events that are detailed in the following two case studies.

6.3. Study I: Price shocks from 2016 pipeline disruption

On September 9, 2016, a major pipeline leak was discovered on Line 1 of the Colonial Pipeline in Shelby County (see Figure 7), and a partial shutdown of that segment of the pipeline immediately followed (ICF 2016). The shutdown lasted until September 21, when the pipeline resumed full operating capacity. On October 31, 2016, a deadly pipeline explosion in Shelby County caused a partial shutdown of the pipeline, and operations were restarted on November 8. These dates are shaded in the inset of Figure 7.

The inset of Figure 7 shows the trajectory of the congestion surcharge for cities that experienced significant surcharges immediately following the pipeline leak. Consistent with our theory, these

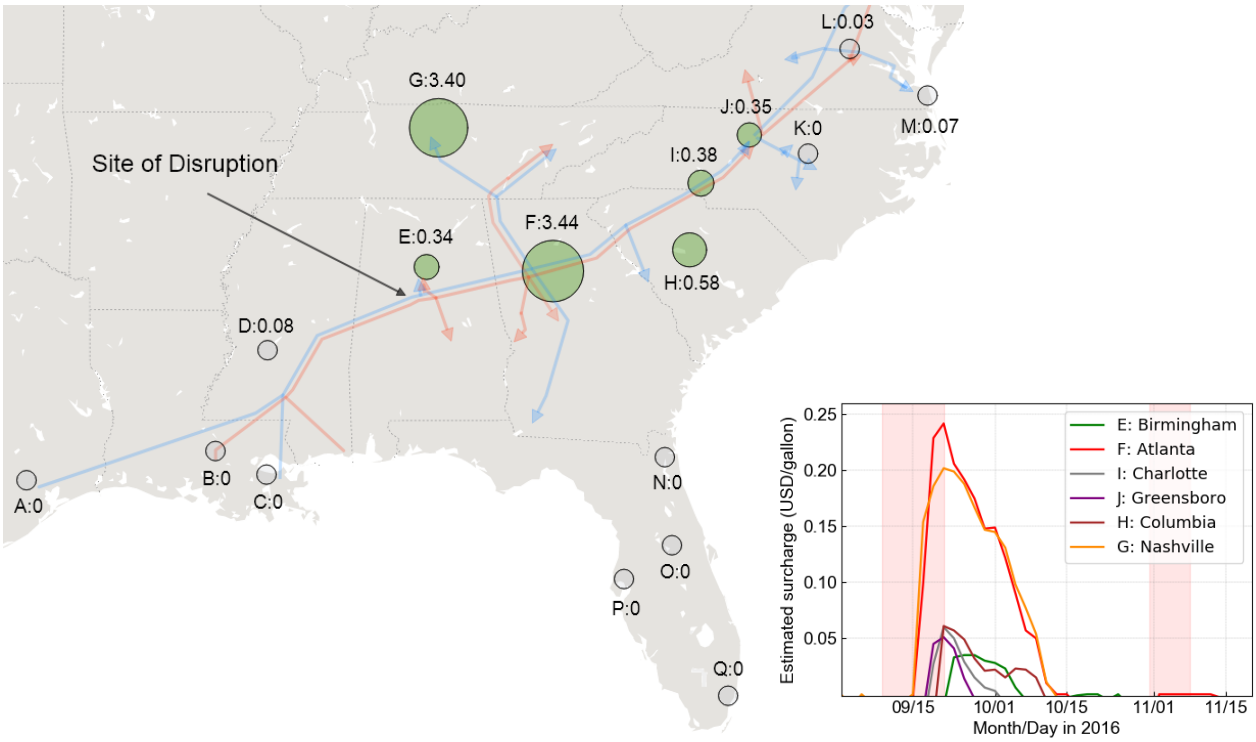

Figure 7 Map: cumulative estimated surcharges (with $\beta = 0.08$) resulting from 2016 pipeline disruption. Inset: estimated per-period surcharge from September to November 2016.

cities are all downstream from the leak. The map in Figure 7 shows the *cumulative* congestion surcharge (area under the curves in the inset) for each city in the region. First, we note that there is a lag between when the first disruption occurs and when the surcharge is observed. This observation captures the effects of the use of stored inventory, well-documented around this period (EIA 2016b), to mitigate potential price shocks. Second, we note that the congestion surcharge over all locations begins to decrease immediately after the pipeline is restored. However, it takes three full weeks for the price differences to completely disappear, highlighting the effects of high demand for pipeline capacity (to replenish inventories) following the disruption. Interestingly, we do not identify any congestion surcharge associated with the second disruption, which is likely due to a restocking of inventories to much higher levels following the previous disruption (EIA 2016a).

We observe that the locations immediately downstream of the site of disruption, in particular, Atlanta (label F) and Nashville (label G), observed the most significant price increases. Furthermore, it is easy to observe, especially in the case of Nashville, that there are no alternative sources of pipeline transportation other than the Colonial Pipeline. Interestingly, we find that all estimated surcharge values are roughly bounded from above by the surcharge incurred at Atlanta, which is the first downstream node of the disrupted pipeline that could not be easily supplied by other means of short-haul transportation such as trucks. This is consistent with our theoretical analysis,

namely Theorem 3 and Example 3, which show that in a setting with a single disrupted link: 1) the location most directly downstream of the link incurs the most significant surcharge, equal to the shadow price of the link and 2) other locations will incur a surcharge bounded by this shadow price, with the magnitude dependent on the availability of alternative transportation resources.

The effect of alternative transportation resources is further highlighted by noting that the estimated surcharge values in the cities immediately downstream of Greensboro (label J), which serves as a major junction point where the pipelines can unload their supply, are essentially negligible. During the pipeline disruption, which occurred in Line 1 of the Colonial Pipeline, a second pipeline running in parallel (Line 2) that is typically used to transport heating oil, diesel and jet fuel, was temporarily used to transport gasoline to the eastern cities (EIA 2016b). The low levels of estimated surcharge in the cities downstream of Greensboro provide strong evidence that this rerouting of gasoline (on Line 2) and the existence of other transportation resources (the Plantation Pipeline) is crucial in mitigating price increases during network disruptions.

Finally, we note that congestion surcharge was not identified in any of the cities in Florida. This is consistent with what we would expect, since gasoline is delivered to these cities through tankers rather than by the Colonial and Plantation pipeline network.

6.4. Study II: Price shocks from 2017 hurricane season

The southeastern U.S. witnessed a catastrophic hurricane season in the fall of 2017. Most notable were Hurricanes Harvey and Irma, which together caused nearly $200 billion dollars worth of damage in this region alone. These two hurricanes also created large disruptions and logistical challenges in the petroleum supply chain. Hurricane Harvey, which made landfall in Texas and Louisiana during the last week of August 2017, resulted in the closure of refineries, docking of tankers, and the Colonial Pipeline being shut down on August 30 for one week, resuming operations at limited capacity on September 6 (EIA 2017a). Immediately following Harvey, Hurricane Irma, traveling from the Caribbean up to Florida, forced ports in Florida to close in the first two weeks of September. The reduced operations at ports around Texas, Louisiana and Florida resulted in significantly reduced product delivery to Florida during the hurricane season, lasting approximately from August 25 to September 13 (EIA 2017b).

First, we find that the estimated surcharge values of the cities labeled D, E, G, F, H, I (and highlighted in green) coincide with and spike immediately following the closure of the Colonial Pipeline. These cities lie directly on the Colonial Pipeline between the refineries and the Greensboro (label J) junction point. Note that in the 2016 pipeline disruption, only the cities F, G, H, I were identified as having a significant surcharge, whereas in the 2017 hurricane season, a significant surcharge is also identified in the cities of Jackson (label D) and Birmingham (label E). The result shows the

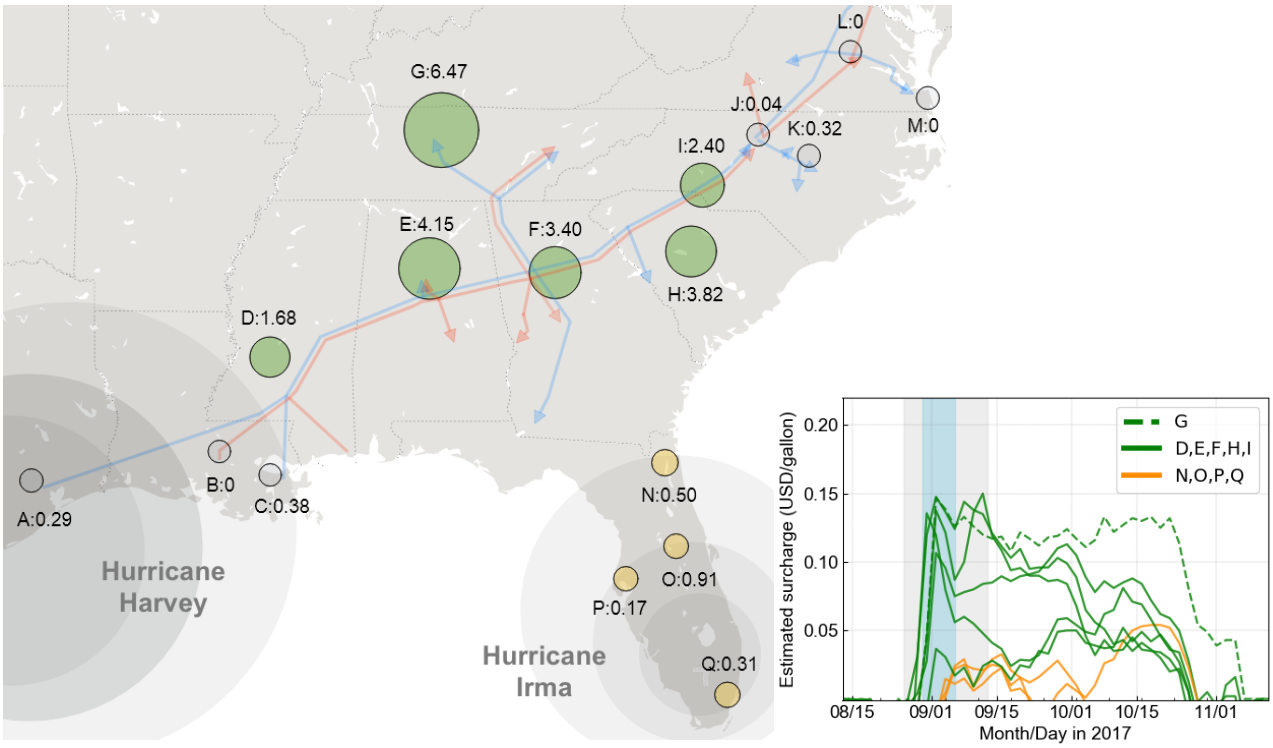

Figure 8 Map: cumulative estimated surcharges (with $\beta = 0.22$) resulting from 2017 hurricane season. Inset: estimated per-period surcharge from August to November 2017. The time periods for the Colonial Pipeline closure (Aug. 30th to Sept. 6th) is highlighted in blue and that of port closures (Aug. 25th to Sept. 13th) is highlighted in gray.

distinction between the effects of a disruption at a precise point in the pipeline, which we examined in the previous subsection, versus the effects of the closure of the entire pipeline, which is what occurred in 2017. Like the 2016 pipeline disruption, however, we find that all locations downstream of the Greensboro junction experience relatively negligible amounts of surcharge, highlighting the importance of the Plantation Pipeline which remained fully operational during this period.

Second, we find that surcharges in Florida again reflect what is expected from the pipeline disruptions. In particular, the surcharge values do not rise at the same time and as sharply as the other cities, highlighting that the inland pipeline disruptions do not have a major impact on prices in Florida. However, unlike 2016, we do find positive congestion surcharges, reflecting the reduced port operations and limited marine movements of tankers due to Hurricane Harvey and Hurricane Irma. Nonetheless, the estimated surcharge values among the cities in Florida are much lower in magnitude than those of the more inland cities. This is likely a result of the flexibility of tankers and marine transport, where capacity can ramp up quickly and deliveries from other coastal ports can be accommodated (EIA 2017b).

Third, unlike the surcharge estimates in the previous section, the surcharges in 2017 appear more erratic, which may reflect the many simultaneously occurring disruptions in the network.

A sustained surcharge estimate is observed in Nashville (label G) over this period, which is the only city served by pipelines but which does not lie in proximity to the Plantation Pipeline. The surcharge values in all other cities exhibit a generally decreasing trend after the onset of the disruptions, lasting for roughly two full months before dissipating.

Interestingly, the magnitude of surcharge estimates found during the 2017 disruptions appear lower in magnitude than those from the 2016 disruption, despite the far more disastrous consequences of the hurricane season. Examining the average price of gasoline across all cities offers an explanation. In 2016, the average price of gasoline during the month of the disruption (September 2016) was 2.09 USD/gallon whereas that of the month before (August 2016) was 2.01 USD/gallon. On the other hand, in 2017 the average price of gasoline was 2.55 USD/gallon during the first month of the disruption (September 2017), which is significantly higher than the average price of 2.20 USD/gallon in August 2017. This observation highlights the difference between the effects of market forces (i.e., supply and demand) and the effect of the transportation network. In particular, the major refinery (supply) disruptions resulted in the average price of gasoline increasing over *all* the cities while the disruption of the transportation network resulted in the estimated price disparities *between* the cities.

7. Conclusion

The impact of transportation constraints in commodity markets has become increasingly relevant as many markets experience demand and supply growth that continues to outpace the growth in transportation infrastructure. Large spatial price discrepancies resulting from transportation bottlenecks are increasingly documented not only in markets with inland transportation but also in markets serviced by shipping networks (e.g., ICIS 2020). Furthermore, major disruptions to these networks, whether from natural weather events or more recent examples of cyberattacks (EIA 2021), can have consequential effects on commodity prices. For these reasons, new theoretical models and empirical methods which explicitly account for transportation capacity constraints are needed for the analysis of these markets.

In this paper, we provide a systematic way of analyzing the distribution of spatial prices in competitive commodity markets with capacity constrained transportation networks. We establish several important results which relate prices and price differences to network structure, costs, and transportation constraints. We examine how congestion surcharges arising from network bottlenecks can propagate through the network and develop a discrete optimization methodology to estimate capacity effects using only price data. Through a case study of the U.S. gasoline market, we show that the methodology is capable of extracting spatiotemporal price effects resulting from capacity constraints in the underlying transportation network. We discuss how these empirical

results reflect the role of network topology, flexible transportation and inventory availability on consumer prices and market integration.

Our empirical methodology provides a principled way of analyzing prices across large geographical regions to estimate the price effects generated by transportation bottlenecks. In practice, these estimates may be of interest to both policy-makers and market participants. For example, the surcharge estimates highlight the potential value (or lack of value) that new transportation infrastructure or trading routes can bring to the market. They also provide an estimate of the value of inventory reserves, which can act as a temporary substitute for transportation services (Oliver et al. 2014). Finally, the estimates can also inform financial or commodity trading practices, such as time arbitrage and locational arbitrage using secondary transportation technologies (Martínez-de Albéniz and Simón 2017).

While our competitive model serves as a good proxy for highly competitive markets, the model may not be suitable in markets where participants have significant market power. Under normal circumstances, the energy markets under consideration in this paper are generally quite competitive, as shown in the literature. However, when capacity constraints arise, transportation bottlenecks may, for example, create opportunities for certain suppliers or those holding limited transportation capacity to raise prices beyond what would otherwise be possible. In such a setting, the estimated congestion surcharges, which reflect price variations exceeding the neutral bands (see Remark 3), may capture not only the shadow price of capacity but also rents associated with transient market power. These rents may potentially obscure the real value of policies, investment decisions, and trading practices.

Disentangling strategic pricing from real congestion surcharges using only price data is difficult, and instead requires more precise modeling and market data. In a market with imperfect competition, the degree of market power that can be exerted (and hence, the rents that are extracted) by different participants depends on various market conditions such as the number and size of local suppliers or the capacity of transportation links, which may vary greatly across locations. Future work should aim to extend models of oligopolistic competition to general networked and capacitated markets, which can then be leveraged to develop new empirical methods capable of differentiating between the price effects from capacity constraints and those from market power. However, these more complex empirical methods will likely require substantial investment in data collection. To this end, the empirical methodology in this paper, which relies only on price data, can also be used to survey and highlight time periods and regions that may warrant further investigation.

References

Abolhassani, Melika, Mohammad Hossein Bateni, Mohammad Taghi Hajiaghayi, Hamid Mahini, Anshul Sawant. 2014. Network cournot competition. *International Conference on Web and Internet Economics*. Springer, 15–29.

Adelman, Morris Albert. 1984. International oil agreements. *The Energy Journal* **5**(3) 1–9.

Al Kathiri, Nader, Yazeed Al-Rashed, Tilak K Doshi, Frederic H Murphy. 2017. "Asian premium" or "North Atlantic discount": does geographical diversification in oil trade always impose costs? *Energy Economics* **66** 411–420.

Avalos, Roger, Timothy Fitzgerald, Randal R Rucker. 2016. Measuring the effects of natural gas pipeline constraints on regional pricing and market integration. *Energy Economics* **60** 217–231.

Balke, Nathan S, Thomas B Fomby. 1997. Threshold cointegration. *International economic review* 627–645.

Barrett, Christopher B, Jau Rong Li. 2002. Distinguishing between equilibrium and integration in spatial price analysis. *American Journal of Agricultural Economics* **84**(2) 292–307.

Bennett, Max, Yue Yuan. 2017. On the price spread of benchmark crude oils: A spatial price equilibrium model. Available at SSRN: https://ssrn.com/abstract=2894389.

Bimpikis, Kostas, Shayan Ehsani, Rahmi Ilkılıç. 2019. Cournot competition in networked markets. *Management Science* **65**(6) 2467–2481.

Birge, John R, Ali Hortaçsu, J Michael Pavlin. 2017. Inverse optimization for the recovery of market structure from market outcomes: An application to the miso electricity market. *Operations Research* **65**(4) 837–855.

Brown, Stephen PA, Mine K Yücel. 2008. What drives natural gas prices? *The Energy Journal* **29**(2) 45–60.

Cremer, Helmuth, Farid Gasmi, Jean-Jacques Laffont. 2003. Access to pipelines in competitive gas markets. *Journal of Regulatory Economics* **24**(1) 5–33.

De Vany, Arthur, W David Walls. 1993. Pipeline access and market integration in the natural gas industry: Evidence from cointegration tests. *The Energy Journal* **14** 1–20.

Dieckhöner, Caroline, Stefan Lochner, Dietmar Lindenberger. 2013. European natural gas infrastructure: the impact of market developments on gas flows and physical market integration. *Applied energy* **102** 994–1003.

Doane, Michael J, Daniel F Spulber. 1994. Open access and the evolution of the us spot market for natural gas. *The Journal of Law and Economics* **37**(2) 477–517.

Dukhanina, Ekaterina, Olivier Massol. 2018. Spatial integration of natural gas markets: a literature review. *Current Sustainable/Renewable Energy Reports* **5**(2) 129–137.

EIA, U.S. Energy Information Administration. 2012. Much of the country's refinery capacity is concentrated along the gulf coast. https://www.eia.gov/todayinenergy/detail.php?id=7170.

EIA, U.S. Energy Information Administration. 2016a. Major gasoline pipeline in southeast disrupted for second time in two months. https://www.eia.gov/todayinenergy/detail.php?id=28632.

EIA, U.S. Energy Information Administration. 2016b. Pipeline disruption leads to record gasoline stock changes in southeast, gulf coast. https://www.eia.gov/todayinenergy/detail.php?id=28172.

EIA, U.S. Energy Information Administration. 2017a. Hurricane harvey caused u.s. gulf coast refinery runs to drop, gasoline prices to rise. https://www.eia.gov/todayinenergy/detail.php?id=32852.

EIA, U.S. Energy Information Administration. 2017b. Hurricanes harvey and irma lead to higher gasoline prices in florida. https://www.eia.gov/todayinenergy/detail.php?id=32932.

EIA, U.S. Energy Information Administration. 2019. Federal and state motor fuels taxes. https://www.eia.gov/petroleum/marketing/monthly/xls/fueltaxes.xls.

EIA, U.S. Energy Information Administration. 2021. Cyberattack halts fuel movement on colonial petroleum pipeline. https://www.eia.gov/todayinenergy/detail.php?id=47917.

Engle, Robert F, Clive WJ Granger. 1987. Co-integration and error correction: representation, estimation, and testing. *Econometrica: journal of the Econometric Society* 251–276.

Gabriel, Steven A, Shree Vikas, David M Ribar. 2000. Measuring the influence of canadian carbon stabilization programs on natural gas exports to the united states via a 'bottom-up' intertemporal spatial price equilibrium model. *Energy Economics* **22**(5) 497–525.

Goodwin, Barry K, Nicholas E Piggott. 2001. Spatial market integration in the presence of threshold effects. *American Journal of Agricultural Economics* **83**(2) 302–317.

Harker, Patrick T. 1986. Alternative models of spatial competition. *Operations Research* **34**(3) 410–425.

Harker, Patrick T, Terry L Friesz. 1985. The use of equilibrium network models in logistics management: with application to the us coal industry. *Transportation Research Part B: Methodological* **19**(5) 457–470.

Hendry, David F, Katarina Juselius. 2000. Explaining cointegration analysis: Part 1. *The Energy Journal* **21**(1) 1–42.

Hogan, William W. 1999. Transmission congestion: the nodal-zonal debate revisited. *Harvard University, John F. Kennedy School of Government, Center for Business and Government. Retrieved August* **29**(4).

Holmberg, Pär, Ewa Lazarczyk. 2015. Comparison of congestion management techniques: Nodal, zonal and discriminatory pricing. *The Energy Journal* **36**(2) 145–166.

Holmes, Mark J, Jesús Otero, Theodore Panagiotidis. 2013. On the dynamics of gasoline market integration in the united states: Evidence from a pair-wise approach. *Energy Economics* **36** 503–510.

ICF. 2016. East coast and gulf coast transportation fuels markets. Tech. rep., U.S. Energy Information Administration.

ICIS, Independent Commodity Intelligence Services. 2020. Lng spot prices surge on outages, shipping congestion. https://www.icis.com/energy-connections/2020/12/lng-spot-prices-surge-on-outages-shipping-congestion/.

Kekatos, Vassilis, Georgios B Giannakis, Ross Baldick. 2014. Grid topology identification using electricity prices. *PES General Meeting— Conference & Exposition, 2014 IEEE*. IEEE, 1–5.

Li, Dong, Anna Nagurney, Min Yu. 2018. Consumer learning of product quality with time delay: Insights from spatial price equilibrium models with differentiated products. *Omega* **81** 150–168.

Li, Raymond, Roselyne Joyeux, Ronald D Ripple. 2014. International natural gas market integration. *The Energy Journal* **35**(4) 159–179.

Lochner, Stefan. 2011. Identification of congestion and valuation of transport infrastructures in the european natural gas market. *Energy* **36**(5) 2483–2492.

Marmer, Vadim, Dmitry Shapiro, Paul MacAvoy. 2007. Bottlenecks in regional markets for natural gas transmission services. *Energy Economics* **29**(1) 37–45.

Martínez-de Albéniz, Victor, Josep Maria Vendrell Simón. 2017. A capacitated commodity trading model with market power. *Real Options in Energy and Commodity Markets*. World Scientific, 31–60.

Massol, Olivier, Albert Banal-Estañol. 2018. Market power and spatial arbitrage between interconnected gas hubs. *The Energy Journal* **39**(SI2) 67–95.

McNew, Kevin, Paul L Fackler. 1997. Testing market equilibrium: is cointegration informative? *Journal of Agricultural and Resource Economics* **22**(2) 191–207.

Micola, Augusto Rupérez, Derek W Bunn. 2007. Two markets and a weak link. *Energy Economics* **29**(1) 79–93.

Mudrageda, Murthy, Frederic H Murphy. 2008. Or practice—an economic equilibrium model of the market for marine transportation services in petroleum products. *Operations Research* **56**(2) 278–285.

Murphy, Frederic H, Murthy V Mudrageda. 1998. A decomposition approach for a class of economic equilibrium models. *Operations Research* **46**(3) 368–377.

Nagurney, Anna, Deniz Besik, June Dong. 2019. Tariffs and quotas in world trade: a unified variational inequality framework. *European Journal of Operational Research* **275**(1) 347–360.

Nagurney, Anna, Dong Li, Ladimer S Nagurney. 2014. Spatial price equilibrium with information asymmetry in quality and minimum quality standards. *International Journal of Production Economics* **158** 300–313.

Negassa, Asfaw, Robert J Myers. 2007. Estimating policy effects on spatial market efficiency: An extension to the parity bounds model. *American Journal of Agricultural Economics* **89**(2) 338–352.

Oliver, Matthew E, Charles F Mason, David Finnoff. 2014. Pipeline congestion and basis differentials. *Journal of Regulatory Economics* **46**(3) 261–291.

Oliver, Matthew E, Charles F Mason, et al. 2018. Natural gas pipeline regulation in the united states: Past, present, and future. *Foundations and Trends in Microeconomics* **11**(4) 227–288.

Oren, Shmuel S. 1997. Economic inefficiency of passive transmission rights in congested electricity systems with competitive generation. *The Energy Journal* **18**(1) 63–83.

Park, Haesun, James W Mjelde, David A Bessler. 2007. Time-varying threshold cointegration and the law of one price. *Applied Economics* **39**(9) 1091–1105.

Parsley, David C, Shang-Jin Wei. 1996. Convergence to the law of one price without trade barriers or currency fluctuations. *The Quarterly Journal of Economics* **111**(4) 1211–1236.

Paul, Rodney J, Dragan Miljkovic, Viju Ipe. 2001. Market integration in us gasoline markets. *Applied Economics* **33**(10) 1335–1340.

Qiu, Yuping. 1991. Solution properties of oligopolistic network equilibria. *Networks* **21**(5) 565–580.

Samuelson, Paul A. 1952. Spatial price equilibrium and linear programming. *The American economic review* **42**(3) 283–303.

Secomandi, Nicola. 2010. On the pricing of natural gas pipeline capacity. *Manufacturing & Service Operations Management* **12**(3) 393–408.

Siliverstovs, Boriss, Guillaume L'Hégaret, Anne Neumann, Christian Von Hirschhausen. 2005. International market integration for natural gas? a cointegration analysis of prices in europe, north america and japan. *Energy Economics* **27**(4) 603–615.

Takayama, Takashi, George G Judge. 1964. Equilibrium among spatially separated markets: A reformulation. *Econometrica* **32**(4) 510–524.

Weiner, Robert J. 1991. Is the world oil market "one great pool"? *The Energy Journal* **12**(3) 95–108.

Appendix. Proofs of Statements

Proof of Lemma 1 We first prove that $\lambda_s \leq \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid \forall k \in \mathcal{K}(s),\ q \in \mathcal{P}_{(k,s)}\},\ \forall s \in \mathcal{S}$. For any node s, there exists a path from every $k \in \mathcal{K}(s)$ to s, by the definition of $\mathcal{K}(s)$. Thus equations (4) and (5) must hold for each of these supply-demand pairs $(k,s),\ \forall k \in \mathcal{K}(s)$, i.e., $\lambda_s \leq \lambda_k + p^q_{ks} + \nu^q_{ks},\ \forall k \in \mathcal{K}(s)\ ,\ \forall q \in \mathcal{P}_{(k,s)}$. This completes the first part of the proof. We now show that when $b_s > 0$ in the optimal allocation, $\lambda_s \geq \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid \forall k \in \mathcal{K}(s),\ q \in \mathcal{P}_{(k,s)}\}\ \forall s \in \mathcal{S}$. We invoke equilibrium condition (3). Since there is positive consumption at the demand node s, there must exist at least one path of positive flow from some supply node $k^* \in K(s)$ to s in an optimal market outcome; we denote one of these paths as q^*. By complementary slackness, $w_{ij} = 0$ for all (i,j) in path q^*. From equation (3), this implies that $\lambda_s - \lambda_{k^*} = p^{q^*}_{k^*s} + \nu^{q^*}_{k^*s}$. Rewriting this equation, we obtain $\lambda_s = \lambda_{k^*} + p^{q^*}_{k^*s} + \nu^{q^*}_{k^*s} \geq \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid q \in \mathcal{P}_{(k,s)},\ k \in \mathcal{K}(s)\}$. □

Proof of Lemma 2 Suppose that the set $\mathcal{S}_I$ represents the set of structurally integrated consumers. By Corollary 1, $\lambda_s = \min\{\lambda_k + p^*_{ks} | k \in \mathcal{K}(s)\},\ \forall s \in \mathcal{S}_I$. Since we are ignoring transportation costs, i.e., $p^*_{ks} = 0\ \forall k \in \mathcal{K}(s)\ \forall s \in \mathcal{S}_I$, then $\lambda_s = \min\{\lambda_k \mid k \in \mathcal{K}(s)\}$. Finally, since $\mathcal{K}(s)$ is the same for all $s \in \mathcal{S}_I$, the equilibrium prices λ_s must all be equal, i.e., $\lambda_s = \lambda'_s\ \ \forall s, s' \in \mathcal{S}_I$. □

Proof of Theorem 1 ($\Rightarrow$) We first show that price differences between s and r cannot exceed this bound for any set of feasible welfare and cost functions. Since s and r are structurally integrated, by definition both s and r have access to the same set of suppliers $k \in \mathcal{K}_{(s)}$. Since we assume that each consumer $s \in \mathcal{S}$ has positive consumption in the optimal market outcome, we let $\hat{k} \in \mathcal{K}_{(s)}$ denote a supply node such that there exists flow from $\hat{k}$ to s in the optimal market outcome. From Equation (3), this implies that $\lambda_s = \lambda_{\hat{k}} + p^*_{\hat{k}s}$. From Corollary 1, this also implies that $\lambda_r \leq \lambda_{\hat{k}} + p^*_{\hat{k}r}$. The equations can be combined into the following inequality: $\lambda_s - \lambda_r \geq p^*_{\hat{k}s} - p^*_{\hat{k}r}$. However, since we do not specify $\hat{k} \in \mathcal{K}_{(s)}$, the following inequality must hold: $\lambda_s - \lambda_r \geq \min\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}(s)\}$. We can now make the same set of logical statements for the node r, thus bounding this inequality from the reverse direction. In doing so, we obtain the inequality $\lambda_s - \lambda_r \leq \max\{p^*_{ks} - p^*_{kr} | k \in \mathcal{K}_{(s)}\}$, which completes the proof.

($\Leftarrow$) We argue by contrapositive; suppose the network is not structurally integrated. We show now that for any non-structurally integrated network, it is possible to find feasible welfare and cost functions such that there does not exist a finite bound between price differences. Without loss of generality, suppose supply node $\hat{k} \in \mathcal{K}(s)$ but $\hat{k} \notin \mathcal{K}_{(r)}$. Let the production cost functions be linear, of the form $W_k(b_k) = y_k b_k,\ \forall k \in \mathcal{K}_{(r)}$ and $W_k(b_{\hat{k}}) = y_{\hat{k}} b_{\hat{k}}$ for supply node $\hat{k}$. When the set of cost function coefficients satisfy $y_{\hat{k}} + p^*_{\hat{k}s} \leq \min\{y_k + p^*_{kr} \mid \forall k \in \mathcal{K}_{(r)}\} + \epsilon$, the prices satisfy $\lambda_s \leq \lambda_r + \epsilon$ in the equilibrium. This is obtained by using Corollary 1, which states that $\lambda_s \leq y_{\hat{k}} + p^*_{\hat{k}s}$

and $\lambda_k = \min\{y_k + p^*_{kr} \mid \forall k \in \mathcal{K}_{(r)}\}$. Thus, by considering cost functions that satisfy increasing values of ϵ, we can increasingly drive apart the values λ_s and λ_r; the price difference thus cannot be bounded. □

Proof of Proposition 1 Let Δ be some arbitrary value within the neutral band for a pair of structurally integrated consumers $s, r \in \mathcal{S}_I$, i.e.,

$$\min\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\} \leq \Delta \leq \max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\}$$

Let $k_1 \in \arg\min\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\}$ and $k_2 \in \arg\max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}_{(s)}\}$ denote two supply nodes shared by $s, r \in \mathcal{S}_I$. Let the production costs at all nodes $k \in \mathcal{K}$ be of the form $W_k(b_k) = y_k b_k$. Let

$$y_{k_2} = y_{k_1} + p_{k_1 s} - p_{k2r} - \Delta,$$

and let

$$y_k \geq \max\{p^*_{ks} + p^*_{kr} \mid k \in \mathcal{K}\} \ \forall k \in \mathcal{K} \backslash \{k^{min}, k^{max}\}.$$

Without loss of generality, we can normalize y_{k_1} by setting it to zero, i.e., $y_{k_1} = 0$. For all production cost functions that satisfy these conditions, all supply nodes except k_1 and k_2 become irrelevant; the production costs at other nodes are simply too high. Under these conditions, we claim that $\lambda_s = p_{k_1 s}$ and $\lambda_r = p_{k_1 s} - \Delta$, implying that $\lambda_s - \lambda_r = \Delta$. First, we need to show that for node s, buying and shipping from k_1 (with cost $p_{k_1 s}$) is less than buying and shipping from k_2 (with cost $y_{k_2} + p_{k_2 s}$). We prove this below:

$$\begin{aligned} y_{k_2} &= p^*_{k_1 s} - p^*_{k2r} - \Delta \\ &\geq p^*_{k_1 s} - p^*_{k_2 r} - p^*_{k_2 s} + p^*_{k_2 r} \quad \text{from def. of } k_2 \\ &= -p^*_{k_2 s} + p^*_{k_1 s}. \end{aligned}$$

This implies that $p^*_{k_1 s} \leq y_{k_2} + p^*_{k_2 s}$ and that in equilibrium, $\lambda_s = p^*_{k_1 s}$. Now, we show that for node r, buying and shipping from k_2 (with cost $y_{k_2} + p_{k_2 r}$) is less than buying and shipping from k_1 (with cost $p_{k_1 r}$). This is shown below:

$$\begin{aligned} y_{k_2} &= p^*_{k_1 s} - p^*_{k2r} - \Delta \\ &\leq p^*_{k_1 s} - p^*_{k_2 r} - p^*_{k_1 s} + p^*_{k_1 r} \quad \text{from def. of } k_1 \\ &= -p^*_{k_2 r} + p^*_{k_1 r}. \end{aligned}$$

Thus, in equilibrium, $\lambda_r = y_{k_2} + p^*_{k_2 r} = p^*_{k_1 s} - \Delta$. We have proved that for any Δ within the neutral band, there exists a set of cost functions for which $\lambda_s - \lambda_r = \Delta$.

Proof of Theorem 2 From equation (7), let

$$w_r - w_s = \max\{\lambda_r - \lambda_k - p^*_{kr} \mid k \in \mathcal{K}\} - \max\{\lambda_s - \lambda_k - p^*_{ks} \mid k \in \mathcal{K}\}.$$

Let $\hat{k} \in \arg\max\{\lambda_r - \lambda_k - p^*_{kr} \mid k \in \mathcal{K}\}$. Then

$$\begin{aligned} w_r - w_s &\le (\lambda_r - \lambda_{\hat{k}} - p^*_{\hat{k}r}) - (\lambda_s - \lambda_{\hat{k}} - p^*_{\hat{k}s}) \\ &= \lambda_r - \lambda_s + p^*_{\hat{k}s} - p^*_{\hat{k}r} \\ &\le \lambda_r - \lambda_s + \max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}\}. \end{aligned}$$

Thus, $\lambda_s - \lambda_r \le w_s - w_r + \max\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}\}$. This completes the one side of the inequality. For the other inequality, let $\hat{k} \in \arg\min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}\}$, and we use equation (8) to obtain

$$\begin{aligned} \lambda_s - \lambda_r &= \min\{\lambda_k + p^*_{ks} \mid k \in \mathcal{K}\} + w_s - \min\{\lambda_k + p^*_{kr} \mid k \in \mathcal{K}\} - w_r \\ &\ge \lambda_{\hat{k}} + p^*_{\hat{k}s} - \lambda_{\hat{k}} + p^*_{\hat{k}r} + w_s - w_r \\ &\ge \min\{p^*_{ks} - p^*_{kr} \mid k \in \mathcal{K}\} + w_s - w_r. \end{aligned}$$

This completes the proof. □

Proof of Proposition 2 Let $\hat{k} \in \mathcal{K}$ be a supply node for which the flow on (i, j) originates. Then

$$\begin{aligned} \lambda_j &= \lambda_{\hat{k}} + p^*_{\hat{k}i} + c_{ij} + \nu_{ij} \\ &\ge \min\{\lambda_k + p^*_{kj} \mid k \in \mathcal{K}\} + v_{ij}. \end{aligned}$$

From lemma 1 we see that

$$\begin{aligned} \lambda_j &= \min\{\lambda_k + p^q_{ks} + \nu^q_{ks} \mid k \in \mathcal{K}, q \in \mathcal{P}_{(k,s)}\} \\ &\le \min\{\lambda_k + p^q_{kj} \mid k \in \mathcal{K}, q \in \mathcal{P}_{(k,j)}\} + \nu_{ij} \\ &= \min\{\lambda_k + p^*_{kj} \mid k \in \mathcal{K}\} + \nu_{ij} \end{aligned}$$

Thus, $\lambda_j = \min\{\lambda_k + p^*_{kj} \mid k \in \mathcal{K}\} + \nu_{ij}$, and $w_j = \nu_{ij}$. □

Proof of Theorem 3 Recall that $\delta^{\min}_e(s) = \min\{\delta_e(k, s) \mid k \in \mathcal{K}\}$ and $\delta^{\max}_e(s) = \max\{\delta_e(k, s) \mid k \in \mathcal{K}\}$, where the value of $\delta_e(k, s)$ measures the cost difference between the minimum-cost path from k to s, and the minimum-cost path from k to s that does not include link e. When $\nu_e < \delta^{\min}_e(s)$, this by definition implies that any flow into s must have traveled on the link e, and hence incurred the full value of ν_e. Similarly, if $\nu_e = \delta^{\min}_e(s)$, then any flow into s must have either

traveled on link e, or an alternative path that does not include link e but is at least as expensive as $p^*_{ks} + \nu_e$, by the definition of $\delta_e(k,s)$. On the other hand, the definition of $\delta^{\max}_e(s)$ and equation (4) imply that for each $k \in \mathcal{K}$, there exists a path q which does not include link e from k to s such that $\lambda_s - \lambda_k \leq p^q_{ks} \leq p^*_{ks} + \delta^{\max}_e(s),\ \forall k \in \mathcal{K}$. Since path q does not include link e and therefore does not have any links with positive shadow price, this bound must hold, which implies that the congestion surcharge of node s with respect to link e must be less than or equal to $\delta^{\max}_e(s)$. □

Proof of Proposition 3 Let $r \in \mathcal{S}$ denote an arbitrary node for which to compare all other nodes with. Let $\eta^t = \lambda^t_r - w^t_r\ \forall t \in \mathcal{T}$, and let $\rho_s = \rho_{rs}$ and $\alpha_s = \alpha_{rs}$ denote the mid-point and half-width of the neutral band, as defined in Section 4.2. By equation (10), all prices in the market over any set of welfare and cost functions can be expressed in the form of equation (13). □

Electronic Companion

EC.1. Supplemental Example for Section 4

We use a simple example to show that network neutral band can be strictly tighter than pairwise bands. Suppose the market is represented by the network in Figure EC.1, where s_1, s_2 denote the two demand nodes in the market and k_1, k_2 denote the supply nodes. Since there are direct links between s_1 and s_2 with transportation cost equal to 1, the pairwise neutral band width has value 1. However, when examining the network holistically, equation (6) implies that the prices at s_1 and s_2 must always be equal. The example highlights the importance of considering the entire network even when analyzing subsets of market participants.

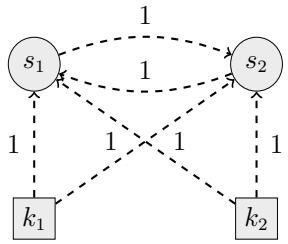

Figure EC.1 An example of a network where consumers are directly connected to each other.

EC.2. Supplemental Table for Section 6

Table EC.1 Summary statistics of daily gasoline prices per gallon for 2016 and 2017

		2016					2017				
Label	City, State	Mean	Std	Min	Max	Range	Mean	Std	Min	Max	Range
A	Houston, TX	1.93	0.17	1.51	2.16	0.65	2.21	0.12	2.04	2.53	0.49
B	Baton Rouge, LA	1.87	0.18	1.47	2.11	0.64	2.13	0.09	1.98	2.33	0.35
C	New Orleans, LA	1.93	0.18	1.52	2.14	0.62	2.18	0.10	2.03	2.42	0.39
D	Jackson, MS	1.89	0.16	1.49	2.12	0.63	2.13	0.12	1.94	2.44	0.50
E	Birmingham, AL	1.91	0.19	1.49	2.17	0.68	2.14	0.14	1.94	2.5	0.56
F	Atlanta,GA	2.21	0.20	1.76	2.62	0.86	2.42	0.15	2.24	2.91	0.67
G	Nashville, TN	2.02	0.21	1.52	2.35	0.83	2.25	0.16	2.07	2.68	0.61
H	Columbia, SC	1.90	0.16	1.54	2.11	0.57	2.11	0.15	1.87	2.56	0.69
I	Charlotte, NC	2.04	0.16	1.67	2.26	0.59	2.28	0.14	2.06	2.65	0.59
J	Greensboro, NC	2.07	0.18	1.66	2.33	0.67	2.30	0.12	2.13	2.63	0.50
K	Raleigh, NC	2.09	0.15	1.73	2.31	0.58	2.30	0.13	2.10	2.62	0.52
L	Richmond, VA	1.94	0.19	1.48	2.17	0.69	2.19	0.11	2.00	2.50	0.50
M	Virginia Beach, VA	1.95	0.19	1.51	2.21	0.70	2.18	0.12	2.00	2.52	0.52
N	Jacksonville, FL	2.05	0.16	1.68	2.35	0.67	2.33	0.14	2.11	2.71	0.6
O	Orlando, FL	2.06	0.18	1.64	2.35	0.74	2.32	0.15	2.05	2.72	0.67
P	Tampa, FL	2.06	0.16	1.66	2.41	0.75	2.32	0.16	2.02	2.73	0.71
Q	Miami, FL	2.22	0.16	1.83	2.46	0.63	2.46	0.14	2.24	2.79	0.55

EC.3. Additional discussion on neutral bands and price dynamics

In this section, we provide further discussion on the neutral band and price decomposition by explicitly mapping out the equilibrium flows that correspond to different price relationships. We use Figure 2 in the paper to illustrate these ideas; the figure is replicated below as Figure EC.2. When the network is uncapacitated, the difference between the equilibrium price of demand nodes s_1 and s_2, i.e., $\lambda_{s_2} - \lambda_{s_1}$, will always be equal to zero in the network of Figure EC.2a, be equal to one in the network of Figure EC.2b, and be within negative one and positive one in the network of Figure EC.2c. We first elaborate on this result below, focusing on a single time period, then extend the discussion to include temporal market dynamics and effects of capacity constraints.

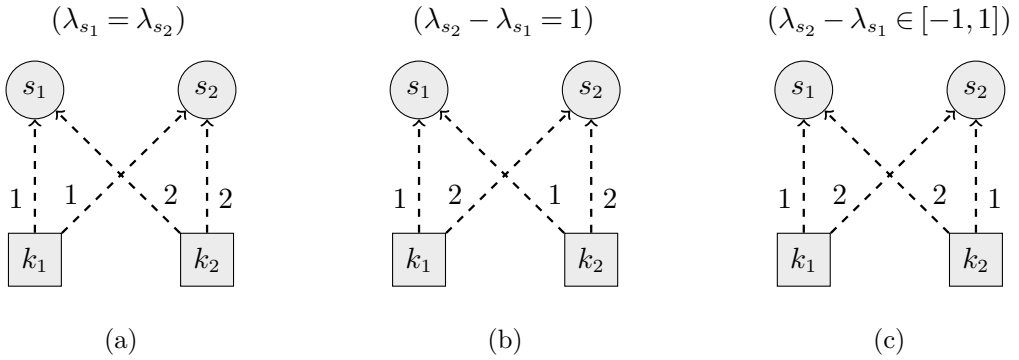

Figure EC.2 Three instances of a four-link, uncapacitated network where the transportation costs on each link are as labeled.

EC.3.1. Neutral bands and transportation bases

To facilitate this discussion, we introduce the concept of a transportation basis, which defines the set of links with positive flow in equilibrium (Murphy and Mudrageda 1998). Using mathematical programming terminology, the set of links within each basis has reduced cost of zero, as given by the complementarity conditions. The set of all possible transportation bases for the nonlinear market allocation problem over the four-link networks of Figure EC.2 is shown in Figure EC.3. Below, we show that the bounds of the neutral band are defined by bases ((a),(c)) and ((b),(d)), whereas bases (e) and (f) capture equilibrium scenarios where prices differences can be strictly within the neutral band.

- **Network EC.2a:** In basis (a), both demand nodes pay a marginal cost of 1 unit to ship from supplier k_1, whereas in basis (b), both demand nodes pay 2 units to ship from supplier k_2. In both cases, the prices at demand nodes s_1 and s_2 will be the same, which implies that this must be true for *any* basis. Specifically, λ_{s_1} and λ_{s_2} must be equal in (c) and (d) because they

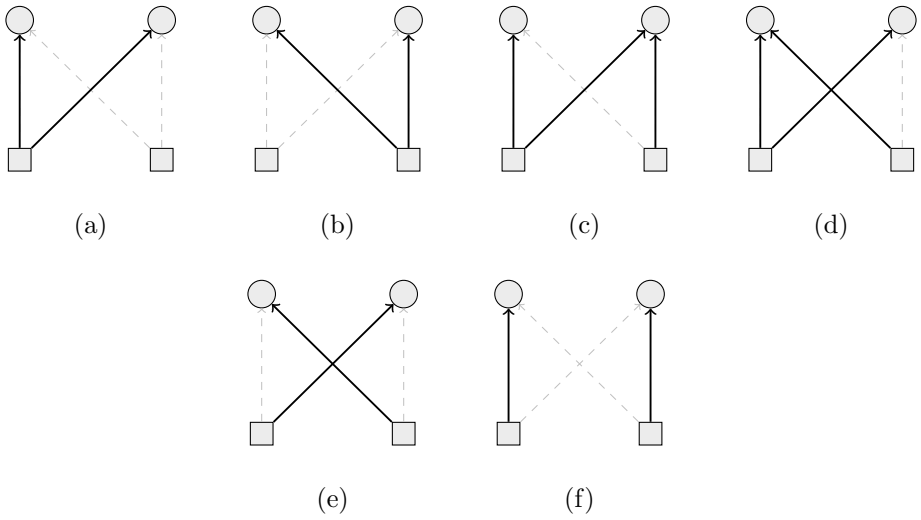

Figure EC.3 The list of feasible transportation bases for a simple four-link network of two supply nodes and two demand nodes. We denote (a)-(e) as connected bases, whereas (f)-(g) are disconnected bases.

contain either (a) or (b), respectively, whereas λ_{s_1} and λ_{s_2} in basis (f) and (g) must be equal because the reduced costs on all links (including those not in the basis) must be zero.

- **Network EC.2b:** In basis (a), demand node s_1 pays a marginal cost of 1 unit to ship from supplier k_1 while s_2 pays 2 units. Similarly, in basis (b), demand node s_1 pays a marginal cost of 1 unit to ship from supplier k_1 while s_2 pays 2 units. Thus, λ_{s_1} will always be 1 unit lower than λ_{s_2}, as the same argument used in the previous bullet follows.
- **Network EC.2c:** In basis (a), demand node s_1 pays exactly 1 unit less than s_2 to ship from k_1, which implies that the equilibrium price λ_1 must be exactly 1 unit lower than λ_2. On the other hand, in basis (b), the opposite is true; i.e., λ_2 is exactly 1 unit lower than λ_1. The same is true for bases (c) and (d) respectfully, which contain (a) and (b). The bases (f) and (g) represent "disconnected" bases in which absolute price differences between s_2 and s_1 can be of different values strictly less than 1 unit, which occurs when the reduced costs of the links not in the bases take on positive values.

EC.3.2. Price dynamics and shifts in the basis

We now introduce the temporal component, specifically, by considering equilibrium conditions over a time horizon with varying welfare and cost functions. We also include additional discussion on the effects of capacity constraints. We use the network in Figure EC.2c as a running example.

Price dynamics in the absence of basis shifts. When we incorporate temporal market dynamics into the model, we can also use the bases to examine price dynamics. For the bases (a)-(d), as long as the basis doesn't change over the time horizon, then the price differences between

the two demand nodes will always be fixed; the prices will rise and fall together and the differences will not change. This is not true however for the disconnected bases (e) and (f). Even if the basis does not shift (i.e., the set of links on which there is positive equilibrium flow does not change), the price differences between the two demand nodes can vary within the neutral band. For example, if we assume that equilibrium conditions at a specific time period over the network in Figure EC.2c is represented by basis (e) or (f), then it is possible for the price at node s_1 to rise or fall by a small incremental amount in the next time period while the price at node s_2 remains constant.

Price dynamics without congestion. In the absence of binding capacity constraints, any demand or supply shocks will at most push the basis from a disconnected one ((e),(f)) to a connected one ((a)-(d)). Specifically, given initially market conditions rendering a disconnected basis, incremental shifts in the welfare or cost functions may lead to non-uniform price changes across the two demand nodes, but only by a small amount (i.e., the value of the reduced costs of unused links) before the equilibrium conditions form a connected basis. Once this occurs, then prices will move synchronously, with the price difference being one of the bounds of the neutral band. Thus, over a time horizon where welfare and cost functions vary, we would expect prices to move in a "loosely" synchronized fashion where price differences may vary independently within a small band but will never exceed the bounds of the band.

Price dynamics with congestion. When there are binding capacity constraints, price differences between the two nodes no longer need to conform to any bound. For example, if there is a bottleneck on the link (k_1, s_2) in Figure EC.2c, then it is possible for s_2 to have a price that is much higher than s_1. A simple illustration of this idea is to keep driving the supply cost of k_1 lower holding everything else constant, which will only reduce the cost for node s_1 but not for node s_2. In this setting, the congestion surcharge on node s_2 (i.e., the reduced cost on link $(k_1,\ s_2)$) will keep increasing (decreasing).

General insights and price dynamics over larger markets. In a large, interconnected market with many demand and supply nodes, fluctuations in market conditions can consistently cause many different shifts in the transportation bases of the underlying market. However, in the absence of binding capacity constraints, these shifts can only render prices across any pair of demand nodes to differ by at most some bound. Furthermore, the distance between these bounds, i.e., the neutral band width, may be small in many settings; in Example 2 we discuss how small neutral band widths are expected when demand or supply nodes are clustered in close proximity and when transportation costs are a function of distance. On the other hand, when there are binding capacity constraints, the price differences are able to exceed these bounds by arbitrary amounts. Furthermore, if common bases have a general chain-like structure across the demand nodes (i.e.,

where there are chains of positive flows across demand nodes), then the effects of a binding capacity constraint can be reflected across many different nodes by causing many equilibrium prices to break from the neutral bands.

EC.4. An Econometric Comparison and Additional Empirical Insights

A large number of different econometric time series methods have been applied to study commodity market prices; we refer to Dukhanina and Massol (2018) for an overview of these methods. The SEM model developed in our paper differentiates itself from this literature primarily in its focus on capturing spatiotemporal price effects, particularly of those stemming from underlying network capacity constraints. The model can easily handle multivariate time series data and is developed directly from price decomposition describing networked and capacitated markets.

In this section, we illustrate the differences and advantage of our SEM model in analyzing capacitated networked markets. We do so by comparing it with two econometric methods applied our data set. We first consider the standard two-stage cointegration approach proposed by Engle and Granger for analyzing pairwise time series data (Engle and Granger 1987). Such methods have been applied in many studies of natural gas and gasoline prices (e.g., De Vany and Walls 1993, Doane and Spulber 1994, Paul et al. 2001, and references in Section 2.2). We then consider a panel regression model with spatial and temporal fixed effects. We show that both approaches are less interpretable, and more importantly, cannot be used to identify the spatiotemporal price effects that reflect bottlenecks in the underlying network.

EC.4.1. Cointegration tests

We follow the standard procedure for the Engle and Granger two-stage cointegration tests. We first use the augmented Dickey-Fuller test to confirm that all time series of prices are not stationary but rather integrated of order one. Given this, we apply the Engle and Granger cointegration model, and our results are illustrated in Figure EC.4. Specifically, Figures EC.4a and EC.4c show the pairs of locations that have a p-value that is less than 0.05, whereas Figures EC.4b and EC.4d show the pairs that have a p-value equal to or greater than 0.05.

As evident through Figure EC.4, cointegration tests produce results that are difficult to interpret over the large geographical region considered. Since the concept of cointegration is not transitive (i.e., two pairs that are cointegrated to a third series may not be cointegrated with each other), it becomes difficult to obtain a holistic perspective on the market. Furthermore, some of the results appear to directly contradict our own findings and what is to be expected of this market. For example, the tests indicate that prices between gasoline producing cities in Texas and Louisiana are not cointegrated. In contrast, our SEM results showed little to no congestion surcharge among

(a) Cointegrated pairs in 2016

(b) Non-cointegrated pairs in 2016

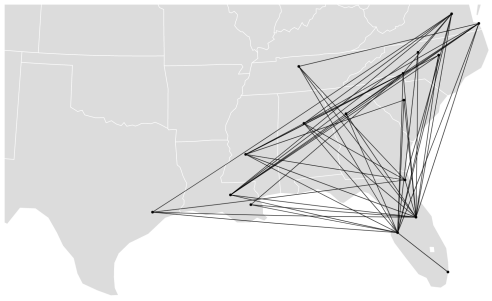

(c) Cointegrated pairs in 2017

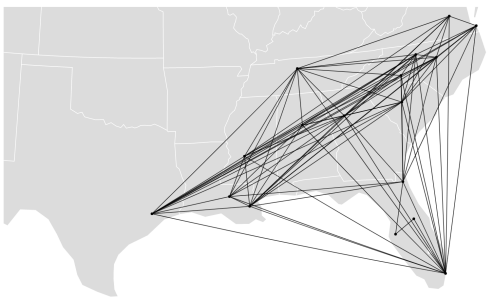

(d) Non-cointegrated pairs in 2017

Figure EC.4 Results from the cointegration tests. Links appearing in Figures EC.4a and EC.4c represent pairs of locations with p-values less than 0.05, whereas Figures EC.4b and EC.4d show pairs of locations that have p-values equal to or exceeding 0.05.

any of these cities. To the best of our knowledge, no major bottlenecks are likely to exist between these geographically-close locations. A more prominent example is that Nashville is shown to be cointegrated with cities in Louisiana and Florida. This again directly contradicts our results, where we find Nashville to have the highest congestion surcharge in both years, a result that is consistent in both time and space with the fact that Nashville relies extensively on pipeline infrastructure that was disrupted in both years. Despite these contradictory results, the network of non-cointegrated pairs in Figure EC.4d and especially Figure EC.4b do seem to be rather dense between inland cities along the pipelines, which suggest some detection of pipeline bottleneck effects.

In summary, while the cointegration tests considered here are popular and easy to implement, they appear to be insufficient in providing a clear picture of the price effects of congestion events. Furthermore, they cannot be used to identify granular spatiotemporal price variations such as the ones identified by the SEM model.

EC.4.2. Panel regression analysis

We consider a simple panel regression model with fixed effects on the price data. Specifically, we consider the following model,

$$y_s^t = \beta_s x_s + \beta_t x_t + \epsilon_s^t, \quad \forall s \in \mathcal{S}, t \in \mathcal{T}, \tag{EC.1}$$

where y_s^t represents the price at time t and location s, and x_s and x_t are dummy variables with a value equal to one when the time and location is t and s, respectively. Note that this panel regression model resembles our price decomposition in equation (13). Specifically, it is equation (13) without η^t (the trend term) and w_s^t (the congestion surcharge terms).

We use the standard ordinary least squares technique to fit equation (EC.1). The model results in an adjusted R^2 value of 0.86. For the sake of space, we do not provide the summary statistics of all 747 coefficients but instead remark that the majority of coefficients, with the except of a few coefficients in January 2016, had a p-value of less than 0.01. More importantly, the residuals of the panel regression provide us with spatiotemporal values with which we can compare the results of the SEM model.

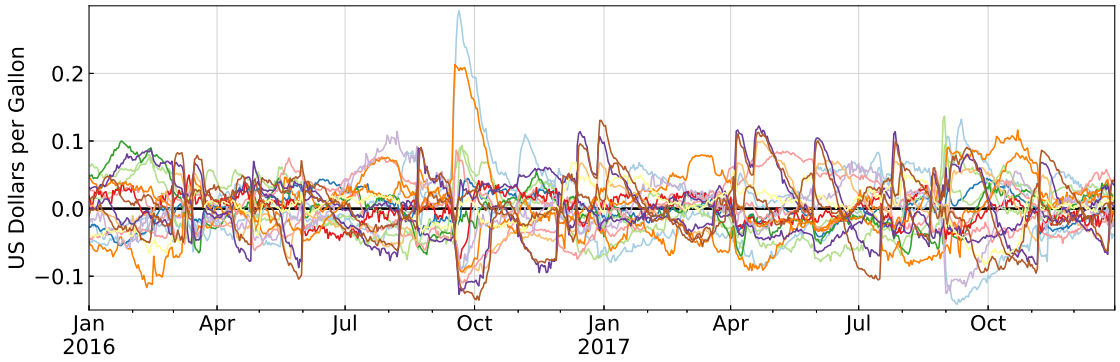

Figure EC.5 Residuals from panel data regression model.

The residuals across all locations are plotted in Figure EC.5. We first make two general observations that relate these residuals to our own results from the case study. First, it is evident from the figure that there are two locations with positive residuals that are much larger in magnitude than those of the other locations around September and October 2016. These large residuals are in line with the results of our SEM model, which detected surcharge estimates following the periods of a major pipeline disruption. Second, the magnitude of the positive residuals over all the other locations in September and October 2016, as well as over any location during the 2017 hurricane season, are similar to the magnitudes of residuals throughout the two year time horizon. Thus, with the exception of two locations in late 2016, the spatiotemporal residuals from the panel regression model during the two specific time windows of the case study (which follow known capacity disruptions) cannot be differentiated from those across the general time horizon.

In Figures EC.6a and EC.6b we further examine the positive residuals over the two time windows of interest. In the figures, we label the same cities that are labeled in the inset figures of Figures 7 and 8 in the paper, which represent cities for which we detect surcharge from the SEM model.

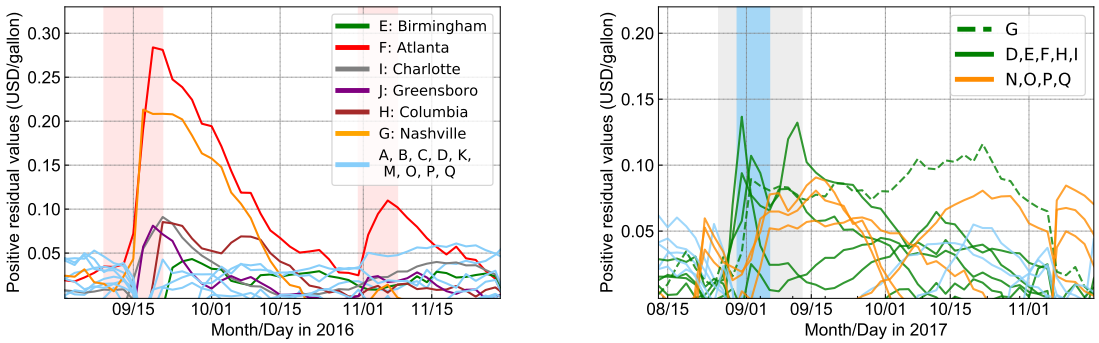

(a) September to November 2016 (b) August to November 2017

Figure EC.6 Plots of the positive residuals from the panel regression, over the two time periods of interest.

The main observation here is that within the two time windows, the magnitudes and trends of the positive residuals for the labeled cities are in fact quite similar to the surcharge estimates we derived from our SEM model.

We examine Figure EC.6a. The residuals over several of the labeled cities increase sharply a few days after the first pipeline disruption. Much like our own surcharge estimates, the residuals peak quickly and then slowly subside. The price effects of the second pipeline disruption is observed, although only in Atlanta and at a much lower magnitude. We also observe that the residuals across all the cities (A - Q) become positive at some point within this time window. These last two observations on price residuals are in contrast to the surcharge estimates of our SEM model. Specifically, we do not detect any surcharge after the second pipeline disruption, and we also do not identify any surcharge except across the labeled cities which are all downstream of the disrupted pipeline segment. One of the reasons behind these differences results from the conservative approach of the SEM model which measures only the values that exceed neutral bands. This can be further noted by observing that the positive residuals are all of higher magnitude than the surcharge estimates derived from the SEM model. Finally, similar observations can also be made when we compare Figure EC.6b to the inset of Figure 8.

The similarity in the trends of the positive residuals and our surcharge estimates can be explained by the idea that in both models, we are using time and space fixed effects, i.e., node-invariant and time-invariant variables, to fit rich spatiotemporal data. The spatiotemporal variation in prices that cannot be captured by these variables are reflected as residuals in the panel regression model, whereas the surcharge estimates from the SEM model reflect the variations which exceed (simultaneously estimated) time-invariant neutral bands. Thus, in both cases, the trends (but not the magnitude) of the spatiotemporal variation would be quite similar.

This econometric comparison enables us to highlight the main advantage of the SEM model. By fitting neutral bands to each location, the SEM model is able to systematically differentiate

surcharge effects from general, idiosyncratic price movements. Specifically, the positive residuals in the panel regression model generally cannot be differentiated from each other over any point in time over the two year time horizon. On the other hand, the SEM is able to identify exactly the spatiotemporal price effects associated with the documented capacity disruptions, and almost nothing more.