On the Price Spread of Benchmark Crude Oils: A Spatial Price Equilibrium Model

Bennett, Max and Yuan, Yue

Washington University in St. Louis, University of Chicago

18 October 2016

On the Price Spread of Benchmark Crude Oils: A Spatial Price Equilibrium Model*

Max S. Bennett
Washington University in St. Louis
max.bennett@wustl.edu

Yue Charles Yuan
University of Chicago
yuany@uchicago.edu

October 18, 2016

Abstract

Benchmark crude oils exhibited dramatic fluctuations in price spreads in the recent decade, a phenomenon that rarely occurred in earlier decades. This paper develops a rational expectations two-period model of spatial price equilibrium, and departs from standard models by assuming increasing marginal costs of transportation and storage. We econometrically validate our model using a dataset that covers an extended time period. The model allows us to determine the underlying causes of the unique phenomenon of drastically changing crude oil price spreads over the past decade.

*We are grateful to Andrew Atkeson, Pablo Fajgelbaum, Sebastian Galiani, François Geerolf, Bruce Petersen, Dorothy Petersen, George S. Tolley, Pierre-Olivier Weill, and seminar participants at Chicago, UCLA, and WUSTL for helpful comments. An earlier version of the paper was circulated under the title "Spatial Price Equilibrium with Convex Marginal Costs of Transportation: Applications to the Brent-WTI Spread."

1. Introduction

Crude oil is critically important for the world economy. To the average consumer of oil, however, it's often easy to get the impression that there is a single global market for crude oil. In reality, there are many different types of crude oil, and there are benchmark oils that serve as references for buyers and sellers of crude oil around the world. Two primary global benchmark oils are West Texas Intermediate (WTI) and Brent Blend. WTI is priced in Cushing, Oklahoma and used primarily in the U.S., whereas Brent is priced in the United Kingdom and used primarily in Europe but also all around the world. In addition, there is another major benchmark oil in the U.S. called Light Louisiana Sweet (LLS), which is priced in St. James, Louisiana near the Gulf Coast.

The most easily refined crude oil, and thus the most valuable, is light sweet crude. WTI, Brent, and LLS are all light sweet crude oils, and are almost identical in physical composition. As such, any substantial deviation in price between these crude oils can only be a consequence of spatial price equilibrium and not differences in intrinsic value.

Historically, Brent, LLS, and WTI have traded with very small price differentials (see Figure 1). Prior to 2011, the price differential between WTI and Brent had never been greater than $5 dollars-per-barrel. Arbitrage between the two markets seemed to ensure that localized supply and demand shocks affected each price relatively equally: "world" oil prices moved in tandem. However, in the beginning of 2011, the Brent-WTI and LLS-WTI spreads started to increase dramatically, and in the months that followed the Brent-WTI spread widened to as much as $25 per barrel, while the Brent-LLS remained relatively small. The dramatic increase in Brent-WTI spread generated a great deal of median attention, most notably because its cause was so unclear. Many media describe this as a result of the "bottleneck" problem in Cushing, OK, which is consistent with some academic studies such as Buyuksahin et al. (2013), but the actual source of shocks that induced the bottleneck problem has not been convincingly identified.

The decoupling of Brent and LLS from WTI is of practical concern for two reasons.

Figure 1: *Price Spreads*

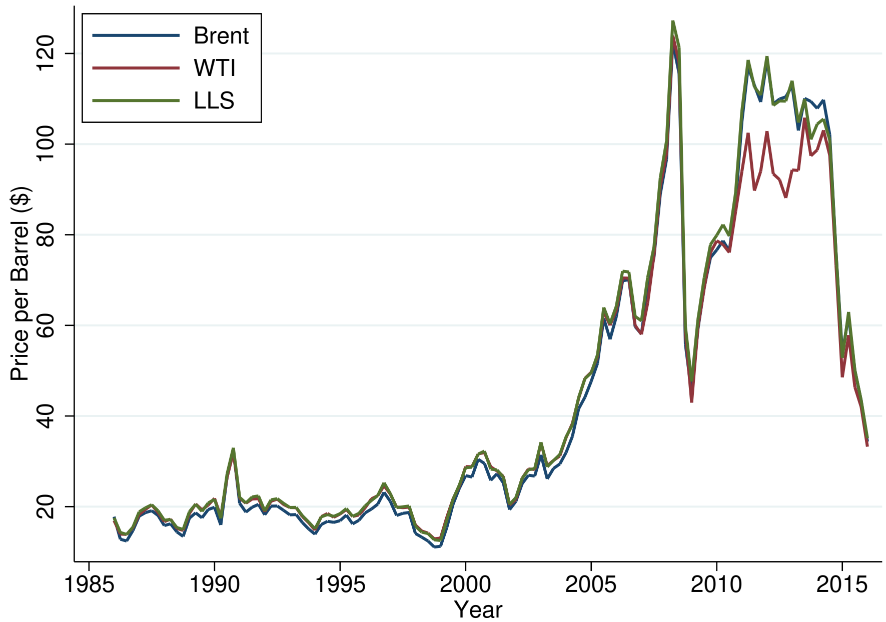

(a) *Quarterly Brent, WTI, and LLS spot prices from 1986Q1 to 2016Q1. 1986 is the first year Bloomberg reports data for all three spot prices.*

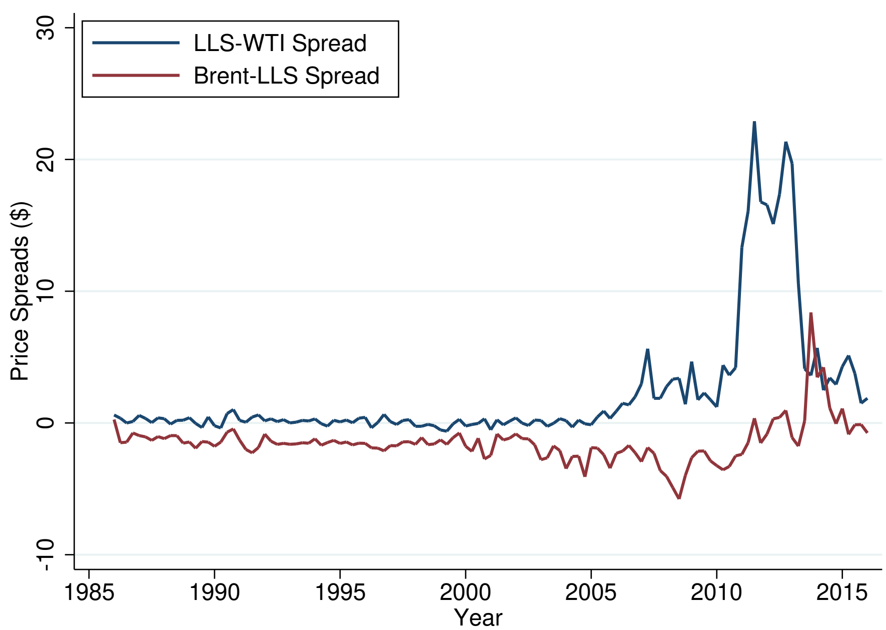

(b) *Quarterly LLS-WTI and Brent-LLS spot price spreads from 1986Q1 to 2016Q1.*

First, market participants in the petroleum industry peg crude oil prices to benchmarks, and if one benchmark is an inaccurate or unrepresentative gauge of the oil prices then it ceases to be useful. Second, WTI and Brent are the underlying crude oils for the New York Mercantile Exchange (NYMEX) and Intercontinental Exchange (ICE) crude oil futures contracts, which are by far the largest crude oil futures markets and are used by economic agents around the world to hedge their exposure to the price fluctuations of oil. If either WTI or Brent ceases to be a good indicator of oil prices, then economic agents will be hindered in their ability to effectively hedge themselves.

Unfortunately, dramatic changes in spatial price spreads of tradable commodities (e.g. oil) cannot be readily explained with standard models in the economic literature. The literature on spatial price equilibrium, whereby no-arbitrage conditions are applied to geographic price differentials of a homogenous commodity, has generally used the existence of transportation costs to rationalize price spreads. Standard models in this literature assume that the marginal cost of transportation is constant. However, for the purposes of this paper, we reevaluate and modify three aspects of the standard: (i) choosing a geographic region where the no-arbitrage condition is applicable; (ii) tying together transportation and storage in the context of the crude oil markets; and (iii) assessing the assumption of the constant marginal cost of transportation.

Applicability of the no-arbitrage condition. The no-arbitrage condition, or the law of one price, should not be taken for granted in the commodities markets. After all, the commodities market has its own idiosyncratic characteristics that may cause the no-arbitrage conditions to fail. Ardeni (1989) uses tests of nonstationarity and co-integration for a group of commodities and show that the law of one price fails as a long-run relationship. He argues that the failure of the law of one price can be rationalized with two factors, namely high costs of arbitrage and institutional barriers.

Other papers have shown results consistent with the arguments of Ardeni (1989). With regard to the costs of arbitrage, Fattouh (2010) studies the dynamics of crude oil price differentials and concludes that crude oil prices are linked and at the general level, the oil market is "one great pool." However, he contends that oil markets are not necessarily integrated in every time period, in the sense that costly and risky arbitrage may cause decoupling of crude oil prices. With regard to institutional barriers, Richardson (1978) documents that the law of one price fails uniformly for commodities arbitrage between the U.S. and Canada. A more recent paper, however, seems to suggest new and extensive interactions of commodity price formulation among different countries. Olsen, Mjelde and Bessler (2015) study the law of one price for 11 natural gas market prices of the U.S. and Canada. They find out that markets geographically adjacent to each other tend to be more highly integrated than markets separated by distance. This higher degree of integration among some cross-border commodity markets can be attributed to deregulation, technological advances, and trade agreements. Such convergence is well documented in other industries for other regions, such as the European car market as discussed in Goldberg and Verboven (2005). Nevertheless, a common theme in the literature is that institutional barriers, especially such factors as exchange rates, still cause the no-arbitrage condition to fail in the commodity markets.

As a result, we would like to focus on geographically adjacent regions within the United States in order to ensure the applicability of the no-arbitrage condition. Importantly, Werner (1987) proves that the existence of a price system that admits no arbitrage opportunity for all consumers is sufficient for the existence of an equilibrium, which serves as an additional assurance for our careful selection of geography as we develop a spatial price equilibrium model. Although the Brent-WTI spread has been the core topic of media attention as discussed earlier, the economically interesting spread in this paper is actually the LLS-WTI spread. Since LLS is unquestionably priced at an intermediate step in the transportation between Brent and WTI, the Brent-WTI spread can be looked at as the summation of the Brent-LLS spread and the LLS-WTI spread. In light of this, we notice from Figure 1b that the Brent-LLS spread (arbitrage between the Gulf Coast and Europe) has been largely unchanged except for a few short time periods, whereas in sharp contrast the LLS-WTI spread (arbitrage between Cushing and the Gulf Coast) widened by $25 dollars-per-barrel between 2011 and 2013. Clearly, the widening of the Brent-WTI spread was due to the widening of the LLS-WTI spread and thus a change in crude oil's transportation and storage mechanisms from Cushing to St. James.

Tying together transportation and storage in the context of the crude oil markets. Classic models of spatial price equilibrium, as exemplified by Samuelson (1952), directly solve for spatial price relations based on localized supply and demand in the presence of transportation. In addition, seminal work by Deaton and Laroque (1992, 1996) lays out the foundation for analyzing commodity price dynamics with competitive storage motives. More recently, Knittel and Pindyck (2016) builds a simple supply and demand model that incorporates storage to study speculation in the crude oil market, in order to check whether speculators are to blame for the wild fluctuations in crude oil prices. These papers, though, do not incorporate both transportation and storage into their models at the same time. For our purposes, it is important to tie together both transportation and storage in order to fully characterize the behavior of the price spread of benchmark crude oils.

Williams and Wright (1991) include storage into a spatial price equilibrium model. This is integral because having the option to store a commodity makes possible a role for expectations, and transportation can be characterized through the spatial price equilibrium. Their approach provides some guidance for studying commodity price spread taking into consideration both transportation and storage.

There are two distinct kinds of transportation "costs" that are described in the literature: traditional costs of transport and time costs of transport. Models with instantaneous transportation can rationally exhibit spatial price differentials without violating no-arbitrage assumptions because of transportation costs; an uncontroversial result of such a model is that locational price spreads are bounded by transportation costs (see Williams and Wright, 1991). A second type of cost, as described by Coleman (2009), arises from the time that transportation takes. Coleman (2009) explains how unbounded price differentials can exist in the short-run without violating no-arbitrage assumptions if transportation is non-instantaneous. However, in the context of the LLS-WTI spread, the applicability of a non-instantaneous transportation model is limited: St. James and Cushing are merely 650 miles apart, and thus the delay in transportation is minimal. We will therefore describe the price spread in the context of an instantaneous transportation model. As such, in our model we are constrained by the no-arbitrage condition that the price spread is bounded by the cost of transportation (see Williams and Wright, 1991).

Assessing the assumptions of the constant marginal costs of transportation and storage. Unlike previous models of instantaneous transport, we will assume that the marginal costs of transportation and storage are not constant, which allows for the possibility of rapidly rising spreads in response to changes in the supply of a commodity. As we argue in this paper, the assumptions of constant marginal costs of transportation and storage greatly restrict their applicability to many spatial price phenomena (most notably, the widening of the LLS-WTI spread in the recent decade) and strongly deviate from the actual marginal cost curves of transportation and storage in many contexts. Furthermore, these assumptions conceals potential determinants of a spatial price spread.

For example, we will show in Section 3.2 that if the marginal cost of transportation is truly constant, then the only explanation of a widening spread is an exogenous upward shift in the marginal cost curve of transportation. This could occur for many reasons, such as the transportation market becoming less competitive, energy becoming more expensive, or transportation infrastructure depreciating and becoming unusable. However, as we will demonstrate in this paper, an exogenous upward shift in the marginal cost curve would cause the observed quantity of transportation between pricing points to decrease, but in regards to the LLS-WTI spread, the quantity of transportation between the regions actually increased along with increases in LLS-WTI spread, contradicting the predictions derived from a model that assumes constant marginal cost of transportation. Furthermore, We will go on to include the possibility of storage which will illuminate the endogeneity of the current spread to expected future costs of transportation, a relationship that has not yet been described in the literature. In sum, there is a theoretical and empirical gap in the literature that we seek to fill: we will build and validate a model of spatial price equilibrium with increasing marginal costs of transportation and storage. Our empirical strategy for validating our theoretical model involves exploiting the autocorrelations of the time series data to address endogeneity in the regressions.

In the end, we would like to apply our theory to identify the dramatic changes of the LLS-WTI spread seen in Figure 1. News centers such as the Wall Street Journal and Bloomberg have written numerous articles regarding the causes of the 2011-2013 widening of the Brent-WTI spread. The articles, often quoting analysis performed by financial institutions and energy consultants, largely describe two hypotheses regarding the causes of the spread:

(1) The widening of the Brent-WTI spread was caused by a negative production shock in the Middle East. In particular, Arab Spring and loss of Libyan Oil put upward pressure on Brent prices while transportation constraints between the Cushing and Europe isolated WTI from this effect.

(2) The widening of the Brent-WTI spread was caused by a positive production shock in the Midwest of the United States, specifically Cushing, Oklahoma. In early 2011, new pipelines were built bringing more Canadian oil into Cushing, and transportation constraints between the Cushing and Europe isolated Brent from this effect.

Our model will provide a direct way to test the efficacy of each of these hypotheses as well as other potential causes of the widening spread.The empirical strategy involves exploiting the observed relationships among inventories, transportation, and the LLS-WTI spread (see Figure 3).

The rest of the paper is organized as follows: Section 2 provides background information on the U.S. crude oil market, with a particular focus on the oil markets at Cushing, Oklahoma, the price settlement point for WTI. Some time series evidence is also presented to characterize the dynamic relationships among spread, storage, and transportation. Section 3 builds a model of spatial price equilibrium that incorporates increasing marginal costs of transportation and storage, starting from a simple baseline model generalized from the standard literature. We then come up with a list of testable predictions, as well as identify potential causes of a changing price spread and their respective impacts on transportation and storage. In Section 4, we describe our dataset, and test the model predictions in order to validate our theory. In Section 5, we apply our theory to identify the causes of the changing LLS-WTI spread over time. Section 6 contains some final remarks.

2. U.S. Crude Oil Market Facts

2.1. Oil markets at Cushing, OK

Cushing, Oklahoma is the delivery point for the NYMEX oil futures contract and therefore refineries, storage facilities, and pipelines have all developed substantial infrastructure in the periphery of the city. As of 2011, it is estimated that storage capacity at Cushing is around 48 million barrels of crude oil, and as much as 600 thousand barrels flow into Cushing daily.

Figure 2: *Geographic separation of LLS and WTI*

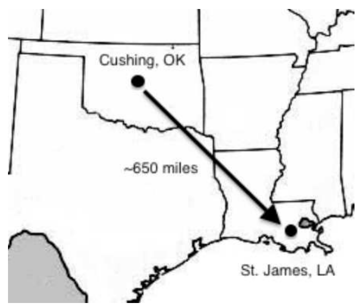

We will use the model developed in this paper to examine the spatial price spread between crude oil at Cushing, Oklahoma (known as West Texas Intermediate or WTI) and that at St. James, Louisiana (known as Light Louisiana Sweet or LLS). As seen in Figure 2, the two locations are approximately 650 miles apart. The transportation infrastructure between Cushing and St. James primarily consists of pipelines. Most pipelines transport liquid at a speed of approximately 3 to 13 miles-per-hour; therefore it should only take between two and nine days to transport crude oil from Cushing to St. James. There are, however, additional modes of transportation between Cushing and St. James: notably, rail and trucking. Transportation from Cushing to St. James is estimated to cost between $7 to $10 per barrel by rail and between $11 and $15 per barrel by trucking, which is much more expensive than the estimated $2 per barrel by pipeline.[1] As noted above, a constant marginal cost curve of transportation has been the standard assumption in the literature,

[1]CommodityOnline, quoting analysis by Bank of American Merrill Lynch.

which is clearly at odds with transportation costs of oil from Cushing to the Gulf Coast if pipeline capacity is insufficient and therefore some oil must be shipped by other, more expensive, modes of transportation.

In late 2010 and early 2011, two large new pipelines that directed oil from Canada to Cushing went online, substantially increasing the availability of crude in Cushing. This occurred over the backdrop of steadily increasing production of crude oil in the Midwest over the past decade. These two forces dramatically increased the quantity of crude flowing into Cushing. Models of spatial price equilibrium with a constant marginal cost of transportation would suggest that this would have no effect on the LLS-WTI spread, as the additional oil would be immediately directed towards the coast via transportation infrastructure. However, we will argue in the model that because of rising marginal costs, a positive supply shock at Cushing is a theoretically sound explanation for the widening of the LLS-WTI spread.

In November 2011, months after the spread initially widened past $10 dollars, a Canadian pipeline company announced that it would reverse a pipeline that transported oil from the Gulf Coast to Cushing.[2] The reversal process can take months, and therefore the announcement represented a decrease in expected future costs of transportation. Nevertheless, the announcement was followed by a $10 dollar decline in the LLS-WTI spread. Previous models cannot rationalize the role of expectations on a spatial price spread. In this paper we will include a storage market to explicitly derive a relationship between the current spatial price spread and expected future costs of transportation. The result will rationalize the empirical behavior of the LLS-WTI spread in response to the change in expectations.

2.2. Time series evidence

We use the impulse response analysis to provide some reduced-form empirical evidence about the dynamic relationships among spread, storage and transportation. In order to

[2] Enbridge was the company that announced it would purchase a pipeline between Cushing and the Coast, and subsequently reverse its flow

do so, we construct a standard orthogonalized VAR system, with quarterly data over the entire sample period from 1986Q1 to 2016Q1. The original VAR includes, in the following order, the LLS-WTI spread, storage in PADD2, storage in the rest of the U.S., and transport from PADD2 to the rest of the U.S., while controlling for exogenous variables including productions and net imports in PADD2 and the rest of the U.S. Section 4.1 has a thorough discussion about the data set that we use, but essentially PADD2 is a proxy for Cushing, Oklahoma, and the rest of the US is a proxy for the Gulf Coast. All the variables in the model are expressed in levels. The VAR is estimated with one lag, chosen according to the Lutkepohl (2005) version of Schwarz' Bayesian Information Criterion (SBIC).

Figure 3 reports the main results of the analysis in the form of dynamic impulse responses over a time period of 10 quarters. A positive shock in the storage levels in Cushing, Oklahoma widens the LLS-WTI spread (Figure 3a), possibly because the WTI oil price becomes depressed as storage builds up in Cushing. A positive shock in transportation from Cushing to the Gulf Coast narrows the LLS-WTI spread (Figure 3b), suggesting that as more oil flows from Cushing to the Gulf Coast, the price levels at the two locations becomes more equalized. On the other hand, a positive shock in the LLS-WTI spread increases the transportation level from Cushing to the Gulf Coast (Figure 3c), suggesting that the presence of higher price spread induces arbitrageurs to engage in more transportation activities. These empirical regularities are consistent with economic intuitions and will guide the construction of the theoretical model in this paper.

Figure 3: *Dynamic relationships among spread, storage, and transport*

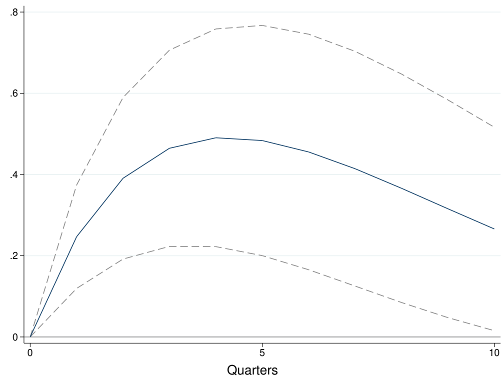

(a) *Impulse: storage level in the Midwest (PADD2), where Cushing, Oklahoma is located. Response: LLS-WTI spread.*

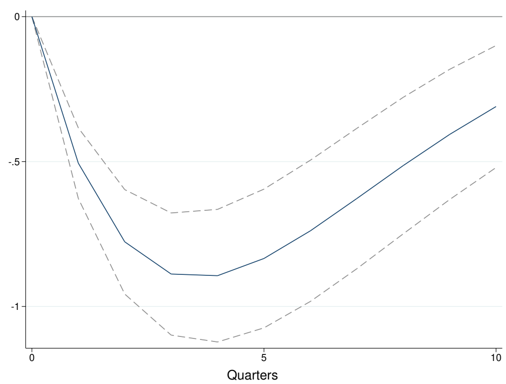

(b) *Impulse: crude oil transportation from the Midwest (PADD2) to the rest of the U.S. Response: LLS-WTI spread.*

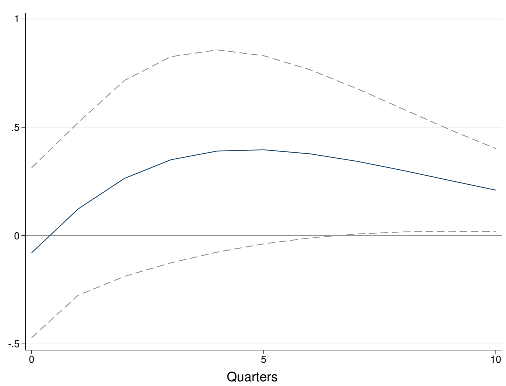

(c) *Impulse: LLS-WTI spread. Response: crude oil transportation from the Midwest (PADD2) to the rest of the U.S.*

3. A Model of Spatial Price Equilibrium

We start with a baseline model that assumes constant marginal cost of transportation and no storage. Then we twist the model by adding in increasing marginal costs of transportation and storage one at a time – two desirable features that build up our new model of spatial price equilibrium. We present several key predictions of the new model, which will be tested against real data and help us eventually identify the causes of the changing LLS-WTI spread.

3.1. The baseline model

Before building a model with increasing marginal costs of transportation and storage, let us first build intuition by describing a model where the marginal cost of transportation between A and B is constant and where there is no option to store commodity x from period t to period $t+1$. These are the assumptions in the standard spatial price equilibrium models pioneered by Samuelson (1952) and Takayama and Judge (1964).

Consider a non-perishable and non-depreciating commodity, called "x", that trades in two locations, point A and point B. Production of x at both A and B during a given period t, denoted Q_t^A and Q_t^B respectively, is assumed to be given as exogenous to arbitrageurs. This is consistent with a commodity whose production occurs in earlier periods, such as oil, or commodities that have perfectly inelastic supply curves. The corresponding price of x at each point is denoted p_t^A, p_t^B. In this model, arbitrageurs are risk-neutral and have the opportunity to transport commodity x from point A to point B at marginal cost $k_{t,AB}^T$, or from point B to point A at marginal cost $k_{t,BA}^T$.

Assumption 1. *Assume, in the baseline model, that the marginal cost of transportation is constant, such that* $k_{t,AB}^T = \bar{k}_{t,AB}^T$ *and* $k_{t,BA}^T = \bar{k}_{t,BA}^T$.

It is clear that in order to maintain no-arbitrage, the spatial price spread, given by $\sigma_t \equiv p_t^B - p_t^A$ must be bounded by the marginal costs of transportation. Explicitly the no-arbitrage condition specifies that

$$-\bar{k}^{T}_{t,BA} \leq \sigma_t \leq \bar{k}^{T}_{t,AB}. \tag{1}$$

Within these bounds, transportation will only occur from A to B if $\sigma_t = \bar{k}^{T}_{t,AB}$ and transportation will only occur from B to A if $\sigma_t = -\bar{k}^{T}_{t,BA}$. If the spread lies within these bounds then transportation between A and B will yield negative profits, and therefore transportation in either direction will not occur.

No-arbitrage condition (1) is well understood in the literature and is the baseline description of the relationship between the spatial price spread and the marginal cost of transportation. However, it says nothing about the quantity of transportation that will occur. In order to understand this we must understand the direct effect that transportation has on the spatial price spread.

We begin by describing the determinants of the absolute price of x at each point. The price of x at point A is given by economic agents' demand for consumption of x at point A. The demand curve for x at A will be denoted by $D^A(N^A_t)$, where N^A_t is the quantity of x available at A for consumption during period t. We will maintain the standard assumption for normal goods that the demand curve is downward sloping. Although the production of x at A during period t is exogenous, the quantity available for consumption is not; arbitrageurs have the option to pull x out of the market at A to transport it to B. This gives a relationship described by $N^A_t = \bar{Q}^A_t - T_t$, where T_t is the quantity of x transported from point A to point B. Arbitrageurs at point A will optimally transport commodity x while reacting to their activities' effect on the price of x at A. It follows that the equilibrium price of x at A would endogenously satisfy[3]:

$$p^A_t = D^A(\bar{Q}^A_t - T_t). \tag{2}$$

[3]Note for the sake of notational cleanness, we don't add $*$ on p^A_t and T_t to indicate equilibrium. This will be the case many times in this paper, and readers should be aware of this notational choice.

The market for x at B is different from the market at A only in that transportation increases the quantity of x available at B, whereas transportation decreases the quantity available at A. Specifically, the quantity of x available at B, denoted N_t^B, is given by the relation $N_t^B = \bar{Q}_t^B + T_t$. Furthermore, just as with the market at A, we assume the demand function is downward sloping. It then follows that the equilibrium price of x at B would satisfy:

$$p_t^B = D^B(\bar{Q}_t^B + T_t). \tag{3}$$

Therefore the equilibrium spatial price spread is determined by what we will refer to as the "spread function," denoted $\sigma_t(T_t, \bar{Q}_t^A, \bar{Q}_t^B)$:

$$\sigma_t(T_t, \bar{Q}_t^A, \bar{Q}_t^B) = D^B(\bar{Q}_t^B + T_t) - D^A(\bar{Q}_t^A - T_t). \tag{4}$$

The spread function is simply the difference in the equilibrium prices of x in the two regions, given supply levels and transportation between the regions. This representation of the spread function illuminates that increases in transportation directly decreases the spatial price spread, holding production exogenous.

We can now explicitly describe the equilibrium level of transportation between A and B. If the spatial price spread without transportation would exceed the constant marginal cost of transportation, then arbitrageurs will increase transportation until the spread decreases to $\sigma_t = \bar{k}_{t,AB}^T$. Additionally, if the spatial price spread without transportation is within the band of marginal costs of transportation between points A and B, then there will be no transportation because it would provide arbitrageurs with negative profits. We formalize this characterization through the following theorem:

Theorem 1. *Under Assumption 1 and the no storage assumption, the spatial price equilibrium is given by the following conditions:*

(i) If $\sigma_t(T_t, ...)|_{T_t=0} \geq \bar{k}_{t,AB}^T$, then $T_t^* \in \mathbb{R}^+$ is such that $D^B(\bar{Q}_t^B + T_t^*) - D^A(\bar{Q}_t^A - T_t^*) = \bar{k}_{t,AB}^T$;

(ii) If $-\bar{k}^T_{t,BA} < \sigma_t(T_t, ...)|_{T_t=0} < \bar{k}^T_{t,AB}$*, then* $T^*_t = 0$.

Figure 4 graphs the equilibrium as described by Theorem 1. Note that supply shocks in either point A or B fail to sufficiently explain the widening of the price spread in the context of the model described here: given that $\sigma_t(T_t, ...)|_{T_t=0} \geq \bar{k}^T_{t,AB}$, changes in production levels can only shift the spread function in or out, after which the quantity of transportation will simply adjust to ensure that the spatial price spread stays at $\bar{k}^T_{t,AB}$ (for an example, see Figure 5a). Therefore, in this model the only way the spatial price spread can increase beyond the initial $\bar{k}^T_{t,AB}$ is if there is an exogenous increase in the marginal cost of transportation (Figure 5b).

Figure 4: *Equilibrium with constant marginal cost of transportation*

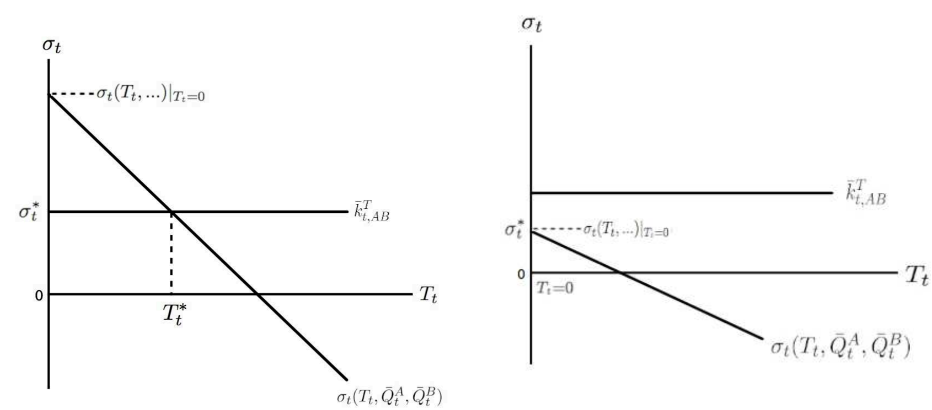

(a) *Graph of condition (i) of Theorem 1. Equilibrium quantity of transportation when* $\sigma_t(T_t, ...)|_{T_t=0} > \bar{k}^T_{t,AB}$.

(b) *Graph of condition (ii) of Theorem 1. Equilibrium quantity of transportation when* $-\bar{k}^T_{t,BA} < \sigma_t(T_t, ...)|_{T_t=0} < \bar{k}^T_{t,AB}$.

Clearly, Theorem 1 implies that the level of transportation is endogenous to the production levels at both A and B, and furthermore, if the marginal cost curve of transportation doesn't shift, the equilibrium spread will never exceed $\bar{k}^T_{t,AB}$.

3.2. Increasing marginal costs of transportation

Introducing increasing marginal costs of transportation changes the economic story and makes it possible for supply shocks at either point A or point B to influence the equilib-

Figure 5: *Illustrations of what can and cannot change the equilibrium price spread in the baseline model*

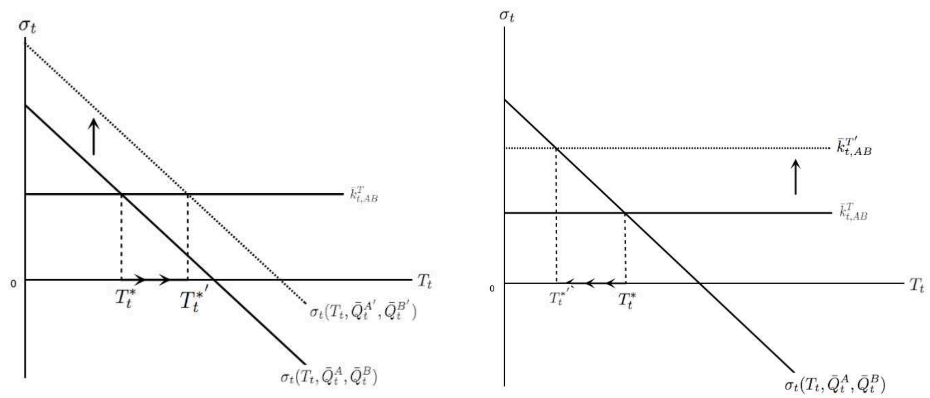

(a) *The spread function can shift out due to a change in production, but equilibrium price spread doesn't change.*

(b) *An exogenous increase in the marginal cost of transportation can increase the equilibrium price spread.*

rium spatial price spread beyond the initial marginal cost of transportation. We adapt Assumption 1 and instead assume increasing marginal cost of transportation as follows:

Assumption 2. *Assume now that the marginal cost of transportation is endogenous to and increasing in the level of transportation. Specifically, $k^T_{t,AB} = k^T_{t,AB}(T_t)$ such that $\partial k^T_{t,AB}(T_t)/\partial T_t > 0$.*

Admittedly, in the real world it is likely that the marginal cost curve of transportation is instead piecewise. There would likely be an initial flat region until the cheapest mode of transportation has reached capacity, and then a jump and another flat region at the second cheapest mode of transportation, and so on, until capacity is reached for all modes of transportation, at which point the marginal cost curve should be perfectly inelastic. However, in order to approximate the effect of multiple modes of transportation we will simply assume that the curve is increasing. This simplifying approximation makes the mathematics easier to deal with, and does not affect the key results of the paper.

The no-arbitrage condition (1), briefly restated, becomes $-k^T_{t,BA}(T_{t,BA}) \leq \sigma_t(T_t, \bar{Q}_t^A, \bar{Q}_t^B) \leq k^T_{t,AB}(T_{t,AB})$, where we will just refer to $T_{t,AB}$, the quantity of transportation from A to B, as T_t. The equilibrium conditions in Theorem 1 are no longer based on constant marginal costs of transportation, and instead are described by the following conditional statements

Theorem 2. *Under Assumption 2 and the no storage assumption, the spatial price equilibrium is given by the following conditions:*

(i) If $\sigma_t(T_t, ...)|_{T_t=0} \geq k^T_{t,AB}(0)$, then $T_t^* \in \mathbb{R}^+$ is such that $D^B(\bar{Q}_t^B + T_t^*) - D^A(\bar{Q}_t^A - T_t^*) = k^T_{t,AB}(T_t^*)$;

(ii) If $-k^T_{t,BA}(0) < \sigma_t(T_t, ...)|_{T_t=0} < k^T_{t,AB}(0)$, then $T_t^* = 0$.

Theorem 2 says that no transportation occurs if and only if the spatial price spread is less than the cost of transporting the first unit of x.

Figure 6: *Equilibrium With Increasing Marginal Costs of Transportation*

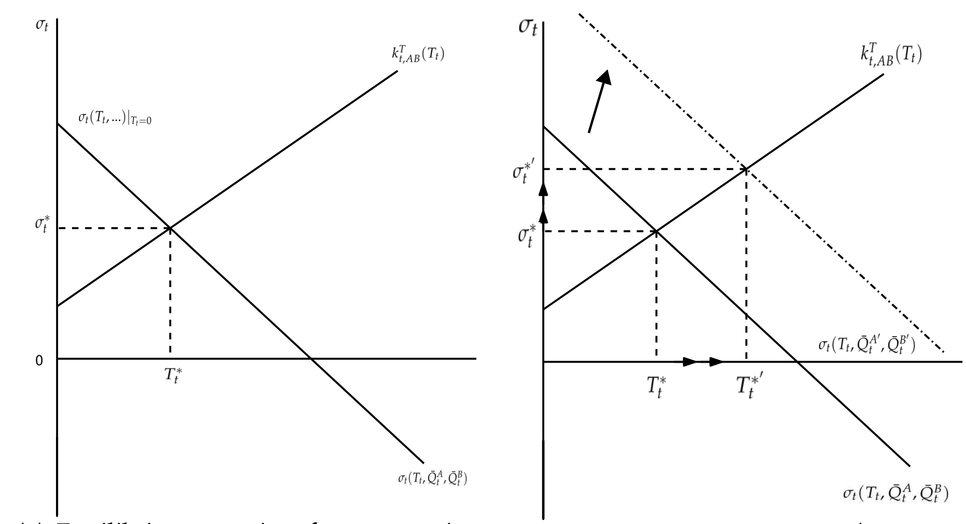

(a) *Equilibrium quantity of transportation and spatial price spread*

(b) *Effects of an increase in Q_t^A or a decrease in Q_t^B*

The key difference in the results of a model with increasing marginal costs of transportation is the relationship between the equilibrium level of transportation and spatial price spread. Specifically, as shown in Figure **??**, production shocks at both A and B can now directly push the equilibrium spatial price spread above the original marginal cost of transportation. We should note, however, that in reality the price spread and transportation may move to the new equilibrium in different rates, and due to the existence of expectations to be discussed in the next section, the price spread may very well move ahead of transportation. We will exploit this feature in Section 4.3 when we test the shape of the marginal cost curve of transportation.

3.3. Storage

The inclusion of a storage market for commodity x significantly complicates the model. The possibility of intertemporal transferring of commodity x makes the spread function sensitive to arbitrageurs' expectations of the future. Although there are a plethora of reasons why economic agents would want to store a commodity, for the sake of the analysis here, we wish to focus on the intertemporal transferring of x performed by speculators planning on taking advantage of the spatial price spread in the future.

A storage market allows a speculator to purchase a unit of x during period t and store it until period $t+1$ at the marginal cost of storage k_t^S. Assume that the marginal cost of storage is as follows:

Assumption 3. *Assume that the marginal cost of storage $k_t^S = k_t^S(S_t)$ is increasing in the storage level, such that $\partial k_t^S(S_t)/\partial S_t > 0$.*

The increasing nature of the marginal cost curve of storage is regularly described in the storage literature. The economic theory suggests that infrastructure constraints make the marginal cost of storage increase as the quantity of storage approaches capacity.

The first way the storage market enters our model is by changing the spread function. As speculators increase storage at A, they directly take units of x out of the market at A, holding all else equal. It follows that the availability of x at A, denoted N_t^A, now equals $\bar{Q}_t^A - \Delta S_t - T_t$, where the change in storage from period $t-1$ to t is denoted as $\Delta S_t = S_t - S_{t-1}$. Our new and final spread function is characterized in Proposition 1.

Proposition 1 (Spread function)**.** *The spread function is $\sigma_t = D^B(\bar{Q}_t^B + T_t) - D^A(\bar{Q}_t^A + S_{t-1} - S_t - T_t)$. Let $N_t^A = \bar{Q}_t^A + S_{t-1} - S_t - T_t$ and $N_t^B = \bar{Q}_t^B + T_t$ be the total available commodity at A and B, respectively, at time t, then $\partial\sigma_t/\partial N_t^A > 0$ and $\partial\sigma_t/\partial N_t^B < 0$.*

The statement of Proposition 1 follows directly from the standard assumption that demand curves are downward sloping. It is more important to point out that the statement is consistent with economic intuition: when there is increasing amount of commodity available at A, the price at A becomes depressed due to the abundance of the commodity there, thus enlarging the price spread; likewise for location B. Note that the variables N_t^A and N_t^B each represent the availability of a commodity at each point, and in some contexts should be exactly equal to consumption at each point. The variables $\bar{Q}_t^B$ and $\bar{Q}_t^A$, although defined as "production", can be more generally thought of as supply in a region derived by means other than drawing down storage and transporting to or from the other considered region, and could thus include net imports into the region.

Risk-neutral speculators seeking to benefit from expectations about the spread in the future will engage in storage if the net present value of the following risk-free transaction is non-negative: buying x at A at price p_t^A, storing it from period t to $t+1$ at marginal cost $k_t^S(S_t)$, transporting it to point B at the expected future marginal cost of transportation $E[k_{t+1}^T]$, and selling the unit of x at the expected future price of x at B during period $t+1$, denoted $E[p_{t+1}^B]$. The net present value of this storage transaction, denoted π_S, is given by

$$\pi_S = \tfrac{1}{1+r}E[p_{t+1}^B] - \tfrac{1}{1+r}E[k_{t+1}^T] - k_t^S(S_t) - p_t^A(\bar{Q}_t^A, S_t, T_t),$$

where r is the 1-period risk-free interest rate. The no-arbitrage hypothesis suggests that the net present value of this risk-free transaction must be non-positive.[4]

Therefore, given small enough values of the first unit of storage $k_t^S(0)$, arbitrageurs will store positive levels of x at A up until they no longer perceive that act of storage to be a positive net present value transaction. This leaves us with the following no-arbitrage condition

$$\tfrac{1}{1+r}E[p_{t+1}^B] = \tfrac{1}{1+r}E[k_{t+1}^T] + k_t^S(S_t) + D^A(\bar{Q}_t^A + S_{t-1} - S_t - T_t).$$

We can now describe general equilibrium in this model with increasing marginal costs

[4]In other words, the following condition must hold: $\frac{1}{1+r}E[p_{t+1}^B] \leq \frac{1}{1+r}E[k_{t+1}^T] + k_t^S(S_t) + D^A(\bar{Q}_t^A + S_{t-1} - S_t - T_t)$

of transportation and storage.

Theorem 3. *Under Assumptions 2 and 3, and given sufficient differentials in $\bar{Q}_t^A$ and $\bar{Q}_t^B$ such that both transportation and storage are positive, the equilibrium quantities of storage and transportation are characterized by the following two conditions:*

(i) No-arbitrage condition in the transportation market: $\sigma_t^* = D^B(\bar{Q}_t^B + T_t^*) - D^A(\bar{Q}_t^A + S_{t-1} - S_t^* - T_t^*) = k_t^T(T_t^*)$;

(ii) No-arbitrage condition in the storage market: $\frac{1}{1+r}E[p_{t+1}^B] = \frac{1}{1+r}E[k_{t+1}^T] + k_t^S(S_t^*) + D^A(\bar{Q}_t^A + S_{t-1} - S_t^* - T_t^*)$.

The no-arbitrage condition in the transportation market in Theorem 3 suggests that, holding constant the storage levels, we can identify the relationship between equilibrium spread and transportation, which should be increasing along the marginal cost curve of transportation. We formalize this idea as follows:

Proposition 2 (Shape of the marginal cost curve of transportation)**.** *Because* $\sigma_t^* = k_t^T(T_t^*)$ *and given Assumption 2, this model predicts that* $\partial\sigma_t^*/\partial T_t > 0$*, holding constant storage* S_t*.*

Again, because of the existence of expectations, transportation is likely to move to the new equilibrium slower than the spread, such that the relationship is in fact $\partial\sigma_{t-j}^*/\partial T_t > 0$ for some $j > 0$. Graphically, the inclusion of a storage market can influence the shape and magnitude of the slope of the spread function, but will not change the sign of its slope. Most importantly however, the storage market adds two exogenous expectational determinants of the spread function: expected future costs of transportation and the expected future price of x at B. This can be seen by combining the two conditions in Theorem 3:

$$\sigma_t^* = \frac{1}{1+r}E[k_{t+1}^T] + k_t^S(S_t^*) + D^B(\bar{Q}_t^B + T_t^*) - \frac{1}{1+r}E[p_{t+1}^B]. \tag{5}$$

Given equation (5), we should observe a positive relationship between expected future costs of transportation between a region and the current spatial price spread. We have not found this relationship in the literature: as expectations of future transportation costs decrease, the current spatial price spread falls. The intuition is relatively simple: if you think that future transportation costs will be low, then you will see more of a benefit from storing during period t and shipping during period $t+1$ than transporting during t; therefore you will decrease transportation during period t, which will drive down the cost of transportation during period t. A supporting evidence for this result is the announcement that Enbridge made in November 2011, when they planned on adding to the pipeline capacity out of Cushing, which should have decreased expected future costs of transportation. Indeed, this announcement drove down the LLS-WTI shortly after.

3.4. Implications for changes in spread

So far, other than the exceptational cause of the change in spread as discussed in the end of Section 3.3, we've identified three key causes for the change in the commodity price spread: (i) a positive supply shock at A, namely an increase in Q_t^A; (ii) a negative supply shock at B, namely a decrease in Q_t^B; and (iii) an upward shift in the price spread function $k_{t,AB}^T(T_t)$. Figure 6b has illustrated that both a positive supply shock at A and a negative supply shock at B will similarly increase the equilibrium price spread and transportation, but we'd like to look further into how the two supply shocks affect the equilibrium quantities of storage at A. We first state the following theorem:

Theorem 4. *Under certain regularity condition, positive supply shocks at A or B will both weakly increase the equilibrium quantities of storage at the exporting region A. In other words, $\partial S_t^*/\partial Q_t^A \geq 0$ and $\partial S_t^*/\partial Q_t^B \geq 0$.*

In order to prove Theorem 4, we first need to clarify the regularity condition as stated in the theorem. The regularity condition requires that the supply shock at B should have limited impacts on storage at A, which conforms to intuition. Formally,

Assumption 4. *Assume that Q_t^B has limited impact on S_t^*, such that*

$$\left|\frac{\partial S_t^*}{\partial Q_t^B}\right| \leq \left|\frac{\partial T_t^*/\partial Q_t^B}{\partial T_t^*/\partial Q_t^A}\right|$$

With Assumption 4, we can prove Theorem 4, as follows:

Proof. The sensitivity of storage to the exogenous productions at both points can be seen by examining equilibrium condition (i) of Theorem 3:

$$D^B(Q_t^B + T_t^*) - D^A(Q_t^A + S_{t-1} - S_t^* - T_t^*) = k_t^T(T_t^*). \tag{6}$$

For notational convenience, recall $N_t^A = Q_t^A + S_{t-1} - S_t^* - T_t^*$ and $N_t^B = Q_t^B + T_t^*$. If we totally differentiate equation (6) with respect to Q_t^B, we have

$$\frac{\partial D^B(N_t^B)}{\partial N_t^B}\left(1 + \frac{\partial T_t^*}{\partial Q_t^B}\right) - \frac{\partial D^A(N_t^A)}{\partial N_t^A}\left(-\frac{\partial S_t^*}{\partial Q_t^B} - \frac{\partial T_t^*}{\partial Q_t^B}\right) = \frac{\partial k_t^T(T_t^*)}{\partial T_t^*}\frac{\partial T_t^*}{\partial Q_t^B}. \tag{7}$$

Note that $\partial S_{t-1}/\partial Q_t^B = 0$ since S_{t-1} is realized before time t; $\partial Q_t^A/\partial Q_t^B = 0$ since production at A is exogenously given and thus there is no contemporaneous effect of a change in Q_t^B on Q_t^A. After re-arranging equation (7) we get

$$\frac{\partial S_t^*}{\partial Q_t^B} = \frac{\frac{\partial k_t^T(T_t^*)}{\partial T_t^*} - \frac{\partial D^B(N_t^B)}{\partial N_t^B} - \frac{\partial D^A(N_t^A)}{\partial N_t^A}}{\frac{\partial D^A(N_t^A)}{\partial N_t^A}}\frac{\partial T_t^*}{\partial Q_t^B} \tag{8}$$

$$\geq 0 \tag{9}$$

The inequality in (9) follows because $\frac{\partial D^B(N_t^B)}{\partial N_t^B}$ and $\frac{\partial D^A(N_t^A)}{\partial N_t^A}$ are weakly negative due to demand curve sloping downward, $\frac{\partial k_t^T(T_t^*)}{\partial T_t^*} \geq 0$ based on Assumption 2, and we already know that $\frac{\partial T_t^*}{\partial Q_t^B} \leq 0$.

On the other hand, if we totally differentiate equation (6) with respect to Q_t^A, we have

$$\frac{\partial D^B(N_t^B)}{\partial N_t^B}\frac{\partial T_t^*}{\partial Q_t^A} - \frac{\partial D^A(N_t^A)}{\partial N_t^A}\left(1 - \frac{\partial S_t^*}{\partial Q_t^A} - \frac{\partial T_t^*}{\partial Q_t^A}\right) = \frac{\partial k_t^T(T_t^*)}{\partial T_t^*}\frac{\partial T_t^*}{\partial Q_t^A}. \tag{10}$$

After re-arranging equation (10) we get

$$1 - \frac{\partial S_t^*}{\partial Q_t^A} = \frac{\frac{\partial D^B(N_t^B)}{\partial N_t^B} + \frac{\partial D^A(N_t^A)}{\partial N_t^A} - \frac{\partial k_t^T(T_t^*)}{\partial T_t^*}}{\frac{\partial D^A(N_t^A)}{\partial N_t^A}} \frac{\partial T_t^*}{\partial Q_t^A} \tag{11}$$

If we combine (8) and (11), we get

$$\left(1 - \frac{\partial S_t^*}{\partial Q_t^A}\right) \Big/ \frac{\partial T_t^*}{\partial Q_t^A} + \frac{\partial S_t^*}{\partial Q_t^B} \Big/ \frac{\partial T_t^*}{\partial Q_t^B} = 0 \tag{12}$$

Re-arranging equation (12) gives us

$$\frac{\partial S_t^*}{\partial Q_t^A} = 1 + \frac{\partial S_t^*}{\partial Q_t^B} \frac{\partial T_t^* / \partial Q_t^A}{\partial T_t^* / \partial Q_t^B} \tag{13}$$

Assumption 4 together with inequality (9) and the fact that $\partial T_t^* / \partial Q_t^B < 0$ and $\partial T_t^* / \partial Q_t^A > 0$ gives us the following relation:

$$0 \leq \frac{\partial S_t^*}{\partial Q_t^B} \leq -\frac{\partial T_t^* / \partial Q_t^B}{\partial T_t^* / \partial Q_t^A}. \tag{14}$$

Re-arranging the second inequality of (14) gives

$$\frac{\partial S_t^*}{\partial Q_t^B} \frac{\partial T_t^* / \partial Q_t^A}{\partial T_t^* / \partial Q_t^B} \geq -1 \tag{15}$$

If we plug (15) into (13), we get

$$\frac{\partial S_t^*}{\partial Q_t^A} = 1 + \frac{\partial S_t^*}{\partial Q_t^B} \frac{\partial T_t^* / \partial Q_t^A}{\partial T_t^* / \partial Q_t^B} \geq 0. \tag{16}$$

Hence the proof, given (9) and (16). □

It follows from Theorem 4 that although a positive supply shock at A or a negative supply shock at B will both similarly increase the price spread, they will have opposite effects on equilibrium quantities of storage: a positive supply shock at A should increase the equilibrium quantity of storage, and a negative supply shock at B should decrease the equilibrium quantity of storage.

As a result, the model allows us to identify different causes of a widening spread by examining distinct relationships among spread, storage, and transportation. The results are summarized in Proposition 3. We will apply this proposition in Section 5 to identify causes of the changing LLS-WTI spread.

Proposition 3 (Causes of a widening spread)**.** *This model predicts generally distinct combinations of effects on equilibrium transportation and storage, given three possible causes of the widening commodity price spread, as follows:*

(i) The first possible cause is an increase in Q_t^A, which will increase both T_t^* and S_t^*;

(ii) The second possible cause is a decrease in Q_t^B, which will increase T_t^* and decrease S_t^*;

(iii) The third possible cause is an upward shift of the $k_{t,AB}^T(T_t)$ curve, which will increase T_t^* and have an ambiguous effect on S_t^*.

4. Testing the Theory

4.1. Data

As discussed in the introduction, to be accurate with our analysis we will use the LLS-WTI spread as the relevant price spread instead of the Brent-WTI spread in order to isolate an individual transportation market.

If we are to apply our model, we must first specify what we mean by "point A" and "point B" in the context of the LLS-WTI spread. Crude oil data from the U.S. Energy Information Administration (EIA) are arguably the most comprehensive and detailed. The EIA reports production, import and export, storage, and transportation data for PADDs, which are five subregions of the United States. PADDs, or Petroleum Administration for Defense Districts, were delineated during World War II to facilitate oil allocation. Since then, they have been used to describe intra-country information on the US crude oil industry. Figure 7 shows the borders of each of the five PADD districts. PADD2 encompasses Cushing, Oklahoma, and therefore is the best proxy for point A, which in the model was the exporting region. LLS is likely a proxy for oil prices in the rest of the United States, and therefore we will use the summations of data for the other four PADDs (or the rest of the US excluding PADD2) as point B, which in the model was the importing region. For brevity, we will henceforth call the regions encompassed by all the PADDs other than PADD2 simply the "US".

Our dataset spans from 1986Q1 to 2016Q1. We compile daily spot prices of Brent, WTI, and LLS from Bloomberg, and then convert them to quarterly time series by computing the average prices within the quarter. 1986 is the first year all three spot prices are reported through Bloomberg, so our dataset covers the most extended time period. We compile the rest of our data from the U.S. Energy Information Administration. Pro_t^{PADD2} and Pro_t^{US} denote crude oil productions in PADD2 and the rest of the U.S. at time t, respectively. Sto_t^{PADD2} and Sto_t^{US} denote total commercial crude oil stocks in PADD2 and and the rest of the U.S. at time t, respectively. We should note that Sto_t^{US} only measures commercial

Figure 7: *Petroleum Administration for Defense Districts (PADD) Map*

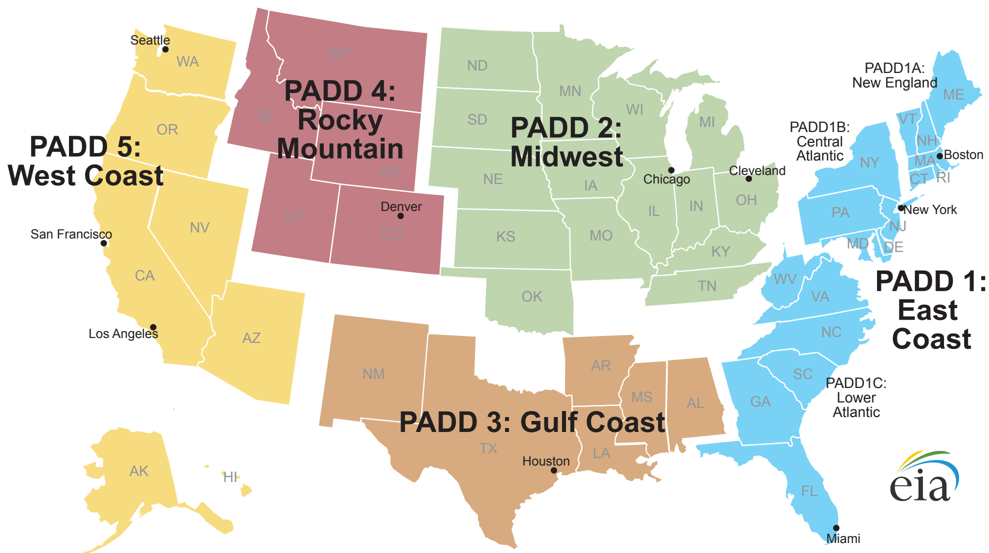

crude oil stock and thus excludes the U.S. Strategic Petroleum Reserve (SPR) maintained in PADD3. Instead, SPR_t denotes the stock of Strategic Petroleum Reserve at time t. Imp_t^{PADD2} and Exp_t^{PADD2} denote imports and exports in and out of PADD2 at time t; likewise, Imp_t^{US} and Exp_t^{US} represent the same variables for the rest of the U.S. Lastly, we also compile detailed crude oil movement data among PADDs. All raw data compiled from the EIA are monthly. We convert them into quarterly data by computing, within the quarter, the averages for the stock variables Sto_t^{PADD2}, Sto_t^{US} and SPR_t, and the sums for all the other flow variables.

For our empirical analysis, we need to construct a number of variables based on the raw data. The $Spread_t$ variable is calculated as the difference between LLS and WTI spot prices. Changes in the stock of commercial crude oil in PADD2 and the rest of the U.S. are defined as:

$$\Delta Sto_t^{PADD2} = Sto_t^{PADD2} - Sto_{t-1}^{PADD2}, \tag{17}$$

$$\Delta Sto_t^{US} = Sto_t^{US} - Sto_{t-1}^{US}. \tag{18}$$

In addition, the variable $Trans_t$, which in the model represents transportation from point A (PADD2) to point B (rest of the U.S.), is computed as the sum of all modes of transportation from PADD2 to the rest of the U.S. at time t. In addition, we define net imports for PADD2 and the rest of the U.S. as

$$NetImp_t^{PADD2} = Imp_t^{PADD2} - Exp_t^{PADD2}, \tag{19}$$

$$NetImp_t^{US} = Imp_t^{US} - Exp_t^{US} - \Delta SPR_t. \tag{20}$$

In essence, net imports are simply defined as the difference between imports and exports, but for $NetImp_t^{US}$ we also subtract out increases (or add back decreases) in the stock of Strategic Petroleum Reserve maintained in the region, since changes in the stock of SPR should not be accounted for in the commercial crude oil activities.

Table 1 lists summary statistics for key variables of the data set we compile. The table provides means of variables from 1986Q1 to 2016Q1, separated by four time periods based on observations of the LLS-WTI spread time series showed in Figure 1b. The following trends are immediately apparent from Table 1. The LLS-WTI has historically been very low from 1998 to 2005, but started to increase from 2006 to 2010; it saw a huge spike from 2011 to 2013, and tapered off from 2014 but remained quite large by historical standards. Transportation, and storage in both PADD2 and the rest of the U.S. increased consistently over time. It should be noted that net imports to PADDs other than PADD2 dropped off sharply starting from 2011, which potentially supports the hypothesis that there was a negative supply shock abroad. Contrastingly however, field productions, particularly that in PADD2, increased substantially from 2011. This leaves room for the possibility that a positive domestic supply shock contributed to the widening LLS-WTI spread. We will econometrically determine the significance of the relationships among these variables in the context of the model developed in this paper, as a two-step process: we shall first validate our theory in Sections 4.2 and 4.3, and then apply our theory to identify the causes of the changing LLS-WTI spread in Section 5.

Table 1: *Summary Statistics of Time-Series Data*

This table presents summary statistics for our quarterly crude oil data set. A number of variables are constructed from the raw data, as discussed in Section 4.1. The $Spread_t$ is constructed based on the raw LLS and WTI spot prices data from Bloomberg; all other variables come from the the U.S. Energy Information Administration (EIA) or are constructed based on the raw data from the agency. Variable superscripts, when applicable, represents the corresponding geographic region. "US" is short for all PADDs other than PADD2. All units are in millions of barrels unless otherwise noted. Data reported in this table are the means of the variables in the time period given in the column titles.

Variable	Description	1986-2005	2006-2010	2011-2013	2014-2016Q1
$Spread_t$	LLS-WTI spread (in dollars)	0.1	2.6	14.8	3.5
$Trans_t$	Transport from PADD2 to the rest of the U.S.	7.2	13.2	65.1	147.1
Sto_t^{PADD2}	Storage in PADD2	68.5	76.1	106.1	122.7
Sto_t^{US}	Storage in the rest of the U.S., excluding SPR	253.5	257.6	259.6	311.8
SPR_t	Storage of Strategic Petroleum Reserve	579.0	705.5	701.6	693.2
Pro_t^{PADD2}	Field production in PADD2	55.4	50.5	101.2	162.4
Pro_t^{US}	Field production in the rest of the U.S.	554.0	424.2	494.4	665.9
$NetImp_t^{PADD2}$	Net imports to PADD2	66.0	104.8	150.8	195.7
$NetImp_t^{US}$	Net imports to the rest of the U.S.	604.7	769.5	611.1	442.9

4.2. Testing the spread function

We will first test the relationship described by the spread function as in Proposition 1, which in our context states that the LLS-WTI spread should be increasing in the amount of available oil in Midwest (PADD2) and decreasing in the amount of available oil in the rest of the US. For testing purposes, we define the two variables denoting the availability of crude oil in PADD2 and the rest of the U.S. at time t as

$$N_t^{PADD2} = Pro_t^{PADD2} + NetImp_t^{PADD2} - Trans_t - \Delta Sto_t^{PADD2}, \tag{21}$$

$$N_t^{US} = Pro_t^{US} + NetImp_t^{US} + Trans_t - \Delta Sto_t^{US}. \tag{22}$$

We include ΔSto_t^{US} because although we assumed away storage at point B for simplicity in the theoretical section of this paper, the assumption departs from the empirics of the US oil market. The inclusion of storage at B simply changes the spread function as seen in the specification above. $NetImp_t^{PADD2}$ and $NetImp_t^{US}$ are included because the US is such a large importer of oil that field production and imports together make up what should be considered exogenous supply of crude oil available for consumption.

The econometric model for the spread function can be written as

$$Spread_t = \beta_0 + \beta_1 N_t^{PADD2} + \beta_2 N_t^{US} + \sum_q \delta_q I_t^q + \varepsilon_t, \tag{23}$$

where N_t^{PADD2} and N_t^{US} are defined by equations (21) and (22), and I_t^q is an indicator variable for quarter q. By including a set of quarter indicator variables we are controlling for the potential seasonal effects of the LLS-WTI spread.

In estimation, we are concerned about the endogeneity of the regressors N_t^{PADD2} and N_t^{US}, which if exists would render parameter estimates inconsistent. Endogeneity typically arises as a result of measurement errors, omitted variables, or simultaneity. Although simultaneity might be less of a concern if we trust the structural model built in Section 3, measurement errors and omitted variables do pose challenges to the consistency of our estimates. For example, we should note that our data on crude oil movements only includes movements reported to the EIA, and therefore may underestimate the actual crude oil movement activities across PADDs, which would cause N_t^{PADD2} and N_t^{US} to be mismeasured; furthermore, N_t^{PADD2} and N_t^{US} are likely not sufficient to fully characterize factors that determine the LLS-WTI spread, so the model as specified in (23) may have the issue of omitted variables. To overcome these potential contaminations, we employ IV to estimate the causal relationships.

To illustrate how we shall implement our IV strategy, we consider the model (23) with the omitted variables problem. The problem of the measurement error, if present, can be tackled in a similar fashion. Because of omitted variables, we interpret the error term ε_t in equation (23) as including the omitted variables. The lagged regressors are often used as potentially valid instruments, since by construction they are not correlated with the contemporaneous error term, but are correlated with the endogenous regressors if the regressors are autocorrelated. However, with time series data, the no serial correlation assumption can often be violated, and the validity of lagged regressors as instruments should be assessed with scrutiny in the presence of serial correlation. Suppose that the error term in our model follows a conventional AR(1) process. In other words, $\varepsilon_t = \rho\varepsilon_{t-1} + u_t$. We should note that the serial correlation could be induced or exacerbated by the autocorrelations of the omitted variables. One conventional strategy is to transform the model in order to correct for serial correlation. However, it is important to point out that a transformation of the model may be able to remove the serial correlation in the error term, but will not remove the contamination. It is therefore still necessary to find proper instruments for endogenous regressors in the transformed models, which can be challenging. We shall discuss the estimation challenges of the transformed models in some more detail.

There are two typical ways of transforming a model like (23). The first is

$$Spread_t = \beta_0(1-\rho) + \rho Spread_{t-1} + \beta_1 N_t^{PADD2} - \beta_1 \rho N_{t-1}^{PADD2} + \beta_2 N_t^{US} - \beta_1 \rho N_{t-1}^{US} + \sum_q \delta_q I_t^q - \sum_q \delta_q \rho I_{t-1}^q + u_t. \quad (24)$$

For this transformed model to be estimated consistently, the variables $Spread_{t-1}$, N_t^{PADD2}, N_{t-1}^{PADD2}, N_t^{US}, N_{t-1}^{US} all need to be instrumented. Such a model with many endogenous variables are difficult to identify, and having correspondingly too many instruments can even be dangerous for inference (Roodman, 2009). In addition, given the transformed model, we need to impose constraints on the relationships among coefficients. For instance, the product of the parameter estimates on $Spread_{t-1}$ and N_t^{PADD2} plus the parameter estimates on N_{t-1}^{PADD2} needs to be zero by construction. These constraints on parameters will further undermine the identification of the model.

Another approach is to write out the transformed model in quasi-differenced form:

$$\widetilde{Spread}_t = \beta_0(1-\rho) + \beta_1 \tilde{N}_t^{PADD2} + \beta_2 \tilde{N}_t^{US} + \sum_q \delta_q \tilde{I}_t^q + u_t, \quad (25)$$

where $\tilde{x}_t = x_t - \rho x_{t-1}$, with x_t representing any variable in general. In this transformed model, an instrument that itself is in a quasi-differenced form would only work if the original error ε_t is uncorrelated with the instrument at times t, $t-1$, and $t+1$. This rules out first lagged regressors at IVs. The second lagged regressors may work as instruments, only if the quasi-differenced regressors and the second lagged quasi-differenced regressors still have correlations strong enough, because otherwise we would run into weak instrument problems. However, quasi-differencing the regressors typically take out most of the autocorrelation in the transformed regressors by construction, and it would be quite rare to see the quasi-differenced regressors having strong autocorrelations at the second order. In sum, it would be very difficult to justify lagged regressors as proper instruments in this type transformed models.

As a result, we shall not conduct our IV estimation through the transformed models. Instead, we resort to IV estimation technique that is robust to serial correlation and that does not require transforming the model. Building upon the works of heteroskedasticity and autocorrelation (HAC) consistent estimation, such as Newey and West (1987), Andrews (1991) and Smith (2005), we can correct for serial correlation by generating an estimate for the covariance matrix using the Bartlett kernel function and an appropriate selection of bandwidth. Baum, Schaffer and Stillman (2003, 2007) have detailed discussions on how an HAC consistent IV estimation should be implemented.

In conclusion, our final model for testing the spread function remains in its original form, as in equation (23):

$$Spread_t = \beta_0 + \beta_1 N_t^{PADD2} + \beta_2 N_t^{US} + \sum_q \delta_q I_t^q + \varepsilon_t,$$

but in light of the concerns of the endogeneity of N_t^{PADD2} and N_t^{US}, we use the method of instrumental variables. Potentially valid instruments include N_{t-j}^{PADD2} and N_{t-j}^{US}, where $j = 1, 2, 3, ...$ In practice, we conduct a series of tests to select an optimal set of valid instruments. Specifically, we start with a set of eight potentially valid instruments, namely N_{t-j}^{PADD2} and N_{t-j}^{US}, where $j = 1, 2, 3, 4$. We test the these instruments one at a time for redundancy. After dropping all the redundant instruments, we also test the set of remaining instruments to avoid possibility of weak identification or underidentification. In the end, we conclude that N_t^{PADD2} and N_t^{US} should be optimally instrumented by N_{t-1}^{PADD2}, N_{t-3}^{PADD2}, N_{t-1}^{US}, and N_{t-2}^{US}.

After we come up with the optimal set of instruments, we also conduct endogeneity tests of regressors N_t^{PADD2} and N_t^{US}. The endogeneity tests suggest that, statistically, N_t^{PADD2} and N_t^{US} can be treated as exogenous. As a result, in addition to running regressions where N_t^{PADD2} and N_t^{US} are instrumented, we also run additional regressions where one or neither of N_t^{PADD2} and N_t^{US} is instrumented.

To address concerns of heteroskedasticity and serial correlation, we run all the aforementioned tests in two versions, one with heteroskedasticity and serial correlation robust standard errors and one without. In all regressions that are not HAC robust, homoskedasticity and serial independence are checked to be violated. As a result, all our reported results are HAC robust. In doing the HAC robust estimations, we use the Bartlett kernel function with a bandwidth of 5 that is optimally chosen based on the number of observations in our data set.

The regression results of key specifications are presented in Table 2.

The table reports three types of regressions: OLS, 2S GMM, and LIML. OLS regressions are run when neither N_t^{PADD2} nor N_t^{US} is instrumented. When any of the regressor is instrumented, we use two-step feasible and efficient GMM (2S GMM) estimator instead of 2SLS because GMM is more efficient when heteroskedasticity is present (Baum, Schaffer and Stillman, 2003). We also use the limited information maximum likelihood (LIML) estimator, first derived by Anderson and Rubin (1949), to replicate all regression specifications under the 2S GMM estimator. This is because even though LIML provides no asymptotic efficiency gains over 2S GMM, recent research suggests that their finite-sample performance may be superior, for example in the presence of weak instruments (Hahn, Hausman and Kuersteiner, 2004).

The regression results indicate that the estimates on N_t^{PADD2} and N_t^{US} are largely consistent across specifications, regardless of whether regressors are instrumented or quarter dummies are included, and across estimators used. In all the regressions where instrumental variables are used, our selected set of instruments are shown to be valid with no concerns of underidentification, weak identification, or overidentification, as can be seen from the test statistics for instrument validity reported in the last three lines of the table. In particular, the null hypothesis test for the underidentification test is that the model is underidentified; the weak identification test reports an F statistic that can be compared to critical values compiled by Stock and Yogo (2005), but the rule of thumb is that the F statistic needs to be greater than 10 not to have concerns about weak identification; the null hypothesis for the overidentification test is that the instruments are valid instruments, i.e.,

Table 2: *Testing the Spread Function, Full Sample*

This table presents IV regression results for the testing the spread function on the full sample from 1986Q1 to 2016Q1. The dataset consists of 121 quarterly observations. All regressions reported in this table are heteroskedasticity and serial correlation robust, using Bartlett kernel function with a bandwidth of 5. Depending on the specification, N_t^{PADD2} and N_t^{US} are only instrumented in certain cases. When they are instrumented, the corresponding instrumental variables for N_t^{PADD2} are N_{t-1}^{PADD2} and N_{t-3}^{PADD2}; the corresponding instrumental variables for N_t^{US} are N_{t-1}^{US} and N_{t-2}^{US}. Section 4.2 has full discussions on the test procedures. Asterisks indicate statistical significance at 1%***, 5%**, and 10%* levels.

	Dependent Variable: LLS-WTI Spread, $Spread_t$										
	OLS		2S GMM						LIML		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Coefficient estimates for regressor N_t^{PADD2}:											
Non-instrumented	0.1041***	0.1033***		0.1041***			0.1077***			0.1040***	
	(0.0365)	(0.0365)		(0.0362)			(0.0359)			(0.0367)	
Instrumented			0.1099***		0.1112***	0.1138***		0.1099***	0.1107***		0.1113***
			(0.0365)		(0.0377)	(0.0374)		(0.0380)	(0.0380)		(0.0377)
Coefficient estimates for regressor N_t^{US}:											
Non-instrumented	−0.0128**	−0.0123**			−0.0155**			−0.0144**			−0.0155**
	(0.0059)	(0.0052)			(0.0066)			(0.0058)			(0.0066)
Instrumented			−0.0147**	−0.0134**		−0.0166**	−0.0155**		−0.0147**	−0.0134**	
			(0.0067)	(0.0063)		(0.0071)	(0.0071)		(0.0067)	(0.0065)	
Quarter dummies?	yes	no	yes	yes	yes	no	no	no	yes	yes	yes
R^2	0.44	0.44	0.44	0.45	0.44	0.44	0.44	0.44	0.44	0.45	0.44
Adj R^2	0.54	0.54	0.55	0.55	0.55	0.54	0.54	0.54	0.55	0.55	0.55
Test statistics for instruments validity (p-values for underidentification and overidentification tests, F-statistic for weak identification test)											
Underidentification	–	–	0.00	0.00	0.02	0.00	0.00	0.02	0.00	0.00	0.02
Weak identification	–	–	276	766	440	162	233	404	276	766	440
Overidentification	–	–	1.00	1.00	0.94	0.85	0.61	0.70	1.00	1.00	0.94

uncorrelated with the error term, and that the excluded instruments are correctly excluded from the estimated equation.

Column (3) reports results obtained from our preferred specification estimated using 2S GMM. In this specification, both N_t^{PADD2} and N_t^{US} are instrumented with the selected set of instruments, and quarter dummies are included to control for seasonal effects. Based on the results of this regression, a ten-million barrel per quarter increase in total available crude oil in PADD2 is estimated to cause a $1.099 increase in LLS-WTI spread, and the estimate is statistically significant at the 1% level; a ten-million barrel per quarter increase in total available crude oil in the rest of the U.S. is estimated to cause a $0.147 decrease in LLS-WTI spread, and the estimate is statistically significant at the 5% level.

For comparison purposes, we take the means of the estimates across all columns (1) to (11). The means of the estimates suggest the following: a ten-million barrel per quarter increase in total available crude oil in PADD2 is estimated to cause a $1.082 increase in LLS-WTI spread, and a ten-million barrel per quarter increase in total available crude oil in the rest of the U.S. is estimated to cause a $0.144 decrease in LLS-WTI spread. These results are very close to those from our preferred regression in column (3).

Because the LLS and WTI spot prices only started to show noticeable divergence from 2006, we also repeat our regressions on the subset that covers 2006Q1-2016Q1. All the regression specifications are kept the same as in the full sample regressions, in order to make sure that the results are comparable. In other words, the set of instruments is not re-optimized for the subsample. Table 3 presents the regression results based on the subsample. Because we do not re-optimize the set of instruments, the first point worth discussing is the tests on instrument validity, reported in the last three lines in the table. As we can see, the test statistics indicate that the instruments are valid in various dimensions, therefore we put away concerns about instruments validity even though they are not re-optimized specifically for the subsample.

Our preferred specification in column (3) reports that for the time period from 2006Q1 to 2016Q1, when LLS-WTI spread departed from historical levels, a ten-million barrel per quarter increase in total available crude oil in PADD2 is estimated to cause a \$0.704 increase in LLS-WTI spread, and a ten-million barrel per quarter increase in total available crude oil in the rest of the U.S. is estimated to cause a \$0.996 decrease in LLS-WTI spread. These two estimates are statistically significant at 5% and 1% levels, respectively.

In comparison, the means for all estimates suggest the following: a ten-million barrel per quarter increase in total available crude oil in PADD2 is estimated to cause a \$0.786 increase in LLS-WTI spread, and a ten-million barrel per quarter increase in total available crude oil in the rest of the U.S. is estimated to cause a \$0.801 decrease in LLS-WTI spread. Again, just like the full sample, these results are not far away from those of our preferred specification.

To sum up, in this section we test the spread function as in Proposition 1, and the results are consistent with the predictions of our theoretical model: an increase in the amount of available oil in the Midwest (PADD2) and in the rest of the U.S. enlarges and narrows, respectively, the LLS-WTI spread. This holds true regardless whether we use the full sample that covers three decades, or the subsample that covers the past decade when the LLS-WTI spread became elevated.

4.3. Testing the marginal cost curve of transportation

Proposition 2 predicts that the short-run equilibrium relation between spread and transport is captured by a strictly increasing marginal cost curve of transportation $k_t^T(\cdot)$, controlling for storage. We should note that Section 3.3 argues that transportation is likely to move to the new equilibrium slower than the spread. From Figure 3c we can see that if we were to allow spread to move ahead of transportation, then the optimal lag choice for spread is 4 quarters. Therefore, we can directly test Proposition 2 using the following model:

$$Trans_t = \beta_0 + \beta_1 Spread_{t-4} + \sum_q \delta_q I_t^q + \gamma_1 Sto_t^2 + \gamma_2 Sto_t^3 + \varepsilon_t. \tag{26}$$

In order to estimate model (26), we are concerned about the endogeneity of $Trans_t$

Table 3: *Testing the Spread Function, Subsample*

This table presents IV regression results for the testing the spread function on the subsample from 2006Q1 to 2016Q1. The dataset consists of 41 quarterly observations. The specifications are exactly the same as corresponding specifications in Table 2. All regressions reported in this table are heteroskedasticity and serial correlation robust, using Bartlett kernel function with a bandwidth of 4. Depending on the specification, N_t^{PADD2} and N_t^{US} are only instrumented in certain cases. When they are instrumented, the corresponding instrumental variables for N_t^{PADD2} are N_{t-1}^{PADD2} and N_{t-3}^{PADD2}; the corresponding instrumental variables for N_t^{US} are N_{t-1}^{US} and N_{t-2}^{US}. Asterisks indicate statistical significance at 1%***, 5%**, and 10%* levels.

	Dependent Variable: LLS-WTI Spread, $Spread_t$										
	OLS		2S GMM						LIML		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Coefficient estimates for regressor N_t^{PADD2}:											
Non-instrumented	0.0782**	0.0772**		0.0587**			0.0772**			0.0804**	
	(0.0353)	(0.0379)		(0.0300)			(0.0374)			(0.0358)	
Instrumented			0.0704**		0.0914**	0.775*		0.0756*	0.0896**		0.0885**
			(0.0296)		(0.0367)	(0.0398)		(0.0410)	(0.0384)		(0.0385)
Coefficient estimates for regressor N_t^{US}:											
Non-instrumented	−0.0805***	−0.0651***			−0.0820***			−0.0606***			−0.0812***
	(0.0309)	(0.0227)			(0.0307)			(0.0215)			(0.0309)
Instrumented			−0.0996***	−0.0978***		−0.0582***	−0.0652***		−0.0953***	−0.0951***	
			(0.0359)	(0.0365)		(0.0195)	(0.0234)		(0.0364)	(0.0368)	
Quarter dummies?	yes	no	yes	yes	yes	no	no	no	yes	yes	yes
R^2	0.40	0.34	0.38	0.37	0.39	0.31	0.31	0.34	0.38	0.39	0.39
Adj R^2	0.70	0.67	0.69	0.69	0.70	0.66	0.65	0.67	0.69	0.69	0.70
Test statistics for instruments validity (p-values for underidentification and overidentification tests, F-statistic for weak identification test)											
Underidentification	–	–	0.07	0.03	0.02	0.08	0.04	0.03	0.07	0.03	0.02
Weak identification	–	–	38	65	183	59	40	157	38	65	183
Overidentification	–	–	0.53	0.27	0.81	0.31	0.15	0.52	0.53	0.27	0.81

similar to when we test for the spread function in Section 4.2. As a result, we apply a similar instrument variable strategy that is heteroskedasticity and serial correlation robust. Using the same test procedures as in Section 4.2, we find that the optimal instrument for $Spread_{t-4}$ is $Spread_{t-5}$. Table 4 reports regression results based on both the full sample that covers 1986Q1-2016Q1 and the subsample that covers 2006Q1-2016Q1. Several observations are immediately clear from the table. First, including quarter dummies or not does not significantly change the coefficient estimate on $Spread_{t-4}$, but it does make a major difference whether or not we include the control variables Sto_t^2 and Sto_t^3. In particular, the R^2 and Adj R^2 drop substantially if we leave out the control variables. This is an indication that the control variables should be included, as consistent with the prediction of Proposition 2. Such an observation provides further support for our theoretical model. Second, the test statistics for instrument validity reported in the last two lines of the table suggest that the instrument that we optimally choose is indeed valid in the regressions. Notice that there is no overidentification test here, because there is only one instrument so the model is exactly identified.

Both full sample and subsample results confirm that equilibrium price spread and transportation have a strictly increasing relationship, corroborating Proposition 2. In particular, the full sample result indicates that a \$1 increase in $Spread_{t-4}$ would cause $Trans_t$ to increase by 2.98 million barrels, according to our preferred specification (1) that includes both quarter dummies and control variables. On the other hand, the subsample result indicates that a \$1 increase in $Spread_{t-4}$ would cause $Trans_t$ to increase by 3.92 million barrels, according to our preferred specification (3) that includes both quarter dummies and control variables. These translate into an elasticity of transportation to price spread of 0.27 for the full sample and 0.43 for the subsample[5], suggesting that starting from 2006, the LLS-WTI spread not only becomes elevated, but its increase also is able to induce a larger amount of transportation activities from the Midwest (PADD2) to the rest of the U.S. compared to earlier decades.

[5]The means for transportation are 24.35 and 57.77 million barrels for the full sample and subsample, respectively. The means for price spread are \$2.24 and \$6.38 for the full sample and subsample, respectively.

Table 4: *Testing the Marginal Cost Curve of Transportation*

This table presents IV regression results for the testing the marginal cost curve of transportation. The full sample consists of 121 quarterly observations; the subsample consists of 41 quarterly observations. All regressions are heteroskedasticity and serial correlation robust with the Bartlett kernel function. The instrumental variable for $Spread_{t-4}$ is optimally chosen to be $Spread_{t-5}$. Section 4.3 has full discussions on the test procedures. Asterisks indicate statistical significance at 1%***, 5%**, and 10%* levels.

Dependent Variable: Transport from PADD2 to the rest of the U.S., $Trans_t$								
	Full Sample, 1986Q1-2016Q1				Subsample, 2006Q1-2016Q1			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Spread_{t-4}$ (instrumented)	2.98*	2.93*	6.11***	6.11***	3.92***	3.78***	4.72***	4.73***
	(1.59)	(1.59)	(1.34)	(1.33)	(1.20)	(1.10)	(1.42)	(1.42)
Control variables:								
Sto_t^{PADD2}	1.01**	1.05**			0.25	0.38		
	(0.49)	(0.50)			(0.43)	(0.38)		
Sto_t^{US}	0.40	0.34			1.43***	1.26***		
	(0.31)	(0.30)			(0.43)	(0.35)		
Quarter dummies?	yes	no	yes	no	yes	no	yes	no
R^2	0.75	0.74	0.36	0.36	0.78	0.75	0.14	0.13
Adj R^2	0.82	0.81	0.53	0.53	0.89	0.88	0.58	0.58
Test statistics for instruments validity (p-value for underidentification test, F-statistic for weak identification test)								
Underidentification	0.03	0.03	0.08	0.08	0.05	0.05	0.06	0.06
Weak identification	135	134	188	189	96	108	168	171

5. Identifying Causes of the Changing LLS-WTI Spread

Now that we've validated our model in Section 4, we would like to apply our model to identify causes of the changing LLS-WTI over time, particularly over the past decade. Proposition 3 states three possible causes of a changing commodity price spread. In our context, each of the three causes represents the following: an change in Q_t^A represents a production shock in the Midwest; a change in Q_t^B represents a production shock abroad; a shift of the $k_{t,AB}^T(T_t)$ represents a structural change in the marginal cost of transportation. These three causes imply different relationships among equilibrium spread, transportation, and storage. Therefore, if spread and transportation are negatively correlated, i.e. $\rho_{\sigma_t,T_t} \leq 0$, then the changing price spread should be attributed to a structural change in the marginal cost of transportation. On the other hand, if spread and transportation are positively correlated, i.e. $\rho_{\sigma_t,T_t} \geq 0$, then there are two further possibilities: if spread and storage are positively correlated, i.e. $\rho_{\sigma_t,S_t} \geq 0$, then the changing spread should be attributed to a production shock in the Midwest; if spread and storage are negatively correlated, i.e. $\rho_{\sigma_t,S_t} \leq 0$, then the changing spread should be attributed to a production shock abroad. I summarize these results in Table 5, assuming a widening price spread. The results for a narrowing price spread are similar.

Given these model implications, we can use correlations to determine if the changing spread was due to a production shock abroad, a production shock in the Midwest, or a structural change in the the marginal cost of transportation. We compute rolling correlations between transportation and the LLS-WTI spread as well as between storage and the LLS-WTI spread. The rolling correlations are computed using a 28-quarter rolling window. The rolling window is chosen in a balanced way, so that it is long enough to uncover patterns in the data and make correlations valid, but also short enough to capture any instability in the rolling correlation trends. We should note that making reasonable changes to the width of the rolling window does not have an impact on our conclusions.

In order to make inferences based on the rolling correlations, we'd like to bound the

Table 5: *Implied relationships among spread, transportation, and storage*

This table relates model predictions to the context of the LLS-WTI price spread. Proposition 3 discusses three causes of the changing price spread. Each of the three causes represents different shocks in the U.S. crude oil market. The last two columns summarize the implied correlations among spread, transportation, and storage based from different causes. Section 5 has a full discussion on these implied correlations.

Cause of widening spread	Model interpretation	Implied correlation between...	
		Spread and transportation	Spread and storage
Positive production shock in Midwest	Increase in Q_t^A	$\rho_{\sigma_t,T_t} \geq 0$	$\rho_{\sigma_t,S_t} \geq 0$
Negative production shock abroad	Decrease in Q_t^B	$\rho_{\sigma_t,T_t} \geq 0$	$\rho_{\sigma_t,S_t} \leq 0$
Structural increase in MC of transportation	Upward shift of $k_{t,AB}^T(T_t)$ curve	$\rho_{\sigma_t,T_t} \leq 0$	ambiguous

correlations by confidence intervals. Unlike confidence intervals around means, confidence intervals around Pearson r's are not symmetrical. We shall very briefly review how to construct a confidence interval on a correlation coefficient estimate ρ. The Fisher's r-to-z transformation gives an expectation of the z statistic of

$$z' = \frac{1}{2} \ln \left(\frac{1+\rho}{1-\rho} \right), \tag{27}$$

with a standard deviation of $\sigma_{z'} = \sqrt{1/(n-3)}$, where n is the sample size. Then a confidence interval in the z-space is simply

$$z' \pm (\text{Z score}) \times \sigma_{z'}, \tag{28}$$

where the Z score for a 95% confidence interval, for example, is 1.96. Denote the upper and lower bounds of the confidence interval in the z-space as z^+ and z^-. Then the upper and lower bounds of the confidence interval, ρ^+ and ρ^-, in the r-space are calculated as

$$\rho^+ = \frac{e^(2z^+ - 1)}{e^(2z^+ + 1)} \tag{29}$$

$$\rho^- = \frac{e^(2z^- - 1)}{e^(2z^- + 1)} \tag{30}$$

Figure 8 shows the rolling correlations between LLS-WTI spread and transport, and between LLS-WTI spread and storage in the Midwest (PADD2), with 95% confidence intervals, starting from 2005. As we can see, before 2013 the rolling correlations between LLS-WTI spread and transport, and between LLS-WTI spread and storage in the Midwest (PADD2), are both positive, suggesting that the increase in price spread during the period should be attributed to positive production shocks in the Midwest. From 2013 and onward, we can see from Figure 8 that both correlations start to turn negative. This is evidence that the dominating cause of the narrowing of the price spread is a structural decrease in the marginal cost of transportation, although a positive production shock abroad could also possibly be a secondary cause that contributes to the negative correlations between spread and storage.

Figure 8: *Rolling correlations among spread, transportation, and storage*

These figures present rolling correlations between the LLS-WTI price spread and transportation from PADD2 to the rest of the U.S., and between the LLS-WTI price spread and storage in PADD2, from 2005. The rolling regression window is chosen to be 28 quarters. The shaded areas indicate 95% confidence intervals.

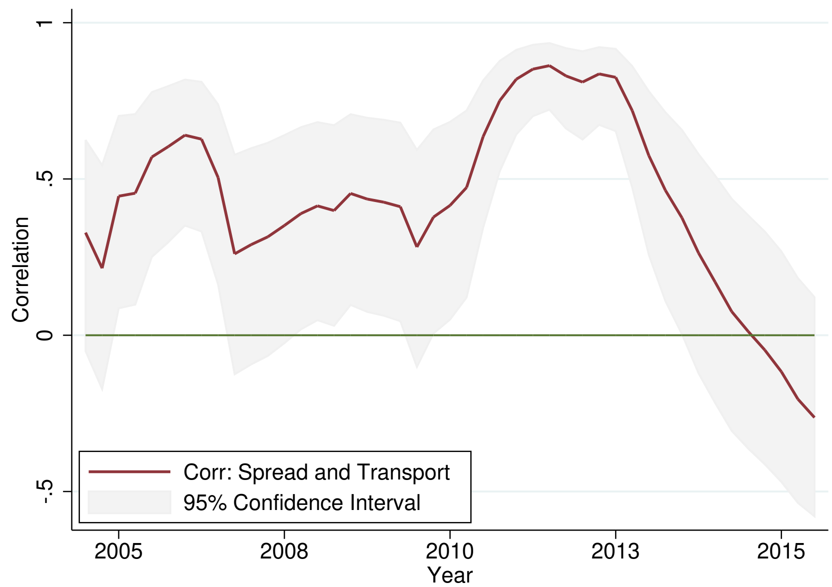

(a) *Rolling correlations between spread and transportation*

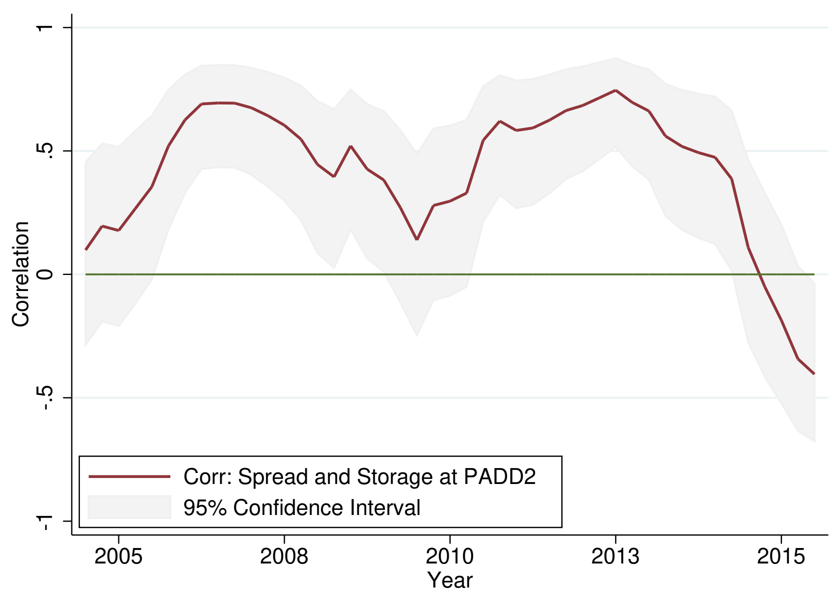

(b) *Rolling correlations between spread and storage*

In sum, we have been able to depict a complete story for the dramatic changes in LLS-WTI price spread over the past decade. From 2006 to 2013, LLS-WTI increased from being almost zero to as high as $25 per barrel, due to a positive production shock in the Midwest. Starting from 2013, the LLS-WTI spread started to narrow, fluctuating around $3.5 per barrel, which was much lower than the peak but remained quite high by historical standards. The primary cause for the narrowing of the spread during the period was a structural decrease in the marginal cost of transportation out of Midwest to the rest of the U.S., possibly due to such factors as new pipeline capacities. A positive production shock abroad, such as the Middle East, was likely a secondary cause for the narrowing of the price spread in recent years, too.

6. Conclusion

Dramatic price spreads among benchmark crude oils is a phenomenon that emerged in the commodities market in the past decade. While this phenomenon has generated a great deal of media attention, there has not been rigorous studies that attempt to identify the causes of the changing price spreads.

The addition that this paper provides to the economic literature is both theoretical and empirical. On the theoretical side, the framework of analysis presented here is a generalized and much clearer version of the standard spatial price equilibrium model first pioneered by Samuelson (1952). We explicitly define a spread function and consider the equilibrium level of transportation as the intersection between the spread function and the marginal cost curve of transportation. This perspective makes clear the key differences between a model with constant marginal costs of transportation versus one with increasing marginal costs of transportation. The addition of storage further allows us to draw predictions from the model that help identify unique causes of the changing price spread.

On the empirical side, we construct a dataset that covers an extended time period of three decades. This comprehensive dataset allows us to uncover patterns of the crude oil market over a long time horizon, which in a way ensures the robustness of our results. We econometrically validate our model using several testable model predictions, and identify the causes of the changing price spread over the past decade by exploiting the relationships among crude oil price spreads, transportation across regions, and crude oil storage levels.

In sum, this paper provides a necessary generalization of the literature that addresses the interconnectedness of spatially separated commodity markets. In fact, the applicability of the generalized model presented here is likely not limited to the LLS-WTI spread, and further research should be performed on the applicability of this model to other spatial price spreads that exhibited similarly abrupt changes.

References

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *Annals of Mathematical Statistics*, 20(1): 46–63.

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3): 817–858.

Ardeni, Pier Giorgio. 1989. "Does the Law of One Price Really Hold for Commodity Prices?" *American Journal of Agricultural Economics*, 71(3): 661–669.

Baum, Christopher F., Mark E. Schaffer, and Steven Stillman. 2003. "Instrumental Variables and GMM: Estimation and Testing." *Stata Journal*, 3(1): 1–31.

Baum, Christopher F., Mark E. Schaffer, and Steven Stillman. 2007. "Enhanced Routines for Instrumental Variables/Generalized Method of Moments Estimation and Testing." *Stata Journal*, 7(4): 465–506.

Buyuksahin, Bahattin, Thomas K. Lee, James T. Moser, and Michel A. Robe. 2013. "Physical Markets, Paper Markets and the WTI-Brent Spread." *Energy Journal*, 34(3): 129–151.

Coleman, Andrew. 2009. "Storage, Slow Transport, and the Law of One Price: Theory with Evidence from Nineteenth-Century U.S. Corn Markets." *Review of Economics and Statistics*, 92(2): 332–350.

Deaton, Angus, and Guy Laroque. 1992. "On the Behaviour of Commodity Prices." *Review of Economic Studies*, 59(1): 1–23.

Deaton, Angus, and Guy Laroque. 1996. "Competitve Storage and Commodity Price Dynamics." *Journal of Political Economy*, 104(5): 896–923.

Fattouh, Bassam. 2010. "The Dynamics of Crude Oil Price Differentials." *Energy Economics*, 32(2): 334–342.

Goldberg, Pinelopi K., and Frank Verboven. 2005. "Market Integration and Convergence to the Law of One Price: Evidence from the European Car Market." *Journal of International Economics*, 65(1): 283–303.

Hahn, Jinyong, Jerry Hausman, and Guido Kuersteiner. 2004. "Estimation with Weak Instruments: Accuracy of Higher-Order Bias and MSE Approximations." *Econometrics Journal*, 7(1): 46–63.

Knittel, Christopher R., and Robert S. Pindyck. 2016. "The Simple Economics of Commodity Price Speculation." *American Economic Journal: Macroeconomics*, 8(2): 85–110.

Lutkepohl, Helmut. 2005. *New Introduction to Multiple Time Series Analysis.* Springer.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Olsen, Kyle K., James W. Mjelde, and David A. Bessler. 2015. "Price Formulation and the Law of One Price in Internationally Linked Markets: An Examination of the Natural Gas Markets in the USA and Canada." *Annals of Regional Science*, 54(1): 117–142.

Richardson, J. David. 1978. "Some Empirical Evidence on Commodity Arbitrage and the Law of One Price." *Journal of International Economics*, 8(2): 341–351.

Roodman, David. 2009. "A Note on the Theme of Too Many Instruments." *Oxford Bulletin of Economics and Statistics*, 71(1): 135–158.

Samuelson, Paul A. 1952. "Spatial Price Equilibrium and Linear Programming." *American Economic Review*, 42(3): 283–303.

Smith, Richard J. 2005. "Automatic Positive Semidefinite HAC Covariance Matrix and GMM Estimation." *Econometric Theory*, 21(1): 158–170.

Stock, James H., and Motohiro Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*. , ed. Donald W. K. Andrews and James H. Stock, Chapter 5. Cambridge University Press.

Takayama, Takashi, and George G. Judge. 1964. "Equilibrium among Spatially Separated Markets: A Reformulation." *Econometrica*, 32(4): 283–303.

Werner, Jan. 1987. "Arbitrage and the Existence of Competitive Equilibrium." *Econometrica*, 55(6): 1403–1418.

Williams, Jeffrey C., and Brian Wright. 1991. *Storage and Commodity Markets.* Cambridge University Press.