

Threshold cointegration and spatial price transmission when expectations matter

Sergio H. Lence^{a,*}, GianCarlo Moschini^b, Fabio Gaetano Santeramo^c

^a*Department of Economics, Iowa State University, Ames, Iowa, USA*

^b*Department of Economics and CARD, Iowa State University, Ames, Iowa, USA*

^c*University of Foggia, Foggia, Italy*

Received 9 March 2016; received in revised form 27 June 2017; accepted 14 July 2017

Abstract

We examine the performance of the threshold cointegration approach, specifically Band-TVECM, to price transmission analysis in an explicit context where trade decisions are made based on expectation of final prices, because trade takes time. We find that, following a standard inference strategy, a large portion of three-regime cases are not identified as such. Results show that transfer costs are systematically underestimated, particularly in three-regime models. The speed of price transmission is also biased in three-regime models. Furthermore, inferences about occurrence of trade are poor, with estimated models suggesting far lower market integration than is true in the data-generating process.

JEL classifications: C32, F17, Q11

Keywords: Band-TVECM; Market integration; Price transmission; Threshold cointegration; Transfer costs

1. Introduction

Prices reflect competing equilibrating forces at work in a market setting, condense complex and decentralized information, and provide vital signals for the allocation of resources by economic agents. Much can be learned from price data, but, because goods and services are heterogeneous with respect to form, space, and time, price analysis can be challenging. In addition to sophisticated econometric techniques, accurate inference may require a clear understanding of industry, and awareness of relevant policies and institutions. The interaction of such factors in the analyst's model is just as important. Efforts to investigate the scope and applicability of the expanding economist's toolkit are essential, as methods and techniques appropriate in one setting might prove inadequate when seemingly minor conditions or assumptions are changed. In this article, we revisit a popular approach to the analysis of price transmission, the threshold autoregressive model, and assess its performance in an explicit spatial equilibrium context.

Spatial links between prices have long been of interest in agricultural economics. Earlier work is summarized and discussed in Fackler and Goodwin (2001). A central concept is

the so-called Law of One Price (LOP), whereby prices in two markets linked by trade tend to be equalized by arbitrage—apart, of course, for the cost of carrying out such arbitrage, which includes transport and other transaction costs (collectively referred to as “transfer costs” hereafter). Important questions regarding market integration, market efficiency, and the effectiveness of policies can relate to how exactly prices are transmitted between markets. Wright and Williams (1989) and Coleman (2009b) highlight the role of storage activities and inventory management in determining the spatial constellation of prices. Shiue (2002) also studies physical arbitrage, with an emphasis on the connection between trade and inter-temporal arbitrage within the framework of spatial price transmission. Coleman (2009a) and Stephens et al. (2011) explore the importance of information on trade capacity constraints and on trade flows in analyses of spatial price transmission. Myers and Jayne (2012) identify the differing impacts on price transmission of private sector trade and government behavior.

The concept of cointegration has played a central role, at least since Ardeni (1989), in the analysis of price transmission (e.g., García-Germán et al., 2016). In particular, following the seminal papers by Balke and Fomby (1997) and Hansen and Seo (2002), many LOP econometric studies have adopted “threshold cointegration” as the benchmark. The essential feature of this approach is to allow for regime switching in the

*Corresponding author. Tel.: +1 (515) 294-8960; fax: +1 (515) 294-0221.
E-mail address: shlence@iastate.edu, (S. H. Lence).

transmission of prices, such that the correction of deviations from the long-run equilibrium displays threshold behavior. Applications of threshold cointegration to price transmission include Goodwin and Piggot (2001), Ben-Kaabia and Gil (2007), and Balcombe et al. (2007). Recent developments have featured increasingly sophisticated econometric specifications, such as polynomial fitting, Markov-switching, and smooth transition models (Goodwin et al., 2011; Hahn et al., 2016; Ihle et al., 2009; Serra et al., 2006).

A feature of the threshold autoregressive models of price transmission is that, under certain assumptions, they can recover transfer costs as estimated parameters. More specifically, in the canonical version of this model for two trading points, such parameters identify a “band of inactivity” such that arbitrage activities (e.g., trade) are triggered when price differences tend to move outside this band. This is attractive because the model offers a coherent way of capturing the nonlinearities implied by the presumption of switching regimes (trade or no trade, depending on whether price differences are smaller or greater than transfer costs). The modeling approach also elegantly deals with the fact that, for most empirical applications of price transmission, transfer costs and trade flows are often not observable.

Despite the wide diffusion of threshold autoregressive price transmission models, concerns have been raised from several points of view. One criticism focuses on the fact that price data alone may provide an insufficient basis to make inference about the performance of markets. Specifically, additional information and data on actual transport costs, quantities transacted and trade flows would be especially helpful (Barrett 2001). Another strand of criticism has focused on the econometric challenges of estimating threshold parameters. A broader cautionary perspective recognizes that applying these price transmission models to real-world data, and interpreting estimation results, requires a lucid grasp of the process being modeled, that is, “. . . a wider understanding of the structure and institutions underlying the price data” (Goodwin and Vavra, 2009, p. 10).

This article contributes to assessing the usefulness of threshold autoregressive price transmission models by reporting the results of a Monte Carlo experiment wherein several performance features of some standard inference procedures are assessed in a context where the data-generating process (DGP) is known. Studies in this setting typically focus on the modeler’s “estimation problem,” wherein one seeks the best procedure to recover the parameters of a model that is held to be true by assumption (e.g., Greb et al. 2013). In the spirit of Goodwin and Vavra’s (2009) thoughtful discussion, however, we propose to focus on what may be termed the modeler’s “specification problem,” whereby a parametric structure is postulated to represent and interpret a real-world setting of interest. As in Fackler and Tastan (2008), we emphasize the need for an explicit consideration of the underlying economic model of price determination. Hence, rather than relying on a reduced-form parameterization, the DGP that we invoke is based on a structural specification in which the notion of price transmission has a clear and explicit economic interpretation.

Specifically, the economic model that generates our data is predicated on the basic competitive model of spatial equilibrium (Takayama and Judge, 1971), but embeds two (related) features that, while plausible attributes of most real world situations, are often neglected in applied work (Fackler and Goodwin, 2001): delivery lags and rational expectations. Within this framework, a primary goal is to explore the empirical relevance of the econometrician’s problem, wherein the analyst typically cannot observe data on transfer costs or trade flows while attempting to provide inference on market integration based on econometric methods. We also study the implications of different attributes of transfer costs, as well as the speed of price transmission.

The novelty of our approach, relative to previous work that has evaluated the performance of threshold autoregressive price transmission models, follows from the motivation for our Monte Carlo experiment articulated in the foregoing discussion. Rather than seeking an improved estimation method, we take a standard inferential strategy as given and ask how well it captures the salient attributes of the underlying economic model. The particular structure of our DGP is predicated on a two-region equilibrium model, in which trade occurs according to rational expectations and the tradeable product is assumed to be perishable. The inferential strategy that we implement is based on testing for unit-roots and estimating a threshold autoregressive model (or a simpler autoregressive specification) in order to quantify the transfer costs. Furthermore, we explore the robustness of our findings through sensitivity analysis.

The remainder of the article is organized as follows. The next section presents the theoretical structure of DGP that embeds delivery lags and (rational) price expectations in a two-region spatial equilibrium model. This is followed by parameterization of the model and a discussion of the implied price dynamics. The various simulation scenarios considered are discussed next, along with the inference strategy used to evaluate the modeling approach. This is followed by a presentation of the results of the analysis, and a discussion of the main conclusions.

2. The DGP: price expectations and spatial arbitrage

The data used to evaluate the estimators of interest is generated from a structural representation of a prototypical market where delivery lags give rise to a meaningful role for price expectations to influence arbitrage decisions. The underlying economic environment is that of a competitive model of spatial equilibrium. To capture the essence of the economic forces at work, we focus on the simplest possible structure: a two-region equilibrium model in which trade occurs according to rational expectations and the tradeable product is assumed to be perishable.

Arbitrage entails that prices in two markets linked by trade tend to be equalized as a consequence of agents pursuing profitable trade opportunities. This results in the LOP noted earlier, according to which prices should differ at most by the transfer costs needed to move the goods from one location to the other



Fig. 1. Sequence of events in period t .

(except for corner solutions). The transfer costs determine the so-called “inactivity band” within which no adjustments take place. In analytical terms, the arbitrage conditions can be stated as follows:

$$|P^A - P^B| \geq T \quad \text{if } X > 0, \tag{1}$$

$$|P^A - P^B| < T \quad \text{if } X = 0, \tag{2}$$

where P^i denotes the price in region $i \in \{A, B\}$, $T > 0$ represents the per-unit transfer costs to ship the good between A and B , and X is the quantity shipped between regions.

The foregoing formulation makes it apparent why one should expect threshold effects in price transmission: the equilibrating role of spatial arbitrage (trade) does not come into play when prices differ by less than transfer costs. A variety of threshold cointegration specifications have been proposed for the econometric analysis of price transmission. Because analyzing all of them is not workable, here we focus our attention on the Band-TVECM proposed by Balke and Fomby (1997, p. 631), arguably the most popular and intuitive specification to model the arbitrage conditions (1) and (2):

$$z_t = \begin{cases} \theta(1 - \rho) + \rho z_{t-1} + e_t & \text{if } z_{t-1} > \theta, \\ z_{t-1} + e_t & \text{if } |z_{t-1}| \leq \theta, \\ -\theta(1 - \rho) + \rho z_{t-1} + e_t & \text{if } z_{t-1} < -\theta. \end{cases} \tag{3}$$

In the above expression, $z_t \equiv P_t^A - P_t^B - P_t^B$ is the spatial price differential, parameter θ determines the equilibrium band $[-\theta, \theta]$ to which variable z tends to return over time, $(1 - \rho)$ reflects the speed of price transmission, and e_t is an independently and identically distributed (i.i.d.) $N(0, \sigma^2)$ error. Within the so-called “inactivity band” $[-\theta, \theta]$, z follows a random walk without drift. If $\theta = 0$, the inactivity band disappears and the price differential follows a process with unconditional mean equal to zero and autocorrelation equal to $\rho \in (0, 1)$. Parameter ρ is inversely related to the speed at which z tends to revert to the band. When $\rho \rightarrow 1$, the speed of price transmission equals zero, and the inactivity band cannot be identified because the price differential behaves as a random walk everywhere.¹

¹ An earlier model assuming a stationary price differential inside the band of inaction was developed by Spiller and Woods (1988a, 1988b). Such assumption is a likely reason why their model has not been widely adopted, as it seems less realistic than the nonstationary no-trade price differential that characterizes the Band-TVECM (3).

Based on competitive spatial equilibrium theory, spatial price differentials z_t that exceed the inactivity band are corrected toward the edges of the band. Within the band, price differentials are not large enough to trigger arbitrage, and behave like a random walk. In terms of trading activities, three regimes can be distinguished based on the value of the price differential z_t , namely: (a) the price differential is larger than the threshold θ , so that trade occurs from B to A ; (b) the price differential is smaller in absolute value than the threshold, implying no trade; and (c) the price differential is smaller than the negative of the threshold (i.e., $P_t^B - P_t^A > \theta$), with trade taking place from A to B .

The performance of the econometric formulation of price transmission in (3), however, is bound to depend on the specific structure that is generating the prices P_t^i . In the experiment reported in this article, we assume a simple but explicit economic model of price determination with stochastic supply and demand conditions in each market, in a setting where agents’ expectations play a meaningful role.

Specifically, each region contains a terminal market where, in each period, price reflects the equilibrium between local demand and total shipments to that market from both regions. In each region production has to be pooled and shipped to either or both of the terminal market locations, and this basic marketing stage takes some time. We abstract from storage and assume that the product is perishable (e.g., vegetables) and must be consumed once it reaches the terminal markets. The two regions are labeled A and B , and the sequence of events is as follows: at the beginning of period t the amount of production S_t^i is realized in region $i \in \{A, B\}$ and becomes available for shipment to the terminal markets of either/both regions. The transportation/marketing stage takes some time, and so shipping decisions have to be based on the agents’ expectation of the equilibrium price that will emerge in the terminal markets later in the period, when consumption will take place. Given the assumption that the product is perishable and storage is not possible, the total amount shipped to either terminal location in period t is available to satisfy demand only in this period. Figure 1 illustrates the presumed timeline, where we adopt the convention of labeling the period by its beginning date.

Let $x_t^{ij} \in [0, 1]$ denote the share of output S_t^i shipped to market j , for $i, j \in \{A, B\}$. The fact that shipments from a region cannot exceed total availability in that region is expressed as the constraint

$$S_t^i \geq x_t^{ii} S_t^i + x_t^{ij} S_t^i, \tag{4}$$

for $i \neq j$ and $i, j \in \{A, B\}$. Similarly, consumption in region $i \in \{A, B\}$ in period t , labeled C_t^i , cannot exceed the amount available, in other words,

$$C_t^i \leq x_t^{ii} S_t^i + x_t^{ji} S_t^j \quad (5)$$

for $j \neq i$ and $j \in \{A, B\}$.

As noted, immediately after output S_t^i is obtained, economic agents decide how much to ship to regions A or B , a decision that is based on the expected gains (given the stochastic elements of the system to be discussed shortly) after paying for transfer costs. Let $T_t^{ij} \geq 0$ denote the transfer cost required to ship one unit of product from region i to region j in period t . Transfer costs may change from period to period, but are known to decision makers at the time of shipping. The cost of shipping between regions is strictly positive ($T_t^{ij} > 0$, $i \neq j$). Furthermore, we maintain the assumption that it is costlier to ship the good to the other region than to the terminal market of the own region (i.e., $T_t^{ij} > T_t^{ii}$, $i \neq j$); thus, without further loss of generality, we set $T_t^{ii} = 0$, $\forall i, t$. Given that the product is perishable, and together with regularity conditions on demand (to be detailed below), it will always be the case that in each period: (a) some of each region's supply is always shipped to the own market (i.e., $x_t^{ii} > 0$ for $i \in \{A, B\}$); (b) the entire supply in each region is shipped (i.e., expression (4) holds as an equality); and (c) in each market the total shipments arriving from both regions will be consumed (i.e., expression (5) holds as an equality). Condition (a) implies that $x_t^{ij} \in [0, 1]$ for $i \neq j$, and condition (b) means that

$$x_t^{ii} + x_t^{ij} = 1, \quad i, j \in \{A, B\}. \quad (6)$$

In a competitive setting where individual agents are atomistic and make shipping decisions so as to maximize expected profits conditional on their information at the time of shipment, *in equilibrium* the shares shipped to the farthest consumption region (\bar{x}_t^{AB} and \bar{x}_t^{BA}) must exhaust *ex ante* arbitrage opportunities, in other words, they must simultaneously satisfy the following conditions:

$$\begin{aligned} E_t(\bar{P}_t^B) - T_t^{AB} - E_t(\bar{P}_t^A) &\leq 0, \quad \bar{x}_t^{AB} \geq 0, \\ [E_t(\bar{P}_t^B) - T_t^{AB} - E_t(\bar{P}_t^A)]\bar{x}_t^{AB} &= 0, \end{aligned} \quad (7a)$$

$$\begin{aligned} E_t(\bar{P}_t^A) - T_t^{BA} - E_t(\bar{P}_t^B) &\leq 0, \quad \bar{x}_t^{BA} \geq 0, \\ [E_t(\bar{P}_t^A) - T_t^{BA} - E_t(\bar{P}_t^B)]\bar{x}_t^{BA} &= 0, \end{aligned} \quad (7b)$$

where the overbar denotes equilibrium values and $E_t(\bullet) \equiv E(\bullet|\Omega_t)$, (i.e., expectations are rational and conditional on the information set Ω_t at the time of shipping). This information set includes all fixed parameters of the model and the values of all random variables that are realized before the time when shipping decisions are made (more details below).

2.1. On the generality of the DGP structure: discussion

The proposed structure of the DGP is admittedly simple. It does not incorporate many real-world features, such as the possibility of storage and/or the existence of trade across many regions. This is a deliberate choice that is meant to make the outcome of our analysis crisper and easier to interpret. If the question of interest is whether a given econometric procedure can uncover certain properties of a DGP, then making the latter arbitrarily complex can only increase the likelihood of a negative answer. Conversely, if one finds that a given procedure performs poorly in a simple setting, the presumption that it would do better in more complex cases is not tenable.

To elaborate further on the advantages of the postulated DGP structure, note that the typical empirical approach consists of employing historical price series to infer transfer costs across regions. Thus, the empirical analysis of a model consisting of two regions relies on two observed price series (one for each region) to recover one transfer cost. If instead one were to analyze spatial equilibrium among many regions, the number of price series would increase linearly but the number of unknown objects would increase quadratically (e.g., for n regions one would use n observed price series to infer $[n(n-1)/2]$ transfer costs). Quite clearly, adding regions to the analysis can only exacerbate the problem of identifying transfer costs from the observed price data.

Explicitly allowing for storage decisions would similarly introduce nontrivial complexities (Coleman, 2009a). To the best of our knowledge, there is no study showing how to extend the standard storage model (Wright and Williams, 1989) to the case where the endogenously-determined equilibrium prices have a unit root, as required by our analysis. For instance, the numerical solution of the standard storage model provided by Miranda and Fackler (2002, pp. 215–217 and 298–301) deals with a setup that yields stationary prices. Nonetheless, we can conjecture some likely effects of storage on price behavior. The first impact is greater autocorrelation, due to intertemporal arbitrage that links the current price with the expected next-period price via storage. The second impact is smaller price changes in response to exogenous (e.g., weather) shocks, as the addition of the demand for stocks makes total demand more elastic. Hence, even though no storage is involved in the proposed DGP, the “Stochastic Supply with High Autocorrelation” and “High Demand Elasticity” scenarios discussed later (which, respectively, yield more autocorrelated prices and more subdued responses to exogenous shocks) may provide valuable insights about the possible impact of introducing storage on inference.

To anticipate some of our findings, results will show that transfer costs are poorly estimated from price data for simplest scenario of $n = 2$ regions and no storage. Hence, one can safely conclude that transfer costs would also be poorly estimated from price data for more complex scenarios with $n > 2$ regions and/or incorporating storage. In short, increasing the complexity of the model can only exacerbate the (lack of) identification problems that we find characterizes the simplest two-region trade model.

3. Parameterization and price dynamics

To make this framework operational we need to specify demand and supply conditions. Demand in each region is assumed to be isoelastic, so that the inverse demand functions are written as

$$P_t^i = \delta_t^i (C_t^i)^{-1/\varepsilon_i}, \quad (8)$$

where P_t^i is the price in region i in period t , $\delta_t^i > 0$ is a demand scaling factor defining the market size, and ε_i is the own-price elasticity of demand. With this demand parameterization, consumers' marginal willingness to pay is bounded away from zero ($P_t^i > 0$ as long as $C_t^i < \infty$), implying that total shipments to each region will be consumed. Together with Eqs. (4)–(6), this implies that equilibrium prices in the two markets are determined along the demand curves, given the quantities shipped to each market in equilibrium:

$$\bar{P}_t^A = \delta_t^A [(1 - \bar{x}_t^{AB}) S_t^A + \bar{x}_t^{BA} S_t^B]^{-1/\varepsilon_A}, \quad (9a)$$

$$\bar{P}_t^B = \delta_t^B [\bar{x}_t^{AB} S_t^A + (1 - \bar{x}_t^{BA}) S_t^B]^{-1/\varepsilon_B}. \quad (9b)$$

From the complementarity slackness conditions in (7a) and (7b), it follows that: (a) the expected spatial price differential is exactly equal to the transfer cost when there is a positive shipment between regions, and (b) between-region shipments are equal to zero when the expected price differential is strictly smaller than the transfer cost. Given the specific form of the assumed demand functions, the expected price differentials cannot exceed the relevant transfer cost between the regions.

In addition to the role of arbitrage, the dynamic behavior of prices in the two markets will depend on the dynamic properties of the exogenous shocks to supply and demand, which we specify as follows. Production in region i is assumed to follow the (covariance) stationary random process

$$\ln(S_t^i) = \mu_i + \gamma_i \ln(S_{t-1}^i) + v_t^i, \quad (10)$$

where μ_i is a scalar that defines the long-term mean supply, $\gamma_i \in [0, 1)$ is another scalar that determines the level of autocorrelation in supply, and v_t^i is an i.i.d. $N(-0.5 \psi_i^2, \psi_i^2)$ random shock. Under the assumed distribution for v_t^i , the random variable $\exp(v_t^i)$ is log-normally distributed, with mean equal to 1 and variance equal to $[\exp(\psi_i^2) - 1]$.²

Because supply is stationary by construction, to generate prices that are $I(1)$ (i.e., integrated of order one), demand (8) is assumed to be subjected to exogenous $I(1)$ shocks:

$$\ln(\delta_t^i) = \ln(\delta_{t-1}^i) + u_t^i, \quad (11)$$

² Therefore, the conditional mean and variance of S_t^i are respectively $E_t(S_t^i) = e^{\mu_i} (S_{t-1}^i)^{\gamma_i}$ and $\text{Var}_t(S_t^i) = (e^{\psi_i^2} - 1)[e^{\mu_i} (S_{t-1}^i)^{\gamma_i}]^2$, whereas the unconditional mean and variance are $E(S_t^i) = \exp[\mu_i/(1 - \gamma_i) - 0.5 \gamma_i \psi_i^2/(1 - \gamma_i^2)]$ and $\text{Var}(S_t^i) = \{\exp[\psi_i^2/(1 - \gamma_i^2)] - 1\} \exp[2 \mu_i/(1 - \gamma_i) - \gamma_i \psi_i^2/(1 - \gamma_i^2)]$, respectively.

where u_t^i is an i.i.d. $N(-0.5 \xi_i^2, \xi_i^2)$ random shock. Given process (11), and given the equilibrium quantity shipments \bar{x}_t^{ij} , the equilibrium price processes (9a) and (9b) can be written as

$$\bar{P}_t^A = \exp(u_t^A) \delta_{t-1}^A [(1 - \bar{x}_t^{AB}) S_t^A + \bar{x}_t^{BA} S_t^B]^{-1/\varepsilon_A}, \quad (12a)$$

$$\bar{P}_t^B = \exp(u_t^B) \delta_{t-1}^B [\bar{x}_t^{AB} S_t^A + (1 - \bar{x}_t^{BA}) S_t^B]^{-1/\varepsilon_B}. \quad (12b)$$

By construction, price series in both regions are $I(1)$ and exhibit threshold cointegration. Consistent with the timeline illustrated in Fig. 1, the realized values of the random variables u_t^i are not part of the information set Ω_t of the conditional expectation $E(\bullet|\Omega_t)$ introduced earlier (i.e., shipping decisions are made without knowing the demand shock that will be realized in the terminal markets). That is, if Ω_0 denotes the set of all fixed parameters of the model, then $\Omega_t = \Omega_{t-1} \cup \{u_{t-1}^A, u_{t-1}^B, v_t^A, v_t^B\}, \forall t > 0$.

4. Simulation scenarios

The purpose of the proposed economic model is to provide a transparent framework to assess the inferences based on threshold cointegration about the true economic structure. To this end, we use the model to simulate data and examine the performance of threshold cointegration inferences for alternative scenarios. The scenarios under study consist of specific parameterizations of the underlying economic model, aimed at obtaining simulated data representative of typical time series observations available to econometricians when conducting price analysis. For each scenario, we simulate 1,000 samples (i.e., sets of time series), each of them comprising either 260 or 520 observations for each of the variables analyzed. The design of the simulations and information about the numerical procedures used to solve the model are provided in the online appendix.

4.1. Transfer costs

A drawback encountered when performing econometric analyses of price transmission is that transfer costs are typically not observable, and usually assumed to be constant or simple functions of time. However, such a simplification might hinder the study of price dynamics because in reality transfer costs may change over time.

In order to understand the empirical relevance of transfer costs, we simulated the model under three alternative assumptions, namely, (a) fixed per-unit transfer costs, (b) exogenous time-varying per-unit transfer costs, and (c) endogenous time-varying per-unit transfer costs. In all instances, it is assumed that transfer costs are known to decision-makers at the time of making shipment decisions, but unknown to the econometrician.

The case of fixed per-unit transfer costs, designated “fixed transfer costs” for simplicity, assumes that $T_t^{ij} = T^{ij}, \forall t$. The case of exogenous time-varying per-unit transfer costs is

Table 1
Parameterizations and observations corresponding to different regimes for the alternative simulation scenarios

Scenario	Observation per sample	Demand elasticities ($\varepsilon_A, \varepsilon_B$)	Supply (S_t^i)	Transfer costs per unit (T_t^{ij})	Percentage of observations in different regimes					
					Spatial price differential			Trade		
					$z_t > \theta$	$ z_t \leq \theta$	$z_t < -\theta$	$x_t^{AB} > 0$	$x_t^{ij} = 0$	$x_t^{BA} > 0$
Baseline	520	0.7, 0.7	1	0.05	26.0	48.3	25.7	47.0	6.0	47.0
High fixed transfer costs	520	0.7, 0.7	1	0.2	19.8	60.2	20.0	38.7	22.3	39.0
Stochastic transfer costs	520	0.7, 0.7	1	$E(T_t^{ij}) = 0.05^a$	26.0	48.2	25.8	47.0	6.0	47.0
High stochastic transfer costs	520	0.7, 0.7	1	$E(T_t^{ij}) = 0.2^b$	19.9	60.1	20.0	38.8	22.2	39.0
Endogenous transfer costs	520	0.7, 0.7	1	$E(T_t^{ij}) = 0.124^c$	24.5	51.1	24.4	47.0	6.0	47.0
Small sample	260	0.7, 0.7	1	0.05	25.9	48.2	25.9	47.1	5.8	47.1
High demand Elasticity	520	1.5, 1.5	1	0.05	26.6	47.1	26.3	47.0	6.0	47.0
Different Demand Elasticities	520	0.7, 1.5	1	0.05	25.5	47.9	26.6	47.0	6.0	47.0
Stoch. supply with low autocorrelation	520	0.7, 0.7	$E(S_t^i) = 1^d$	0.05	26.4	47.5	26.1	47.0	6.0	47.0
Stoch. supply with high autocorrelation	520	0.7, 0.7	$E(S_t^i) = 1^e$	0.05	26.5	47.6	25.9	47.0	6.0	47.0

^aTransfer costs are stochastic, with DGP given by T_t^{ij} i.i.d. Beta(12, 12, 0.025, 0.075).

^bTransfer costs are stochastic, with DGP given by T_t^{ij} i.i.d. Beta(12, 12, 0.1, 0.3).

^cTransfer costs are endogenously determined according to the transportation supply function $T_t^{ij} = 0.05 + 0.26 x_t^{ij} S_t^i$.

^dOutput is stochastic, with DGP given by $\ln(S_t^i) = 0.00485 + 0.2 \ln(S_{t-1}^i) + v_t^i$, where v_t^i i.i.d. $N(-0.0291, 0.0582)$.

^eOutput is stochastic, with DGP given by $\ln(S_t^i) = 0.002728 + 0.9 \ln(S_{t-1}^i) + v_t^i$, where v_t^i i.i.d. $N(-0.0057595, 0.011519)$.

Note: Each scenario involved 1,000 simulated samples, each of them consisting of the reported number of observations.

obtained by drawing i.i.d. values of T_t^{ij} from a Beta probability distribution. To simplify the presentation, this case is labeled “stochastic transfer costs” in the remainder of the study. Finally, we consider a scenario with endogenous time-varying per-unit transfer costs, or “endogenous transfer costs” for short. International trade studies provide a variety of perspectives on how such transfer costs can be modeled (Behrens and Picard, 2011), but have consistently found that endogenous freight rates respond to trade imbalances (best exemplified, perhaps, by the so-called “backhaul problem”). Thus, we postulate that larger shipment volumes would need to bid away transportation resources from other uses, thereby increasing the unit transfer cost. Hence, for this case per-unit transfer costs are represented by

$$T_t^{ij} = \tau_0 + \tau_1 x_t^{ij} S_t^i, \quad (13)$$

where τ_0 and τ_1 are strictly positive parameters. Unlike the cases of fixed or stochastic transfer costs, endogenous transfer costs imply that the supply of transfer services is not infinitely elastic. Importantly, per-unit transfer costs are endogenous in equilibrium, because the equilibrium share of output shipped \bar{x}_t^{ij} that determines T_t^{ij} must simultaneously solve the no-arbitrage conditions (7a) and (7b).

4.2. Parameter values

The parameterization adopted for the baseline scenario consists of 520 observations per sample, fixed and symmetric transfer costs with $T_t^{ij} = 0.05$, a constant and unitary supply in each region ($\mu_i = \gamma_i = \psi_i = 0$), inelastic demand functions with $\varepsilon_i = 0.7$, and demand shocks with a variance of $\xi_i^2 = 0.00120$.

The latter value is the weekly equivalent of a 25% annual price volatility under autarky, which is a representative volatility for many agricultural commodities (FAO, 2004 p. 51; Gilbert, 2006 p. 51).³ This means that the simulations are aimed at mimicking samples of weekly time series observations. Thus, in the case of the baseline, each sample (520 observations) represents ten years of data.

The baseline transfer cost value of $T_t^{ij} = 0.05$ was determined by calibration. It was obtained by simulating the model so that approximately half of the spatial price differential observations fall inside the transfer cost band (i.e., $\text{Prob}(|z_t| \leq 0.05) = 0.5$), and the other half are evenly distributed above and below such band (i.e., $\text{Prob}(z_t > 0.05) = \text{Prob}(z_t < -0.05) = 0.25$). This was done to ensure a sufficiently large number of observations pertaining to each price differential regime to facilitate their identification. That is, if anything, the baseline simulated series are slanted in favor of identification.

The first four columns of Table 1 show the parameterizations of the alternative scenarios analyzed relative to the baseline. The “high fixed transfer cost” scenario is aimed at uncovering the impact of higher fixed transfer costs. The “stochastic transfer cost” and “high stochastic transfer cost” scenarios are the stochastic counterparts of the baseline and the high fixed transfer cost scenarios. Both of the stochastic transfer cost scenarios assume a four-parameter symmetric β probability distribution with a coefficient of variation equal to 10%, and a lower bound on the transfer cost equal to half of the mean. The “endogenous transfer cost” scenario corresponds to the

³ To see this, note that an annual variance of $(0.25)^2$ is equal to a weekly variance of $(0.25)^2/52 = 0.00120$.

situation where transfer costs are modeled as in Eq. (13), with $\tau_0 = 0.05$ and $\tau_1 = 0.26$. Setting $\tau_0 = 0.05$ ensures that trade under “endogenous transfer cost” occurs if and only if trade in the baseline occurs. Parameter $\tau_1 = 0.26$ was set by calibration, so that the unconditional mean transfer cost is approximately the average of the transfer costs in the baseline and the “high transfer cost” scenarios.

The “small sample” scenario is characterized by simulated samples with half of the observations assumed for the baseline, each of them representing five years of weekly data. The “high demand elasticity” scenario imposes elastic rather than inelastic demand in both regions, whereas the “different demand elasticities” scenario assumes that demand is inelastic in region A and elastic in region B. Finally, the “stochastic supply” scenarios with “low autocorrelation” and “high autocorrelation” are generated by setting, respectively, $[\gamma_i, \mu_i, \psi_i^2] = [0.2, 0.00485, 0.0582]$ and $[\gamma_i, \mu_i, \psi_i^2] = [0.9, 0.002728, 0.011519]$ in the production process (10). Both parameterizations imply the same unconditional expected output as in the baseline, and an unconditional coefficient of variation equal to $[\text{Var}(S_i^i)]^{0.5}/E(S_i^i) = 25\%$. However, the weekly autocorrelations are respectively $\gamma_i = 0.2$ and $\gamma_i = 0.9$ for the stochastic supply scenarios with “low” and “high” autocorrelation.

5. Inference strategy

The estimation strategy that we evaluate is aimed at inferring aspects of the true DGP from the simulated data. The procedure consists of two main tasks. The first one is classifying each sample according to the inferred number of spatial price differential z_t regimes comprised in it. The second task consists of estimating each sample’s DGP parameters, based on the results from the first task.

5.1. Classification of samples according to the number of spatial price differential regimes

The strategy to infer the number of z_t regimes in each sample is summarized in Fig. 2. Following standard practice, the first step consists of conducting an augmented Dickey–Fuller (ADF) unit root test (Dickey and Fuller, 1979) on each sample’s individual price series, followed by an ADF unit root test on the spatial price differential.⁴ The ADF test assumes $I(1)$ under the null hypothesis. The sample is excluded from further consideration if the null of a unit root is rejected for P_t^A or P_t^B , as cointegration requires both series to be $I(1)$.

The samples that are kept but fail to reject the ADF test null of a unit root for z_t are subjected to a Seo test of the

null hypothesis of linear no cointegration (Seo, 2006). We proceeded in this manner because the ADF test has much lower power than the Seo test against the alternative hypothesis of threshold cointegration (Seo, 2006), and therefore might incorrectly fail to reject the null. No additional tests are performed on samples that fail to reject Seo’s null, as they are inferred to consist entirely of $I(1)$ spatial price differentials. In other words, such samples are inferred to only contain observations pertaining to a single z_t regime, consisting of “within band” realizations.

All samples for which the null that z_t is $I(1)$ is rejected are assumed to consist of z_t observations that belong to at least one $I(0)$ regime. That is, such samples are taken to have some z_t realizations outside the no-arbitrage band. To determine whether a particular sample contains z_t realizations for a single regime (i.e., all z_t s are either below or above the band), two regimes (i.e., all z_t s are either below and within the band, or above and within the band), or all three regimes, we apply Hansen’s sequential testing procedure (Hansen, 1999). Hansen’s F_{12} (F_{13} , F_{23}) test postulates a null hypothesis of one (one, two) regime(s) versus an alternative of two (three, three) regimes. Hansen’s F_{12} test is conducted first. If F_{12} ’s null is rejected, favoring two regimes rather than one, the sample is subjected to the F_{23} test. In this manner, rejecting (failing to reject) F_{23} ’s null allows us to conclude that the sample contains z_t observations pertaining to three (two) regimes. If F_{12} ’s null is not rejected, the F_{13} test is conducted to discriminate between samples involving one and three regimes. We infer that the sample comprises observations for three regimes if F_{13} ’s null is rejected, and for only one regime otherwise.

The ADF test was performed using the “adftest” function in the MATLAB program. The tsDyn package (Stigler, 2013) written in the R program was used to conduct Seo’s and Hansen’s tests. The trimming for the tests was set at 10% (i.e., each regime was assumed to comprise at least 10% of the observations in any of the tested samples).

5.2. Estimation of transfer costs and speed of price transmission for each sample

Once it has been determined whether a sample comprises one $I(0)$, two, or three price differential regimes, one can proceed to estimate its corresponding transfer costs and speed of price transmission. In the case of samples inferred to have one $I(0)$ z_t regime, the outer equations in expression (3) imply that transfer costs can be estimated by running the unrestricted autoregression

$$z_t = \alpha + \rho z_{t-1} + e_t. \quad (14)$$

The estimated speed of price transmission is $(1 - \hat{\rho})$, and the transfer cost estimate is recovered as $|\hat{\alpha}/(1 - \hat{\rho})|$, where hats represent the respective parameter estimates.

For samples inferred to comprise two z_t regimes, the economic model consists of a system of two equations representing

⁴ Even though individual price series are $I(1)$ and the spatial price differential z_t is $I(0)$ by construction, in a particular sample any P_t^A , P_t^B , and z_t could be inferred to be $I(0)$ or $I(1)$. However, for a sufficiently large sample with frequent enough spatial arbitrage, one would most likely infer P_t^A , P_t^B , and z_t to be $I(1)$, $I(1)$, and $I(0)$, respectively.

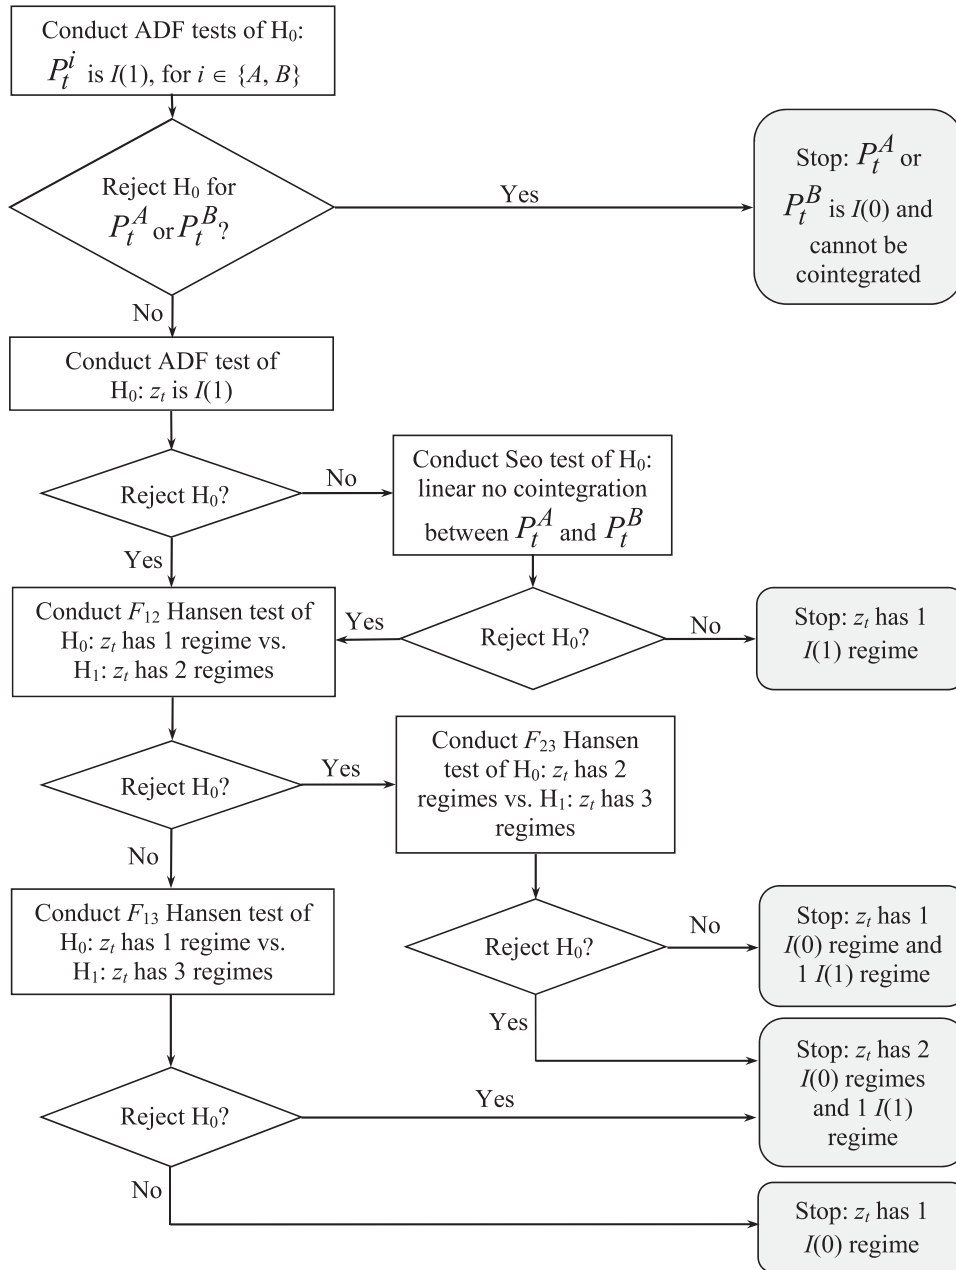


Fig. 2. Econometric strategy to classify individual simulated samples.

the middle and one of the outer regimes in expression (3). Similarly, samples inferred to have three z_t regimes are assumed to span all of the three regimes. All models are estimated by concentrating the likelihood function on the threshold θ (Hansen and Seo, 2002). Succinctly, we set a 10% trim and perform a grid search over the potential threshold values. For each threshold value θ_i on the grid, we obtain the concentrated likelihood $L(\theta_i)$ (i.e., the value of the likelihood function maximized with respect to other parameters of the model). The estimated transfer cost $\hat{\theta}$ is the value θ_i that results in the largest value for $L(\theta)$ over the grid.

6. Results

Table 1 summarizes the features of the various parameterizations of the scenarios analyzed, as well as the percentages of observations belonging to each price differential regime and to each trading regime. As explained earlier, transfer costs in the baseline scenario have been calibrated so that observations on price differentials z_t are about evenly distributed inside and outside the band of transfer costs. Even though almost half of the observations on price differentials are inside the band, only 6% of the observations involve no trade. This seemingly

Table 2
True and inferred number of samples of different types in the baseline scenario

		True number of z_t regimes in the sample			Row total	
		One		Two		
		$I(1)$	$I(0)$			
Inferred number of z_t regimes in the sample	P_t^A or P_t^B are $I(0)$	0	0	33	71	104
	One $I(1)$	0	0	127	113	240
	One $I(0)$	0	0	20	325	345
	Two	0	0	17	190	207
	Three	0	0	1	103	104
Column total		0	0	198	802	1000

Note: There were 65 samples for which the numerical procedures for the Seo test failed. Consistent with the stated intent of favoring correct inferences, such samples were classified as having cointegrated prices (i.e., it was assumed that for them the Seo test rejected the null of linear no cointegration between P^A and P^B).

counterintuitive result is actually an important feature of the problem at hand. Recall that the arbitrage conditions in (7) depend on expected prices, and they imply that it is the expected price differential for any observation involving positive trade that must be on the boundary of the transfer cost band. Realized prices differ from expected prices by the unanticipated realized shock. In fact, conditional on the equilibrium trade flows \bar{x}_t^j , the three regimes of the DGP can be represented as follows

$$z_t = \begin{cases} \theta + e_t & \text{if } \bar{x}_t > 0, \\ z_{t-1} + e_t & \text{if } \bar{x}_t = 0, \\ -\theta + e_t & \text{if } \bar{x}_t < 0, \end{cases} \quad (15)$$

where $\bar{x}_t \equiv \bar{x}_t^{AB} - \bar{x}_t^{BA}$. Hence, actual price differentials when trade occurs fall outside the band with about 50% probability. Thus, the 94% of observations where trade occurs imply about 47% of realized price differentials outside the transfer cost band. More generally it is worth noting that, with the explicit structural DGP that we postulated, realized price differentials outside of the band are fully consistent with equilibrium and arbitrage.

Except for the scenarios with large transfer costs, the percentages of observations in different price differential regimes are similar across scenarios. The same is true regarding the percentages of observations in different trading regimes. As expected, scenarios with larger transfer costs yield larger percentages of observations inside the transfer cost band and involving no trade. Importantly, however, all of the scenarios include a relatively large percentage of observations for each of the price differential regimes. This implies that, if anything, the simulated data unduly favor the identification of regimes.

Results for the baseline scenario are discussed in greater depth in the next subsections. Results for the alternative scenarios are largely consistent with the baseline results; hence, in the interest of space, they are reported and discussed in the online appendix.

6.1. Baseline scenario: inference about price differential regimes

Results regarding inferences on the number of price differential regimes in each sample are reported in Table 2. Column totals show that 198 out of the 1,000 samples contain price differentials that are inside and below or above the transfer cost band (i.e., belong to two regimes), and that the other 802 samples comprise price differentials inside, below, and above the band (i.e., belong to three regimes). In contrast, the inference procedure yields 207 two-regime samples and only 104 three-regime samples.

There are several factors explaining the striking differences between the true and the inferred number of regimes. First, 10.4% of the samples are incorrectly inferred to involve at least one stationary price series. Such a percentage is somewhat large given the 5% significance level adopted for the Dickey–Fuller tests. Overrejection arises from using price levels rather than price logarithms to perform the tests, as tests conducted in logarithms yield percentages of unit-root rejections consistent with the 5% significance level.⁵

Second, 24.0% of the samples are incorrectly categorized as having price differentials which are all inside the transfer cost band, because they are inferred to consist only of non-cointegrated non-stationary prices. The flow chart in Fig. 2 reveals that both the Dickey–Fuller test and the Seo test are the sources of these inferential errors. Importantly, if anything, Table 2 understates the actual number of incorrect inferences, because the 6.5% of the samples that failed the numerical computation of the Seo test were assumed to have cointegrated prices to favor the odds of obtaining correct inferences.

Third, 34.5% of the samples were erroneously inferred to consist of cointegrated price differentials in only one regime

⁵ Most practitioners use prices in logarithms when defining price differentials. However, as shown in the theoretical model section, consistency with trade models requires that price differentials in the empirical specification be constructed using price levels, as opposed to logarithms. Price levels lead to overrejection at the first of the estimation procedure, but they most definitely favor identification at the following stages, because in the true model the no-arbitrage restriction involves price levels.

(i.e., price differentials either only above or only below the transfer cost band). This inferential mistake accounts for the largest number of misclassifications, and is attributable to Hansen’s F_{12} and F_{13} tests (see Fig. 2).

The final major source of faulty inferences is the large number (190) of three-regime samples misclassified as having only two regimes. This error is worth noting because the similarity between the inferred and true numbers of two-regime samples (207 vs. 198) might be misleading: in fact, only 17 out of the 207 samples inferred to have only two regimes actually do so. According to Fig. 2, these inferential mistakes arise from Hansen’s F_{13} and F_{23} tests.

Before turning to the estimated transfer costs, it is worth emphasizing the poor performance of the Seo test, especially given the heavy computational efforts it requires. In the case of the tsDyn package in R, conducting the Seo test on a single sample took about 50 times as long as performing Hansen’s F_{12} and F_{13} tests together. Yet, the numerical computation of the Seo test failed in 65 out of the 305 tested samples, and it did not reject the null of no cointegration in any of the other 240 samples. The present results suggest that, from a practitioner’s perspective, carrying out the Seo test may simply not be worth the effort.

6.2. Baseline scenario: estimated transfer costs

Summary statistics regarding estimated transfer costs for the baseline scenario are reported in the second column of Table 3. Point estimates of transfer costs based on the inferred one-regime samples (mean and median equal to 0.0477 and 0.050, respectively) are quite similar to the true transfer costs of 0.050. Interestingly, one-regime estimates have a distribution noticeably skewed to the left, as evidenced by the 5% and 95% quantiles of 0.025 and 0.0055, respectively. Transfer cost estimates based on the inferred two-regime samples exhibit

similar characteristics, except for a larger standard deviation and a lesser left skew (5% and 95% quantiles of 0.030 and 0.060, respectively).

Estimates based on the inferred three-regime samples are the worst ones. The mean estimate (0.038) and the median (0.044) are respectively 24% and 12% smaller than the true transfer cost (0.050). Further, the estimates exhibit a distribution highly skewed to the left: The 5% quantile equals only 0.011, and the 95% quantile of 0.051 barely includes the true transfer cost. The poor performance of the model at estimating the band parameter is consistent with the Monte Carlo results reported in Greb et al. (2013) (whose DGP, however, is different from the one studied in this article, as explained earlier).

As discussed in the previous subsection, inferences about the number of regimes are incorrect for most of the samples. Hence, it is of interest to uncover the impact of such errors on the estimation of transfer costs. To this end, the third column in Table 3 provides summary statistics for estimates based on the true number of regimes in each sample. Because none of the samples truly consists of z_t observations that belong to one $I(0)$ regime, the corresponding cells are labeled “not applicable.”

In the case of transfer costs based on true two-regime samples, point estimates (mean and median equal to 0.0505 and 0.050, respectively) are almost identical to the true transfer costs. Transfer costs are estimated much more precisely when using true as opposed to inferred two-regime samples, as the standard deviation for the former (equal to 0.0019) is almost one order of magnitude smaller than the standard deviation for the latter (equal to 0.015). Another major difference between the two types of estimates is that the ones based on true two-regime samples have a fairly symmetric distribution, with 5% and 95% quantiles of 0.047 and 0.054, respectively.

The substantially lower quality of the estimates obtained from the inferred two-regime samples should not be surprising, given the inferences shown in Table 2. This table shows that only 17

Table 3
Estimated transfer costs for the baseline scenario^a

	Transfer cost estimates based on	
	Inferred sample types	True sample types
	Mean (standard deviation) [5%, 50%, 95%] Quantiles	Mean (standard deviation) [5%, 50%, 95%] Quantiles
One $I(0)$ z_t regime samples	0.0477 (0.0097) [0.025, 0.050, 0.055]	Not applicable Not applicable Not applicable
Two z_t regime samples	0.048 (0.015) [0.030, 0.049, 0.060]	0.0505 (0.0019) [0.047, 0.050, 0.054]
Three z_t regime samples	0.038 (0.014) [0.011, 0.044, 0.051]	0.031 ^b (0.015) ^b [0.011, 0.029, 0.051] ^b

^aTrue transfer costs equal 0.05.

^bEstimates based on true three-regime samples are obtained using only samples which have both (i) at least 10% of the price differential observations greater than transaction costs, and (ii) at least 10% of the price differential observations smaller than transfer costs.

out of the 207 samples inferred to have two regimes truly have two regimes, whereas the other 190 samples actually have three regimes. In other words, 190 out of the 207 inference-based cost estimates are obtained by applying two-regime methods to three-regime samples.

Surprisingly, the transfer cost estimates based on true three-regime samples are of lesser quality than their counterparts based on inferred three-regime samples, as the mean and median exhibit a greater downward bias and the standard deviation is slightly higher. This counterintuitive result is driven by the fact that 103 out of the 104 samples inferred to have three regimes indeed have three regimes, and 699 out of the 802 true three-regime samples are inferred to not contain three regimes (see Table 2). Thus, the inference-based three-regime cost estimates are obtained from the subset of true three-regime samples with the strongest three-regime features.

Given the structural representation in (15) noted earlier, it is clear that if shipment data were available, one could use them not only to determine the number of regimes in a sample, but also to estimate transfer costs. Importantly, shipment data would obviate the need to perform inferences on the number of price differential regimes. In fact, if shipment data were available, transfer costs could be estimated by simply fitting linear regressions like (14) using only the price differentials corresponding to nonzero shipments. Unfortunately, implementing this strategy may not be feasible in most empirical applications, as the econometrician may not be able to obtain shipment data.

6.3. Baseline scenario: estimated speed of price transmission

The second column of Table 4 reports summary statistics for the estimated speeds of price transmission in the baseline scenario. Point estimates from the inferred one- and two-regime samples are centered around the true value of one (e.g., the

respective medians are 0.993 and 1.00), whereas those from the inferred three-regime samples are downwardly biased (mean and median equal to 0.90 and 0.92, respectively). Surprisingly, the one- and two-regime estimates have opposite skewness, the former to the left and the latter to the right. The precision of the estimates worsens substantially as one moves from the inferred one-regime samples to the three-regime samples (standard deviations are respectively 0.078, 0.17, and 0.37 for the one-, two-, and three-regime estimates). Overall, the three-regime samples yield very noisy estimates, with 5% and 95% quantiles equal to 0.25 and 1.52, respectively.

Summary statistics for speed of price transmission estimates based on the true number of regimes in each sample are shown in the third column of Table 4. Compared to the inference-based estimates, the ones based on the true regimes are more accurate and, in the case of the three-regime samples, less biased (e.g., the mean is 1.04 and the median equals 0.97).

6.4. Baseline scenario: inferences about individual observations

In practice, it is not uncommon to use the estimated transfer costs to make inferences as to whether (i) the individual price differential observations in the sample are within or outside the cost band, or (ii) whether trade occurred at the times corresponding to the sample observations. We provide information about inferences (i) and (ii) for the baseline scenario in Tables 5 and 6, respectively. The reported results correspond to samples inferred to have two and three z_t regimes (panels A and B, respectively), because the empirical threshold cointegration literature concentrates its attention on multi-regime cases.

According to Table 5, inferences about price differentials being within or outside the band are reasonably accurate for samples inferred to have two z_t regimes, and less so for their

Table 4
Estimated speed of price transmission for the baseline scenario^a

	Speed of price transmission estimates based on	
	Inferred sample types	True sample types
	Mean (standard deviation) [5%, 50%, 95%] Quantiles	Mean (standard deviation) [5%, 50%, 95%] Quantiles
One $I(0)$ z_t regime samples	0.981 (0.078) [0.820, 0.993, 1.071]	Not applicable Not applicable Not applicable
Two z_t regime samples	1.02 (0.17) [0.80, 1.00, 1.29]	1.03 (0.11) [0.82, 1.04, 1.20]
Three z_t regime samples	0.90 (0.37) [0.24, 0.92, 1.52]	1.04 ^b (0.29) ^b [0.70, 0.97, 1.60] ^b

^aTrue speed of price transmission equals 1.00.

^bEstimates based on true three-regime samples are obtained using only samples which have both (i) at least 10% of the price differential observations greater than transfer costs, and (ii) at least 10% of the price differential observations smaller than transfer costs.

Table 5
Percentages of observations in alternative spatial price differential regimes for the baseline scenario

		True z_t regime ^c		Row total
		Within band	Outside band	
A. Conditional on samples inferred to have two z_t regimes. ^a				
Inferred z_t Regime ^b	Within band	45.6	1.9	47.5
	Outside band	4.7	47.8	52.5
Column total		50.3	49.7	100.0
B. Conditional on samples inferred to have three z_t regimes. ^d				
Inferred z_t Regime ^c	Within band	40.3	0.3	40.6
	Outside band	15.6	43.8	59.4
Column total		55.9	44.1	100.0

^aThe table only classifies 20.7% of the total simulated observations, corresponding to the 207 baseline samples inferred to have two z_t regimes.

^bFor each sample i ($i = 1, \dots, 207$), z_t observations are inferred to be within the band if $|z_t| \leq \hat{\theta}_i$, and outside the band otherwise, where $\hat{\theta}_i$ is the transfer cost estimated using only sample i data.

^cThe true z_t regime is within the band if $|z_t| \leq 0.05$, and is outside the band otherwise.

^dThe table only classifies 10.4% of the total simulated observations, corresponding to the 104 baseline samples inferred to have three z_t regimes.

^eFor each sample i ($i = 1, \dots, 104$), z_t observations are inferred to be within the band if $|z_t| \leq \hat{\theta}_i$, and outside the band otherwise, where $\hat{\theta}_i$ is the transfer cost estimated using only sample i data.

Table 6
Percentages of observations in alternative trading regimes for the baseline scenario

		True trading regime ^c		Row total
		No trade	Trade	
A. Conditional on samples inferred to have two z_t regimes ^a				
Inferred trading regime ^b	No trade	6.5	41.0	47.5
	Trade	2.3	50.2	52.5
Column total		8.8	91.2	100.0
B. Conditional on samples inferred to have three z_t regimes ^d				
Inferred trading regime ^e	No trade	12.9	27.7	40.6
	Trade	5.1	54.3	59.4
Column total		18.0	82.0	100.0

^aThe table only classifies 20.7% of the total simulated observations, corresponding to the 207 baseline samples inferred to have two z_t regimes.

^bFor each sample i ($i = 1, \dots, 207$), “no trade” is inferred if $|z_t| \leq \hat{\theta}_i$, and “trade” is inferred otherwise, where $\hat{\theta}_i$ is the transfer cost estimated using only sample i data.

^cThe true regime is “no trade” if trade did not occur, and “trade” otherwise.

^dThe table only classifies 10.4% of the total simulated observations, corresponding to the 104 baseline samples inferred to have three z_t regimes.

^eFor each sample i ($i = 1, \dots, 104$), “no trade” is inferred if $|z_t| \leq \hat{\theta}_i$, and “trade” is inferred otherwise, where $\hat{\theta}_i$ is the transfer cost estimated using only sample i data.

three z_t regime counterparts. The results in panel A indicate that whenever an inferred two-regime sample yields transfer cost estimates considerably different from the true transfer costs,

it contains few price differential observations in the interval determined by the estimated and the true transfer costs.

Table 5 reveals that most of the misclassifications consist of price differentials which are actually within the band but are inferred to be outside the band, especially for inferred three regime samples. This finding is consistent with the fact that transfer cost estimates are downwardly biased, especially for inferred three-regime samples (see Table 3).

Studies of price transmission by threshold cointegration analysis are often interested in assessing market integration and market efficiency. Market integration is a concept typically associated with the ease of tradability (Baulch, 1997). For our two-region model, an obvious metric to measure the extent of market integration is the probability (or frequency) of trade taking place, which itself depends on the size of transfer costs. If transfer costs were prohibitively high, no trade would ever take place, and segmented equilibria would result wherein each market absorbs its own supply and demand shocks. Conversely, if transfer costs were zero, trade would always take place (demand and supply shocks are fully shared between the regions), and we would say markets are perfectly integrated. As transfer costs increase from zero, the frequency of trade decreases, and hence we can measure market integration by how often the two regions trade with one another.

The results in Table 6 suggest that inferences about trade are quite poor, with the main problem being that “no trade” is inferred much more often than warranted (e.g., 47.5% vs. 8.8% for inferred two-regime samples). This result emerges even though, somewhat paradoxically, the model also tends to underestimate the transfer costs (band) parameter, a fact underscoring the relevance of the economic model that we have postulated as the true DGP. The threshold cointegration approach that we have evaluated, therefore, tends to suggest a much lower degree of market integration than is true in the DGP. In light of the previously discussed downwardly biased transfer cost estimates, this finding is seemingly paradoxical; however, it can be easily explained by the structure of the economic model. Whereas the occurrence of trade implies that the expected price differential must be on the boundary of the transfer cost band, it is fully consistent with realized price differentials both inside and outside the band (see discussion of expression (15)).

7. Conclusions

The Band-TVECM specification is a popular model to study spatial price transmission. It is rooted in the econometrics of time-series cointegration, and the threshold framework it implements is an appealing way to capture the fact that transfer costs are real impediments to the arbitrage role of trade. In this model, spatial price differences that exceed transfer costs disappear over time, as implied by market equilibrium. This appears *prima facie* consistent with the notion that market equilibrium cannot absorb shocks instantaneously and that time is required for trade to perform its arbitrage role. How well Band-TVECM

parameterizations perform at representing the dynamics of spatial market equilibrium, however, has remained an open question.

Previous work on this topic has privileged the study of estimation challenges inherent with this econometric model, while maintaining the validity of the model's parametric structure (e.g., Greb et al., 2013). In this article, we take an alternative tack, and ask how well the Band-TVECM formulation works at capturing the arbitrage conditions of spatial equilibrium. As noted, trade shipments inevitably take time, suggesting that expectations are bound to play a key role in arbitrage decisions. How market equilibrium adjusts to shocks, therefore, will depend on how such expectations affect relevant decisions. Because this is likely to be context-specific, an explicit economic model of price transmission is necessary in order to evaluate the econometric model of interest. Consequently, in this article we use an explicit economic model—a simple two-region model for a perishable product, a setting where the competitive equilibrium critically depends on agents' expectations—to generate the data that is then used to evaluate the performance of standard inference strategies with the Band-TVECM approach.

We found remarkable differences between the true and the inferred number of price differential regimes in each sample. Even though most samples truly comprise three price differential regimes, the majority of the samples are inferred to contain only one regime. This finding suggests that only a small proportion of real-world price series could end up being analyzed by threshold cointegration. The inferential errors regarding the number of regimes in each sample emerges notwithstanding our reliance on the canonical tests proposed by Dickey and Fuller, Hansen, and Seo. Importantly, the poor performance of the Seo test suggests it may not be worthwhile to conduct it, especially in view of the computational effort that it requires.

Transfer costs estimates based on the inferred three-regime samples were found to be severely downwardly biased. Importantly, they are better than the transfer costs estimated from the true three-regime samples. This counterintuitive result is explained by the fact that the Band-TVECM does not exactly match the structure of the economic model underlying the data. Worryingly, it suggests that improved inference regarding the number of regimes in the sample may not lead to better transfer cost estimates.

The speed of price transmission estimated from the inferred tree-regime samples is downwardly biased and has a large standard deviation. The scenario with endogenous transfer costs and the two scenarios with high transfer costs exhibit the largest bias, with mean estimates 30–40% smaller than the true speed of price transmission.

Inferences about individual observations of price differentials being within or outside the band proved to be reasonably accurate, especially for samples inferred to have two regimes. In contrast, inferences as to whether individual observations involved trade or not were quite poor, with “no trade” being inferred far more often than warranted. The result is seemingly paradoxical, as it suggests that threshold cointegration may mis-

lead the practitioner into underestimating both the magnitude of transfer costs and the extent of market integration. This underscores the importance of the actual economic model, which in our case yields about half of the observations on price differentials outside the transfer costs band, even though trade occurs most of the time.

Based on the tests that most parsimoniously mimic the structure of the underlying trading model, our results indicate that the Band-TVECM typically used to analyze price transmission, while intuitively appealing, fails to capture essential features of the dynamics concerning expectations and trade. Even though it is beyond the scope of the present analysis to examine all of the tests that could be used (e.g., Enders and Granger, 1998; Kapetanios and Shin, 2006), our results suggest that, to the extent that they only rely on price data for inference, they are likely to face the same limitations as the tests used here. As such, our study suggests that results from applying the Band-TVECM should be more strongly qualified, and that it would be valuable to develop threshold cointegration models whose structure is better aligned with the economic fundamentals of trade.

References

- Ardeni, P.G., 1989. Does the Law of One Price really hold for commodity prices? *Am. J. Agr. Econ.* 71, 661–669.
- Balcombe, K., Bailey, A., Brooks, J., 2007. Threshold effects in price transmission: The case of Brazilian wheat, maize and soya prices. *Am. J. Agr. Econ.* 89, 308–323.
- Balke, N.S., Fomby, T.B., 1997. Threshold cointegration. *Int. Econ. Rev.* 38, 627–645.
- Barrett, C.B., 2001. Measuring integration and efficiency in international agricultural markets. *Rev. Agr. Econ.* 23, 19–32.
- Baulch, B., 1997. Transfer costs, spatial arbitrage, and testing for food market integration. *Am. J. Agr. Econ.* 79, 477–487.
- Ben-Kaabia, M., Gil, J.M., 2007. Asymmetric price transmission in the Spanish lamb sector. *Eur. Rev. Agr. Econ.* 34, 53–80.
- Behrens, K., Picard, P.M., 2011. Transportation, freight rates, and economic geography. *J. Int. Econ.* 85, 280–291.
- Coleman, A., 2009a. A model of spatial arbitrage with transport capacity constraints and endogenous transport prices. *Am. J. Agr. Econ.* 91, 42–56.
- Coleman, A., 2009b. Storage, slow transport, and the Law of One Price: Theory with evidence from nineteenth-century U.S. corn markets. *Rev. Econ. Stat.* 91, 332–350.
- Dickey, D.A., Fuller, W.A., 1979. Distribution of the estimators for autoregressive time series with a unit root. *J. Am. Stat. Assoc.* 74, 427–431.
- Enders, W., Granger, C., 1998. Unit-root tests and asymmetric adjustment with an example using the term structure of interest rates. *J. Bus. Econ. Stat.* 16, 304–311.
- Fackler, P.L., Goodwin, B.K., 2001. Spatial price analysis. In: Gardner, B.L., and Rausser, G.C., (Eds.), *Handbook of Agricultural Economics*, Vol. 1, Part 2, pp. 971–1024. Elsevier, Amsterdam.
- Fackler, P.L., Tastan, H., 2008. Estimating the degree of market integration. *Am. J. Agr. Econ.* 90, 69–85.
- Food and Agriculture Organization of the United Nations (FAO), 2004. *The State of Agricultural Commodity Markets 2004*. FAO, Rome.
- García-Germán, S., Bardají, I., Garrido, A., 2016. Evaluating price transmission between global agricultural markets and consumer food price indices in the European Union. *Agr. Econ.* 47, 59–70.
- Gilbert, C.L., 2006. Trends and volatility in agricultural commodity prices. In: Sarris, A., and Hallam, D., (Eds.), *Agricultural Commodity Markets and*

- Trade: New Approaches to Analyzing Market Structure and Instability, Ch. 2 pp. 31–60. Edward Elgar Publishing, Northampton, MA.
- Goodwin, B.K., Holt, M.T., Prestemon, J.P., 2011. North American oriented strand board markets, arbitrage activity, and market price dynamics: A smooth transition approach. *Am. J. Agr. Econ.* 93, 993–1014.
- Goodwin, B.K., Piggott, N., 2001. Spatial market integration in the presence of threshold effects. *Am. J. Agr. Econ.* 83, 302–317.
- Goodwin, B.K., Vavra, P., 2009. What can we learn from spatial and vertical price transmission studies? Empirical examples from U.S. meat markets. Courant Research Centre ‘Poverty, Equity and Growth’ Inaugural Conference, Georg-August-Universität Göttingen, Germany.
- Greb, F., von Cramon-Taubadel, S., Krivobokova, T., Munk, A., 2013. The estimation of threshold models in price transmission analysis. *Am. J. Agr. Econ.* 95, 900–916.
- Hahn, W., Stewart, H., Blayney, D.P., Davis, C.G., 2016. Modeling price transmission between farm and retail prices: A soft switches approach. *Agr. Econ.* 47, 193–203.
- Hansen, B.E., 1999. Testing for linearity. *J. Econ. Surveys* 13, 551–576.
- Hansen, B.E., Seo, B., 2002. Testing for two-regime threshold cointegration in vector error-correction models. *J. Econometr.* 110, 293–318.
- Ihle, R., von Cramon-Taubadel, S., Sergiy, Z., 2009. Markov-switching estimation of spatial maize price transmission processes between Tanzania and Kenya. *Am. J. Agr. Econ.* 91, 1432–1439.
- Kapetanios, G., Shin, Y., 2006. Unit root tests in three-regime SETAR models. *Econometr. J.* 9, 252–278.
- Miranda, M.J., Fackler, P.L., 2002. *Applied Computational Economics and Finance*. MIT Press, Cambridge, MA.
- Myers, R.J., Jayne, T.S., 2012. Multiple-regime spatial price transmission with an application to maize markets in Southern Africa. *Am. J. Agr. Econ.* 94, 174–188.
- Seo, M., 2006. Bootstrap testing for the null of no cointegration in a threshold vector error correction model. *J. Econometrics* 134, 129–150.
- Serra, T., Gil, J.M., Goodwin, B.K., 2006. Local polynomial fitting and spatial price relationships: Price transmission in EU pork markets. *Eur. Rev. Agr. Econ.* 33: 415–436.
- Shiue, C.H., 2002. Transport costs and the geography of arbitrage in eighteenth-century China. *Am. Econ. Rev.* 92, 1406–1419.
- Spiller, P.T., Wood, R.O., 1988a. Arbitrage during the dollar-sterling gold standard, 1899–1908: An econometric approach. *J. Polit. Econ.* 96, 882–892.
- Spiller, P.T., Wood, R.O., 1988b. The estimation of transaction costs in arbitrage models. *J. Econometrics* 39, 309–326.
- Stephens, E.C., Mabaya, E., von Cramon-Taubadel, S., Barrett, C., 2011. Spatial price adjustment with and without trade. *Oxf. B. Econ. Stat.* 74, 453–469.
- Stigler, M., 2013. Threshold cointegration: Overview and implementation in R. Working Paper. Accessed June 2017, available at <https://cran.r-project.org/web/packages/tsDyn/vignettes/ThCointOverview.pdf>
- Takayama, T., Judge, G.G., 1971. *Spatial and Temporal Price and Allocation Models*. North-Holland, Amsterdam.
- Wright, B.D., Williams, J.C., 1989. A theory of negative prices for storage. *J. Fut. Markets* 9, 1–13.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher’s website:

- Table B.2.1. True and inferred number of samples of different types in the High Fixed Transfer Costs scenario.
- Table B.2.2. True and inferred number of samples of different types in the Stochastic Transfer Costs scenario.

- Table B.2.3. True and inferred number of samples of different types in the High Stochastic Transfer Costs scenario.
- Table B.2.4. True and inferred number of samples of different types in the Endogenous Transfer Costs scenario.
- Table B.2.5. True and inferred number of samples of different types in the Small Sample scenario.
- Table B.2.6. True and inferred number of samples of different types in the High Demand Elasticity scenario.
- Table B.2.7. True and inferred number of samples of different types in the Different Demand Elasticities scenario.
- Table B.2.8. True and inferred number of samples of different types in the Stochastic Supply with Low Autocorrelation scenario.
- Table B.2.9. True and inferred number of samples of different types in the Stochastic Supply with High Autocorrelation scenario.
- Table B.3.1. True and estimated transfer costs for the alternative scenarios, based on the inferred number of z_t regimes in the sample.
- Table B.3.2. True and estimated transfer costs for the alternative scenarios under two and three regimes, based on the true number of z_t regimes in the sample.
- Table B.4.1. Estimated speed of price transmission for the alternative scenarios, based on the inferred number of z_t regimes in the sample.
- Table B.4.2. Estimated speed of price transmission for the alternative scenarios under two and three regimes, based on the true number of z_t regimes in the sample.
- Table B.5.1. Percentages of observations in alternative spatial price differential regimes for the High Fixed Transfer Costs scenario.
- Table B.5.2. Percentages of observations in alternative spatial price differential regimes for the Stochastic Transfer Costs scenario.
- Table B.5.3. Percentages of observations in alternative spatial price differential regimes for the High Stochastic Transfer Costs scenario.
- Table B.5.4. Percentages of observations in alternative spatial price differential regimes for the Endogenous Transfer Costs scenario.
- Table B.5.5. Percentages of observations in alternative spatial price differential regimes for the Small Sample scenario.
- Table B.5.6. Percentages of observations in alternative spatial price differential regimes for the High Demand Elasticity scenario.
- Table B.5.7. Percentages of observations in alternative spatial price differential regimes for the Different Demand Elasticities scenario.
- Table B.5.8. Percentages of observations in alternative spatial price differential regimes for the Stochastic Supply with Low Autocorrelation scenario.
- Table B.5.9. Percentages of observations in alternative spatial price differential regimes for the Stochastic Supply with High Autocorrelation scenario.

Table B.6.1. Percentages of observations in alternative trading regimes for the High Fixed Transfer Costs scenario.

Table B.6.2. Percentages of observations in alternative trading regimes for the Stochastic Transfer Costs scenario.

Table B.6.3. Percentages of observations in alternative trading regimes for the High Stochastic Transfer Costs scenario.

Table B.6.4. Percentages of observations in alternative trading regimes for the Endogenous Transfer Costs scenario.

Table B.6.5. Percentages of observations in alternative trading regimes for the Small Sample scenario.

Table B.6.6. Percentages of observations in alternative trading regimes for the High Demand Elasticity scenario.

Table B.6.7. Percentages of observations in alternative trading regimes for the Different Demand Elasticities scenario.

Table B.6.8. Percentages of observations in alternative trading regimes for the Stochastic Supply with Low Autocorrelation scenario.

Table B.6.9. Percentages of observations in alternative trading regimes for the Stochastic Supply with High Autocorrelation scenario.