WP/10/222

IMF Working Paper

How Commodity Price Curves and Inventories React to a Short-Run Scarcity Shock

Shaun Roache and Neşe Erbil

INTERNATIONAL MONETARY FUND

 WP/10/222

IMF Working Paper

Research

How Commodity Price Curves and Inventories React to a Short-Run Scarcity Shock

Prepared by Shaun Roache and Nese Erbil[1]

Authorized for distribution by Thomas Helbling

September 2010

Abstract

How does a commodity market adjust to a temporary scarcity shock which causes a shift in the slope of the futures price curve? We find long-run relationships between spot and futures prices, inventories and interest rates, which means that such shocks lead to an adjustment back towards a stable equilibrium. We find evidence that the adjustment is somewhat consistent with well-known theoretical models, such as Pindyck (2001); in other words, spot prices rise and then fall, while inventories are used to absorb the shock. Importantly, the pace and nature of the adjustment depends upon whether inventories were initially high or low, which introduces significant nonlinearities into the adjustment process.

JEL Classification Numbers: D53, G12

Keywords: Financial markets, Commodity markets, Asset pricing

Author's E-Mail Address: sroache@imf.org, nerbil@imf.org.

[1] We would like to thank Thomas Helbling and Research Department seminar participants at the IMF for useful comments and suggestions. The usual disclaimer applies.

I. INTRODUCTION

The relationship between commodity spot and futures prices reflects, in part, the perception of short-term physical scarcity and the prevailing level of inventories. The slope of the futures curve, measured here as the difference between the price of a futures contract at some given maturity and the spot price, can thus provide information on whether market participants anticipate relative abundance (an upward sloping curve) or scarcity (a flat or downward sloping curve) in the physical market. Price curves also provide incentives for market participants to change their exposure to commodity prices.

A greater understanding of futures price curves and inventory dynamics can help market participants plan their responses to supply and demand shocks. It may also enrich the information that can be obtained from commodities futures markets, providing for a more informed interpretation of price developments. In this context, this paper seeks to fill a gap in the existing literature and asks three key questions: In the event of a short-run shock, is there such a thing as a "normal" commodity market back towards which spot and futures prices and inventories adjust over time? How does a commodity market adjust to a temporary scarcity shock which moves the price curve away from this equilibrium? How quickly do inventories and prices adjust following such shocks?

Our interest is in temporary shocks to physical market balances which cause changes in the futures curve slope; in contrast, permanent or long-lasting shocks should have an impact across the futures curve, which may change the level of the curve but leave the slope little changed. The slope of the futures curve can change for one of three reasons: a change in interest rates; a change in physical storage costs; or a change in the market's perception of short-term scarcity and a compensating adjustment in the utility afforded by holding inventories. (Shifts in the risk premium afforded by holding commodity futures or in expected future spot prices should lead to a shift across the curve.)

Large changes in the futures curve slope are rarely caused by the first two explanations. Although both interest rates and storage costs can move significantly over time, the very large discrete changes required to steepen or flatten futures curves sharply and rapidly are unlikely. Our analysis thus pertains to the effects of changes in actual or expected scarcity over short horizons. In most cases, these shocks will reflect actual or expected supply disruptions, as most demand shocks exhibit a higher degree of persistence.

We compare the adjustment of spot and futures prices and inventories to the predictions from a theoretical model (Pindyck, 2001). We also focus on the possible existence of asymmetries in market reactions to temporary shocks. A strong clue about the nature of base metal price adjustment to shocks in different market states (usually defined by the inventory cycle) is provided by the interest rate-adjusted difference between spot and futures prices (which is a measure of the slope of the futures curve). As Figure 1 shows for six metals, this variable appears to be stationary over time. There also appears to be a degree of nonlinearity when the variable deviates from its average value, with large but short-lived spikes higher (backwardation) coexisting with more sustained, yet less dramatic, declines (contango).

Figure 1. Base Metals: Interest rate-adjusted Basis and Inventories 1/

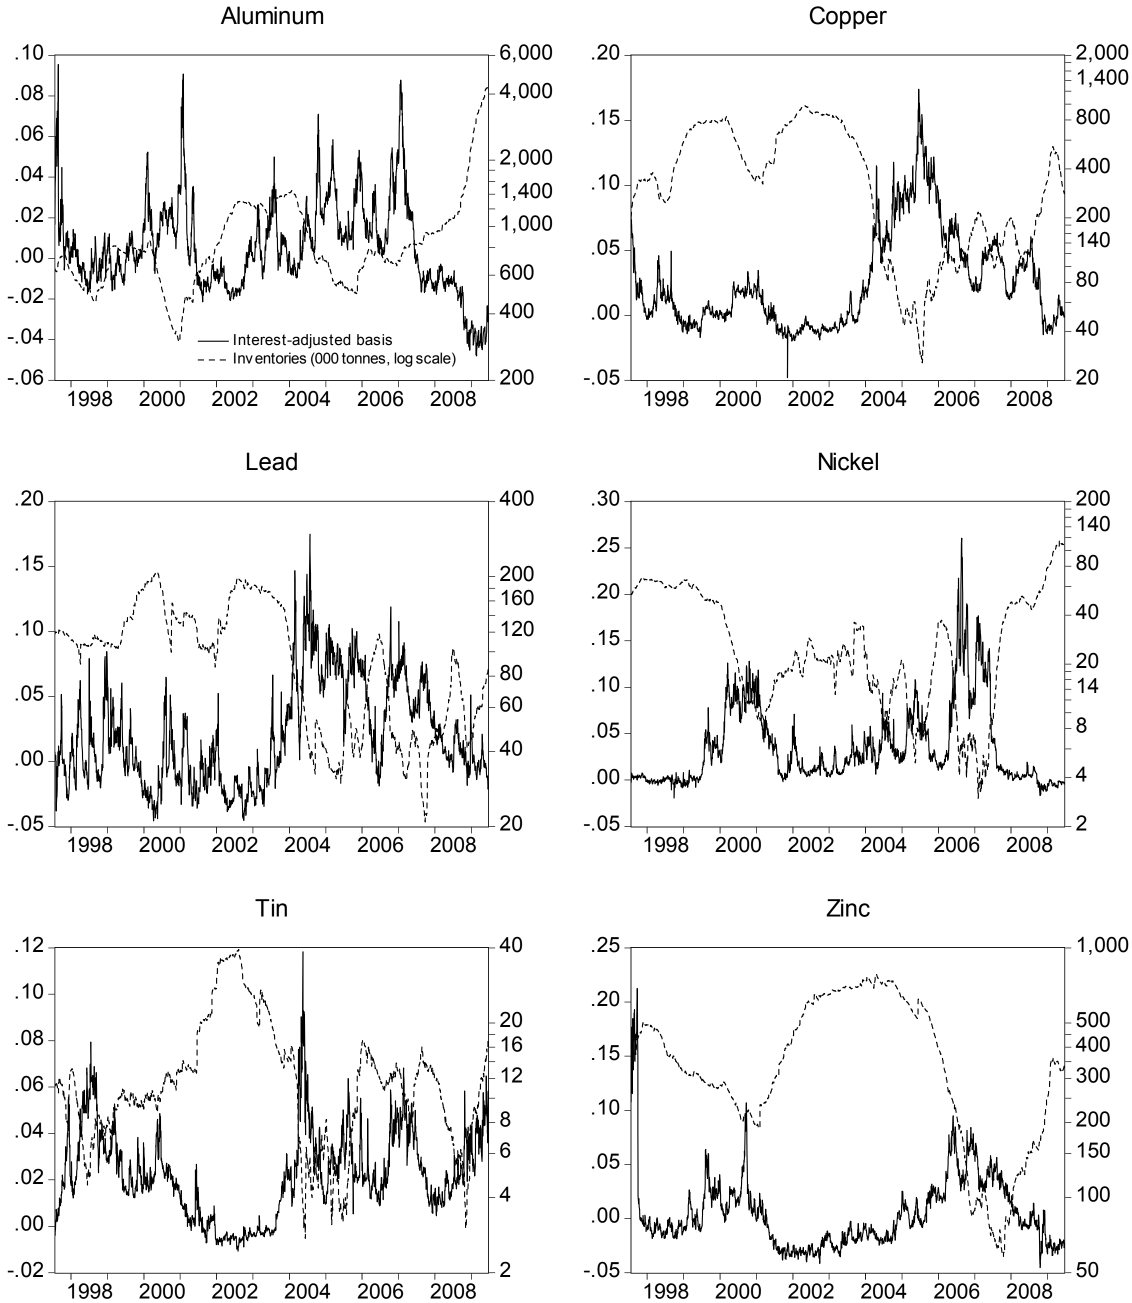

Sources: London Metal Exchange, Bloomberg L. P. and authors' estimates.

1/ The interest-adjusted basis is calculated as $s(t) - f(t,T) + r(t,T)$, where s is the log spot price, $f(t,T)$ is the log price of a futures contract at maturity T at period t, and $r(t,T)$ is the interest rate for over the period $T - t$.

Theory also guides us to expect asymmetric adjustment in the relationship between these variables. To address possible nonlinearity in commodity market adjustment, we apply self-exciting threshold vector autoregression models, in which the thresholds are determined by the nature of the disequilibrium; in practice, this is largely determined by the slope of the futures curve itself, which often reflects current and expected levels of inventories, and is thus closely related to the inventory cycle.

Many earlier empirical studies of the relationship between commodity spot and futures prices ignored non-stationarity, even though it is now generally accepted that commodity prices are I(1) processes, at least over shorter-run samples. More recently, research has focused on the existence of cointegrating relationships between spot prices, futures prices, and inventories; see for example Heaney (1998), Watkins and McAleer (2006), and Crompton and Xiarchos (2008). The studies closest to ours include McMillan (2005) and Kouassie (2008). McMillan (2005) applied an asymmetric threshold model to three base metals and found nonlinear cointegrating relationships between spot and futures prices, with more rapid adjustment towards equilibrium occurring when futures prices are greater than spot prices. However, interest rate and inventory levels were not included in this model. Kouassie (2008) examined the interaction between inventories and prices of six metals by using threshold models and found cointegration with asymmetric adjustment between prices and inventories.

This paper provides a technical contribution to the literature in a number of ways. First, we adopt a comprehensive self-exciting threshold approach which includes both inventories and interest rates along with spot and futures prices, building on previous work which has omitted one or more of these variables. Second, we compare the results from an empirical model to the predictions of a widely accepted theoretical model. Finally, we use higher frequency daily data, which should provide important insights given the relatively high liquidity and rapid adjustment patterns of the major commodity markets. Until now, there have been few studies on the interaction between inventories and commodity prices at a daily frequency, partly due to poor and infrequent data. Base metals, which include aluminum, copper, lead, nickel, tin and zinc, provide the richest data set, particularly as trading of spot, futures, and options is concentrated on the London Metals Exchange (LME).

The plan of the paper is as follows. In the next section, we will discuss the methodology and set up the model to be tested. Then we will discuss the data in section III and present and discuss the results in section IV and finally conclude in section V.

II. Methodology

Spot and futures price arbitrage for financial assets

Before analyzing the dynamics of market adjustment, it is worthwhile reviewing the theory behind the relationship between spot and futures prices. In particular, the well-known arbitrage condition that determines the relationship between spot and futures prices for financial assets rarely holds for commodities. The role of commodities as consumption and processing goods, and the pivotal importance of physical inventories, among other factors, lead to a more complex and dynamic relationship. To elaborate, we first present the cost-of-carry relationship for a financial asset, ignoring coupon or dividend payments, in a market without frictions. This states that the price of a futures contract at time *t* which specifies delivery at $T > t$ denoted by *F(t,T)* is equal to the current spot price *S(t)* multiplied by the continuously compounded interest rate *r* for the period *t* to *T*:

$$F(t,T) = S(t)\exp[r(T-t)] \tag{1}$$

Commodity market arbitrage between spot and futures
This relationship tends not to hold for commodities, for two main reasons. Spot and futures prices must take into account the costs of holding physical inventory, e.g., warehousing and insurance, which increases the "cost of carry" (which for financial assets includes only the interest rate). Also, market participants may hold physical inventory of a commodity for its value as a consumption good, rather than as a financial asset. The benefit that accrues to the inventory holder is often referred to as the "convenience yield".

We incorporate the physical storage costs, denoted by *k,* as a constant proportion of the spot price and this serves to create a small and—assuming that storage costs do not vary too much—a fairly stable wedge between the two sides of equation (1). The inclusion of the convenience yield for the marginal unit of inventory, denoted by ψ, leads to more profound changes. A number of theoretical models indicate that there should be a strong and non-linear relationship between the current and expected future level of inventories, which we denote by *N*, and the value of ψ. This nonlinearity reflects a declining marginal utility of inventories. In particular, as the level of current and expected future inventories falls, the probability of experiencing a physical "stock out" increases, and ψ should rise, at an increasing rate as inventories fall towards their zero bound (e.g., Deaton and Laroque, 1992 and Williams and Wright, 1992). Stock outs can be very costly for the producers and users of the physical commodity as it can interfere with production and customer delivery schedules. This nonlinear relationship between convenience yield and current and future inventory levels (denoted by *N*) is summarized by Figure 3. This shows that the effect of a given change in inventories on the convenience yield is dependent upon the starting level of inventories:

Figure 2. Commodity Convenience Yield and Current and Expected Future Inventories

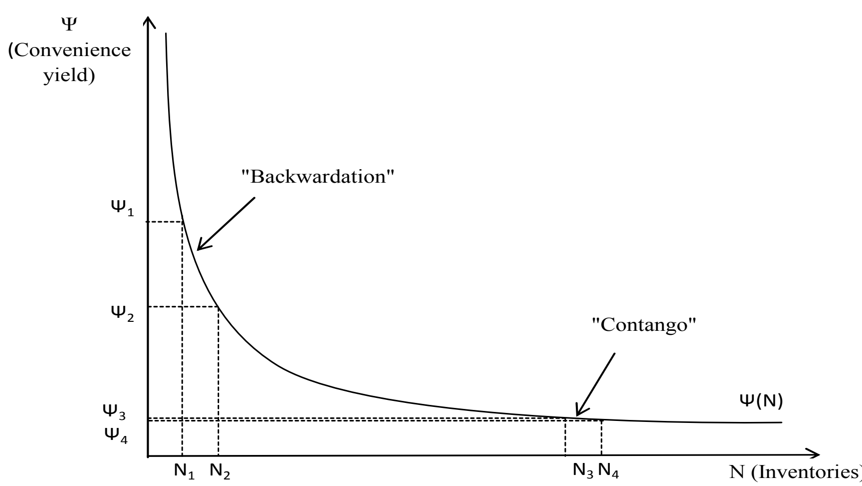

Source: Pindyck (2001) and Authors.

Incorporating these two features of commodity markets, storage costs and marginal convenience yields, into equation (1) obtains an arbitrage condition that is written as:

$$F(t,T) = S(t)\exp[(r + k - \psi(N))(T - t)] \tag{2}$$

As described by Markert and Zimmermann (2006), these features of the commodity markets imply that there exists an arbitrage "upper bound" for commodity futures, when expressed in the form of standard financial asset interest rate arbitrage condition:

$$F(t,T) \le S(t)\exp[r(T - t)] \qquad \text{where} \qquad \psi(N) - k \ge 0 \tag{3}$$

Price and inventory adjustment in backwardation

Equation (3) shows that current and expected future inventories play a key role in determining the shape of the futures curve. When inventories *N* are very low or expected to decline significantly and the probability of a physical stock out is relatively high, the marginal convenience yield ψ will be high. In other words, commodity producers, consumers, and processors will value highly the marginal unit of physical inventory. In this case, the futures curve may be in "backwardation", with spot prices *S* higher than futures prices *F*. (When the spot price is higher than the futures price, the market is defined as "strongly backwardated". When the spot price is below the futures price, but equation (3) holds as a strict inequality—i.e. the spot price is higher than the discounted value of the futures price—the market is defined as "weakly backwardated".)

In a backwardated market, it is often assumed that the sensitivity of the marginal convenience yield net of storage costs (ψ-k) to changes, or expected changes, in inventories is very high (i.e. from Figure 2 we can see that $\partial\psi/\partial N$ is large). Starting from an initial condition of a tight physical market, with the inventory cycle near its low point (consistent with a backwardated market), small changes in expectations regarding the future path of inventories should have very large effects on the shape of the futures curve. In the case of a temporary supply shock which causes a large rise in the convenience yield and the spot price, expectations of a return to more "normal" physical market conditions should lead to a rapid reversal in these moves.

Price and inventory adjustment in contango

When inventories are abundant and, conditional on demand projections, the probability of a physical stock out is low, the net marginal convenience yield (ψ-k) will be very low. In this case, assuming that physical storage costs are not too large, the spot price *S* will be lower than the futures price *F* and the futures curve will be in contango. As ψ reaches its zero lower bound, then interest rate arbitrage forces will limit the steepness of the curve. In particular, if *S* is lower than the discounted value of *F*, then the incentive and capacity to place an interest rate-based arbitrage trade will exist. Arbitrageurs will be able to take a long position in the spot market with a cost-of-carry equal to $r + k$ and take an offsetting short position in the futures markets, earning risk-free profits.[2]

In a contangoed market, the sensitivity of the net marginal convenience yield to changes, or expected changes, in inventories is very low (i.e. from Figure 2 we can see that the first derivative $\partial\psi/\partial N$ is small). Starting from an initial condition of a well-supplied physical market, with the inventory cycle near its high point (consistent with steep contango), even large changes in expectations regarding the future path of inventories should have relatively little effect on the shape of the futures curve. In this case, a temporary supply disruption would have only a modest effect on convenience yields and spot prices, as inventories would be sufficient to absorb the shock. In turn, the adjustment back towards equilibrium should also be more gradual.

A Nonlinear Price-Inventory Adjustment Model

In this section, we will propose an empirical model for this system of variables. The endogenous variables in our system, spot and futures prices and inventories, are all jointly determined and reflect current physical market conditions, but also expectations for the future. A natural question to ask is whether these variables share a stable long-run relationship; to put it differently, is there a futures market curve and level of inventories that together reflect "normal" or steady-state market conditions? During a normal market, it might be presumed that inventories, or stock-use ratios, are close to their average levels and the futures market curve reflects a steady-state perspective on the outlook, in particular with regard to the evolution of inventories in future periods. In other words, the system is anchored over the long-run by a steady-state level of inventories.

We can test the hypothesis of long-run equilibrium existence by assessing evidence for a long-run cointegrating relationship between the endogenous variables S, F, and N and the exogenous variable that theory suggests should also determine curve slope, the interest rate r. We can write such a hypothesized relationship using the log-levels of each of these variables (with the exception of interest rates which are in levels) and normalizing with respect to the spot price s as:

$$s_t = \beta_1 + \beta_2 f_{t,T} + \beta_3 n_t + \beta_4 r_{t,T} + z_t \tag{4}$$

In (4), s_t is the log of the spot price, $f_{t,T}$ is the log of the futures price, n_t is the log of inventories, and r_t is the interest rate level. The constant β_1 can be interpreted approximately as a mean net marginal convenience yield and the unconditional expectation of z is zero. The β_1 can only be interpreted as an approximate mean net marginal convenience yield due to the presence of inventories on the right-hand side of equation (4).

[2] Mabro (2009) provides an example in oil markets where ample oil supplies in August 1997 and in 2008 moved the term structure of futures prices into a very steep contango. The increasing differential between spot and futures contracts gave sufficient incentives for traders to buy physical oil to add to inventories and sell a futures contract. This resulted in an inventory build-up subsequently pressuring prices flattening the term structure in 1998 and 2009, respectively.

How can we interpret the residual z? Comparing (4) with (2), it can be seen that the residual z is closely related to the concept of the net marginal convenience yield and, in our specification with current inventories as an explanatory variable, it approximately represents deviations from the average convenience yield. Because we have included current inventories on the right-hand side as an endogenous variable, this deviation in convenience yield arises from unobserved variables, principally expectations of future inventories and spot price volatility. The absolute values of the estimated coefficients may also diverge from unity, which means that a direct comparison with convenience yields cannot be made.

Equation (4) has implicitly assumed that the convenience yield over the long-run is a linear function of inventories; the nonlinearity in our model will instead emerge from the short-run dynamics of adjustment back to equilibrium. In particular, we test whether the initial level of z captures information regarding expected inventories and whether this in turn will affect how z converges back to zero. In other words, z is the self-exciting threshold process in this model.

Empirical specification of the nonlinear adjustment

One way to model the potential nonlinearity of the process $\{z\}$ from equation (4) is suggested by Figure 2; that is, to follow Martens, Kofman, and Vorst (1998) and hypothesize that it might be described by a threshold autoregressive model, with the speed of adjustment determined by lagged values, such as:

$$z_t = \phi_0^{(j)} + \sum_{i=1}^{p} \phi_i^{(j)} z_{t-i} + \varepsilon_t^{(j)} \qquad k_{j-1} < z_{t-d} \leq k_j \tag{5}$$

In this model, j indexes the thresholds separating regimes, d is the lag of z which determines the threshold, and k denotes the value of the thresholds. This specification provides a linear approximation for the adjustment process in each regime; in other words, in each regime it assumes a different linear adjustment process, characterized by varying autoregressive parameters ϕ and speeds of adjustment.

One of the hypotheses we wish to test is whether there is a band around the equilibrium (i.e. a "middle regime") which is characterized by a random walk and two outer regimes in which the process converges back towards equilibrium. This is a common finding for financial assets and is often interpreted as a mispricing that is not sufficiently large and profitable to arbitrage away due to transaction costs (e.g., Martens, Kofman, and Vorst, 1998).

A priori, we can only hypothesize about the value of the thresholds and the nature of the regimes; the thresholds will be estimated from the data. For now we assume that the forward curve process summarized by $\{z\}$ from equation (4) is subject to two thresholds and three regimes which are determined by the value of z itself (i.e., the threshold process is self-exciting). We hypothesize that these regimes are: (i) a lower regime, in which z is negative and the forward curve is upward sloping and steep (contango); (ii) a middle regime in which z is close to zero and the system is close to equilibrium; and (iii) an upper regime in which z

is high and the forward curve is either relatively flat or inverted (backwardation). We can now write this model as:

$$z_t = \begin{cases} \phi_0^{(l)} + \phi_1^{(l)} z_{t-i} + \varepsilon_t^{(l)} & z_{t-1} < \underline{z} \\ \phi_0^{(m)} + \phi_1^{(m)} z_{t-i} + \varepsilon_t^{(m)} & \underline{z} \le z_{t-1} \le \overline{z} \\ \phi_0^{(u)} + \phi_1^{(u)} z_{t-i} + \varepsilon_t^{(u)} & z_{t-1} > \overline{z} \end{cases} \tag{6}$$

Where each of the regimes is listed respectively.

Applying a Threshold Vector Error-Correction Model

The system described by (6) allows us to test formally for the existence of nonlinearities, but it does not provide any guidance related to which variable (spot or futures prices or inventories) takes the burden of adjustment when the system deviates from equilibrium. A natural way to understand these features is by applying a threshold vector error model (T-VECM). In this model, global behavior is defined by the cointegrating vector and its residual *z*. Local behavior (or short-run dynamics) is described by the adjustment coefficient on the cointegrating vector and the coefficients on the lagged first-differences. This model can then be written as:

$$\Delta \mathbf{X}_t = \mathbf{A}_0^{(j)} + \sum_{k=1}^{K} \mathbf{A}_k^{(j)} \Delta \mathbf{X}_{t-k} + \boldsymbol{\beta}^{(j)} z_{t-1} + \sum_{l=0}^{L} \mathbf{C}_l^{(j)} r_{t-l} + \mathbf{E}_t^{(j)} \tag{7}$$

where: the superscript j indexes the regime, **X** is the (3x1) vector of endogenous variables, including the spot price, futures price, and inventories; Δ is the first difference operator; **β** is the (3x1) vector of adjustment coefficients; z_{t-1} is the lagged value of the variable described in equation (4); **C** is a (3x1) vector of coefficients on the exogenous interest rate r; and **E** is the (3x1) vector of reduced form residuals.

III. DATA

The source of the base metals spot prices, futures prices, and inventories data is the London Metal Exchange (LME) which accounts for the largest share of trading in base metal spot and futures markets. For example, about 95 percent of the total world trade in copper futures occurs through the LME. The LME also provides storage facilities to make it possible for market participants to take or make physical delivery of metals. For futures prices we use two different contract maturities: three and six months. The interest rates are the three-month and six-month London interbank offered rates (Libor), as calculated by the British Bankers' Association. The sample period spans from July 23, 1997 to June 19, 2009. The start date is determined by the availability of all the data series at a daily frequency. Summary statistics are provided in Table 1.

Table 1. Base Metal Price and Inventory Series: Summary Statistics

	Levels				First differenced logs			
	Mean	Standard deviation	Skew	Kurtosis	Mean	Standard deviation	Skew	Kurtosis
Aluminum								
spot price	1,787	526	1.0	2.7	0.00	0.01	-0.30	5.92
futures price (3-month)	1,804	529	1.1	2.7	0.00	0.01	-0.30	6.28
futures price (6-month)	1,811	529	1.1	2.8	0.00	0.01	-0.38	6.42
inventories	935,013	597,629	3.3	15.9	0.00	0.01	4.63	47.40
Copper								
spot price	3,368	2,309	1.1	2.7	0.00	0.02	-0.12	7.84
futures price (3-month)	3,350	2,280	1.1	2.7	0.00	0.02	-0.13	8.01
futures price (6-month)	3,316	2,241	1.1	2.7	0.00	0.02	-0.13	8.40
inventories	402,435	285,122	0.4	1.8	0.00	0.01	5.18	91.29
Lead								
spot price	976	739	1.8	5.9	0.00	0.02	-0.18	6.84
futures price (3-month)	975	732	1.8	5.8	0.00	0.02	-0.21	7.53
futures price (6-month)	966	721	1.9	5.8	0.00	0.02	-0.25	8.12
inventories	101,671	50,669	0.3	1.9	0.00	0.02	10.04	240.87
Nickel								
spot price	13,515	10,151	1.7	5.6	0.00	0.02	-0.12	6.97
futures price (3-month)	13,367	9,817	1.6	5.2	0.00	0.02	-0.16	7.09
futures price (6-month)	13,123	9,449	1.6	5.0	0.00	0.02	-0.16	7.16
inventories	32,674	23,950	1.0	3.6	0.00	0.02	1.19	23.95
Tin								
spot price	8,022	4,522	1.6	5.1	0.00	0.02	-0.09	11.81
futures price (3-month)	8,011	4,511	1.6	5.1	0.00	0.02	-0.12	12.13
futures price (6-month)	7,979	4,478	1.7	5.2	0.00	0.02	-0.13	12.16
inventories	13,069	7,988	1.6	5.1	0.00	0.03	2.96	38.18
Zinc								
spot price	1,493	896	1.6	4.6	0.00	0.02	-0.37	7.19
futures price (3-month)	1,502	887	1.6	4.5	0.00	0.02	-0.32	7.15
futures price (6-month)	1,500	862	1.6	4.4	0.00	0.02	-0.30	7.04
inventories	385,746	208,802	0.3	1.8	0.00	0.01	6.21	92.51

Sources: London Metal Exchange; authors' calculations

As the initial step of our analysis, we assess the order of integration of the endogenous variables and interest rates using the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root tests. There are no theoretical reasons to suggest that the endogenous variables should be stationary. Interest rates should be expected to remain bounded over the very long run, but it is widely reported in the literature that interest rates follow integrated processes, a result which may be due in part to the low power of standard unit root tests in small samples (see Wu and Zhang, 1997). Overall, we were unable to reject the null that all of the log levels of each of these series, including interest rates, contain a unit root at standard levels of confidence, although there was clear evidence of stationarity for the first-differenced data (see Appendix Table A1). These results are consistent with those for base metal data as presented in McMillan (2005), Kouassie (2008), and Watkins and McAleer (2006).

IV. Estimation and Results

Ideally, testing for cointegration and thresholds could be achieved using a single, consistent approach. However, as pointed out by Balke and Formby (1997), the threshold variable itself is determined by the cointegrating vector, which itself must be estimated. In other words, some of the alternative hypotheses have "nuisance" parameters—namely the thresholds—that do not form part of the null of no cointegration/ linearity, which results in a nonstandard inference problem. Consequently, we follow their suggestion to approach the analysis in two stages: first, an assessment of global behavior, with tests for cointegration; and second, assessing the local behavior of the time series, which tests for nonlinearity.

A. Testing for Cointegration

An inspection of Figure 2 suggests that a reasonable starting point is to assume that the spot price, futures price, relevant financing interest rate, and possibly the level of inventories, are cointegrated, in the form of equation (4). Although the relationship may change based on the shape of the futures curve—which could influence the local behavior of the equilibrium error z—the system should be eventually drift back towards equilibrium. There is no well developed model that would suggest inventories should share a long-run linear relationship with the other variables; rather, a short-run nonlinear relationship is often suggested. For now, we include inventories in all of the cointegration tests to ensure our specification does not suffer from omitted variable bias.

We test for cointegration using the Philips-Perron and Engle-Granger tests based on the residuals from equation (4). Although the Johansen (1988) test has become the standard procedure for multivariate systems, Balke and Formby (1997) present evidence that this procedure may have particularly low power for asymmetric systems when compared to the Philips-Perron test. Enders and Siklos (2001) construct a direct test for asymmetry in a cointegrated system, but they acknowledge that this also suffers from particularly low power for a standard Threshold Autoregressive Model (TAR) model.

Strong evidence in favor of cointegration

In fact, regardless of the tests used, we find strong evidence for cointegration (Appendix Tables A2). The estimates of equation (4) obtain coefficients that, in most cases, are statistically significant and close to what theory predicts based on equation (2) (Table 2). For example, after normalizing on the spot price the coefficients on the futures prices are close to one. The coefficients on interest rates are negative, but for the most liquid contracts (aluminum and copper) they are greater than one in absolute value. Inventory coefficients are negative, indicating that higher physical stocks are consistent with lower spot prices (keeping all else constant), although the economic significance of these coefficients is very low.

The constants in all cases (β_1 from equation (4)), except zinc are positive, which means the futures curve is flatter (steeper) when in contango (backwardation) than standard interest rate arbitrage relationships would predict. This reflects the constant component of the non-zero average convenience yield (net of storage costs) which is independent of inventories; even when inventories are high, there remains a non-zero probability of stockouts and together with other utilities obtained from holding physical stocks.

Table 2. Long-Run Cointegrating Relationships Between the System Variables
(Significance at the 1, 5, and 10 percent levels denoted by ***, **, and * respectively)

	Aluminum	Copper	Lead	Nickel	Tin	Zinc
			Three-month			
Constant	0.13 ***	0.31 ***	0.51 ***	0.14 ***	0.12 ***	-0.23 ***
Futures	1.01 ***	0.99 ***	0.99 ***	1.00 ***	1.00 ***	1.02 ***
Inventories	-0.01 ***	-0.02 ***	-0.04 ***	-0.02 ***	-0.01 ***	0.01 ***
Interest rate	-1.44 ***	-1.12 ***	-1.16 ***	0.61 ***	-0.76 ***	0.64 ***
			Six-month			
Constant	0.30 ***	0.70 ***	0.98 ***	0.28 ***	0.20 ***	-0.58 ***
Futures	1.02 ***	0.98 ***	0.98 ***	1.01 ***	1.00 ***	1.05 ***
Inventories	-0.03 ***	-0.04 ***	-0.07 ***	-0.04 ***	-0.02 ***	0.02 ***
Interest rate	-1.38 ***	-0.97 ***	-0.93 ***	0.61 ***	-0.78 ***	0.91 ***

B. Testing for and Locating Thresholds

Does the adjustment depend on the initial slope of the curve? In particular, we are interested in the speed with which the market adjusts to scarcity shocks, and which variables take the burden of adjustment. To achieve this, we implement the ordered autoregression described by Tsay (1989) in which the cases are arranged according to the values of a particular regressor.

To identify the appropriate autoregressive structure, we assessed both partial autocorrelation functions and information criteria and find that the AR process is strongly determined by the first lag (see Appendix Tables A4 and A5). This is not surprising given that commodity markets assimilate all new information quickly, weakening the influence of lags greater than one period. As a result, we assume no delay and a simple AR(1) process, which involves ordering the regressions according to the values of last period's equilibrium error, or z_{t-1}. This allows us ensure that the observations in a particular group follow the same AR process, conditional upon accurate identification of the thresholds.[3]

Strong evidence in favor of nonlinearity
We conduct two tests for non-linearity of the ordered autoregressions: the Andrews-Quandt breakpoint test and the Tsay (1989) test.[4] We find clear evidence of at least one break in all

[3] We also ran the threshold identification procedure for AR(3) processes, consistent with Aikaike information criteria, and found that the results were mostly identical (or very close) to those obtained from an AR(1).

[4] We supplement this approach less formal analysis of AR(1) coefficient t-ratio scatter plots, obtained from recursive least squares regressions. This yielded less clear-cut conclusions, but tended to support the number and location of the thresholds obtained from the formal methods described. Details are available on request from the authors.

of the commodity relationships and proceeding iteratively, we find evidence for two or more breaks (see Appendix for detailed results). The evolution of the equilibrium error over time, and the location of the thresholds is shown in Figure 3 for the six-month relationships. These breaks indicate that the speed of adjustment back towards equilibrium is conditional upon the size and sign of the deviation itself (Table 3).

The threshold locations are similar across each of the metals and correlated significantly with the slope of the futures curve (Table 3). When the equilibrium error *z* is below the lower threshold, the market is very likely to be in contango, with the spot price below the futures price. In contrast, when *z* is above the upper threshold, the market is much more likely to be in backwardation, with spot prices above futures prices.

We do not find evidence of a "no arbitrage" band in metals markets similar to that typically found for financial assets. The autoregressive structure of *z* is stable in all three regimes and the speed of convergence is typically faster in the middle regime which is closest to equilibrium. This indicates that transaction costs do not play a significant role in the commodity market adjustment. For the remainder of the analysis, we continue to divide the adjustment paths based on these three regimes, however. Assessing whether adjustment to temporary shocks is significantly different when the market is initially close to equilibrium will provide an important insight into the nature of commodity market dynamics.

Table 3. Threshold Values and Curve Slope: Percent of Time in Contango 1/

	When equilibrium error *z* is:			
	below lower threshold	between upper and lower	above upper threshold	Correlation of z with curve slope
Aluminum	100.0	94.7	22.1	0.91
Copper	74.5	59.4	48.1	0.42
Lead	88.3	69.9	13.8	0.65
Nickel	56.9	67.1	22.7	0.64
Tin	76.7	48.9	0.0	0.77
Zinc	94.6	88.2	46.9	0.89
Aluminum	96.6	55.6	8.2	0.85
Copper	65.3	56.8	42.7	0.36
Lead	73.2	35.3	3.2	0.61
Nickel	53.5	6.1	24.7	0.55
Tin	69.0	4.8	0.0	0.71
Zinc	88.3	77.7	48.8	0.76

1/ Defined as the futures price at the given maturity being above the spot price.

Figure 3. Deviations from Equilibrium and Identification of the Thresholds

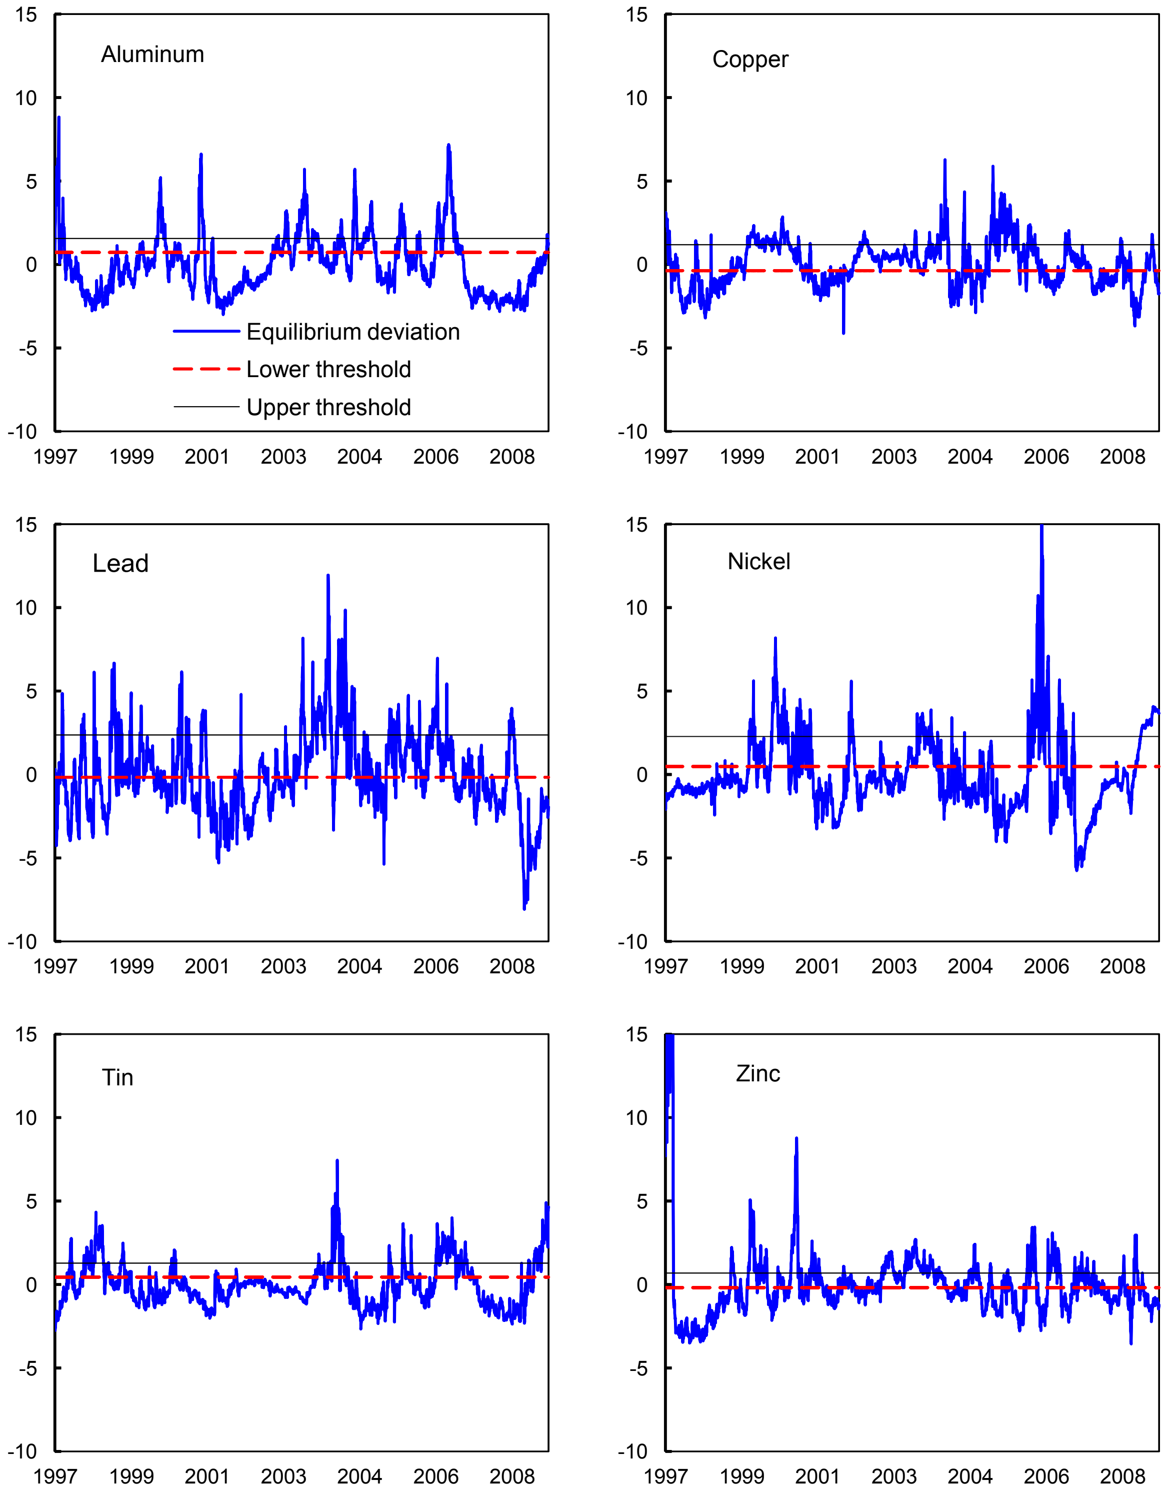

Source: Authors’ calculations.

C. Adjustment to temporary shocks

We assess how the system of spot and futures prices and inventories responds to temporary shocks by estimating separate VECM models for each of the three regimes. The cointegrating vector for each commodity (i.e., the coefficients on the level variables in the VECM) is the same in each regime, with the estimates taken from equation (4) we used to test for cointegration. We proceed to estimate the system (7), using the optimal lag lengths identified in Appendix Table A6. The results that follow use the six-month futures contract; the results with the three-month contract are qualitatively similar but less pronounced.

Describing the curve shock—a short-term spot price shock

Our main interest is in the behavior of the three variables when the relationship between them is in a disequilibrium caused by a temporary shock in the physical market; to create this condition, we apply a shock to the slope of the futures curve for the VECM in each regime. As discussed above, the most likely cause of a sharp and rapid change in the slope of the futures curve is a shift in actual or perceived short-term physical scarcity and a corresponding change in the marginal convenience yield. Interest rates and storage costs, the other two factors which explain the gradient, are unlikely to experience discrete jumps sufficient to match the observed volatility in curve slopes.

Why should the change in scarcity premium be confined to the short-term? If expectations of changing long-term scarcity emerged, then we should expect to see a permanent change in spot prices and, as a result, a shift across the entire futures curve. This would leave the slope constant (or at least little changed) and the system would remain in, or very close to, equilibrium.

The implication of these arguments is that these short-term futures curve shocks are most likely characterized by spot price shocks. Futures prices will be anchored by expectations that the supply shock should eventually dissipate and that market participants will be able to smooth the adjustment over time, in part by managing their inventories. Ideally, we would want to impose a structural spot price shock to the VECM system. One method would be to impose a Choleski ordering on the system, but this approach runs into some difficulties due to the challenges in disentangling spot and futures price shocks. The contemporaneous correlation between log changes in spot and futures prices is very high (above 0.9 for all metals), while the correlation when one of the price changes is lagged one period is very low and statistically insignificant. The correlation between the reduced form VECM residuals for the spot and futures price equations is also above 0.9 for all metals.

Alternative restrictions are suggested by theory and the nature of the data. For the system with four variables, we require at least $(n^2-n)/2 = 6$ restrictions. The first set of restrictions we apply is that interest rates are exogenous to all other variables in the system. The second set of restrictions is that there are no contemporaneous effects from inventories to prices (or vice versa). Inventory data, which may affect prices, are only available from the LME with a one day lag, while movements in physical inventories are unlikely to respond to price signals during the same day, in large part due to logistical constraints. One final restriction that we apply is that the futures curve moves in parallel in response to a futures price shock; in other words, the contemporaneous coefficient on the futures price change in the spot price equation is 1. The justification for this restriction is that the futures price can rise or fall as a result of changes in either the expected future spot price or the risk premium, which compensates the holder of the futures contract for holding the exposure to commodity price volatility. Arbitrage then links the spot price to the futures price, ensuring that these changes are reflected one-for-one in the spot price. For example, if the risk premium declines, leading to higher futures prices, then for unchanged carrying costs and convenience yield, the spot price must also increase by the same amount. This is because at the time of the futures contract's specified physical delivery; there is no difference between holding the spot or the future. Today's spot price can then be discounted back from the futures price by the carry cost and the convenience yield.

We apply a 1 percent positive shock to the spot price which, given the restrictions described above, implies shocks to the reduced form residuals in the spot and futures price equations of *1/(1-b)* and *b/(1-b)*. The parameter *b* is the contemporaneous coefficient of the log change futures price on the log change spot price. For all commodities and all regimes, the estimate of this coefficient from the estimated VECM is between zero and one, implying that futures prices respond positively, but less than one-for-one to spot price shocks. In almost every case, the sensitivity of the futures to the spot price is highest in regime 1 when the curve is upward sloping (average 0.6) and lowest in regime 3, when the typical curve is backwardated (average 0.3). This difference likely reflects the dominant effects of the convenience yield on spot prices in backwardation.

A theoretical framework to assess the empirical results

In this section, we compare the dynamics of adjustment as implied from our empirical approach to those of the theoretical model outlined by Pindyck (2001). Pindyck characterizes commodity market equilibrium as the outcome of interactions in the cash and storage markets. Total demand (denoted by *Q*) in the cash market is a function of the spot price (*P*), other demand shift variables z_Q (e.g., the effect of macroeconomic policies), and random shocks ε_Q (e.g., tastes and technologies). The supply of a commodity in the cash market (denoted by *X*) is also a function of the spot price, other variables affecting supply z_X (e.g., input costs), as well as random shocks ε_X, such as strikes or other unexpected supply disruptions. In equilibrium, net demand, which is the demand for production in excess of consumption, must equal the change in inventories ΔN by identity, so that we can write the cash market equilibrium as:

$$\Delta N_t = X\left(P_t; z_{Xt}, \varepsilon_{Xt}\right) - Q\left(P_t; z_{Qt}, \varepsilon_{Qt}\right) \tag{8}$$

The inverse net demand function can then be written as:

$$P_t = f\left(\Delta N_t; z_{Xt}, z_{Qt}, \varepsilon\right) \tag{9}$$

The inverse net demand function is upward sloping in ΔN; in other words, an increasing rate of inventory accumulation requires higher spot prices to increase supply and reduce demand.

In storage market equilibrium, as described by Figure 2 above, the marginal convenience yield ψ is a function of inventory levels N and other variables, including price volatility σ, future and current consumption rates z_3 and random shocks ε_3. This can be written as:

$$\psi = g\left(N;\sigma, z_{\psi}, \varepsilon_{\psi}\right) \tag{10}$$

Given the values for σ and z_{ψ} equilibrium in the storage market gives ψ_t and N_t. Then given the values and N_{t-1}, z_X, and z_Q, we can find ΔN and solve for P.

What does the model predict in the event of a temporary supply shock? We will assume that a particular metals market is in a steady-state equilibrium with $\Delta N = 0$. Now consider that a the effects of an unanticipated strike at a particularly large mine. This will decrease supply X and cause the net demand function to shift upwards and the spot price to rise (see Figure 6). Because the shock is seen as temporary, inventories will be run down, limiting the increase in spot prices, and the marginal convenience yield will increase. Futures prices will likely rise, but by less than the spot price, which will flatten or even invert the futures curve. Once the strike ends, the net demand curve will shift lower, but until the marginal convenience yield returns back to ψ_0, spot prices will fall but remain above the initial level to ensure that production exceeds consumption and inventories are rebuilt. The futures curve will move back towards its equilibrium slope as spot prices fall by more than futures prices.

Theoretical predictions versus empirical results

How well does this theory predict the effects of short-term supply shocks in metal markets? Figure 6 presents the cumulative impulse responses from the estimated VECM models in each regime. A 1 percent spot price shock leads to a change in the log level of spot prices by more than the initial shock. This is because the spot price contemporaneously affects the futures prices, which in turn has feedback effects on the spot price, and so on. Initially, the curve flattens or inverts as spot prices increase by more than futures in each case, but the dynamics thereafter contrast sharply in each regime.

In many cases, the increase in spot and futures prices is gradually and partially reversed over time, as predicted by Pindyck's model. This pattern is strongest in regimes 2 and 3, where the market started out close to equilibrium or was already in a state of relative short-term scarcity. Inventories also tend to fall, as predicted, albeit more gradually than prices, as market participants run down stocks in response to scarcity in the physical market. In some respects, the empirical results confirm the predictions of the model, but there are three important discrepancies: the behavior of prices in steep contango (regime 1); the permanence of the effects on price and inventory levels of an initial spot price shock; and different outcomes for specific metals.

What can explain the apparent permanent change in spot and futures prices and inventories following a spot price shock? In our framework, we have interpreted an identified spot price shock as the result of a scarcity shock. This is an intuitive approach consistent with the predictions of most theoretical commodity price models for the instant response to a supply disruption. However, our results indicate that some of the effects from a spot price shock are permanent. In particular, in contango for many metals, a spot price shock leads to: a permanent shift higher across the futures curve; a modest flattening in the futures curve, with spot prices relatively higher than futures prices as compared to before the shock; and a compensating permanent decline in inventories. (In the long-run equations, the coefficient on inventories is small, which means that a relatively large decline is accompanied by only a small change in the slope of the futures curve.)

These results suggest that some spot price shocks have a permanent effect, perhaps due to learning over time. An initial supply disruption may, over time, be recognized as a more persistent impairment of supply capacity. Examples might include deteriorating ore quality in well established mines or strikes which persist for months rather than weeks. In these cases, the market would learn gradually about the new supply environment, preventing a decline in prices to the levels which prevailed before the shock. This suggests that alternative identification methods may also be useful in exploring the effects of temporary scarcity shocks, including Blanchard-Quah decompositions.

When the futures curve is steeply upward sloping (regime 1), a spot price shock has relatively large effects on the futures price (albeit less than one-for-one). This means that given the same initial shock to the curve originating in the spot price, the entire curve shifts much higher than in other regimes. This suggests that in a market with abundant inventories and steep contango, markets perceive that spot price shocks are more likely to reflect longer-lasting changes in market conditions.

Figure 4. Effect of Temporary Demand Shock

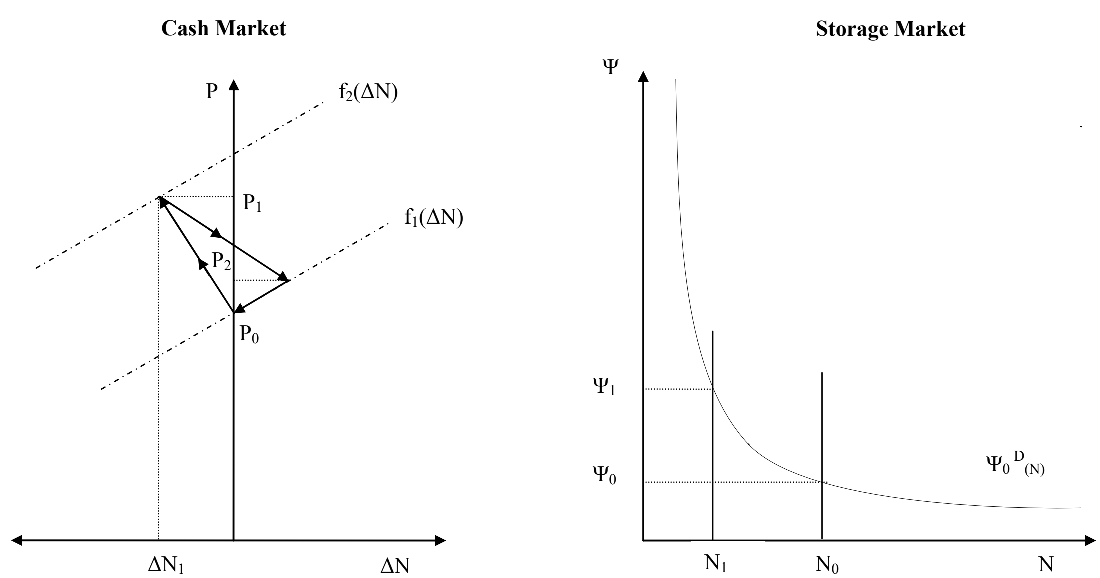

Source: Pindyck (2001)

Figure 5. Impulse Responses from a Spot Price Shock

(equivalent to a one percentage point futures price curve shock)

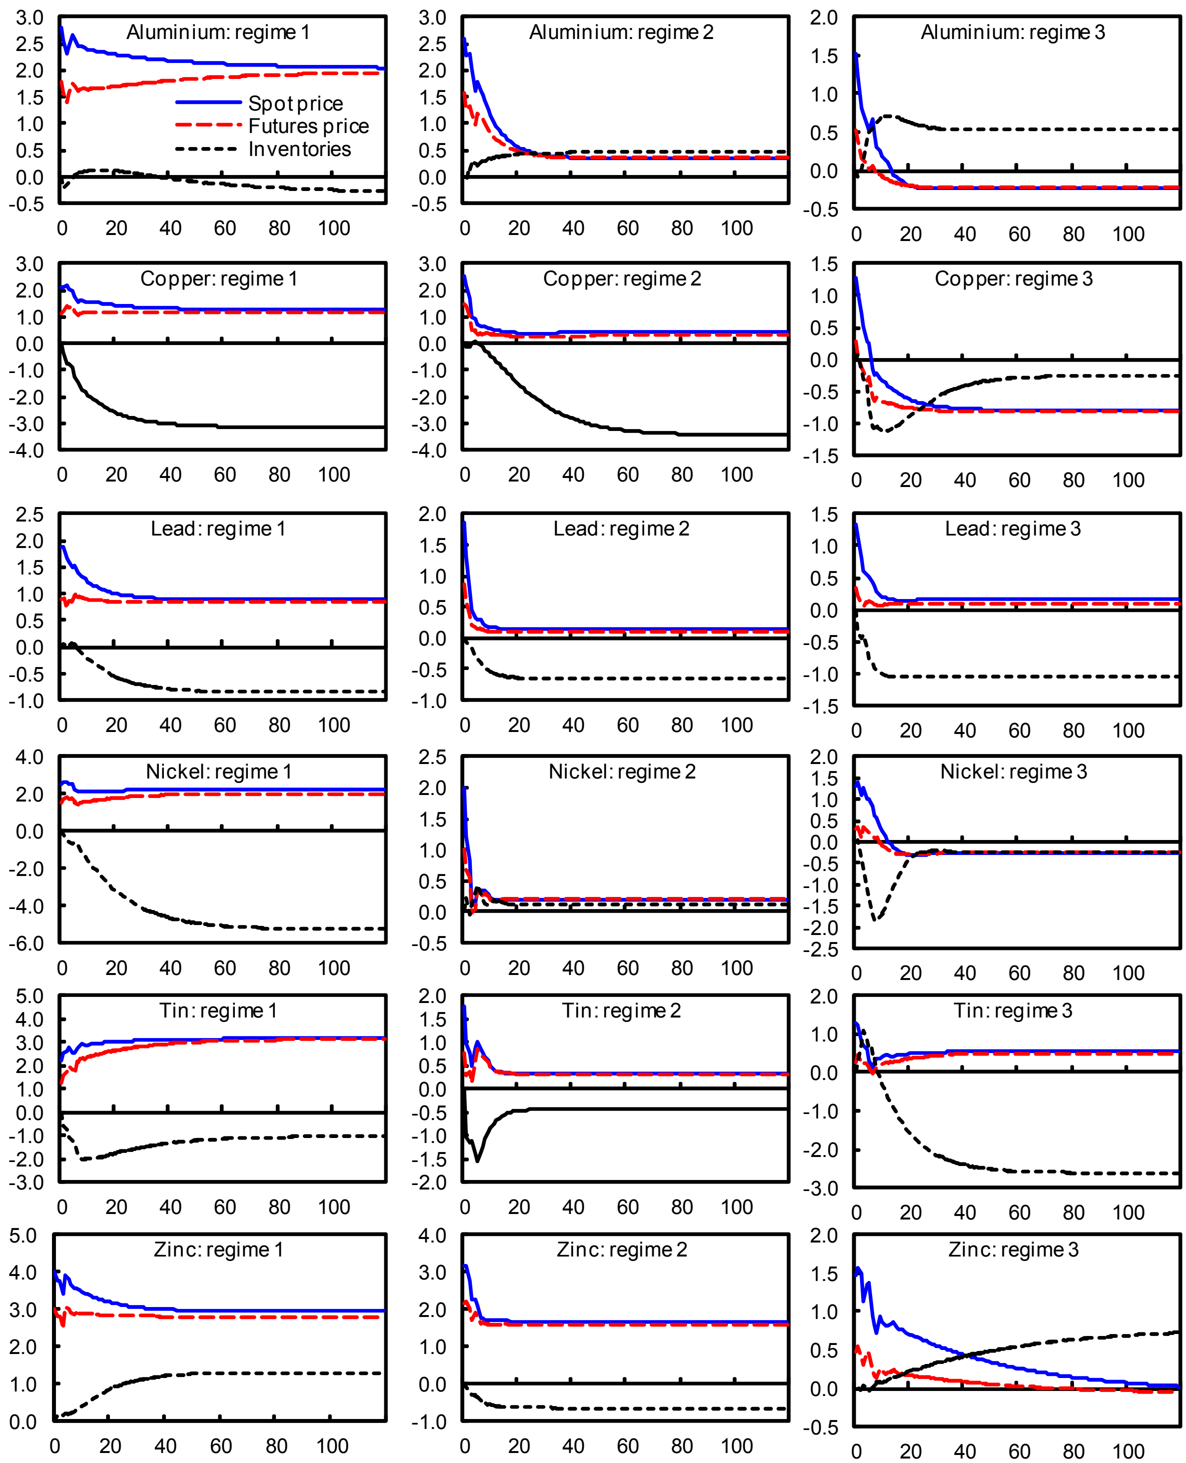

Source: Authors' estimates

Second, the confidence intervals around impulse response estimates are much wider in regime 1 (contango) compared to regimes 2 and 3 (close to equilibrium and backwardation, respectively) for all commodities (Appendix figures A1 through A6). To generate standard errors for the impulse responses, we bootstrapped the residuals from each sample, produced 500 replications, and then calculated the standard deviation of the impulse responses from these estimations. These findings are less easy to interpret, but to some extent, they may reflect the large sample sizes for regime 1, with perhaps a greater range of conditions in this sample as compared to regime 3. A more detailed discussion of this particular results lies outside the scope of this paper.

Different pace of adjustment in each regime

A key result from this paper is that the adjustment path back towards equilibrium, for a given percentage point shock to spot prices, is generally more gradual when the futures curves is in contango and steeply upward sloping (regime 1). The adjustment is more rapid when the futures curve is relatively flat or inverted (regime 3). The most rapid adjustment occurs when the system is close to equilibrium (regime 2). This result holds especially for nickel in regime 3. The only metal for which this result does not hold is zinc (Figure 5).

In most cases, spot prices share much of the burden of adjustment and in backwardation, spot prices adjust particularly rapidly. This likely reflects the convexity of the marginal convenience yield with respect to inventories; in other words, when inventories are already low, the effect of supply disruptions on the marginal utility of inventories is significantly higher. A 1 percentage point shock to spot prices may reflect only a small supply disruption in terms of actual quantities, given the much higher sensitivity of the system to spot prices in backwardation. As inventories are rapidly drawn down, expectations for a more stable path for inventories allows the spot price to fall quickly and the futures curve to return to a more “normal” slope.

In contrast, a 1 percentage point spot price shock in a contangoed market may represent a very significant supply disruption since it will have little effect on the marginal convenience yield as inventories are already abundant. Inventories are drawn down more gradually and the price adjustment is slower. Mechanically, the adjustment coefficients for the VECMs in contangoed markets (regime 1) are much lower than in backwardation, which leads to a much more gradual error-correction process.

Figure 6. Adjustment Back to Equilibrium Following a 1 Percent Spot Price Shock

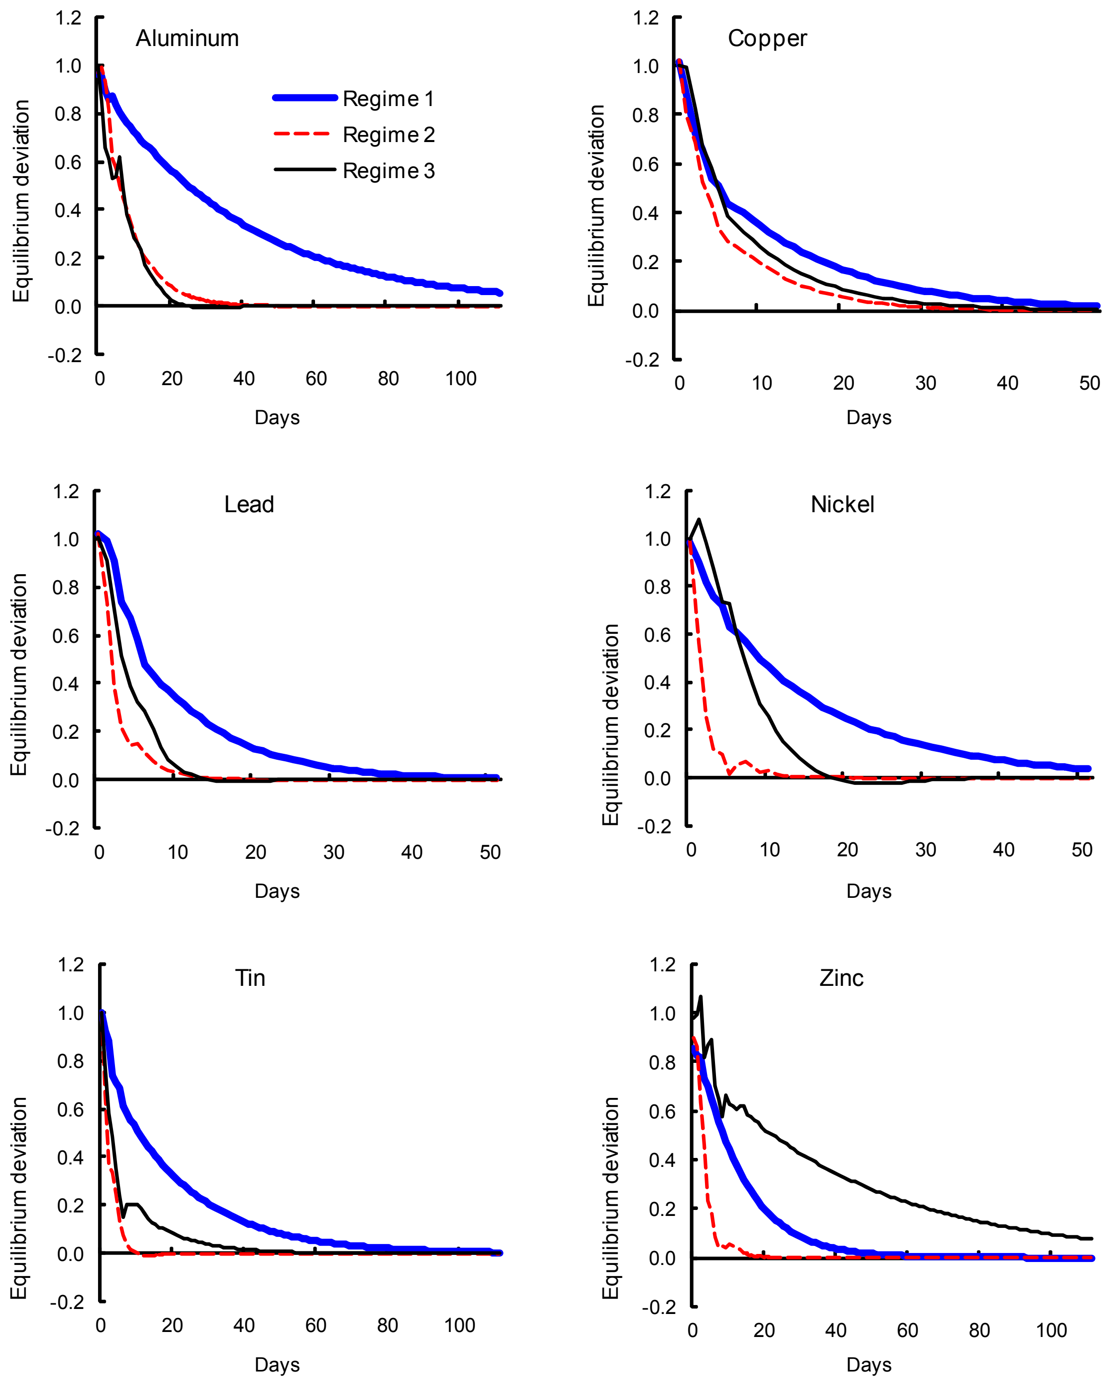

Source: Authors’ estimates.

V. CONCLUDING REMARKS

In this paper, we ask three questions: Is there such a thing as a "normal" commodity market, in which the relationship between spot and futures prices and inventories settles down to a long-run stable equilibrium? How does a commodity market adjust to a temporary scarcity shock which moves the price curve away from this equilibrium? How quickly do inventories and prices respond to such shocks?

Our answer to the first question is "yes". We find that the relationship between base metal spot prices, futures prices, inventories, and interest rates is cointegrating; to put it another way, it is possible to consider whether a commodity market is in "equilibrium" based on the relative values of each of these variables. When the system is away from equilibrium in response to a temporary shock, we should expect it to adjust back towards the steady state over time. The dynamics of this adjustment, however, vary across metals and depend on the initial state of the market.

To the second question, we find some evidence that a temporary scarcity shock, modeled as a spot price shock which changes the slope of the futures curve, does cause a reaction in commodity markets somewhat consistent with a theoretical model, such as Pindyck (2001). In particular, inventories are drawn down and spot prices gradually fall back towards their initial level. However, the initial state of the market is an important conditioning factor for the subsequent adjustment. In a contangoed market with abundant inventories, spot price shocks produce a much more gradual inventory response, while the effect on price levels can be permanent. In contrast, in a backwardated market the inventory drawdown occurs much faster and the rise in both spot and futures prices are temporary.

Our answer to the final question is that the adjustment of prices and inventories back towards equilibrium is much more gradual in a contangoed market. This may reflect the diminishing marginal utility of inventories and the resulting sensitivity of spot prices to supply disruptions in different initial states. For example, a 1 percentage point shock to spot prices may reflect only a small supply disruption in a tight, backwardated market, but a significant disruption when inventories are abundant and spot prices are much less sensitive to perceptions of scarcity. In summary, in a tight physical market, even a small supply disruption can have large price effects, but these typically prove to be short-lived.

These results are important for consumers, producers and inventory holders of commodities. In particular, they suggest that market participants should condition their response to market signals during periods of unusual conditions—or disequilibrium as we have defined it in this paper—on the state of the inventory cycle, which is typically reflected in the slope of the futures curve.

VI. References

Balke, N.S. and T.B. Fomby, 1997, "Threshold cointegration", International Economic Review, 38, pp. 627-645.

Crompton, P. and I.Xiarchos, 2008, "Metal Prices and the Supply of Storage", Commodity modeling and pricing: methods for analyzing resource market behavior, edited by Schaeffer, Peter, p 108-116.

Deaton A. and G. Laroque, 1997, "Competitive Storage and Commodity Price Dynamics", Journal of Political Economy, Vol. 104 No. 5.

Enders**,** W. and P. Siklos, 2001, "Cointegration and Threshold Adjustment," Journal of Business & Economic Statistics, American Statistical Association, vol. 19(2), pp. 166-76.

Heaney**,** R. A**,** 1998, **"**A test of the cost-of-carry relationship using the London Metal Exchange lead contract", The Journal of Futures Markets 18, pp. 177–200.

Johansen**,** S., 1988, 'Statistical Analysis of Cointegration Vectors,**'** _Journal of Economic Dynamics and Control,_v. 12, pp. 231-254.

Kouassi, Eugene, 2008, "Testing for Temporal Asymmetry In the Metal Price-Stock Relationship", Commodity modeling and pricing : methods for analyzing resource market behavior, edited by Schaeffer, Peter, p 118-134

Mabro, R., "The Oil Price Crises of 1998-9 and 2008-9", Oxford Energy Forum, May 2009.

Martens**,** M**.,** Kofman**,** P. and Vorst**,** A.C.F., 1998**, "**A Threshold Error Correction Model for Intraday Futures and Index Returns," Journal of Applied Econometrics, 13, pp. 245-263.

Markert, V. and H. Zimmermann, 2008, Chapter 5, "The Relationship Between Risk Premium and Convenience Yield Models' in Handbook of Commodity Investing by Dieter G. Kaiser, Frank J. Fabozzi, Roland Fuss

McMillian, David, 2005, "Threshold adjustment in spot-futures metals prices", Applied Financial Economics Letters, 1, pp 5-8.

Tsay, R. S., 1989, "Testing and modeling threshold autoregressive processes", Journal of the American Statistical Association, 84, pp. 231-240, 1989.

Pindyck, R., 2001, "The dynamics of commodity spot and futures markets: a primer", Energy Journal, Vol 22, No. 3:1-29.

Watkins, C. and M. McAleer, 2006. "Pricing of non-ferrous metals futures on the London Metal Exchange," Applied Financial Economics, Taylor and Francis Journals, vol. 16(12), pp 853-880.

Williams, J and B. D. Wright, Storage and commodity markets, Cambridge ; New York : Cambridge University Press, 1991.

Wu, Yangru, and Hua Zhang, 1997, "Do Interest Rates Follow Unit-Root Processes? Evidence from Cross-Maturity Treasury Bill Yields," *Review of Quantitative Finance & Accounting*, Vol. 8, Issue 1, pp. 69-81.

VII. APPENDIX

Table A1. Unit Root Tests 1/

	Log levels				First differenced logs			
	ADF		PP		ADF		PP	
	t-statistic	p-value	t-statistic	p-value	t-statistic	p-value	t-statistic	p-value
Libor interest rates								
3-month	-0.68	0.85	-0.54	0.88	-14.90	0.00	-43.89	0.00
6-month	-0.61	0.87	-0.61	0.86	-16.78	0.00	-48.70	0.00
Aluminum								
spot price	-1.38	0.59	-1.34	0.61	-57.49	0.00	-57.49	0.00
futures price (3-month)	-0.45	0.90	-0.52	0.89	-59.54	0.00	-59.41	0.00
futures price (6-month)	-0.60	0.87	-0.61	0.87	-53.92	0.00	-53.90	0.00
inventories	-1.14	0.70	-1.13	0.70	-55.37	0.00	-55.37	0.00
Copper								
spot price	-0.45	0.90	-0.39	0.91	-53.71	0.00	-53.72	0.00
futures price (3-month)	-1.12	0.71	-1.07	0.73	-57.64	0.00	-57.68	0.00
futures price (6-month)	-1.34	0.61	-1.34	0.61	-57.90	0.00	-57.90	0.00
inventories	-0.43	0.90	-0.51	0.89	-59.74	0.00	-59.60	0.00
Lead								
spot price	-0.51	0.89	-0.54	0.88	-54.20	0.00	-54.19	0.00
futures price (3-month)	-1.12	0.71	-1.11	0.71	-55.93	0.00	-55.95	0.00
futures price (6-month)	-0.39	0.91	-0.35	0.92	-53.77	0.00	-53.77	0.00
inventories	-1.09	0.72	-1.04	0.74	-57.94	0.00	-58.00	0.00
Nickel								
spot price	-1.26	0.65	-1.24	0.66	-57.63	0.00	-57.63	0.00
futures price (3-month)	-0.40	0.91	-0.49	0.89	-60.05	0.00	-59.89	0.00
futures price (6-month)	-0.44	0.90	-0.46	0.90	-54.29	0.00	-54.28	0.00
inventories	-1.08	0.73	-1.06	0.74	-56.01	0.00	-56.02	0.00
Tin								
spot price	-0.45	0.90	-0.39	0.91	-54.01	0.00	-54.03	0.00
futures price (3-month)	-1.03	0.74	-0.97	0.77	-57.85	0.00	-57.93	0.00
futures price (6-month)	2.27	1.00	1.56	1.00	-14.28	0.00	-50.45	0.00
inventories	-1.22	0.67	-1.23	0.66	-11.98	0.00	-57.17	0.00
Zinc								
spot price	-1.51	0.53	-1.71	0.42	-20.33	0.00	-51.57	0.00
futures price (3-month)	-1.18	0.69	-1.26	0.65	-20.01	0.00	-47.20	0.00
futures price (6-month)	-1.91	0.33	-1.90	0.33	-21.69	0.00	-48.60	0.00
inventories	-0.83	0.81	-0.90	0.79	-15.96	0.00	-53.62	0.00

1/ ADF denotes Augmented Dickey-Fuller test and PP denotes the Philips-Perron test. Various lag lengths were used for the ADF tests and Table A1 show the results from tests with a lag length of six.

A. Cointegration Tests

For the Engle-Granger procedure, we tested the null hypothesis of no cointegration by estimating the following regression, using the residuals from equation (6):

$$\Delta\hat{\varepsilon}_t = a_1\hat{\varepsilon}_{t-1} + u_t \tag{A1}$$

Table A2 presents the t-statistics from these regressions for each commodity using the 3-month and 6-month futures contract and interest rate, together with the 5 percent Engle-Granger critical values. (Johansen test results available on request.) In all cases, we were able to reject the null of no cointegration.

Table A2. Engle-Granger tests of Cointegrating Residuals 1/

	Critical values	Test statistic 3-month model	Test statistic 6-month model
Aluminum	-4.12	-9.11	-4.88
Copper	-4.12	-12.73	-6.37
Lead	-4.12	-10.09	-6.91
Nickel	-4.12	-9.52	-6.71
Tin	-4.12	-8.25	-5.53
Zinc	-4.12	-8.00	-5.94

1/ The null hypothesis is for a unit root in the residuals of the equation and no cointegration. The test statistic is calculated using the Philips-Perron procedure and the critical values are taken from MacKinnon (1991).

For the VECM estimations, we use the lag length identified by standard selection criteria for the VAR in log-levels for each commodity and based on the variables in equation (4) (Table A3).

Table A3. Vector Autoregression Lag Length Tests 3/

	Akaike	Schwarz-Bayes	Hannan-Quinn
	Three-month model		
Aluminum	**7**	3	6
Copper	**8**	3	7
Lead	**7**	2	3
Nickel	**7**	2	4
Tin	**8**	2	3
Zinc	**11**	2	3
	Six-month model		
Aluminum	**7**	2	5
Copper	**7**	3	7
Lead	**7**	2	3
Nickel	**7**	2	3
Tin	**7**	2	3
Zinc	**7**	2	7

Source: Authors' calculations.
3/ Information criteria include Aikaike (AIC), Schwarz-Bayesian (SIC), and Hannan-Quin (HQ). We base our decisions on the AIC.

B. Tsay's Test for Threshold Nonlinearity

The first stage is to assess the autoregressive structure of the equilibrium error. We find that partial autocorrelations decline rapidly after the first lag, although they remain statistically significant (Table A4). Although information criteria indicate that the optimal AR order varies between 2 and 5 (Table A5), running the threshold tests on AR(1) or these optimal AR orders produced either identical or very similar results, underscoring the dominant influence of the first lag.

Table A4. Partial Autocorrelation Functions for the Equilibrium Errors 4/

	Lag order				
	1	2	3	4	5
			3-month		
Aluminum	**0.94**	**0.07**	**0.14**	0.03	**-0.06**
Copper	**0.89**	**0.09**	**0.11**	**0.04**	0.01
Lead	**0.93**	**-0.05**	**0.07**	0.03	0.04
Nickel	**0.94**	0.00	**0.05**	**0.05**	**-0.04**
Tin	**0.95**	0.02	**0.09**	0.01	0.00
Zinc	**0.97**	**-0.06**	**-0.14**	0.00	-0.03
			6-month		
Aluminum	**0.98**	**-0.04**	**0.06**	0.00	0.02
Copper	**0.97**	0.00	0.03	**0.04**	0.03
Lead	**0.97**	**-0.10**	**0.05**	**0.06**	0.03
Nickel	**0.97**	**-0.10**	0.01	0.00	0.00
Tin	**0.97**	0.03	**0.05**	0.01	**0.04**
Zinc	**0.98**	**-0.15**	**-0.12**	**0.05**	0.00

4/ Autocorrelations significant at the 95 percent level.

Table A5. Information Criteria for Equilibrium Error AR(p) Equations 5/

	AR order				
	1	2	3	4	5
			3-month		
Aluminum	-8.2896	-8.2942	-8.3125	-8.3123	**-8.3151**
Copper	-8.5228	-8.5309	-8.5431	**-8.5439**	-8.5430
Lead	-7.6642	-7.6653	-7.6689	-7.6691	**-7.6694**
Nickel	-7.9656	-7.9647	-7.9658	-7.9676	**-7.9682**
Tin	-9.2667	-9.2663	**-9.2733**	-9.2725	-9.2716
Zinc	-8.2470	-8.2505	**-8.2685**	-8.2682	-8.2685
			6-month		
Aluminum	-8.7211	-8.7229	-8.7271	-8.7271	**-8.7276**
Copper	-8.6499	-8.6498	-8.6526	**-8.6536**	-8.6534
Lead	-7.3551	-7.3640	-7.3653	-7.3683	**-7.3680**
Nickel	-7.7603	**-7.7702**	-7.7696	-7.7686	-7.7677
Tin	-8.7312	-8.7312	**-8.7330**	-8.7322	-8.7315
Zinc	-8.2937	-8.3254	-8.3421	-8.3624	**-8.3723**

Source: Authors' calculations
5/ Minimum criteria values in bold.

As a result, we arrange the data such that it is increasing in the value of the AR(1) regressor, in our case the equilibrium error in the previous period z_{t-1}. The least squares estimates of the AR(1) regressor in equation (5) will be consistent for each set of cases, if the value of the thresholds were known. Since the value of the thresholds is unknown, we proceed sequentially. The predictive residuals from equation (5) will be white noise asymptotically and orthogonal to the regressor until z_{t-1} reaches a threshold, at which point the predictive residual will be biased and a function of the regressor. To test this, we obtain the standardized predictive residuals from an ordered autoregression, where π_i is the time index of the tth smallest observation, and run the least squares regression:

$$\varepsilon_{\pi_t+1} = \omega_0 + \omega_1 z_{\pi_t} + v_{\pi_t+1} \tag{A1}$$

We do this for all sample periods $i = k + 1,\ldots,T - 1$, where k is the number of explanatory variable (on our case 1) and compute Tsay's statistic, which is the F-statistic of the resulting regression:

$$F(p,d) = \frac{\left(\sum \hat{\varepsilon}_t^2 - \sum \hat{v}_t^2\right)/(p+1)}{\sum \hat{v}_t^2 / (T-d-k-p-h)} \tag{A2}$$

In equation (A2), d is the delay parameter (in our case 1), p is the order of the autoregression, and h is obtained from $\max\{1, p + 1\}$.This test statistic follows an F distribution with $p + 1$ and $T - d - k - p - h$ degrees of freedom. Implementing this procedure on the residuals from the cointegrating equations for each commodity obtains the following test statistics and p-values for the null hypothesis that the standardized recursive residuals are not a function of the regressor. Table A6 shows the results of these tests. In all cases (with the exception of 6-month zinc), it was possible to reject the null hypothesis of linearity at the 1 percent level of confidence. Although 6-month zinc was an exception, other tests for structural breaks suggested that there is a significant degree of nonlinearity.

C. Quandt-Andrews Tests for Structural Breaks

This procedure performs a Chow test at every observation between two dates, or observations. We then identify the maximum Wald F statistic from each individual Chow test and assess whether it is possible to reject the null of no structural break at the 95 percent confidence level. We then perform the same procedure for the largest remaining sub-sample to check for another breakpoint. In all cases, we found it was possible to reject the null hypothesis of no structural break at the 1 percent level for the overall sample and the sub-sample constructed by removing the smallest section of the sample split by the first threshold. We used these F statistic maxima to identify the threshold locations. Details are available from authors on request.

Table A6. Tsay's Nonlinearity Test Results 6/

	3-month		6-month	
	Test statistic	p-value 1/	Test statistic	p-value 1/
Aluminum				
full sample	72.16	0.0000	32.2	0.0000
sub-sample	98.37	0.0000	161.4	0.0000
Copper				
full sample	300.36	0.0000	174.4	0.0000
sub-sample	11.52	0.0000	324.89	0.0000
Lead				
full sample	83.94	0.0000	29.31	0.0000
sub-sample	50.53	0.0000	16.41	0.0000
Nickel				
full sample	53.39	0.0000	13.07	0.0000
sub-sample	122.05	0.0000	177	0.0000
Tin				
full sample	27.64	0.0000	32.36	0.0000
sub-sample	14.16	0.0000	197.02	0.0000
Zinc				
full sample	6.01	0.0025	0.46	0.6313
sub-sample	74.2	0.0000	122.98	0.0000

Source: Authors' calculations
6/ Probability that the null hypothesis of linearity (no thresholds) is true.

Figure A1. Aluminum: Impulse Responses to 1 percent Spot Price Shock

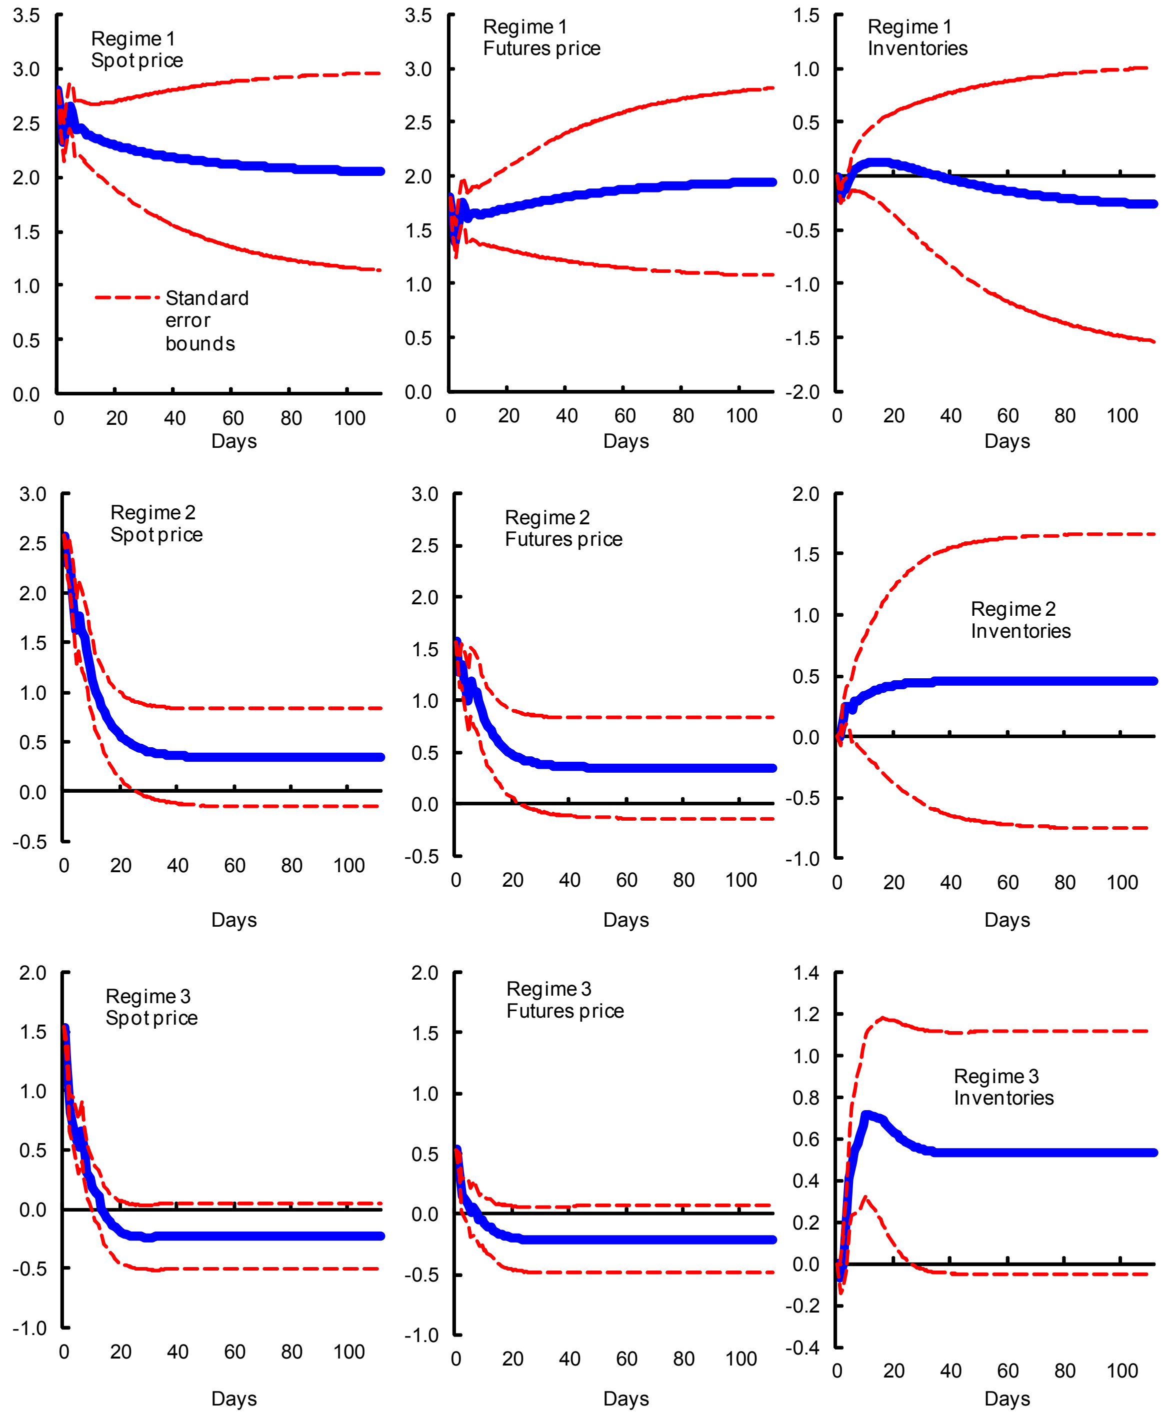

Source: Authors' estimates

Figure A2. Copper: Impulse Responses to 1 percent Spot Price Shock

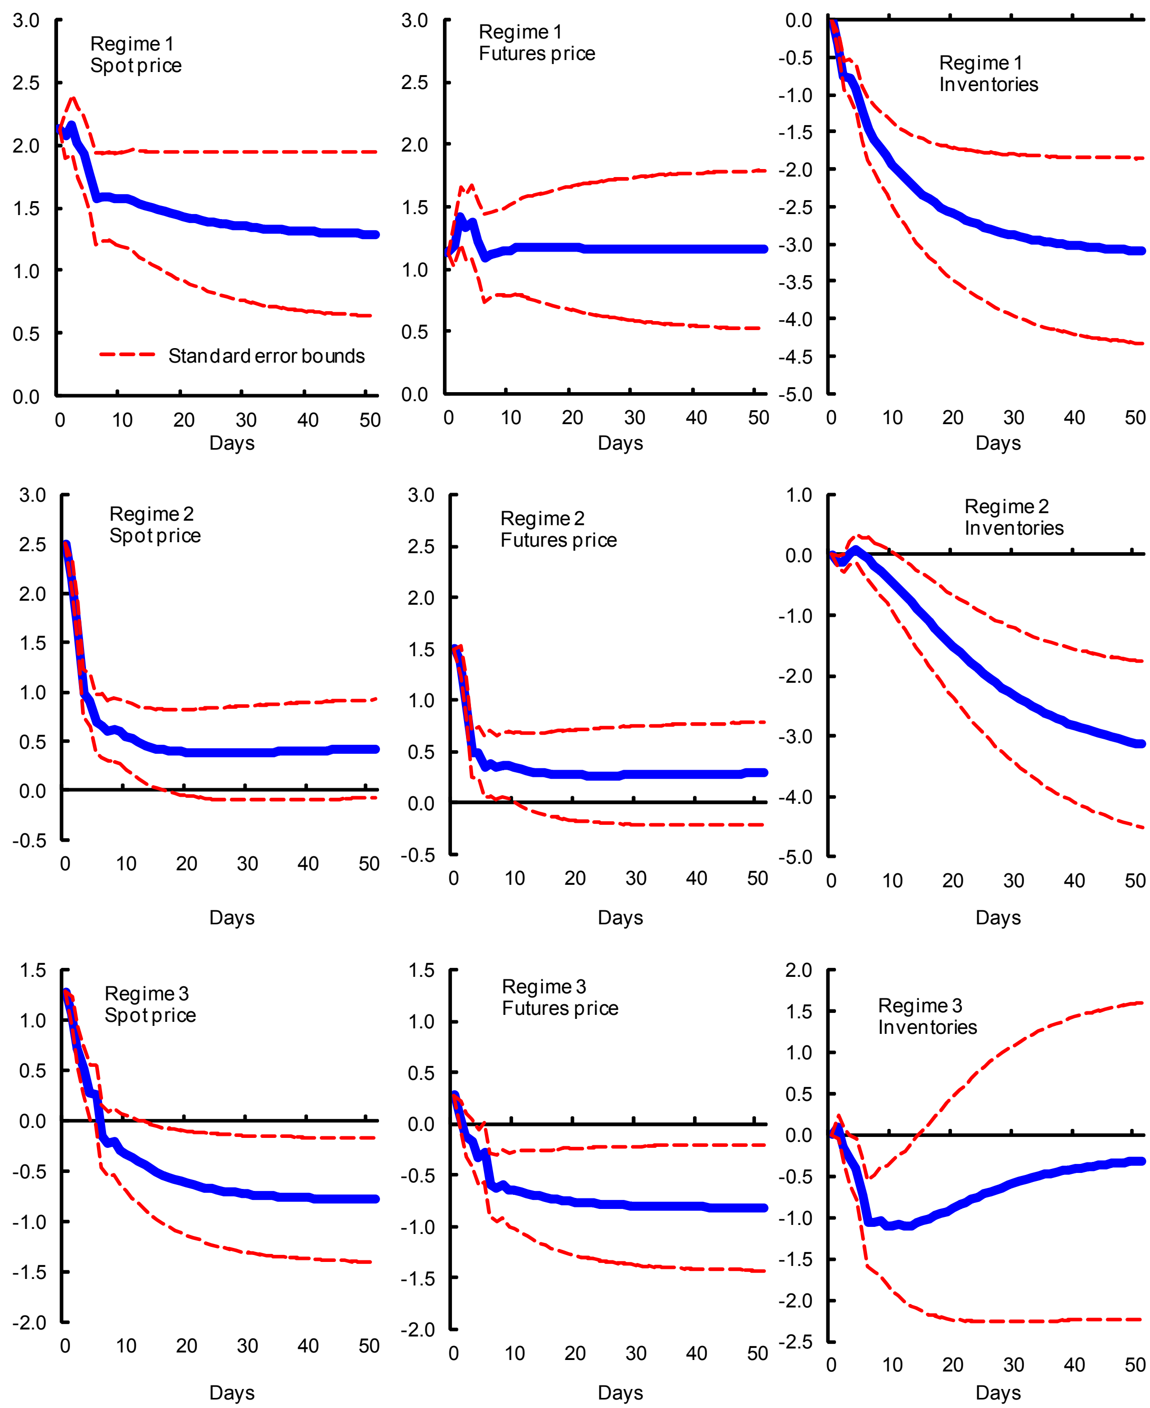

Source: Authors' estimates

Figure A3. Lead: Impulse Responses to 1 percent Spot Price Shock

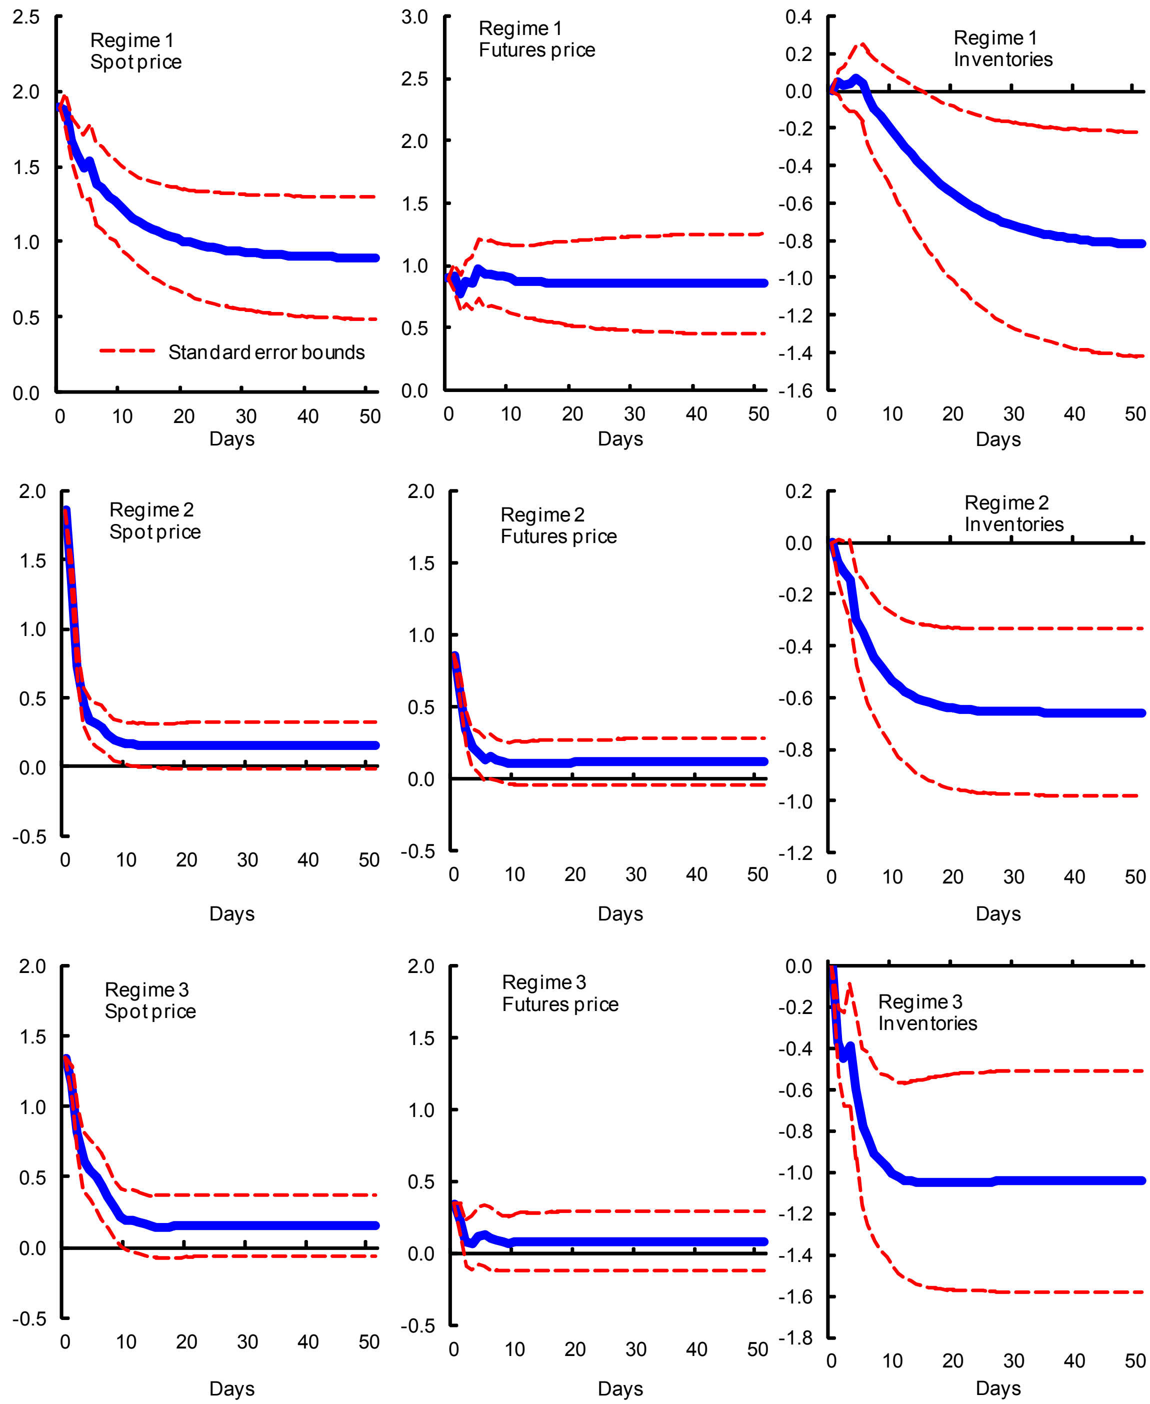

Source: Authors' estimates

Figure A4. Nickel: Impulse Responses to 1 percent Spot Price Shock

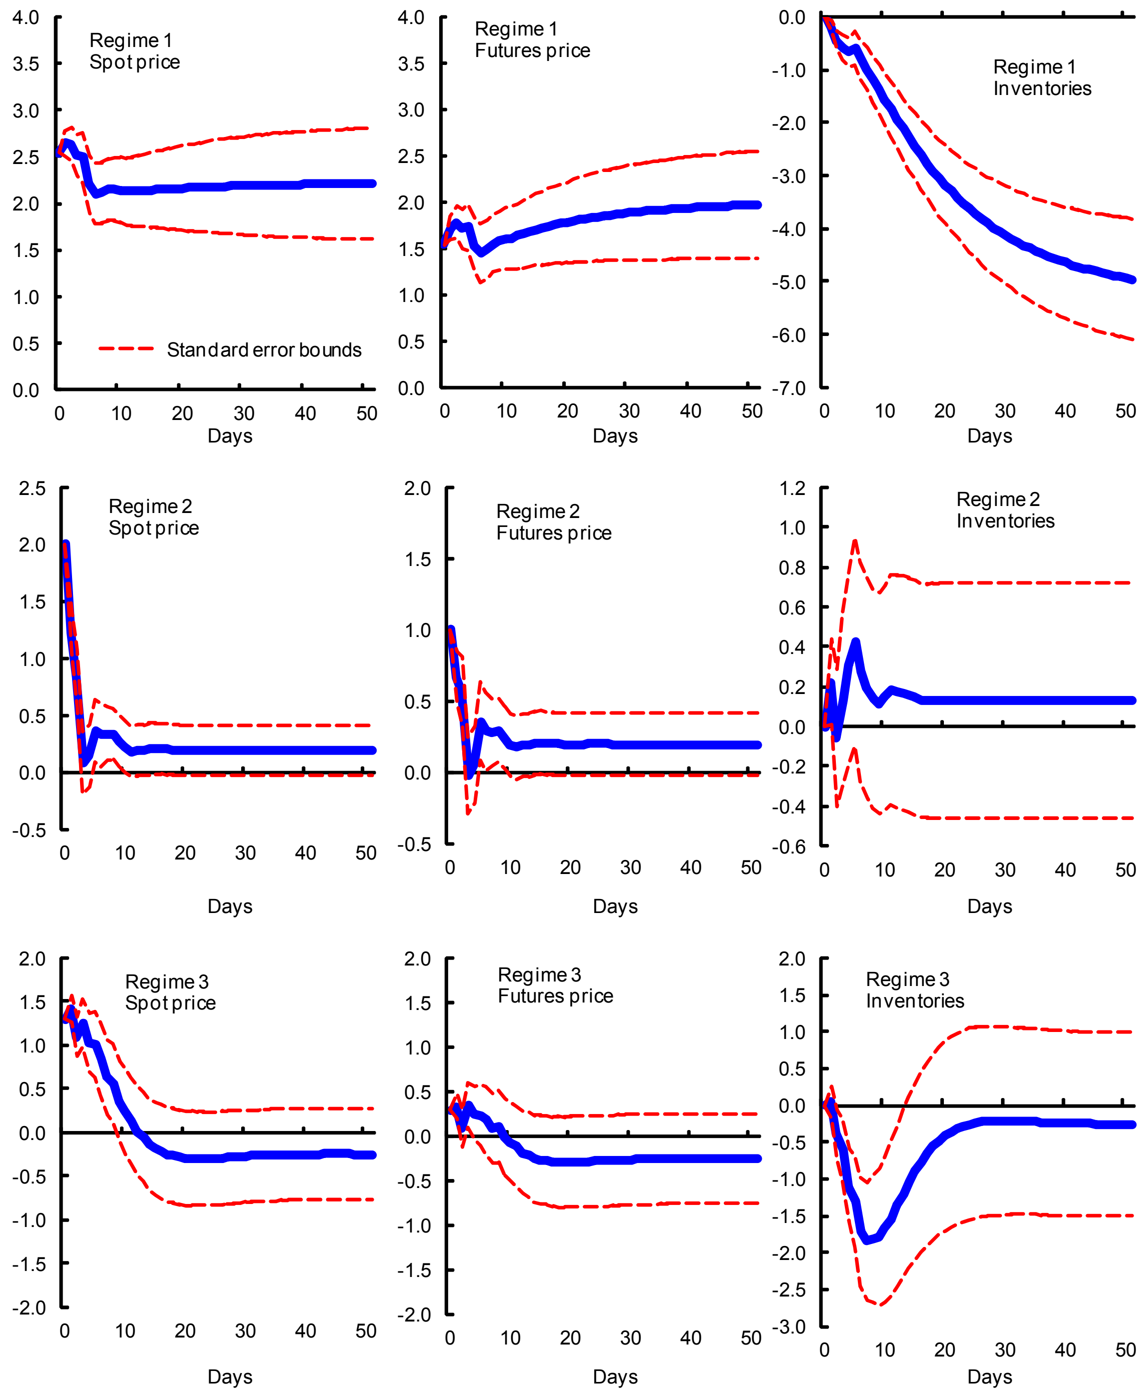

Source: Authors' estimates

Figure A5. Tin: Impulse Responses to 1 percent Spot Price Shock

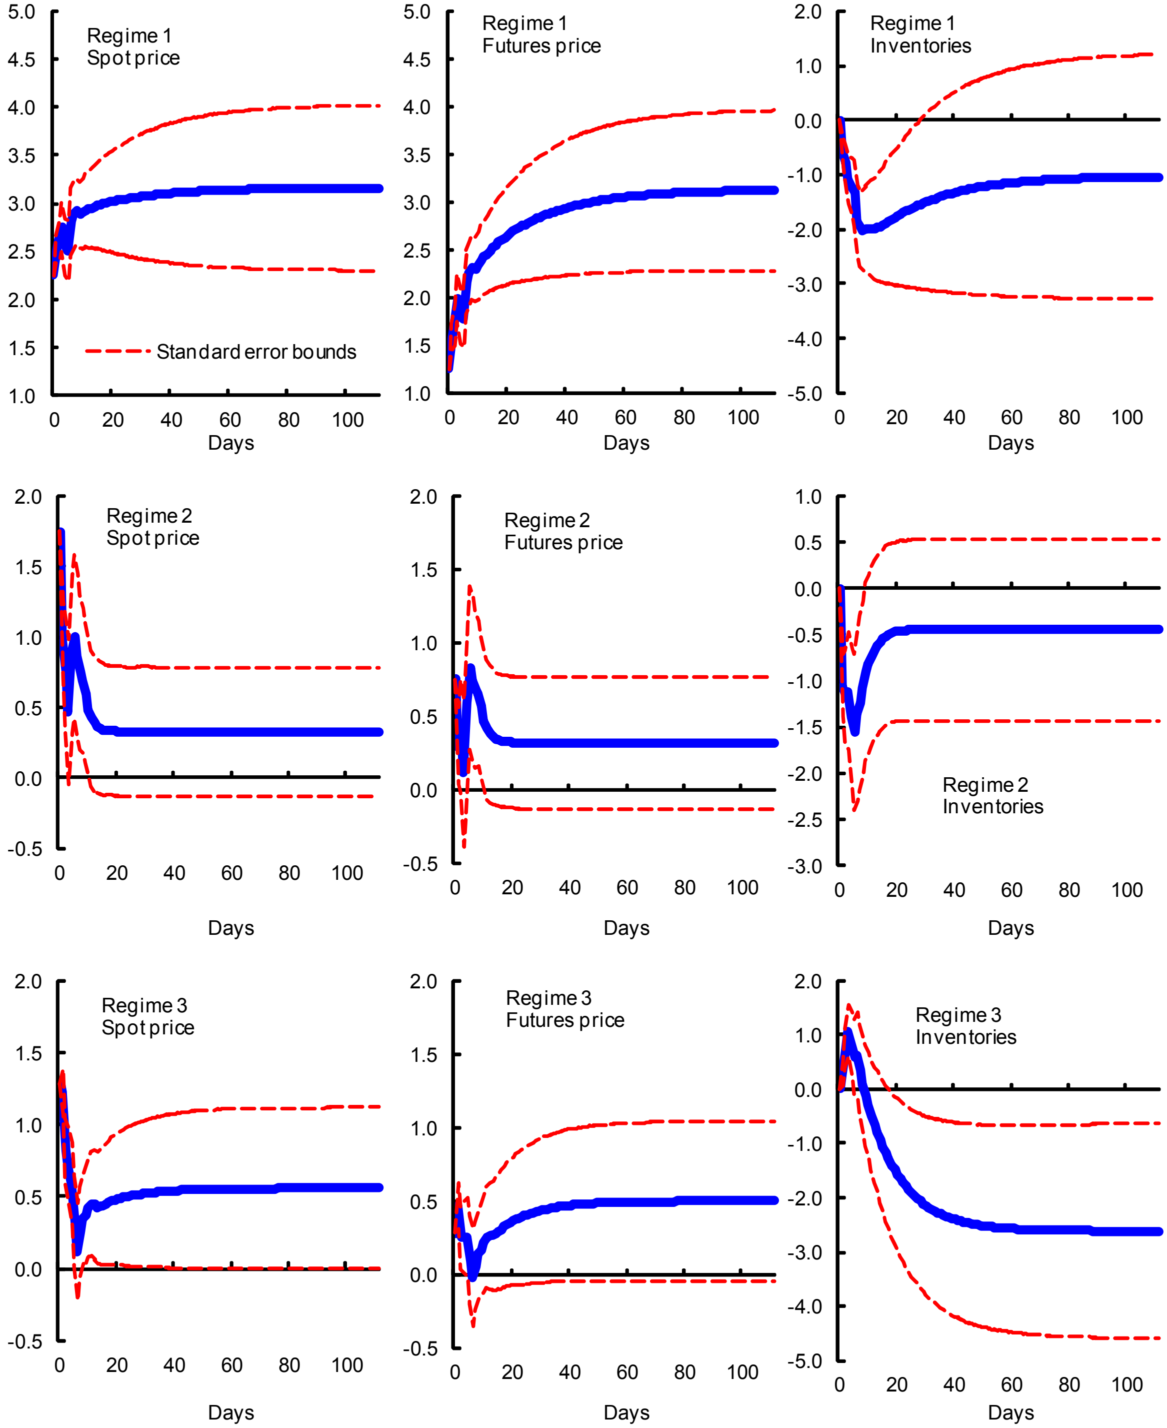

Source: Authors' estimates

Figure A6. Zinc: Impulse Responses to 1 percent Spot Price Shock

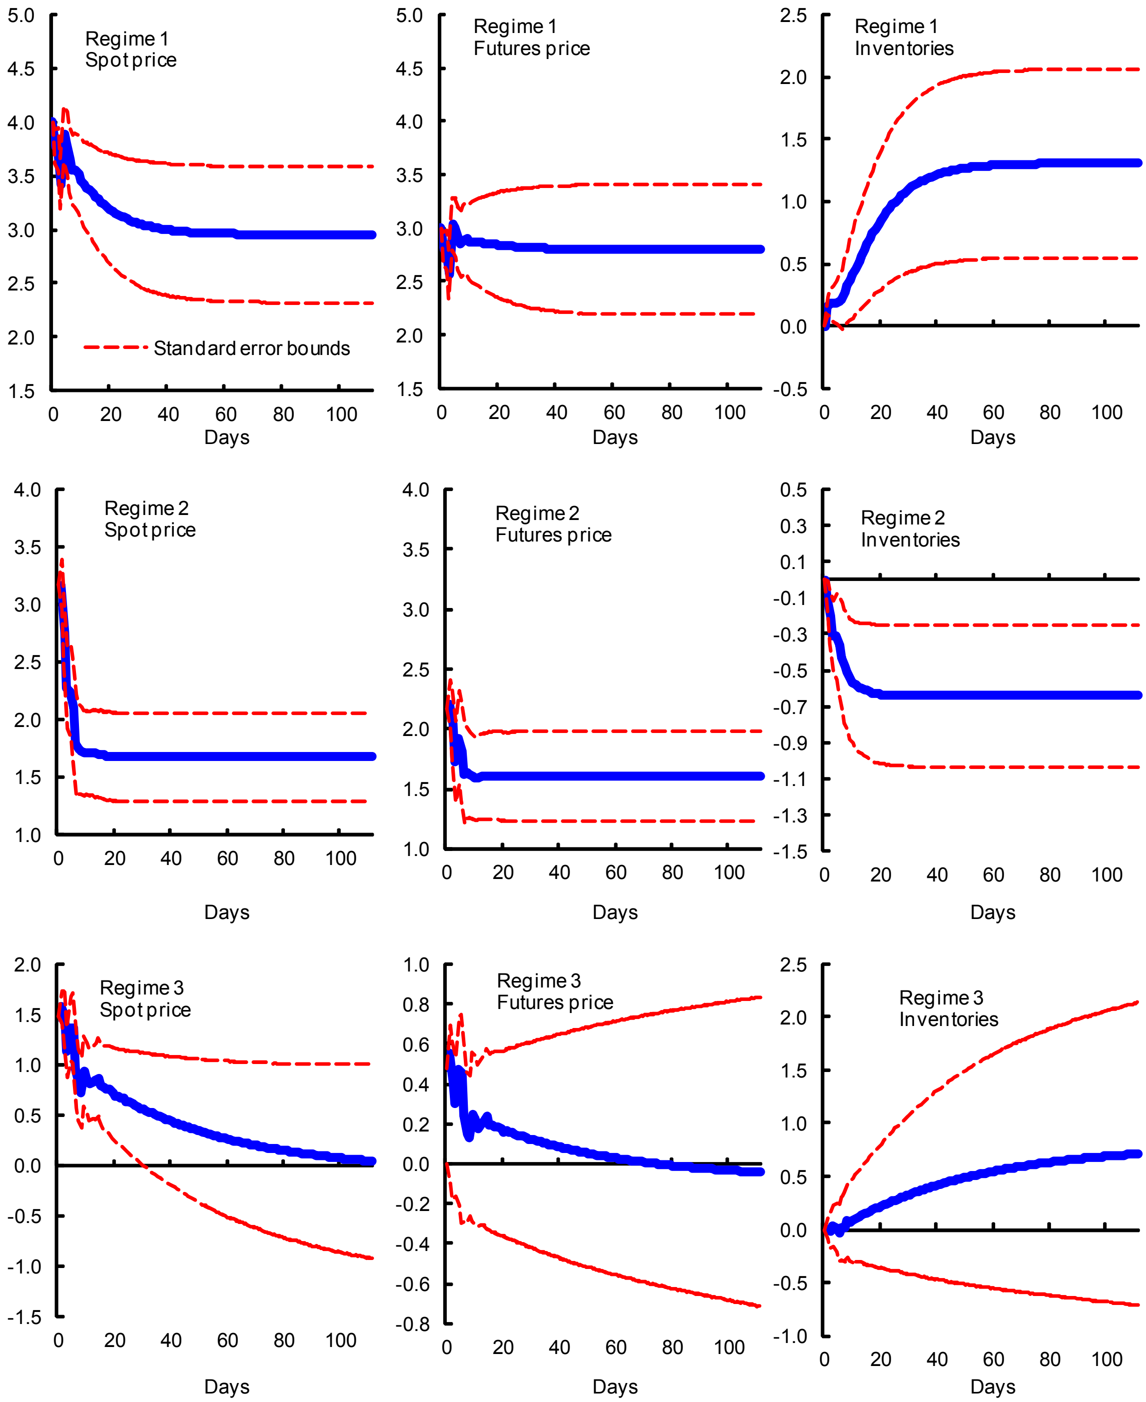

Source: Authors' estimates