THE UNIVERSITY OF SYDNEY

Economics Working Paper Series

2016 – 10

The Role of El Niño Southern Oscillation in Commodity Price Movement and Predictability

David Ubilava

August 3, 2017

The Role of El Niño Southern Oscillation in Commodity Price Movement and Predictability*

David Ubilava†

School of Economics, University of Sydney

August 3, 2017

Abstract

El Niño Southern Oscillation (ENSO) impacts the supply and, to some extent, the demand for primary commodities. But what are the consequences of this climatic phenomenon for these commodity price dynamics? I consider monthly series of 43 primary commodity prices and sea surface temperature (SST) anomalies in the `Niño3.4` region from January 1980 to December 2016. The SST anomalies serve as a proxy for ENSO, as persistent positive (negative) SST deviations from their historical mean are associated with El Niño (La Niña) events. I apply a family of time-varying smooth transition autoregressive models to account for potentially complex dynamic relationships between SST anomalies and prices. Overall, the estimated nonlinearities bring out more amplified price responses during El Niño events, and at the onset of the ENSO cycle. I find statistically significant linkages between SST anomalies and a subset of agricultural commodity prices. This in-sample fit manifests in a forecasting environment for the commodities produced in the tropics. While I also find some in-sample evidence for prices of selected nonagricultural commodities, e.g., timber and metals, the ability of SST anomalies to predict these commodity prices in an out-of-sample setting is lacking. These findings carry important welfare implications, especially for developing economies that have been historically linked to the behavior of primary commodity prices, and offer valuable insights to policy makers working in areas related to economic growth and foreign aid programs, as well as those concerned with issues of farm income and rural poverty.

Keywords: Commodity Prices, El Niño Southern Oscillation, Nonlinear Dynamics, Time-Varying Smooth Transition Autoregression

JEL Codes: C51, E31, Q54

---

*An earlier version the manuscript was circulated under the title *"Rises and Falls in Primary Commodity Prices: Blame it on ENSO or Leave Them Kids Alone?"* From the earlier version to the current, there are important changes in research execution and results. I would like to thank Peter Exterkate, Aaron Smith, Jesse Tack, and Nelson Villoria for their valuable comments and suggestions. Errors are mine.

†E-Mail: david.ubilava@sydney.edu.au

During 2015–2016, in the wake of one of the strongest El Niño events in recent history, news concerning possible global food shortage and commodity prices spikes headlined the popular media and intensified academic discussions. The concern was not unfounded, as some of the most devastating weather events have been associated with extreme episodes of the El Niño Southern Oscillation (ENSO) (Pielke Jr and Landsea, 1999; Davis, 2002). This climate phenomenon is assessed more scrupulously now than ever before as a result of its improved predictability (Stone et al., 1996; Ludescher et al., 2014), as well as a better general understanding of its global economic consequences (Brunner, 2002; Cashin et al., 2017). Because of apparent linkages between ENSO and weather, a large number of studies have examined the production effect of this climate phenomenon, particularly in the agricultural sector (e.g., Iizumi et al., 2014; Hsiang and Meng, 2015). But nontrivial linkages between ENSO and a broader range of primary commodity prices require further investigation. To the extent that the economies of developing countries are affected by primary commodity price behavior, the insights from this article are of importance to policy makers within local governments, as well as those in international organizations, who rely on short- and intermediate-term forecasts of commodity prices.

This article contributes to the literature in several ways. First, it incorporates a comprehensive set of commodity spot prices to gain insight into the ENSO effect across different commodity groups. Second, it relies on the time-varying smooth transition autoregressive (TV-STAR) modeling framework of Lundbergh et al. (2003) to account for any parameter nonlinearity or structural change in the ENSO–price relationship. Third, it applies the highest density regions (HDRs) of Hyndman (1995, 1996) to the generalized impulse–response functions (GIRs) of Koop et al. (1996) to offer a more complete illustration of nonlinear dynamics in commodity prices as a result of ENSO shocks. Finally, this article extends the modeling framework to a forecasting exercise to assess the informational content of ENSO in the out-of-sample predictability, i.e., causality in the Granger sense, of commodity prices.

El Niño is one of the two extreme phases of the ENSO cycle, the other being La Niña. When neither of these is present, ENSO is said to be in a neutral phase. Figure 1 illustrates some notable features of the weather effects of the ENSO phases. First, as a result of the presence of so-called *teleconnections*, ENSO-related weather shocks not only manifest in regions tangential to the Pacific (where the climate anomaly is observed), but also radiate to more distant regions of the globe (Rasmusson, 1991). Second, regions within the tropical band appear to be affected more than those outside of this band (Hsiang and Meng, 2015). Third, the two phases do not necessarily result in weather conditions that are mirror images

of each other (Hoerling et al., 1997; Zhang et al., 2014).

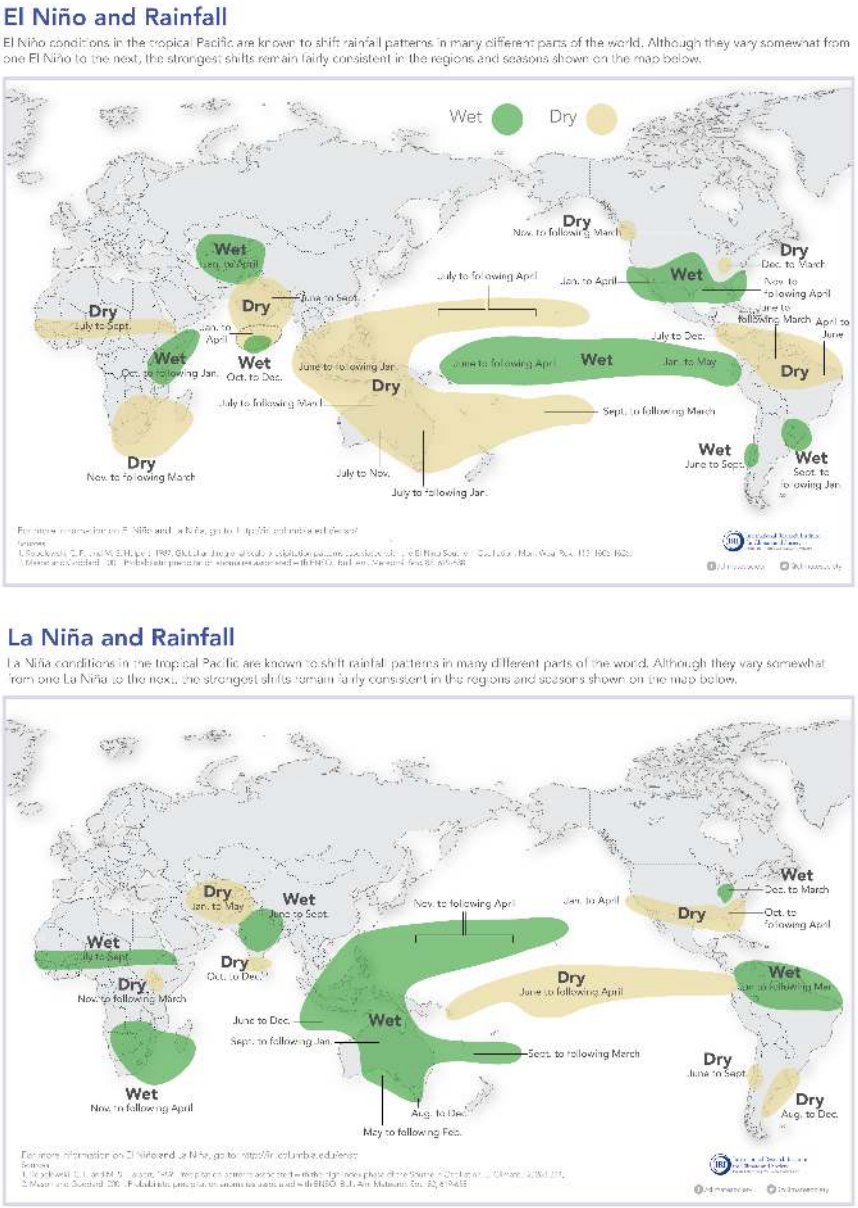

Figure 1: Global weather effects associated with ENSO phases

Source: International Research Institute for Climate and Society at Columbia University, based on Ropelewski and Halpert (1989) and Mason and Goddard (2001).

Production is a crucial link connecting ENSO phases to commodity prices because ENSO-induced weather anomalies directly affect regions that supply a wide range of agricultural commodities. For example, the two largest palm oil producing countries, Malaysia and

Indonesia, account for approximately 90% of world exports.[1] Peru and Chile, the region where the unusual warming of ocean waters associated with the El Niño phase was first observed by local fishermen, account for nearly two-thirds of world fishmeal exports. Vietnam and Indonesia produce and export more than 60% of internationally traded *Robusta* coffee, and Brazil is responsible for nearly half of the world's *Arabica* coffee exports.

While the agricultural sector is most susceptible to ENSO shocks, nonagricultural commodities may also be influenced by this climate anomaly. For example, ENSO-related droughts can affect mining activities that rely on hydropower, while excessive precipitation can flood roads and sites, thereby disrupting the supply of metals and minerals (World Bank Group, 2015; Cashin et al., 2017). ENSO-induced heat and drought can cause wildfires that may affect the availability of timber (Nepstad et al., 1999; Gan, 2006). As is the case with agricultural commodities, the production of some nonagricultural commodities is geographically concentrated in ENSO-affected regions. While China leads all other countries in the production of many metals, South Africa is the world's largest platinum supplier, with approximately 70% of global market share, Chile's share of world copper production is 35%, Mexico's share of world silver production is nearly 20%, and Australia, Indonesia, and Peru are noteworthy suppliers of a range of metals as well. Compared with agricultural commodities, however, the effect of ENSO shocks on production in the nonagricultural sector is likely to be weak, both statistically and economically. First, mining centers are dispersed throughout Southeast Asia and various countries in Africa and the Americas, and not all of these areas are simultaneously affected by ENSO. Second, in the nonagricultural sector, ENSO-induced weather anomalies can only disrupt production in the short run, compared with a more prolonged effect in the case of agricultural commodities.

The supply side effect is not the sole channel through which the ENSO cycle can impact commodity price movements. To the extent that ENSO events affect economies of developed as well as developing countries (e.g., Brunner, 2002; Laosuthi and Selover, 2007; Cashin et al., 2017; Smith and Ubilava, 2017), this climate phenomenon can impact the prices of commodities such as timber and metals, which are known to be linked to business cycles. Alternatively, storms and hurricanes, the frequency of which can be attributed to distinct ENSO phases (e.g., during the El Niño phase, "named storms" and hurricanes are less likely in the Caribbean and Atlantic, but more likely in the eastern Pacific), can constitute yet another channel through which this climate phenomenon can affect the demand for, and

[1]The statistics presented here and elsewhere, unless otherwise stated, are the author's calculations, based on various online publications of the U.S. Department of Agriculture and the U.S. Geological Survey.

thus the prices of commodities used in construction. Moreover, extreme ENSO episodes that adversely affect household incomes, particularly those in rural areas, can result in weaker derived demand and downward pressure on precious metal prices. However, these effects are unlikely to be strong or economically meaningful.

The foregoing discussion provides examples of the numerous ways in which ENSO cycles can affect commodity price behavior. Here, I do not aim to investigate every channel that can potentially impact price behavior in a structural framework. That is, the goal is not to identify pathways through which ENSO impacts price movements. Rather, I am concerned with examining the overall impact of ENSO on commodity prices in a reduced-form framework. Put differently, the overarching aim of this study is to examine the relationship between ENSO phases and price dynamics, regardless of what the causal mechanism might be, and to assess the ability of ENSO-related sea surface temperature anomalies to predict commodity price behavior. Thus, the findings of this study will provide important benefits, particularly in terms of facilitating good decision-making, by producing more accurate commodity price predictions.

Nonlinear Modeling and Testing Frameworks

In this article, I relax the linearity assumption in relation to the ENSO–price nexus. That is, I allow for the possibility that price responses to ENSO shocks can be episodic, or regime-dependent. This is facilitated by several underlying factors. First, the ENSO cycle is asymmetric, as El Niño events tend to develop somewhat abruptly, whereas La Niña events often evolve following El Niño events (e.g., Hall et al., 2001; Ubilava and Helmers, 2013). Second, as noted above, weather anomalies associated with El Niño and La Niña phases are not mirror images. Therefore, it is possible that the two phases can affect market conditions in different ways. Moreover, observed weather patterns associated with ENSO events, i.e., *teleconnections*, are likely to be more pronounced during the most extreme episodes of this climate phenomenon. Therefore, large deviations in the ENSO cycle may result in price changes of disproportionate magnitudes, as compared with those associated with more moderate shocks. Finally, the very nature of the production process of some commodities may result in asymmetric price dynamics. For example, one can quickly reduce stocks of crops or livestock in response to new information in the market, but restocking can take an extended period of time (e.g., Holt and Craig, 2006).

In addition to the regime-dependent nonlinearities, the question of structural change in

the ENSO–price relationship has recently gained attention. There are two key reasons for this. First, the increasing knowledge about this climatic phenomenon may be causing different reactions from economic agents. Second, factors other than ENSO (e.g., technological advancements and new policies), which have changed primary commodity price dynamics over recent decades (Enders and Holt, 2012), may have altered the linkages between ENSO and prices. In both cases, it is vital to account for parameter non-constancy in the regression settings to facilitate accurate identification of the relationship between ENSO and commodity prices.

While there may be several options for nonlinear modeling, a convenient econometric approach for the current analysis is the TV-STAR modeling framework of Lundbergh et al. (2003). The concept of smooth transition regressions was pioneered by Bacon and Watts (1971), but Chan and Tong (1986) were the first to introduce a time series variant, the smooth *threshold* autoregression. This was subsequently popularized as the smooth *transition* autoregression (STAR) model. In a series of papers, Luukkonen et al. (1988); Teräsvirta and Anderson (1992); Teräsvirta (1994); Eitrheim and Teräsvirta (1996) formalized the STAR modeling and testing frameworks.

An important characteristic of a family of smooth transition models is that they relax the assumption of instantaneous switching between regimes. That is, these models allow a continuum of points, or thresholds, over which the transition between regimes occurs. Such smoothness usually has an economic interpretation. For example, it can be the result of heterogeneity among economic agents with different degrees of risk aversion who face different transaction costs, or possess varying levels of ability (or willingness) to process ENSO-related information. The structural change can also be gradual—rather than abrupt—given the prolonged nature of technological change and adaptation. Thus, a TV-STAR model can be seen as a generalized framework that nests an array of nonlinear models, as well as a basic linear autoregression, as special cases. To the extent that the exact nature of the ENSO–price relationship is *a priori* unknown, and moreover given that these relationships are likely to vary across different commodity groups, the aforementioned flexibility, afforded by the TV-STAR framework, is a particularly attractive feature of the current modeling exercise. It is worth noting that TV-STAR models, or their restricted variants, have already been successfully applied, albeit in different contexts, to analyzing commodity price behavior (Holt and Craig, 2006; Craig and Holt, 2008; Balagtas and Holt, 2009; Ubilava and Holt, 2013; Hood and Dorfman, 2015). In what follows, I provide a brief description of the econometric specifications and testing framework of the time-varying and regime-dependent models.

A Time Varying Smooth Transition Autoregressive Model

To begin, consider an additive nonlinear (i.e., piecewise linear) univariate time series model:

$$x_t = \theta_{00} + \sum_{i=1}^{p} \theta_{0i} x_{t-i} + \sum_{k=1}^{K-1} \left(\theta_{k0} + \sum_{i=1}^{p} \theta_{ki} x_{t-i} \right) G\left(s_{kt}; \boldsymbol{\vartheta}_k\right) + \varepsilon_t, \tag{1}$$

where x_t is a realization of a random variable at time t; p is the autoregressive lag length; K is the total number of regimes in the model; $G\left(s_{kt}; \boldsymbol{\vartheta}_k\right)$ is a transition function, by construction bounded between zero and one, where s_{kt} is the transition variable, and $\boldsymbol{\vartheta}_k$ is the set of parameters associated with (and defining the type or the curvature of) the transition function; $\varepsilon_t \sim iid(0, \sigma_\varepsilon^2)$ is the white noise process; and the rest are parameters to be estimated. Equation (1) can also incorporate a set of deterministic or exogenous variables, but these have been omitted here, for the sake of brevity.

A set of restrictions can transform equation (1) into various well-known autoregressive models. For example, if $G\left(s_{kt}; \boldsymbol{\vartheta}_k\right) = 0, \ \forall \ k, t$, equation (1) becomes a basic linear autoregression:

$$x_t = \theta_0 + \sum_{i=1}^{p} \theta_i x_{t-i} + \varepsilon_t. \tag{2}$$

If $K = 2$ and $G\left(s_{kt}; \boldsymbol{\vartheta}\right) = I\left(s_{kt}; \boldsymbol{\vartheta}\right)$, i.e., when the transition function is a Heaviside indicator, equation (1) becomes a threshold autoregression (TAR) of Tong and Lim (1980).

Alternatively, if $K = 2$ and $G\left(s_{kt}; \boldsymbol{\vartheta}\right) \in [0, 1]$, i.e., the transition function takes on a continuum of values between zero and one, equation (1) becomes a smooth transition autoregression (STAR) of Luukkonen et al. (1988) and Teräsvirta (1994):

$$x_t = \theta_{00} + \sum_{i=1}^{p} \theta_{0i} x_{t-i} + \left(\theta_{10} + \sum_{i=1}^{p} \theta_{1i} x_{t-i} \right) G\left(s_t; \gamma, c\right) + \varepsilon_t, \tag{3}$$

where $G\left(s_t; \gamma, c\right)$ can take several possible forms. Two commonly applied logistic and exponential transition functions—forming *logistic* STAR (LSTAR) and *exponential* STAR (ESTAR) models, respectively—are given by:

$$G\left(s_t; \gamma, c\right) = \left\{ 1 + \exp\left[-\gamma \left(\frac{s_t - c}{\sigma_s} \right) \right] \right\}^{-1} \qquad \text{(logistic)} \tag{4}$$

$$G\left(s_t; \gamma, c\right) = \left\{ 1 - \exp\left[-\gamma \left(\frac{s_t - c}{\sigma_s} \right)^2 \right] \right\} \qquad \text{(exponential)}. \tag{5}$$

In these smooth transition functions, the parameter vector consists of the smoothness pa-

rameter $\gamma > 0$, and the centrality parameter, c, the latter usually constrained by $[\kappa_s, 1 - \kappa_s]$, where κ_s is a truncation factor, within the range of the transition variable, which ensures sufficient degrees of freedom in each regime of the estimated STAR model. Finally, σ_s is the standard deviation of the transition variable. Its role in these expressions is to normalize the smoothness parameter, γ, to make it effectively unit-free.

A practical benefit of working with STAR models is that they embed previously defined AR and TAR models as special cases. For example, the LSTAR converges to the AR as $\gamma \to 0$, and to the TAR as $\gamma \to \infty$, while the ESTAR converges to the AR both as $\gamma \to 0$ and as $\gamma \to \infty$. Furthermore, if the transition variable is set to be a function of time, e.g., $s_t \equiv t^* = t/T$, equation (3) becomes a time-varying autoregression (TVAR).

Finally, if $K = 3$, i.e., if a model contains two transition functions, of which one is a function of time, then equation (3) becomes a TV-STAR:

$$\begin{aligned} x_t = \theta_{00} + \sum_{i=1}^{p} \theta_{0i} x_{t-i} + \left(\theta_{10} + \sum_{i=1}^{p} \theta_{1i} x_{t-i} \right) G\left(t^*; \gamma_\tau, \tau\right) \\ + \left(\theta_{20} + \sum_{i=1}^{p} \theta_{2i} x_{t-i} \right) G\left(s_t; \gamma_c, c\right) + \varepsilon_t, \end{aligned} \tag{6}$$

where, similar to c, $\tau \in [\kappa_{t^*}, 1 - \kappa_{t^*}]$, and the rest of the variables are as defined above. The TV-STAR framework can be suitable for numerous applications, as the dynamic properties of many economic variables, which include commodity price series, are often simultaneously characterized by time-varying parameter nonconstancy as well as regime-dependent parameter nonlinearity.

Testing Linearity and Parameter Constancy

The question of whether a time-varying or regime-dependent nonlinearity is truly an underlying feature of the data, is empirical and should be tested. But the problem is that the adequate hypothesis, i.e., $H_0 : \gamma = 0$, cannot be examined directly because of unidentified nuisance parameters under the null hypothesis. A case in point is the so-called Davies' problem (Davies, 1977, 1987). For example, consider a STAR model given by equation (3), where $s_t = x_{t-\ell}$ and $\ell \leq p$ is a positive integer known as the delay factor. This equation can be reduced to a linear AR model either by imposing a restriction on the transition parameter, i.e. $\gamma = 0$, or by imposing a restriction on the autoregressive parameters associated with the additive regime of the model, i.e. $\theta_{10} = \theta_{11} = \ldots = \theta_{1p} = 0$. Therefore, the standard test statistics are no longer applicable. As it turns out, the problem can be avoided by replacing

the transition function, $G(s_t;\gamma,c)$, with a third order Taylor series approximation around $\gamma = 0$ (see Luukkonen et al., 1988; Teräsvirta, 1994, for details). The result is a testable auxiliary regression:

$$x_t = \sum_{j=0}^{3}\left(\varphi_{j0} + \sum_{i=1}^{p}\varphi_{ji}x_{t-i}\right)s_t^j + \xi_t, \tag{7}$$

where $\varphi_{ji}, \quad \forall \quad j,i$, are functions of the parameters of the original STAR model, and ξ_t combines the original error term, ε_t, and the approximation error resulting from the Taylor series expansion. The linearity test is now equivalent to testing the null hypothesis of $H_0^{'}$: $\varphi_{ji} = 0, \forall\ j = 1,\ldots,3,\ i = 0,\ldots p$. Additionally, tests against LSTAR and ESTAR models are also embedded in the testing framework. In particular, $H_{03}^{'}$: $\varphi_{3j} = 0$ and $H_{01}^{'}$: $\varphi_{1j} = 0 \mid \varphi_{kj} = 0, \forall\ k = 2,3;\ j = 1,\ldots,p$, are tests against LSTAR; while $H_{02}^{'}$: $\varphi_{2j} = 0 \mid \varphi_{3j} = 0,\ j = 1,\ldots,p$, is a test against ESTAR. The appropriate model is selected based on the probability values of the above hypotheses (e.g., Teräsvirta, 1994).

Several features associated with nonlinear model selection need to be noted. First, the transition variable, s_t, is often *a priori* unknown. In such instances, a set of candidate transition variables is considered, and the most appropriate one is selected based on probability values associated with the null hypothesis for linearity. In conjunction with this, one may also estimate several candidate models, and decide on a suitable transition variable and type of nonlinear function based on the model fit (e.g., Akaike Information Criterion), as well as the remaining nonlinearity test results. Second, the parameter constancy test is a special case of the linearity test, where s_t is substituted by t^* in equation (7). Typically, parameter nonconstancy is addressed—i.e., a TVAR is estimated where applicable—before moving on to testing and estimating a STAR or a TV-STAR. Finally, if a TV-STAR is a suspect, the so-called *specific-to-general* approach can be implemented. This involves estimating the multiple-regime model (e.g., van Dijk and Franses, 1999), whereby both time-varying and regime-dependent components are incorporated as transition functions (see, e.g., Lundbergh et al., 2003; Holt and Craig, 2006).

Data

In this study, I use monthly series of a measure of ENSO intensity and primary commodity prices spanning the period from January 1980 to December 2016. Two better-known indices depicting ENSO cycles are sea surface temperature (SST)-based measures in the equatorial Pacific region, and the air-pressure-based Southern Oscillation Index (SOI) calculated from

measures near Tahiti and Darwin. Previous studies have used either or both of these indices (e.g., Brunner, 2002; Ubilava and Holt, 2013; Cashin et al., 2017), but the SST anomalies derived from observations collected in the `Niño3.4` region—a rectangular area between 5°N and 5°S and between 170°W and 120°W—has become a more commonly applied measure of ENSO intensity in the climatology and climate economics literature (e.g., Hsiang et al., 2011; Hsiang and Meng, 2015). In this study, I also rely on the SST-based measure,[2] which depicts deviations from the average historic temperatures in a given month over the centered 30-year base periods, and is tabulated by the Climate Prediction Center at the National Oceanic and Atmospheric Administration (see Appendix figure A1 for the visual comparison of the SST and SOI measures).

The ENSO cycle is said to be in an El Niño or a La Niña phase, if a three-month running mean of the SST anomalies—a measure referred to as the Oceanic Niño Index (ONI)—exceeds (in absolute terms) a threshold of ±0.5°C for a minimum of five consecutive periods. Thus, persistent positive deviations in the ONI constitute an El Niño event, and persistent negative deviations amount to a La Niña episode. These events, along with the monthly SST anomalies, are shown in figure 2.

Primary commodity price series are obtained from World Bank and International Monetary Fund publications. These are spot prices (FOB or CIF), and are indicative of world prices for the various commodities. I consider several important commodity groups, including cereal grains, forestry, farms and fishery, vegetable oils and meals, and industrial and rare metals. See table 1 for a complete list and a brief description of commodities examined in this article. For the purposes of the analysis, the nominal spot prices are deflated using the U.S. producer price index, obtained from the U.S. Bureau of Labor Statistics, and then transformed to natural logarithms. Hereafter, unless otherwise stated, *prices* refer to the natural logarithm of real commodity prices.

Model Selection and Estimation

The primary building block of most nonlinear models, and certainly of the one used in this study, is a basic linear specification. Let z_t denote the measure of ENSO intensity, i.e., the SST anomaly, in period t; and let y_t denote the log-transformed real price of a commodity

[2]The Pearson correlation coefficient between the SST and SOI is approximately 0.8, suggesting that the two indices will possibly yield comparable, but not identical results. See the Appendix for discrepancies in model selection.

Table 1: Description and Origin of the Considered Commodity Prices

Commodity	Description and Origin
Aluminum	99.5% minimum purity, LME spot price, CIF UK ports
Barley	Canadian No.1 Western Barley, spot price
Beef	Australian and New Zealand 85% lean fores, CIF U.S. import price, (c/lb)
Chicken	Whole bird spot price, Ready-to-cook, whole, iced, Georgia docks, (c/lb)
Cocoa	International Cocoa Organization cash price, CIF US and European ports
Coconut Oil	Philippines/Indonesia, in bulk, CIF Rotterdam
Coffee (Arabica)	International Coffee Organization New York cash price, ex-dock New York, (c/lb)
Coffee (Robusta)	International Coffee Organization New York cash price, ex-dock New York, (c/lb)
Copper	Grade A cathode, LME spot price, CIF European ports
Copra	Philippines/Indonesia, in bulk, CIF NW European ports
Cotton	Cotton Outlook 'A Index', Middling 1-3/32 inch staple, CIF Liverpool, (c/lb)
Fishmeal	Peru pellets, 65% Protein, CIF
Gold	99.5% refined, London afternoon fixing
Groundnut Oil	Any origin, CIF Rotterdam
Hard Logs	Best quality Malaysian Meranti, import price Japan, ($/m^3)
Hard Sawnwood	Dark Red Meranti, select and better quality, C&F U.K port, ($/m^3)
Hides	Heavy native steers, over 53 pounds, wholesale dealer's price, FOB Chicago, (c/lb)
Lamb	Frozen carcass, Smithfield London, (c/lb)
Lead	99.97% pure, LME spot price, CIF European Ports
Maize	U.S. No.2 Yellow, FOB Gulf of Mexico
Nickel	Melting grade, LME spot price, CIF European ports
Olive Oil	Extra virgin less than 1% free fatty acid, ex-tanker price U.K.
Palm Oil	Malaysia/Indonesia, in Bulk, 5% FFA, CIF NW European Ports
Platinum	99.9% refined, London afternoon fixing
Pork	51-52% lean Hogs, U.S. price (c/lb)
Rapeseed Oil	Crude, FOB Rotterdam
Rice	Thailand 5% broken milled white rice, FOB Bangkok
Salmon	Farm Bred Norwegian Salmon, export price, ($/kg)
Silver	99.9% refined, London afternoon fixing
Soft Logs	Average export price, U.S. for Douglas Fir, ($/m^3)
Soft Sawnwood	Average export price, U.S. for Douglas Fir, ($/m^3)
Sorghum	U.S. No.2 milo yellow, FOB Gulf ports
Soybean Meal	Argentine 45/46% extraction (after January 1990); U.S. 44%, CIF Rotterdam
Soybean Oil	Any origin, crude, FOB ex-mill Netherlands
Soybeans	U.S. No.2 Yellow, CIF Rotterdam
Sugar	ISA daily price, raw, FOB and stowed at greater Caribbean Ports ($/kg)
Sunflowerseed Oil	European Union, FOB NW European ports
Tea	Kenya auction price (after July 1998); London auctions, CIF U.K. warehouses, (c/kg)
Tin	Standard grade, LME spot price
Tobacco	Any origin, unmanufactured, general import, CIF U.S.
Wheat	U.S. No.1 Hard Red Winter, Ordinary Protein, FOB Gulf of Mexico
Wool	Coarse, 23 micron, Australian Wool Exchange spot quote (c/kg)
Zinc	High grade 98% pure

Note: the commodity prices are denominated in US$ per metric ton, unless otherwise specified. The series were sourced from World Bank and International Monetary Fund online data portals.

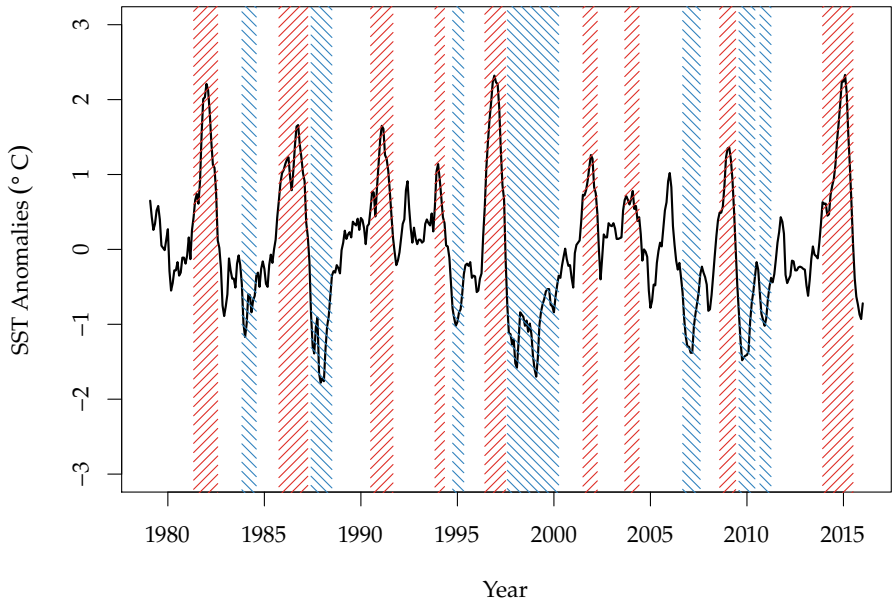

Figure 2: SST anomalies and ENSO events

Note: Regions shaded with red (upward-sloping) lines represent the observed El Niño events; regions shaded with blue (downward-sloping) lines represent the observed La Niña events.

in the same period.[3] The first step is to identify the degree of integration of the time series under consideration. Both the augmented Dickey–Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests suggest that z_t is stationary in levels, i.e., integrated of order zero. So, a linear specification for the SST anomaly is an autoregressive process, augmented with seasonal binary variables:

$$z_t = \alpha_0 + \sum_{s=1}^{11} \alpha_s D_{st} + \sum_{i=1}^{r} \beta_i z_{t-i} + \nu_t, \tag{8}$$

where D_{st} is a monthly binary variable, and r is selected using the bottom-up sequential testing approach to mitigate any serial correlation in ν_t, so that $\nu_t \sim iid\,(0, \sigma_\nu^2)$.

As for commodity prices, the tests suggest that some are stationary in levels, i.e., integrated of order zero, while others follow a unit root process, i.e., are integrated of order one (see Appendix table A1 for details). However, the ADF and KPSS tests rely on the assumption that nonlinearity or parameter nonconstancy are not a feature of the data (e.g., Balagtas and Holt, 2009). Indeed, when applying the Zivot–Andrews (ZA) test, which allows

[3]Here, and in what follows, the notation omits the commodity-specific subscripts for simplicity.

for a structural break in the time series, one can reject the null hypothesis of a unit root for a subset of commodity prices that previously appeared to be integrated of order one. But this hardly addresses all the peculiarities associated with unit root testing under nonlinearity or parameter non-constancy—an area that is still evolving (for details, and a brief review of the literature, see Demetrescu and Kruse, 2013).

This study abstracts from the foregoing complexity by simply allowing the possibility of unit roots as well as stationarity to be a characteristic feature of commodity prices. That is, while the prices of some commodities (e.g., metals) are likely to be better characterized by a unit root process, both theory and empirical evidence support the notion of stationarity for other groups of commodities (see, e.g., Wang and Tomek, 2007; Enders and Holt, 2012). Therefore, I apply the following decision rule. For a given commodity price series, if either the ADF or the ZA test rejects the null hypothesis of a unit root, the autoregressive process is modeled in levels, otherwise—in first-differences. That is, for a stationary time series the following linear autoregressive model is applied:

$$\Delta y_t = \delta_0 + \sum_{s=1}^{11} \delta_s D_{st} + \sum_{j=0}^{q} \eta_j z_{t-j} + \sum_{i=1}^{p-1} \phi_i \Delta y_{t-i} + \rho y_{t-1} + \varepsilon_t, \tag{9}$$

where Δ is a first-difference operator; p is selected using the bottom-up sequential testing approach to mitigate any serial correlation in ε_t, so that $\varepsilon_t \sim iid\,(0, \sigma_\varepsilon^2)$; and q is selected based on the sample-size-corrected Akaike information criterion (AICc). For difference-stationary price series, the unit root is imposed by setting $\rho = 0$.

Note that equation (9) is, in essence, an autoregressive distributed lag model, ARDL(p,q), augmented with seasonal dummy variables. Notably, this nests an autoregressive model, AR(p), as a special case. Alternatively, equations (8) and (9) together can be seen as a restricted bivariate vector autoregressive process, wherein the SST anomaly is assumed to be weakly exogenous to prices—an assumption that is hardly controversial and is supported by empirical evidence (see, e.g., Brunner, 2002). Consistent with the exogeneity assumption—which, moreover, serves as an identification condition in the bivariate system of equations—the SST shocks and price innovations are uncorrelated, i.e., $Cov(\varepsilon_t, \upsilon_t) = 0$.

Notably, a number of factors that are not included in the currently specified model are correlated with commodity prices, and may also be correlated with the SST anomaly. These omitted factors may be *covariates* or *mediators*, but are certainly not *confounders*. That is, the omitted variables are either uncorrelated with SST (covariates) or are caused by SST (mediators). Therefore, a causal mechanism, if any, from the SST anomaly to prices, is

unidirectional, passing through a (set of) mediator(s). To the extent that the SST anomaly is exogenous to mediators that are not included in the equation, the model does not suffer from omitted variable issues (e.g., Edelstein and Kilian, 2009). As for covariates, it may be that they spuriously correlate with the ENSO cycles. Given the frequency of the data used, this could happen if, for example, a covariate has a seasonal pattern (as does the measure of ENSO intensity). The currently applied modeling framework addresses this issue by incorporating seasonal binary variables in the equations. Alternatively, the correlation can take an "episodic" form, wherein an episode is a single event rather than an (irregularly) repeated event. But this should not be a cause for concern in relation to long time series, such as those applied in the current study.

Once the linear models, as specified in equations (8) and (9), are identified, we can turn to linearity and parameter constancy testing using the auxiliary regression framework as outlined above. The nonlinear model identification proceeds as follows: (i) Test the null hypotheses of linearity and parameter constancy based on equation (7). If the null of parameter constancy is rejected, estimate a time-varying autoregressive (distributed lag) model using a nonlinear least squares method. (ii) Test the null hypothesis of no remaining nonlinearity. If, at this point, the null hypothesis of remaining nonlinearity is rejected, estimate a time-varying smooth transition autoregressive (distributed lag) model. Alternatively, if in (i), we fail to reject the null hypothesis of parameter constancy, but reject the null hypothesis of linearity, estimate a smooth transition autoregressive (distributed lag) model. See Appendix figure A2 for an illustration of the model selection algorithm.

Following the aforementioned steps, one can identify and estimate a range of linear and nonlinear models, which can be summarized by the following equation:

$$\begin{aligned}\Delta y_t = {} & \delta_0 + \sum_{s=1}^{11} \delta_{0s} D_{st} + \sum_{j=0}^{q} \eta_{0j} z_{t-j} + \sum_{i=1}^{p-1} \phi_{0i} \Delta y_{t-i} + \rho_0 y_{t-1} \\ & + \left[\delta_1 + \sum_{s=1}^{11} \delta_{1s} D_{st} + \sum_{j=0}^{q} \eta_{1j} z_{t-j} + \sum_{i=1}^{p-1} \phi_{1i} \Delta y_{t-i} + \rho_1 y_{t-1} \right] G(t^*; \gamma_\tau, \tau) \\ & + \left[\delta_2 + \sum_{s=1}^{11} \delta_{2s} D_{st} + \sum_{j=0}^{q} \eta_{2j} z_{t-j} + \sum_{i=1}^{p-1} \phi_{2i} \Delta y_{t-i} + \rho_2 y_{t-1} \right] G(s_t; \gamma_c, c) + \varepsilon_t,\end{aligned} \tag{10}$$

where $G(t^*, \gamma_\tau, \tau)$ and $G(s_t, \gamma_c, c)$ respectively are the logistic or exponential transition functions of $t^* = t/T$ and $s_t = z_{t-\ell}$, and $0 \leq \ell \leq q$. In the case of logistic transition functions, the smoothness parameter, γ, is bounded by values of two and 100 (a value less than two typically results in dynamics that can be indistinguishable from that of a linear model, while

a value that is close to or greater than 100, results in an instantaneous switch between the regimes, irrespective of the actual estimated value of the smoothness parameter), while for exponential transition functions, γ is bounded by values of one and 10. As for the location parameters, in the case of logistic functions, c is bounded by the 15th and 85th percentiles of the transition variable, while in exponential functions, c is bounded by the 25th and 75th percentiles of the transition function. Equation (10) is a TV-STARDL representation, within which is nested a STARDL model when $G(t^*, \gamma_\tau, \tau) = 0, \ \forall\, t$, a TVARDL model when $G(s_t, \gamma_c, c) = 0, \ \forall\, t$, and an ARDL model when both transition functions are set to zero.

Note that the transition function, $G(s_t, \gamma_c, c)$, is conditioned on the state of ENSO, rather than on (a function of) lagged prices (as in, e.g., Ubilava, 2012; Ubilava and Holt, 2013). Thus, the current specification allows for changes in ENSO intensity to manifest into commodity price dynamics through linear—(i.e., the distributed lag component)—as well as nonlinear—i.e., the transition function—channels. Moreover, the time-varying component allows for potentially smooth transition between the regimes as a result of structural change. This modeling approach is similar to that applied by Enders and Holt (2012). But unlike their modeling framework, the current specification also accounts for possible changes in autoregressive as well as deterministic components of the price dynamics.

Empirical Results and Discussion

Using the strategy outlined in the previous section, it was found that the dynamics of the SST anomalies are best characterized by a logistic STAR specification. As for the commodity prices under consideration, a total of 12 nonlinear models were estimated. In seven of these models there is evidence of linkages between SST anomalies and commodity prices (the other five models identify structural changes in the price dynamics). Furthermore, in an additional set of 12 commodities, there is an indication of linear linkages between ENSO and prices. Table 2 summarizes the estimated models and the associated diagnostic test results.

Several general tendencies are apparent from these results. First, as expected, the prices of agricultural commodities are most affected by SST anomalies. This is particularly true for vegetable oils and oilseeds produced in the tropics. The prices of tropically grown beverages, including coffee varieties and cocoa, are also affected by SST anomalies, as are fishmeal and salmon prices. Notably, there is no evidence of any effect on the prices of cereal grains. There are a couple of possible reasons for this. First, cereal grains are predominantly produced in temperate regions, where the weather effects of ENSO are minimal (Hsiang and Meng,

Table 2: Selected Models and Residual Diagnostics

Commodity	Model	p	q	ℓ	d	k	$\hat{\gamma}_c$	$\hat{c}$	$\hat{\gamma}_\tau$	$\hat{\tau}$	p_{t^*}	p_{AC}	p_{ARCH}
Aluminum	ARDL	2	1		0	16					0.253	0.226	<0.000
Barley	AR	2			0	14					0.117	0.506	<0.000
Beef	AR	3			1	14					0.507	0.532	0.592
Chicken	LTVARDL	2	2		0	36			2.00 (1.16)	0.15 (0.21)	0.017	0.363	<0.000
Cocoa	ARDL	3	0		1	15					0.541	0.570	0.016
Coconut Oil	ARDL	2	5		1	19					0.610	0.967	0.422
Coffee (Arabica)	ARDL	3	1		0	17					0.944	0.690	<0.000
Coffee (Robusta)	ARDL	2	1		1	15					0.732	0.477	0.024
Copper	AR	2			1	13					0.945	0.059	0.003
Copra	ARDL	2	5		1	19					0.194	0.537	0.923
Cotton	AR	2			0	14					0.859	0.255	<0.000
Fishmeal	ARDL	2	0		1	14					0.126	0.144	0.072
Gold	AR	2			1	13					0.013	0.163	0.055
Groundnut Oil	AR	2			0	14					0.065	0.838	0.138
Hides	LSTARDL	2	8	6	1	46	33.0 (33.5)	− 0.48 (0.03)			0.508	0.073	0.001
Lamb	LTVAR	3			1	30			2.00 (1.28)	0.15 (0.21)	0.073	0.729	0.101
Lead	AR	2			1	13					0.583	0.960	<0.000
Logs (Hard)	LTVAR	2			0	30			100 (271)	0.22 (0.01)	0.116	0.066	<0.000
Logs (Soft)	AR	2			1	13					0.204	0.553	0.003
Maize	AR	2			0	14					0.936	0.790	0.827
Nickel	AR	2			1	13					0.932	0.305	0.980
Olive Oil	LTVAR	2			1	28			13.6 (10.6)	0.85 (0.02)	0.622	0.053	0.001
Palm Oil	ARDL	5	6		0	24					0.383	0.975	0.551
Platinum	AR	2			1	13					0.204	0.765	0.002
Pork	LTVAR	1			0	28			3.66 (1.66)	0.27 (0.05)	0.775	0.221	<0.000
Rapeseed Oil	ESTARDL	2	6	6	0	44	10.0 (4.53)	− 0.17 (0.03)			0.206	0.656	<0.000
Rice	AR	3			0	15					0.060	0.985	<0.000
Salmon	LTVARDL	2	4		1	38			3.72 (1.80)	0.78 (0.06)	0.766	0.445	0.079
Sawnwood (Hard)	LSTARDL	3	7	5	1	46	89.9 (258)	0.91 (0.02)			0.227	0.149	0.458
Sawnwood (Soft)	ESTARDL	3	0	0	1	32	1.00 (0.42)	− 0.04 (0.07)			0.007	0.292	0.107
Silver	ARDL	2	0		1	14					0.044	0.372	0.035
Sorghum	AR	2			0	14					0.159	0.845	0.928
Soybean Meal	AR	2			0	14					0.366	0.389	0.082
Soybean Oil	ARDL	3	6		1	21					0.129	0.971	0.063
Soybeans	ARDL	2	1		0	16					0.453	0.276	0.048
Sugar	AR	3			0	15					0.358	0.183	0.148
Sunflowerseed Oil	LSTARDL	3	5	1	0	44	3.98 (1.98)	− 0.82 (0.20)			0.601	0.223	0.754
Tea	LTVAR	2			0	30			3.71 (2.55)	0.38 (0.06)	0.673	0.222	0.937
Tin	AR	3			1	14					0.713	0.673	0.542
Tobacco	AR	2			1	13					1.000	0.159	<0.000
Wheat	AR	2			1	13					0.797	0.235	0.001
Wool	ARDL	2	1		1	15					0.232	0.112	0.531
Zinc	AR	2			0	14					0.945	0.715	0.655

Note: p indicates the autoregressive lag length; q indicates the distributed lag length; ℓ denotes the delay factor of the transition variable in a regime-dependent model; d indicates the order of integration of the price series; k indicates the total number of estimated parameters; $\hat{\gamma}_c$, $\hat{c}$, $\hat{\gamma}_\tau$, and $\hat{\tau}$ are estimated smoothness and location parameters (values in parentheses are standard errors); p_{t^*}, p_{AC}, and p_{ARCH} denote probability values associated with hypotheses of (remaining) parameter nonconstancy, residual autocorrelation, and autoregressive conditional heteroskedasticity.

2015). Second, as a result of the regional diversification of crop production and the effects of trade, losses in one region are more or less offset by gains in another region (see, e.g., Lybbert et al., 2014; Glauber and Miranda, 2016). While the modeling framework used in this study cannot differentiate between factors that could mitigate price responses to SST shocks, this finding points to the important observation that the prices of cereal grains are not as responsive to ENSO shocks as commonly thought.

Of the nonagricultural commodities, timber and a small subset of metals appear to be affected by SST anomalies. For timber, the impact is dynamically complex. That is, regime-dependent nonlinearities are the underlying features of the dynamic process. Unlike agricultural commodities, where the price effect is primarily supply-driven, for these industrial commodities demand-side factors, in addition to supply-side shocks, such as production disruptions resonated from ENSO-induced heat and drought, may also play a role in this relationship.

Figure 3 shows the estimated transition functions associated with nonlinear models. In the majority of cases, the transition between regimes appears to be gradual. In regime-dependent models, because the transition variable is the SST anomaly, the estimated location parameter, $\hat{c}$, indicates the point of inflection (or switching point) between the two ENSO phases. The estimated smoothness parameter, $\hat{\gamma}_c$, indicates how quickly the model dynamics change as one ENSO phase diminishes and the other begins to emerge. In time-varying models, the estimated location parameter, $\hat{\tau}$, identifies the point in time around which the price dynamics have altered, whereas the smoothness parameter, $\hat{\gamma}_\tau$ defines the time frame for the structural change. Table 2 presents these parameter estimates.

The estimated parameters of a nonlinear model, other than those of a transition function, cannot be interpreted directly. This is because nonlinear models are not invariant to idiosyncratic shocks that may alter the underlying dynamics of a stochastic process. This also implies that the so-called naïve extrapolation, which is used in linear models to generate impulse–response functions at horizons greater than one, yields biased results, and is not valid in the case of nonlinear models. Analytical expressions of impulse–response functions are infeasible for horizons greater than one. To circumvent this issue, Koop et al. (1996) proposed a numerical approximation technique that produces a generalized impulse–response function (GIR) for a given history and shock:

$$\pi_x\left(h, \upsilon, \omega_{t-1}\right) = \mathbb{E}\left(x_{t+h}|\upsilon, \omega_{t-1}\right) - \mathbb{E}\left(x_{t+h}|\omega_{t-1}\right), \quad h = 0, 1, \ldots, \tag{11}$$

where $\pi_x\left(h, \upsilon, \omega_{t-1}\right)$ is a GIR for variable x at horizon h; $\omega_{t-1} \in \Omega_{t-1}$ denotes the history,

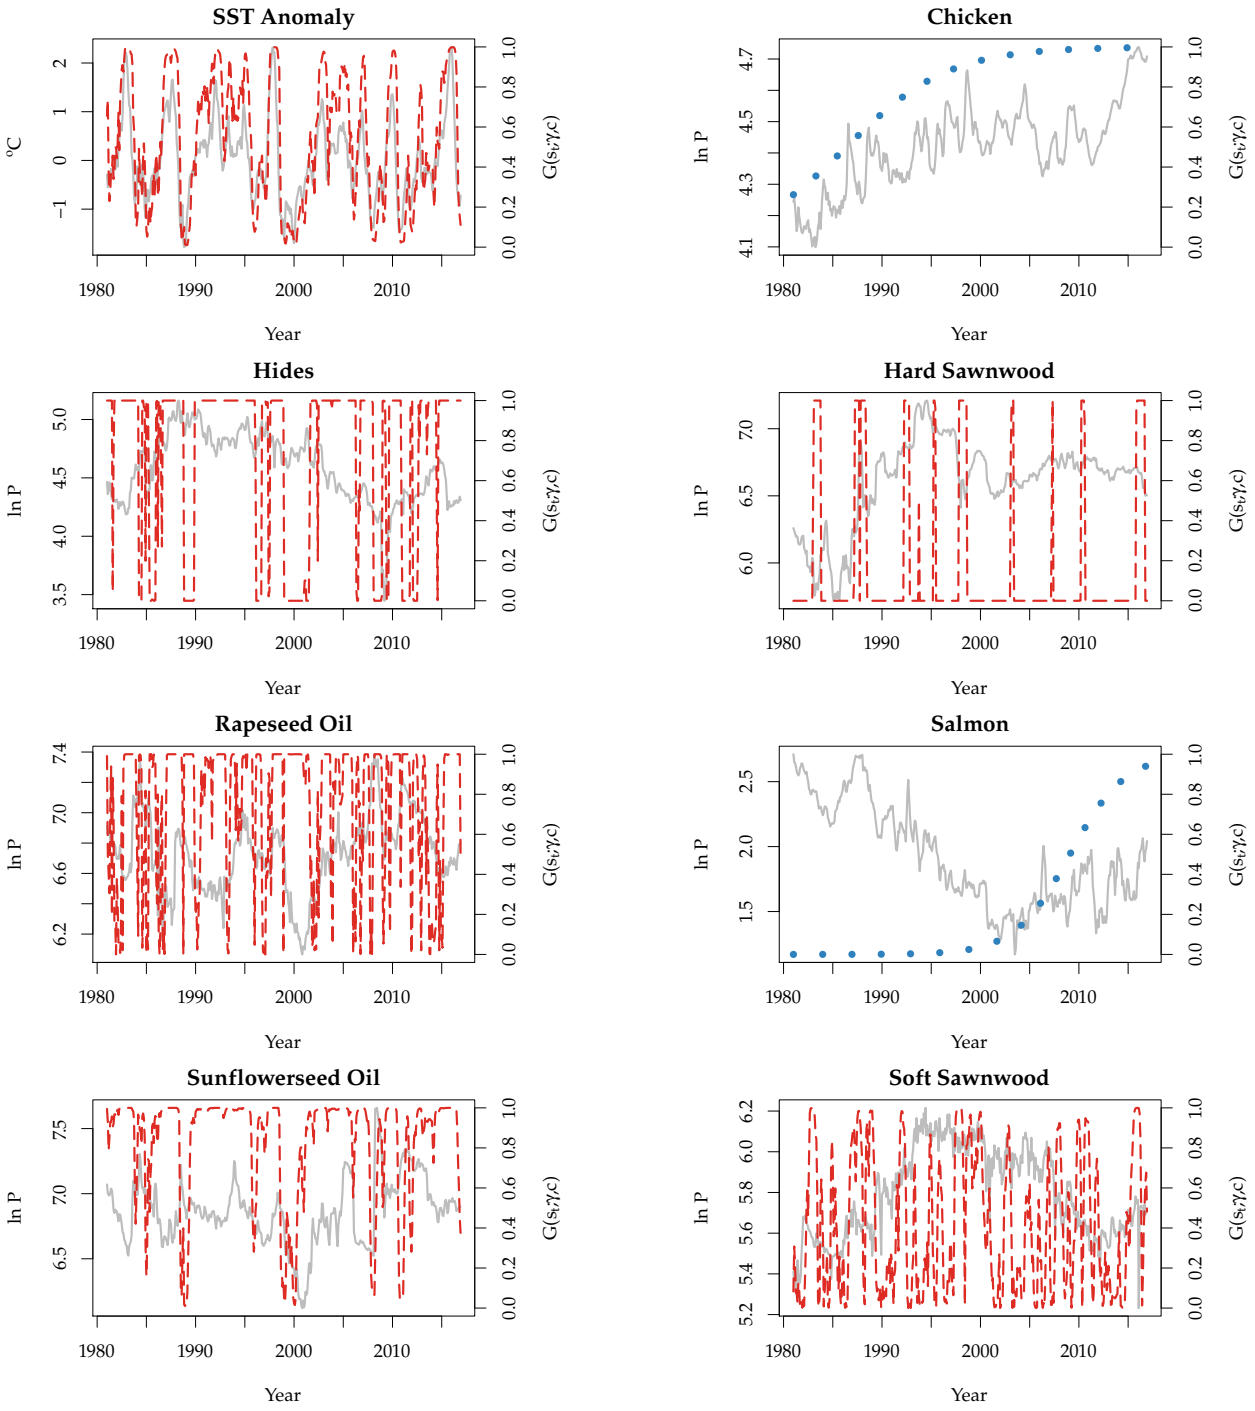

Figure 3: Estimated transition functions of ENSO and price equations

Note: These plots feature the SST anomaly (°C) and the commodity price series (natural logarithm), along with the associated estimated transition functions. The SST anomaly and the commodity price series are depicted by solid gray lines; the regime-dependent transition functions, i.e., $G(s_t; \gamma_c, c)$, are depicted by dashed red lines; and the time-varying transition functions, i.e., $G(t^*; \gamma_\tau, \tau)$, are depicted by blue dots.

i.e., available information at a time when a forecast is made, and $\upsilon \in \Upsilon$ denotes an impulse, i.e., the realization of an initial shock, from the distribution of shocks under consideration; $\mathbb{E}$ is an expectation operator. To the extent that υ and ω_{t-1} are realizations of the random variables Υ and Ω_{t-1}, the GIR given by $\pi_x(h, \upsilon, \omega_{t-1})$, which is the difference between two random variables, also represents a realization of the random variable:

$$\pi_x(h, \Upsilon, \Omega_{t-1}) = \mathbb{E}(x_{t+h}|\Upsilon, \Omega_{t-1}) - \mathbb{E}(x_{t+h}|\Omega_{t-1}), \quad h = 0, 1, \ldots. \tag{12}$$

In principle, Ω_{t-1} can contain every history from the available index set, but a subset of histories, $\Omega'_{t-1} \subseteq \Omega_{t-1}$, can also be applied to obtain conditional expectations. For example, if one is interested in price dynamics during El Niño conditions, only the histories associated with this phase are sampled. A similar logic applies to Υ. For example, one may be interested in price dynamics after a shock that exceeds one or two standard deviations of the available set of impulses.

The Price Effects of SST Anomalies across the ENSO Regimes

I apply a bootstrap resampling algorithm to approximate GIRs.[4] To illustrate the regime-dependent nonlinearities/asymmetries, two subsets of histories, each of size 50, are sampled from the El Niño and La Niña regimes, as depicted in figure 2. For a randomly sampled history from each of these subsets, 100 bootstrap projections of the SST anomaly are computed with and without an initial shock at $h = 0$. This initial shock is randomly sampled from a uniform distribution bounded by $0.5\hat{\sigma}_\nu$ and $1.5\hat{\sigma}_\nu$, where $\hat{\sigma}_\nu$ is the residual standard deviation from the estimated STAR model of ENSO. By using a distribution of shocks—rather than a single value—the GIRs account for nonlinearities in the underlying autoregressive process, so that for $h > 0$ the distribution of the GIRs may have an asymmetric or multimodal form. At each horizon $h > 0$, the aforementioned projections are disturbed by idiosyncratic shocks (randomly sampled from residuals of the estimated STAR model of ENSO). The averages of these bootstrap projections form the conditional expectations of the SST anomaly, $\mathbb{E}(z_{t+h}|\upsilon, \omega_{t-1})$ and $\mathbb{E}(z_{t+h}|\omega_{t-1})$, at each horizon. The difference between these two expectations yields a realization of the GIR, $\pi_z(h, \upsilon, \omega_{t-1})$.

The aforementioned conditional expectations of the SST anomaly are then incorporated into the commodity price equations to generate their GIRs (using an approach similar to

[4]See Koop et al. (1996) for more details on generalized impulse–response functions. See the Appendix for the specific algorithm used in this analysis.

that outlined above, but with innovations randomly sampled from residuals of the estimated model of commodity prices). A total of 4000 GIRs are generated for the SST anomaly and prices. Finally, because prices are modeled as first-differences, their GIRs are integrated over the length of the horizon to obtain the effect of ENSO shocks on log-levels of real commodity prices. That is:

$$\pi_y\left(h, \upsilon, \omega_{t-1}\right) = \sum_{s=0}^{h} \pi_{\Delta y}\left(s, \upsilon, \omega_{t-1}\right). \tag{13}$$

As noted above, the characteristic features of GIRs from regime-dependent nonlinear models are asymmetry and multimodality. For multimodality in particular, the conventional confidence intervals may not be informative. Instead, the HDRs of Hyndman (1995, 1996) are better suited to illustrating these distributional peculiarities.[5] Here, I consider the 50% and 90% HDRs to compactly illustrate the distribution of GIRs for ENSO and commodity prices.

[5]The HDR, or more precisely a $100\left(1-\alpha\right)\%$ HDR, is the subset, $R\left(f_\alpha\right)$, of the sample space of the random variable X, such that $R\left(f_\alpha\right) = \{x : f(x) \geq f_\alpha\}$, where $f(x)$ is the density function of X, and f_α is the largest constant that satisfies $P\left[X \in R\left(f_\alpha\right)\right] \geq 1-\alpha$. For further details, see Hyndman (1996). For an illustration, see the Appendix.

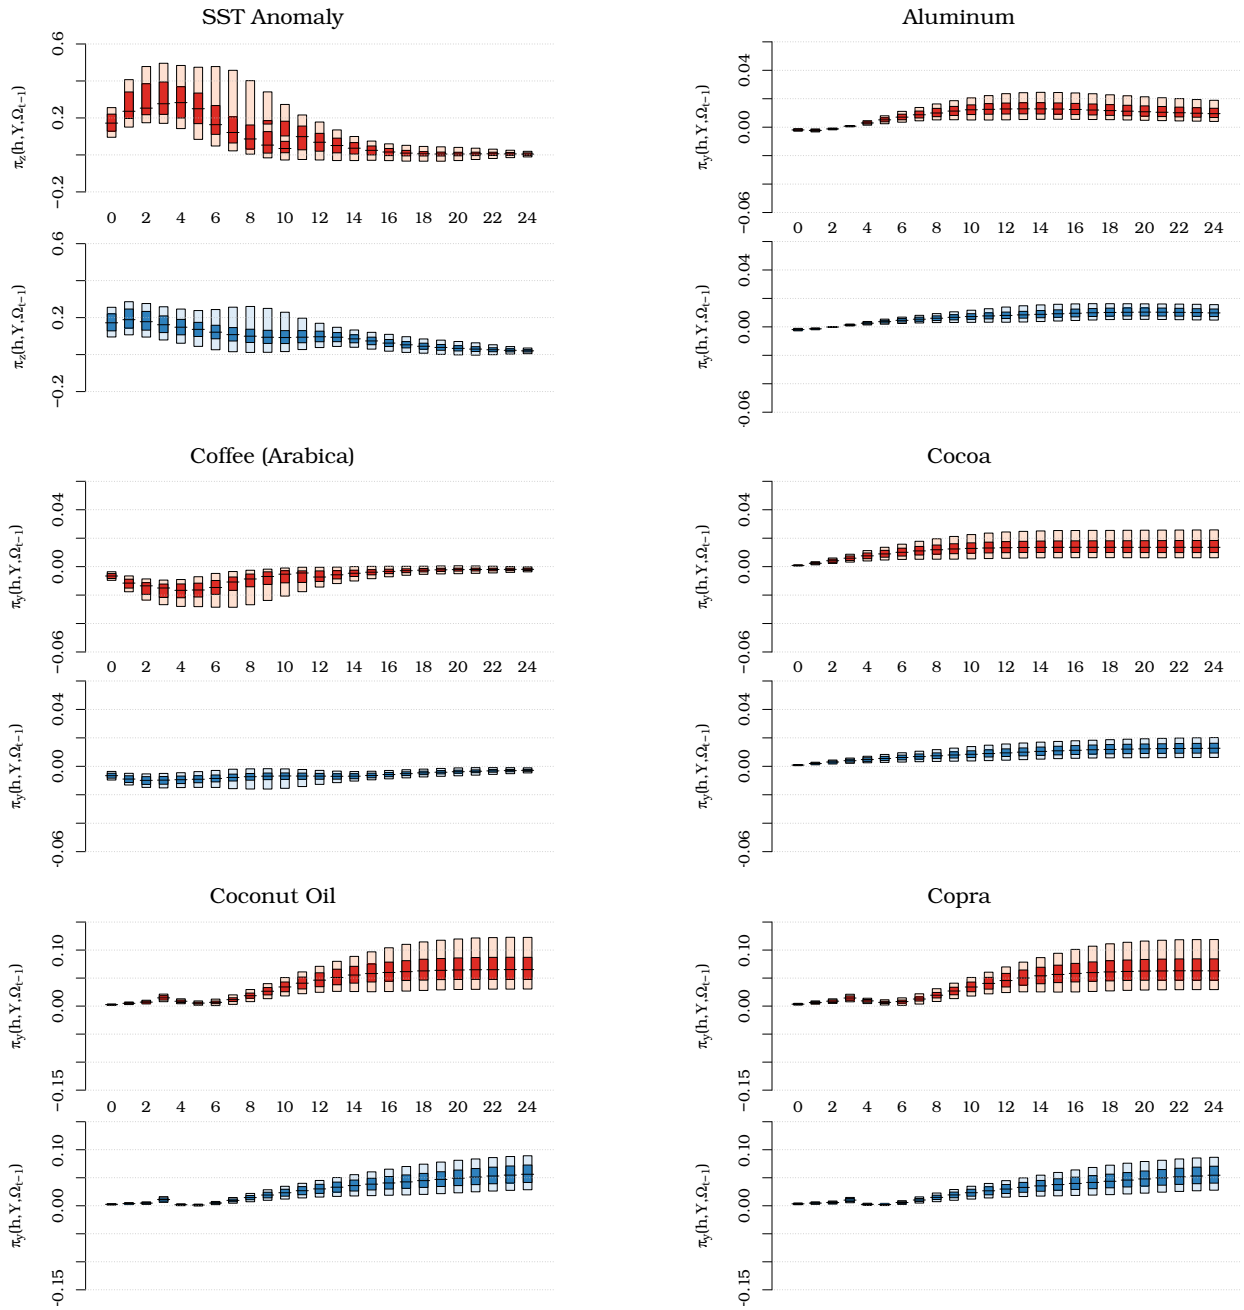

Figure 4: Regime-dependent asymmetries in SST and selected commodity prices

Note: These plots feature 50% (darker shade) and 90% (lighter shade) highest density regions (HDRs) of the generalized impulse–response functions (GIRs) at horizon h, $\forall\ h = 0, 1, \ldots, 24$. In each plot, the GIRs are associated with an average one standard deviation SST shock during the El Niño regime (top panel, shades of red) and the La Niña regime (bottom panel, shades of blue).

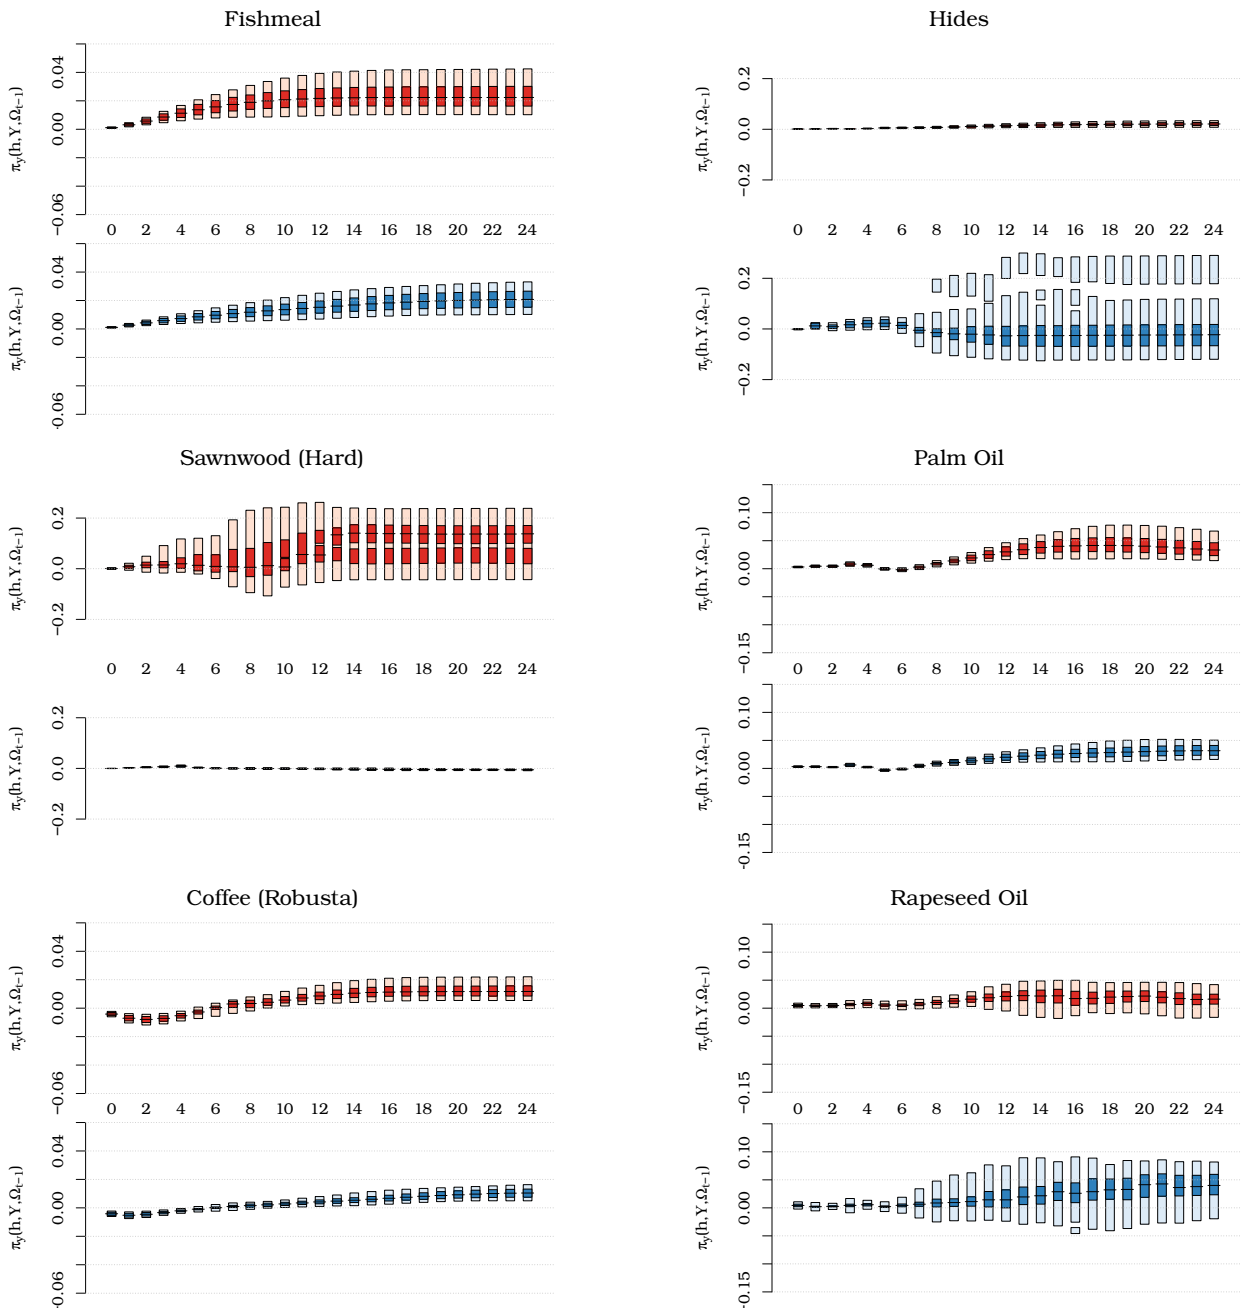

Figure 4: continued from the previous page

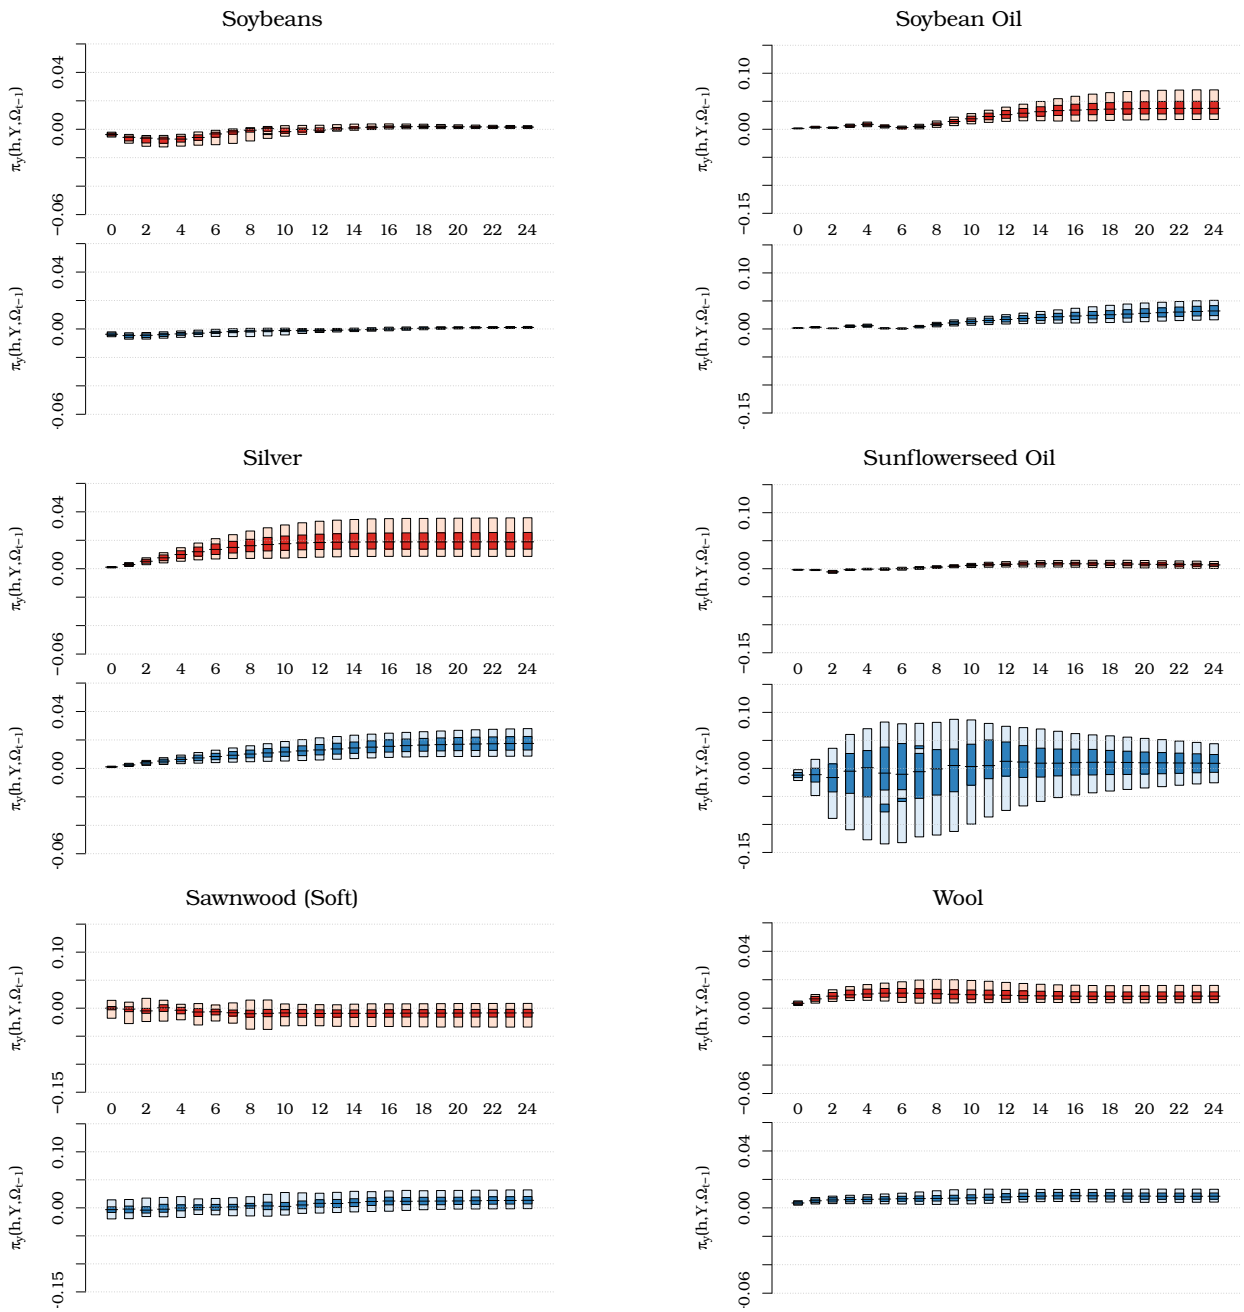

Figure 4: continued from the previous page

These HDRs represent distributions of GIRs that result from an unexpected shock to SST anomalies where the shock is, on average, equivalent to one standard deviation of the residuals of the estimated ENSO model.[6] As for magnitude of impulses and responses, the

[6]Note that these distributions of impulse–responses are not due to parameter uncertainty, but rather

mean value of SST shocks used in this analysis is approximately 0.18°C; the price responses are given in relative terms (i.e., on a log-scale). Several general observations stem from this figure. First, for the vast majority of the commodities under consideration, a positive SST shock tends to increase prices. Second, the effect is more pronounced in response to shocks during an El Niño regime, and is predominantly driven by more amplified GIRs of SST anomalies in this regime compared with those during a La Niña regime. A one standard deviation positive SST shock results in a two-to-ten percent increase in the prices of vegetable oils and oilseeds in the long run. This effect is more apparent for vegetable oils produced in the tropics. The price effect is in the range of one-to-three percent for most other affected agricultural commodities. The majority of nonagricultural commodities, with the exception of timber varieties, aluminum and silver, are unaffected by SST deviations. With regard to the aforementioned exceptions, for timber, the effect is nonlinear in that it varies across ENSO regimes. In the case of the metals, the effect appears to be economically and statistically negligible. This is revisited below in the context of forecast accuracy assessment.

The Price Effects of SST Anomalies Within a Calendar Year and Over Time

For further insights, I now turn to history-specific GIRs. First, consider subsets of histories during the May–June–July (MJJ) trimester and the November–December–January (NDJ) trimester. The former is the onset of the ENSO cycle, while the latter represents the peak of the ENSO season. Additional features of interest emerge from these GIRs (see Figure 5). First, seasonal discrepancies in the ENSO effect are apparent. The price impact is larger for shocks that occur during the MJJ trimester compared with those that occur during the NDJ trimester. This is expected for agricultural commodities with an annual production cycle. But in general, this finding implies that the information content of SST anomalies is greater when an ENSO event starts to form, rather than when the event has already realized. Second, seasonal differences are also apparent for commodities that are characterized by linear ARDL models, suggesting that a key driver of seasonal discrepancies is the nonlinear ENSO cycle. For example, a one standard deviation shock that occurs during the MJJ trimester results in more amplified SST dynamics during the following six months, than a similar shock occurring during the NDJ trimester, which is soon followed by the "spring barrier."

they illustrate history- and shock-specific discrepancies in the nonlinear model. That is, in the absence of nonlinearities (i.e., if both the ENSO equation as well as the price equations were given by linear models), for any given shock the GIRs across different histories would follow exactly the same path, while for an array of symmetrically distributed shocks, the GIRs would also be distributed symmetrically about a single mode.

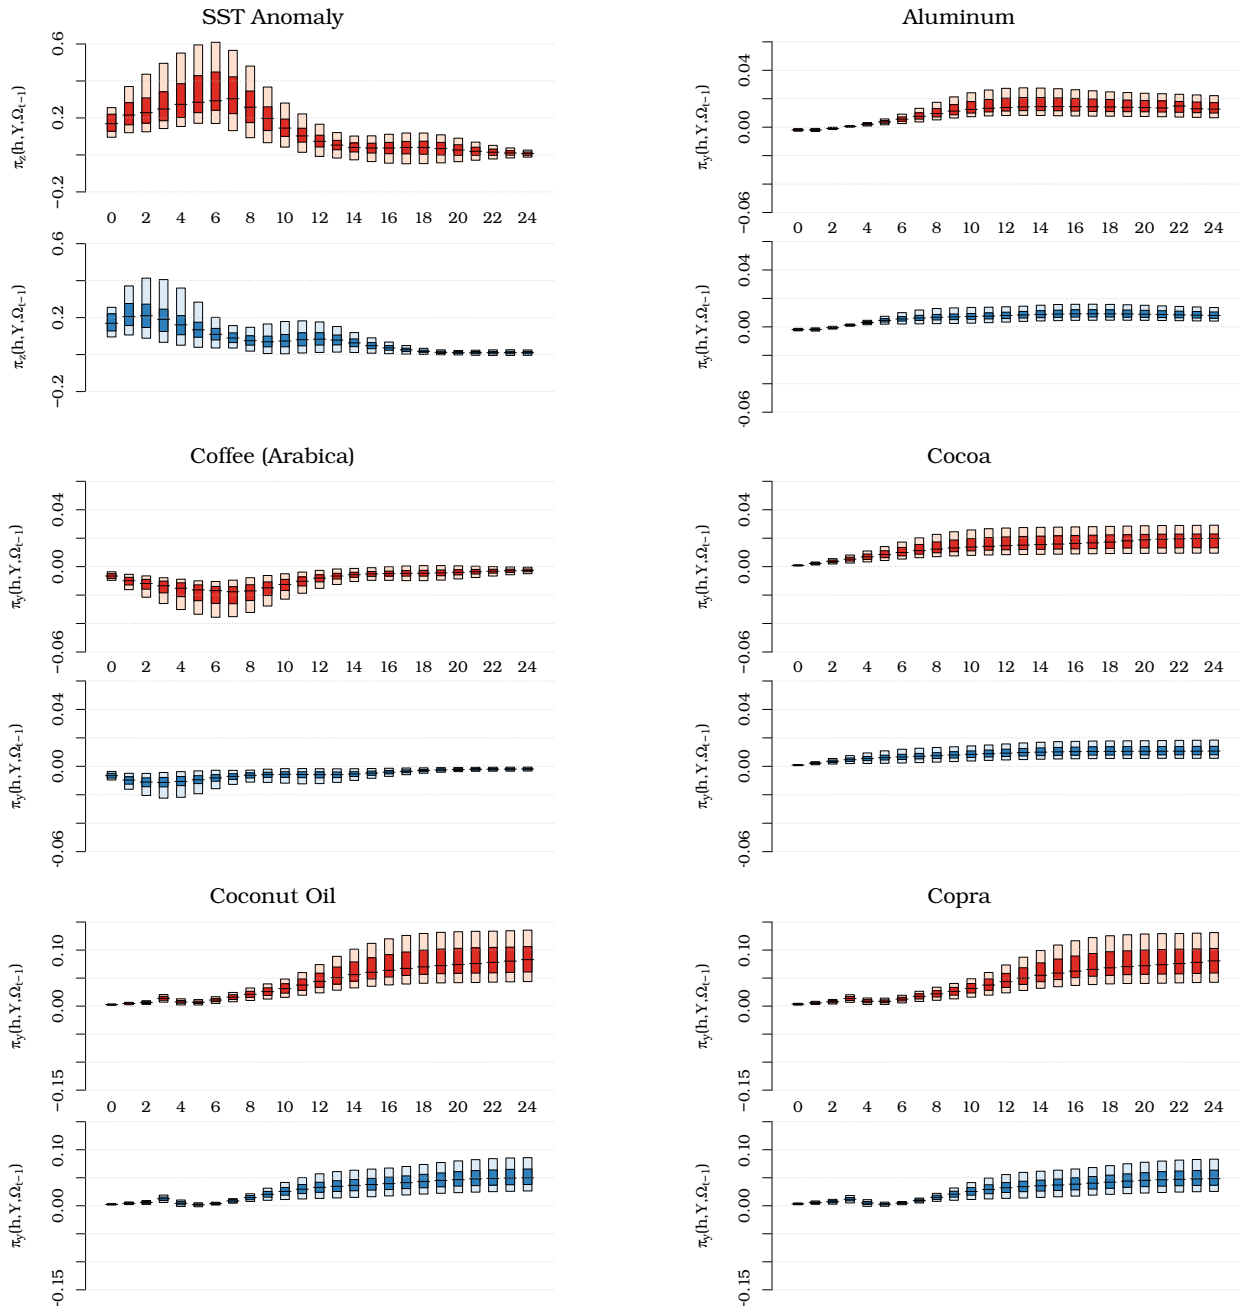

Figure 5: Seasonal asymmetries in SST and selected commodity prices

Note: These plots feature 50% (darker shade) and 90% (lighter shade) highest density regions (HDRs) of the generalized impulse–response functions (GIRs) at horizon h, $\forall\ h = 0, 1, \ldots, 24$. In each plot, the GIRs are associated with an average one standard deviation SST shock during the MJJ trimester (top panel, shades of red) and the NDJ trimester (bottom panel, shades of blue).

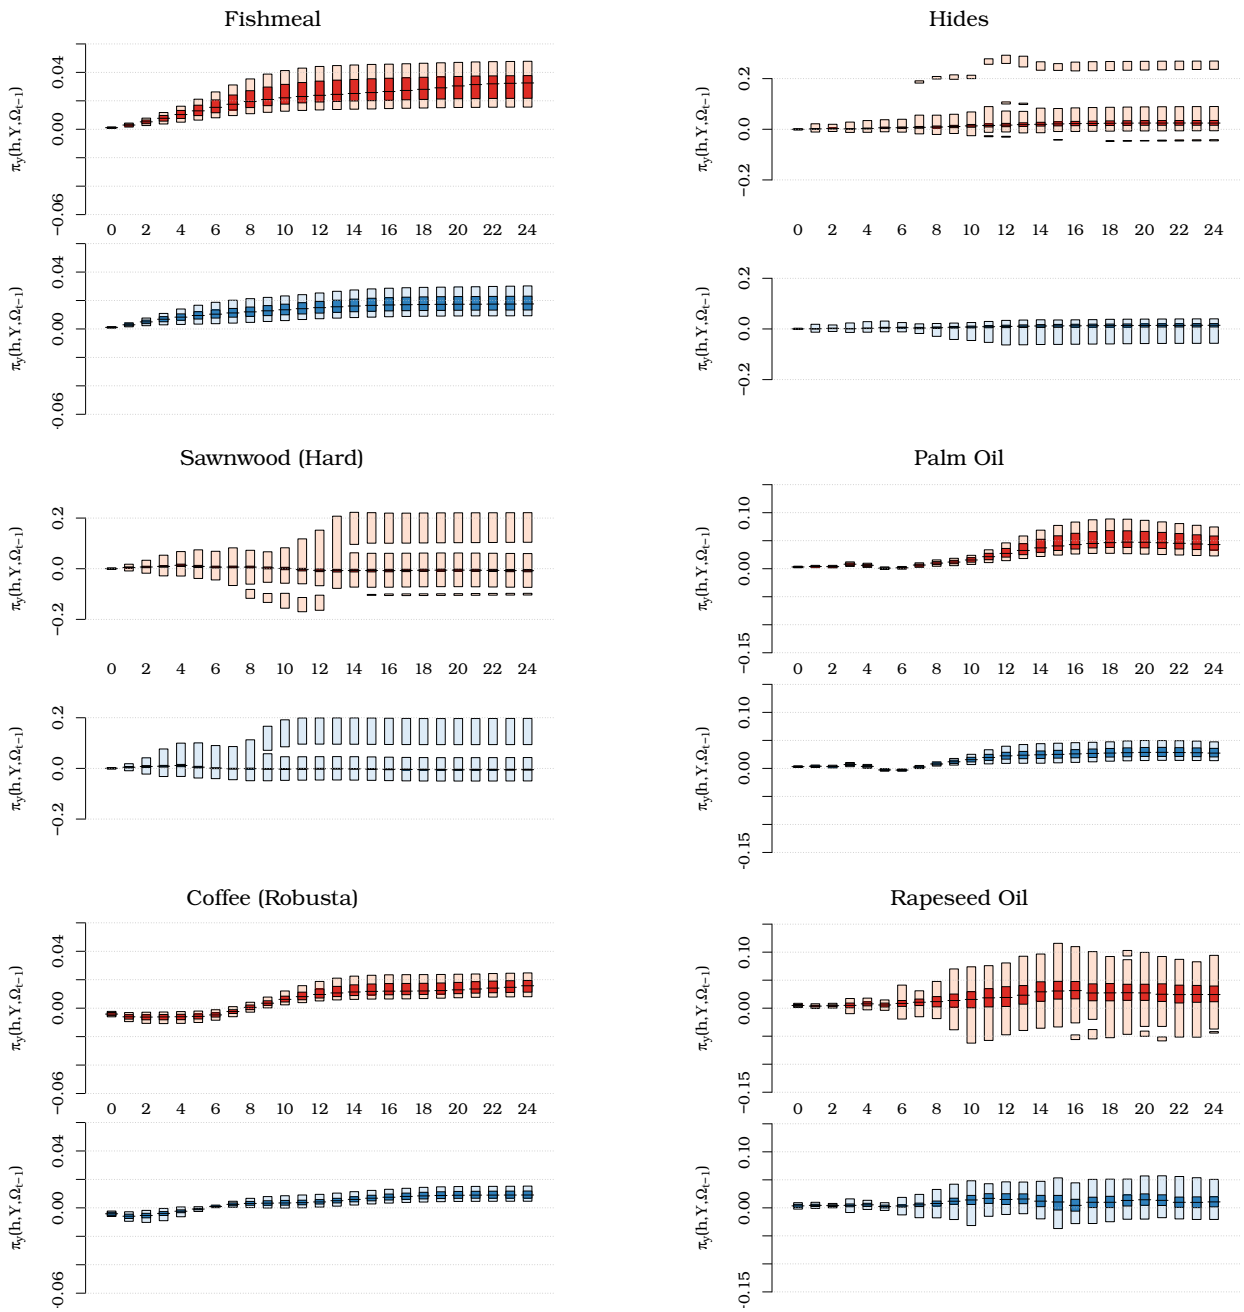

Figure 5: continued from the previous page

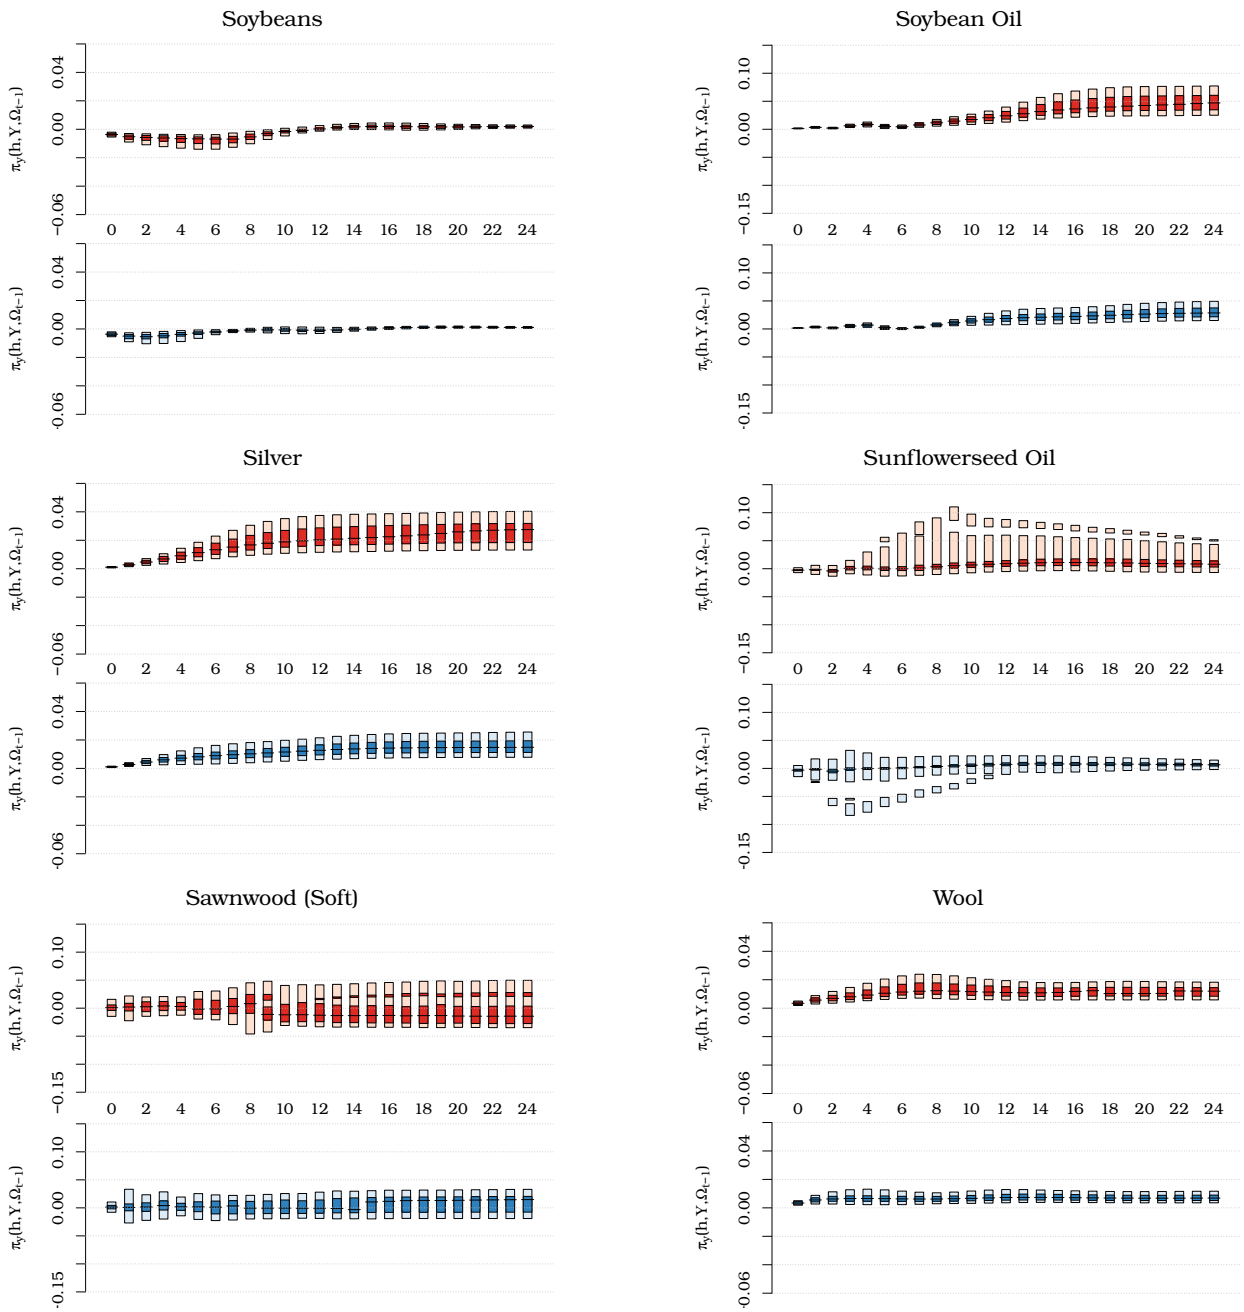

Figure 5: continued from the previous page

To better illustrate the changes in the dynamics of these commodity prices characterized by time-varying parameters—i.e., salmon and chicken—I generate and compare GIRs based on histories sampled before and after the estimated structural change. These GIRs are illustrated in figure 6. Since the early-to-mid-2000s, salmon prices have become more responsive to ENSO shocks. Notably, this period coincides with stronger correlations between the cost

of input to the price of output in the salmon industry (Asche and Oglend, 2016). So, to the extent that fishmeal is a key factor in salmon production, the relationship between ENSO and fishmeal prices appears to be playing an important role in salmon price behavior. In the case of chicken, the structural change has reversed the price impact of ENSO shocks. It is difficult to pinpoint the reason(s) for this effect. Nonetheless, the role of ENSO appears to be negligible in relation to out-of-sample predictability of chicken prices (see the section for more details). Hence, chicken price dynamics belong to the same discussion as those of other meat products. That is, the applied modeling framework reveals the gradual structural change in meat industries, which began in the early 1990s, but there is little evidence to suggest that ENSO has had any influence on the movements of these prices. This is not surprising, given the negligible impact of SST shocks on cereal grain prices.

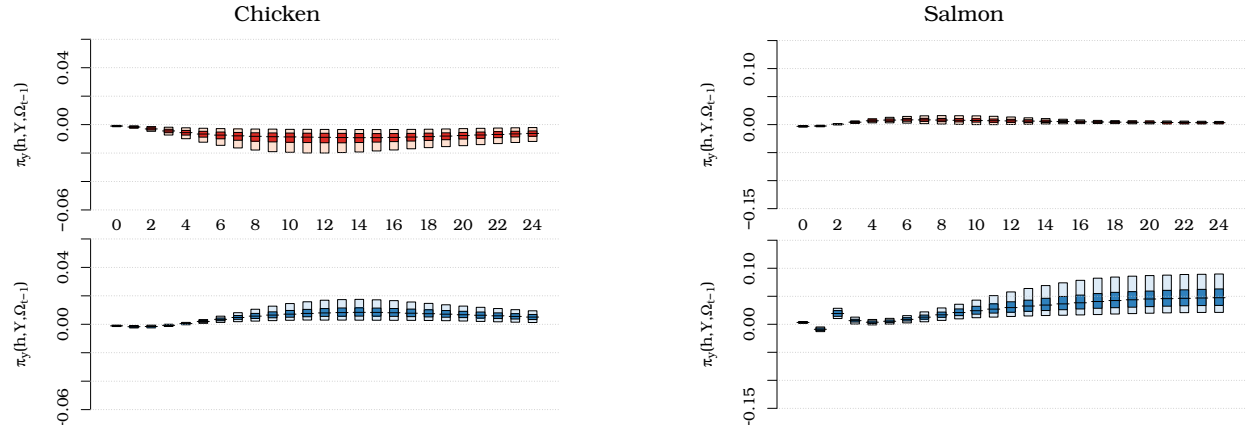

Figure 6: Time-varying dynamics of selected commodity prices

Note: These plots feature 50% (darker shade) and 90% (lighter shade) highest density regions (HDRs) of the generalized impulse–response functions (GIRs) at horizon, h, for all $h = 0, 1, \ldots, 24$. In each plot, the GIRs are associated with an average one standard deviation SST shock during the 1981–1985 period (top panel, shades of red) and the 2010–2014 period (bottom panel, shades of blue).

Forecasting Commodity Prices Using SST Anomalies

The foregoing exercise relies on in-sample goodness-of-fit measures, e.g., the Akaike Information Criterion, to identify the most suitable model. However, a better in-sample fit does not necessarily guarantee improved out-of-sample predictability. More to the point, even if the SST anomalies appear to be Granger causing price movements based on in-sample statistics, their role in the out-of-sample predictability of prices may be limited or statisti-

cally insignificant. In turn, out-of-sample predictability, or forecasting, is at the core of time series analysis. In this section, I undertake a pseudo-forecasting exercise to obtain insights into commodity price predictability in relation to SST anomalies. In particular, I apply the so-called rolling window approach to generate a sequence of pseudo-forecasts from linear or nonlinear ARDL models, i.e., those with the SST anomaly in the information set, and their parsimonious AR variants, i.e., the information set that excludes the SST anomaly.

To begin, consider an h-step-ahead point forecast, given by:

$$\hat{y}_{t+h|t} = f\left(\mathcal{F}_t; \hat{\boldsymbol{\theta}}_t\right), \tag{14}$$

where $f(\cdot)$ is the functional form of the estimated model, $\hat{\boldsymbol{\theta}}_t$ is a set of parameter estimates, and $\mathcal{F}_t$ is the information set available at the time of forecast. Thus, for unrestricted models, $\mathcal{F}_t = (y_{t-1}, \ldots, y_{t-p}, z_t, \ldots, z_{t-q})'$, while for restricted models, $\mathcal{F}_t = (y_{t-1}, \ldots, y_{t-p})'$. The goal is to compare the two forecasts, one of which is based on the information set with the SST anomaly, while the other excludes the SST anomaly.

As per convention, I assess forecast accuracy assuming quadratic loss. To that end, the out-of-sample measure of the "departure from the perfect fit," i.e., the root mean square forecast error (RMSFE), is comparable to the usual in-sample measure, the residual standard deviation. Let $\hat{e}_{t+h} = y_{t+h} - \hat{y}_{t+h|t}$ be the forecast error at horizon h, where y_{t+h} is the actual realization of the variable of interest in the forecast period. Then the associated RMSFE is given by:

$$\text{RMSFE} = \sqrt{\frac{1}{P}\sum_{t=R}^{T-h} \hat{e}_{t+h}^2}, \tag{15}$$

where P is the total number of out-of-sample forecasts, and R denotes the estimation window size, i.e., the subset of observations used to obtain a set of parameter estimates, such that $R + P = T - h$, where T denotes the total sample size. If a loss differential between the two models, of which the former nests the latter, is statistically significantly different from zero, the SST anomaly will be said to cause commodity prices in the Granger sense.[7]

Here, I set $R \approx 0.75T$. Thus, the first estimation window starts in January 1981 and ends in December 2007, and the associated forecast window spans the period from Jan-

[7]The most popular metric for the assessment of forecast accuracy is the Diebold and Mariano (1995) (DM) statistic, which was designed to test the predictive accuracy of two non-nested models. However, when one of the competing models nests another, the standard asymptotic critical values are no longer applicable (Clark and McCracken, 2001). This issue can be circumvented by incorporating an adjustment factor into the forecast accuracy statistic as per Clark and West (2007). See Clark and McCracken (2001); Clark and West (2007), and references therein, for further details.

uary 2008 to December 2009. The successive estimation and forecast windows are rolled over by one month, resulting in a total of 108 h-step-ahead forecasts. For each commodity price under consideration—i.e., those that are caused by SST anomalies based on in-sample examination—the specification identified in the modelling stage of this exercise is fitted in each rolling window. That is, the model type, as well as the autoregressive and distributed lag orders, are maintained throughout the rolling windows, but new set of parameters are estimated in each rolling window. To the extent that the forecasting exercise makes use of some out-of-sample information, a more accurate description of this procedure would be *quasi-out-of-sample forecasting* (see, e.g., Rothman et al., 2001).

Note that for one-step-ahead forecasts, or for any h-step-ahead forecast where the underlying model specification is linear, a naïve, or so-called *skeleton extrapolation* will suffice. But in the case of nonlinear models, the skeleton extrapolation will yield biased forecasts for horizons greater than one, for reasons similar to those discussed above in the context of impulse–response functions. To circumvent this issue, I simulate 1000 projections using bootstrapped residuals of the estimated model within a given rolling window,[8] and then take the average of these projections to approximate an h-step-ahead forecast for all $h > 1$. Using these forecasts and actual realizations of commodity prices, I obtain the relative forecast accuracy measures, and their statistical significance, which I present in table 3.

These results confirm some of the previously reported findings by emphasizing the relevance of SST anomalies in predicting the prices of some oilseeds and vegetable oils. The economically meaningful and statistically significant effect is particularly apparent in the intermediate term. Of the remaining agricultural commodities, SST anomalies help to improve the predictability of *Robusta* coffee prices, but this effect is not statistically significant. For the other agricultural commodities, the informational content of SST anomalies is not sufficient to offer improved out-of-sample predictability, despite the favorable in-sample fit. Notably, beyond the agricultural sector, SST anomalies maintain relevance in the intermediate-term predictability of aluminum prices, but are unable to aid in predicting the prices of other commodities, such as timber. Another important observation here is that SST anomalies cannot improve the predictability of those commodity prices that are characterized by nonlinear dynamics with regard to the ENSO regimes. This is not surprising, as out-of-sample Granger causality between SST anomalies and these commodity prices is likely of episodic nature. For example, SST anomalies may help to forecast prices only in extreme

[8]The approach is similar to that used to generate expected forecast path (without an initial shock) in generalized impulse–response functions. I shall thus refrain from further details here. But a reader may refer to Teräsvirta et al. (2010) for an in-depth discussion on this matter.

Table 3: Out-of-Sample Relative Predictive Accuracy of Commodity Prices

Commodity	$h = 1$	$h = 2$	$h = 3$	$h = 6$	$h = 12$	$h = 18$	$h = 24$
Aluminum	1.01	1.02	1.02	1.01	0.96^{*}	0.92^{**}	0.98
Chicken	1.01	1.01	1.01	1.00	1.01	0.99	0.99
Cocoa	1.01	1.01	1.02	1.02	1.01	1.01	1.02
Coconut Oil	0.97^{**}	0.95^{**}	0.92^{**}	0.88^{**}	0.83^{***}	0.83^{***}	0.88^{**}
Coffee (Arabica)	0.99	1.00	1.00	0.99	0.99	0.99	1.00
Coffee (Robusta)	1.01	1.03	1.04	1.08	1.01	0.95	0.96
Copra	0.97^{**}	0.95^{**}	0.93^{**}	0.90^{**}	0.84^{***}	0.84^{***}	0.89^{**}
Fishmeal	0.99^{*}	0.99^{*}	0.98^{*}	0.99	1.01	1.03	1.09
Hides	0.98	1.02	1.03	1.07	1.11	1.06	1.01
Palm Oil	1.00	0.99	0.98	0.96	0.99	1.04	1.10
Rapeseed Oil	1.10	1.10	1.13	1.17	1.08	1.15	1.16
Salmon	1.00	1.01	1.02	1.03	0.99	0.99	1.00
Sawnwood (Hard)	1.11	1.20	1.22	1.17	1.17	1.16	1.13
Sawnwood (Soft)	1.02	1.05	1.04	1.06	1.05	1.10	1.10
Silver	1.01	1.01	1.02	1.02	1.01	1.00	1.01
Soybean Oil	1.00	1.01	1.00	0.98	0.92^{**}	0.91^{**}	0.95^{*}
Soybeans	1.01	1.01	1.02	1.02	1.00	1.00	1.01
Sunflowerseed Oil	1.01	1.02	1.06	1.13	1.11	1.00	0.90
Wool	1.00	1.00	1.00	1.01	1.01	1.00	1.00

Note: the table entries are the ratios of the RMSFE from two competing models, where the restricted model serves as denominator (so a value of less than one suggests a better out-of-sample fit by a model with the SST anomaly). ***, **, and *denote statistical significance at $\alpha = 0.01$, 0.05, and 0.10 levels, based on Clark and West (2007) critical values.

ENSO events. But the phase-specific relative forecast accuracy measures are not presented here, given the relatively small number of periods that were classified as El Niño or La Niña regimes during the 2008–2016 forecasting period.

Conclusion

The literature on commodity price behavior has provided consensus on some of their well-established characteristics. Commodity price series are found to be highly persistent (Cashin et al., 2000; Ghoshray, 2013), with occasional spikes, and possibly nonlinear dynamics (Tomek, 2000; Cashin et al., 2002; Enders and Holt, 2012). Even so, the question of *why prices move as they do* continues to challenge economists. In this study, I address the aforementioned question by examining the extent to which an exogenous climatic factor, ENSO, may Granger cause primary commodity price movement. I do this by applying a nonlinear modeling framework to better approximate complex linkages between SST anomalies and an extensive list of primary commodity prices, and then testing the predictability of prices in relation to this climate phenomenon in a pseudo-forecasting environment.

I find that SST anomalies facilitate price movements in selected primary commodities, particularly those in the agricultural sector. But the ability of this climate phenomenon to "genuinely" forecast prices is limited to just a few commodities. Vegetable oils and protein meals—particularly those produced in the western region of the Pacific—represent a key group of commodities that respond most robustly to ENSO shocks. For a range of other commodities, such as beverages and timber, SST anomalies contribute to improved in-sample fit of the data, but this fit does not translate into more accurate forecasts. For timber in particular, the observed relationship between SST anomalies and prices is episodic, i.e., the effect is apparent when nonlinearities in the ENSO–price linkage are accounted for. Thus, lack of overall predictability can be attributed to possibly episodic out-of-sample causality. Contrary to expectations, I found no evidence of any price effect of SST anomalies for cereal grains. This may be the result of the limited exposure of temperate regions to ENSO shocks, as well as the north–south diversification of these crops, and the resultant buffer provided by intra-annual supply responses and international trade.

The findings of this study add considerably to existing knowledge of the ENSO–price nexus, and carry important socio-economic implications, particularly for developing nations. These counties are major producers of many of the primary commodities considered in this study, and their economic growth has historically been affected by price movements in these commodities (Deaton, 1999; Chinn and Coibion, 2014). Therefore, the ENSO–price linkage can be seen as an important channel through which climate shocks can affect the well-being of citizens in these lower income countries. In conclusion, disentangling the supply- and demand-side channels through which ENSO may impact price movements is a worthwhile subject for future research, and one for which this study offers important insights.

References

Asche, F. and A. Oglend (2016). The Relationship Between Input-factor and Output Prices in Commodity Industries: The Case of Norwegian Salmon Aquaculture. *Journal of Commodity Markets 1*(1), 35 – 47.

Bacon, D. and D. Watts (1971). Estimating the Transition between Two Intersecting Straight Lines. *Biometrika 58*(3), 525–534.

Balagtas, J. V. and M. T. Holt (2009). The Commodity Terms of Trade, Unit Roots, and Nonlinear Alternatives: A Smooth Transition Approach. *American Journal of Agricultural Economics 91*(1), 87–105.

Brunner, A. (2002). El Nino and World Primary Commodity Prices: Warm Water or Hot Air? *Review of Economics and Statistics 84*(1), 176–183.

Cashin, P., H. Liang, and C. J. McDermott (2000). How Persistent Are Shocks to World Commodity Prices? *IMF Staff Papers 42*(2), 177–217.

Cashin, P., C. J. McDermott, and A. Scott (2002). Booms and Slumps in World Commodity Prices. *Journal of Development Economics 69*(1), 277–296.

Cashin, P. A., K. Mohaddes, and M. Raissi (2017). Fair Weather or Foul? The Macroeconomic Effects of El Niño. *Journal of International Economics 106*, 37–54.

Chan, K. and H. Tong (1986). On Estimating Thresholds in Autoregressive Models. *Journal of Time Series Analysis 7*(3), 179–190.

Chinn, M. D. and O. Coibion (2014). The Predictive Content of Commodity Futures. *Journal of Futures Markets 34*(7), 607–636.

Clark, T. and M. McCracken (2001). Tests of Equal Forecast Accuracy and Encompassing for Nested Models. *Journal of Econometrics 105*(1), 85–110.

Clark, T. and K. West (2007). Approximately Normal Tests for Equal Predictive Accuracy in Nested Models. *Journal of Econometrics 138*(1), 291–311.

Craig, L. A. and M. T. Holt (2008). Mechanical Refrigeration, Seasonality, and the Hog-Corn Cycle in the United States: 1870-1940. *Explorations in Economic History 45*(1), 30–50.

Davies, R. (1977). Hypothesis Testing when a Nuisance Parameter is Present only under the Alternative. *Biometrika 64*(2), 247–254.

Davies, R. (1987). Hypothesis Testing when a Nuisance Parameter is Present only under the Alternative. *Biometrika 74*(1), 33–43.

Davis, M. (2002). *Late Victorian Holocausts: El Niño Famines and the Making of the Third World*. Verso.

Deaton, A. (1999). Commodity Prices and Growth in Africa. *The Journal of Economic Perspectives 13*(3), 23–40.

Demetrescu, M. and R. Kruse (2013). The Power of Unit Root Tests Against Nonlinear Local Alternatives. *Journal of Time Series Analysis 34*(1), 40–61.

Diebold, F. and R. Mariano (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics 13*(3), 253–263.

Edelstein, P. and L. Kilian (2009). How Sensitive are Consumer Expenditures to Retail Energy Prices? *Journal of Monetary Economics 56*(6), 766–779.

Eitrheim, O. and T. Teräsvirta (1996). Testing the Adequacy of Smooth Transition Autoregressive Models. *Journal of Econometrics 74*(1), 59–75.

Enders, W. and M. T. Holt (2012). Sharp Breaks or Smooth Shifts? An Investigation of the Evolution of Primary Commodity Prices. *American Journal of Agricultural Economics 94*(3), 659–673.

Gan, J. (2006). Causasilty among Wildfire, ENSO, Timber Harvest, and Urban Sprawl: The Vector Autoregression Approach. *Ecological modelling 191*(2), 304–314.

Ghoshray, A. (2013). Dynamic Persistence of Primary Commodity Prices. *American Journal of Agricultural Economics 95*(1), 153–164.

Glauber, J. W. and M. J. Miranda (2016). The effects of southern hemisphere crop production on trade, stocks, and price integration. In M. Kalkuhl, J. von Braun, and M. Torero (Eds.), *Food Price Volatility and Its Implications for Food Security and Policy*, pp. 83–100. Switzerland: Springer International Publishing.

Hall, A., J. Skalin, and T. Teräsvirta (2001). A Nonlinear Time Series Model of El Niño. *Environmental Modelling & Software 16*(2), 139–146.

Hoerling, M. P., A. Kumar, and M. Zhong (1997). El Niño, La Niña, and the Nonlinearity of Their Teleconnections. *Journal of Climate 10*(8), 1769–1786.

Holt, M. T. and L. A. Craig (2006). Nonlinear Dynamics and Structural Change in the U.S. Hog-Corn Cycle: A Time-Varying STAR Approach. *American Journal of Agricultural Economics 88*(1), 215–233.

Hood, H. B. and J. H. Dorfman (2015). Examining Dynamically Changing Timber Market Linkages. *American Journal of Agricultural Economics 97*(5), 1451–1463.

Hsiang, S., K. Meng, and M. Cane (2011). Civil Conflicts are Associated with the Global Climate. *Nature 476*(7361), 438–441.

Hsiang, S. M. and K. C. Meng (2015). Tropical Economics. *American Economic Review: Papers & Proceedings 105*(5), 257–261.

Hyndman, R. J. (1995). Highest–Density Forecast Regions for Nonlinear and Non-normal Time Series Models. *Journal of Forecasting 14*(5), 431–441.

Hyndman, R. J. (1996). Computing and Graphing Highest Density Regions. *American Statistician 50*(2), 120–126.

Iizumi, T., J.-J. Luo, A. J. Challinor, G. Sakurai, M. Yokozawa, H. Sakuma, M. E. Brown, and T. Yamagata (2014). Impacts of El Niño Southern Oscillation on the Global Yields of Major Crops. *Nature Communications 5*(3712).

Koop, G., M. Pesaran, and S. Potter (1996). Impulse Response Analysis in Nonlinear Multivariate Models. *Journal of Econometrics 74*(1), 119–147.

Laosuthi, T. and D. D. Selover (2007). Does El Niño Affect Business Cycles? *Eastern Economic Journal 33*(1), 21–42.

Ludescher, J., A. Gozolchiani, M. I. Bogachev, A. Bunde, S. Havlin, and H. J. Schellnhuber (2014). Very Early Warning of Next El Niño. *Proceedings of the National Academy of Sciences 111*(6), 2064–2066.

Lundbergh, S., T. Teräsvirta, and D. Van Dijk (2003). Time–Varying Smooth Transition Autoregressive Models. *Journal of Business & Economic Statistics 21*(1), 104–122.

Luukkonen, R., P. Saikkonen, and T. Teräsvirta (1988). Testing Linearity Against Smooth Transition Autoregressive Models. *Biometrika 75*(3), 491–499.

Lybbert, T. J., A. Smith, and D. A. Sumner (2014). Weather Shocks and Inter-hemispheric Supply Responses: Implications for Climate Change Effects on Global Food Markets. *Climate Change Economics 5*(4), 1450010.

Mason, S. J. and L. Goddard (2001). Probabilistic Precipitation Anomalies Associated With ENSO. *Bulletin of the American Meteorological Society 82*(4), 619–638.

Nepstad, D. C., A. Verssimo, A. Alencar, C. Nobre, E. Lima, P. Lefebvre, P. Schlesinger, C. Potter, P. Moutinho, E. Mendoza, M. Cochrane, and V. Brooks (1999). Large-scale Impoverishment of Amazonian Forests by Logging and Fire. *Nature 398*(6727), 505–508.

Pielke Jr, R. and C. Landsea (1999). La Niña, El Niño, and Atlantic Hurricane Damages in the United States. *Bulletin of the American Meteorological Society 80*(10), 2027–2034.

Rasmusson, E. (1991). *Teleconnections Linking Worldwide Climate Anomalies*, Chapter Observational Aspects of ENSO Cycle Teleconnections, pp. 309–343. Cambridge University Press, New York.

Ropelewski, C. and M. Halpert (1989). Precipitation Patterns Associated With the High Index Phase of the Southern Oscillation. *Journal of Climate 2*(3), 268–284.

Rothman, P., D. van Dijk, and P. H. Franses (2001). Multivariate STAR Analysis of Money–Output Relationship. *Macroeconomic Dynamics 5*(4), 506–532.

Smith, S. C. and D. Ubilava (2017). The El Niño Southern Oscillation and Economic Growth in the Developing World. *Global Environmental Change 45*, 151–164.

Stone, R., G. Hammer, and T. Marcussen (1996). Prediction of Global Rainfall Probabilities Using Phases of the Southern Oscillation Index. *Nature 384*(6606), 252–255.

Teräsvirta, T. (1994). Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models. *Journal of the American Statistical Association 89*(425), 208–218.

Teräsvirta, T. and H. Anderson (1992). Characterizing Nonlinearities in Business Cycles using Smooth Transition Autoregressive Models. *Journal of Applied Econometrics 7*(S1), S119–S136.

Teräsvirta, T., D. Tjøstheim, and C. W. J. Granger (2010). *Modelling Nonlinear Economic Time Series*. Advanced Texts in Econometrics. Oxford University Press.

Tomek, W. (2000). Commodity Prices Revisited. *Agricultural and Resource Economics Review 29*(2), 125–137.

Tong, H. and K. S. Lim (1980). Threshold Autoregression, Limit Cycles and Cyclical Data. *Journal of the Royal Statistical Society. Series B (Methodological) 42*(3), 245–292.

Ubilava, D. (2012). El Niño, La Niña, and World Coffee Price Dynamics. *Agricultural Economics 43*(1), 17–26.

Ubilava, D. and C. Helmers (2013). Forecasting ENSO with a Smooth Transition Autoregressive Model. *Environmental Modelling & Software 40*(1), 181–190.

Ubilava, D. and M. Holt (2013). El Niño Southern Oscillation and its Effects on World Vegetable Oil Prices: Assessing Asymmetries using Smooth Transition Models. *Australian Journal of Agricultural and Resource Economics 57*(2), 273–297.

van Dijk, D. and P. Franses (1999). Modeling Multiple Regimes in the Business Cycle. *Macroeconomic Dynamics 3*(3), 311–340.

Wang, D. and W. G. Tomek (2007). Commodity Prices and Unit Root Tests. *American Journal of Agricultural Economics 89*(4), 873–889.

World Bank Group (2015). Commodity Markets Outlook, October 2015. Technical report, World Bank, Washington, DC.

Zhang, T., J. Perlwitz, and M. P. Hoerling (2014). What is Responsible for the Strong Observed Asymmetry in Teleconnections between El Niño and La Niña? *Geophysical Research Letters 41*(3), 1019–1025.

Appendix

A1 Generalized Impulse–Responses: A Bootstrap Resampling Algorithm

1. Estimate the parameters of a suitable model; denote the estimated parameters by $\hat{\boldsymbol{\theta}}$, and the residuals by $\hat{\varepsilon}_t$.

2. Identify a subset of histories, Ω_{t-1}, that satisfy a condition of interest (e.g., coinciding with a strong El Niño occurrence).

3. Identify a subset of initial shocks, Υ, that satisfy a condition of interest (e.g., exceeding one standard deviation of the estimated residuals).

4. Sample (randomly with replacement) a history, ω_{t-1}, and an initial shock, υ, where $\omega_{t-1} \in \Omega_{t-1}$ and $\upsilon \in \Upsilon$, obtain a GIR as follows:

   (a) Generate a set of idiosyncratic disturbances, $\{\varepsilon^i_{t+h} : h = 0, 1, \ldots\}$, by randomly sampling (with replacement) from $\{\hat{\varepsilon}_t : t = 1, \ldots, T\}$.

   (b) Construct a forecast path (without the initial shock) iteratively:

   $$x^i_{t+h} = f(x^i_{t+h-1}, x^i_{t+h-2}, \ldots; \hat{\boldsymbol{\theta}}) + \varepsilon^i_{t+h}.$$

   (c) Construct another forecast path (with the initial shock) iteratively:

   $$x^i_{t+h|\upsilon} = f(x^i_{t+h-1|\upsilon}, x^i_{t+h-2|\upsilon}, \ldots; \hat{\boldsymbol{\theta}}) + \varepsilon^i_{t+h}, \quad \text{where } \varepsilon^i_t = \upsilon.$$

   (d) Repeat steps 4a – 4c B times, where B is an integer denoting the total number of bootstrap iterations, and obtain the averages at each horizon for each of the two forecast paths:

   $$\mathbb{E}(x_{t+h}|\omega_{t-1}) = B^{-1} \sum_{i=1}^{B} x^i_{t+h},$$

   $$\mathbb{E}(x_{t+h}|\upsilon, \omega_{t-1}) = B^{-1} \sum_{i=1}^{B} x^i_{t+h|\upsilon}.$$

   (e) Take the difference between the two averages to obtain an estimate of GIR:

   $$\hat{\pi}_x(h, \upsilon, \omega_{t-1}) = \mathbb{E}(x_{t+h}|\upsilon, \omega_{t-1}) - \mathbb{E}(x_{t+h}|\omega_{t-1}).$$

5. Repeat step 4 a sufficiently large number of times to generate empirical distributions of the GIR at each horizon.

Table A1: Unit Root Test Results

Commodity	ADF_y	$ADF_{\Delta y}$	ZA_y	$ZA_{\Delta y}$	$KPSS_y$	$KPSS_{\Delta y}$
Aluminum	-3.51	-10.23	-4.18	-10.59	2.83	0.08
Barley	-3.28	-13.23	-4.07	-13.37	2.91	0.03
Beef	-2.40	-11.63	-4.44	-11.96	3.82	0.20
Chicken	-2.57	-10.94	-5.38	-11.06	8.76	0.04
Cocoa	-2.82	-10.33	-4.41	-10.63	2.77	0.19
Coconut Oil	-2.82	-9.87	-3.82	-10.02	0.94	0.16
Coffee (Arabica)	-3.36	-10.43	-4.21	-10.70	3.03	0.07
Coffee (Robusta)	-2.47	-10.15	-3.51	-10.46	6.80	0.18
Copper	-1.89	-11.59	-4.20	-11.78	5.77	0.13
Copra	-2.73	-9.56	-3.65	-9.68	1.02	0.17
Cotton	-3.06	-9.49	-4.86	-9.76	7.71	0.05
Fishmeal	-2.35	-11.27	-3.95	-11.66	2.99	0.21
Gold	-1.12	-12.27	-3.97	-12.87	4.45	0.62
Groundnut Oil	-3.66	-9.69	-4.45	-9.88	0.81	0.04
Hard Logs	-3.24	-11.00	-5.12	-11.28	1.87	0.04
Hard Sawnwood	-1.97	-10.46	-5.37	-10.87	3.68	0.07
Hides	-2.41	-12.71	-4.49	-13.10	5.09	0.03
Lamb	-2.70	-10.07	-4.73	-10.44	4.32	0.04
Lead	-1.92	-12.03	-4.71	-12.57	5.60	0.26
Maize	-3.20	-10.67	-4.81	-10.83	1.74	0.04
Nickel	-2.68	-10.73	-3.50	-11.03	3.00	0.05
Olive Oil	-2.28	-11.44	-3.59	-11.77	1.66	0.06
Palm Oil	-3.64	-7.73	-4.57	-7.88	0.54	0.08
Platinum	-2.00	-11.34	-4.04	-11.72	6.56	0.21
Pork	-2.58	-15.07	-6.27	-15.29	15.38	0.02
Rapeseed Oil	-3.12	-12.50	-4.07	-12.65	1.35	0.06
Rice	-3.18	-10.83	-4.18	-11.13	1.64	0.08
Salmon	-2.46	-12.34	-4.24	-12.54	10.50	0.15
Silver	-2.17	-12.77	-3.98	-13.27	3.32	0.60
Soft Logs	-1.67	-14.12	-3.56	-14.51	2.91	0.07
Soft Sawnwood	-2.05	-12.45	-4.37	-12.86	2.55	0.05
Sorghum	-3.58	-11.38	-5.15	-11.53	1.70	0.04
Soybean Meal	-3.19	-11.00	-4.96	-11.12	1.83	0.07
Soybean Oil	-2.85	-9.09	-3.76	-9.25	0.80	0.09
Soybeans	-3.19	-12.26	-5.18	-12.48	2.04	0.06
Sugar	-4.25	-10.34	-4.93	-10.97	0.62	0.10
Sunflowerseed Oil	-3.69	-10.15	-5.11	-10.28	1.02	0.03
Tea	-4.23	-12.93	-5.69	-13.02	4.53	0.04
Tin	-2.35	-9.29	-4.41	-9.75	2.37	0.40
Tobacco	-2.23	-9.56	-3.36	-9.81	4.62	0.13
Wheat	-2.76	-11.29	-3.73	-11.59	1.81	0.07
Wool	-2.44	-10.96	-3.77	-11.13	2.14	0.19
Zinc	-2.88	-10.63	-3.87	-10.85	1.20	0.05

Note: Augmented Dickey–Fuller (ADF), Zivot–Andrews (ZA), and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests are applied to levels and first-differences of log-transformed real commodity prices. The 5% critical vales for ADF, ZA, and KPSS are: -2.87, -4.80, and 0.46, respectively.

Table A2: Selected models using alternative measures of ENSO intensity

	SST					SOI				
Commodity	Model	p	q	ℓ	AIC_c	Model	p	q	ℓ	AIC_c
Aluminum	ARDL	2	1		0.211	ARDL	2	0		0.207
Barley	AR	2			0.585	LTVARDL	2	8		0.558
Beef	AR	3			-0.539	LSTARDL	3	2	0	-0.540
Chicken	LTVARDL	2	2		-1.835	LTVARDL	2	2		-1.836
Cocoa	ARDL	3	0		0.355	ARDL	3	0		0.354
Coconut Oil	ARDL	2	5		0.810	ARDL	2	6		0.840
Coffee (Arabica)	ARDL	3	1		0.885	AR	3			0.890
Coffee (Robusta)	ARDL	2	1		0.623	ARDL	2	0		0.623
Copper	AR	2			0.335	AR	2			0.335
Copra	ARDL	2	5		0.743	ARDL	2	6		0.759
Cotton	AR	2			-0.021	AR	2			-0.021
Fishmeal	ARDL	2	0		-0.053	LTVARDL	2	8		-0.077
Gold	AR	2			-0.397	LTVARDL	2	0		-0.409
Groundnut Oil	AR	2			0.087	LTVARDL	2	6		-0.002
Hides	LSTARDL	2	8	6	0.430	AR	2			0.441
Lamb	LTVAR	3			-0.410	LTVARDL	3	1		-0.435
Lead	AR	2			0.706	AR	2			0.706
Logs (Hard)	LTVAR	2			0.264	LTVAR	2			0.264
Logs (Soft)	AR	2			0.591	AR	2			0.591
Maize	AR	2			0.235	AR	2			0.235
Nickel	AR	2			0.963	AR	2			0.963
Olive Oil	LTVAR	2			-0.074	LTVAR	2			-0.074
Palm Oil	ARDL	5	6		0.609	ARDL	5	6		0.628
Platinum	AR	2			0.181	LSTARDL	2	7	6	0.167
Pork	LTVAR	1			1.412	LTVAR	1			1.412
Rapeseed Oil	ESTARDL	2	6	6	0.911	LTVARDL	2	6		0.790
Rice	AR	3			0.247	ESTARDL	3	1	1	0.225
Salmon	LTVARDL	2	4		0.187	LTVAR	2			0.199
Sawnwood (Hard)	LSTARDL	3	7	5	-0.068	AR	3			-0.018
Sawnwood (Soft)	ESTARDL	3	0	0	0.598	LSTARDL	3	5	2	0.622
Silver	ARDL	2	0		0.651	LSTARDL	2	0	0	0.637
Sorghum	AR	2			0.392	AR	2			0.392
Soybean Meal	AR	2			0.180	AR	2			0.180
Soybean Oil	ARDL	3	6		0.192	ARDL	3	6		0.205
Soybeans	ARDL	2	1		0.000	AR	2			0.001
Sugar	AR	3			1.079	AR	3			1.079
Sunflowerseed Oil	LSTARDL	3	5	1	0.641	LSTARDL	3	1	1	0.561
Tea	LTVAR	2			0.781	LTVAR	2			0.781
Tin	AR	3			0.202	AR	3			0.202
Tobacco	AR	2			-1.878	ARDL	2	10		-1.896
Wheat	AR	2			0.261	AR	2			0.261
Wool	ARDL	2	1		-0.041	AR	2			-0.040
Zinc	AR	2			0.420	AR	2			0.420

Note: the estimated models are given by equations (9) and (10), or their restricted (i.e., without a measure of ENSO intensity in the equation) variants; p indicates the autoregressive lag length; q indicates the distributed lag length; ℓ denotes the delay factor of the transition variable in a nonlinear model. AIC_c is the small sample adjusted Akaike information criterion.

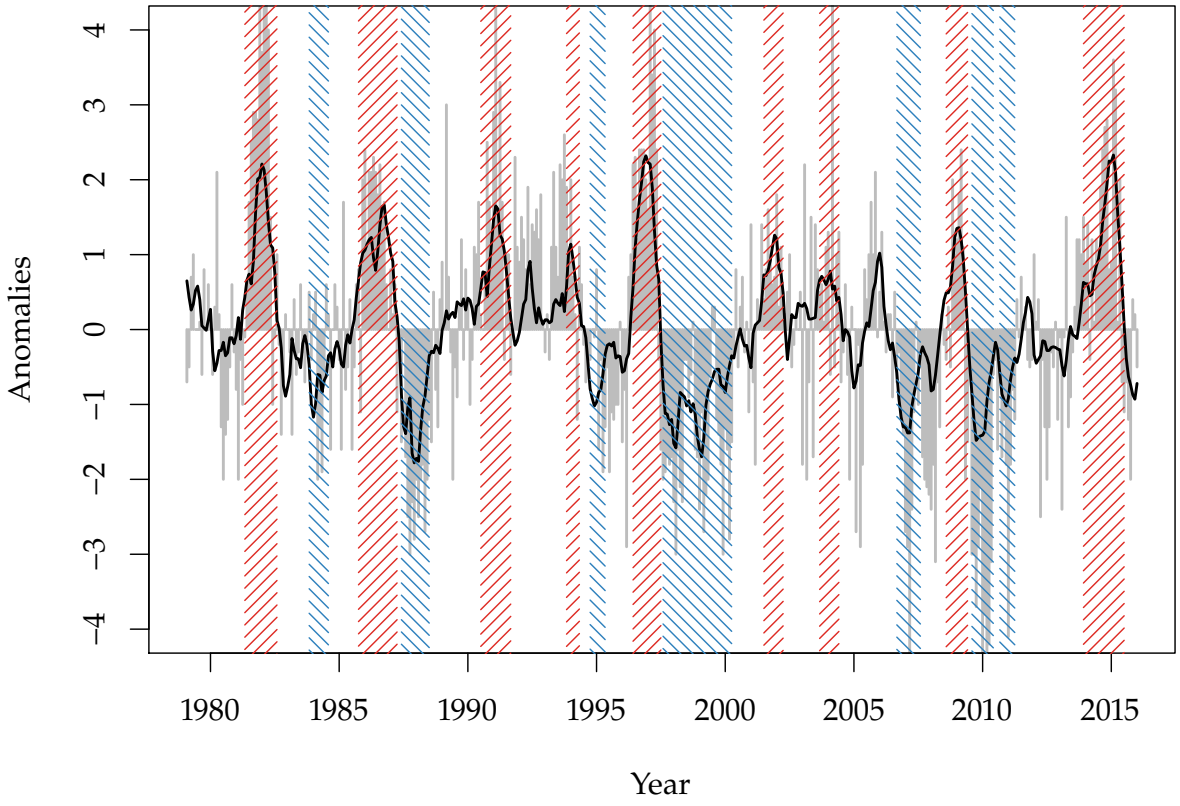

Figure A1: SST and SOI anomalies and ENSO events

Note: the black curve depicts the sea surface temperature (SST) deviations from centered 30-year base periods, in the `Nino3.4` region; the gray bars depict the Southern Oscillation index (SOI), calculated from standardized Tahiti and Darwin air-pressures, relative to the 1981–2010 base period; the SOI is multiplied by negative one to facilitate comparison with SST; regions shaded with (upward-sloping) red lines represent the El Niño events, and regions shaded with (downward-sloping) blue lines represent the La Niña events; the El Niño and La Niña events are identified based on the Oceanic Niño Index (ONI), which is a three-month running mean of the SST anomaly.

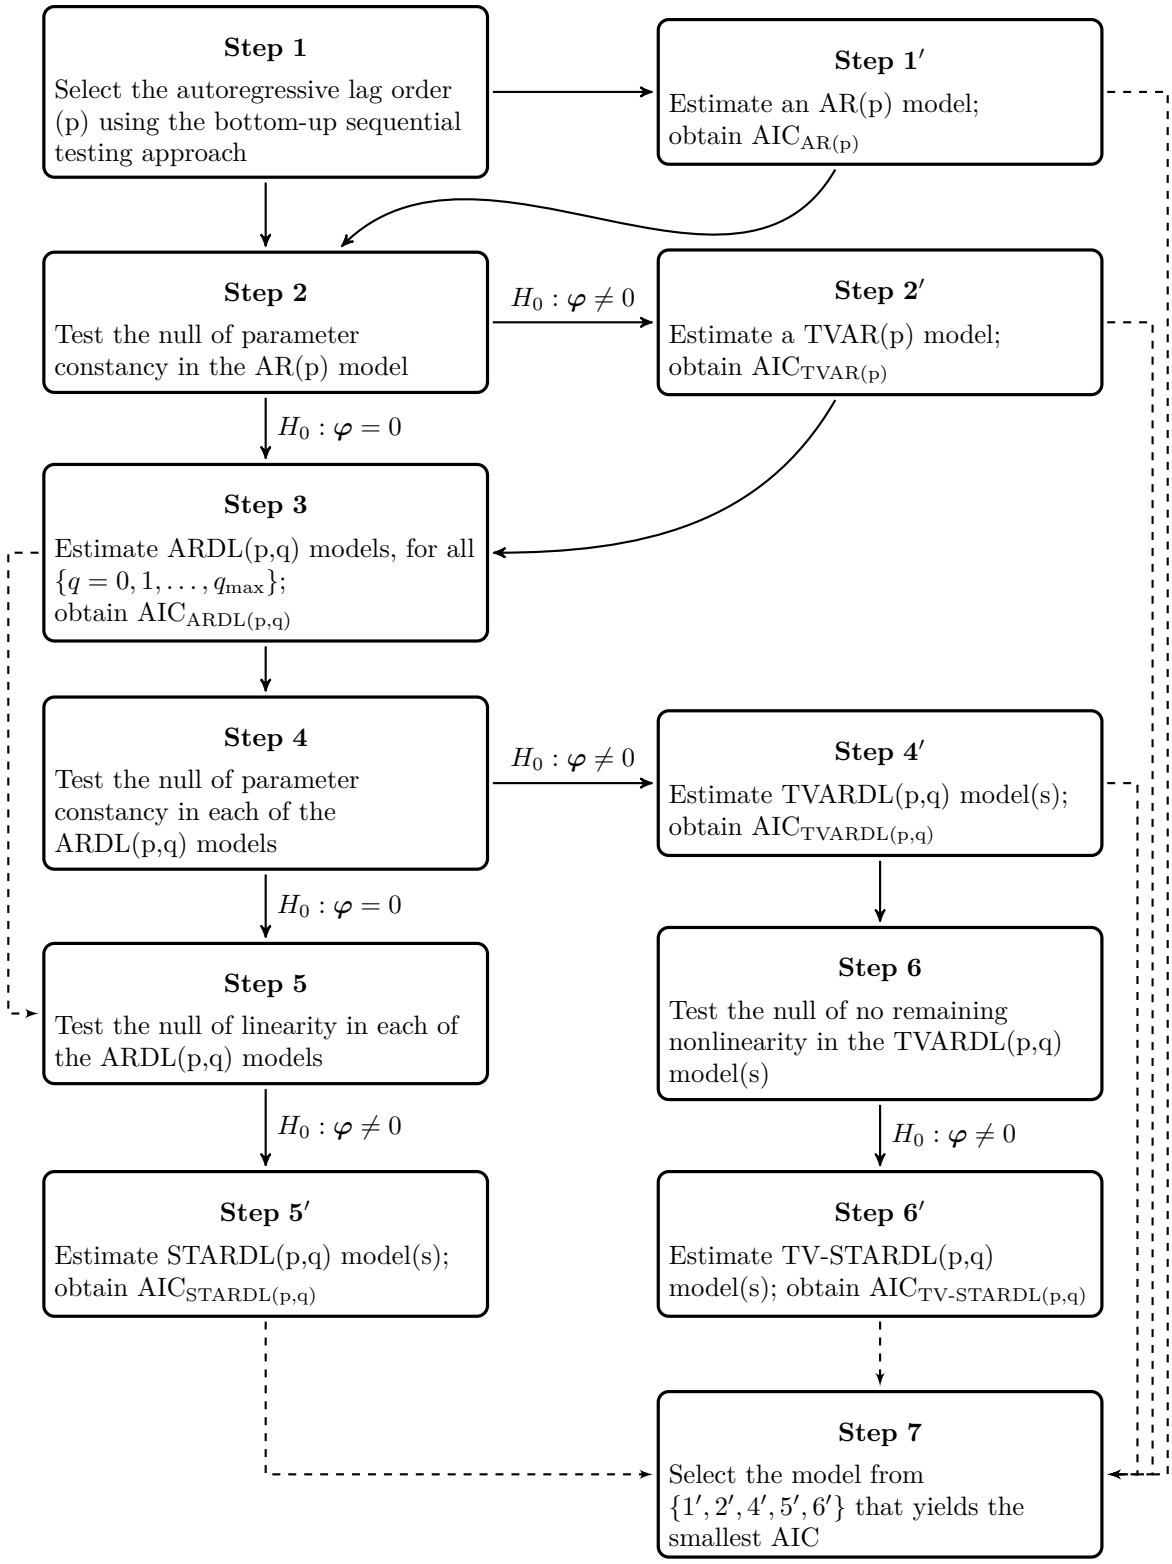

Figure A2: Model selection algorithm: a graphical illustration

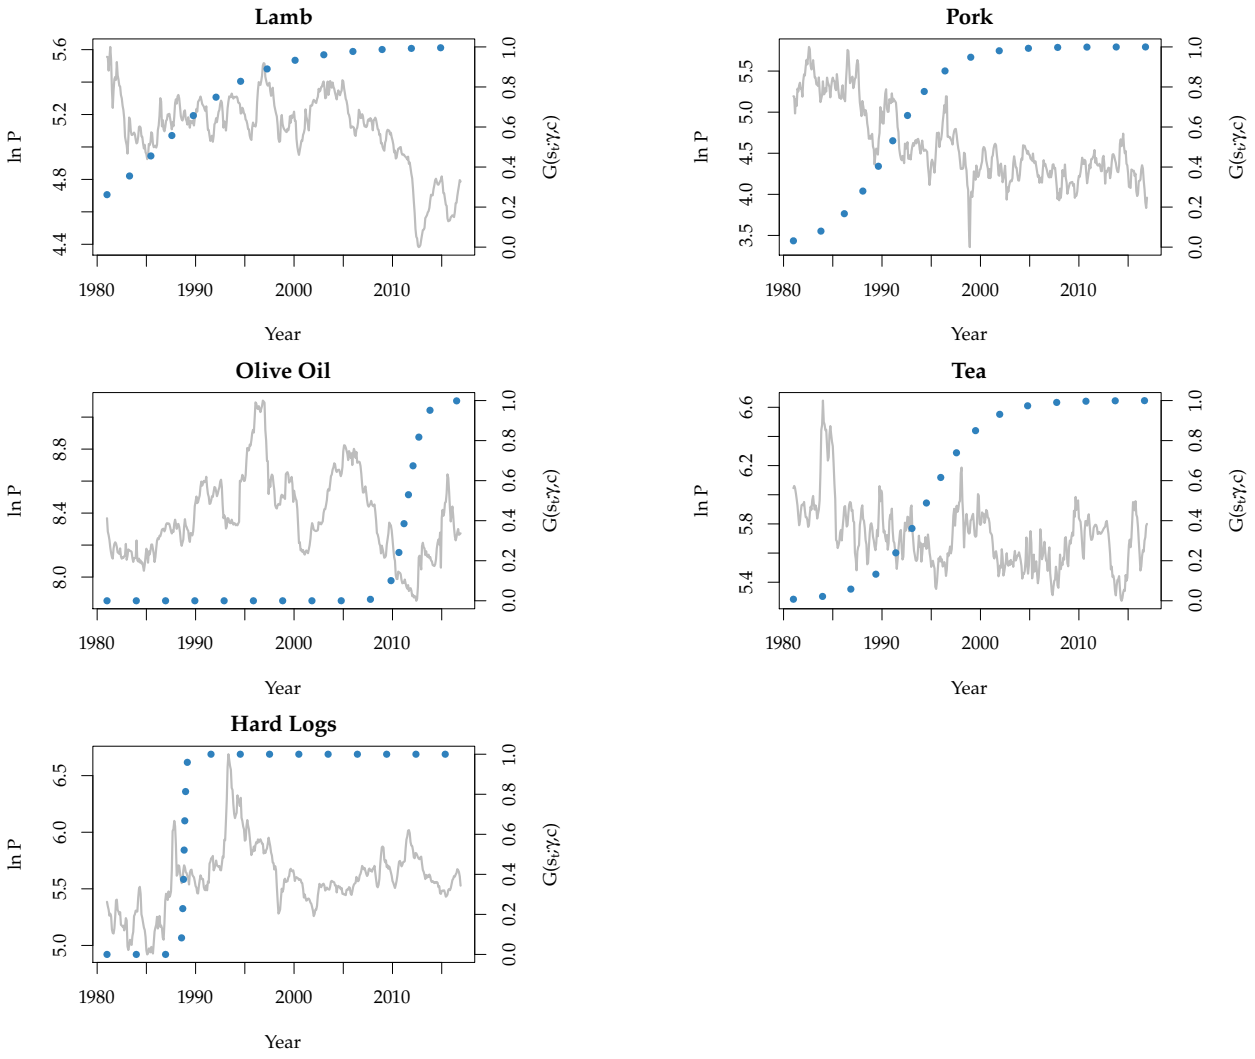

Figure A3: Estimated transition functions of time-varying price equations

Note: These plots feature commodity price series (in gray), as well as the estimated transition functions, $G(t^*; \gamma_\tau, \tau)$, depicted by blue dots.

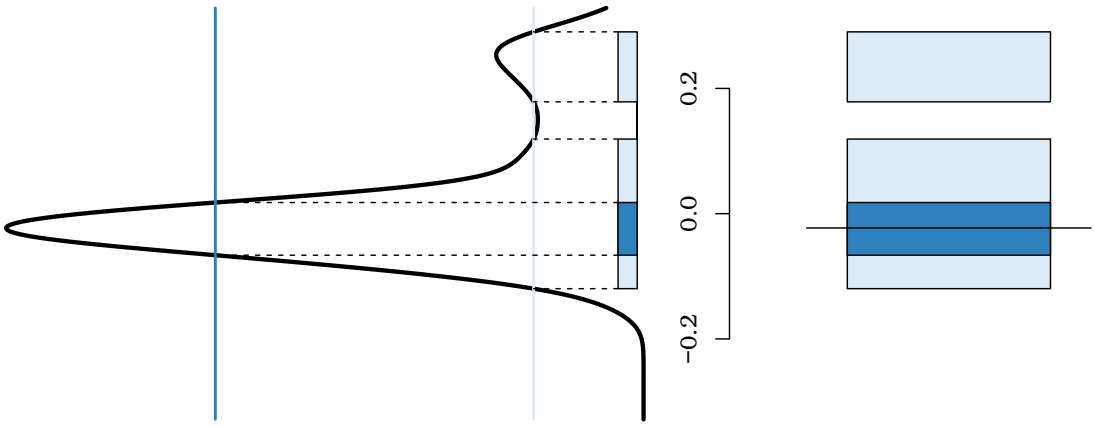

Figure A4: Highest density regions: an illustration