

NBER WORKING PAPER SERIES

COMMODITY-PRICE COMOVEMENT AND GLOBAL ECONOMIC ACTIVITY

Ron Alquist
Olivier Coibion

Working Paper 20003
<http://www.nber.org/papers/w20003>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2014

The authors are grateful to Yuriy Gorodnichenko, Olivier Blanchard, John Bluedorn, Zeno Enders, Julian di Giovanni, Lutz Kilian, Serena Ng, Benjamin Wong, and seminar participants at the Bank of France, the Bundesbank, the Board of Governors, the Centre for Applied Macro and Petroleum Economics conference “Oil and Macroeconomics,” the European Central Bank, the Norges Bank, the Toulouse School of Economics, and UC Irvine for helpful comments. Data for this project were kindly provided by Christiane Baumeister, Lutz Kilian and the trade associations of the aluminum (EEA), copper (ICSG), tin (ITRI), and nickel (INSG) industries. The first draft of this paper was written while Coibion was a visiting scholar at the IMF, whose support was greatly appreciated. The paper was previously distributed under the title “Commodity Price Comovement: Sources and Implications.” The views expressed in the paper are those of the authors and should not be interpreted as reflecting the views of the Bank of Canada, its Governing Council, the International Monetary Fund, the National Bureau of Economic Research, or any other institution with which the authors are or have been affiliated.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Ron Alquist and Olivier Coibion. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Commodity-Price Comovement and Global Economic Activity
Ron Alquist and Olivier Coibion
NBER Working Paper No. 20003
March 2014, Revised August 2014
JEL No. E3,F4

ABSTRACT

Guided by a macroeconomic model in which non-energy commodity prices are endogenously determined, we apply a new factor-based identification strategy to decompose the historical sources of changes in commodity prices and global economic activity. The model yields a factor structure for commodity prices and identification conditions through which factors have economic interpretations: one factor captures the combined contribution of shocks that affect commodity markets only through general equilibrium forces. Applied to a cross-section of commodity prices since 1968, the theoretical restrictions are consistent with the data and yield structural interpretations of the common factors in commodity prices. Commodity-related shocks have contributed modestly to global economic fluctuations.

Ron Alquist
Bank of Canada
234 Wellington Street
Ottawa, ON K1A 0G9
Canada
ralquist@bank-banque-canada.ca

Olivier Coibion
7600 Pitter Pat
Austin, TX 78736
and NBER
ocoibion@gmail.com

1 Introduction

Between January 2003 and July 2008, the prices of most major commodities grew rapidly: wheat by 120%, corn by 122%, copper by 363%, aluminum by 100%, and nickel by 138%. Many observers concluded that the simultaneous rise in prices across such a broad cross-section of commodities reflected a common cause: an increase in the global demand for commodities due to growth in emerging Asia and especially China. Other episodes of widespread comovement in commodity prices have similarly suggested global demand as a common source of commodity price movements, such as those in the early 1970s or in the late 1990s. But this need not be the only explanation: exogenous changes in the prices of oil and other energy products could simultaneously drive the prices of many non-energy commodities because of the important role played by transportation costs in their distribution. What is more, changing preferences on the part of consumers could shift the demand for commodity-intensive products, as could technological changes that affect the relative importance of raw materials in the production of consumption goods. Decomposing the sources of commodity price comovement is therefore inextricably linked to identifying the sources of global business cycle fluctuations.

In this paper, we develop and implement a novel methodology for decomposing the sources of commodity price comovement and global business cycle fluctuations. Underlying this methodology is a general equilibrium model of global business cycles with commodities that predicts a factor structure for real commodity prices. The predicted factor structure decomposes the sources of global business cycle fluctuations and commodity-price movements, and the theory suggests several ways to recover a structural interpretation to the common factors extracted from commodity prices. In other words, this methodology provides a way to use the comovement in commodity prices to disentangle the simultaneous determination of commodity prices and business cycles.

The factor structure in commodity prices predicted by the model separates exogenous forces (or “shocks”) into two types. The first set of shocks are those that directly shift the supply and demand curves for commodities and thus affect commodity prices even in the absence of general equilibrium changes in aggregate income. We refer to these factors as direct factors. They potentially reflect a variety of common shocks to the prices of inputs used to produce commodities, such as labor or energy, common productivity shocks, or demand factors such as changes in the relative need for commodities to produce final consumption goods. The second set of shocks are those that affect commodity prices only indirectly through their effects on aggregate output. We refer to them as indirect factors. The indirect effects can come through two channels. One is the standard demand channel. When aggregate economic activity is high, the demand for commodities used to produce the final good is also high, thereby raising the prices of all commodities. There also can be a supply-side channel. When aggregate income is high, agents may be less willing to supply the inputs used to produce commodities because of income effects, thereby pushing up the prices of commodities. Both channels induce positive comovement in the prices of commodities.

The theory predicts a new result about indirect shocks. Because their effects on commodity prices are summarized entirely by their effects on aggregate output, *each of the shocks induces the same comovement among commodity prices*. As a result, their combined effect on commodity prices can be aggregated into a single factor. Furthermore, this factor has a precise structural interpretation in the model. It corresponds to the counterfactual level of global economic activity that would have obtained in the absence of direct commodity

shocks. Hence, identifying this factor provides a new way to recover historical changes in global economic activity and commodity prices that reflect endogenous responses to non-commodity-related shocks. By contrast, direct shocks induce additional shifts in the supply of or demand for commodities, above and beyond the general equilibrium effects on output, and thus imply a different pattern of comovement among commodities. As a result, each of the direct commodity shocks implies the existence of a distinct factor.

However, because standard empirical factor decompositions identify factors only up to a rotation, one cannot immediately recover the indirect common factor from a simple factor decomposition of commodity prices. The second element of our approach is to impose identification conditions, again grounded in the predictions of the theoretical model, to recover the direct and indirect factors underlying commodity-price movements. The theoretical model provides two ways of doing so: sign restrictions on factor loadings of the indirect common factor and orthogonality conditions given instruments for either the direct or indirect factors. Using a cross-section of forty non-energy commodity prices available since 1968, we apply both identification strategies to identify the indirect factor and find similar results across specifications, indicating that the results are robust to the choice of identification strategy and instruments.

Our main empirical finding is that the vast majority of historical commodity-price movements are associated with the indirect factor, i.e., broad-based changes in commodity prices can largely be attributed to a general equilibrium response to aggregate non-commodity shocks rather than direct shocks to commodity markets. While there are a number of historical episodes during which direct shocks to commodity markets played some role in accounting for commodity-price movements and changes in global production (e.g., 1979-1980, the run-up in commodity prices in the 2000s, and the decline in prices in 2008-09), the primary source of commodity-price movements is their endogenous response to non-commodity-related shocks, as argued in Kilian (2009) in the case of oil prices.

Our approach is related to the literature on the macroeconomic effects oil and commodity price shocks (Bosworth and Lawrence 1982; Hamilton 1983; Barsky and Kilian 2002; Hamilton 2009; and Blinder and Rudd 2012) as well as a growing body of recent research on identifying the sources of oil price movements such as Kilian (2009), Lombardi and Van Robays (2012), Kilian and Murphy (2013), Kilian and Lee (2013). However, we differ from this line of research in a number of ways. First, whereas previous work has focused primarily on oil prices, we focus on a broader range of non-energy commodities, which is essential to implement our identification strategy. Second, our identification strategy is new. Whereas previous work has relied on structural VARs of individual commodity markets or estimated DSGE models, we apply factor methods that decompose the comovement across different commodity prices. We then exploit the predictions about this decomposition from a micro-founded model to identify the structural sources of fluctuations in commodity prices and aggregate output. Third, while identification in structural VARs of commodity markets typically decomposes shocks into supply and demand shocks, our general equilibrium model allows for the fact that exogenous forces can have both supply and demand effects. For example, an increase in productivity in the production of final goods will raise the demand for commodities but may also lower their supply if income effects induce households to restrict the supply of inputs used in the production of commodities. To the extent that income effects are small empirically, the resulting identification of the indirect common factor could be interpreted as primarily reflecting global demand forces, but this interpretation is not something that is imposed in our identification.

We are not the first to apply factor methods to commodity prices. Some papers have examined whether there is excess comovement among unrelated commodities – that is, price comovement in excess of what one would expect conditional on macroeconomic fundamentals (Pindyck and Rotemberg 1990; Deb, Trivedi, and Varangis 1996; and Ai, Chatrath, and Song 2006). Other papers have investigated the forecasting performance of the common factor in metals prices for individual metals prices (West and Wong 2014) and commodity convenience yields for inflation (Gospodinov and Ng 2013). But there has been little attempt at interpreting the resulting factors in a structural sense.

Our model provides a structural interpretation to a factor representation for commodity prices along with the requisite identification conditions, so that we are able to disentangle the different *economic* channels underlying commodity-price movements. In this respect, our approach is related to work that uses economic theory to assign factors an economic interpretation. For example, Forni and Reichlin (1998) impose constraints guided by economic theory on common factors to identify technological and non-technological shocks (see also Gorodnichenko 2006). Another set of papers has identified the factors driving macroeconomic aggregates common to all countries and specific subsets of countries (Stock and Watson 2003; and Kose et al. 2012). This approach has also been used to identify relative price changes for specific goods and the absolute price changes common to all goods (Reis and Watson 2010) and the relative importance of aggregate and sector-specific shocks that have driven U.S. industrial production (Foerster et al. 2011). Our paper differs from this line of research in that we use commodity price dynamics to identify the sources of global business cycle fluctuations and in our identification strategy, which relies on the use of sign restrictions and orthogonality conditions rather than zero restrictions on the factor loadings.

We also show that our factor-based method can help with forecasting commodity prices. Using recursive out-of-sample forecasts, we find that a bivariate factor-augmented VAR (FAVAR) that includes each commodity's price and the first common factor extracted from the cross-section of commodities generates improvements in forecast accuracy relative to the no-change forecast, particularly at short (1-, 3-, and 6-month) horizons. This result extends to broader commodity price indices, such as the CRB spot index, the World Bank non-energy index, and the IMF index of non-energy commodity prices. We also find that the indirect common factor extracted from the cross-section of commodity prices helps to predict real oil prices, again with the largest gains being at short horizons (e.g., 20% reductions in the MSPE at the 1-month horizon). These improvements in oil forecasting accuracy are similar in size to those obtained using oil-market VARs in Baumeister and Kilian (2012) and Alquist et al. (2013). But unlike the monthly oil-market VARs, our approach relies only on a cross-section of commodity prices that can readily be updated at monthly or quarterly frequencies. This is an important advantage because production and inventory data for commodities are often unavailable at these frequencies. Our factor-based approach thus provides a unified framework to forecast commodity prices and a structural interpretation to these forecasts.

The rest of the paper is organized as follows. Section 2 presents a general equilibrium business cycle model with commodities and shows how the model can be used to assign the common factors in commodity prices a structural interpretation. The section also shows how the model permits an econometrician to recover the economic factors from typical factor decompositions through identification restrictions. Section 3 applies these results to a historical cross-section of commodity prices. Section 4 considers the implications of

commodity storage while section 5 uses the indirect common factor in a recursive out-of-sample forecasting exercise. Section 6 concludes.

2 The Sources of Commodity Price Comovement: Theory

In this section, we present a model that characterizes the sources of commodity-price comovement. In particular, we show that the model yields a tractable factor structure for a cross-section of commodity prices and that permits an economic interpretation of the factors.

2.1 Model of commodity prices

The baseline model consists of households, a continuum of heterogeneous primary commodities, a sector that aggregates these commodities into a single intermediate commodity input, and a final goods sector that combines commodities, labor and technology into a final good.

The Household

A representative consumer maximizes expected discounted utility over consumption (C), labor supply (N^S), and the amount of another input supplied to each commodity sector ($L^S(j)$) as follows

$$\max E_t \sum_{i=0}^{\infty} \beta^i \left[\frac{C_{t+i}^{1-\sigma}}{1-\sigma} - e^{-\varepsilon_{t+i}^n} \varphi_n \frac{N_{t+i}^S \frac{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} - \varphi_L e^{-\varepsilon_t^L} \frac{\int_0^1 L_{t+i}^S(j)^{1+\frac{1}{\nu}} dj}{1+\frac{1}{\nu}} \right]$$

where β is the discount factor. We refer to the input supplied to each commodity sector as “land” simply to differentiate it by name from the labor input provided to the final goods sector. But to be clear, “land” is simply a label: the variable is an input into the production process for primary commodities and can be interpreted in many different ways. For example, one could interpret the input as another form of labor that cannot be reallocated across sectors. In this case, one can think of N^S as the supply of labor to manufacturing or service sectors, whereas L^S could be thought of as the supply of labor to the mining and agricultural sectors. Alternatively, one can interpret the input literally as land. In that case, the use of land generates direct benefits to the household and hence is included in the utility function, but it can also be provided to commodity producers (e.g., for farming or mining) in exchange for a rental payment. The assumption that this input enters the utility function, along with the introduction of the preference shifter ε_t^L , is a reduced form way of generating an upward-sloping supply curve for the input into the commodity production process, but the specific mechanism used does not play an important role in the analysis. The same qualitative results would apply if this input did not enter into the utility function so that the household supplied its total endowment each period. With $\varphi_n > 0$ and $\varphi_L > 0$, welfare is decreasing in hours worked and the amount of land supplied to commodity sectors. The $e^{\varepsilon_t^n}$ term is an exogenous shock to the disutility of hours worked while $e^{\varepsilon_t^L}$ is an exogenous shock to the disutility of supplying land.

The household pays a price P_t for the consumption good, receives wage W_t for each unit of labor supplied, and is paid a rental rate of land $R_t^L(j)$ for each unit of land supplied to the primary commodity sector j . The household also can purchase risk-free bonds B_t that pay a gross nominal interest rate of R_t . The budget constraint is

$$P_t C_t + B_t = B_{t-1} R_{t-1} + W_t N_t^s + \int_0^1 R_t^L(j) L_t^S(j) dj + T_t$$

where T_t represent payments from the ownership of firms.

The Primary Commodity-Production Sector

Each primary commodity j is produced by a representative price-taking firm who uses land ($L_t^d(j)$) to produce a quantity $Q_t(j)$ of good j given a production function

$$Q_t(j) = A_t(j) L_t^d(j)^{\alpha_j}$$

where $A_t(j)$ is the exogenously determined level of productivity for commodity j and $0 < \alpha_j < 1$ is the commodity-specific degree of diminishing returns to land. Given the price of commodity j $P_t(j)$, and the rental rate of land $R_t^L(j)$ specific to commodity j , the firm chooses the amount of land input to maximize profits

$$\max P_t(j) Q_t(j) - R_t^L(j) L_t^d(j).$$

This yields the following demand curve for land for each commodity j :

$$R_t^L(j)/P_t = \alpha_j \left(\frac{P_t(j)}{P_t} \right) A_t(j) L_t^d(j)^{\alpha_j - 1}.$$

We assume that the steady-state level of productivity $\overline{A(j)}$ is such that the steady-state level of production in each sector is equal. Equilibrium in the market for land requires

$$L_t^S(j) = L_t^d(j)$$

for each sector j .

The Intermediate Commodity

A perfectly competitive sector purchases $Y_t(j)$ of each primary commodity j and aggregates them into an intermediate commodity Q_t^C using the Dixit-Stiglitz aggregator

$$Q_t^C = \left(\int_0^1 Y_t^j \frac{\theta_c - 1}{\theta_c} dj \right)^{\frac{\theta_c}{\theta_c - 1}}$$

which yields a demand for each commodity j of

$$P_t(j)/P_t^C = (Y_t(j)/Q_t^C)^{-1/\theta_c}$$

where θ_c is the elasticity of substitution across commodities and the price of the intermediate commodity aggregate is given by $P_t^C = \left(\int_0^1 P_t(j)^{1-\theta_c} dj \right)^{\frac{1}{1-\theta_c}}$. Market-clearing for each commodity sector j requires

$$Q_t(j) = Y_t(j).$$

Note that the setup implicitly assumes that no storage of commodities takes place since all commodities produced must be used contemporaneously. We discuss the rationale for this assumption and its implications in more detail in section 4.

The Final Goods Sector

A perfectly competitive sector combines purchases of the intermediate commodity good Y_t^C and labor N_t^d according to the Cobb-Douglas production function

$$Y_t = A_t Y_t^C \alpha_t N_t^d^{1-\alpha_t}$$

to maximize profits

$$P_t Y_t - W_t N_t^d - P_t^C Y_t^C$$

taking as given all prices and where A_t is an exogenously determined aggregate productivity process. This yields the following demand for each input

$$\begin{aligned}\alpha_t &= (P_t^C/P_t)(Y_t^C/Y_t) \\ 1 - \alpha_t &= (W_t/P_t)(N_t^d/Y_t)\end{aligned}$$

Since all of the final good is purchased by the household, equilibrium in the final goods market requires $C_t = Y_t$. The fact that α_t is potentially time-varying allows for exogenous variation in the relative demand for commodities and labor in the production of the final good.

The Linearized Model

A detailed solution of the model is provided in Appendix A. We assume that exogenous processes are stationary around their steady-state levels, so that all real variables are constant in the steady-state. Lower-case letters denote log-deviations from steady-state (e.g., $c_t \equiv \log C_t - \log \bar{C}$), and we normalize the nominal variables by the final goods price level (e.g., $p_t(j) \equiv \log P_t(j)/P_t - \log(\bar{P}(j)/\bar{P})$). We normalize commodity productivity shocks as $v_t(j) \equiv a_t(j) \left(1 + (\varepsilon_j \theta_c)^{-1}\right)^{-1} (1 + \varepsilon_j^{-1})$ to simplify the aggregation across commodities, where $\varepsilon_j \equiv (v^{-1} + 1 - \alpha_j)/\alpha_j$.¹ We assume that productivity shocks to each commodity sector have an idiosyncratic component and a common component such that $v_t(j) = v_t^a + v_t^j$, which implies that the sum of productivity across commodities is $v_t \equiv \int_0^1 v_t(j) dj = v_t^a$. The idiosyncratic shocks are orthogonal across commodity sectors, such that $E[v_t^j v_t^k] = 0 \forall j \neq k$ and $E[v_t] = 0$. The log-deviation of α_t from its steady-state value of α is denoted by $\check{\alpha}_t$.

The aggregate level of production of final goods follows

$$y_t = \varphi_y [a_t + \kappa_L \varepsilon_t^L + \kappa_n \varepsilon_t^n + \kappa_v v_t + \kappa_\alpha \check{\alpha}_t] \quad (1)$$

where $\varphi_y \equiv \left(1 - \alpha \left[\frac{(1-\sigma)\theta_c \varphi}{1+(\theta_c-1)\varphi}\right] - (1 - \alpha) \left[\frac{1-\sigma}{1+\eta^{-1}}\right]\right)^{-1}$, $\kappa_L \equiv \frac{\alpha \theta_c \varphi}{1+(\theta_c-1)\varphi}$, $\kappa_n \equiv \frac{1-\alpha}{1+\eta^{-1}}$, $\kappa_v \equiv \frac{\alpha}{1+(\theta_c-1)\varphi}$, $\kappa_\alpha \equiv \varphi_\alpha + \frac{\alpha \varphi \theta_c}{1+\varphi(\theta_c-1)} - \frac{\alpha}{1+\eta^{-1}}$, $\varphi_\alpha \equiv \alpha(\ln \bar{Y}^c - \ln \bar{N})$, and $\varphi \equiv \int_0^1 (1 + \varepsilon_j \theta_c)^{-1} dj$. Output is rising with aggregate productivity, positive shocks to the household's willingness to supply land and labor, and a positive sum over commodity-specific productivity shocks. Whether output rises when the relative demand for commodities increases ($\check{\alpha}_t$) depends on specific parameter values.

2.2 Comovement in Commodity Prices

We now consider the determinants of commodity prices. First, the supply of commodity j is given by

$$p_t(j) = \varepsilon_j y_t(j) - \frac{(1+\varepsilon_j \theta_c)}{\theta} (v_t^a + v_t^j) + \sigma y_t - \varepsilon_t^L \quad (2)$$

where ε_j is the elasticity of commodity supply with respect to its price. First, changes in aggregate output shift the supply curve when income effects on the input are present ($\sigma > 0$). This implies that all macroeconomic

¹ The rescaling of the commodity-specific productivity shock ensures that a 1% increase in productivity in each commodity sector raises the equilibrium level of production of that commodity by equal amounts for each commodity. This would not be the case without the rescaling because each primary commodity sector's supply curve has a different slope.

shocks that affect aggregate production in the model cause an equal upward or downward shift in the supply of every commodity in general equilibrium. Hence, all shocks in the model are, in a sense, supply shocks to commodities. Second, the supply of commodity j increases whenever its productivity level rises, which can reflect common productivity shocks (v_t^a) or idiosyncratic shocks (v_t^j). Finally, shocks to the household's willingness to supply land to the commodity sector directly affect the supply curve. Thus, we can write the supply curve of commodity j more succinctly as

$$p_t(j) = S_j(y_t(j); y_t(a_t, \varepsilon_t^n, \varepsilon_t^L, \check{\alpha}_t, v_t^a); v_t^a, \varepsilon_t^L, v_t^j) \quad (2')$$

which captures the fact that some shocks affect the supply of commodity j indirectly through general equilibrium effects captured by aggregate output, some shocks affect supply directly by shifting the curve holding aggregate output fixed, and some shocks do both.

The demand for commodity j is

$$p_t(j) = -\frac{1}{\theta_c} y_t(j) + \left(\frac{1+(\theta_c-1)\sigma\varphi}{1+(\theta_c-1)\varphi} \right) y_t - \frac{\varphi(\theta_c-1)}{1+(\theta_c-1)\varphi} \varepsilon_t^L - \frac{(\theta_c-1)}{1+(\theta_c-1)\varphi} \left(\frac{1}{\theta_c} \right) v_t^a + \frac{1}{1+\varphi(\theta_c-1)} \check{\alpha}_t. \quad (3)$$

Demand for commodity j is increasing with aggregate output, which reflects the role of commodities as an input into the production of final goods. This term therefore captures general equilibrium demand effects, and all macroeconomic shocks that affect aggregate production in the model result in an equal upward or downward shift in the demand for each commodity. Thus, all shocks in the model other than idiosyncratic shocks are both demand and supply shocks. But in addition to these general equilibrium shifts in commodity demand, the demand for commodity j rises with changes in the relative demand for commodities ($\check{\alpha}_t$), holding aggregate output constant. It also shifts, holding aggregate output constant, with exogenous changes in the household's willingness to supply land and with exogenous common commodity productivity shocks. While the latter two would more commonly be thought of as supply shocks, the effect that they have on all commodities implies that they affect equilibrium prices and quantities of the intermediate commodity bundle and thus the demand for each commodity through the CES structure. We can again write the demand curve of commodity j more succinctly as

$$p_t(j) = D_j(y_t(j); y_t(a_t, \varepsilon_t^n, \varepsilon_t^L, \check{\alpha}_t, v_t^a); v_t^a, \varepsilon_t^L, \check{\alpha}_t) \quad (3')$$

to highlight that some shocks affect the demand for commodity j indirectly through general equilibrium effects on output, some shocks shift the demand for each commodity j directly holding aggregate output constant, and some do both.

In this setting, there are no well-defined supply and demand shocks to a given commodity, so identification procedures that rely on supply and demand characterizations may be ill-defined. However, the comovement across commodities can help to resolve this identification problem. Consider, for example, the effect of an aggregate productivity shock (a_t) on commodity prices. Such a shock affects both supply and demand for every commodity, but it does so *only through its equilibrium effects on aggregate output*. A positive productivity shock in this setting would increase output and thereby increase the demand for each commodity j and decrease its supply through income effects. Both effects tend to increase the prices of all commodities. While the size of the effect differs across commodities with the slopes of their supply curves (which, in turn, depend on the α_j 's), there is necessarily positive price comovement implied by such shocks.

This point is illustrated visually in the left graph of Panel A in Figure 1, which shows the price implications of an increase in aggregate productivity for a commodity with relatively elastic supply $S_E(y(a))$

and one with relatively inelastic supply $S_j(y(a))$. The graph on the right plots the set of prices for the two commodities that result from different levels of aggregate productivity, denoted by $R(a_t)$. Higher levels of productivity increase the prices of both commodities, so that $R(a_t)$ is upward-sloping. This example illustrates the positive commodity-price comovement that results from productivity shocks.

Importantly, *any* shock that affects commodity prices only through its effects on aggregate output induces the *same* relative comovement of commodity prices as productivity shocks. In the model, shocks to the household's willingness to supply labor (ε^n) also affect commodity prices only through y_t and delivers the same pattern of comovement among commodities as an aggregate productivity shock, i.e. $R(a_t) = R(\varepsilon^n)$. While there are only two exogenous variables in the model that affect commodity prices solely through general equilibrium effects, one could readily integrate a wider set of such forces into a more complex model. For example, if differentiated forms of labor were used in the production of final goods, then variation in households' willingness to supply each form of labor would generate the same comovement. Another example is if the final good were produced under imperfect competition, exogenous variation in the desired markups would again generate the same pattern of comovement in commodity prices.

By contrast, any shock that directly (i.e., holding aggregate output constant) affects the supply or demand of a commodity induces different comovement among commodities. This point is illustrated in Panel B of Figure 1 for the case of a decrease in the relative demand for commodities (from $\check{\alpha}_t$ to $\check{\alpha}'_t$) that is then assumed to raise aggregate output. (The covariance of $\check{\alpha}_t$ and y_t in the model depends on specific parameter values). The decline in $\check{\alpha}_t$ has two effects on the supply and demand for commodities. The first effect is the indirect general equilibrium effect operating through aggregate activity. Given our assumption that the decline in $\check{\alpha}_t$ raises y_t , this effect shifts the demand and supply of commodities in exactly the same way as an increase in aggregate productivity. The second effect is the direct decrease in the demand for commodities, illustrated graphically by $D(y(\check{\alpha}_t), \check{\alpha}'_t)$, so that the combined effect on demand for commodities is given by the demand curve $D(y(\check{\alpha}'_t), \check{\alpha}'_t)$. As a result of these shifts, the prices of both commodities are again higher, but the price of the elastically supplied commodity increases by more than that of the inelastically supplied commodity, yielding a different pattern of commodity-price comovement. The latter is illustrated graphically on the right-hand side of Panel B in Figure 1. $R(\check{\alpha}_t)$, the set of possible prices of the two commodities for different levels of $\check{\alpha}_t$, is flatter than what obtains for changes in aggregate productivity or changes in the household's willingness to supply labor. Indeed, any shock that has both direct and indirect effects on commodity markets leads to a different pattern of comovement among commodities than shocks that have only indirect effects.

2.3 The Factor Structure in Commodity Prices

To solve for commodity prices, we combine equations (2) and (3) yielding

$$p_t(j)(1 + \varepsilon_j \theta_c) = \left[\sigma + \frac{\varepsilon_j \theta_c (1 + (\theta_c - 1) \sigma \varphi)}{1 + (\theta_c - 1) \varphi} \right] y_t - \left[\frac{\varepsilon_j \theta_c \varphi (\theta_c - 1)}{1 + (\theta_c - 1) \varphi} + 1 \right] \varepsilon_t^L \\ - \frac{1}{\theta_c} \left(1 + \varepsilon_j \theta_c + \frac{\varepsilon_j \theta_c (\theta_c - 1)}{1 + (\theta_c - 1) \varphi} \right) v_t^a + \frac{\varepsilon_j \theta_c}{1 + \varphi (\theta_c - 1)} \check{\alpha}_t - \frac{1}{\theta_c} (1 + \varepsilon_j \theta_c) v_t^j(j).$$

Because aggregate output y_t is a function of all aggregate shocks in the model, we can decompose it as follows

$$y_t = y_t^{nc}(a_t, \varepsilon_t^n) + \varphi_y [\kappa_L \varepsilon_t^L + \kappa_v v_t^a + \kappa_\alpha \check{\alpha}_t]$$

where $y_t^{nc} = \varphi_y [a_t + \kappa_n \varepsilon_t^n]$ is the level of aggregate output coming exclusively from changes in aggregate productivity and changes in the household's willingness to supply labor to the final goods sector. We can then rewrite the equilibrium price of commodity j as

$$p_t(j) = \underbrace{\lambda_j^y y_t^{nc}(a_t, \varepsilon_t^n)}_{\text{indirect (IC)}} + \underbrace{\lambda_j^L \varepsilon_t^L + \lambda_j^v v_t^a + \lambda_j^\alpha \check{\alpha}_t}_{\text{direct (DC)}} - \underbrace{\frac{1}{\theta_c} v_t^j(j)}_{\text{idiosyncratic}} \quad (4)$$

$$= \lambda_j F_t + \xi_t^j$$

where $\lambda_j^y \equiv (1 + \theta \varepsilon_j)^{-1} \left[\sigma + \frac{\varepsilon_j \theta (1 + (\theta - 1) \sigma \varphi)}{1 + (\theta - 1) \varphi} \right]$, $\lambda_j^L \equiv \varphi_y \kappa_L \lambda_j^y - \left[\frac{1}{1 + \varepsilon_j \theta_c} + \frac{\varepsilon_j \theta_c}{1 + \varepsilon_j \theta_c} (\theta_c - 1) \right]$, $\lambda_j^v \equiv \varphi_y \kappa_v \lambda_j^y - \left[\frac{1}{\theta_c} + \frac{\varepsilon_j \theta_c}{1 + \varepsilon_j \theta_c} \left(\frac{\varphi (\theta_c - 1)}{\varphi} \right) \right]$, $\lambda_j^\alpha \equiv \varphi_y \kappa_\alpha \lambda_j^y + \frac{\varepsilon_j \theta_c}{1 + \varepsilon_j \theta_c} \left(\frac{1}{1 + \varphi (\theta_c - 1)} \right)$, $\lambda_j \equiv [\lambda_j^y \lambda_j^L \lambda_j^v \lambda_j^\alpha]$, $F_t \equiv [y_t^{nc} \varepsilon_t^L v_t^a \check{\alpha}_t]$, and $\xi_t^j \equiv -\frac{1}{\theta_c} v_t^j$.

Equation (4) provides a factor structure for real commodity prices with three distinct and orthogonal components. The last term on the right-hand side reflects idiosyncratic shocks to commodity j that have no aggregate real effects. The second term on the right-hand side consists of a factor for each shock that has direct effects on commodity market (i.e., that shifts the supply or demand for commodities holding aggregate output constant). For this reason, we refer to these factors as “direct common” (DC) factors. In this setup, there are three such factors: common shocks to the input used in the production of commodities, a common productivity shock, and a shock to the relative demand for commodities in the production of final goods. Each enters as a separate factor because each shifts supply and demand curves in different ways and therefore has distinct implications for the price of a single commodity. Because these forces have both direct and indirect effects on the market for commodity j , there is in general no guarantee that their respective loadings have the same signs across commodities.

The most interesting component of the factor structure is the first term on the right-hand side of (4), which reflects the *combined* contribution on the price of commodity j from all shocks whose effects on commodity prices operate only indirectly through aggregate output (i.e., only through general equilibrium effects). We refer to this common factor as the “indirect common” (IC) factor. It captures the fact that, because some shocks affect commodity markets only through changes in aggregate output, they all have identical implications for the price of a given commodity conditional on the size of their effect on aggregate output and induce the same comovement across different commodity prices. As a result, they can be represented as a single factor. Furthermore, this factor has a well-defined interpretation: it is the level of global output that would have occurred *in the absence of any direct commodity shocks*. Thus, this common factor represents a way to reconstruct the counterfactual history of aggregate output without direct commodity shocks, as well as to decompose historical commodity-price changes into those components reflecting direct commodity shocks versus all other aggregate economic forces captured by the IC factor. Unlike the DC factors, another key characteristic of the IC factor is that the loadings on this factor must all be positive ($\lambda_j^y > 0 \forall j$). This prediction reflects the fact that the shocks incorporated in the IC factor raise commodity demand when the shock is expansionary and simultaneously restrict the commodity supply through income effects, which unambiguously increases commodity prices. Finally, in the absence of income effects on the common input into the production of commodities, the IC could be interpreted as capturing exogenously driven global

demand for commodities. In short, this factor decomposition provides a new way to separate causality in the presence of simultaneously determined prices and production levels.

2.4 Recovering a Structural Interpretation of the Factors

A key limitation of factor structures is that, empirically, factors are only identified only up to a rotation. For example, if one estimated a factor structure on commodity prices, one could not directly associate the extracted factors with the structural interpretation suggested by (4). However, the theory developed in this section has implications that can be used to identify the unique rotation consistent with those predictions and permits us to assign the factors driving commodity prices an economic interpretation.

To see this, suppose that as in the theory above, the N variables in vector X_t (N by 1) of real commodity prices have a factor structure:

$$X_t = LF_t + \varepsilon_t$$

where F_t is a K by 1 vector of unobserved variables, and L is an N by K matrix of factor loadings. Let the variance of ε_i be given by φ_i and the covariance matrix of ε_i be $cov(\varepsilon) = diag(\varphi_i) = \Psi$ such that the ε_i s are uncorrelated with one another. We make the typical assumptions underlying factor analysis: a) $E(F) = 0$, b) $E(\varepsilon_i) = 0$, c) $E(F\varepsilon_i) = 0$, and d) $cov(F) = I$ so that the factors are orthogonal to one another and have variance normalized to one. Then, letting $\Sigma \equiv cov(X)$ be the covariance matrix of X , it follows that $\Sigma = LL' + \Psi$. The identification problem is that for any K by K orthogonal matrix T such that $TT' = I$, we can define $\tilde{L} = LT$ and $\tilde{F}_t = T'F_t$ such that

$$X_t = \tilde{L}\tilde{F}_t + \varepsilon_t.$$

As a result, an empirical estimate of the factors underlying X_t do not, in general, permit the economic identification of the factors F_t but rather some rotation \tilde{F}_t .

However, the model provides additional restrictions on the factor structure that can be used to assign the factors an economic interpretation and recover the “structural” factors F_t from the estimated factors \tilde{F}_t . For example, consider the factor structure of equation (4) in section 2.3 in which real commodity prices reflect two underlying factors: a common commodity-related shock (ε_t^L) and the level of aggregate production that would have occurred in the absence of this shock (y_t^{nc}), so $F_t = [F_t^1 F_t^2]' = [y_t^{nc} \varepsilon_t^L]'$. As we discuss below, this two-factor structure is the most empirically relevant case. A factor decomposition of commodity prices would yield some rotation of these factors \tilde{F}_t such that

$$F_t = T'\tilde{F}_t = \begin{bmatrix} t_{11} & t_{12} \\ t_{21} & t_{22} \end{bmatrix}' [\tilde{F}_t^1 \tilde{F}_t^2]' = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} [\tilde{F}_t^1 \tilde{F}_t^2]'$$

where the last equality reflects the properties of rotation matrices. Recovering the “structural” factors F_t corresponds to identifying the parameter θ and the rotation matrix T such that $F_t = T'\tilde{F}_t$.

The theory imposes three types of conditions that potentially can be used to identify θ . The first is that y_t^{nc} (the IC factor) is orthogonal to commodity-related shocks (DC factors). Hence, if one had a 1 by S vector of instruments z_t that is correlated with the commodity-related shocks ε_t^L , the orthogonality of y_t^{nc} would deliver S moment conditions $E[y_t^{nc} z_t] = 0$. The conditions can be rewritten as

$$E[y_t^{nc} z_t] = E[(\tilde{F}_t^1 \cos \theta + \tilde{F}_t^2 \sin \theta) z_t] = 0. \quad (5)$$

If $S = 1$, then θ would be uniquely identified. If $S > 1$, then θ is overidentified, and one could estimate it using standard GMM methods by writing the moment conditions as

$$J(\theta) = E[y_t^{nc} z_t] W E[y_t^{nc} z_t]', \quad (6)$$

where W is a weighting matrix, such that $\hat{\theta} = \text{argmin} J(\theta)$. Letting W be the inverse of the variance-covariance matrix associated with the moment conditions, standard GMM asymptotic results apply including standard errors for θ and tests of the over-identifying conditions for N and T large enough for the factors to be considered as observed variables rather than generated (e.g., Stock and Watson 2002; and Bai and Ng 2002).

A second approach would be to make use of the theoretical prediction that y_t^{nc} is a linear combination of exogenous variables that have only indirect effects on the commodity sector such as the productivity shocks or labor supply shocks considered in the model. If one had a vector of S by 1 instruments z_t for each period correlated with one or more of these exogenous drivers, then another set of orthogonality conditions imposed by the theory would be $E[F_t^2 z_t] = 0$. As in the previous case, one could estimate θ using GMM given these orthogonality conditions and test over-identifying restrictions if $S > 1$.

In both of these cases, the econometrician must take a stand on whether the chosen instruments should be correlated with “commodity-related” shocks or with y_t^{nc} . While economic theory may provide clear guidance in some cases, this choice may be problematic when one is interested in whether an exogenous variable affects commodities only through general equilibrium effects or more directly. Within our framework, this question amounts to whether the exogenous variable should be considered part of y_t^{nc} or one of the “commodity-related” shocks. For example, in the case of commodity prices, monetary policy shocks could potentially have direct effects on commodity markets in the presence of storage motives but would otherwise not be expected to have direct effects on commodity markets if the speculative channel is absent or sufficiently small. We return to this particular point in section 4.

A third approach is to make use of sign restrictions on the loadings. The theory predicts that the loadings on y_t^{nc} must all be positive (since $\lambda_j^y > 0 \forall j$ in equation (4)). Letting \tilde{L} be the N by 2 matrix of unrotated factor loadings, the rotated or “structural” loadings are $L = \tilde{L}T = [\tilde{L}^1 \tilde{L}^2]T$. The loadings on the first rotated factor (corresponding to y_t^{nc}) are then $L^1 = \tilde{L}^1 \cos \theta + \tilde{L}^2 \sin \theta$. Imposing that all of the elements of L^1 be positive would therefore correspond to identifying the range of values of θ such that $\min(\tilde{L}^1 \cos \theta + \tilde{L}^2 \sin \theta) > 0$. In general, this leads only to a set of admissible values of θ and associated rotation matrices without uniquely identifying the rotation matrix. This approach would be akin to the weak identification of VAR’s by sign restrictions (Uhlig 2002).

In short, the theoretical model of commodity prices yields not only a factor structure for commodity prices but also a set of conditions that can be used to identify (or, in the case of sign restrictions, limit the set of) the rotation matrix necessary to recover the underlying factors. Furthermore, the factors have economic interpretations. The IC factor corresponds to the level of production and income net of commodity-related shocks, while other factors correspond to one or more of these commodity-related shocks. The identification of the rotation matrix, and thus the underlying economic factors, follows from orthogonality conditions implied by the model, as well as sign restrictions on the loadings predicted by the theory. The implied factor structure of the model combined with the ability to recover an economic interpretation of the factors thus provides a new method for separating fluctuations in aggregate output into those driven by commodity-related shocks and those driven by non-commodity-related shocks.

3 The Sources of Commodity Price Comovement: Empirical Evidence

In this section, we implement the factor decomposition of real commodity prices suggested by the theory. We first construct a historical cross-section of real commodity prices for the commodities that conform to the theoretical structure of the model along several dimensions. We then implement a factor decomposition and identify the factors suggested by the theory. After considering a wide range of robustness checks, we argue that commodity-related shocks have contributed only modestly to fluctuations in global economic activity.

3.1 Data

The selection of the commodities used in the empirical analysis is guided by the theoretical model. In particular, we use four criteria to decide which commodities to include in the data set and which to exclude. First, commodities must not be vertically integrated. Second, the main use of commodities must be directly related to the aggregate consumption bundle, and they should not be primarily used for the purposes of financial speculation. Third, commodities must not be jointly produced. Finally, the pricing of commodities must be determined freely in spot markets and must not display the price stickiness associated with the existence of long-term contractual agreements.

The first criterion, that commodities must not be vertically integrated, conforms to the structure of the model in which the only direct interaction between commodities is through their use in the production of the aggregate consumption good. Vertically integrated commodities would introduce the possibility of price comovement due to idiosyncratic shocks to one commodity, thereby affecting prices in other commodities through the supply chain. For example, an exogenous shock to the production of sorghum would affect the price of non-grass-fed beef because sorghum is primarily used as feed. Thus, this shock could ultimately affect the price of milk and hides as well. To satisfy this condition, we exclude from the sample a number of commodities that are frequently incorporated in commodity price indices. For example, we exclude prices of non-grass-fed cattle, poultry (broilers), milk, hogs, lard, pork bellies, eggs, tallow, and hides. In the same spirit, we exclude energy commodities and any fertilizer products.² In addition, when commodities are available in closely related forms (e.g., soybeans, soybean meal and soybean oil), we use at most one of the available price series. The second criterion ensures that the primary forces driving the prices of the included commodities are related to the production and consumption of each commodity. Some commodities, such as precious metals, have long been recognized as behaving more like financial assets than normal commodities (Chinn and Coibion 2013). Thus, we exclude gold, silver, platinum and palladium from the cross-section of commodities.

The third criterion reflects the fact that some commodities are derivative products of the production of other commodities. This is particularly the case for minerals, which are commonly recovered during the mining for metal commodities. For example, antimony and molybdenum are derivatives of copper mining while cadmium is recovered during mining for zinc. For this type of commodities, the assumption of orthogonal productivity shocks is clearly inapplicable. The fourth criterion is that the prices of commodities be primarily determined in spot markets rather than through contractual agreements or government regulations. While many commodities have long been traded on liquid international spot markets, this is not the case for

² Another reason to exclude energy prices is that, in the model, it is assumed that each commodity is too small for its idiosyncratic shocks to have aggregate implications. This condition would almost certainly not apply to energy commodities.

other commodities. For example, the price measure of tung oil (primarily used for wood-finishing) tracked by the Commodity Research Bureau Statistical Yearbooks varies little over time and is often fixed for periods lasting as long as one year. Because we want to focus on commodities whose prices reflect contemporaneous economic conditions, we exclude commodities such as tung oil that systematically display long periods of price invariance. For some commodities in the sample, prices were not determined in flexible markets until much later than others; for these commodities we treat early price data as missing values (e.g., aluminum prior to 1973). For mercury, the reverse is true as its use declined over time and its price begins to display long periods with no price changes starting in 1995. We treat its prices after March 1995 as missing. Appendix B provides more details on these adjustments.

Applying these criteria leaves us with forty commodities in the sample. It includes twenty-two commodities that we refer to as agricultural or food commodities: apples, bananas, barley, cocoa, coffee, corn, fishmeal, grass-fed beef, hay, oats, onion, orange juice concentrate, pepper, potatoes, rice, shrimp, sorghum, soybeans, sugar, tea, tobacco, and wheat. The data set also includes five food oils: coconut, groundnuts (peanut), palm, rapeseed (canola), and sunflower (safflower). Finally, we have 13 industrial commodities: aluminum, burlap, cement, cotton, copper, lead, mercury, nickel, rubber, tin, wool, and zinc. We compiled these monthly data from January 1957 to January 2013 (as available) from a number of sources including the CRB Statistical Yearbooks, the CRB InfoTech CD, the World Bank (WB) GEM Commodity Price Data, the International Monetary Fund's (IMF) Commodity Price data, and the Bureau of Labor Statistics. While most of the data are consistently available from 1968:1 until 2013:1, in some cases there are a number of missing observations in the underlying data, as well as periods when we treat the available data as missing because spot trading was limited. Appendix C provides details on the construction of each series, their availability, and any periods over which we treat the data as missing because of infrequent price changes. Furthermore, while we can construct price data going back to at least 1957 for many commodities, we restrict the empirical analysis to the period since 1968, in light of the numerous price regulations and government price support mechanisms in place during the earlier period.

There is wide geographic variation in where commodities are produced. This point is illustrated in Appendix D, which presents information on the primary producing countries for each commodity in 1990, the middle of the sample period, as well as information on the common uses of each type of commodity. While some countries are consistently among the major producers of many commodities due to their size (e.g., the USSR, China and India), the geographic variation is nonetheless substantial and reflects the disproportionate influence of some smaller countries on the production of individual commodities. For example, while the former USSR was the primary producer of potatoes and sunflower oil in 1990, Poland was second for potatoes, accounting for 13% of global production, while Argentina was second in sunflower oil, accounting for 17% of global production. Among industrial commodities, Chile is well-known as one of the world's largest producers of copper. But production of other commodities is also quite geographically differentiated. For example, Uzbekistan was the third largest producer of cotton (14% of global production), Bangladesh accounted for 30% of global production of jute/burlap, while Australia and New Zealand were the largest producers of wool, jointly accounting for nearly 50% of world production. This geographic variation in the production of commodities has also been used in other contexts (e.g., Chen, Rogoff and Rossi 2010).

The table also describes some of the uses of each commodity, again primarily as reported by the CRB statistical yearbooks and the UN FAO. It is important to recognize that while we group commodities into three categories (“agriculture”, “oils”, and “industrial”) in the same way as the IMF, the World Bank, and the CRB, these groupings are somewhat arbitrary. While they are based on end-use (e.g., cotton is used primarily in textiles, hence is considered industrial), most commodities are used in a variety of ways that make such a classification problematic. For example, many of the “agricultural/food” commodities also have industrial uses or serve as inputs into the production of refined products that require significant additional value-added: potatoes and grains are used in significant quantities for distillation, pepper and soybeans can be made into oils that have medical, cosmetic, or industrial uses, corn and sugar are increasingly used as fuel, and so on. Similarly, the oils in the sample are well-known for their use in cooking but some (such as palm and coconut oil) also have a number of important industrial uses.

3.2 Common Factors in Commodity Prices

Prior to conducting the factor analysis, we normalize each price series by the U.S. CPI, so that the analysis is in terms of real commodity prices. Second, we take logs of all series. Third, we normalize each series by its standard deviation. Because there are missing observations in the data, we use the expectation-maximization (EM) algorithm of Stock and Watson (2002).³ We follow Kilian (2009) in focusing on the (log) level of real commodity prices, but document in our robustness checks that our results are qualitatively unchanged if we first difference real commodity prices or use linearly detrended series.

We consider several metrics to characterize the contribution of the first five factors in accounting for commodity-price movements, summarized in Table 1.⁴ The first row presents the sum of eigenvalues associated with each number of factors normalized by the sum across all eigenvalues, a simple measure of variance explained by common factors. In addition, we present additional metrics based on R^2 s that explicitly take into account missing values associated with some commodities. For example, the second row presents the average across the individual R^2 s computed for each commodity (excluding commodity-specific imputed values) for the numbers of factors ranging from one to five. The next row presents the median across these same commodity-specific R^2 s, while the following row presents the R^2 constructed across all commodities (again omitting imputed values). Because different commodities have different time samples, the R^2 s are not directly comparable across commodities, but they nonetheless provide a useful metric for evaluating the importance of common factors to the comovement of commodity prices.

The key result from this table is that the first common factor explains a large share of the price variation across commodities, ranging from 60-70% depending on the specific measure used. By contrast, all of the additional factors explain smaller fractions of the variance of commodity prices. The second factor, for example, accounts for between 6% and 10%, while the third factor contributes another 5% of the variance. Thus, the first two factors jointly account for approximately 70-75% of the variance in commodity prices. The

³ Specifically, we first demean each series and replace missing values with zeroes before recovering the first K factors. We use these K factors to impute the value of missing observations, then re-do the factor analysis, iterating on this procedure until convergence. We use $K=5$ factors for the imputation, but the results are not sensitive to the specific number of factors used.

⁴ Following Connor and Korajczyk (1993) and Bai and Ng (2002), we use principal components on the variance-covariance matrix of commodity prices to estimate the approximate factors.

next three factors jointly bring the combined variance up to 85%. Given these contributions to variance, statistical tests of the number of factors point toward parsimonious factor specifications. For example, the PC2 and IC2 criteria of Bai and Ng (2002), each select one factor. The same result obtains using the test suggested by Onatski (2010) or the two criteria proposed in Ahn and Horenstein (2013).⁵

The ability of the first two factors, and the first common factor in particular, to account for so much of the variance holds across commodity groups. Table 1 includes the contribution of different factors to explaining the variance across the three subsets of commodities in the sample: agricultural/food, oils and industrials. Differences across subsets of commodities are quite small: the contribution of the first factor ranges from 55% (pooled R^2 across all commodities in this subset) for industrial commodities to 64% for agricultural commodities and 72% for oils. The differences are largely driven by a few commodities within each grouping for which the first factor accounts for a much smaller share of the historical real price variation than others (Appendix E). Among agricultural commodities, apples, bananas, onions, pepper and shrimp have much smaller R^2 s than most other commodities, likely reflecting the fact that these are the agricultural commodities for which non-industrial uses are least important. Among industrial commodities, nickel and cement are the two commodities for which the first common factor accounts for the smallest share of the variance. But with the exception of these few commodities, the decomposition does not suggest that one needs different factors for different types of commodities. This point is worth stressing because a common concern with factor analysis is that different factors are needed to explain different subsets of the data. For example, Blanchard (2009) notes that the macroeconomics factor literature has yielded a puzzling need for separate factors to explain real, nominal, and financial variables. In our context, one might be concerned that a factor decomposition of real commodity prices across a wide set of commodities may lead to separate factors being needed for industrial and agricultural commodities. As illustrated in Table 1, this is not the case.

3.3 Identification of the Rotation Matrix and the Underlying Economic Factors

To implement a structural interpretation of the factors as suggested by the model, we interpret the results of Table 1 as indicating that a two-factor representation adequately characterizes the data. First, additional factors beyond the first two add relatively little in explanatory power and can be omitted. Second, under the null of the model, it is *a priori* unlikely for there to be fewer than two factors. Indeed, such a finding would imply that there are *no* shocks that directly affect commodity prices and that all commodity-price movements reflect either the level of aggregate economic activity or idiosyncratic commodity factors. We can rule this argument out immediately because there exists at least one common shock to the supply of commodities: exogenous energy price movements. Because commodities require energy in production and distribution, exogenous shocks to energy prices necessarily induce some comovement in commodity prices since commodities are produced in different parts of the world but consumption occurs disproportionately in advanced economies, thereby generating significant shipping and distribution costs. As a result, energy can be interpreted as a common input into the production of commodities in the same spirit as the “land” in the model of section 2.

To assess whether exogenous energy shocks feed through to other commodity prices, we regress each commodity’s real price on lags of itself as well as contemporaneous and lagged values of Kilian’s (2008)

⁵ These information criteria, however, can be sensitive to the sample period. For example, the Onatski (2010) test picks three factors instead of one when we start the sample period just one year earlier, in January 1967 instead of 1968.

measure of exogenous OPEC production shocks that has been updated through January 2013 by Bastianin and Manera (2014).⁶ Following Romer and Romer (2004), we use two years of lags for the autoregressive component and three years of lags for the exogenous variable (OPEC production shocks). From the impulse responses implied by the estimates, we find that we can reject the null hypothesis of no response to an OPEC production shock for 20 (14) commodities at the 10% (5%) level. This evidence suggests that exogenous oil production shocks tend to affect commodity prices and therefore that there is at least one source of direct commodity shocks. Thus, we focus on the two-factor representation of real commodity prices.

To estimate the rotation matrix, our baseline is to impose orthogonality conditions on the indirect common factor F_t^1 . Specifically, we take ε_t^{oprec} the measure of OPEC production shocks from Bastianin and Manera's updated version of the Kilian (2008) series and define the orthogonality conditions as $E[F_t^1 z_t]$ where $z_t \equiv [1 \ \varepsilon_t^{oprec} \ \dots \ \varepsilon_{t-L}^{oprec}]$ is the vector of instruments that consists of a constant, the contemporaneous value of the production shock series as well as L lags of the shock. The IC factor F_t^1 (y_t^{nc} in the model) is a rotation over the two estimated factors \hat{F}_t^1 and \hat{F}_t^2 , i.e., $F_t^1 = t_{11}\hat{F}_t^1 + t_{21}\hat{F}_t^2$ where the orthogonal rotation parameters t_{11} and t_{21} can be expressed as a function of a single underlying rotation parameter θ such that $t_{11} = \cos \theta$ and $t_{21} = \sin \theta$. Given that there are more moment conditions ($L+2$) than parameters (θ), we can estimate the rotation parameter θ using GMM by minimizing $J(\theta)$

$$J(\theta) = \left[\frac{1}{T} \sum_t (F_t^1(\theta) z_t) \right] W \left[\frac{1}{T} \sum_t (F_t^1(\theta) z_t) \right]'. \quad (7)$$

As noted before, many commodity prices respond significantly to exogenous OPEC oil production shocks. Furthermore, the second unrotated factor is significantly impacted by OPEC production shocks, with peak effects obtaining 16 months after the shock and declining gradually thereafter, and we can reject the null that OPEC production shocks have no effect on the unrotated second factor at the 5% level.⁷ Thus, the orthogonality condition of the instrument follows from the theory, and this empirical evidence suggests that the exogenous OPEC production shocks have clearly discernible effects on commodity prices, justifying their use as instruments. We set $L=36$ months for the baseline estimation to capture the fact that the OPEC production shocks have long-lived effects on commodity prices, although the results are robust to both shorter and longer lag specifications as well, as we document below. W is the Newey-West HAC estimate of the inverse of the variance covariance matrix of moment conditions, and we iterate over minimizing $J(\theta)$ then computing the implied weighting matrix until the estimate of θ has converged ($W=I$ in the first step). Table 2 presents the resulting estimate of θ and its associated standard error. With $\hat{\theta} = -0.10$ and a standard error of 0.31, we cannot reject the null hypothesis that $\theta = 0$. From this estimate of θ , we construct estimates of the rotation parameters t_{11} and t_{21} : t_{11} is close to 1, while we cannot reject the null hypothesis that $t_{21} = 0$, so the

⁶ We thank Lutz Kilian and Andrea Bastianin for providing us with the original and updated OPEC production shock data. The correlation between the original Kilian (2008) series and Bastianin and Manera's updated series is 0.99. There are small numerical differences between the two series that are likely due to the fact that Bastianin includes Ecuador in the set of OPEC countries. Ecuador rejoined OPEC in 2007, which is after the first draft of the Kilian paper was written. Kilian's (2008) measure of OPEC production shocks is available on a monthly basis from January 1968 until August 2004, but the first production shock does not occur until November 1973. To extend the series back to January 1968, we set the Kilian production shock series to zero prior to 1973.

⁷ Specifically, we regress the unrotated second factor on a constant, the contemporaneous OPEC production shock, and 24 lags of the OPEC production shock and test the null hypothesis that all coefficients on OPEC production shocks are zero.

estimated rotation matrix is not statistically different from the identity matrix. Furthermore, the over-identification conditions cannot be rejected.

The results are insensitive to many of the specific choices made for the estimation of θ . For example, we report in Table 2 the results from using fewer moment conditions ($L = 12$ and 24 months) as well as more moment conditions ($L = 48$ months). Neither changes the estimates by much. With fewer lags, the standard errors get somewhat larger. This finding reflects the fact that OPEC production shocks have only gradual effects on commodity prices, so that moment conditions at shorter lag lengths are only weakly informative. Similarly, we redid the GMM estimates using a 2-step procedure, in which θ is first estimated using a weighting matrix equal to the identity matrix with no subsequent iterations after updating the weighting matrix as well as continuously updated GMM in which we minimize over θ and W jointly until convergence. In both cases, the results are qualitatively similar. Finally, because non-linear GMM can be sensitive to normalizations, we replicate the baseline estimation after rewriting moment conditions as $E \left[(\hat{F}_t^1 + \hat{F}_t^2 \frac{\sin \theta}{\cos \theta}) z_t \right] = 0$, and the results are again qualitatively unchanged.

The reason why the estimated rotation matrix is close to the identity matrix is that while the first unrotated factor is largely uncorrelated with OPEC production shocks, this condition is not satisfied for the second unrotated factor. Because the unrotated factors are already largely consistent with the theoretically predicted orthogonality conditions (namely, that the first factor is orthogonal to commodity shocks, but the second is not), the estimation procedure yields only a slight rotation of the original factors.

While the fact that we cannot reject the over-identifying conditions is consistent with the theory, we can further assess the extent to which the estimated rotation satisfies the theoretical predictions of the model. For example, an additional theoretical prediction is that the loadings on the indirect factor all be of the same sign. To assess this prediction, we present in Table 3 the estimated factor loadings for each rotated factor. The loadings on the IC factor are positive for all commodities, as predicted by the theory. By contrast, loadings on the commodity-related factor are of mixed signs. There are no systematic patterns across commodity groups which again confirms that the factors explaining commodity prices are common across commodity subsets. Without imposing any restrictions on the loadings as part of the identification strategy for the rotation matrix, we find that the estimated rotation satisfies theoretical predictions on the factor loadings and those implied by the over-identifying restrictions.

Given the estimate of θ and the rotation matrix, we construct the rotated factor F_t^1 that, according to the model, corresponds to the level of aggregate output and income that would have occurred in the absence of commodity-related shocks. This factor is presented in Figure 2 after HP-filtering with $\lambda=129,600$, the typical value for monthly data, to highlight variation at business cycle frequencies. In addition, we draw from the estimated distribution of θ , construct F_t^1 for each new draw, and use this distribution to characterize the 99% confidence interval of the HP-filtered factor.

This factor displays a sharp rise in 1973-1974 before falling sharply during the 1974-1975 U.S. recession. It is followed by a progressive increase over the course of the mid to late 1970s, peaking in 1979 before falling sharply during each of the “twin” recessions of 1980-1982, and then rebounding sharply after the end of the Volcker disinflation. Thus, over the course of the 1970s, this structural factor displays a clear cyclical pattern. During the mid-1980s, the factor drops sharply before rebounding in the late 1980s, then falls

gradually through the 1990 U.S. recession before rebounding through the mid-1990s. It experiences a large decline in the late 1990s, prior to the 2000-2001 U.S. recession and then rebounds shortly thereafter. After a brief decline in the mid-2000s, the factor displays a sharp increase from 2005 to 2008, the period during which many commodity prices boomed, then falls sharply in late 2008 and 2009 before rebounding strongly in 2010. In short, there is a clear procyclical pattern to the IC factor relative to U.S. economic conditions, a point to which we return in greater detail in section 3.4.

To assess the sensitivity of our results, we consider an alternative identification strategy suggested in section 2.4, namely to exploit the theoretical predictions for signs of factor loadings: loadings on the IC factor should all be positive. Thus, one can characterize the set of admissible rotation matrices by restricting them to be consistent with the sign restrictions implied by the theory, in the spirit of Uhlig (2002). In our case, this procedure consists of identifying the set of θ such that $\min(\tilde{L}^1 \cos \theta + \tilde{L}^2 \sin \theta) > 0$, where \tilde{L}^i for $i = \{1, 2\}$ are the loading vectors associated with the unrotated factors and \min is with respect to the elements of L^1 . We consider values of $\theta \in [-\pi, \pi]$ (at increments of 0.001) and for each θ determine whether the restriction is satisfied. This yields a set of admissible rotation matrices and thereby a set of possible IC factors. We HP-filter each of these and plot the resulting minimum and maximum values for each month in Panel B of Figure 2, along with the 99% confidence interval for the rotated IC factor from the baseline GMM estimation. There is significant overlap between the two approaches, with the minimum and maximum values from the sign restriction typically being within the 99% confidence interval of the GMM-estimated IC factor. Thus, despite the fact that the two identification strategies are quite different, they point toward a remarkably consistent characterization of the non-commodity-related structural factor.

We verify that our results are not unduly sensitive to specific commodities or groups of commodities within the cross-section (Appendix F). For example, the sample includes five closely-related grains (barley, hay, oats, sorghum, wheat), which out of a cross-section of forty commodities could lead to the appearance of more general comovement if these specific commodities were affected by a common shock. Keeping only wheat out of the grains makes little difference for our results as does keeping only palm oil out of the five oils. One might also be concerned about too much overlap in how some commodities are used. We replicate our results dropping either all commodities whose primary (60% or more in Table 1) use is as food or as feed. Another concern is that while there is significant geographic variation among the primary producing countries of different commodities, it is still the case that the former U.S.S.R., China and India accounted for a large proportion of the production of many commodities during our sample period. But we find that dropping all commodities for which the primary producing country in 1990 was the former U.S.S.R (8 commodities) increases the uncertainty around the estimated IC factor, primarily in the 1970s, but leaves the estimate of θ and the underlying unrotated factors unchanged. Dropping all of the commodities for which either China or India were the primary producers in 1990 (thirteen commodities) yields results almost identical to our baseline. Thus, we conclude that the baseline estimation of the common factors in commodity prices is robust to the choice of commodities included in the cross-section.

We also assess whether the results are sensitive to statistical considerations. For example, our baseline approach uses the (log) level of real commodity prices. While there is little visual evidence of commodity prices exhibiting pronounced trends over this period (see Appendix G), we want to ensure that the results are not driven by spurious correlations from trends. We replicate our analysis after linearly detrending each series

prior to extracting factors and also consider an alternative specification in which real commodity prices are first-differenced in the factor analysis. Neither yields qualitatively different results, thereby ensuring that our results are robust to alternative assumptions about stationarity of commodity prices.

We consider two final checks on the results. First, we drop all commodities for which some significant imputations had to be done (e.g., commodities with more than a few missing observations at the end of the sample), or 7 commodities in total. Omitting these series again has almost no effect on the results. Thus, our findings are insensitive to the imputation of commodity prices. Second, we implement the initial factor analysis by decomposing the correlation matrix of commodity prices rather than the covariance matrix, again finding little difference relative to the baseline. In short, the estimates of the IC factor are quite robust to commodity selection issues, treatment of trends in the data, the imputation of commodity prices, and the identification procedure used to recover the rotation matrix.

The robustness of the results reflects two features of the data. First, the initial factor decomposition, and particularly the first unrotated common factor, is largely insensitive to the specific set of commodities used or econometric details such as the treatment of trends or the specific method used to decompose the data. This reflects the fact that there is widespread and persistent comovement in real commodity prices, most of which is captured by a single factor. Second, this first unrotated factor already satisfies the theoretical restrictions implied by the theory: the factor is largely orthogonal to exogenous OPEC production shocks and its factor loadings are all of the same sign. Thus, when imposing these theoretical restrictions implied by the model to identify the rotation matrix, we cannot reject the null hypothesis that the rotation matrix is equal to the identity matrix. Almost all subsequent sensitivity found in robustness checks reflects variation in the standard errors of the GMM estimate of the rotation parameter, not variation in the underlying factor decomposition or the point estimate of the rotation matrix.

3.4 The Contributions of the Factors to Commodity Prices, Comovement and Global Real Activity

The model presented in section 2 suggests that one of the common factors driving real commodity prices can be interpreted as the level of global economic activity that would have prevailed in the absence of commodity-related shocks. Furthermore, the theory provides guidance on how one can identify this factor from the data, and the previous sections have shown how to implement this identification procedure. In this section, we construct historical decompositions of commodity-price movements and global economic activity following the structural interpretation suggested by the theory.

3.4.1 Sources of Average Commodity Price Movements

For prices, we decompose the average annual percentage change in commodity prices into the components driven by indirect and direct common factors. The decomposition follows directly from the rotated factor structure, yielding

$$\overline{p_t - p_{t-12}} = \overline{L^{IC}}(F_t^{IC} - F_{t-12}^{IC}) + \overline{L^{DC}}(F_t^{DC} - F_{t-12}^{DC}) + (\overline{\varepsilon_t - \varepsilon_{t-12}})$$

where the bar denotes averages across all commodities in the cross-section. The first term on the right-hand side represents the contribution of the IC factor to average commodity-price movements, the second represents the contribution of the DC factor, and the third reflects average idiosyncratic effects. We focus on annual changes in prices to abstract from higher frequency commodity-price changes.

The results of this decomposition are presented in the top panel of Figure 3, in which we plot the contributions from the IC and DC factors each month as well as the actual annual average price change across commodities. The IC factor explains the vast majority of historical commodity-price changes. Thus, historical changes in commodity prices have primarily reflected endogenous responses to non-commodity shocks. To the extent that income effects on inputs into the production of commodities are likely weak, the IC factor could be interpreted as primarily reflecting changing demand for commodities related to changes in global economic activity. During the commodity boom of 1973-74, for example, indirect shocks to commodity markets accounted for almost all of the rise in commodity prices, with the remainder reflecting direct commodity-related shocks. This pattern appears to be the historical norm: every major historical episode of large average commodity price movements is accounted for by the indirect factor, i.e. as an endogenous response of commodity prices to global business cycle conditions not driven by commodity-related shocks.

3.4.2 Sources of Commodity Price Comovement

We can also quantify how changes in each factor have contributed to the time-variation in *comovement* among commodity prices. Specifically, we can decompose, each month, annual changes in real commodity prices as:

$$p_t(j) - p_{t-12}(j) = \lambda_j^{IC} [F_t^{IC} - F_{t-12}^{IC}] + \lambda_j^{DC} [F_t^{DC} - F_{t-12}^{DC}] + \varepsilon_t^j - \varepsilon_{t-12}^j.$$

From this, we can construct the R^2 each month coming from both factors (i.e., the ability of changes in both factors to explain commodity price movements through common forces) as well as the partial R^2 coming from the IC factor. These series are plotted in Panel B of Figure 3. There is significant variation over time in the overall comovement of commodity prices, as captured by both factors, with the highest degrees of comovement in commodity price changes occurring between 1973 and 1975, in the early to mid-1980s, in the late 1990s, and in the mid to late 2000s continuing to 2013. Most of the time-variation in comovement can again be explained by changes in the indirect factor, implying that periods in which commodity prices comove most strongly have also been periods in which commodity price changes have been driven by the endogenous response of commodity prices to non-commodity shocks.

3.4.3 Sources of Global Business Cycle Fluctuations

We next assess the contribution of each factor to global economic activity. To do so, we rely on a measure of global industrial production constructed by Baumeister and Peersman (2011), who collected the industrial production data in the United Nations' *Monthly Bulletin of Statistics* from 1947Q1 until 2008Q3 and aggregated individual country industrial production measures into a global measure of industrial production. The series was extended from 2008Q3 until 2010Q4 using only advanced economy industrial production.

Unlike with commodity prices, the factor structure does not immediately lend itself to a decomposition of historical changes in global industrial production. To do so, we first rely on the theory of section 2 in which the IC factor corresponds to the level of global activity that would have occurred in the absence of direct commodity shocks (y_t^{nc}). Thus, changes in the IC factor can be directly interpreted as changes in aggregate output driven by indirect shocks. Because the scale of the IC factor is not identified, we normalize it such that the standard deviation of quarterly changes in the IC is equal to the standard deviation of quarterly percent changes in global IP and treat the resulting historical changes in the IC as the contribution of indirect shocks to global IP. The difference between the demeaned quarterly growth rate of global IP and the demeaned change in the IC (which we define as δ_t , where $\delta_t \equiv \Delta y_t - \Delta y_t^{nc}$) should reflect the contribution of direct commodity

shocks, potentially omitted factors, as well as mismeasurement in global production levels. To evaluate the contribution of direct commodity shocks to global IP, we estimate

$$\delta_t = c + \sum_{j=1}^4 \beta_j \delta_{t-j} + \sum_{j=1}^8 \gamma_j F_{t-j}^{DC} + \varepsilon_t.$$

such that the direct factor can have dynamic effects on global IP. Unlike the IC factor, the DC factor not only reflects the contribution of direct commodity shocks to aggregate production but also the effects of such shocks on commodity markets through direct shifts in supply or demand. Such shifts have effects above and beyond the general-equilibrium effects of the direct commodity shocks on aggregate output. Estimated at a quarterly frequency, we allow for one year of autoregressive lags and two years of lags of the DC factor to capture potentially dynamic effects of commodity-related shocks on global IP. From this specification, we construct the contribution of the DC factor to global IP net of the contribution of the IC factor. Note that this approach leaves a component of global activity unaccounted for. This component can be interpreted as reflecting measurement error, omitted variables or model misspecification.

We plot the resulting contributions of the IC and DC factors to global IP growth in Panel C of Figure 3, again showing only the annual changes to filter out the high-frequency variation in the measurement of global IP. The correlation between changes in the IC factor and annual changes in global IP is high (0.59) so that historical changes in global IP are primarily attributed to indirect non-commodity shocks. This is particularly true from the early 1970s through the mid-1980s, although commodity-related shocks deepened the decline in global IP during late 1974 and early 1975. As was the case with the decomposition of commodity prices, the decline in economic activity during the Volcker disinflation is accounted for by the IC factor. The dynamics of global activity from the late 1980s to mid-1990s are also largely attributed to the IC factor, although actual changes in global IP exceeded those predicted by the two factors. Growth in the IC factor during the 2000s also coincides with the growth in global IP during this time period, while commodity-shocks in the DC contributed modest downward pressure on economic activity in 2002 and 2003, then again in 2007-2010. To the extent that the DC factor reflects exogenous energy price fluctuations, the negative contribution of the DC factor from late 2007 through 2010 (subtracting 1-2% from the annual growth rate of global IP) is broadly consistent with Hamilton (2009), who argues that oil-price shocks contributed to the severity of the Great Recession. Nonetheless, the decomposition suggests that most of the decline in the growth rate of global IP from late 2007 to the depth of the recession can be attributed to declines in the growth rate of the IC factor.

4 Storage

The model in section 2 yields a factor structure of commodity prices whose properties conform closely to the data and permit us to make causal inferences about the relationship between global real activity and commodity-related shocks. The key to the identification in the factor structure is that all indirect shocks to commodity markets (i.e., all shocks that affect commodity prices through the general equilibrium response of output) are aggregated into a single factor, the IC factor. This conclusion relies on the premise that indirect shocks all induce identical comovement of commodity prices.

This aggregation property of the factor structure can be broken in the presence of storage. To see why, suppose that we extend the model to include a perfectly competitive storage sector for each primary

commodity j that purchases or sells that commodity on the spot market, leading it to hold inventories in the steady-state. As illustrated in Deaton and Laroque (1992), the key determinant of whether the storage sector increases or decreases its inventories is the expected path of prices of the commodity. If a current increase in prices is not expected to persist, then the storage sector sells a positive amount of its inventories on the spot market today when prices are high and rebuilds inventories in future periods when prices are lower. This behavior increases the contemporaneous supply of the good and reduces it in the future. By contrast, if the shock is expected to generate a persistent increase in prices, the storage sector does not have an incentive to change its stock of inventories and therefore is not a net purchaser of the good. Thus, the persistence of the driving process affects the size of net purchases by the storage sector through its effect on the path of expected prices. For example, if aggregate productivity shocks in the model were highly persistent while labor supply shocks were less persistent, the presence of storage would lead these shocks to have different supply responses depending on the size of the storage sector's net purchases. The comovement in commodity prices would then not necessarily be the same across the two shocks, potentially breaking the aggregation result.

In practice, there are three reasons to think that this issue is unlikely to be quantitatively important. First, if the aggregation of indirect shocks into a single IC factor were broken, we would expect a factor decomposition of commodity prices to indicate that many factors were required to explain the comovement of commodity prices, since there likely are a number of different aggregate structural shocks affecting commodity prices through the indirect channel of global activity, such as financial shocks, markup shocks, and fiscal shocks, in addition to the productivity and labor supply shocks that we explicitly model. But as documented in Table 1, the comovement of commodity prices is well-characterized by two factors, with any additional factors adding little explanatory power. This finding suggests either that different indirect shocks have common effects on expected price paths of commodity prices (such that the response of storage is similar across all indirect shocks and, therefore, that the aggregation of indirect shocks still holds) or that the effects of net purchases for the storage motive are second-order in affecting commodity prices.

The second reason why storage is unlikely to be important is precisely because the effects of net purchases for storage motives appear to be second-order for most commodities. To examine this claim, suppose again that we integrated a storage sector for each primary commodity into the model, in which firms purchase or sell the commodity on the spot market as well as store it. The storage sector would therefore affect spot markets through its forward-looking net purchases, defined as $NP_t(j)$ at time t for commodity j . The market clearing condition in the presence of an additional storage sector would then be given by $Q_t(j) = Y_t(j) + NP_t(j)$ such that high (low) net purchases by the storage sector to accumulate (draw down) inventories would increase (decrease) the demand for commodity j at time t holding all else constant. Allowing for trend growth in production such that Y/Q and NP/Q are stationary along the balanced growth path, the log-linearized version of this equation is

$$(Y/Q - 1)np_t(j) = (Y/Q)y_t(j)$$

where the terms in parentheses are BGP ratios. For the storage sector to have first-order effects on equilibrium outcomes (including prices), it must be the case that net purchases are different from zero on average, or equivalently that the ratio of consumption to production (Y/Q) of the commodity is different from one.

Table 4 presents estimates of the mean annual ratio of consumption to production (minus 1) for commodities for which such data could be collected: $r_t = Y_t/Q_t - 1$.⁸ Out of thirty two commodities, we reject the null that $r_t = 0$ on average for only nine: apples, bananas, onions, potatoes, rice, sugar, tea, palm oil, and safflower oil. Note that four of these are highly perishable commodities (apples, bananas, onions, and potatoes), so one would expect some fraction of the goods to spoil while being transported from production to retail facilities. But even in the case of these highly perishable goods, the implied gaps between consumption and production are small -- less than 1% per year. Furthermore, in the case of potatoes the rejection of the null has the wrong sign (i.e., consumption is larger than production on average). Among the less perishable agricultural commodities (e.g., grains), there is little evidence that consumption is significantly less than production on average, with most of the point estimates being less than 1%. This conclusion also applies to industrial commodities, which are highly storable and for which one would expect inventory motives to be potentially important. In fact, there is little evidence of non-zero net purchases by the storage sector. Thus, with the exception of a few commodities it is difficult to reject the null that speculative motives through storage have only second-order effects on prices.⁹ Furthermore, the failure to reject the null does not typically reflect large standard errors. Rather, the point estimates of the net ratio are typically smaller than 1%, which suggests that net flows to the storage sector are small on average. Finally, if we replicate our baseline factor analysis using only the commodities for which we cannot reject the null of zero net-purchases on average, it has little effect on the estimated IC factor (Appendix H).

A third way to assess the possibility that the effects of storage could break the aggregation of indirect commodities into a common IC factor is to note that, in the presence of storage motives, interest rates would play an important role in affecting commodity prices (Deaton and Laroque 1992; and Frankel 2008). As a result, the logic of the model in section 2 would imply that monetary policy shocks would *directly* affect commodity prices through changes in desired inventories. Therefore, in a factor decomposition these monetary policy shocks would not be incorporated into the *indirect* factor. Hence, a testable implication of a quantitatively important storage motive is that monetary policy shocks should not affect the IC factor.

To test this prediction, we identify U.S. monetary policy shocks using a time-varying-coefficients (TVC) Taylor rule

$$i_t = c_t + \varphi_t^\pi F_t \pi_{t+1,t+2} + \varphi_t^{gy} F_t g y_t + \varphi_t^x F_t x_t + \rho_t i_{t-1} + \varepsilon_t^{mp} \quad (8)$$

in which the central bank responds to real-time forecasts (F_t) of average inflation over the next two quarters ($\pi_{t+1,t+2}$), the current quarter's output growth ($g y_t$), the current quarter's output gap (x_t), and the previous period's interest rate as in Kozicki and Tinsley (2009) and Coibion and Gorodnichenko (2011). We assume

⁸ We use measures of consumption and production of commodities from the CRB. When these are not available, we rely on measures from the UN FAO for agricultural and oil commodities, from the U.S. Department of Agriculture's Food and Agricultural Services (USDA FAS), and from trade associations. Aluminum data were provided to us by the European Aluminum Association (EAA), data for copper is from the International Copper Study Group (ICSG), data for tin were provided by the International Tin Research Institute (ITRI), nickel data are from International Nickel Study Group (INSG), while data for zinc and lead were tabulated from the International Lead and Zinc Study Group's *Monthly Bulletin*. For many commodities, we were able to construct global production and consumption data going back to 1968. There are only eight commodities for which we could not compile consumption and production data: beef, hay, orange juice, shrimp, cement, lumber, mercury and wool.

⁹ This evidence is also consistent with the well-documented inconsistencies between the standard storage model and the observed data (see, among others, Ng 1996).

that each of the TVCs follows a random walk, including the intercept that captures changes in the central bank's target levels of macroeconomic variables and the natural rate of interest. Following Orphanides (2003) and Romer and Romer (2004), we use the Greenbook forecasts prepared by the staff of the Federal Reserve prior to each FOMC meeting to characterize the FOMC's real-time beliefs about current and future macroeconomic conditions. The TVCs allow us to distinguish between systematic changes in the monetary policy rule from transitory deviations captured by the residuals. We estimate this rule using data at the frequency of FOMC meetings from March 1969 until December 2008. Because Greenbook data are not available after 2007, we use Blue Chip Economic Indicator forecasts. The sample ends in December 2008 when the zero-bound on interest rates was reached. We then define the residuals from estimated equation (8) as monetary policy shocks and construct a monthly series from the FOMC-frequency dated series of shocks.

To quantify the effects of monetary policy shocks on the indirect common factor, we use a vector autoregressive representation of macroeconomic dynamics with four variables: our measure of monetary policy shocks, the log of U.S. industrial production, the log of the U.S. Consumer Price Index (CPI), and the IC factor. We order the monetary policy shock first given that it should already incorporate the most recent economic information obtained from the Greenbook forecasts and to allow other variables to respond on impact to this shock. We use data from 1969:3 until 2008:12 to estimate the VAR with 18 months of lags, midway between the 12 month lag specifications typical of monetary VAR's and the 24 month lag specification used by Romer and Romer (2004). We then plot in Figure 4 the impulse responses of industrial production, the CPI, and the IC factor to a monetary policy innovation.

An expansionary monetary policy shock in the VAR leads to higher industrial production, with peak effects happening one to two years after the shock. The CPI rises moderately but persistently around 6 months after the shock, consistent with the delayed effect on prices of monetary policy shocks long observed in the empirical monetary policy literature (e.g., Christiano et al. 1999). The indirect factor rises much more rapidly, within the first 3 months, but does not peak until nearly two years after the shock before gradually declining back toward zero. The responses are significantly different from zero at the 5% level for the first twenty months and briefly at the 1% level.¹⁰ Thus, we can statistically reject the null hypothesis that monetary policy shocks have no effect on the IC factor. In addition, the quantitative contribution of U.S. monetary policy

¹⁰ Note that the reported standard errors do not account for the fact that the IC factor is a generated regressor, and they therefore may understate the true uncertainty around the point estimates. However, there are at least two reasons to suspect that this is not quantitatively important. First, one could also test the null that monetary policy shocks have no effect on the IC factor by regressing it on current and lagged monetary shocks, i.e., $F_t^{IC} = c + \sum_{i=0}^l \beta_i \varepsilon_{t-i}^{mp} + v_t$, setting $l=36$ months to account for the gradual effects of monetary policy shocks on macroeconomic variables. From this procedure, we can reject the null hypothesis that monetary policy shocks have no effect on the IC factor (i.e., $\hat{\beta}_i = 0 \forall i$) with a p -value of 0.019. The generated regressor issue is not binding in this case since the IC factor is only on the left-hand side and the null hypothesis is that the coefficients on monetary policy shocks are zero, so asymptotic (Newey-West) standard errors are valid (Pagan 1984). The advantage of the VAR specification is that it also purges the monetary policy shocks of potentially remaining predictability from macroeconomic variables and is in this respect a more conservative approach. Second, given that we cannot reject the null of the rotation matrix being equal to the identity matrix, one can use the unrotated first common factor in the VAR in lieu of the rotated one. Since the unrotated factor can be treated as observable following Bai and Ng (2002) and Stock and Watson (2002) for large enough cross-sections and time samples, the corresponding standard errors are valid. The results from this alternative specification are almost identical, and we can reject the null of no response at the same confidence level.

shocks to the indirect factor is relatively large, accounting for much of the sustained increase in the IC factor from late 1975 until 1980 and around two-thirds of the subsequent decline from 1980 to 1982.

In short, while the presence of commodity storage could potentially break the aggregation of indirect shocks into a common IC factor, there is little quantitative evidence in favor of this claim.¹¹ First, the fact that the comovement in commodity prices is well-characterized by a small number of factors is difficult to reconcile with the aggregation result failing to hold. Second, for most commodities we cannot reject the null that storage has only second-order effects on commodity prices. And third, monetary policy shocks have both statistically and economically significant effects on the IC factor, which suggests that the factor decomposition is not treating them as a direct commodity-related shock as would be the case if speculative considerations were economically important. While storage motives are nonetheless likely to play a role in commodity prices in periods when inventory constraints are close to binding, the results suggest that, on average, the aggregation result from section 2 provides a succinct and adequate characterization of the data.

5 Forecasting Applications

The model presented in Section 2 predicts that the level of real commodity prices and total demand for commodities are endogenous and jointly determined. Furthermore, the empirical evidence in section 3 documented that a large proportion of commodity-price movements are systematically related to one another and can be interpreted as reflecting aggregate shocks that are not specific to the commodity sector. Guided by this insight, we examine whether the common factor identified from the cross-section of commodity prices contains information relevant for predicting real commodity prices in a recursive out-of-sample forecasting exercise. In the out-of-sample forecasting exercise, we examine the ability of the common commodity factor to forecast not just the set of commodities in the data set but also commonly used commodity indices and the real price of oil.

5.1 Forecasting Model

The forecasting model is a linear bivariate FAVAR(p) model for the real price of commodity j and the IC factor:

$$x_{t+1} = A(L)x_t + e_{t+1} \quad (9)$$

where $x_t = [rpc_{jt}, IC_t]'$, rpc_{jt} denotes the log of real price of commodity j , IC_t is the IC factor extracted from the cross-section of real commodity prices, e_{t+1} is the regression error, and $A(L) = A_1 + A_2L + A_3L^2 + \dots + A_pL^{p-1}$. In the forecasting exercise, the lag length p is chosen recursively using the Bayesian information criterion (BIC).

¹¹ Another reason why one might be skeptical of the quantitative importance of the storage mechanism is recent work examining the role of speculative shocks in oil markets has found little evidence that these have contributed in economically significant ways to historical oil-price fluctuations, either in statistical VAR models such as Kilian and Murphy (2013) and Kilian and Lee (2013) or DSGE models such as Unalmis et al. (2012). While little evidence exists on this question for other commodities, one would expect that oil markets would be most likely to display sensitivity to speculation given the relative ease with which oil can be stored (both underground and in above-ground storage facilities) and the potentially large convenience yields to refineries associated with holding oil as inventories. The fact that storage shocks are not quantitatively important of course does not imply that storage has no effects on the response of prices to other shocks, but it is consistent with this result.

All of the nominal commodity prices are deflated by U.S. CPI. In addition to the cross-section of 40 commodity prices used to compute the IC factor, we examine the ability of the IC factor to forecast three widely used commodity price indices – the CRB spot index, the World Bank non-energy index, and the International Monetary Fund non-fuel index.¹² The indices are also deflated by U.S. CPI. The real price of oil used in the forecasting exercise is the U.S. refiner’s acquisition cost of imported oil, which is a good proxy for the international price of crude oil (see Alquist et al. 2013). We apply the EM algorithm recursively to fill in the missing observations and estimate the common factor at each point in time. We appeal to the fact that, in section 3, we are unable to reject the null that the rotation matrix equals the identity matrix and therefore use the unrotated first factor in the forecasting exercises. The rationale is the well-known sensitivity of GMM in short-samples and the related concern that small-sample considerations may induce significant variation in the estimate of the rotation matrix across periods.

The forecast performance of the FAVAR is evaluated over two periods. In the first case, the forecast evaluation period depends on the commodity. It begins either in 1968:1 or at the earliest date subject to the condition that the initial estimation window contains at least 48 observations (see Appendix Table I.1). The second forecast evaluation period begins in 1984:1 and ends in 2012:12, with the initial estimation window ending in 1983:12. We again impose the condition that the initial estimation period contains at least 48 observations. These constraints reduce the total number of commodities that we can consider in the common forecast evaluation period from 40 to 28. We evaluate the recursive MSPE of the FAVAR-based forecast at the 1-, 3-, 6-, and 12-month horizons. All forecast accuracy comparisons are conducted relative to the no-change benchmark. Multistep-ahead forecasts are computed iteratively using the FAVAR.

5.2 Forecasting Results

Table 5 summarizes the results obtained from the forecasting exercise for the commodity-specific and common sample periods. The first column of Table 5 shows the aggregate MSPE ratio, which is defined as:

$$\text{Aggregate MSPE Ratio} \equiv \frac{\sum_{j=1}^N MSPE_j^{FAVAR}}{\sum_{j=1}^N MSPE_j^{RW}}$$

where $MSPE_j^{VAR}$ is the mean-squared prediction error of the FAVAR-based forecast for commodity j ; $MSPE_j^{RW}$ is the mean-squared prediction error of the random walk forecast for commodity j . Thus, the aggregate MSPE ratio summarizes the performance of the all of the forecasting models for a given horizon. For both the commodity-specific and the common forecast evaluation periods, common-factor based forecasts generate improvements in forecast accuracy relative to the no-change forecast up to the 6-month horizon. In the commodity-specific period, the improvements range between about 6-8%. In the common forecast evaluation period, the improvements are smaller and lie in the 2% to 7% range at horizons up to 6 months. But these summary statistics mask the heterogeneity in the ability of the FAVAR to produce more accurate forecasts than the no-change forecast. Table 5 also reports the distribution of the MSPE ratios for each forecast evaluation period. In the commodity-specific period, there are 32 (out of 40) commodities at the 1-month

¹² The IMF non-fuel commodity price index available from Haver Analytics begins in 1980:2. The price index was backcast to 1957:1 using the IMF agricultural raw, beverage, food, and metals sub-indices using the weights obtained from regressing the non-fuel index on the individual sub-indices. Over the sample period during which the indices overlap, a regression of the non-fuel index on the sub-indices yields an R^2 in excess of 0.99999.

horizon and 20 (out of 40) commodities at the 3-month horizon for which the FAVAR-based forecasts are more accurate than the no-change forecast. The performance of the FAVAR deteriorates as the forecast horizon lengthens. Similar results obtain in the common forecast evaluation period. There are 22 (out of 28) commodities at the 1-month horizon and 18 (out of 28) commodities at the 3-month horizon for which the VAR-based forecasts are more accurate than the no-change forecast. In addition, at the 6- and 12-month horizons the FAVAR generates superior forecasts relative to the no-change forecast for about half of the commodities in the sample.

For the commodity-specific sample period, the common factor-based forecasts of the real commodity price indices achieve improvements in forecast accuracy relative to the no-change forecast at the 1-month horizon. The FAVAR does best at predicting the World Bank non-energy index and the IMF non-fuel index, with forecast accuracy improvements in the 11-13% range. The accuracy of the FAVAR-based forecast diminishes at the 3-month horizon, with a maximum improvement in forecast accuracy of about 1% for the IMF non-fuel index. Over the common forecast evaluation period, the FAVAR does somewhat better at forecasting the price indices compared to the no-change forecast. Again, the largest improvements in forecast accuracy are obtained for the World Bank and IMF commodity-price indices, with improvements of at most 14% relative to the no-change forecast. At the 3-month horizon, the FAVAR is more accurate than the no-change forecast, but the improvements are smaller (i.e., at most about 7%). The FAVAR model also does well at forecasting the real price of oil at short horizons. For both forecast evaluation periods, it produces improvements in forecast accuracy of about 20% at the 1-month horizon. The 3-month ahead forecasts are 3-6% more accurate than the random walk forecast. The forecasts based on the FAVAR become less accurate as the forecast horizon lengthens.

Appendix Tables I.1 and I.2 report the forecast accuracy results for the individual commodities for the commodity-specific and common sample periods. First, the FAVAR-based forecasts generate improvements in forecast accuracy for some agricultural commodities and oils up to 12 months ahead. For example, 12 (out of 15) agricultural commodities and 2 (out of 3) oils achieve improvements in forecast accuracy at the 12-month horizon. For the agricultural commodities, the improvements in forecast accuracy relative to the random walk forecast range between about 4% for cocoa to 41% for hay. For oils, the gains are about 32% for groundnut oil and about 4% for palm oil at the 12 month horizon. Second, the improvements in forecast accuracy in the industrial commodities are concentrated at the 1- and 3-months horizons. Appendix Table I.2 shows that the improvements in forecast accuracy range between about 22% for cotton to less than 1% for lead at the 1-month horizon; and between about 11% for tin to around 1% for aluminum at the 3-month horizon.

Additional results on the ability of the commodity-price factor to forecast the real price of oil are reported in Appendix Table I.3. The table compares the bivariate FAVAR with a standard VAR model of the global oil market that has been shown to perform well at forecasting the real price of oil out-of-sample (Baumeister and Kilian 2012; and Alquist et al. 2013).¹³ Due to constraints on the availability of oil-market data, the start date for the exercise is January 1973. The IC factor based model does well relative to the oil-

¹³ We thank Christiane Baumeister for sharing the real-time data set for the oil-market model. The variables in the oil-market VAR include the percent change in global crude oil production, the global real activity index constructed in Kilian (2009), the log of the real price of oil, and a proxy for the change in global above-ground crude oil inventories. For further discussion of these data, see Kilian and Murphy (2013).

market VAR model at the 1- and 3-month horizons when the BIC is used.¹⁴ On the other hand, the IC factor-based model is dominated by the oil-market model when a fixed lag length of 12 is used, although the IC factor model still delivers improvements in forecast accuracy up to about 14% relative to the no-change forecast.¹⁵ This evidence suggests that the IC factor contains information relevant for forecasting the real price of crude oil at short horizons. It also underscores the similarities between the economic models underlying the two forecasting models and, in particular, the important role that demand plays in forecasting not only the real price of oil but also the real prices of other agricultural and industrial commodities.

Taken together, these findings indicate that the prices of internationally traded commodities are, to some extent, forecastable in a way suggested by the model presented in section 2. The improvements in forecast accuracy can be substantial, particularly at short horizons, and agricultural commodities and oils tend to be more predictable than industrial commodities. These results show that a FAVAR can be used to generate accurate forecasts of real commodity prices relative to the no-change benchmark. Thus, the factor structure in commodity prices can serve a dual purpose for policymakers and practitioners – providing a structural decomposition of the forces driving commodity prices while also helping to forecast commodity-price movements within a common framework.

6 Conclusion

In this paper, we propose a new empirical strategy, grounded in a micro-founded business cycle model with commodities, to identify the driving forces of global economic activity and commodity prices. First, the model predicts the existence of a factor structure for commodity prices that has a direct economic interpretation. The first component of the factor structure captures idiosyncratic price movements, the second one captures global economic forces, and the third one is related to commodity-specific shocks. The IC factor is of particular interest because it represents a precise counterfactual: the level of global economic activity that would have prevailed in the absence of any contemporaneous commodity-related shocks. Thus, the factor structure of commodity prices predicted by theory suggests a way that the IC factor can help to resolve the identification problem associated with the joint determination of global economic activity and commodity prices. We also show how the model's predictions can be used to identify the rotation matrix that recovers the underlying economic factors implied by the theory, including the IC factor, from a standard empirical factor decomposition of commodity prices. This point addresses the central problem of factor analysis – namely, that it is problematic to assign the factors an economic interpretation. However, the theory provides a set of orthogonality conditions and sign restrictions that can each be used to identify the parameters of the rotation matrix consistent with a structural interpretation of the factors.

Applying these methods to a broad cross-section of commodity prices, the IC factor that we identify accounts for about 60-70% of the variance in commodity prices, and this finding is not sensitive to using two alternative identification strategies. In addition, we are unable to reject the theoretical restrictions implied by the model. The IC factor is highly correlated with independently computed measures of global economic

¹⁴ During the 1984:1-2012:8 forecast evaluation period, for example, the IC factor based model achieves an improvement in forecast accuracy of about 21% relative to the no-change forecast whereas the oil-market fundamental model's improvement in forecast accuracy is about 17% at the 1-month horizon.

¹⁵ The oil-market fundamentals model generates forecast-accuracy improvements up to about 16% compared to the no-change forecast.

activity at business cycle frequencies. Its behavior during the 1970s and 1980s suggest that the macroeconomic fluctuations observed during that era were not driven primarily by commodity-related shocks. Nevertheless, there are episodes during which the direct commodity shocks contributed negatively to global economic activity, particularly in the early 1990s and again during the Great Recession.

Finally, we show that the IC factor is useful for forecasting real commodity prices, some widely used commodity-price indices, and the real price of crude oil. A recursive out-of-sample forecasting exercise shows that a simple bivariate FAVAR that includes the IC factor and the real commodity price can generate economically large improvements in forecast accuracy relative to a no-change benchmark. Because our identification strategy relies only on commodity prices, it can be implemented in real time. Hence, our approach provides a unified framework to forecast a wide range of commodity prices in real time and to assign them a structural interpretation.

In sum, we provide a new conceptual framework for identifying the sources and implications of commodity-price comovement and its relationship to global macroeconomic conditions. The framework suggests a way of interpreting the common factors driving commodity prices and offers a fresh perspective on the historical behavior of a broad cross-section of internationally traded commodities since the early 1970s.

References

- Ahn, Seung C. and Alex R. Horenstein, 2013. "Eigenvalue Ratio Test for the Number of Factors," forthcoming: *Econometrica*.
- Ai, Chungrong, Arjun Chatrath, and Frank Song, 2006. "On the Comovement of Commodity Prices," *American Journal of Agricultural Economics*, 88(3), 574-588.
- Alquist, Ron, Lutz Kilian, and Robert Vigfusson, 2013. "Forecasting the Price of Oil," forthcoming in: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland.
- Aruoba, S. Boragan, Francis X. Diebold, M. Ayhan Kose, Marco E. Terrones, 2011. "Globalization, the Business Cycle, and Macroeconomic Monitoring," IMF Working Paper 11/25.
- Bai, Jushan and Serena Ng, 2002. "Determining the Number of Factors in Approximate Factor Models," *Econometrica*, 70, 191-221.
- Barsky, Robert B. and Lutz Kilian, 2002. "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in B. Bernanke and K. Rogoff (eds.), *NBER Macroeconomics Annual 2001*, May 2002, 137-183.
- Bastianin, Andrea and Manera, Matteo. 2014. "How does stock market volatility react to oil shocks?" FEEM Working Paper No. 70.2014
- Baumeister, Christiane and Geert Peersman, 2011. "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market," Bank of Canada Working Paper 11-28.
- Baumeister, Christiane and Lutz Kilian, 2012. "Real-Time Forecasts of the Real Price of Oil," *Journal of Business and Economic Statistics*, 30(2), 326-336.
- Blanchard, Olivier. 2009. "The State of Macro," *Annual Review of Economics*, 1(1), 209-228.
- Blinder, Alan S. and Jeremy B. Rudd, 2012. "The Supply Shock Explanation of the Great Stagflation Revisited," NBER Chapters, in: *The Great Inflation: The Rebirth of Modern Central Banking* National Bureau of Economic Research, Inc.
- Bosworth, Barry P. and Robert Z. Lawrence, 1982. *Commodity Prices and the New Inflation*, The Brookings Institution, Washington D.C.
- Chen, Yu Chin, Kenneth Rogoff, and Barbara Rossi, 2010. "Can Exchange Rates Forecast Commodity Prices?" *Quarterly Journal of Economics* 125(3), 1145-1194.
- Chinn, Menzie and Olivier Coibion, 2013. "The Predictive Content of Commodity Futures," forthcoming in the *Journal of Futures Markets*.

- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 1999. "Monetary policy shocks: What have we learned and to what end?," in: J. B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, edition 1, volume 1, chapter 2, pages 65-148 Elsevier.
- Coibion, Olivier and Yuriy Gorodnichenko, 2011. "Monetary Policy, Trend Inflation and the Great Moderation: An Alternative Interpretation," *American Economic Review*, 101(1), 341-370.
- Connor, Gregory, and Robert A. Korajczyk, 1993. "A Test for the Number of Factors in an Approximate Factor Model," *Journal of Finance* 48(4), 1263-1291.
- Deaton, Angus and Guy Laroque, 1992. "Behavior of Commodity Prices," *Review of Economic Studies* 59(1), 1-23.
- Deb, Partha, Pravin K. Trivedi, and Panayotis Varangis, 1996. "The Excess Comovement of Commodity Prices Reconsidered," *Journal of Applied Econometrics*, 11, 275-291.
- Foerster, Andrew, Pierre-Daniel G. Sarte, and Mark Watson, 2011. "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy* 119(1), 1-38.
- Forni, Mario and Lucrezia Reichlin, 1998. "Let's Get Real: A Factor Analytical Approach to Disaggregated Business Cycle Dynamics," *Review of Economic Studies* 65(3), 453-473.
- Frankel, Jeffrey A., 2008. "The Effect of Monetary Policy on Real Commodity Prices," in John H. Campbell (ed.) *Asset Prices and Monetary Policy*, NBER Books.
- Gorodnichenko, Yuriy, 2006. "Reduced-Rank Identification of Structural VARs," mimeo, U.C. Berkeley.
- Gospodinov, Nikolay and Serena Ng, 2013. "Commodity Prices, Convenience Yields, and Inflation," *The Review of Economics and Statistics* 95(1), 206-219.
- Hamilton, James D., 1983. "Oil and the Macroeconomy since World War II," *Journal of Political Economy* 91(2), 228-248.
- Hamilton, James D. 2009. "Causes and Consequences of the Oil Shock of 2007-2008," *Brookings Papers on Economic Activity* 2009(2): 215-259.
- Kilian, Lutz. 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90(2), 216-240.
- Kilian, Lutz. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99(3), June 2009, 1053-1069.
- Kilian, Lutz and Thomas K. Lee. 2013. "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," manuscript.
- Kilian, Lutz and Daniel P. Murphy. 2013. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," forthcoming: *Journal of Applied Econometrics*.
- Kose, M. Ayan, Christopher Otrok, and Eswar Prasad, 2012. "Global Business Cycles: Convergence or Decoupling?" *International Economic Review* 53(2), 511-538.
- Kozicki, Sharon and Tinsley, Peter A., 2009. "Perhaps the 1970s FOMC did what it said it did," *Journal of Monetary Economics*, vol. 56(6), pages 842-855.
- Lombardi, Marco J. and Ine Van Robays, 2012. "Do Financial Investors Destabilize the Oil Price?" ECB Working Paper 1346.
- Ng, Serena (1996). "Looking for Evidence of Speculative Stockholding in Commodity Markets," *Journal of Economic Dynamics and Control* 20, 123-143.
- Newey, Whitney K; West, Kenneth D (1987). "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". *Econometrica* 55 (3): 703-708.
- Onatski, Alexei. 2010. "Determining the Number of Factors From Empirical Distribution of Eigenvalues," *Review of Economics and Statistics*, 92(4), 1004-1016.
- Orphanides, Athanasios, "Historical Monetary Policy Analysis and the Taylor Rule", *Journal of Monetary Economics*, 50(5), 983-1022, July 2003.
- Pagan, Adrian. "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economics Review*, 25(1), 221-247.
- Pindyck, Robert S., Julio J. Rotemberg (1990). "The Excess Comovement of Commodity Prices," *The Economic Journal*, 100(403), 1173-1189.
- Reis, Ricardo and Mark W. Watson, 2010. "Relative Goods' Prices, Pure Inflation, and the Philipilps Correlation," *American Economic Journal: Macroeconomics* 2 (July), 128-157.

- Romer, Christina D. and David H. Romer, 2002. "The Evolution of Economic Understanding and Postwar Stabilization Policy," in *Rethinking Stabilization Policy*, Federal Reserve Bank of Kansas City.
- Romer, Christina D. and David H. Romer. 2004. "A New Measure Of Monetary Shocks: Derivation And Implications," *American Economic Review*, 2004, v94(4,Sep), 1055-1084.
- Stock, James and Mark W. Watson. 2002. "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147-162.
- Stock, James and Mark W. Watson. 2005. "Understanding Changes in International Business Cycle Dynamics," *Journal of the European Economic Association* 3(5), 968-1006.
- Uhlig, Harald, 2002. "What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics* 52(2), 381-419.
- Unalmis, D., I. Unalmis and D.F. Unsal, 2012. "On Oil Price Shocks: The Role of Storage," *IMF Economic Review*, 60, 505-532.
- West, Kenneth D. and Ka-Fu Wong. 2014. "A Factor Model for Co-movements of Commodity Prices," *Journal of International Money and Finance*, 42, 289-309.

TABLE 1: CONTRIBUTION OF COMMON FACTORS TO COMMODITY PRICES

Number of common factors:	Cumulative Variance Explained by Common Factors				
	1	2	3	4	5
<i>Complete Sample:</i>					
Cumulative eigenvalue shares	0.59	0.69	0.75	0.79	0.82
Mean across commodity-specific R^2 s	0.60	0.69	0.74	0.78	0.81
Median across commodity-specific R^2 s	0.70	0.76	0.78	0.84	0.85
R^2 across all commodities	0.62	0.71	0.75	0.79	0.82
<i>Subset of Commodities:</i>					
R^2 across food/agric. commodities	0.64	0.72	0.75	0.77	0.80
R^2 across oils	0.72	0.74	0.76	0.82	0.85
R^2 across industrial commodities	0.55	0.68	0.75	0.80	0.83

Note: The table provides metrics of the cumulative variance associated with using additional factors as indicated by each column. The first row provides the cumulative sum of eigenvalues associated with each factor normalized by the sum of all eigenvalues. The second row provides the mean across the R^2 of each commodity for each given factor, using the specific sample associated with each commodity. The third row provides the median R^2 across all commodity-specific R^2 s. The fourth row provides the joint R^2 constructed using all commodities. In addition, the top panel presents joint R^2 s for subsets of commodities (as defined in Table 1). Each of the R^2 omits imputed values. See section 3.2 for details.

TABLE 2: GMM ESTIMATES OF ROTATION MATRIX

	GMM Estimates of Rotation Parameter				Implied Rotation Coefficients			
	θ	$se(\theta)$	$p(\text{over-id})$	N	t_{11}	95% CI(t_{11})	t_{21}	95% CI(t_{21})
Baseline GMM Estimates: (Iterative GMM, $L=36$)	-0.10	(0.31)	1.00	505	1.00	[0.75 1.00]	-0.10	[-0.65 0.49]
Robustness of GMM Estimates:								
More moments: ($L=48$)	-0.15	(0.27)	1.00	493	0.99	[0.77 1.00]	-0.15	[-0.63 0.39]
Fewer moments: ($L=24$)	-0.13	(0.35)	1.00	517	0.99	[0.67 1.00]	-0.13	[-0.73 0.54]
Fewer moments: ($L=12$)	-0.23	(0.50)	1.00	529	0.97	[0.32 1.00]	-0.23	[-0.94 0.69]
Two-step GMM	-0.10	(0.31)	1.00	505	1.00	[0.75 1.00]	-0.10	[-0.65 0.47]
Continuous GMM	-0.07	(0.31)	1.00	505	1.00	[0.76 1.00]	-0.07	[-0.62 0.52]
Alternative normalization	-0.08	(0.31)	1.00	505	1.00	[0.75 1.00]	-0.08	[-0.64 0.50]

Notes: The table presents nonlinear GMM estimates of parameter θ from (7) in the text, along with Newey-West (1987) standard errors ($se(\theta)$), the p -value for over-identifying restrictions ($p(\text{over-id})$), and the number of observations used in the estimation (N). The panel on the right presents the implied parameters of the first row of the rotation matrix, along with the 95% confidence interval implied from the estimated distribution of θ . The baseline estimates are based on iterative GMM until convergence, using a constant as well as the contemporaneous value and 36 lags of OPEC production shocks for moment conditions. Subsequent rows present robustness to using more or fewer lags of OPEC production shocks as moment conditions, a 2-step GMM procedure, a continuously-updated GMM procedure, and an alternative normalization of moment conditions. See section 3.3 for details.

TABLE 3: ROTATED COMMODITY-SPECIFIC FACTOR LOADINGS

Commodity	Factor Loadings		Commodity	Factor Loadings	
	IC	DC		IC	DC
<i>Agr./Food Commodities</i>			<i>Oils</i>		
Apples	0.46	0.13	Coconut oil	0.82	0.02
Bananas	0.57	0.22	Groundnut oil	0.86	0.13
Barley	0.75	0.41	Palm oil	0.89	0.13
Beef	0.87	-0.09	Rapeseed oil	0.53	0.39
Cocoa	0.89	-0.12	Sun/Safflower oil	0.83	0.22
Coffee	0.85	-0.17			
Corn	0.95	0.09	<i>Industrial Commodities</i>		
Fishmeal	0.91	0.15	Aluminum	0.80	0.05
Hay	0.86	-0.04	Burlap	0.85	-0.00
Oats	0.88	0.11	Cement	0.21	0.06
Orange Juice	0.74	-0.22	Copper	0.60	0.69
Onions	0.53	-0.39	Cotton	0.92	-0.20
Pepper	0.56	-0.62	Lead	0.73	0.58
Potatoes	0.73	-0.05	Lumber	0.53	-0.23
Rice	0.93	0.09	Mercury	0.46	0.75
Shrimp	0.44	-0.75	Nickel	0.20	0.74
Sorghums	0.95	0.08	Rubber	0.79	0.45
Soybeans	0.95	0.02	Tin	0.90	0.18
Sugar	0.78	0.11	Wool	0.87	0.16
Tea	0.87	-0.22	Zinc	0.60	0.36
Tobacco	0.84	-0.33			
Wheat	0.92	0.13			

Note: The table presents the rotated loadings from factor analysis using the GMM estimates of the rotation matrix. See section 3.3 for details.

TABLE 4: TESTING THE NULL OF ZERO NET PURCHASES BY STORAGE SECTOR

Number of Factors:	Estimates of Mean Ratio of Consumption to Production - 1				
	\hat{c}	$se(\hat{c})$	N	Sample	Source
<i>Agr./Food Commodities</i>					
Apples	-0.007***	(0.003)	42	1968-2009	UN FAO
Bananas	-0.008**	(0.004)	42	1968-2009	UN FAO
Barley	0.001	(0.005)	33	1979-2011	CRB
Beef					
Cocoa	-0.009	(0.010)	43	1968-2010	CRB
Coffee	0.016	(0.011)	41	1968-2009	UN FAO
Corn	0.004	(0.005)	32	1980-2011	CRB
Fishmeal	-0.014	(0.016)	45	1968-2012	USDA-FAS
Hay					
Oats	0.002	(0.004)	45	1968-2012	USDA-FAS
Orange Juice					
Onions	-0.007***	(0.001)	42	1968-2009	UN FAO
Pepper	-0.000	(0.018)	42	1968-2009	UN FAO
Potatoes	0.005**	(0.002)	42	1968-2009	UN FAO
Rice	-0.010**	(0.005)	45	1968-2012	USDA-FAS
Shrimp					
Sorghums	0.010	(0.009)	28	1983-2011	CRB
Soybeans	-0.002	(0.006)	42	1968-2009	UN FAO
Sugar	-0.020***	(0.005)	45	1968-2012	USDA-FAS
Tea	-0.022***	(0.005)	42	1968-2009	UN FAO
Tobacco	0.004	(0.015)	37	1968-2004	USDA-FAS
Wheat	0.000	(0.006)	34	1978-2011	CRB
<i>Oils</i>					
Coconut oil	0.003	(0.009)	42	1968-2009	UN FAO
Groundnut oil	-0.003	(0.004)	41	1971-2011	USDA-FAS
Palm oil	-0.045**	(0.017)	42	1968-2009	UN FAO
Rapeseed oil	-0.007	(0.005)	45	1968-2012	USDA-FAS
Sun/Safflower oil	-0.024**	(0.010)	41	1972-2012	USDA-FAS
<i>Industrial Commodities</i>					
Aluminum	-0.007	(0.005)	45	1968-2012	TA
Burlap	0.020	(0.012)	42	1968-2009	UN FAO
Cement					
Copper	0.001	(0.005)	45	1968-2011	TA
Cotton	0.001	(0.010)	43	1968-2010	CRB
Lead	-0.001	(0.004)	39	1972-2012	TA
Lumber					
Mercury					
Nickel	-0.009	(0.008)	45	1968-2012	BREE
Rubber	0.001	(0.004)	43	1968-2010	CRB
Tin	0.011	(0.012)	45	1968-2012	TA
Wool					
Zinc	-0.007	(0.006)	39	1972-2012	TA

Note: The table presents the average ratio of consumption to production (minus one) for each commodity and associated Newey-West standard errors. Data on global consumption and production are from Commodity Research Bureau (CRB), trade associations (TA), the UN Food and Agriculture Organization (FAO), the Food and Agricultural Services of the U.S. Department of Agriculture (USDA-FAS), or the Bureau of Resources and Energy Economics of the Australian Government (BREE). See section 4 for details and additional information on specific trade organizations. Series left blank are those for which consumption and production data are unavailable.

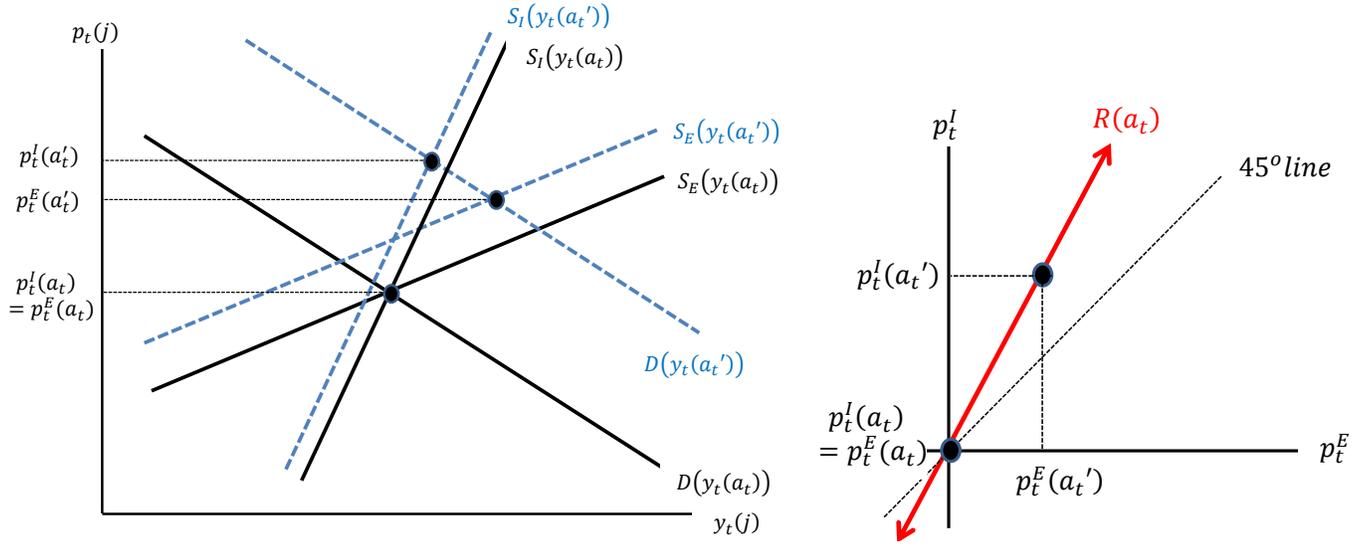
TABLE 5: SUMMARY OF RECURSIVE FORECAST ACCURACY DIAGNOSTICS FOR REAL COMMODITY PRICES

<u>Forecast Evaluation Period: Commodity-Specific</u>											
	<u>Aggregate MSPE Ratio</u>	<u>Distribution of MSPE Ratios</u>					<u>CRB</u>	<u>WB</u>	<u>IMF</u>	<u>Crude Oil</u>	
		<u>[0,0.9)</u>	<u>[0.9,0.95)</u>	<u>[0.95,1)</u>	<u>[0,1)</u>	<u>[1,∞)</u>					
1 month	0.921	10	11	11	32	8	0.974	0.834	0.874	0.805	
3 months	0.922	4	5	11	20	20	1.057	1.022	0.990	0.972	
6 months	0.938	5	4	4	13	27	1.127	1.245	1.072	1.141	
12 months	1.096	5	6	5	16	24	1.187	1.214	1.155	1.318	
No. of commodities	40						24 (15)	39 (17)	45(17)		
<u>Forecast Evaluation Period: 1984:1-2012:12</u>											
	<u>Aggregate MSPE Ratio</u>	<u>Distribution of MSPE Ratios</u>					<u>CRB</u>	<u>WB</u>	<u>IMF</u>	<u>Crude Oil</u>	
		<u>[0,0.9)</u>	<u>[0.9,0.95)</u>	<u>[0.95,1)</u>	<u>[0,1)</u>	<u>[1,∞)</u>					
1 month	0.930	8	7	7	22	6	0.964	0.863	0.888	0.790	
3 months	0.946	7	4	7	18	10	0.991	0.982	0.928	0.947	
6 months	0.984	8	3	3	14	14	1.068	1.106	1.008	1.111	
12 months	1.106	9	3	5	17	11	1.128	1.256	1.112	1.308	
No. of commodities	28						24 (15)	39 (17)	45 (17)		

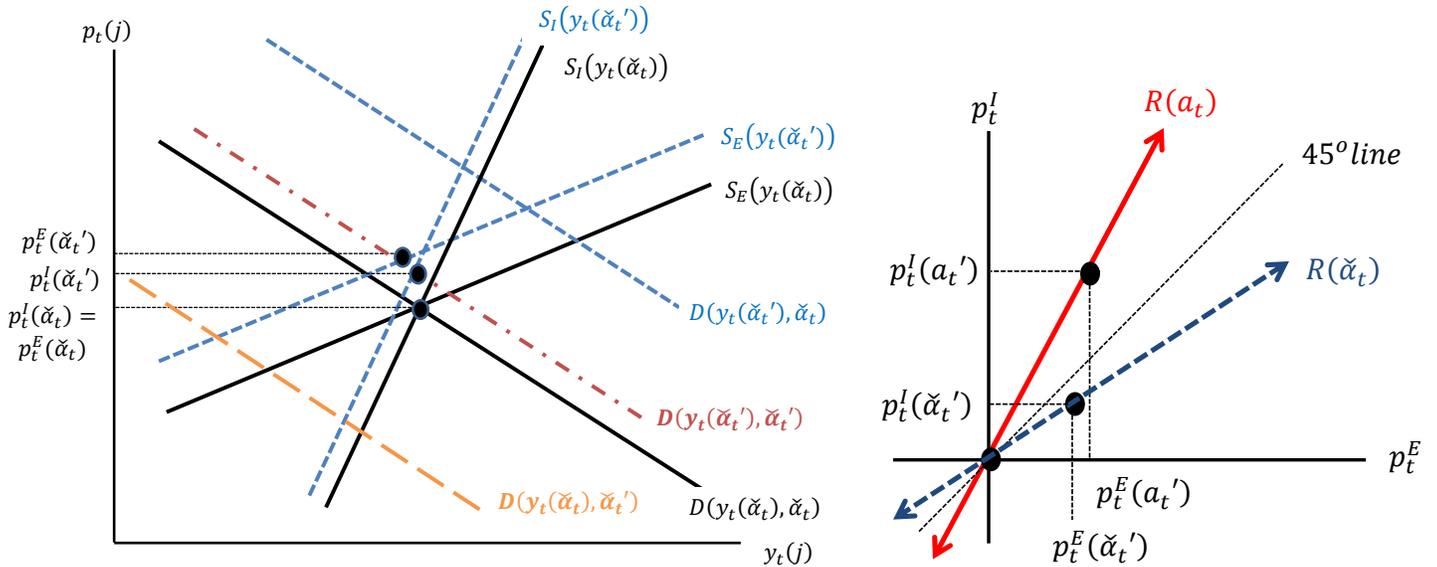
Notes: For the commodity-specific forecast evaluation period, the initial estimation window depends on the commodity. It begins either in 1968:1 or at the earliest date such that the initial estimation window contains at least 48 observations. The maximum length of the recursive sample is restricted by the end of the data and the forecast horizon. The “Aggregate MSPE Ratio” is the ratio of the sum of the MSPEs for the bivariate FAVAR forecasts of the real commodity prices relative to the sum of the MSPEs for the no-change forecast. The MSPE ratios of the individual real-commodity price forecasts are also computed relative to the benchmark no-change forecast. For the FAVAR-based forecasts, the lag length is chosen recursively using the BIC. The number of commodities included in the commodity-price indices but not in the cross-section of 40 commodities used to extract the factor is in parentheses.

FIGURE 1: COMPARATIVE STATICS AND COMMODITY-PRICE COMOVEMENT ACROSS SHOCKS

Panel A: Expansionary Change in Aggregate Productivity



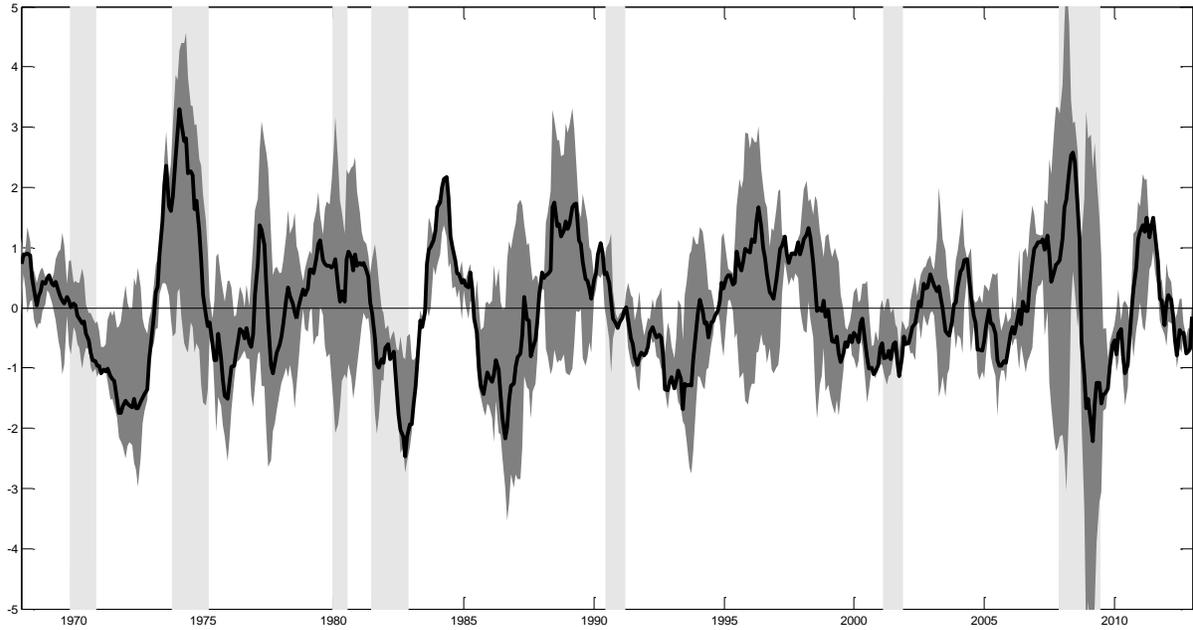
Panel B: Expansionary Change in Relative Demand for Commodities



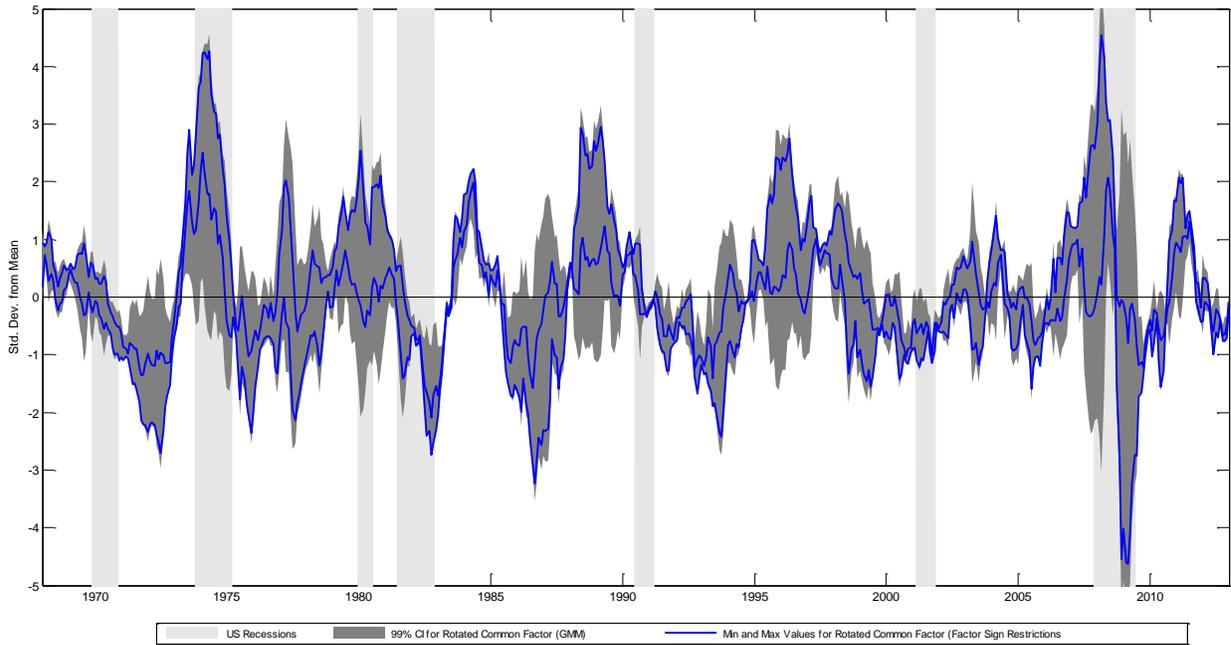
Notes: The two figures in Panel A plot the effects of a change in aggregate productivity from a_t to a_t' on commodity prices. In the graph on the left, S_E and S_I are supply curves for relatively elastically and inelastically supplied commodities, D denotes demand curves. In the graph on the right, $R(a)$ shows the set of prices of the two commodities that may arise as a result of productivity changes. The two figures in Panel B plot the equivalent comparative statics for a decrease in the relative demand for commodities (\check{a}_t), which is assumed to raise aggregate production y by the same amount as the increase in productivity in Panel A. See section 2.2 in the text for details.

FIGURE 2: INDIRECT COMMON FACTOR IN COMMODITY PRICES

Panel A: Indirect Common Factor (GMM Approach)



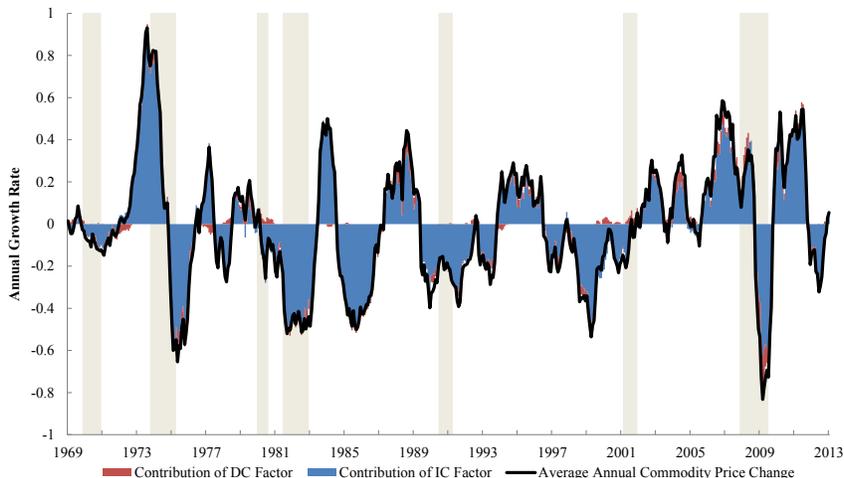
Panel B: Indirect Common Factor (Factor Loading Sign Restrictions)



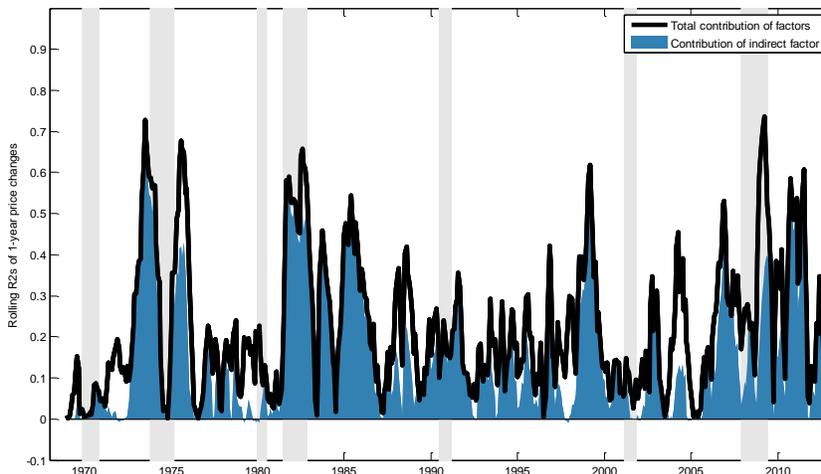
Note: The top figure presents the IC factor from the factor analysis in section 3.3. The IC factor is HP-filtered ($\lambda=129,600$) in the figure. The light grey shaded areas are NBER-dated recessions. The grey shaded areas are 99% confidence intervals of HP-filtered rotated factors constructed from the estimated distribution of rotation parameters. The bottom figure plots the 99% confidence interval of the IC factor as estimated by GMM (shaded areas) and the minimum and maximum range for admissible values of the IC factor using sign restrictions on factor loadings (solid blue lines). See sections 3.3 in the text for details.

FIGURE 3: THE CONTRIBUTION OF INDIRECT AND DIRECT FACTORS TO COMMODITY-PRICE CHANGES

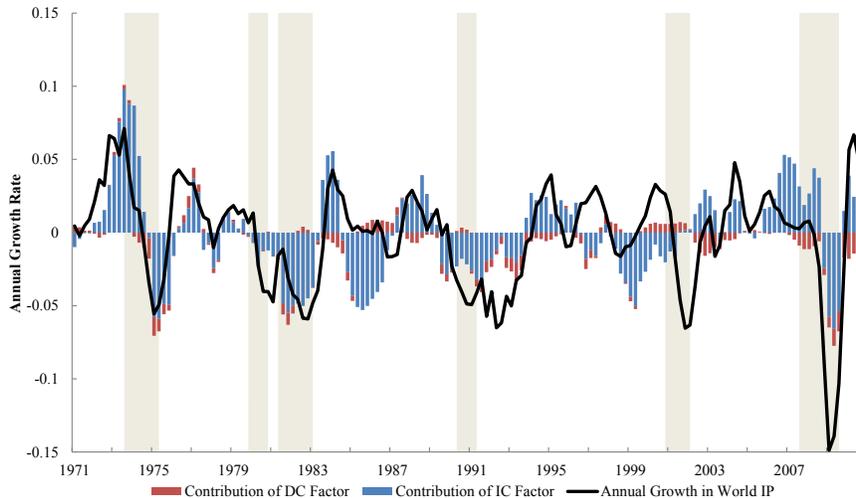
Panel A: Contributions to Average Annual Commodity-Price Changes



Panel B: Contributions to Comovement in Commodity Price Changes

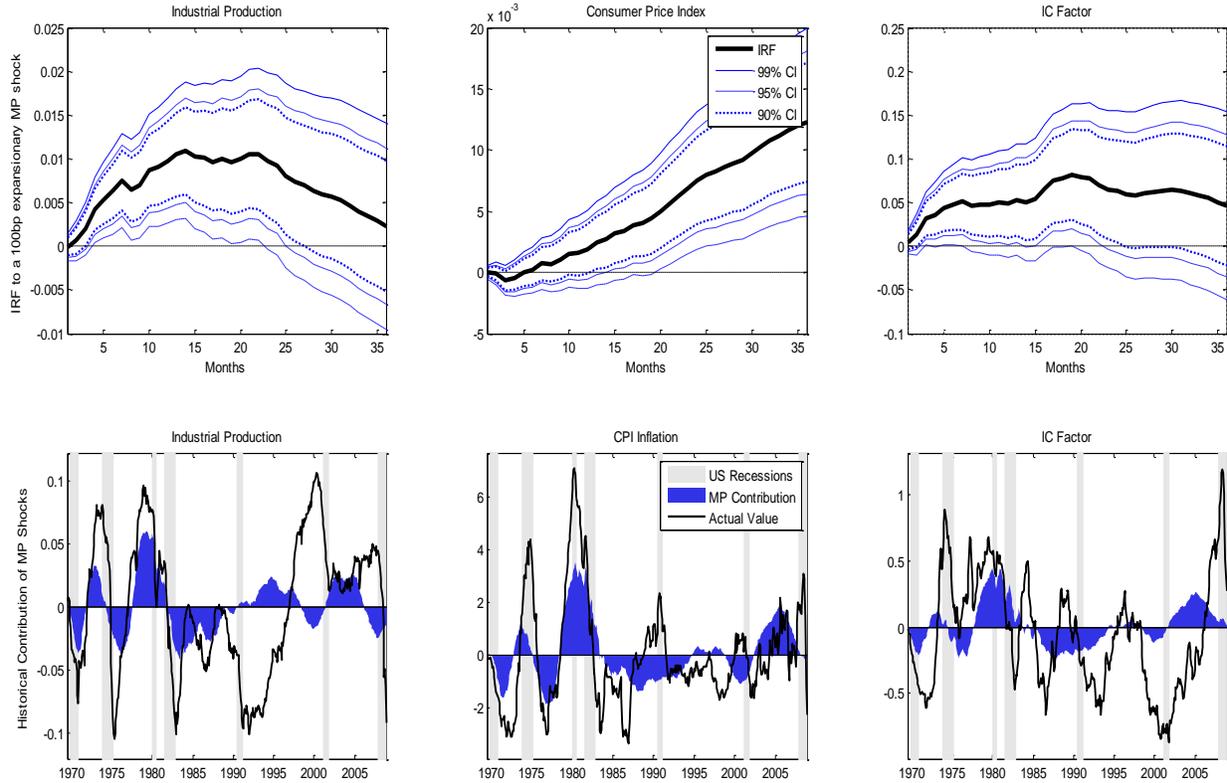


Panel C: Contributions to Annual Changes in Global Industrial Production



Note: Panel A plots the contributions of the direct and indirect factors (DC and IC, respectively) to the average annual commodity price change across all commodities. Panel B plots the contribution of the two factors to cross-sectional variation in 1-year commodity price changes (black line) and that coming solely from IC factor (blue shaded area). See section 3.4 for details. Panel C plots the equivalent contributions to the annual growth rate of global industrial production.

FIGURE 4: EFFECTS OF MONETARY POLICY SHOCKS ON THE INDIRECT COMMON FACTOR



Note: The figures in the top row present estimated impulse responses of U.S. industrial production, the U.S. consumer price index, and the IC factor to a 100 basis point expansionary monetary policy shock using the VAR described in section 4. Confidence intervals are constructed from the distribution of impulse responses generated by drawing 2000 times from the estimated distribution of VAR parameters. The bottom row presents actual values of each variable normalized by the predicted values from the VAR given initial conditions and no subsequent shocks (solid black line), U.S. recessions (light grey shaded areas), and the estimated contribution of monetary policy shocks to historical variation in each variable (blue areas). For the CPI, the bottom figure presents year-on-year inflation rates. See section 4 in the text for details.

NOT FOR PUBLICATION
APPENDICES

Appendix A: Model Derivations

The Household

A representative consumer maximizes expected discounted utility over consumption (C), labor supply (N^S), and the amount of another input supplied to each commodity sector ($L^S(j)$) as follows

$$\max E_t \sum_{i=0}^{\infty} \beta^i \left[\frac{C_{t+i}^{1-\sigma}}{1-\sigma} - e^{-\varepsilon_{t+i}^n} \varphi_n \frac{N_{t+i}^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}} - \varphi_L e^{-\varepsilon_t^L} \frac{\int_0^1 L_{t+i}^S(j)^{1+\frac{1}{\nu}} dj}{1+\frac{1}{\nu}} \right]$$

where β is the discount factor. The $e^{\varepsilon_t^n}$ term is an exogenous shock to the disutility of hours worked while $e^{\varepsilon_t^L}$ is an exogenous shock to the disutility of supplying land.

The household pays a price P_t for the consumption good, receives wage W_t for each unit of labor supplied, and is paid a rental rate of land $R_t^L(j)$ for each unit of land supplied to the primary commodity sector j . The household also can purchase risk-free bonds B_t that pay a gross nominal interest rate of R_t . The budget constraint is

$$P_t C_t + B_t = B_{t-1} R_{t-1} + W_t N_t^S + \int_0^1 R_t^L(j) L_t^S(j) dj + T_t$$

where T_t represent payments from the ownership of firms. Assuming that the household takes all prices as given, its first-order conditions are

$$\varphi_n C_t^\sigma N_t^{S\eta^{-1}} = e^{\varepsilon_t^n} W_t / P_t \quad (\text{A.1})$$

$$\varphi_L C_t^\sigma L_t^S(j)^{\nu^{-1}} = e^{\varepsilon_t^L} R_t^L(j) / P_t \quad (\text{A.2})$$

$$C_t^{-\sigma} = E_t \beta \left[C_{t+1}^{-\sigma} R_t \frac{P_t}{P_{t+1}} \right]. \quad (\text{A.3})$$

The Primary Commodity-Production Sector

Each primary commodity j is produced by a representative price-taking firm who uses land ($L_t^d(j)$) to produce a quantity $Q_t(j)$ of good j given a production function

$$Q_t(j) = A_t(j) L_t^d(j)^{\alpha_j} \quad (\text{A.4})$$

where $A_t(j)$ is the exogenously determined level of productivity for commodity j and $0 < \alpha_j < 1$ is the commodity-specific degree of diminishing returns to land. Given the price of commodity j $P_t(j)$, and the rental rate of land $R_t^L(j)$ specific to commodity j , the firm chooses the amount of land input to maximize profits

$$\max P_t(j) Q_t(j) - R_t^L(j) L_t^d(j).$$

This yields the following demand curve for land for each commodity j :

$$R_t^L(j) / P_t = \alpha_j \left(\frac{P_t(j)}{P_t} \right) A_t(j) L_t^d(j)^{\alpha_j - 1}. \quad (\text{A.5})$$

We assume that the steady-state level of productivity $\bar{A}(j)$ is such that the steady-state level of production in each sector is equal. Equilibrium in the market for land requires

$$L_t^S(j) = L_t^d(j) \quad (\text{A.6})$$

for each sector j .

The Intermediate Commodity

A perfectly competitive sector purchases $Y_t(j)$ of each primary commodity j and aggregates them into an intermediate commodity Q_t^C using the Dixit-Stiglitz aggregator

$$Q_t^C = \left(\int_0^1 Y_t^j \frac{\theta_c - 1}{\theta_c} dj \right)^{\frac{\theta_c}{\theta_c - 1}} \quad (\text{A.7})$$

which yields a demand for each commodity j of

$$P_t(j) / P_t^C = (Y_t(j) / Q_t^C)^{-1/\theta_c} \quad (\text{A.8})$$

where θ_c is the elasticity of substitution across commodities and the price of the intermediate commodity aggregate is given by $P_t^C = \left(\int_0^1 P_t(j)^{1-\theta_c} dj \right)^{\frac{1}{1-\theta_c}}$. Market-clearing for each commodity sector j requires

$$Q_t(j) = Y_t(j). \quad (\text{A.9})$$

The Final Goods Sector

A perfectly competitive sector combines purchases of the intermediate commodity good Y_t^C and labor N_t^d according to the Cobb-Douglas production function

$$Y_t = A_t Y_t^C \alpha_t N_t^d^{1-\alpha_t} \quad (\text{A.10})$$

to maximize profits

$$P_t Y_t - W_t N_t^d - P_t^C Y_t^C$$

taking as given all prices and where A_t is an exogenously determined aggregate productivity process. This yields the following demand for each input

$$\alpha_t = (P_t^C / P_t) (Y_t^C / Y_t) \quad (\text{A.11})$$

$$1 - \alpha_t = (W_t / P_t) (N_t^d / Y_t) \quad (\text{A.12})$$

Since all of the final good is purchased by the household, equilibrium in the final goods market requires that $C_t = Y_t$. The fact that α_t is potentially time-varying allows for exogenous variation in the relative demand for commodities and labor in the production of the final good.

The Linearized Model

We assume that exogenous processes are stationary around their steady-state levels, so that all real variables are constant in the steady-state. Letting lower-case letters denote log-deviations from steady-state (e.g., $c_t \equiv \log C_t - \log \bar{C}$) and normalizing all nominal variables by the final goods price level (e.g., $p_t(j) \equiv \log P_t(j) / P_t - \log(\bar{P}(j) / \bar{P})$), the first-order conditions from the household's problem are

$$\sigma y_t + \frac{1}{\eta} n_t = w_t + \varepsilon_t^n \quad (\text{A.13})$$

$$\sigma y_t + \frac{1}{\nu} l_t(j) = r_t^L(j) + \varepsilon_t^L \quad (\text{A.14})$$

$$y_t = E_t \left[y_{t+1} - \frac{1}{\sigma} r_t \right] \quad (\text{A.15})$$

where we have imposed the market-clearing conditions $C_t = Y_t$ and $N_t^d = N_t^s \equiv N_t$ and defined r_t as the log-deviation of the gross real interest rate from its steady-state value.

Each primary commodity-producing sector is summarized by the following equations

$$r_t^L(j) = p_t(j) + a_t(j) - (1 - \alpha_j) l_t(j) \quad (\text{A.16})$$

$$y_t(j) = a_t(j) + \alpha_j l_t(j) \quad (\text{A.17})$$

where we have imposed the market clearing conditions $L_t^s(j) = L_t^d(j) \equiv L_t(j)$ and $Q_t(j) = Y_t(j)$. The intermediate commodity sector is given by

$$y_{c,t} = \int_0^1 y_t(j) dj \quad (\text{A.18})$$

$$p_t(j) = p_{c,t} - \frac{1}{\theta_c} (y_t(j) - y_{c,t}). \quad (\text{A.19})$$

Finally, letting α be the steady-state value of α_t and the log-deviation of α_t from its steady-state value of α be $\check{\alpha}_t$, the final goods sector follows

$$y_t = a_t + \alpha y_{c,t} + (1 - \alpha) n_t + \varphi_\alpha \check{\alpha}_t \quad (\text{A.20})$$

$$p_{c,t} = y_t - y_{c,t} + \check{\alpha}_t \quad (\text{A.21})$$

$$w_t = y_t - n_t - \frac{\alpha}{1-\alpha} \check{\alpha}_t. \quad (\text{A.22})$$

where $\varphi_\alpha \equiv \alpha (\ln \bar{Y}^C - \ln \bar{N})$.

Equilibrium Dynamics

Labor market equilibrium for primary commodity j requires

$$l_t(j) = \left(\frac{1}{v} + 1 - \alpha_j\right)^{-1} [p_t(j) + a_t(j) - \sigma y_t + \varepsilon_t^L]$$

so production of commodity j is given by

$$y_t(j) = a_t(j)(1 + \varepsilon_j^{-1}) + \varepsilon_j^{-1} p_t(j) - \sigma \varepsilon_j^{-1} y_t + \varepsilon_j^{-1} \varepsilon_t^L \quad (\text{A.23})$$

where $\varepsilon_j \equiv \left(\frac{1}{v} + 1 - \alpha_j\right) / \alpha_j$. Substituting in the relative demand for commodity j yields

$$y_t(j) = v_t(j) + \left(1 + \frac{1}{\varepsilon_j \theta_c}\right)^{-1} \left[\varepsilon_j^{-1} (p_{c,t} + \frac{1}{\theta_c} y_{c,t}) - \sigma \varepsilon_j^{-1} y_t + \varepsilon_j^{-1} \varepsilon_t^L\right]. \quad (\text{A.24})$$

where $v_t(j) \equiv a_t(j) \left(1 + \frac{1}{\varepsilon_j \theta_c}\right)^{-1} (1 + \varepsilon_j^{-1})$ is a rescaled version of each commodity's productivity.¹⁶

Then the aggregate supply of commodities follows from aggregating (24) across all j

$$p_{c,t} = \frac{1}{\theta_c} \left(\frac{1}{\varphi} - 1\right) y_{c,t} + \sigma y_t - \frac{1}{\varphi \theta_c} v_t - \varepsilon_t^L \quad (\text{A.25})$$

where $\varphi \equiv \int_0^1 (1 + \varepsilon_j \theta_c)^{-1} dj$ s.t. $0 < \varphi < 1/2$ and $v_t \equiv \int_0^1 v_t(j) dj$ is the aggregate over the rescaled productivity shocks in all commodity sectors. The aggregate output level on the right-hand side of (A.25) reflects income effects to the supply of land on the part of the household, which lower the aggregate supply of commodities when income is high. The supply of commodities also shifts with the aggregated commodity productivity level and shocks to the household's willingness to supply land.

With the demand for the commodity bundle given by $p_{c,t} = y_t - y_{c,t} + \check{\alpha}_t$, equilibrium production of the intermediate commodity bundle is given by

$$y_{c,t} = \frac{(1-\sigma)\theta_c\varphi}{1+(\theta_c-1)\varphi} y_t + \frac{1}{1+(\theta_c-1)\varphi} v_t + \frac{\theta_c\varphi}{1+(\theta_c-1)\varphi} \varepsilon_t^L + \frac{\theta_c\varphi}{1+(\theta_c-1)\varphi} \check{\alpha}_t. \quad (\text{A.26})$$

Whether equilibrium total commodity production rises or falls with income (holding v and ε^L constant) depends on the strength of the income effect, which here is captured by σ . If $\sigma < 1$, then commodity production comoves positively with total production.

Equilibrium in the labor market is given by

$$n_t = \frac{1-\sigma}{1+\eta^{-1}} y_t + \frac{1}{1+\eta^{-1}} \varepsilon_t^n - \frac{1}{1+\eta^{-1}} \check{\alpha}_t \quad (\text{A.27})$$

Therefore, the aggregate level of production of final goods follows from the production function

$$y_t = \varphi_y [a_t + \kappa_L \varepsilon_t^L + \kappa_n \varepsilon_t^n + \kappa_v v_t + \kappa_\alpha \check{\alpha}_t] \quad (\text{A.28})$$

where $\varphi_y \equiv \left(1 - \alpha \left[\frac{(1-\sigma)\theta_c\varphi}{1+(\theta_c-1)\varphi}\right] - (1-\alpha) \left[\frac{1-\sigma}{1+\eta^{-1}}\right]\right)^{-1}$, $\kappa_L \equiv \frac{\alpha\theta_c\varphi}{1+(\theta_c-1)\varphi}$, $\kappa_n \equiv \frac{1-\alpha}{1+\eta^{-1}}$, $\kappa_v \equiv \frac{\alpha}{1+(\theta_c-1)\varphi}$, and $\kappa_\alpha \equiv \varphi_\alpha + \frac{\alpha\varphi\theta_c}{1+\varphi(\theta_c-1)} - \frac{\alpha}{1+\eta^{-1}}$. Output is rising with aggregate productivity, positive shocks to the household's willingness to supply land and labor, a positive average over commodity-specific productivity shocks. Whether output rises when the relative demand for commodities increases ($\check{\alpha}_t$) depends on specific parameter values.

Comovement in Commodity Prices

We assume that productivity shocks to each commodity sector have an idiosyncratic component and a common component such that $v_t(j) = v_t^a + v_t^j$, which implies that the aggregate across commodities is $v_t = v_t^a$. The idiosyncratic shocks are orthogonal across commodity sectors, such that $E[v_t^j v_t^k] = 0 \forall j \neq k$ and $E[v_t] = 0$.

We now consider the determinants of individual commodity prices. First, the supply of commodity j follows from equations (A.19), (A.21) and (A.26) and is given by

¹⁶ The rescaling of the commodity-specific productivity shock ensures that a 1% increase in productivity in each commodity sector raises the equilibrium level of production of that commodity by equal amounts for each commodity. This would not be the case without the rescaling because each primary commodity sector's supply curve has a different slope. The rescaling simplifies the aggregation across commodity sectors.

$$p_t(j) = \varepsilon_j y_t(j) - \frac{(1+\varepsilon_j\theta_c)}{\theta} (v_t^a + v_t^j) + \sigma y_t - \varepsilon_t^L \quad (\text{A.29})$$

where ε_j is the elasticity of commodity supply with respect to its price. We can write the supply curve of commodity j as in the text

$$p_t(j) = S_j(y_t(j); y_t(a_t, \varepsilon_t^n, \varepsilon_t^L, \check{\alpha}_t, v_t^a); v_t^a, \varepsilon_t^L, v_t^j) \quad (\text{A.30})$$

which captures the fact that some shocks affect the supply of commodity j indirectly through general equilibrium effects captured by aggregate output, some shocks affect supply directly by shifting the curve holding aggregate output fixed, and some shocks do both.

The demand for commodity j comes from combining equation (A.19) with (A.21) and (A.25) yielding

$$p_t(j) = -\frac{1}{\theta_c} y_t(j) + \left(\frac{1+(\theta_c-1)\sigma\varphi}{1+(\theta_c-1)\varphi} \right) y_t - \frac{\varphi(\theta_c-1)}{1+(\theta_c-1)\varphi} \varepsilon_t^L - \frac{(\theta_c-1)}{1+(\theta_c-1)\varphi} \left(\frac{1}{\theta_c} \right) v_t^a + \frac{1}{1+\varphi(\theta_c-1)} \check{\alpha}_t \quad (\text{A.31})$$

We can again write the demand curve of commodity j more succinctly as

$$p_t(j) = D_j(y_t(j); y_t(a_t, \varepsilon_t^n, \varepsilon_t^L, \check{\alpha}_t, v_t^a); v_t^a, \varepsilon_t^L, \check{\alpha}_t) \quad (\text{A.32})$$

to highlight that some shocks affect the demand for commodity j indirectly through general equilibrium effects on output, some shocks shift the demand for each commodity j directly holding aggregate output constant, and some do both.

The Factor Structure in Commodity Prices

To solve for commodity prices, we combine equations (A.29) and (A.31) yielding

$$p_t(j)(1 + \varepsilon_j\theta_c) = \left[\sigma + \frac{\varepsilon_j\theta_c(1+(\theta_c-1)\sigma\varphi)}{1+(\theta_c-1)\varphi} \right] y_t - \left[\frac{\varepsilon_j\theta_c\varphi(\theta_c-1)}{1+(\theta_c-1)\varphi} + 1 \right] \varepsilon_t^L - \frac{1}{\theta_c} \left(1 + \varepsilon_j\theta_c + \frac{\varepsilon_j\theta_c(\theta_c-1)}{1+(\theta_c-1)\varphi} \right) v_t^a + \frac{\varepsilon_j\theta_c}{1+\varphi(\theta_c-1)} \check{\alpha}_t - \frac{1}{\theta_c} (1 + \varepsilon_j\theta_c) v_t^j. \quad (\text{A.33})$$

Because aggregate output y_t is itself a function of all aggregate shocks in the model, we can decompose it as follows

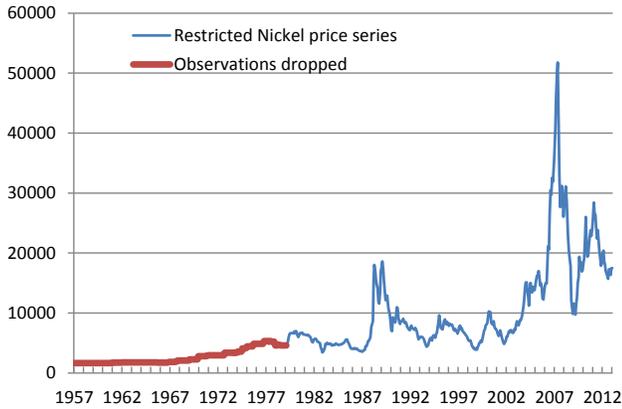
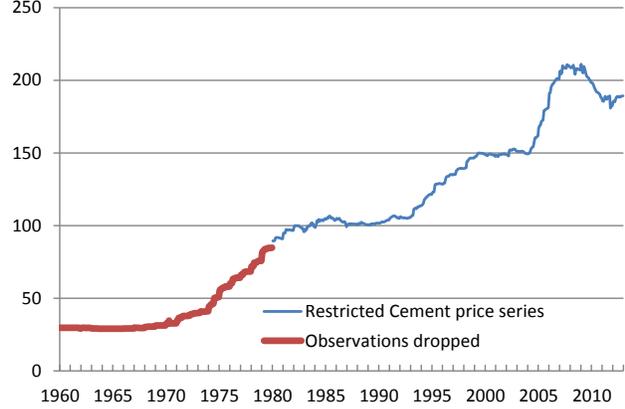
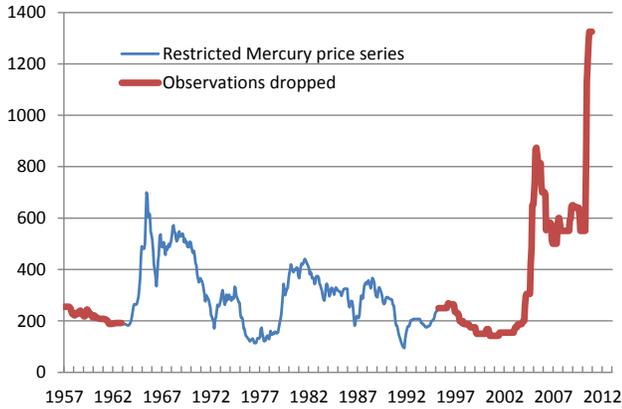
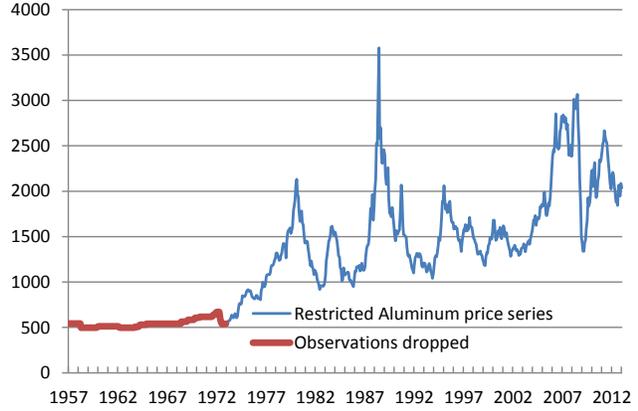
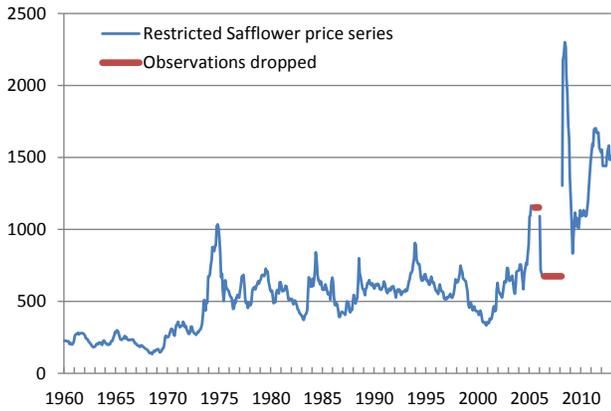
$$y_t = y_t^{nc}(a_t, \varepsilon_t^n) + \varphi_y[\kappa_L \varepsilon_t^L + \kappa_v v_t^a + \kappa_\alpha \check{\alpha}_t]$$

where $y_t^{nc} = \varphi_y[a_t + \kappa_n \varepsilon_t^n]$. Given this decomposition, we can rewrite the equilibrium price of commodity j as

$$p_t(j) = \underbrace{\lambda_j^y y_t^{nc}(a_t, \varepsilon_t^n)}_{\text{indirect (IC)}} + \underbrace{\lambda_j^L \varepsilon_t^L + \lambda_j^v v_t^a + \lambda_j^\alpha \check{\alpha}_t}_{\text{direct (DC)}} - \underbrace{\frac{1}{\theta_c} v_t^j}_{\text{idiosyncratic}} = \lambda_j F_t + \xi_t^j \quad (\text{A.34})$$

where $\lambda_j^y \equiv (1 + \theta \varepsilon_j)^{-1} \left[\sigma + \frac{\varepsilon_j\theta(1+(\theta-1)\sigma\varphi)}{1+(\theta-1)\varphi} \right]$, $\lambda_j^L \equiv \varphi_y \kappa_L \lambda_j^y - \left[\frac{1}{1+\varepsilon_j\theta_c} + \frac{\varepsilon_j\theta_c}{1+\varepsilon_j\theta_c} (\theta_c - 1) \right]$, $\lambda_j^v \equiv \varphi_y \kappa_v \lambda_j^y - \left[\frac{1}{\theta_c} + \frac{\varepsilon_j\theta_c}{1+\varepsilon_j\theta_c} \left(\frac{\varphi(\theta_c-1)}{\varphi\theta} \right) \right]$, $\lambda_j^\alpha \equiv \varphi_y \kappa_\alpha \lambda_j^y + \frac{\varepsilon_j\theta_c}{1+\varepsilon_j\theta_c} \left(\frac{1}{1+\varphi(\theta_c-1)} \right)$, $\lambda_j \equiv \left[\lambda_j^y \lambda_j^L \lambda_j^v \lambda_j^\alpha \right]$, $F_t \equiv [y_t^{nc} \varepsilon_t^L v_t^a \check{\alpha}_t]$, and $\xi_t^j \equiv -\frac{1}{\theta_c} v_t^j$ which is the expression in the text.

Appendix B: Price Observations Dropped



Note: Each figure presents the price series used in the empirical analysis (light blue line: “Restricted X price series”) and the observations dropped (thick red line: “Observations dropped”).

Appendix C: Notes on Commodity Price Data

Commodity	Sources	Description	Available Sample	Additional Notes
Apples	CRB	Wholesale price of (delicious) apples in U.S. until 1978:12, apple price received by growers starting 1979:1	1957:1-2011:12	Data from 1979:1 is apple price received by growers. Data prior to that is wholesale price of (delicious) apples in U.S., rescaled by average price ratio of two series from 1979:1-1980:12. Data prior to 1979 has numerous missing values.
Bananas	WB	Bananas (Central & South America), major brands, U.S. import price, free on truck (f.o.t.) U.S. Gulf ports	1960:1-2013:1	
Barley	CRB/WB	WB: Barley (Canada), feed, Western No. 1, Winnipeg Commodity Exchange, spot, wholesale farmers' price. CRB: No. 3 straight Barley, Minneapolis Exchange.	1957:1-2013:1	Data from 1957:1-1959:12 is CRB series. Data from 1960:1-2013:1 is WB series rescaled by ratio of the two series in 1960:1.
Beef	IMF	Australian and New Zealand, frozen boneless, 85 percent visible lean cow meat, U.S. import price FOB port of entry	1957:1-2013:1	
Cocoa	IMF	International Cocoa Organization cash price. Average of the three nearest active futures trading months in the New York Cocoa Exchange at noon and the London Terminal market at closing time, CIF U.S. and European ports.	1957:1-2012:12	
Coffee	IMF	International Coffee Organization; cash prices for 4 kinds of beans: Brazilian unwashed Arabica, Columbian mild Arabica, other mild Arabica, and Robustas.	1957:1-2012:12	Value for 1957:1 is average across all four types of coffee beans. Subsequent values are equally-weighted average of percent change in price of each kind of bean times previous period's price.
Corn	IMF	U.S. No. 2 yellow, prompt shipment, FOB Gulf of Mexico ports (USDA, Grain and Feed Market News, Washington, D.C.).	1957:1-2012:12	
Fishmeal	IMF	Peru Fish meal/pellets, 65% protein, CIF United Kingdom (DataStream)	1957:1-2012:12	
Hay	CRB	Mid-month price received by farmers for all hay (baled) in the U.S., dollars per ton	1957:1-2012:2	
Oats	CRB CD		1957:1-2010:11	
Orange Juice	CRB CD	Orange Juice Frozen Concentrate: nearest-	1967:1-	

		term futures contract traded on ICE.	2012:10	
Onions	CRB	Average price received by farmers.	1957:1-2011:12	
Pepper	CRB	1- Average black pepper (Brazilian) arriving in NY. 2- Average black pepper (Lampong) arriving in NY.	1957:1-2007:6	From 1984:1-2007:6, we use Brazilian pepper price. Prior to 1984, we use Lampong price rescaled by ratio of two prices in 1984:1.
Potatoes	CRB	Average price received by farmers	1957:1-2011:12	
Rice	IMF	Thai, white milled, 5 percent broken, nominal price quotes, FOB Bangkok (USDA, Rice Market News, Little Rock, Arkansas).	1957:1-2012:12	
Shrimp	IMF	Mexican, west coast, white, No. 1, shell-on, headless, 26 to 30 count per pound, wholesale price at New York	1957:1-2013:1	
Sorghums	CRB/WB	CRB: average price of no. 2, yellow, at Kansas City, \$/100 pounds, WB: no. 2 milo yellow, f.o.b. Gulf ports	1957:1-2013:1	From 1960:1-2013:1, we use the WB series. Prior to 1960:1, we use the CRB series rescaled by the ratio of the two series in 1960:1.
Soybeans	CRB CD	No. 1 yellow, Chicago Board of Trade.	1959:7-2012:9	
Sugar	IMF	CSCE contract No. 11, nearest future position (Coffee, Sugar and Cocoa Exchange, New York Board of Trade).	1957:1-2012:12	
Tea	IMF	Mombasa auction price for best PF1, Kenyan Tea. Replaces London auction price beginning July 1998	1957:1-2013:1	
Tobacco	WB	Tobacco (any origin), unmanufactured, general import, cif, U.S.	1968:1-2013:1	
Wheat	IMF	U.S. No. 1 hard red winter, ordinary protein, prompt shipment, FOB \$/Mt, Gulf of Mexico ports (USDA, Grain and Feed Market News).	1957:1-2012:12	
Coconut oil	CRB	Avg price of coconut oil (crude) at Pacific Coast of U.S. and Avg price of coconut oil (crude) tank cars in NY	1965:1-2010:12	Data from 1965:1-1980:12 is Pacific Coast, data from 1981:1-2010:12 is NY. Series have identical prices in overlapping months: 1980:1-1980:12.
Groundnut oil	WB	Groundnut oil (any origin), c.i.f. Rotterdam	1960:1-2013:1	
Palm oil	IMF	Crude Palm Oil Futures (first contract forward) 4-5 percent FFA, Bursa Malaysian Derivatives Berhad.	1957:1-2013:1	

Rapeseed oil	IMF	Crude, fob Rotterdam (Datastream)	1980:1-2013:1	
Sun/Safflower oil	IMF	Sunflower Oil, crude, U.S. export price from Gulf of Mexico (DataStream)	1960:1-2013:1	Data from 2005:7-2005:12 and data from 2006:6-2008:2 treated as missing because of no price variation.
Aluminum	IMF	London Metal Exchange, standard grade, spot price, minimum purity 99.5 percent, CIF U.K. ports (Wall Street Journal, New York and Metals Week, New York). Prior to 1979, U.K. producer price, minimum purity 99 percent.	1957:1-2013:1	Data from 1957:1-1972:12 treated as missing because of infrequent price variation.
Burlap	CRB CD	Original source of data is USDA.	1957:1-2012:9	
Cement	BLS	BLS PPI Index Industry (series PCU32731-32731) Cement Manufacturing.	1965:1-2012:12	Data prior to 1980:1 is treated as missing because of infrequent price variation.
Copper	IMF	London Metal Exchange, grade A cathodes, spot price, CIF European ports (Wall Street Journal, New York and Metals Week, New York). Prior to July 1986, higher grade, wire bars, or cathodes.	1957:1-2012:12	
Cotton	IMF	Middling 1-3/32 inch staple, Liverpool Index "A", average of the cheapest five of fourteen styles, CIF Liverpool (Cotton Outlook, Liverpool). From January 1968 to May 1981 strict middling 1-1/16 inch staple. Prior to 1968, Mexican 1-1/16.	1957:1-2012:12	
Lead	IMF	London Metal Exchange, 99.97 percent pure, spot price, CIF European ports	1957:1-2012:12	
Lumber	CRB/IMF	CRB: Douglas-fir softwood lumber 2x4 dried, S4S, IMF: Average export price of Douglas-fir, Western hemlock and other sawn softwood exported from Canada.	1957:1-2012:12	From 1975:1-2012:12, we use the IMF series. Prior to 1975:1, we use the CRB series rescaled by the ratio of the two price series in 1975:1.
Mercury	CRB	Average cash price in N.Y. for flask of 76 pounds.	1957:1-2010:12	Only data from 1962:12-1995:3 is used, other periods display infrequent price adjustment.
Nickel	IMF	London Metal Exchange, melting grade, spot price, CIF Northern European ports (Wall Street Journal, New York and Metals Week, New York). Prior to 1980 INCO, melting grade, CIF Far East and American ports (Metal Bulletin, London).	1957:1-2013:1	Data prior to 1979:3 is treated as missing because of infrequent price variation.

Rubber	CRB	Average spot crude rubber prices (smoked sheets, no 1, ribbed, plantation rubber) in NY, cents per pound	1957:1-2010:12
Tin	IMF	London Metal Exchange, standard grade, spot price, CIF European ports (Wall Street Journal, New York, New York). From Dec. 1985 to June 1989 Malaysian, straits, minimum 99.85 percent purity, Kuala Lumpur Tin Market settlement price. Prior to November 1985, London Metal Exchange	1957:1-2012:12
Wool	IMF	23 micron (AWEX, Australian Wool Exchange) Sidney, Australia	1957:1-2012:12
Zinc	IMF	London Metal Exchange, high grade 98 percent pure, spot price, CIF U.K. ports (Wall Street Journal and Metals Week, New York). Prior to January 1987, standard grade.	1957:1-2012:12

Appendix D: The Production and Use of Commodities

	Largest Producers	Primary Uses
<i>Agr./Food Commodities</i>		
Apples (1990-91)	U.S. (0.21), Germany (0.10), Italy (0.10)	Food (0.86), beverage, feed.
Bananas* (1990)	India (0.15), Brazil (0.12), Ecuad. (0.07)	Food (0.84), feed, other.
Barley (1990-91)	USSR (0.28), Germany (0.08)	Feed (0.73), distillation, food
Beef		Food
Cocoa (1990-91)	Ivory Coast (0.32), Brazil (0.25)	Food (0.96)
Coffee (1990-91)	Brazil (0.31), Columbia (0.14)	Food/beverages (0.98)
Corn (1990-91)	U.S. (0.42), Brazil (0.05)	Feed (0.62), food (0.16), adhesives
Fishmeal* (1984)	Japan (0.21), Chile (0.17), Peru (0.08)	Feed (0.90)
Hay		Feed
Oats (1990-91)	USSR (0.39), U.S. (0.13)	Food (0.74), feed (0.09), ref. solvent
Orange Juice (1990-1)	Oranges: Brazil (0.35), Spain (0.07)	Beverage (pulp for feed, oil)
Onions* (1990)	China (0.16), India (0.10)	Food (0.91)
Pepper (1990)	Main exporters: Indonesia, India	Food (0.96), oil (medical, perfumes)
Potatoes* (1990)	USSR (0.24), Poland (0.13)	Food (0.52), distillation, feed (0.19).
Rice (1990-91)	China (0.36), India (0.21)	Food (0.84), distillation, other.
Shrimp		Food
Sorghum* (1990)	U.S. (0.26), India (0.21), Mex. (0.11)	Food (0.39), feed (0.52)
Soybeans (1990-91)	U.S. (0.50), Brazil (0.15)	Food/feed (0.11), (paints, plastics)
Sugar (1990-91)	India (0.12), Brazil (0.07), Cuba (0.07)	Food/beverages (0.96), fuel.
Tea (1990)	India (0.29), China (0.21), S. Lank (.09)	Beverage (0.98)
Tobacco (1990)	China (0.37), U.S. (0.10)	Smoking
Wheat (1990-91)	USSR (0.17), China (0.17), U.S. (0.13)	Food (0.65), feed (0.22)
<i>Oils</i>		
Coconut oil (1990-91)	Philippines (0.41), Indonesia (0.27)	Food (0.57), cosmetics, synth. Rubber
Groundnut oil* (1990)	India (0.45), China (0.22), Nigeria (.09)	Food (0.98)
Palm oil (1990-91)	Malaysia (0.55), Indonesia (0.25)	Food (0.57), soaps, machine lubricants
Rapeseed oil (1990)	China (0.28), India (0.20), Canada (.13)	Food (0.82), inks, pharma, cosmetics
Sun/Safflower oil (90-1)	USSR (0.29), Argentina (0.17)	Food (0.90), fuel.
<i>Industrial Commodities</i>		
Aluminum (1990)	U.S. (0.22), USSR (0.12), Canada (0.09)	Transportation, containers
Burlap* (1990)	India (0.52), Bangladesh (0.30)	Fabric
Cement (1990)	China (0.18), USSR (0.12), Japan (0.07)	Construction
Copper (1990)	Chile (0.18), U.S. (0.18)	Electrical (0.75), construction
Cotton (1990-91)	China (0.24), U.S. (0.18), Uzb. (0.14)	Clothing, furnishings, medical
Lead (1990)	U.S. (0.23), Kazakhstan (0.12)	Construction, lining, batteries
Lumber	Russia (0.39), Canada (0.39)	Construction, industrial uses
Mercury (1990)	China (0.22), Russia (0.18)	Batteries, paints, dental
Nickel (1990)	USSR (0.24), Canada (0.22)	Coins, batteries, electronics
Natural Rubber (1990)	Malaysia (0.25), Thailand (0.24)	Household to industrial
Tin (1990)	China (0.19), Brazil (0.18)	Industrial uses
Wool (1990-91)	Australia (0.35), New Zealand (0.12)	Clothing/furnishing, insulation
Zinc (1990)	USSR (0.13), Japan (0.10), Can. (0.08)	Coating, alloy, batteries, medical

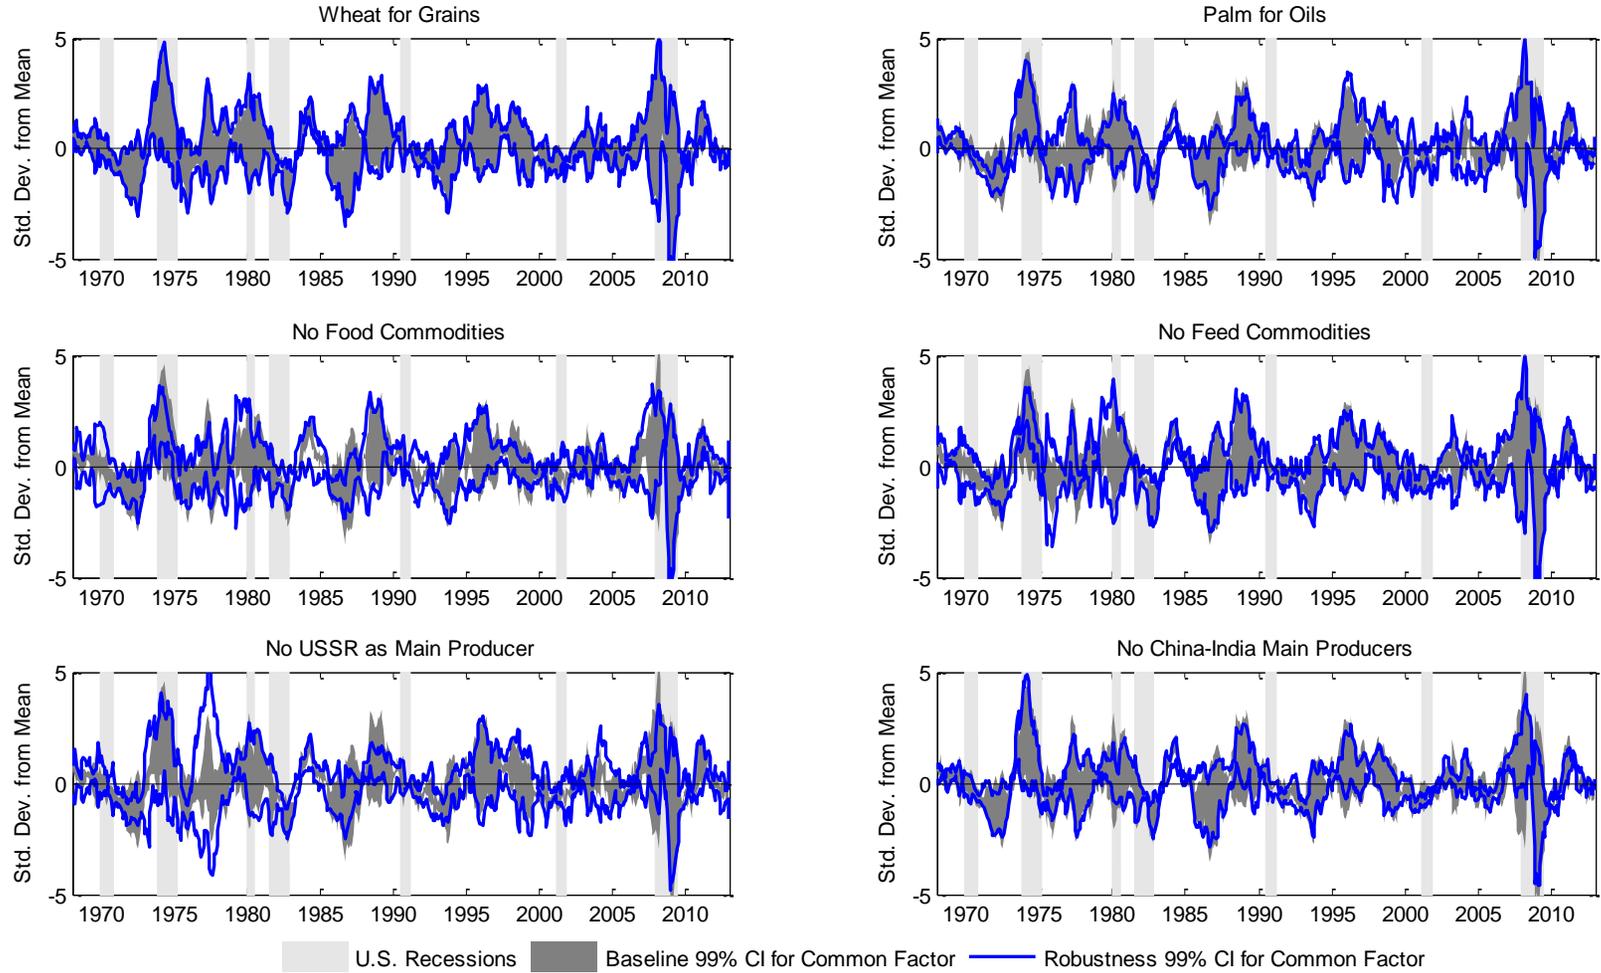
Note: The table presents information on the largest-producing countries for each type of commodity in 1990 or as available. These data comes from the CRB or the FAO (when marked with a *). The third column presents the most common uses of each commodity, as reported by the CRB (for industrials) or by the FAO in 1990 for all others.

Appendix E: Contribution of Common Factors to Individual Commodity Prices

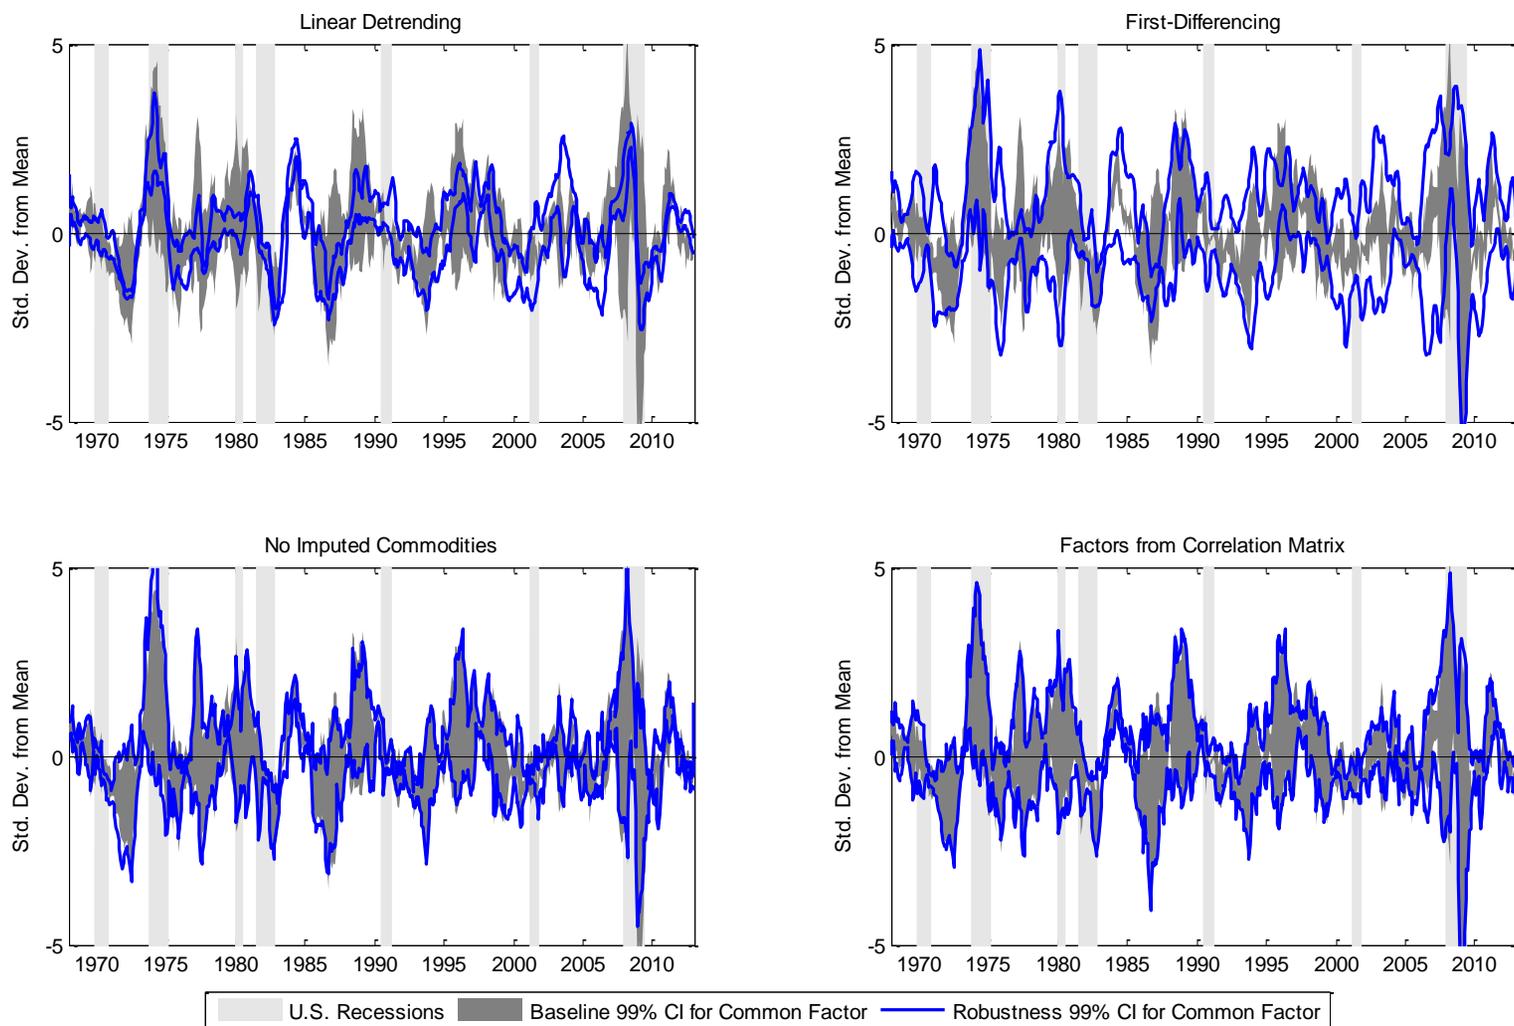
Number of Factors:	Cumulative R^2 from Common Factors				
	1	2	3	4	5
<i>Agricultural/Food</i>					
Apples	0.20	0.22	0.22	0.23	0.38
Bananas	0.34	0.37	0.43	0.43	0.63
Barley	0.62	0.73	0.78	0.86	0.86
Beef	0.74	0.77	0.77	0.85	0.85
Cocoa	0.76	0.80	0.88	0.89	0.90
Coffee	0.69	0.75	0.86	0.87	0.87
Corn	0.91	0.91	0.93	0.94	0.94
Fishmeal	0.85	0.85	0.85	0.86	0.86
Hay	0.73	0.75	0.76	0.84	0.87
Oats	0.82	0.82	0.82	0.84	0.84
Orange Juice	0.51	0.59	0.64	0.73	0.78
Onions	0.24	0.43	0.46	0.47	0.53
Pepper	0.25	0.50	0.52	0.59	0.59
Potatoes	0.54	0.55	0.64	0.64	0.69
Rice	0.87	0.87	0.89	0.89	0.89
Shrimp	0.14	0.76	0.79	0.79	0.80
Sorghums	0.90	0.90	0.93	0.93	0.93
Soybeans	0.91	0.91	0.93	0.93	0.93
Sugar	0.61	0.62	0.71	0.73	0.75
Tea	0.71	0.80	0.82	0.83	0.84
Tobacco	0.65	0.82	0.82	0.83	0.84
Wheat	0.87	0.87	0.89	0.90	0.90
<i>Oils</i>					
Coconut oil	0.71	0.71	0.71	0.71	0.79
Groundnut oil	0.75	0.75	0.78	0.83	0.86
Palm oil	0.81	0.81	0.81	0.85	0.90
Rapeseed oil	0.46	0.63	0.71	0.85	0.85
Sun/Safflower oil	0.73	0.76	0.78	0.84	0.85
<i>Industrials</i>					
Aluminum	0.62	0.62	0.68	0.78	0.79
Burlap	0.72	0.72	0.73	0.81	0.85
Cement	0.14	0.14	0.79	0.79	0.80
Copper	0.44	0.83	0.85	0.92	0.93
Cotton	0.80	0.88	0.89	0.89	0.89
Lead	0.60	0.86	0.87	0.87	0.87
Lumber	0.25	0.33	0.53	0.64	0.76
Mercury	0.25	0.49	0.51	0.73	0.77
Nickel	0.13	0.70	0.70	0.84	0.87
Rubber	0.71	0.84	0.86	0.86	0.86
Tin	0.84	0.85	0.92	0.93	0.93
Wool	0.78	0.79	0.79	0.79	0.79
Zinc	0.39	0.48	0.54	0.54	0.65

Note: The table presents the R^2 associated with the cumulative number of factors across columns for each commodity. Imputed values are not included in R^2 calculations. See section 3.2 in the text for details.

Appendix F: Robustness Checks

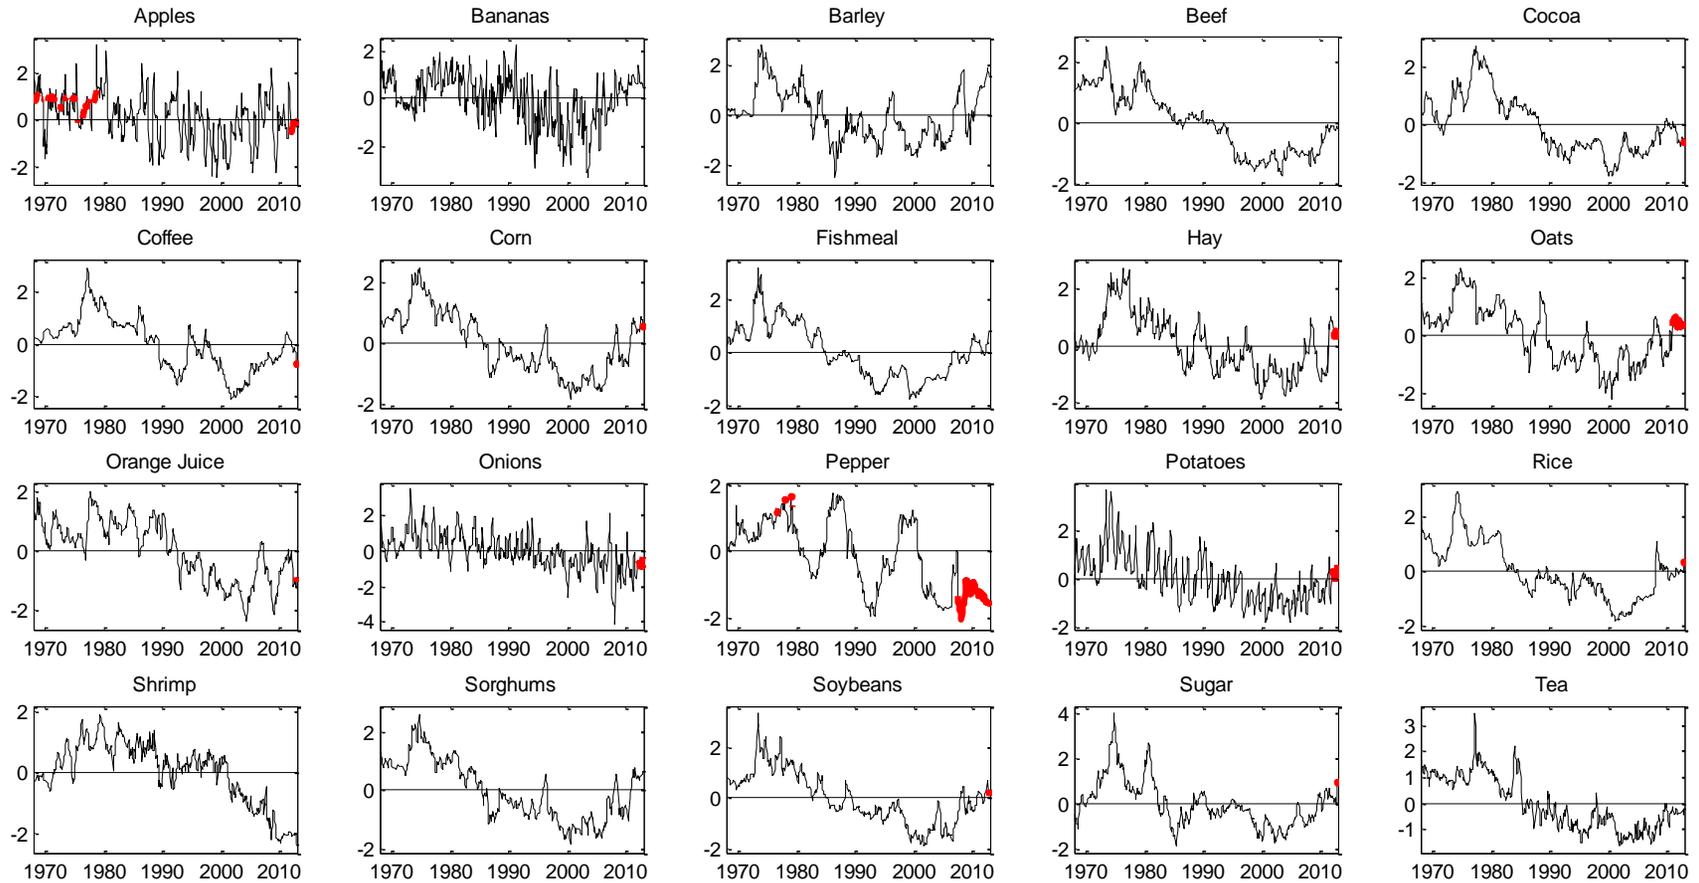


Note: The figures present the baseline 99% confidence interval for the (HP-filtered) IC factor (grey shaded area) and the 99% confidence intervals for the HP-filtered IC factor for subsets of commodities (areas between blue lines). In the top two panels, we drop barley, hay, oats, and sorghums from the cross-section of commodities (left figure) and coconut oil, peanut oil, rapeseed oil, and safflower oil (right figure). In the two middle panels, we drop all commodities for which food is the primary use (left figure) and all commodities for which feed is the primary use (right figure). In the bottom two figures, we drop all commodities for which the former USSR was the primary producer in 1990 (8 commodities) and all commodities for which China or India were primary producers (13 commodities, right figure). See section 3.3 for details.

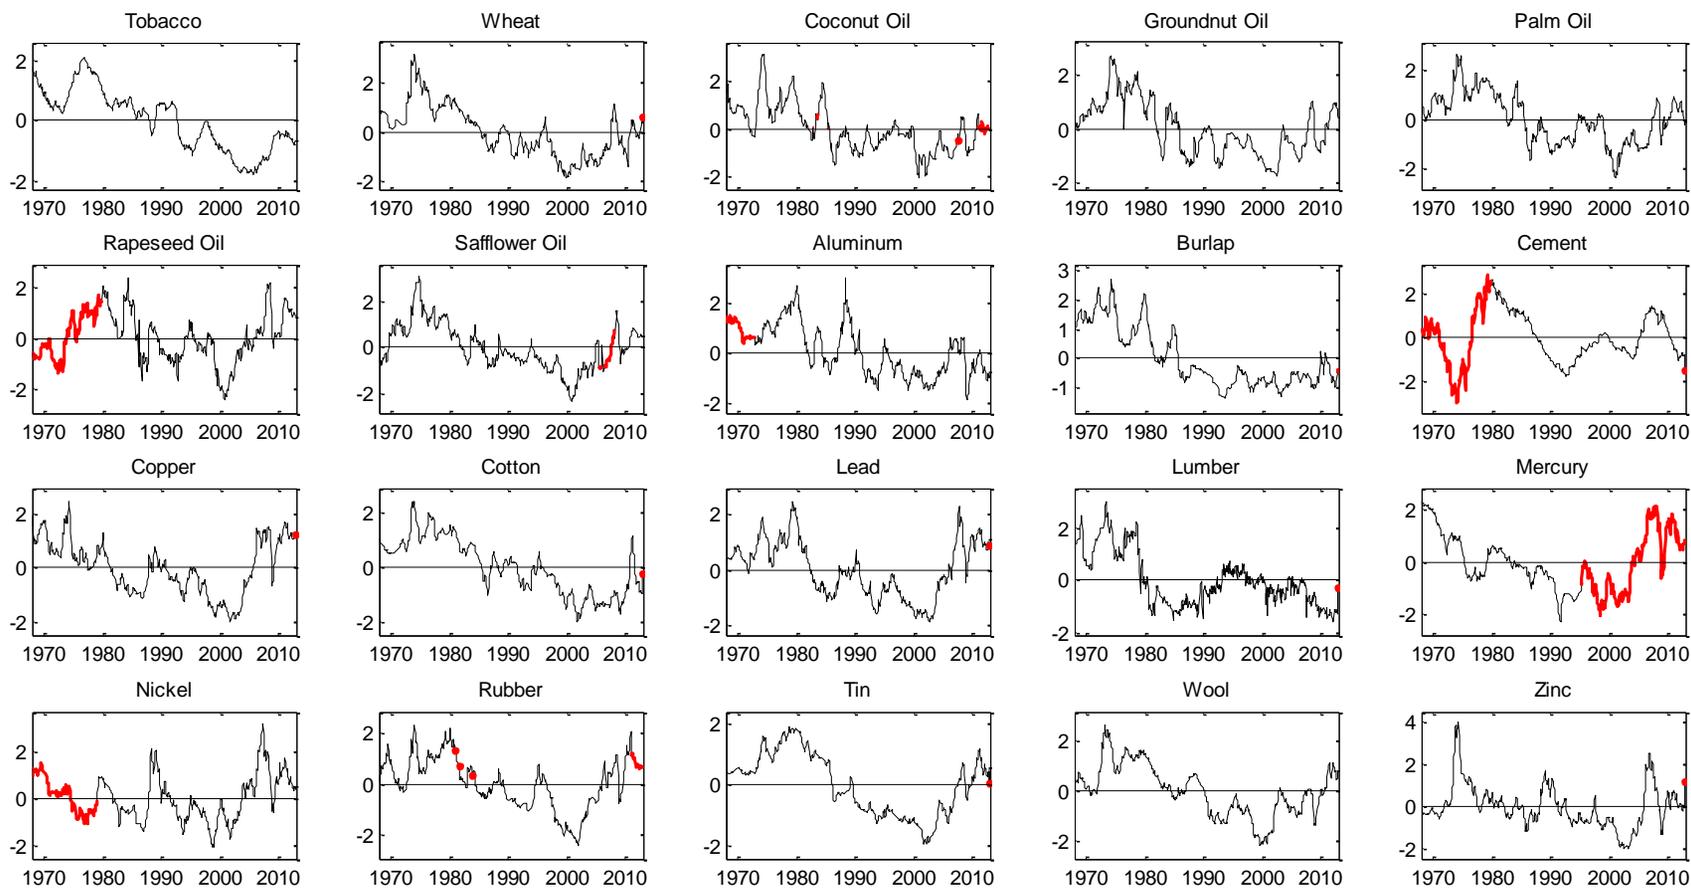


Note: The figures present the baseline 99% confidence interval for the (HP-filtered) IC factor (grey shaded area) and the 99% confidence intervals for the HP-filtered IC factor under alternative conditions (areas between blue lines). In the top left figure, we linearly detrend each real commodity price series prior to factor analysis. In the top right figure, we implement factor analysis in first-differences. In the bottom left figure, we include only commodities for which no imputation was necessary prior to 2010. In the bottom right figure, we extract factors from the correlation matrix of the cross-section of real commodity prices rather than the covariance matrix. See section 3.3 in the text for details.

Appendix G: Time Series of (Log) Real Commodity Prices and Imputed Values

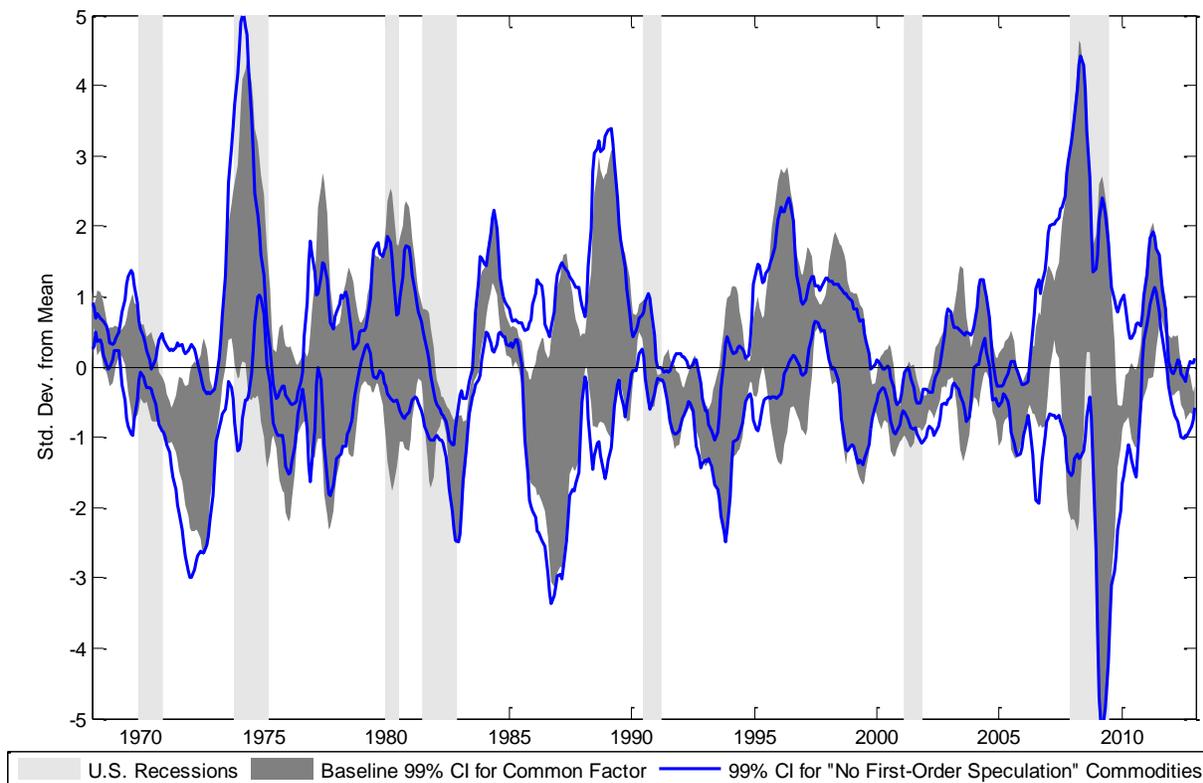


Note: The figure plots real commodity prices (black lines) and imputed values (bold red values) from the EM algorithm of Stock and Watson (2002).



Note: The figure plots real commodity prices (black lines) and imputed values (bold red values) from the EM algorithm of Stock and Watson (2002).

Appendix H: Robustness to Dropping Commodities for which the Null of “No First Order Speculation” Is Rejected



Note: The figure presents the baseline 99% confidence interval of the (HP-filtered) IC factor from the factor analysis on the full cross-section of commodities in section 3.3 using the estimated rotation parameters from GMM estimates (grey shaded area). The blue lines correspond to the 99% confidence interval for the equivalent factor using only those commodities for which we cannot reject the null of no first-order speculative price effects in Table 4. Confidence intervals are 3-month moving averages. See section 4 for details.

Appendix I: Additional Tables on Out-of-Sample Forecasting

APPENDIX TABLE I.1: RECURSIVE FORECAST ERROR DIAGNOSTICS FOR REAL COMMODITY PRICES

	$h = 1$	$h = 3$	$h = 6$	$h = 12$	Forecast Evaluation Period
<i>Agr./Food Commodities</i>					
Apples	0.886	0.738	0.598	0.703	1982:11-2011:12
Bananas	0.898	0.726	0.659	0.929	1968:1-2013:1
Barley	0.973	0.975	1.002	0.986	1968:1-2013:1
Beef	1.138	1.261	1.359	1.367	1968:1-2013:1
Cocoa	0.933	1.020	1.039	1.032	1968:1-2012:12
Coffee	0.959	0.986	1.072	1.088	1968:1-2012:12
Corn	0.904	0.943	0.924	0.910	1968:1-2012:12
Fishmeal	1.025	1.167	1.108	1.078	1968:1-2013:1
Hay	1.026	0.953	0.909	0.878	1968:1-2013:3
Oats	0.932	0.965	0.937	0.955	1968:1-2010:11
Orange juice	0.967	1.023	1.045	0.967	1971:2-2012:10
Onions	0.886	0.762	0.618	0.623	1968:1-2011:12
Pepper	0.906	1.073	1.197	1.375	1983:6-2007:6
Potatoes	0.816	0.799	0.701	0.947	1968:1-2011:12
Rice	0.873	0.961	1.025	1.115	1968:1-2012:12
Shrimp	1.029	1.100	1.136	1.256	1968:1-2013:1
Sorghum	0.930	0.997	0.988	0.982	1968:1-2013:1
Soybeans	0.936	1.016	1.053	1.078	1968:1-2012:9
Sugar	0.937	0.999	1.025	1.038	1968:1-2012:12
Tea	1.042	1.193	1.237	1.313	1968:1-2013:1
Tobacco	0.894	0.912	0.904	0.873	1968:1-2013:1
Wheat	0.970	1.049	0.997	0.947	1968:1-2012:12
<i>Oils</i>					
Coconut	0.988	0.984	0.964	0.914	1989:7-2010:12
Groundnut	0.993	0.937	0.893	0.773	1968:1-2013:1
Palm	0.915	1.071	1.072	1.036	1968:1-2013:1
Rapeseed	1.030	0.992	1.028	0.963	1984:1-2013:1
Sunflower	0.946	1.028	1.057	1.106	1968:1-2005:6
<i>Industrial Commodities</i>					
Aluminum	0.999	1.004	1.058	1.155	1977:1-2013:1
Burlap	0.880	1.050	1.068	1.054	1968:1-2012:9
Cement	1.028	1.075	1.148	1.200	1984:1-2012:12
Copper	0.887	1.006	1.072	1.104	1968:1-2012:12
Cotton	0.762	0.927	1.000	0.950	1968:1-2012:12
Lead	0.964	1.034	1.084	1.092	1968:1-2012:12
Lumber	1.005	1.127	1.149	1.172	1968:1-2012:12
Mercury	0.884	1.077	1.198	1.419	1968:1-1995:3
Nickel	0.955	1.157	1.444	2.422	1983:3-2013:1
Rubber	0.952	0.989	1.054	1.117	1968:1-2010:12
Tin	0.915	0.922	0.991	1.068	1968:1-2012:12
Wool	0.967	0.987	1.034	1.096	1968:1-2013:1
Zinc	0.936	1.030	1.101	1.339	1968:1-2012:12

Notes: The forecast evaluation period depends on the commodity. It begins either in 1968:1 or at the earliest date such that the initial estimation window contains at least 48 observations. The maximum length of the recursive sample is restricted by the end of the data and the forecast horizon. All forecasts are obtained from a bivariate VAR that includes the level of the real commodity price and the first principal component extracted from the cross-section of real commodity prices. The lag length of the VAR is chosen recursively using the BIC. The MSPE of the VAR forecast is expressed as a ratio relative to that of the no-change forecast. Entries smaller than 1 indicate that the VAR forecast is superior to the no-change forecast and are shown in boldface.

APPENDIX TABLE I.2: RECURSIVE FORECAST ERROR DIAGNOSTICS FOR REAL COMMODITY PRICES

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
<i>Agr./Food Commodities</i>				
Bananas	0.880	0.698	0.625	0.842
Barley	0.956	0.955	0.994	0.931
Beef	1.048	1.207	1.475	1.787
Cocoa	0.972	0.996	1.002	0.964
Coffee	0.963	0.948	0.987	0.954
Corn	0.874	0.870	0.838	0.769
Fishmeal	0.968	1.104	1.199	1.319
Hay	0.951	0.829	0.697	0.588
Rice	0.847	0.885	0.838	0.758
Shrimp	1.030	1.079	1.081	1.187
Sorghum	0.908	0.911	0.863	0.813
Sugar	0.942	1.010	0.994	0.922
Tea	0.958	0.980	0.946	0.941
Tobacco	0.858	0.876	0.831	0.726
Wheat	0.921	0.919	0.826	0.750
<i>Oils</i>				
Groundnut	0.877	0.891	0.825	0.679
Palm	0.914	1.088	1.042	0.962
Rapeseed	1.008	1.007	1.077	1.006
<i>Industrial Commodities</i>				
Aluminum	1.000	0.985	1.000	1.020
Cement	1.023	1.057	1.128	1.190
Copper	0.865	0.980	1.015	1.063
Cotton	0.784	0.913	1.019	0.972
Lead	0.995	1.050	1.080	1.123
Lumber	1.040	1.052	1.083	1.242
Nickel	0.948	1.147	1.453	2.504
Tin	0.893	0.889	0.947	0.969
Wool	0.924	0.961	1.015	1.079
Zinc	0.923	0.960	0.929	0.872

Notes: The forecast evaluation period is 1984:1-2012:12. The initial estimation window begins at the earliest date such that it contains at least 48 observations. The maximum length of the recursive sample is restricted by the end of the data and the forecast horizon. All forecasts are obtained from a bivariate VAR that includes the level of the real commodity price and the first principal component extracted from the cross-section of real commodity prices. The lag length of the VAR is chosen recursively using the BIC. The MSPE of the VAR forecast is expressed as a ratio relative to that of the no-change forecast. Entries smaller than 1 indicate that the VAR forecast is superior to the no-change forecast and are shown in boldface.

APPENDIX TABLE I.3: SUMMARY OF RECURSIVE FORECAST ACCURACY DIAGNOSTICS FOR THE REAL PRICE OF OIL

<u>Forecast Evaluation Period: 1984:1-2012:8</u>				
	BIC		12 lags	
	<u>FAVAR</u>	<u>VAR</u>	<u>FAVAR</u>	<u>VAR</u>
1 month	0.790	0.825	0.858	0.843
3 months	0.947	1.047	1.037	1.028
6 months	1.111	1.268	1.224	1.206
12 months	1.308	1.501	1.419	1.427

<u>Forecast Evaluation Period: 1992:1-2012:8</u>				
	BIC		12 lags	
	<u>FAVAR</u>	<u>VAR</u>	<u>FAVAR</u>	<u>VAR</u>
1 month	0.832	0.846	0.904	0.857
3 months	0.980	1.016	1.105	0.960
6 months	1.182	1.174	1.329	1.115
12 months	1.459	1.336	1.524	1.172

Notes: The oil-market data are from Baumeister and Kilian (2012) and span the period 1973:1-2012:8. “FAVAR” refers to the bivariate factor-augmented VAR forecasting model that includes the commodity-price factor and the real price of oil. “VAR” refers to the four-variable oil-market VAR, as described in the text. “BIC” indicates that the lag length is chosen recursively using the BIC. “12 lags” indicates that the lag length is fixed at 12. The MSPE ratios of the real-oil price forecasts are computed relative to the benchmark no-change forecast. Entries smaller than 1 indicate that the model-based forecast is superior to the no-change forecast and are shown in boldface.