IMF Working Paper

The Comovement in Commodity Prices: Sources and Implications

Ron Alquist and Olivier Coibion

IMF Working Paper

Research

The Comovement in Commodity Prices: Sources and Implications

Prepared by Ron Alquist and Olivier Coibion

Authorized for distribution by Prakash Loungani

Lwpg 2013

Abstract

We present a simple macroeconomic model with a continuum of primary commodities used in the production of the final good, such that the real prices of commodities have a factor structure. One factor captures the combined contribution of all aggregate shocks which have no direct effects on commodity markets other than through general equilibrium effects on output, while other factors represent direct commodity shocks. Thus, the factor structure provides a decomposition of underlying structural shocks. The theory also provides guidance on how empirical factors can be rotated to identify the structural factors. We apply factor analysis and the identification conditions implied by the model to a cross-section of real non-energy commodity prices. The theoretical restrictions implied by the model are consistent with the data and thus yield a structural interpretation of the common factors in commodity prices. The analysis suggests that commodity-related shocks have generally played a limited role in global business cycle fluctuations.

1   Introduction

Macroeconomists have long struggled to identify exogenous sources of business cycle fluctuations. Thirty years after Kydland and Prescott (1982), the notion that negative technology shocks can explain recessions remains difficult for many to accept. Other potential sources of business cycle fluctuations also generate raised eyebrows: negative shocks to the discount factor, shocks to households' willingness to supply labor, markup shocks, and so on. But one possible source of exogenous economic volatility that raises fewer questions than the rest is commodity-price shocks. From droughts in the American Midwest to labor strikes in the mines of South America to geopolitical instability in the Middle East, large commodity-price fluctuations seem easier to accept as exogenous changes because they frequently can be reconciled with visible news events. Furthermore, large commodity-price movements have often coincided with periods of macroeconomic volatility, lending support to arguments that attribute causality from commodity-price shocks to aggregate macroeconomic fluctuations. The stagflation of the 1970s, for example, is frequently attributed to oil and food price shocks of that era (Blinder and Rudd 2012). Hamilton (1983) noted that almost all U.S. recessions had been preceded by oil price increases, a feature which continued with the recent oil- and commodity-price boom immediately preceding the Great Recession.

Despite this, the direction of causality is difficult to ascertain. Bosworth and Lawrence's (1982) efforts to separate supply and demand forces across many commodities in the 1970s did not point to clear evidence of supply-side factors driving commodity-price movements, Barsky and Kilian (2002) argued that monetary forces caused both the Great Inflation and the commodity price booms of the 1970s, and Kilian (2009) found that much of the oil price run-up in the 2000s could be attributed to rising global demand. Indeed, a growing body of work focusing specifically on oil markets increasingly emphasizes the role of demand forces in accounting for historical oil-price fluctuations and calls into question the role of commodity-price shocks as a central driver of global business cycles. But comparatively little effort has been devoted to explaining common price movements among commodities, which were a defining feature of both the 1970s and the 2000s commodity booms.

We build on this literature in several ways. First, we develop a simple general equilibrium macroeconomic model with commodities that yields a tractable factor structure for real commodity prices. This factor structure decomposes each commodity's price into three distinct components. The first component captures idiosyncratic price movements. The second component captures shocks (such as aggregate productivity shocks and shocks to labor supply) that are not directly related to commodity demand and supply. We refer to this factor as the indirect aggregate common (IAC) factor. The third component, which we label the direct aggregate common (DAC) factor(s), represents the shocks that directly affect commodity supply and demand, holding aggregate output constant.

The IAC factor is of particular interest. First, this factor reflects the combined contribution of all "indirect" aggregate shocks, i.e. shocks for which the effects on commodity prices reflect only general equilibrium changes in aggregate output. Thus, this factor can be thought of as capturing the effects of "non-commodity shocks". Second, the IAC factor implied by the theory is the level of aggregate output that would have occurred in the absence of "direct" commodity shocks, i.e. it represents the

counterfactual history of global activity driven only by non-commodity shocks. Hence, the factor structure of commodity prices from the model provides a mechanism for decomposing the endogenous responses of both commodity prices and global activity into components reflecting commodity-related (or "direct") shocks from other ("indirect") shocks.

The intuition behind this structural decomposition is straightforward. Most aggregate shocks should affect individual commodity prices only by changing the total demand for commodities or through supply-side effects stemming from income effects on common inputs, i.e. their effects on commodity prices are summarized by their effects on aggregate output. As a result, conditional on the effects of aggregate shocks on aggregate production, their implications for commodity prices are identical across shocks. Thus, these "indirect" shocks will all generate a single pattern of comovement among commodities which can be aggregated into a single factor within a factor decomposition of historical commodity price movements. In contrast, commodity-related shocks directly affect the supply and demand for commodities above and beyond their general equilibrium effects on aggregate output. The additional shifts in supply and/or demand curves associated with direct commodity shocks generate a different pattern of commodity comovement than indirect shocks and this heterogeneity in comovement will be captured by the factor structure. Thus, the fact that all indirect shocks induce a common comovement in commodity prices whereas direct shocks generate different comovement patterns underlies the decomposition of shocks into direct and indirect factors within the factor structure.

Applying factor methods to commodity prices is common because these methods are well-suited to characterizing comovement across many variables (e.g., Byrne, Fazio and Fiess 2011, Gospodinov and Ng 2013, West and Wong 2012). But a key challenge in this literature is the inability to interpret the factors in a structural manner because they are only identified up to a rotation. Our second contribution is to show how one can use the theoretical predictions of the model to identify the rotation matrix that maps empirically estimated factors into the structural factors implied by the model, including the indirect aggregate common factor. Specifically, the theory provides two potential identification strategies. The first is to impose orthogonality conditions on structural factors with respect to a set of instruments. For example, the IAC factor should be orthogonal to commodity-related shocks such as exogenous oil price shocks (since energy is a common input in the production and distribution of commodities). The parameters of the rotation matrix can then be estimated by GMM based on these orthogonality conditions. A second strategy is to impose sign restrictions predicted by the theory, as in Uhlig (2005). Each commodity should have a positive loading on the IAC factor because higher global activity both increases the demand for all commodities and potentially lowers their supply if income effects on inputs into the production process of commodities are present. Both of these forces induce commodity prices to rise with global activity, so commodity loadings on the IAC factor should all be positive. Thus, one can recover a structural interpretation from the empirical factor analysis, guided by the theory.

Our third contribution is to apply this methodology to a cross-section of historical commodity prices. We select commodities that conform to the structure and assumptions of the model. We thus exclude commodities that are vertically integrated. For example, we include sorghums in the data set but because they are used as feed, we exclude poultry and eggs. By the same logic, we exclude all energy

commodities because they are used as inputs into the production and distribution of commodities. Second, we exclude commodities that are commonly used for financial speculation, such as precious metals. Third, we exclude commodities whose production is a derivative of the production of another commodity (e.g., many minerals). This leaves a cross-section of forty commodities, ranging from food and agricultural products (e.g., wheat and coffee), oils (e.g., coconut oil), and industrials (e.g., copper and rubber). We document the geographic variation in the production of these commodities as well as the diversity (and sometimes lack thereof) in their uses.

We extract the underlying factors from this cross-section of commodities and implement our identification procedure to recover the structural factors. The baseline identification strategy is to estimate the rotation parameter by GMM, using as moment conditions the orthogonality of the IAC factor with respect to the OPEC oil production shocks from Kilian (2008). The overidentification conditions cannot be rejected and the implied rotation matrix satisfies the sign restrictions of factor loadings implied by the theory. A notable result is that we cannot reject the null that the rotation matrix is equal to the identity matrix, such that one could directly interpret the unrotated factors in a structural manner. This result is insensitive to several robustness checks, including the subset of commodities used, econometric issues associated with the identification of the rotation matrix, the initial factor decomposition, and the treatment of trends. We also show that the set of rotations selected using the alternative sign restriction approach lie within the 99% confidence interval of the GMM estimate of the rotation matrix. Hence, the two identification strategies yield remarkably consistent results and suggest that our baseline results are robust to the identification procedure and the choice of instruments.

The IAC factor accounts for approximately 60-70% of the variance in real commodity prices overall and much of the historical changes in commodity prices since the early 1970s. However, direct commodity shocks have also played a nontrivial role in accounting for some commodity price movements, including during 1979-1980, in the late 1980s, as well as during the run-up in commodity prices in the 2000s and their subsequent decline in 2008-2009. Nonetheless, the IAC accounts for much more of the historical variation in commodity prices. Similar results obtain when decomposing historical changes in global economic activity. While there have been periods in which direct commodity price shocks contributed to changes in global activity (such as in 1974-1975 or during the Great Recession), we again find that the IAC has contributed much more to historical global activity, even in those periods when commodity price shocks are often cited as a prominent source of economic volatility.

A potential caveat to our results is that the aggregation result underlying our suggested factor structure can be broken in the presence of storage. Specifically, if indirect shocks have different implications from one another for the path of expected prices, as would be the case if their persistence is very different, then speculators would pursue different inventory management strategies for each indirect shock. In this case, the effect of an aggregate shock on output would no longer be a sufficient statistic for its effect on commodity prices and the aggregation of all indirect shocks into a single factor would no longer follow. There are at least three reasons to be skeptical of this argument. First, the fact that commodity prices are well-characterized empirically by very few factors strongly suggests that some form of aggregation does indeed hold. Second, for most commodities in our sample, we cannot reject the

null that storage motives have only second-order effects on prices, based on the average net purchases of the storage sector implied by global consumption and production data for commodities. Third, in the presence of storage, interest rates would have a direct effect on commodity markets and therefore, by the logic of the factor structure, should not be incorporated into the IAC. Using a measure of monetary policy shocks constructed from a time-varying coefficients Taylor rule estimated with real-time forecasts of the Federal Reserve, we incorporate our IAC factor into a monetary VAR and find that we can reject the hypothesis that monetary shocks have no effect on the IAC factor. In fact, we document that monetary policy shocks can account for much of the variation in the IAC factor during the 1970s, consistent with the monetary interpretation of the Great Inflation suggested by Barsky and Kilian (2002). By contrast, exogenous monetary policy shocks have contributed little to global economic volatility since the early 1980s, including during the run-up in commodity prices of the mid- to late-2000s. Thus, all three results suggest that our factor structure is an adequate and succinct representation of the historical determinants of commodity prices.

Finally, we assess whether the IAC factor can be used to forecast commodities in real-time. Using pseudo out-of-sample forecasts, we find that a bivariate VAR that includes each commodity's price and the 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 IAC 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 to in size those obtained using oil-market VARs in Baumeister and Kilian (2012) and Alquist et al. (forthcoming). But unlike the oil-market VARs that use production and inventory data that are usually unavailable for most non-oil commodities, our approach relies only on a cross-section of commodity prices that can readily be tracked and updated at high frequencies. Thus, the factor structure of commodity prices can be used not only to provide a structural interpretation of commodity-price movements but also to improve the quality of forecasts for policy-makers and practitioners.

The structure of the paper is as follows. Section 2 presents our macroeconomic model with commodity prices and shows how the latter will have a factor structure. The section also shows how the structure of the model permits the econometrician to recover the structural factors from typical factor decompositions of a cross-section of commodity prices. 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 aggregate common factor in a pseudo out-of-sample forecasting exercise. Section 6 concludes.

2 The Sources of Commodity Price Comovement: Theory

In this section, we present a simple model that characterizes the sources of commodity-price comovement. In particular, we show that the model yields a convenient and tractable factor structure for a cross-section of commodity prices in that the factors have a structural interpretation.

2.1 Model of commodity prices

Our 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 "land" supplied to each commodity sector ($L^s(j)$) as follows

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

where β is the discount factor. 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 therefore

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

where T_t represent payments from the ownership of firms.

Assuming that the household takes all prices as given, its first-order conditions are

$$\varphi_n C_t^\sigma N_t^{s\eta^{-1}} = e^{\varepsilon_t^n} W_t / P_t \tag{1}$$

$$\varphi_L C_t^\sigma L_t^s(j)^{v^{-1}} = e^{\varepsilon_t^L} R_t^L(j) / P_t \tag{2}$$

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

This setup is standard, with the exception of the "land" provided by the household. This variable is an input into the production process for primary commodities and can be interpreted in several ways. Referring to this input as "land", for example, follows from the notion that the use of land yields direct benefits to the household (and hence is included in the utility function) but that it can be provided to commodity producers in exchange for a rental payment. This conveniently yields a traditional supply curve for this input. One could equivalently interpret it as a form of labor that cannot be reallocated across sectors. In this case one could 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 mining and agricultural sectors. 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 utilized 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.

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

$$Q_t(j) = A_t(j)L_t^d(j)^{\alpha_j} \tag{4}$$

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 $P_t(j)$, the price of commodity 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:

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

We assume WLOG 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

$$L_t^s(j) = L_t^d(j) \tag{6}$$

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 per a Dixit-Stiglitz aggregator

$$Q_t^c = \left(\int_0^1 Y_t^{j\frac{\theta_c-1}{\theta_c}} dj\right)^{\frac{\theta_c}{\theta_c-1}} \tag{7}$$

which yields a demand for each commodity j of

$$P_t(j)/P_t^c = (Y_t(j)/Q_t^C)^{-1/\theta_c} \tag{8}$$

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

$$Q_t(j) = Y_t(j). \tag{9}$$

Note that this setup implicitly assumes that no storage of commodities takes place since all commodities produced must be used contemporaneously. We discuss the motivation for this assumption, as well as 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

$$Y_t = A_t Y_t^{C^{\alpha_t}} N_t^{d^{1-\alpha_t}} \tag{10}$$

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

$$\alpha_t = (P_t^C/P_t)(Y_t^C/Y_t) \tag{11}$$

$$1 - \alpha_t = (W_t/P_t)(N_t^d/Y_t) \tag{12}$$

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

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., $p_t(j) \equiv \log P_t(j)/P_t - \log(\overline{P(j)/P})$), the first-order conditions from the household's problem are

$$\sigma y_t + \frac{1}{\eta} n_t = w_t + \varepsilon_t^n \tag{13}$$

$$\sigma y_t + \frac{1}{v} l_t(j) = r_t^L(j) + \varepsilon_t^L \tag{14}$$

$$y_t = E_t \left[y_{t+1} - \frac{1}{\sigma} r_t \right] \tag{15}$$

where we've 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

$$r_t^L(j) = p_t(j) + a_t(j) - (1 - \alpha_j) l_t(j) \tag{16}$$

$$y_t(j) = a_t(j) + \alpha_j l_t(j) \tag{17}$$

where we've 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

$$y_{c,t} = \int_0^1 y_t(j) dj \tag{18}$$

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

Finally, letting α be the steady-state value of α_t, the final goods sector follows

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

$$p_{c,t} = y_t - y_{c,t} + \check{\alpha}_t \tag{21}$$

$$w_t = y_t - n_t - \frac{\alpha}{1-\alpha} \check{\alpha}_t. \tag{22}$$

where $\varphi_\alpha \equiv \alpha\left(\ln \overline{Y_t^c} - \ln \overline{N_t}\right)$.

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

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

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

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

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.[1]

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

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

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 RHS corresponds to 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

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

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 will comove positively with total production.

Equilibrium in the labor market is given by

$$n_t = \frac{1-\sigma}{1+\eta^{-1}} y_t + \frac{1}{1+\eta^{-1}} \varepsilon_t^n - \frac{1}{1+\eta^{-1}} \check{\alpha}_t \tag{27}$$

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

$$y_t = \varphi_y[a_t + \kappa_L \varepsilon_t^L + \kappa_n \varepsilon_t^n + \kappa_v v_t + \kappa_\alpha \check{\alpha}_t] \tag{28}$$

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

---

[1] The rescaling of the commodity-specific productivity shock ensures that a 1% increase in TFP 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.

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 ($\breve{\alpha}_t$) depends on specific parameter values.

2.2   Comovement in Commodity Prices

For simplicity, we assume that productivity shocks to each commodity sector have an idiosyncratic component and a common component such that $v_t(j) = v_t^a + v_t^j(j)$, which implies that the average across commodities is simply $v_t = v_t^a$. Idiosyncratic components are orthogonal across commodity sectors, such that $E[v_t(j)v_t(k)] = 0 \; \forall j \neq k$ such that $v_t = 0$.

We then consider the determinants of individual commodity prices. First, the supply of commodity j follows from equations (14), (16) and (17) and is given by

$$p_t(j) = \varepsilon_j y_t(j) - \frac{(1+\varepsilon_j \theta_c)}{\theta}\left(v_t^a + v_t^j(j)\right) + \sigma y_t - \varepsilon_t^L \tag{29}$$

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 shocks that affect aggregate production in the model will 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 will increase whenever its productivity level rises, which can reflect common productivity shocks (v_t^a) or idiosyncratic shocks ($v_t^j(j)$). Finally, shocks to the household's willingness to supply land to the commodity sector will directly affect the supply curve. Thus, we can write the supply curve of commodity j more succinctly as

$$p_t(j) = S_j \left(y_t(j); \; y_t(a_t, \varepsilon_t^n, \varepsilon_t^L, \breve{\alpha}_t, v_t^a); \; v_t^a, \varepsilon_t^L, v_t^j(j)\right) \tag{30}$$

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 (19) with (21) and (25) yielding

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

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 will result in an equal upward or downward shift in the demand for every commodity. Thus, all shocks in the model other than idiosyncratic shocks are demand shocks as well as supply shocks. But in addition to these general equilibrium shifts in commodity demand, the demand for commodity j will rise with changes in the relative demand for commodities ($\breve{\alpha}_t$), holding aggregate output constant. It will also shift, 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 fact that they affect all commodities implies that they affect equilibrium prices and quantities of the intermediate commodity bundle, and therefore affect the demand for each

commodity through the CES structure. We can again write the demand curve of commodity j more succinctly as

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

to highlight the fact that some shocks affect the supply of commodity j indirectly through general equilibrium effects on output, some shocks can 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 break this impasse. 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 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 will tend to raise prices of all commodities. While the magnitude of the effect will differ across commodities depending on the slopes of their respective supply curves (which depend on α_j's), there will necessarily be positive comovement implied by such shocks. This is illustrated visually in Panel A of Figure 1, which shows the price implications of an increase in aggregate productivity for a commodity with elastic supply $S_E(y(a))$ and one with inelastic supply $S_{INE}(y(a))$. The red lines with arrows ($P_E(a)$ and $P_{INE}(a)$) then show the set of possible prices for each commodity for different values of a, and therefore the implied comovement, which can result from productivity shocks.

Importantly, *any* shock that only affects commodity prices through its effects on aggregate output will induce the same relative comovement of commodity prices as productivity shocks. In our model, shocks to the household's willingness to supply labor also affect commodity prices only through y_t and therefore will deliver the exact same pattern of comovement among commodities as an aggregate productivity shock. By contrast, any shock that directly (i.e., holding aggregate output constant) affects the supply and/or demand of a commodity will induce a different comovement across commodities. This is illustrated in Panel B of Figure for the case of an increase in the relative demand for commodities $\check{\alpha}_t$ that is then assumed to lower aggregate output. The rise in $\check{\alpha}_t$ directly shifts up the demand for each commodity, holding output constant, then shifts both the supply and demand curves through the effect of aggregate output on each. As a result of these shifts, the price of the inelastically supplied commodity rises while that of the elastically supplied commodity falls. The red lines with arrows ($P_E(\check{\alpha}_t)$ and $P_{INE}(\check{\alpha}_t)$) again trace out all possible price paths for each commodity conditional on changes in $\check{\alpha}_t$. The fact that this change in relative demand leads to negative comovement is meant only for illustrative purposes, since whether the price of the elastically supplied good rises or falls will depend on parameter values. The important point is that shocks that have both direct and indirect effects on commodity markets will lead to different comovement among commodities than those shocks that have only indirect effects.

2.3 The Factor Structure in Commodity Prices

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

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

$$- \frac{1}{\theta_c}\left(1 + \varepsilon_j \theta_c + \frac{\varepsilon_j \theta_c (\theta_c-1)}{1+(\theta_c-1)\varphi}\right) v_t^a + \frac{\varepsilon_j \theta_c}{1+\varphi(\theta_c-1)} \breve{\alpha}_t - \frac{1}{\theta_c}\left(1 + \varepsilon_j \theta_c\right) v_t^j(j). \tag{33}$$

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

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

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

$$p_t(j) = \underbrace{\lambda_j^y y_t^{nc}(a_t, \varepsilon_t^n)}_{\text{indirect } (IAC)} + \underbrace{\lambda_j^L \varepsilon_t^L + \lambda_j^v v_t^a + \lambda_j^\alpha \breve{\alpha}_t}_{\text{direct } (DAC)} - \underbrace{\frac{1}{\theta_c} v_t^j(j)}_{\text{idiosyncratic}} \tag{34}$$

where $\lambda_j^y = (1 + \theta \varepsilon_j)^{-1}\left[\sigma + \frac{\varepsilon_j \theta(1+(\theta-1)\sigma\varphi)}{1+(\theta-1)\varphi}\right]$, $\lambda_j^L = \varphi_y \kappa_L \lambda_j^y - \left[\frac{1}{1+\varepsilon_j \theta_c} + \frac{\varepsilon_j \theta_c}{1+\varepsilon_j \theta_c}(\theta_c - 1)\right]$, $\lambda_j^v = \varphi_y \kappa_v \lambda_j^y - \left[\frac{1}{\theta_c} + \frac{\varepsilon_j \theta_c}{1+\varepsilon_j \theta_c}\left(\frac{\varphi(\theta_c-1)}{\varphi\theta}\right)\right]$, and $\lambda_j^\alpha = \varphi_y \kappa_\alpha \lambda_j^y + \frac{\varepsilon_j \theta_c}{1+\varepsilon_j \theta_c}\left(\frac{1}{1+\varphi(\theta_c-1)}\right)$.

Equation (34) provides a factor structure for real commodity prices with three distinct and orthogonal components. The last term on the RHS reflects idiosyncratic shocks to commodity j that have no aggregate real effects. The second term on the RHS consists of a factor for each shock that has direct effects on commodity markets, e.g., that shifts the supply and/or demand for commodities holding aggregate output constant. We therefore refer to these factors as "direct aggregate commodity" (DAC) factors. In this specific 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 separately as a 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 shocks have both direct and indirect effects on the market for commodity j, there is in general no guarantee that their respective loadings will have the same signs. For example, a common shock can raise the price of one commodity and lower the price of another, as illustrated in Panel 2 of Figure 1.

The most interesting component of the factor structure however is the first term on the RHS, 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. This common factor is the "indirect aggregate commodity" (IAC) factor. It captures the fact that, because some shocks will affect commodity markets only through changes in aggregate output, they will all have identical implications for commodity prices conditional on the size of their effect on aggregate output. As a result, they will be integrated into 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 IAC factor. Another notable characteristic of the IAC is that, unlike for DAC factors, the loadings on this factor must all be positive ($\lambda_j^y > 0 \ \forall j$). This reflects the fact that any of the shocks incorporated in the IAC factor raise commodity demand when the shock is expansionary and simultaneously restrict the commodity supply through income effects, with both effects unambiguously pushing commodity prices up. Finally, in the absence of income effects on the common input into the production of commodities, the IAC could be interpreted as capturing exogenously-driven global demand for commodities. In short, this factor decomposition can provide a novel way to separate causality in the presence of simultaneously determined prices and production levels.

2.4   Recovering the Structural Factors

A key limitation of factor structures is that, empirically, factors are identified only up to a rotation. For example, if one estimated a factor structure on commodity prices determined by (34), one could not directly associate the extracted factors with the structural interpretation suggested by (34). However, the theory developed in this section has implications that can be used to identify the unique rotation consistent with those predictions, and therefore allows us to recover the structural factors driving commodity prices.

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

$$X_t - \mu = LF_t + \varepsilon_t$$

where μ is vector of means, 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 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

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

As a result, an empirical estimate of the factors underlying X_t will in general not recover the structural factors F_t but rather some rotation $\tilde{F}_t$.

However, the theory above provides additional restrictions on the factor structure that can be used to identify the underlying structural factors. For simplicity (as well as the fact that empirically this will be the most relevant case), consider the factor structure of equation (34) in section 2.3 in which real commodity prices reflect two 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^n), so $F_t = [F_t^1 \ F_t^2]' = [y_t^n \ \varepsilon_t^L]'$. A factor decomposition of commodity prices would yield some rotation of these factors $\tilde{F}_t$ such that

$$F_t = T'\tilde{F}_t = \begin{bmatrix} t_{11} & t_{12} \\ t_{21} & t_{22} \end{bmatrix}'[\tilde{F}_t^1 \ \tilde{F}_t^2]' = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}[\tilde{F}_t^1 \ \tilde{F}_t^2]' \qquad (35)$$

where the second equality reflects the properties of rotation matrices. Finding the rotation matrix that recovers the structural factors therefore corresponds to identifying the parameter θ.

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

$$E[y_t^n z_t] = E\big[(\tilde{F}_t^1 \cos\theta + \tilde{F}_t^2 \sin\theta)z_t\big] = 0. \tag{36}$$

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

$$J(\theta) = E[y_t^{nc}z_t]WE[y_t^{nc}z_t]' \tag{37}$$

where W is a weighting matrix such that $\hat{\theta} = argmin\, J(\theta)$. Letting W be the inverse of the variance-covariance matrix associated with the moment conditions, standard GMM asymptotic results would apply including standard errors for θ and tests of the over-identifying conditions assuming that N and T are 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 unrelated to the commodity sector such as the productivity shocks or labor supply shocks considered in our 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 $S > 1$.

Note that in both of these cases, the econometrician must take a stand as to 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 (as we suggest in section 3), this choice may be problematic in other cases, such as when one is interested in whether an exogenous variable affects commodities only through aggregate production and income levels or more directly through other channels, i.e., whether this 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 would be to make use of the sign restrictions on the loadings. The theory predicts that the loadings on y_t^{nc} must all be positive. Letting $\tilde{L}$ be the N by 2 matrix of unrotated factor loadings, the rotated loadings are $L = \tilde{L}T = [\tilde{L}^1\ \tilde{L}^2]T$ so that the loadings on the first rotated factor (corresponding to y_t^{nc}) are $L^1 = \tilde{L}^1 \cos\theta + \tilde{L}^2 \sin\theta$. 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 will lead only to a set of admissible values of θ and not uniquely identify the rotation

matrix. Thus, this approach would be akin to the weak identification of VAR's by sign restrictions (as in Uhlig 2005) in which one may identify a wide range of models for which the restrictions hold.

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 structural factors. Furthermore, these factors will have clear interpretations: one will correspond to the level of production and income net of commodity-related shocks (i.e., the IAC factor), while other factors would correspond to one or more of these commodity-related shocks. The identification of the rotation matrix, and therefore the structural 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 the structural factors empirically therefore provide a novel identification method for separating fluctuations in aggregate output into those driven by commodity-related shocks versus 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 those commodities that conform along a number of dimensions to the theoretical structure of the model. We then implement a factor decomposition and apply identification conditions to recover the underlying structural factors as suggested by the theory. After considering a wide range of robustness checks, we argue that exogenous commodity-related shocks have contributed relatively little to global economic volatility.

3.1 Data

The choice of commodities in the empirical analysis is guided by the theoretical model. In particular, there are four guidelines that we follow. First, commodities must not be vertically integrated. Second, the primary use of commodities must be directly related to the aggregate consumption bundle, and not for financial speculation motives. Third, commodities must not be jointly produced. Finally, the pricing of commodities must be determined freely in spot markets and must not display "stickiness" in prices reflecting long-term contractual agreements or menu costs.

The first condition, 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 sorghums would, because sorghums is primarily used as feed, then affect the price of non-grass-fed beef, and subsequently the price of milk and hides. To satisfy this condition, we deliberately exclude from our cross-section 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 deliberately do not include any energy commodity or 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 condition is to ensure that the primary forces driving prices of the included commodities are the production conditions of each commodity as well as the demand for these commodities reflecting the level of economic activity. Some commodities, in particular 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 our cross-section of commodities as well.

The third condition 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 would clearly not be applicable.

The fourth condition 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 many of the commodities for which one can find data. 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 who systematically display long periods of price invariance. For some commodities in our 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 1 provides more details on these adjustments.

These criteria for exclusion leave us with forty commodities in our sample. These include twenty-two commodities that we will 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, sorghums, soybeans, sugar, tea, tobacco, and wheat. Our data set also includes five 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 this data monthly 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 data are consistently available from 1968:1 until 2013:1, there are nonetheless a number of missing observations in the underlying data, as well as periods when we treat

---

[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.

the available data as missing when spot trading was limited. Appendix 1 provides details on the construction of each series, and 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 subsequent empirical analysis to the period since 1968, in light of the numerous price regulations and government price support mechanisms in place during this earlier period.

Table 1 presents information on the primary producing countries for each commodity in 1990, the middle of our sample, as well as information on the common uses of each type of commodity. The data on production come first from the CRB Statistical Yearbook, when available, and otherwise from other sources such as the United Nation's Food and Agricultural Organization (FAO). The table documents the wide regional variation in production patterns across commodities. While some countries are consistently and not surprisingly among the major producers of many commodities due to their size (e.g., USSR, China and India), the geographic variation is nonetheless quite 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 in 1990, Poland was second, accounting for thirteen percent of global production. Similarly, while the former USSR was also the largest producer of sunflower oil in 1990 with 29% of global production, Argentina was the second largest, accounting for 17% of global production. Among industrial commodities, Chile is well-known to be one of the world's largest producers or 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), Canada was the second largest producer of nickel (22%), 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 should be emphasized 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 in fact used in a variety of ways that makes 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 our sample are well-known for their use in cooking but some (like palm and coconut oil) also have a number of important industrial uses.

3.2 Reduced Form Common Factors in Commodity Prices

Prior to the factor analysis, we normalize each price series by the US 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 consider several metrics to characterize the contribution of the first five factors to accounting for commodity-price movements, summarized in Table 2.[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^2s that explicitly take into account missing values associated with some commodities. For example, the second row presents the average across the individual R^2s 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^2s, while the following row presents the R^2 constructed across all commodities (again omitting imputed values). Note that because different commodities have different time samples, the R^2s are not directly comparable across commodities, but this nonetheless provides 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 additional factors explain much 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 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 sparse 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 2 includes the contribution of different factors to explaining the variance across the three subsets of commodities in our sample: agricultural/food, oils and industrials. Differences across subsets of commodities are quite small: 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. Classical likelihood methods for estimating factors yield indistinguishable results.

[5] These information criteria for the optimal number of factors, 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.

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.[6] Among agricultural commodities, apples, bananas, onions, pepper and shrimp have much smaller R^2s 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 is worth emphasizing 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 2, this is not the case.

3.3   Identification of the Rotation Matrix and Structural Factors

To implement a structural interpretation of the factors as suggested by the model, we interpret the results of Table 2 as indicating that a two-factor representation is a reasonable one. First, additional factors beyond the first two add relatively little in explanatory power and therefore can safely be omitted. Second, it is, under the null of the model, *a priori* unlikely for there to be less than two structural factors. Indeed, such a finding if taken literally would imply that there are *no* shocks which directly affect commodity prices and therefore that all commodity-price movements reflect either the level of aggregate economic activity or idiosyncratic commodity factors. We can rule this out immediately because there exists at least one common shock to the supply of commodities: exogenous energy price movements. Because most commodities require energy in production and distribution, exogenous shocks to energy prices will necessarily induce some comovement in commodity prices since, as illustrated in Table 2, 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 do indeed 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) measure of exogenous OPEC production shocks. Following Kilian (2008), we use one year of lags for the autoregressive component and two years of lags for OPEC 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 suggests that oil production shocks do tend to affect commodity prices and therefore that there exists at least one source of direct

---

[6] Appendix Table 2 presents R^2s for each commodity from each factor.

commodity shocks. Thus, we will focus on the two-factor representation of real commodity prices from this point on.

To estimate the rotation matrix, our baseline approach is to impose orthogonality conditions on the indirect aggregate common factor F_t^1. Specifically, we take ε_t^{opec} the measure of OPEC production shocks from Kilian (2008) and define the orthogonality conditions as $E[F_t^1 z_t]$ where $z_t \equiv \begin{bmatrix} 1 & \varepsilon_t^{opec} & \dots & \varepsilon_{t-L}^{opec} \end{bmatrix}$ is the vector of instruments that consists of a constant, the contemporaneous value of Kilian's (2008) OPEC production shock, as well as L lags of the shock. The IAC 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 we have more moment conditions (L+2) than parameters (θ), we can estimate the rotation parameter θ using GMM by minimizing $J(\theta)$

$$J(\theta) = \left[\tfrac{1}{T}\sum_t(F_t^1(\theta)z_t)\right] W \left[\tfrac{1}{T}\sum_t(F_t^1(\theta)z_t)\right]'. \tag{38}$$

Kilian's (2008) measure of OPEC production shocks is available on a monthly basis from January 1968 until August 2004 although the first production shock does not occur until November 1973.[7] As noted before, many commodity prices respond significantly to OPEC oil production shocks. Furthermore, the second unrotated factor is significantly impacted by OPEC shocks, with peak effects obtaining 15 months after the shock and declining gradually thereafter. We can reject the null that OPEC shocks have no effect on the unrotated second factor at the 1% level using anywhere between 18 and 36 lags of OPEC shocks.[8] Thus, the orthogonality condition of the instrument follows from the theory and this empirical evidence suggests that OPEC shocks have clearly discernible effects on commodity prices, justifying their use as instruments. We set L=36 months for our baseline estimation to capture the fact that OPEC shocks have long-lived effects on commodity prices, although as we document below, the results are robust to both shorter and longer lag specifications as well. 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 3 presents the resulting estimate of θ and its associated standard error. With $\hat{\theta} = -0.24$ and a standard error of 0.20, 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 $t_{21} = 0$. As a result, the 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 3 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 reflects the fact that OPEC shocks have only gradual

---

[7] We are grateful to Lutz Kilian for providing us with the monthly series underlying the quarterly data used in his (2008) paper. We extend the series back to January 1968, with zero shocks to the series prior to 1973.

[8] Specifically, we regress the unrotated factor on a constant, the contemporaneous OPEC shock, and L lags of the OPEC shock and test the null hypothesis that all coefficients on OPEC shocks are zero.

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 qualitatively unchanged.

The fact that the estimated rotation matrix is close to the identity matrix reflects the fact that while the first unrotated factor is largely uncorrelated with OPEC production shocks, this is not the case 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 our 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 4 the estimated factor loadings for each rotated factor. The loadings on the IAC 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 (i.e., agricultural versus industrial) which again confirms that the factors explaining commodity prices are common across commodity subsets. Thus, and despite not imposing any restrictions on the loadings as part of the identification strategy for the rotation matrix, the estimated rotation satisfies theoretical predictions on factor loadings as well as overidentifying restrictions.

Given our estimate of θ and therefore 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 structural factor is presented in Figure 2 after HP-filtering with λ=129,600, the typical value for monthly data, to highlight cyclical variation. In addition, we repetitively 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) structural 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 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 procylical pattern to the IAC factor relative to U.S. economic conditions, a point to which we return in greater detail in section 3.5.

3.4   Robustness Analysis of the Estimated Indirect Aggregate Common Factor

In this section, we investigate the sensitivity of the estimated IAC factor to a number of potential issues. These include the identification strategy, the choice of commodities, the treatment of trends in the data, the imputation procedure for missing values, and the initial factor decomposition method.

First, we consider an alternative identification strategy for the rotation matrix. Our baseline approach is to impose orthogonality conditions, namely that the non-commodity related shock be orthogonal to OPEC oil production shocks, but GMM estimates can be notoriously sensitive. An alternative approach described in section 2.5 is to use theoretical predictions on signs of factor loadings: loadings on the IAC 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 consists of identifying the set of θ such that $\min\left(\tilde{L}^1 \cos\theta + \tilde{L}^2 \sin\theta\right) >$ 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 IAC 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 IAC factor from the baseline GMM estimation. There is tremendous 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 IAC 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.

Second, we want to verify that the results are not unduly sensitive to specific commodities or groups of commodities within our cross-section. 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. In the top left panel of Figure 3, we reproduce the 99% confidence interval from the GMM estimation of the rotation matrix when we keep only wheat out of the grains and replicate the factor analysis and rotation estimation. There is, qualitatively, little difference between our baseline result and this alternative. In the same spirit, we reproduce our results in the top right panel of Figure 3 keeping only palm oil out of the five oils in our cross-section. Again, this changes little other than to increase the confidence intervals in a few periods, such as 1975-1976 and 1995-2000.

One might also be concerned about too much overlap in how some commodities are used. For example, in Table 1, we documented that many of the agricultural commodities and oils are primarily used as feed or food. In the two middle panels of Figure 3, we replicate our results dropping either all commodities whose primary (60% or more in Table 1) use is as food (left panel) or as feed (right panel).

Despite the fact that this implies dropping many commodities in each case (16 in the case of food, 6 in the case of feed), the results are again quite similar to our baseline case.

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 stand out as accounting for a large proportion of many of the commodities. As a result, country-specific shocks could potentially induce comovement within the subset of commodities primarily produced in that country. To assess this possibility, we consider two additional exercises. First, we drop all commodities for which the primary producing country in 1990 (or as available in Table 1) was the former U.S.S.R. Results from this robustness exercise, which entails dropping 8 commodities out of the 40 in our cross-section, are in the bottom left panel of Figure 3. As can immediately be seen, there is now much more uncertainty around the estimated IAC factor than in our baseline. Interestingly, the increase in uncertainty primarily occurs in the 1970s, not in the 1990s after the collapse of the Soviet Union. Furthermore, the increase in uncertainty primarily reflects an increase in the standard error of the estimated rotation parameter, while the actual estimates of θ and the underlying unrotated factors are almost identical to those in the baseline. Thus, this is not suggestive of country-specific driven comovement in commodity prices. In the bottom right panel of Figure 3, we perform a similar robustness check dropping all of the commodities for which either China or India were the primary producers in 1990, or thirteen commodities in total. In this case, the results are almost identical to those generated by the baseline and, if anything, add some precision over the course of the late 2000s. Thus, we conclude that the baseline estimation of the common factors in commodity prices is robust to the choice of commodities included in our cross-section.

We are also interested in assessing whether our results are sensitive to more statistical considerations. For example, we perform factor analysis using the level of real commodity prices. As can be seen in Appendix Figure 2, there is little visual evidence of commodity prices exhibiting pronounced trends over this period. Nonetheless, we want to ensure that our results are not driven by spurious correlations from trends. We address this in two ways. First, we replicate our analysis after linear detrending each series prior to extracting factors. Results from this alternative approach are presented in the top left panel of Figure 4. The point estimates of the indirect aggregate common factor are similar, although the uncertainty surrounding these estimates follows a different pattern than in the baseline: the confidence intervals are much narrower through much of the sample but wider in the early 2000s. An alternative is to perform the factor analysis using the first-difference of real commodity prices. We present the IAC factor (accumulated in levels) from this additional specification in the top right panel of Figure 4. The uncertainty surrounding the estimates is now much higher in a number of periods, but there are few qualitative differences between the two sets of estimates. Thus, we conclude that our baseline results are not overly sensitive to the assumptions made about underlying trends in the data.

We consider two final checks on our results. First, we drop all commodities for which some significant imputations had to be done (i.e., more than a few missing observations at the end of the sample), or 7 commodities in total). As shown in the bottom left panel of Figure 4, this has almost no effect on the results. Thus, our findings are not sensitive 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 our baseline, as shown in the bottom right panel of Figure 4. In short, our estimates of the IAC 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 any 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 OPEC 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.

3.5   The Contributions of the Factors to Historical Commodity Prices and Global Activity

The theory of section 2 suggests that one can interpret one of the common factors among real commodity prices as the level of global activity that would have obtained absent any commodity-related shocks. Furthermore, the theory provides guidance as to how one can identify this specific factor from the data, and the previous sections have shown how one can implement this identification procedure empirically. In this section, we construct historical decompositions of commodity price movements and global economic activity following the structural interpretation suggested by the theory.

For prices, we decompose the average (across commodities) annual percentage change in commodity prices into those components driven by "indirect" shocks versus "direct" commodity shocks. This follows directly from the rotated factor structure, yielding

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

where the bar denotes that these are averages across all commodities in the cross-section. The first term on the RHS therefore constitutes the contribution of the IAC factor to average commodity price movements, the second constitutes the contribution of the DAC factor, and the third term reflects average idiosyncratic effects. We focus on annual changes in prices to abstract from very high-frequency volatility in commodity price changes.

The results of this decomposition are presented in the top panel of Figure 5, in which we plot the contributions from the IAC and DAC factors each month as well as the actual annual average price change across commodities (the idiosyncratic component contributes little, so we omit it from the figure). The IAC factor explains most of the historical commodity price changes. Thus, historical changes in commodity prices have primarily reflected endogenous responses to non-commodity shocks, i.e., shocks

that do not directly affect commodity markets other than through their general equilibrium effects on global economic activity. To the extent that income effects on inputs into the production of commodities are most likely weak, the IAC factor could then be interpreted as primarily reflecting changing global demand for commodities in light of global economic activity. During the commodity boom of 1973-74, for example, indirect shocks to commodity markets accounted for over two-thirds of the rise in commodity prices, with the remainder reflecting direct commodity-related shocks. Similarly, the fall in commodity prices during the Volcker era of the early 1980s is attributed almost entirely to a decline in the IAC factor.

The second commodity boom of the 1970s, however, suggests a more nuanced interpretation. While the rise in commodity prices in 1976 reflected rising levels of global activity, the IAC factor contributed much less to rising commodity prices during the second half of 1978 and was actually pushing toward lower commodity prices for most of 1979. Despite this downward pressure from non-commodity shocks, direct commodity shocks pushed real commodity prices higher during 1979 and did not weaken until very early 1980. Thus, while the brunt of the second commodity boom of the 1970s can be interpreted as an endogenous response of commodity prices to non-commodity shocks, commodity-related shocks played an important role in extending this period of rising commodity prices into early 1980.

The decomposition of commodity prices since the early 2000s also presents a mixed interpretation. While much of the rise in commodity prices since 2003 is accounted for by the IAC factor, direct commodity shocks account for much of the surge in prices during 2004 and approximately 30% of the rise from early 2006 to late 2007. The majority of the subsequent decline in commodity prices between October of 2008 and March of 2009 is also accounted for by direct commodity shocks, while the indirect factor accounts for most of the continuing decline after March 2009. Out of the total decline in commodity prices between October of 2008 and October of 2009, over half (56%) was due to direct commodity shocks. By contrast, the resurgence in commodity prices since the end of 2009 primarily reflects non-commodity shocks as measured by the IAC factor. In short, the decomposition suggests that while most historical price movements in commodity prices have been endogenous responses to non-commodity shocks, affecting commodity prices through the effects of these shocks on global activity, there have been a number of episodes in which direct shocks to commodity markets have played quantitatively important roles, including during the second commodity price boom of the 1970s as well as the run-up in commodity prices from 2003 to 2008 and their subsequent decline in 2008 and 2009.

We can also 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 used 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 IAC 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 IAC factor can be directly interpreted as those changes in aggregate output driven by indirect shocks. Because the scale of the IAC factor is not identified, we normalize it such that the standard deviation of quarterly changes in the IAC is equal to the standard deviation of quarterly percent changes in global IP and then treat the resulting historical changes in the IAC 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 IAC (which we define as δ_t where $\delta_t \equiv \Delta y_t - \Delta y_t^{nc}$) should therefore 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 then estimate

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

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

We plot the resulting contributions of the IAC and DAC factors to global IP growth in the bottom panel of Figure 5, again showing only the annual changes to filter out the significant high-frequency variation in the measurement of global IP. The correlation between changes in the IAC factor and annual changes in global IP is quite high (0.55) 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 IAC factor. The dynamics of global activity from the late 1980s to mid 1990s are also largely attributed to the IAC factor, although actual changes in global IP exceeded those predicted by the two factors. As was also the case with commodity prices, growth in the IAC factor during the 2000s coincides with the growth in global IP during this time period, whereas commodity-shocks in the DAC contributed modest downward pressure on economic activity in 2002 and 2003, then again in 2007-2010. To the extent that the DAC factor reflects exogenous energy price fluctuations, the negative contribution of the DAC factor from late 2007 through 2010 (subtracting 1-2% from the annual growth rate of global IP) is consistent with Hamilton (2009) who argued that oil shocks contributed to the severity of the Great Recession. Nonetheless, our decomposition suggests that

approximately two-thirds 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 IAC factor, i.e., non-commodity shocks.

Commodity shocks as captured by the DAC factor also had non-trivial consequences for global IP growth in several historical episodes. From mid 1985 to late 1986, for example, the DAC factor contributed an extra 1 percentage point to global growth, likely reflecting the concurrent large declines in oil prices. From 1991 to 1994 after the Iraq war, this pattern was reversed with the DAC factor subtracting between one half to one percentage point from the growth rate of global production. Thus, the decomposition does point to some historical role for exogenous commodity shocks in affecting global production. But the key take-away from this decomposition is that this contribution has generally been dwarfed by other economic shocks by the IAC factor.

4   Storage

The model in section 2 yields a tractable factor structure of commodity prices whose properties, as documented in section 3, conform closely to the data and permit us to make causal inferences about the relationship between global activity and commodity-related shocks. The key to the identification in the factor structure is the fact 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 IAC factor. This reflects the fact 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 in the model of section 2 productivity shocks are highly persistent while shocks to the household's willingness to supply labor are transitory. We can 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 and has access to a storage technology, leading it to hold inventories in the steady-state. As illustrated in Deaton and Laroque (1992), the key determinant of whether the storage sector will increase or decrease its inventories is the expected path of prices of the commodity. If a current increase in prices is expected to be temporary, then the storage sector will sell a positive amount of its inventories on the spot market today when prices are high and rebuild inventories in future periods when prices are lower. This will increase the contemporaneous supply of the good and reduce it in the future. By contrast, if the shock is expected to increase prices permanently, the storage sector has no incentive to change its stock of inventories and therefore will not be a net purchaser of the good. Thus, the persistence of the driving process (through its effect on the path of expected prices) affects the size of net purchases made by the storage sector. This in turn alters the contemporaneous response of prices. Thus, if aggregate productivity shocks in our model were highly persistent while labor supply shocks were transitory (or vice-versa), the presence of storage would lead these shocks to have different supply responses depending on the size of the resulting net purchases of the storage sector and therefore the comovement in commodity prices would not be the same across the two shocks. This feature would break the aggregation of indirect shocks into a single IAC factor.

In practice, there are three reasons to think that this is unlikely to be quantitatively important. First, if storage played an important role in the determination of commodity prices and the indirect shocks did affect the paths of expected prices differently such that the aggregation of indirect shocks into a single IAC 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. This is because there are a number of different aggregate structural shocks affecting commodity prices through the indirect channel of global activity, such as financial shocks and fiscal policy shocks, in addition to the productivity and labor supply shocks that we explicitly model. But as documented in Table 2, the comovement of commodity prices is well-characterized by two factors, with any additional factors adding little explanatory power. This 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 our aggregation of indirect shocks 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 that the effects of net purchases for storage motives do indeed appear to be second-order for most commodities. To examine this claim, suppose 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 subject to some depreciation, costs, and convenience yield. In the presence of adjustment costs associated with changing inventory holdings, net purchases of the storage sector would reflect expected price changes, interest rates and the current stock of inventories. 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 net purchases by the storage sector to accumulate inventories would increase 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

$$\left(\frac{Y}{Q} - 1\right) np_t(j) = \left(\frac{Y}{Q}\right) y_t(j)$$

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.

We investigate whether ratios of consumption to production of commodities are empirically significantly different from one. Specifically, for each commodity we construct a time series of the ratio of consumption to production and test the null hypothesis that the mean is different from one. Because of the limited availability of historical consumption and production data, the analysis is done at the annual frequency. When available, 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 US Department of Agriculture's Food and Agricultural Services (USDA FAS),

and from trade associations. For example, 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.

With annual time series for global consumption and production of commodities, we define $r_t = \frac{Y_t}{Q_t} - 1$ for each commodity where Y is global consumption and Q is global production of the commodity. The difference between Y and Q reflects the net purchases of the storage sector, i.e., the change in the stock of inventories after depreciation. We then regress the net ratio r_t on a constant. Results from these regressions are presented in Table 5 for all commodities for which the data are available. Out of thirty two commodities, we reject the null that $r_t = 0$ on average for only eight: 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 go bad 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, on average, 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 is also the case for 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. 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 storage sector are on average small. Finally, if we replicate our analysis using only the commodities for which we cannot reject the null of zero net-purchases on average, this has little effect on our estimate of the IAC factor (Appendix Figure 3).

A third way to assess the possibility that the effects of storage could break the aggregation of indirect commodities into a common IAC factor is to note that, in the presence of storage motives, interest rates would play an important role in affecting commodity prices as an opportunity cost to holding inventories (Deaton and Laroque 1992, 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 IAC factor.

To assess this prediction, we identify monetary policy shocks using a time-varying-coefficients Taylor rule

$$i_t = c_t + \varphi_t^{\pi} F_t \pi_{t+1,t+2} + \varphi_t^{gy} F_t gy_t + \varphi_t^x F_t x_t + \rho_t i_{t-1} + \varepsilon_t^{mp} \tag{39}$$

in which the central bank responds to its real-time forecasts (F_t) of each of average inflation over the next two quarters ($\pi_{t+1,t+2}$), the current quarter's output growth (gy_t), and the current quarter's output gap (x_t) as well as the previous period's interest rate as in Kozicki and Tinsley (2009) and Coibion and Gorodnichenko (2011). We assume that each of the time-varying coefficients follows a random walk, including the intercept which 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 time-varying coefficients 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 (39) 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 aggregate common factor, we use a vector autoregressive representation of macroeconomic dynamics. Specifically, we estimate a VAR with four variables: our measure of monetary policy shocks, the log of US industrial production, the log of the U.S. Consumer Price Index (CPI), and the IAC factor. We order the monetary policy shock first given that it should already incorporate the most recent economic information via Greenbook forecasts and to allow other variables to respond on effect of 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 6 the impulse responses of industrial production, the CPI, and the IAC 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 starts rising 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, Eichenbaum and Evans 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.[9] Thus, we can

---

[9] Note that these standard errors do not account for the fact that the IAC factor is a generated regressor, and they therefore understate the true uncertainty around the point estimates. This will be addressed in a future version of the paper. However, there are at least two reasons to suspect that this result will not change. First, one could also test the null that monetary policy shocks have no effect on the AIC factor by regressing it on current and lagged monetary shocks, i.e. $f_t^{iac} = c + \sum_{i=0}^{I} \beta_i \varepsilon_{t-i}^{mp} + v_t$, setting I=36 months to account for the gradual effects of

statistically reject the null hypothesis arising from speculative motives that monetary policy shocks have no effect on the IAC factor.

In addition, we plot the historical contribution of monetary policy shocks to each of these variables in the bottom panel of Figure 6. Monetary policy shocks can account for much of the historical variation in industrial production and CPI inflation over the course of the 1970s and early 1980s, consistent with the "stop-go" description of monetary policy during this time period in Romer and Romer (2002). By contrast, monetary policy shocks have accounted for little of the macroeconomic volatility since the mid-1980s, consistent with Coibion (2012). For the indirect aggregate common factor, we find that monetary policy shocks can account for much of the sustained increase in the IAC factor from late 1975 until 1980, and approximately two-thirds of the subsequent decline from 1980 to 1982, which is consistent with the monetary interpretation of the mid-1970s suggested by Barsky and Kilian (2002). However, exogenous U.S. monetary policy shocks appear to have contributed little to common commodity prices in other periods, including during the first large run-up in commodity prices in 1973-1974 as well as during the more recent run-up from 2003-2008. Thus, neither episode can be directly attributed to US monetary policy according to the VAR.

In short, while the presence of commodity storage could potentially break the aggregation of indirect shocks into a common IAC factor, there is little quantitative evidence for this possibility.[10] First, the fact that the comovement is 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 IAC 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

---

monetary policy shocks on macroeconomic variables. From this procedure, we can reject the null hypothesis that monetary policy shocks have no effect on the global factor (i.e. $\hat{\beta}_i = 0 \ \forall i$) with a *p*-value of 0.013. The generated regressor issue is not binding in this case since the global factor is only on the LHS 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 purifies the monetary policy shocks of potentially remaining predictability from macroeconomic variables and is therefore in this sense 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.

[10] 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 (2012) 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.

binding (i.e. when log-linear approximations are less binding), our 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 mutually endogenous and jointly determined. For example, a positive technology shock increases total income and, all else equal, increases the demand for commodities and hence their real prices. Furthermore, the empirical evidence in section 3 documented that a large proportion of commodity-price movements are systematically related 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 pseudo out-of-sample forecasting exercise.

        While we restricted the cross-section of commodities in section 3 to conform to the theoretical structure of the model for the purposes of better identifying the structural factors, from a forecasting perspective we can use our identified factors to forecast a much broader set of commodities. For example, we ruled out vertically integrated commodities such as soybeans and soybean meal to avoid identifying spurious comovement reflecting idiosyncratic shocks to the soybean sector. But once we have recovered the IAC factor from the restricted cross-section, it should be able to help predict *all* commodities, not just those in our sample. Thus, in the out-sample forecasting exercise, we assess the ability of the common commodity factor to forecasted not just the set of commodities in our data set but also commonly used commodity indices frequently and the real price of oil.

5.1   Forecasting Model

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

$$x_{t+1} = A(L)x_t + e_{t+1} \qquad (40)$$

where $x_t = [rpc_{jt}, \ iac_t]'$, rpc_{jt} denotes the log of real price of commodity j, iac_t is the IAC factor extracted from the cross-section of real commodity prices, e_{t+1} is the regression error, and $A(L) = A_1 + A_2 L + A_3 L^2 + \cdots + A_p L^{p-1}$. In the forecasting exercise, the lag length p is chosen recursively using the Bayesian information criterion (BIC).

        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 IAC factor, we examine the ability of the IAC 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.[11] The indices are also deflated by U.S. CPI. Finally, we evaluate the ability of the bivariate VAR to forecast the real price of crude oil given the evidence that

---

[11] 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.

VAR-based models of oil-market fundamentals can generate economically large improvements in forecast accuracy (Baumeister and Kilian 2012; and Alquist et al. forthcoming). The real price of oil used in the forecasting exercise is the refiner's acquisition cost of imported oil, which is a good proxy for the international price of crude oil (see Alquist et al. forthcoming).

We apply the EM algorithm recursively to fill in the missing observations and estimate the common factor at each point in time (Stock and Watson 2002). 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 VAR 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 3). 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 MSPE of the VAR-based forecast the real commodity price at the 1-, 3-, 6-, and 12-month horizons, which correspond to the horizons of interest for policy makers. All forecast accuracy comparisons are conducted relative to the no-change benchmark. Multistep-ahead forecasts are computed iteratively using the VAR.

5.2   Forecasting Results

Table 6 summarizes the results obtained from the forecasting exercise for the commodity-specific and common sample periods. The first column of Table 6 shows the aggregate MSPE ratio, which is defined as:

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

where $MSPE_j^{VAR}$ is the mean-squared prediction error of the VAR-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 1% to 7% range. These forecast-accuracy improvements are modest in economic terms.

But these summary statistics mask the heterogeneity in the ability of the VAR to produce more accurate forecasts than the no-change forecast. Table 6 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 horizon and 20 (out of 40) commodities at the 3-month horizon for which the

VAR-based forecasts are more accurate than the no-change forecast. The performance of the VAR deteriorates as the forecast horizon lengthens. Similar, if not somewhat stronger, 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 VAR 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 VAR 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 VAR-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 VAR 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 VAR is more accurate than the no-change forecast, but the improvements are smaller (i.e., at most about 7%).

The VAR model also does well at forecasting the real price of oil at short horizons. For both forecast evaluation periods, it is able to produce improvements in forecast accuracy of about 20% at the 1-month horizon. The 3-month ahead forecasts are about 3-6% more accurate than the random walk forecast. The forecasts based on the VAR become less accurate as the forecast horizon lengthens.

Appendix Tables 3 and 4 report the forecast accuracy results for the individual commodities for the commodity-specific and common sample periods. Several things stand out about these results. First, the VAR-based forecasts generate improvements in forecast accuracy for some agricultural commodities and oils up to 12 months ahead. For example, Appendix Table 4 shows that 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 4 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 5. That table compares the bivariate VAR with a standard 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).[12] The first column of Appendix Table 5 shows

---

that the IAC factor based model does well relative to the oil-market model at the 1- and 3-month horizons when the BIC is used.[13] On the other hand, the IAC factor-based model is dominated by the oil-market model when a fixed lag length of 12 is used, although the IAC factor model still delivers improvements in forecast accuracy up to about 14% relative to the no-change forecast.[14] This evidence suggests that the IAC factor contains some information relevant for forecasting the real price of crude oil at short horizons.

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. This evidence is important from a practical perspective: data on market fundamentals at the relevant frequency for many of the commodities included in our data set are unavailable, unlike the case of crude oil, making the construction of forecasting models challenging. These results show that a VAR whose construction is guided by an economic model that focuses on the role of demand in driving commodity prices can be used to generate accurate forecasts of real commodity prices relative to the no-change benchmark. Moreover, to the extent that commodity prices do not adjust instantaneously to news about global demand conditions, we expect that the indirect aggregate common factor will contain some predictive power for real commodity prices. This model-based intuition is validated by the forecasting exercise. 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 examine the long-standing question of what drives the comovement among commodity prices and propose a way of interpreting this comovement in economic terms. We reach four main conclusions. First, the model suggests that the each commodity price can be decomposed into three distinct components. The first component captures the idiosyncratic price movements, the second one captures global economic forces, and the third one is related to commodity-specific shocks. In terms of our subsequent analysis, the IAC factor is of particular interest because it represents 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 IAC factor could help address the identification problem surrounding global economic conditions and commodity prices by exploiting the common comovement induced by indirect shocks captured by the IAC.

---

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 (forthcoming).

[13] During the 1984:1-2012:8 forecast evaluation period, for example, the IAC 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.

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

Second, we show how the model's predictions can be used to identify the rotation matrix that recovers a set of structural factors implied by the theory, including the IAC factor. This point addresses the central problem of factor analysis, namely that it is difficult to interpret the factors in economic terms. However, the theory provides a set of orthogonality conditions and sign restrictions that can be used to identify the parameters of the rotation matrix consistent with a structural interpretation of the factors.

Third, we apply these methods to a broad cross-section of commodity prices that, in some cases, extend back to the late 1950s to understand better the key drivers of the commodity-price comovement over long period of time. The IAC factor that we identify accounts for about 60-70% of the variance in commodity prices. In addition, we are unable to reject the theoretical restrictions implied by the model. The IAC factor is highly correlated with independently computed measures of global economic activity at business cycle frequencies. Its behavior during the 1970s and 1980s suggest that the macroeconomic volatility observed during that era was not driven primarily by commodity-related shocks, consistent with the interpretation of Barsky and Kilian (2002). 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 IAC factor is useful for forecasting real commodity prices, some widely used commodity-price indices, and the real price of crude oil. A pseudo out-of-sample forecasting exercise shows that a simple bivariate model that includes the IAC factor and the real commodity price can generate economically important improvements in forecast accuracy relative to a no-change benchmark. The results of this exercise are also consistent with the macroeconomic model that we use to interpret commodity-price comovement. Insofar as commodity prices do not respond instantaneously to macroeconomic news, the theory predicts that the IAC factor should have some predictive power for real commodity prices. Thus, our approach can be used to simultaneously forecast a wide range of commodity prices within a common framework while also providing a structural interpretation to those forecasts.

In sum, we provide a new model for thinking about the sources and implications of commodity-price comovement and show that several implications of the model are consistent with the data. This framework suggests a way of interpreting the common factors driving commodity prices and offers a new 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*.

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.

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.

Byrne, Joseph P., Giorgio Fazio, and Norbert Fiess, 2011. "Primary Commodity Prices: Co-Movements, Common Factors and Fundamentals," World Bank Working Paper 5578.

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. 2012. "Are the effects of monetary policy shocks big or small?" *American Economic Journal: Macroeconomics*, 4(2), 1-32.

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, A. and G. Laroque, 1992. "Behaviour of Commodity Prices," *Review of Economic Studies* 59(1), 1-23.

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.

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. 2012. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," forthcoming: *Journal of Applied Econometrics*.

Kydland, Finn E. and Edward C. Prescott, 1982. "Time to Build and Aggregate Fluctuations," *Econometrica* 50(6), 1345-1370.

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

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.

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, 2002.

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 Watson. 2002 "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147-162.

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, 2012. "Using Co-Movements to Forecast Commodity Prices," Manuscript.

Table 1: The Production and Usage of Commodities

	Largest Producers	Primary Uses
Agr./Food Commodities		
Apples (1990-91)	US (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)	US (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), US (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
Sorghums* (1990)	US (0.26), India (0.21), Mex. (0.11)	Food (0.39), feed (0.52)
Soybeans (1990-91)	US (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), US (0.10)	Smoking
Wheat (1990-91)	USSR (0.17), China (0.17), US (0.13)	Food (0.65), feed (0.22)
Oils		
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.
Industrial Commodities		
Aluminum (1990)	US (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), US (0.18)	Electrical (0.75), construction
Cotton (1990-91)	China (0.24), US (0.18), Uzb. (0.14)	Clothing, furnishings, medical
Lead (1990)	US (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.

Table 2: Contribution of common factors to commodity prices

Number of common factors:	Cumulative Variance Explained by Common Factors				
	1	2	3	4	5
Complete Sample:					
Cumulative eigenvalue shares	0.59	0.69	0.75	0.79	0.82
Mean across commodity-specific R^2s	0.60	0.69	0.74	0.78	0.81
Median across commodity-specific R^2s	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
Subset of Commodities:					
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^2s. The fourth row provides the joint R^2 constructed using all commodities. In addition, the top panel presents joint R^2s for subsets of commodities (as defined in Table 1). Each of the R^2 omits imputed values. See section 3.2 for details.

Table 3: GMM Estimates of Rotation Matrix

	GMM Estimates of Rotation Parameter				Implied Rotation Coefficients			
	θ	se(θ)	p(over-id)	N	t_{11}	95% CI(t_{11})	t_{21}	95% CI(t_{21})
Baseline GMM Estimates:	-0.24	(0.20)	1.00	405	0.97	[0.81 1.00]	-0.24	[-0.59 0.14]
(Iterative GMM, L=36)								
Robustness of GMM Estimates:								
More moments: (L=48)	-0.25	(0.19)	1.00	393	0.97	[0.82 1.00]	-0.25	[-0.58 0.12]
Fewer moments: (L=24)	-0.24	(0.21)	1.00	417	0.97	[0.78 1.00]	-0.24	[-0.62 0.18]
Fewer moments: (L=12)	-0.31	(0.30)	1.00	429	0.95	[0.61 1.00]	-0.31	[-0.79 0.28]
Two-step GMM	-0.25	(0.20)	1.00	405	0.97	[0.80 1.00]	-0.25	[-0.59 0.13]
Continuous GMM	-0.26	(0.19)	1.00	405	0.97	[0.80 1.00]	-0.25	[-0.60 0.12]
Alternative normalization	-0.22	(0.20)	1.00	405	0.98	[0.82 1.00]	-0.22	[-0.58 0.17]

Notes: The table presents nonlinear GMM estimates of parameter θ from (44) in the text, along with Newey-West (1987) standard errors (se(θ)), the p-value for over-identifying restrictions (p(over-id)), and the number of observations used in the estimation. The panel on the right presents the implied reduced-form 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 4: Rotated Commodity-Specific Factor Loadings:

	Factor Loadings			Factor Loadings	
Commodity	IAC	DAC	Commodity	IAC	DAC
Agr./Food Commodities			*Oils*		
Apples	0.43	0.20	Coconut oil	0.81	0.14
Bananas	0.53	0.30	Groundnut oil	0.83	0.25
Barley	0.68	0.51	Palm oil	0.86	0.26
Beef	0.88	0.03	Rapeseed oil	0.47	0.46
Cocoa	0.90	0.01	Sun/Safflower oil	0.79	0.34
Coffee	0.87	-0.05			
Corn	0.93	0.23	*Industrial Commodities*		
Fishmeal	0.88	0.28	Aluminum	0.79	0.17
Hay	0.85	0.08	Burlap	0.84	0.12
Oats	0.86	0.24	Cement	0.20	0.09
Orange Juice	0.76	-0.11	Copper	0.50	0.77
Onions	0.58	-0.31	Cotton	0.94	-0.07
Pepper	0.64	-0.54	Lead	0.63	0.68
Potatoes	0.73	0.05	Lumber	0.55	-0.15
Rice	0.91	0.22	Mercury	0.35	0.81
Shrimp	0.55	-0.68	Nickel	0.09	0.76
Sorghums	0.92	0.22	Rubber	0.72	0.56
Soybeans	0.94	0.15	Tin	0.87	0.30
Sugar	0.75	0.22	Wool	0.84	0.29
Tea	0.89	-0.09	Zinc	0.54	0.44
Tobacco	0.88	-0.21			
Wheat	0.90	0.26			

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 5: 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
Agr./Food Commodities					
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
Oils					
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
Industrial Commodities					
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 US 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 6: Summary of Recursive Forecast Accuracy Diagnostics for Real Commodity Prices

Forecast Evaluation Period: Commodity-Specific

	Aggregate MSPE Ratio	Distribution of MSPE Ratios					CRB	WB	IMF	Crude Oil
		[0,0.9)	[0.9,0.95)	[0.95,1)	[0,1)	[1,∞)				
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)	

Forecast Evaluation Period: 1984:1-2012:12

	Aggregate MSPE Ratio	Distribution of MSPE Ratios					CRB	WB	IMF	Crude Oil
		[0,0.9)	[0.9,0.95)	[0.95,1)	[0,1)	[1,∞)				
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 VAR 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 VAR-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 Comovement across Shocks

Panel A: Aggregate Productivity Shock


Panel B: Relative Demand Shock for Commodities
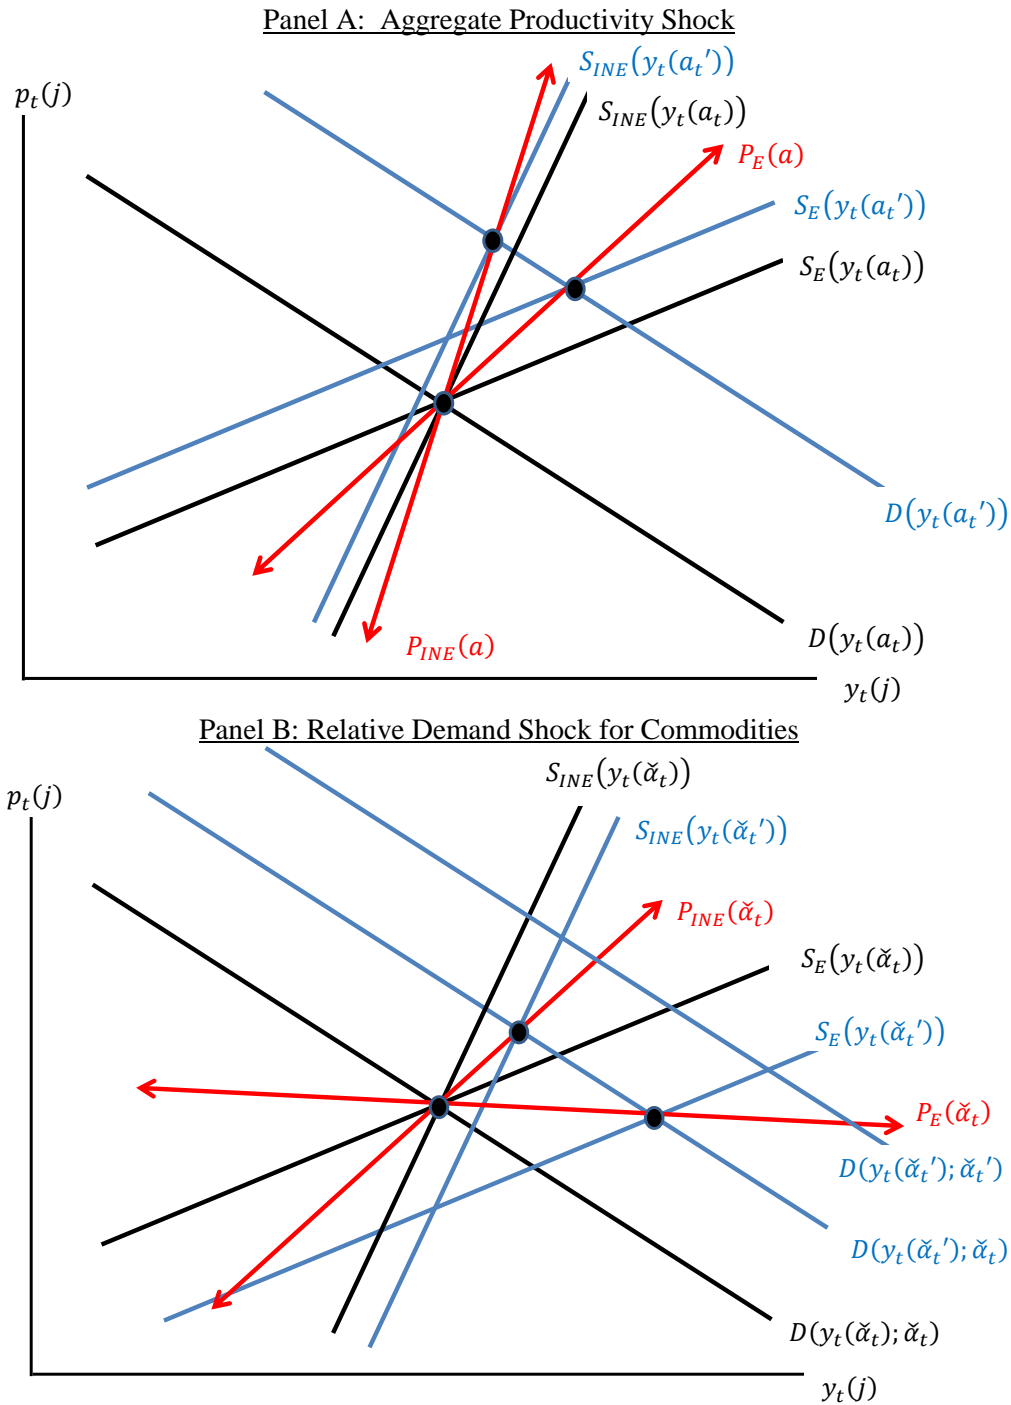
Notes: The top figure plots the effects of a change in aggregate TFP from a_t to a_t' on commodity prices. S_E and S_{INE} are supply curves for relatively elastically and inelastically supplied commodities, D denotes demand curves. Black lines correspond to curves before the change in TFP, blue lines denote curves after the change in TFP. Red lines with arrows denote the set of all possible prices that can obtain from changes in TFP of different sizes, with P_E being for elastically supplied commodity and P_{INE} for inelastically supplied commodity. Bold black dots identify equilibrium outcomes for given change in TFP. The bottom panel plots equivalent comparative statics for an increase in the relative demand for commodities ($\breve{\alpha}_t$).

Figure 2: Indirect Aggregate Common Factor in Commodity Prices

Panel A: Indirect Aggregate Common Factor (GMM Approach)

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

| US Recessions | 99% CI for Rotated Common Factor (GMM) | Min and Max Values for Rotated Common Factor (Factor Loading Sign Restrictions) |

Note: The top figure presents the IAC factor from the factor analysis in section 3.3. The IAC factor is HP-filtered (λ=129,600) in the figure. The light grey shaded areas are NBER-dated recessions. The dark 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 IAC factor as estimated by GMM (dark shaded areas) and the minimum and maximum range for admissible values of the IAC factor using sign restrictions on factor loadings (solid blue lines). See sections 3.3 and 3.4 in the text for details.

Figure 3: Robustness of Indirect Aggregate Common Factor using Subsets of Commodities
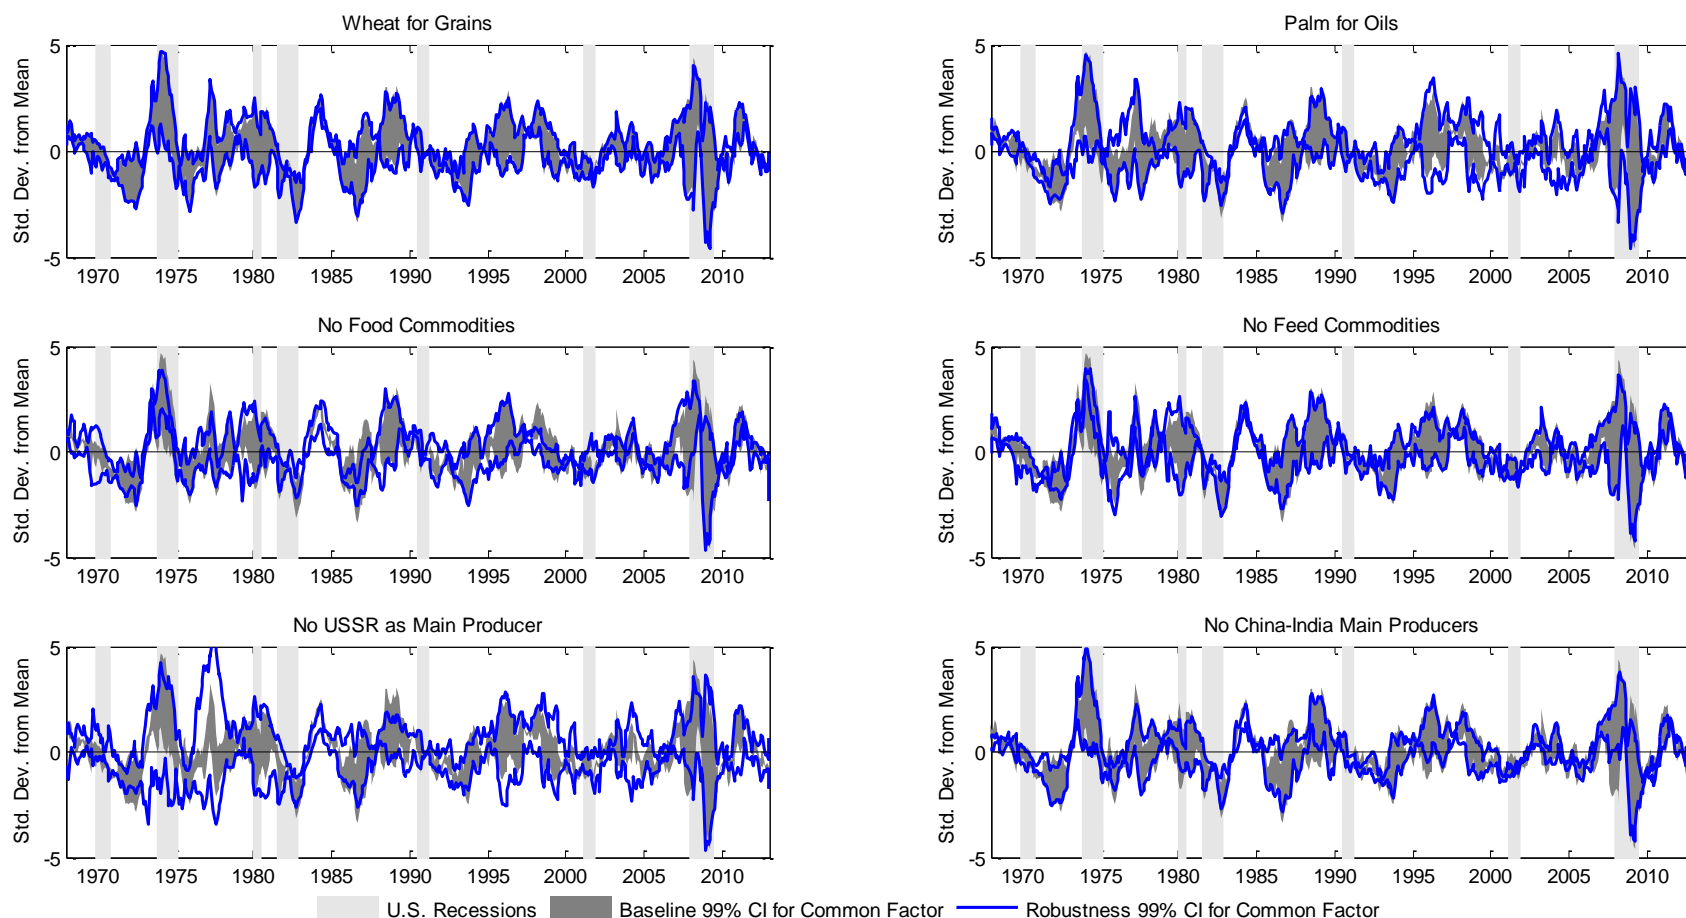
Note: The figures present the 99% confidence interval for the (HP-filtered) IAC factor (dark grey shaded area) and the 99% confidence intervals for the HP-filtered IAC factor for subsets of commodities (areas between blue lines). In the top two panels, we drop barley, hay, oats, and sorghums from the 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 as measured in Table 1 (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) as measured in Table 1 (left figure) and all commodities for which China or India were primary producers (13 commodities, right figure). See section 3.4 for details.

Figure 4: Additional Robustness Checks of Indirect Aggregate Common Factor
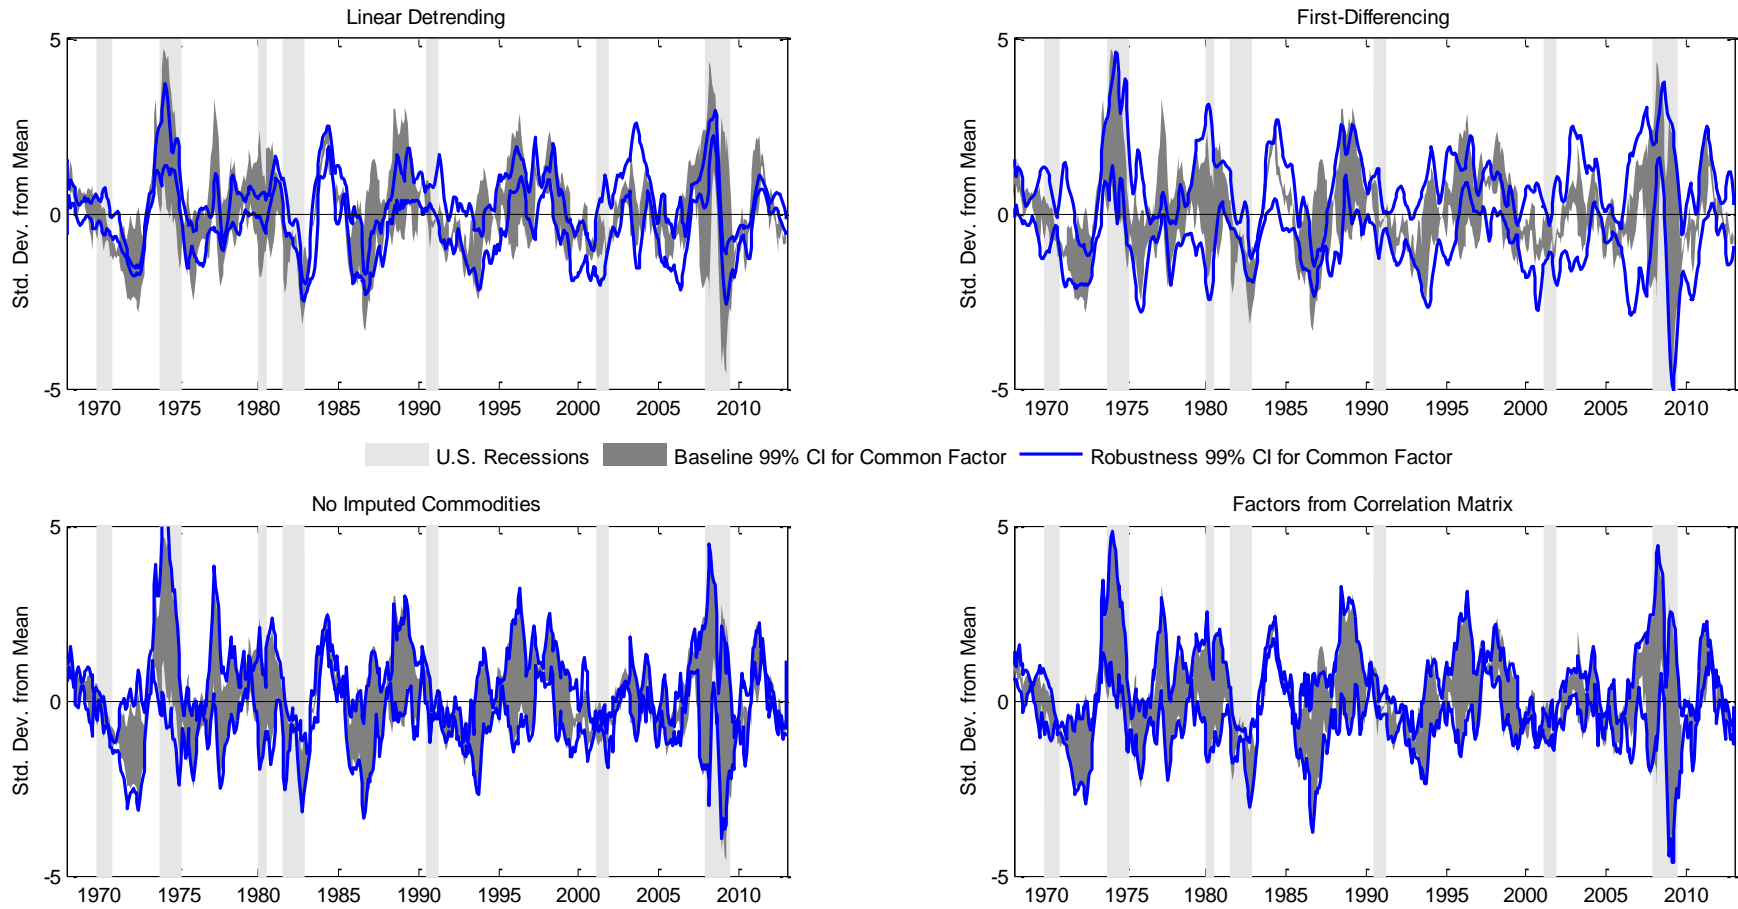
0

Note: The figures present the 99% confidence interval for the (HP-filtered) IAC factor (dark grey shaded area) and the 99% confidence intervals for the HP-filtered IAC 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 3.4 in the text for details.

Figure 5: The Contribution of "Indirect" and "Direct" Factors to Commodity Price Changes

Panel A: Contributions to Average Annual Commodity Price Changes

Panel B: Contributions to Annual Changes in Global Industrial Production
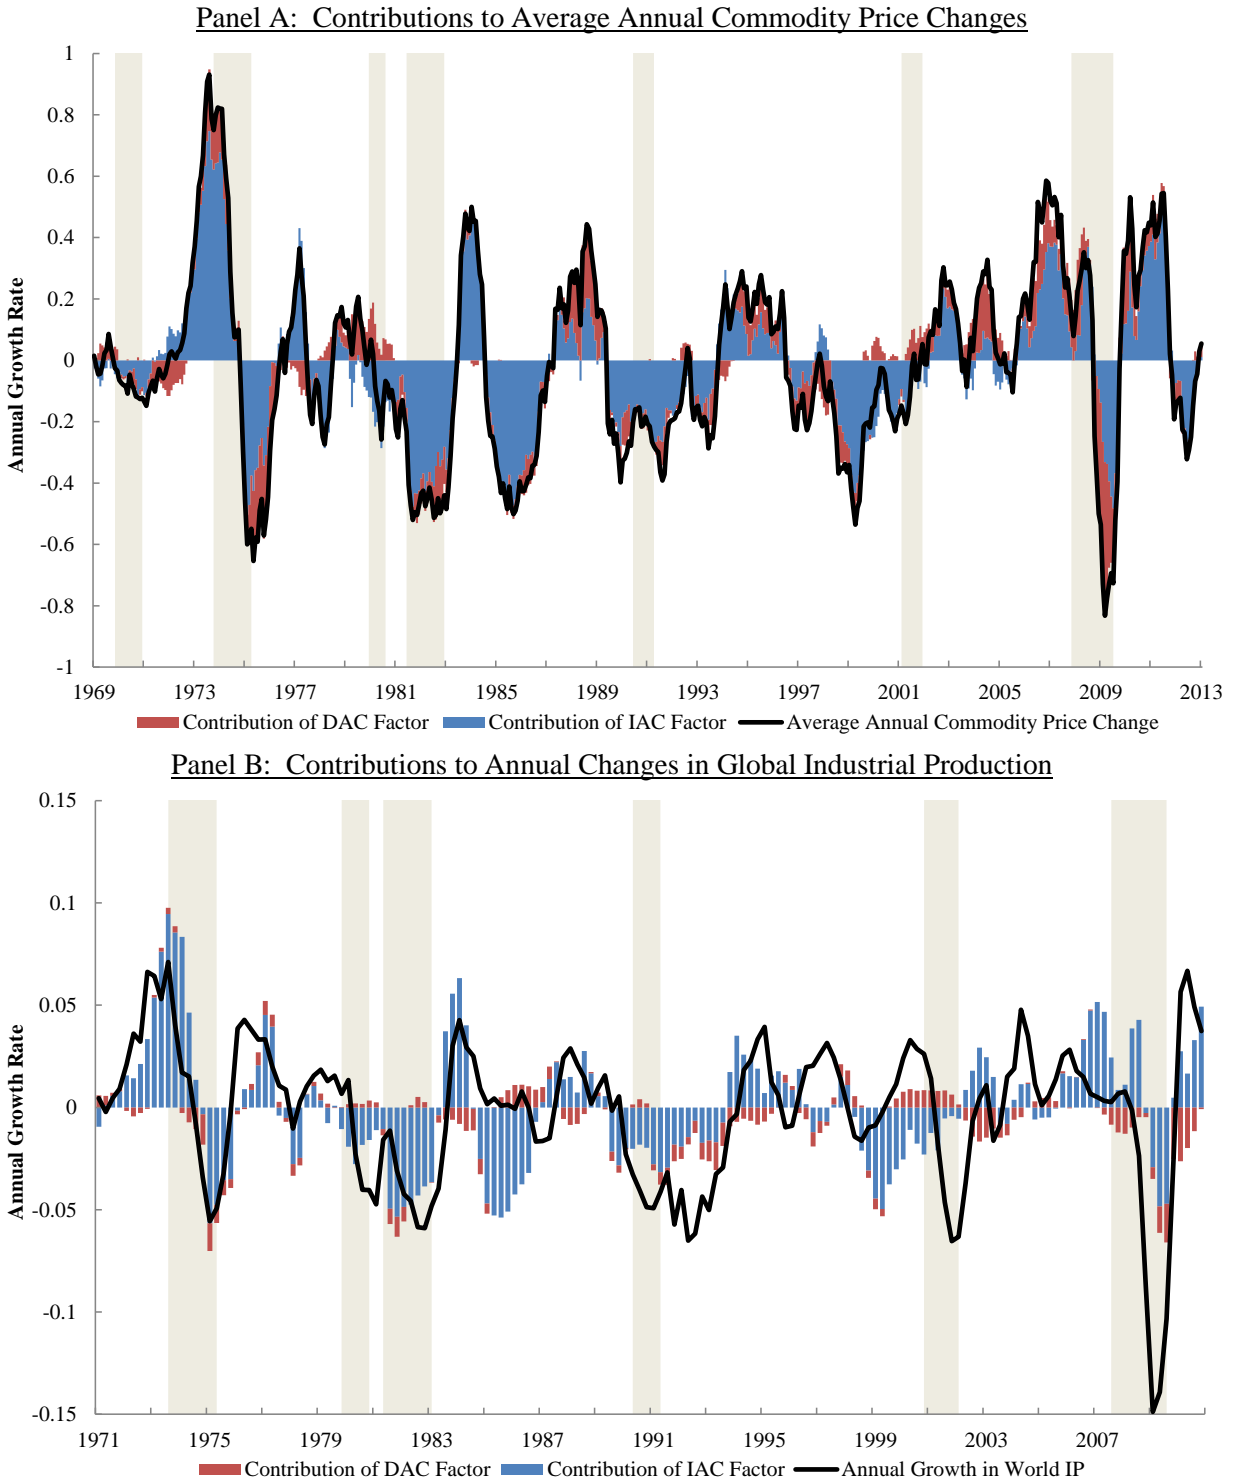
Note: The two figures plot the contributions of the "direct" and "indirect" factors (DAC and IAC respectively) to the average (across all commodities in our sample) annual commodity price change (top panel) and the annual growth rate of global industrial production (bottom panel). Data is monthly in the top panel and quarterly in the bottom panel. See section 3.5 for details.

Figure 6: Effects of Monetary Policy Shocks on the Indirect Aggregate Common Factor
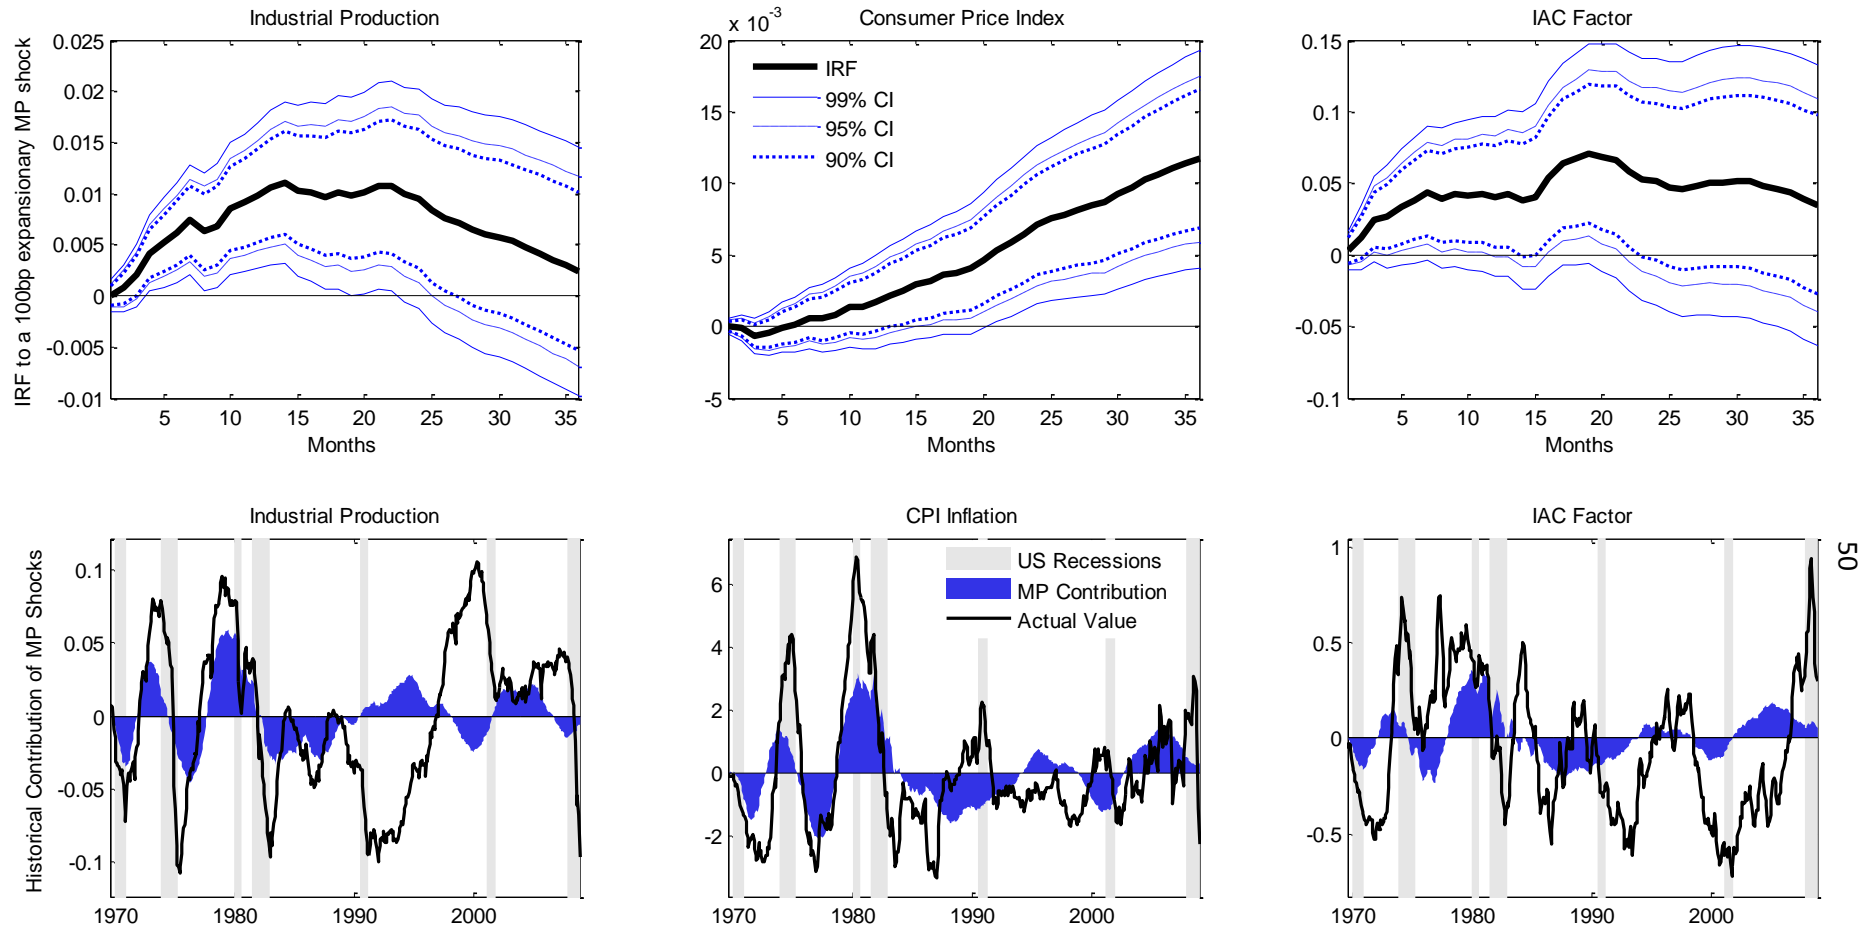
Note: The figures in the top row present estimated impulse responses of U.S. industrial production, the U.S. consumer price index, and the IAC factor to a 100b.p. expansionary monetary policy shock using the VAR described in section 4. Confidence intervals are constructed from the distribution of impulse 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.

Appendix Table 1: 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) applies 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, US import price, free on truck (f.o.t.) US 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 US, 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, US	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 US 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, US 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 Table 2: Contribution of Common Factors to Individual Commodity Prices

Number of Factors:	Cumulative R^2 from Common Factors				
	1	2	3	4	5
Agricultural/Food					
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
Oils					
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
Industrials					
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.
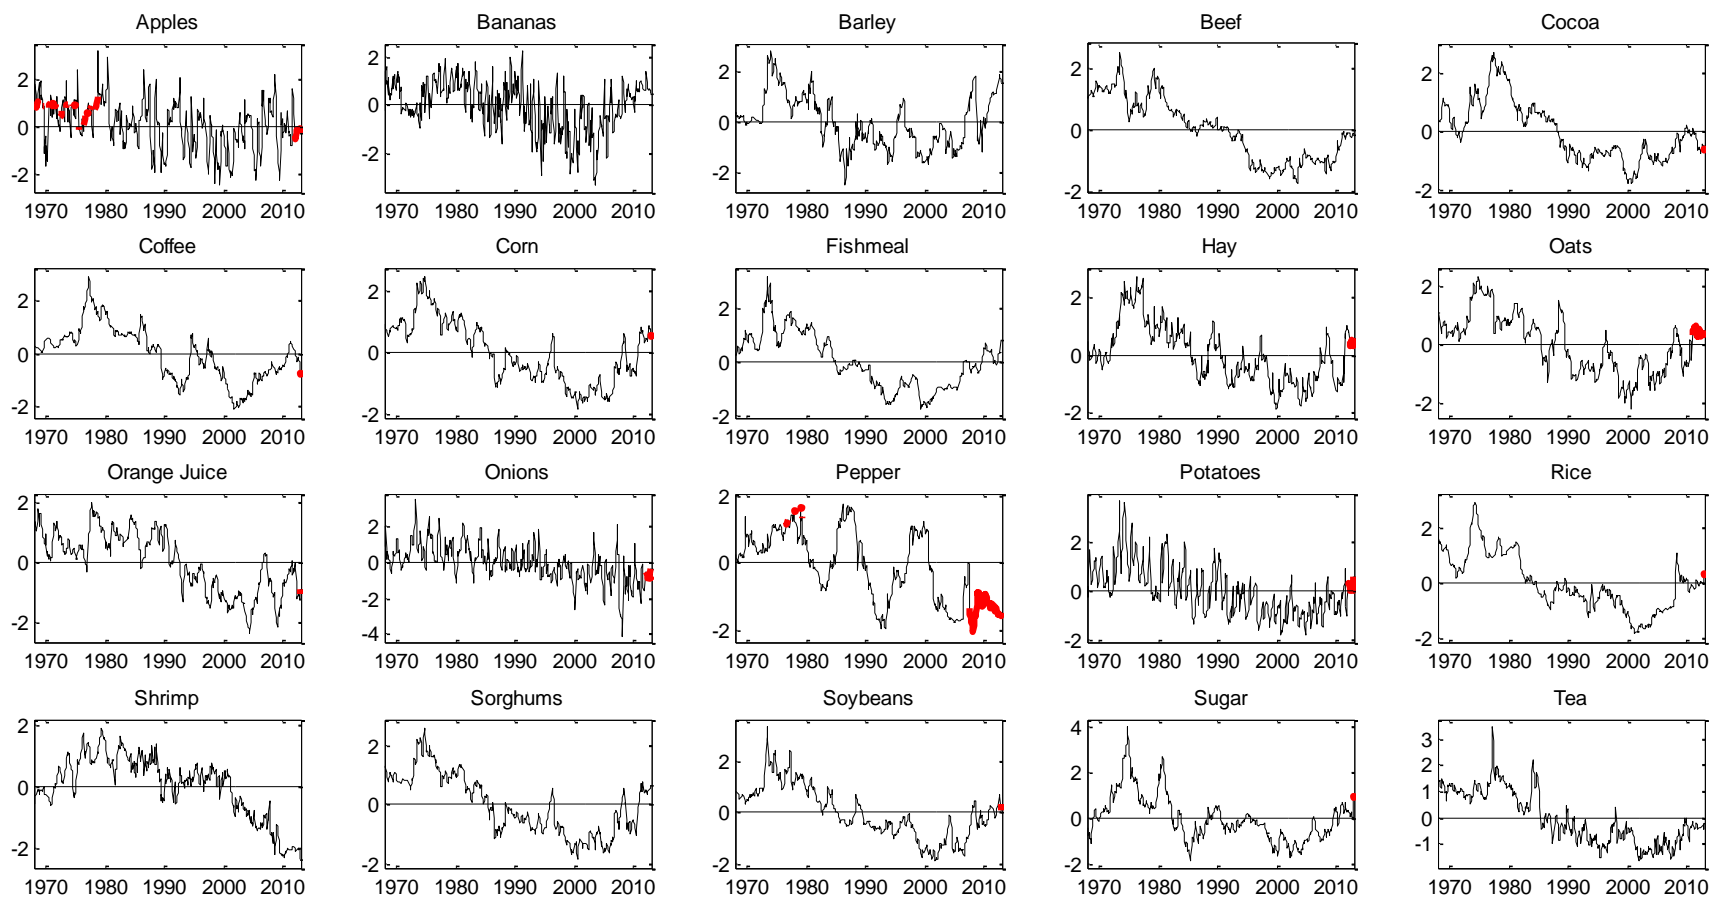
Appendix Figure 1: Price Observations Dropped
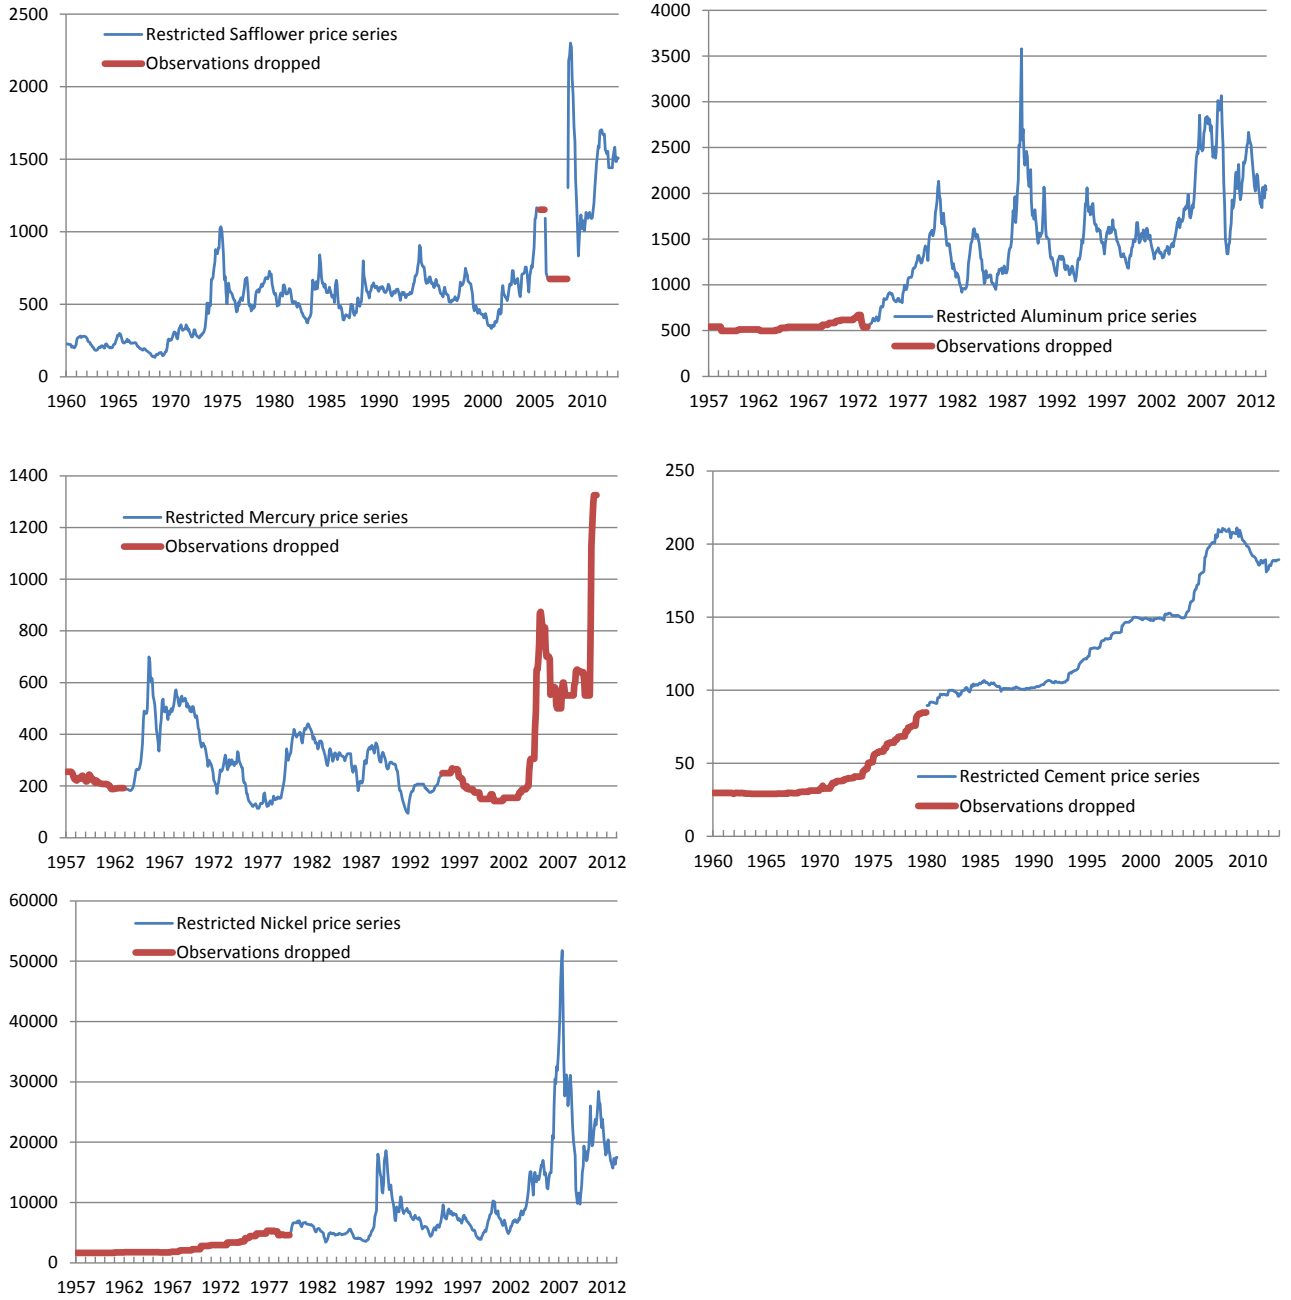
Note: Each figure presents the price series used in the empirical analysis (light blue line: "Restricted X price series") and the observations dropped (thick red line: "Observations dropped").

Appendix Figure 2: 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).

Appendix Figure 2 (continued): Real Commodity Prices and Imputed Values
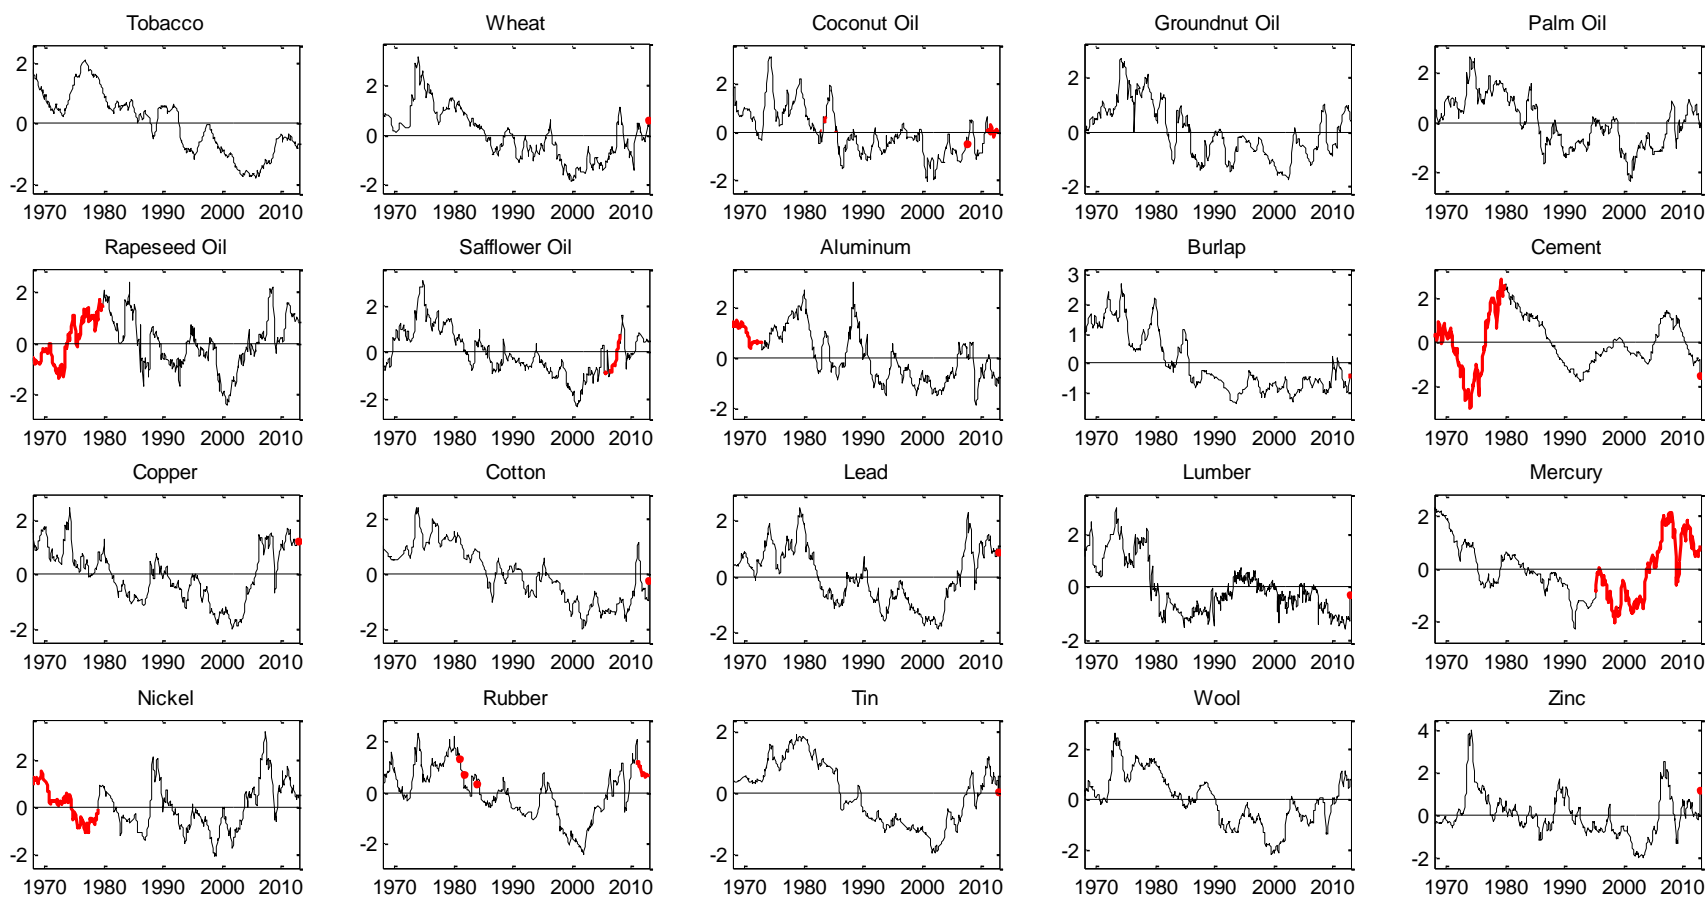
Note: The figure plots real commodity prices (black lines) and imputed values (bold red values) from the EM algorithm of Stock and Watson (2002).

Appendix Figure 3: Indirect Aggregate Common Factor from Subset of Commodities with "No First Order Speculation"

Note: The figure presents the 99% confidence interval of the (HP-filtered) IAC factor from the factor analysis on the full cross-section of commodities in section 3.3 using the estimated rotation parameters from GMM estimates (dark 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 5. Confidence intervals are 3-month moving averages. See section 4 for details.

Appendix Table 3: Recursive Forecast Error Diagnostics for Real Commodity Prices

	$h = 1$	$h = 3$	$h = 6$	$h = 12$	Forecast Evaluation Period
Agr./Food Commodities					
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
Oils					
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
Industrial Commodities					
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 4: Recursive Forecast Error Diagnostics for Real Commodity Prices

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Agr./Food Commodities				
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**
Oils				
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
Industrial Commodities				
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 5: Summary of Recursive Forecast Accuracy Diagnostics for the Real Price of Oil

Forecast Evaluation Period: 1984:1-2012:8

| | BIC | | 12 lags | |
	Factor	Oil Market	Factor	Oil Market
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

Forecast Evaluation Period: 1992:1-2012:8

| | BIC | | 12 lags | |
	Factor	Oil Market	Factor	Oil Market
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. "Factor" refers to the bivariate VAR forecasting model that includes the commodity-price factor and the real price of oil. "Oil Market" 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 VAR forecast is superior to the no-change forecast and are shown in boldface.