World Shocks, World Prices, And Business Cycles: An Empirical Investigation*

Andrés Fernández† Stephanie Schmitt-Grohé‡ Martín Uribe§

June 10, 2016

Abstract

There is no consensus on the importance of world prices in explaining aggregate fluctuations in individual countries. Existing studies, both theoretical and empirical, concentrate on single measures of the terms of trade constructed in different ways. This paper presents an empirical framework in which, for each country, agricultural, metal, and fuel commodity prices enter separately as mediators of world disturbances. We find that jointly, these three commodity prices explain on average 30 percent of aggregate fluctuations in a group of 138 countries. This contribution lies in between the range of values obtained by existing empirical and theoretical studies.

JEL classification: F41.

Keywords: Commodity prices, Terms of Trade, Business cycles.

*We thank Santiago Téllez for excellent research assistance.

†Inter American Development Bank. E-mail: andresf@iadb.org.

‡Columbia University, CEPR, and NBER. E-mail: stephanie.schmittgrohe@columbia.edu.

§Columbia University and NBER. E-mail: martin.uribe@columbia.edu.

1 Introduction

The conventional wisdom is that disturbances in the terms of trade represent a major source of aggregate fluctuations in both developed and developing countries. This view is, to a large extent, based on the predictions of calibrated open economy real business-cycle models (Mendoza, 1995; Kose, 2002). However, recent empirical work based on structural vector autoregression (SVAR) models suggests that terms of trade shocks explain on average only 10 percent of variations in output and other indicators of aggregate activity in poor and emerging countries (Schmitt-Grohé and Uribe, 2015). One criticism of this approach is that standard measures of the terms of trade are based on weighted averages of all export and import prices. At business cycle frequency, the argument goes, commodity-price fluctuations are a more relevant source of disturbance than broad price indices (Fernández, González, and Rodríguez, 2015; Shousha, 2015). An implication of this argument is that a better measure of a country's terms of trade should consist of primary commodity prices. The literature that has adopted this approach has limited attention to a small group of countries in which commodities represent a major share of exports, which is not suitable for deriving general conclusions on the role of world price shocks as drivers of business cycles.

The present study investigates the joint contribution of agricultural, metal, and fuel commodity prices on aggregate activity. For that purpose, we construct a comprehensive dataset of country-specific national income and product account (NIPA) variables together with indices of world prices of these three commodity goods spanning the period 1960 to the present and covering 138 countries, thereby allowing us to cover a large number of developed, emerging and poor countries. We then use this data to estimate country-specific SVAR models with which we obtain new estimates of the contribution of world relative prices to business cycles. After accounting for potential short sample biases, we find that commodity prices account for about one third of movements in aggregate activity. This number is in between the large role assigned to world prices in theoretical studies and the low values obtained by using measures of the terms of trade in empirical studies.

We conduct several robustness and extensions checks. First, we find that results hold across levels of development, that the preponderance of world shocks is inversely related to the country size, and increases for countries that are net commodity importers. Results, however, do not seem to change when categorizing countries only based on whether they export or import oil. Cluster analysis shows that there is, nonetheless, considerable heterogeneity across countries, with three distinctive groups of countries with roughly similar number in each one: those for which world price shocks account for less than 10 percent, an intermediate group for which shocks account for about a third, and a group of countries where the contribution rises to almost two thirds. Moreover, when we expand the set of world price shocks beyond commodity prices to include also world interest shocks, the total contribution of world price shocks to total variance increases to about half for the case of real income. Similar estimates are also found when we run an analysis based on quarterly data. We also compare results when terms of trade are used instead of commodity prices and find that the role assigned to world price shocks is considerably reduced. And qualitatively similar results in terms of the important role of world prices for business cycles are obtained when we consider alternative filtering techniques.

We conduct two last extensions. First, we run a counterfactual simulation where we investigate what would have been the behavior of the domestic NIPA variables if only world shocks had been observed. We find the behavior of the artificial time series to track pretty closely that of the observed variables, particularly since the 1990s. Lastly, we zoom in on the last decade or so that included not only the Great Recession but also the period sometimes referred to as the financialization of commodity markets. We find evidence of an increase in the role of commodity price shocks when accounting for business cycles in that period.

The rest of this work is divided into 8 sections including this Introduction. Section 2 provides more details of the construction of the dataset. Section 3 presents some empirical regularities of the commodity price data. Sections 4 and 5 describe our empirical approach for modeling, respectively, the commodity processes and their interaction with domestic macro variables in the countries we consider. Section 6 contains our benchmark results in terms of the relevance of world shocks for business cycles. Section 7 presents several robustness and extensions. Section 8 concludes.

2 The Data

In our work we use data on commodity prices, world interest rates, terms of trade, and country-specific national income and product accounts (NIPA). Data on commodity prices come from the World Bank's Pink Sheet. This is a public dataset that produces monthly series based on nominal US dollars since January 1960.[1] In particular, we focus on three aggregate commodity price indices: Fuel, Agriculture, and Metals and Minerals. The fuel index is constructed using a weighted average of spot prices of coal, crude oil and natural gas. The agricultural index is also a weighted average of, in turn, three subindices on beverages (including prices on cocoa, coffee and tea), food (computed using other indices on oils & meals, grains, and others), and raw materials (including timber, and other raw materials). The metals and minerals index is based on the spot prices of aluminum, copper, iron ore,

[1] See http://www.worldbank.org/en/research/commodity-markets

lead, nickel, steel, tin, and zinc. All three indices are deflated by the monthly US CPI, and are then annualized by taking simple averages over the twelve months of the year.

In terms of country-specific variables, we focus on annual data on terms of trade and four NIPA variables: GDP, private consumption, investment, and the trade balance. NIPA variables are already in constant local currency units. The sources for these variables are either World Bank's World Development Indicators (WDI) or IMF's World Economic Outlook (WEO) which contain annual information on a wide variety of countries since the early 1960s.[2] When selecting among the two sources we proceed as follows. First, for each country available in each source we form an unbalanced panel with the five variables using all available information. Then, for each of the two sources, we form the longest balanced panel (within country) with the five variables and compare the range across the two sources. We discard countries for which no balanced panel can be formed of a minimum of 25 observations in either of the two sources. Then, for each country that is left, we select the source that delivers the longest balanced panel between the two. If the range happens to be identical in the two, we use WDI as default. This delivers a sample of 138 countries with an average panel per country of 38 annual observations.

In one robustness extension, we use world interest rates that we proxy with the real exante 3-Month US TBills rate. Monthly data on (annualized) yields since January 1960 are deflated by taking an average of the US inflation rate in the previous 12 months, and are then annualized by taking simple averages over the twelve months of the year. In another robustness test we use quarterly data on real GDP. For that we use IMF's International Financial Statistics, OECD (when applicable), Latin Macro Watch, and national sources (central bank websites). The data is already in constant local currency units and no splicing is done. We only ran a standard seasonal adjustment when needed. This delivers a sample of 86 countries with an average time length per country of 116 quarterly observations.

3 Commodity Prices: Some Empirical Regularities

The left panel of figure 1 displays the level of the real price of three groups of commodities, agricultural, fuels, and metals. All prices are deflated using the U.S. CPI index, expressed in logs, and normalized to 1960=0. The three commodity prices display a similar pattern. In particular, all experience large increases in the mid 1970s. Fuels experienced the largest price increases during this period, ending the decade 7 times as expensive as at the beginning

[2]WDI is publicly available at the web (see http://data.worldbank.org). IMF's WEO is also publicly available but not for all NIPA variables. We use an Appendix of the WEO that the IMF shares with other multilateral organizations. The latest WEO dataset starts in 1980 for most countries. We spliced backwards the data using historical WEO reports starting in 1970.

Figure 1: Real Commodity Prices: Level and Cyclical Component, 1960-2014

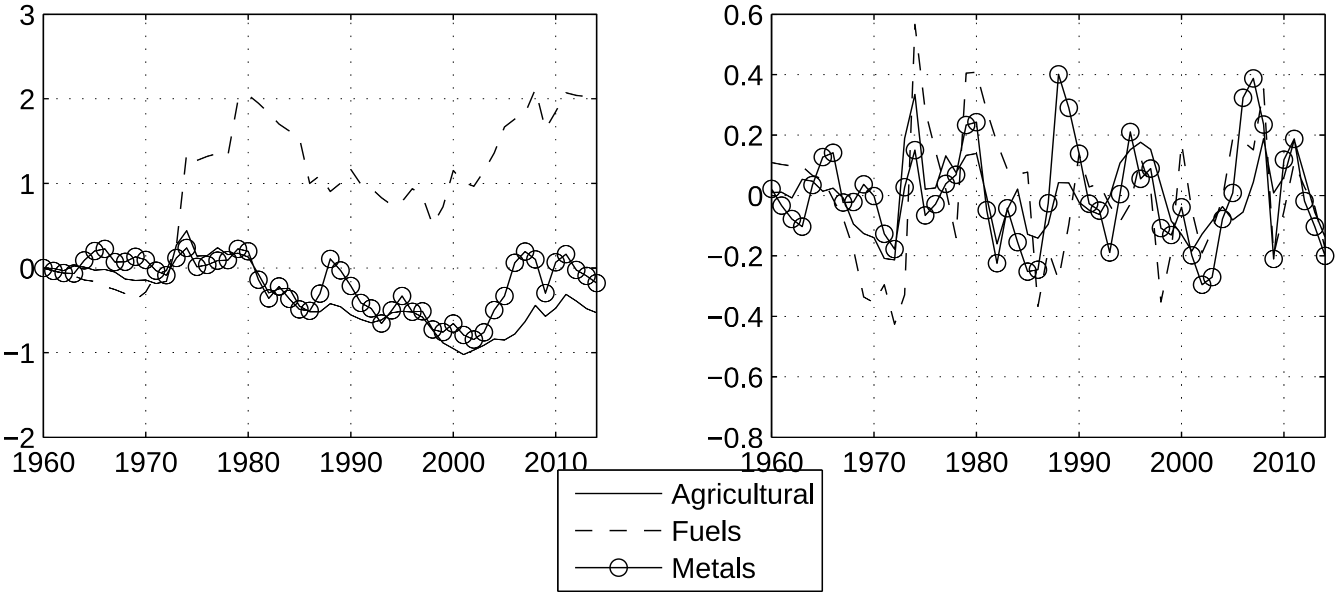

Note. The left panel displays the log of real commodity prices normalized to 1960=0. The right panel displays the cyclical component of the log of real commodity prices as measured by the HP filter with smoothing parameter 100.

thereof. The 1980s and 1990s were characterized by gradual declines. Then, in the early 2000s all three prices recovered vigorously until the Great Recession, which was accompanied by widespread declines in commodity prices. Prices in all three commodity goods recovered vigorously but for a brief period only to fall again.

The right panel of figure 1 displays the cyclical component of the real price of three groups of commodities, as captured by the HP filter with smoothing parameter 100. Two characteristics stand out. First, real commodity prices are highly volatile, especially metals and fuel, with percent deviations from trend in the interval ±40 percent. Second, the three real prices display positive comovement. These features are confirmed in Table 1, which shows a number of empirical second moments. The standard deviation of deviations of real commodity prices from trend ranges from 12 to 21 percent. This means that commodity prices are between 3 and 5 times as volatile than the average country in our sample of 138 countries. In the table, positive comovement is measured by the cross correlation across prices, which is about 0.45 on average. Finally, movements in commodity prices are moderately persistent, with a serial correlation of about 0.35.

Table 1: Commodity Prices: Second Moments

Statistic	p^a	p^f	p^m
$\sigma(p)$	0.12	0.21	0.17
$\rho(p)$	0.57	0.47	0.52
$\rho(p^a, p)$	1.00	0.49	0.59
$\rho(p^f, p)$	0.49	1.00	0.35
$\rho(p^m, p)$	0.59	0.35	1.00
$\sigma(p)/\sigma(GDP)$	2.69	4.98	3.90

Note. All prices are deflated by the U.S. CPI index, in logs, and HP filtered with smoothing parameter 100. The relative standard deviation with respect to GDP is an average over the 138 countries in the sample. Annual data from 1960 to 2014.

4 The Commodity-Price Process

We assume that world commodity prices are exogenous to each specific country. We therefore formulate a VAR specification for the joint evolution of agricultural, fuel, and metal commodity prices that is independent of domestic macroeconomic indicators in individual countries. Formally, let

$$p_t = \begin{bmatrix} p_t^a \\ p_t^f \\ p_t^m \end{bmatrix}$$

where p_t^a, p_t^f, and p_t^m denote the cyclical component of the logarithm of real world prices of agricultural, metal, and fuel commodities, respectively. We assume that p_t evolves according to the following autoregressive system

$$p_t = Ap_{t-1} + \mu_t, \tag{1}$$

where μ_t is a random vector with mean zero and variance-covariance Σ_μ. We interpret the vector μ_t as representing a combination of world shocks affecting commodity prices. The present investigation is not concerned with the identification of specific world shocks (such as, for example, shocks to the world supply of oil). Instead, our focus is to ascertain what fraction of business-cycle fluctuations in individual countries is due jointly to world shocks and mediated by the world commodity prices included in p_t. That is, we are interested in estimating the joint contribution of μ_t to domestic business cycles in individual countries. For this purpose, no further identification assumptions on the above system are required.

We estimate the VAR system by ordinary least squares (OLS) equation by equation using annual data from 1960 to 2014. Prior to the estimation, we detrended the logs of real commodity prices using the HP filter with smoothing parameter 100. The estimates of A and Σ_μ are

$$A = \begin{bmatrix} 0.64 & -0.14 & 0.07 \\ 0.58 & 0.29 & 0.11 \\ 0.03 & -0.21 & 0.61 \end{bmatrix}; \quad \Sigma_\mu = \begin{bmatrix} 0.0084 & 0.0063 & 0.0073 \\ 0.0063 & 0.0312 & 0.0091 \\ 0.0073 & 0.0091 & 0.0190 \end{bmatrix},$$

$$R^2 = \begin{bmatrix} 0.38 & 0.32 & 0.33 \end{bmatrix}.$$

R^2 statistics indicate that about 80 percent of movements in prices is explained by contemporaneous disturbances and the remaining 20 percent by the autoregressive component.

5 The Domestic Block

Let Y_t denote a vector of domestic macroeconomic indicators of country i, such as output. We assume that Y_t evolves according to the expression

$$Y_t = Bp_{t-1} + CY_{t-1} + Dp_t + \epsilon_t, \tag{2}$$

where ϵ_t is an innovation with mean 0 and variance-covariance matrix Σ_ϵ. Note that because p_t appears contemporaneously on the right-hand side of this expression, the innovation ϵ_t is independent of the innovation μ_t. We interpret ϵ_t as a vector of country-specific shocks. This interpretation is based on the fact that the typical country in our sample of 138 countries is a small open economy. As such, world shocks can have an effect only through world prices, such as commodity prices or the world interest rate. For now, we leave the world interest rate out of the system, but will incorporate it later as part of a robustness analysis.

We estimate the domestic block (2) by OLS for each of the 138 countries in the sample. We consider four domestic macroeconomic indicators, output, consumption, investment, and the trade-balance-to-output ratio, denoted y_t, c_t, i_t, and tby_t, respectively. All variables are expressed in real terms and detrended using the HP filter with smoothing parameter 100. All variables are expressed in logs before detrending, except for tby_t. The longest balanced sample of these 4 indicators contains 55 annual observations from 1960 to 2014 and occurs in 5 countries. The shortest balanced sample contains 25 observations and occurs in 7 countries. And the average balanced sample has 38 observations from 1977 to 2013.

Combining (1) and (2) we obtain the following autoregressive representation for the joint behavior of p_t and Y_t

$$\begin{bmatrix} p_t \\ Y_t \end{bmatrix} = F \begin{bmatrix} p_{t-1} \\ Y_{t-1} \end{bmatrix} + G \begin{bmatrix} \mu_t \\ \epsilon_t \end{bmatrix}, \tag{3}$$

where

$$F = \begin{bmatrix} A & \emptyset \\ DA+B & C \end{bmatrix}; \quad G = \begin{bmatrix} I & \emptyset \\ D & I \end{bmatrix}; \quad E\begin{bmatrix} \mu_t\mu_t' & \mu_t\epsilon_t' \\ \epsilon_t\mu_t' & \epsilon_t\epsilon_t' \end{bmatrix} = \Sigma \equiv \begin{bmatrix} \Sigma_\mu & \emptyset \\ \emptyset & \Sigma_\epsilon \end{bmatrix} \tag{4}$$

Given estimates of A, B, C, D, Σ_μ, and Σ_ϵ, one can use this representation to obtain an estimate of the contribution of world shocks (μ_t) to movements in domestic macroeconomic indicators (Y_t).

Formally, to compute such contribution we proceed as follows. First, we express the unconditional variance of Y as

$$vec\left[var\left(Y\right)\right] = \left[I - (F \otimes F)\right]^{-1} vec\left(G\Sigma G'\right)$$

To get the variance of $\mathbf{Y}$ in the SVAR conditional on only the commodity shocks we modify Σ as follows

$$\Sigma^{Co} = \begin{bmatrix} \Sigma & \mathbf{0} \\ \mathbf{0} & 0 \end{bmatrix}$$

and compute the (now conditional) variance as

$$vec\left[var\left(Y\right)^{Co}\right] = \left[I - (F \otimes F)\right]^{-1} vec\left(G\mathbf{\Sigma}^{Co}G'\right)$$

Then the share of variance explained by the (reduced form) shocks in commodity prices is

$$\frac{vec\left[var\left(Y\right)^{Co}\right]}{vec\left[var\left(Y\right)\right]}$$

6 How Important Are World Shocks?

When all four domestic macroeconomic indicators (y_t, c_t, i_t, and tby_t) are included in the vector Y_t, each equation of the domestic block contains 11 regressors, namely, 3 contemporaneous commodity prices, 3 lagged commodity prices, 4 lagged domestic indicators, and a constant (not shown in the derivations above). Since the number of observations for the domestic block ranges from 25 to 55 across countries, we have that for some countries including 11 regressors results in a relatively small number of degrees of freedom. For this reason, we estimate the domestic block in two ways. One is to include all four indicators in the vector Y_t, which imposes the maximum strain on the degrees of freedom. The other is to include in Y_t only one domestic indicator at the time, and estimate the domestic block four times per country, once for each indicator. We refer to the first approach as joint estimation and to the second as sequential estimation.

An issue that must be taken into account is the possibility of a small-sample upward bias in the estimation of the contribution of world shocks to the variance of domestic macroeconomic indicators. The fact that the variance is a positive statistic means that any correlation between the vector of commodity prices p_t and the vector of macroeconomic indicators Y_t results in some participation of world shocks in the variance of Y_t. In particular, even if p_t and Y_t were independent random variables, any spurious correlation (positive or negative) in finite sample will result in a positive share of world shocks in the variance of Y_t, creating an upward bias that exaggerates the importance of commodity prices. This bias is increasing in the number of commodities entering p_t and in the variance of p_t and decreasing in the sample size. Correcting this bias is therefore particularly important when one compares one-price SVAR specification (e.g., specifications including only the terms of trade) with multiple-price specifications, like the one studied thus far. In addition, as is well known, OLS estimates of SVAR coefficients are typically biased in short sample. To address these problems, we apply a Monte Carlo approach to correct for bias. The procedure consists of the following steps: (1) Let $\hat{F}$, $\hat{G}$, and $\hat{\Sigma}$ denote the estimates of F, G, and Σ obtained using actual data from a given country. Let $\hat{\sigma}$ denote the associate estimate of the share of the variance of Y_t explained by μ_t. Use these matrices to generate an artificial data set of a desired length (here 250) from the model (3). (2) Let T^p denote the sample size of commodity prices (55 years in our data set) and T^y denote the balanced sample size of the four NIPA indicators for the country in question (in our sample, $T^y \leq T^p$). Then, use the last T^p observations of the artificial data to estimate the p_t block of the SVAR (i.e., the matrices A and Σ_μ. Use the last T^y observations of the artificial time series to estimate the domestic block (i.e., the matrices B, C, D, and Σ_ϵ). (3) steps (1) and (2) yield an estimate of the matrices F, G, and Σ from artificial data. Use this estimate to compute the share of the variance of Y_t explained by μ_t shocks. (5) Repeat (2)-(3) a desired number of times (here 1e3 times) and compute averages of the resulting estimates of F, G, Σ, and σ (the vector of shares of the variances of Y_t explained by μ_t shocks). Denote these averages $\bar{F}$, $\bar{G}$, $\bar{\Sigma}$, and $\bar{\sigma}$. Then the estimated bias in the contribution of world shocks to the variance of Y_t is given by $\bar{\sigma} - \hat{\sigma}$.

Table 2 displays the shares of the variances of output, consumption, investment, and the trade-balance-to-output ratio explained by world shocks. Both the sequential and joint estimation approaches deliver the same message. Before correcting for small-sample bias,

Table 2: Share of Variances Explained by World Shocks

Cross Country Median	y	c	i	tby
Sequential Estimation				
Noncorrected Estimate	0.44	0.34	0.34	0.29
Small-Sample Bias	0.10	0.13	0.12	0.13
Corrected Estimate	0.34	0.22	0.21	0.16
MAD of Corrected Estimate	0.20	0.15	0.15	0.15
Joint Estimation				
Noncorrected Estimate	0.46	0.37	0.39	0.35
Small-Sample Bias	0.11	0.13	0.13	0.14
Corrected Estimate	0.35	0.25	0.26	0.21
MAD of Corrected Estimate	0.20	0.18	0.18	0.15

Note. Mad stands for the cross-country median absolute deviation. Statistics are computed across 138 countries.

world shocks are estimated to explain on average 44 percent of business cycles in real GDP. The small-sample bias median is on about 10 percentage points. Thus, after correcting for small-sample bias, we find that world shocks explain about a third, 0.34, of observed business cycles. This number lies in between the high values of over 50 percent obtained using theoretical models (Mendoza, 1995; and Kose, 2002) and the low values of around 10 percent obtained using empirical models with world prices captured by the terms of trade (Schmitt-Grohé and Uribe, 2015). The estimated contribution of world shocks, however, is far from homogenous across countries. The associated median absolute deviation is 20 percentage points. This means that across countries most of the estimated variance shares lie in an interval ranging from 14 to 54 percent, which comprises high and low values of the type found in the aforementioned sources. The associated variance share for the three remaining NIPA variables is still high although marginally lower than that of GDP, in between 0.16, for trade balance, and 0.22/0.21 for consumption and investment.

7 Extensions and Robustness

7.1 Level of Development

In panel A. of Table 3 we disaggregate our benchmark results by categorizing the 138 countries according to their level of development. In particular we use four levels of development coming from the World Bank's WDI indicators in 2015 (number in parenthesis are the number of countries in each category): Low Income (22); Lower Middle Income (33); Upper Middle Income (31); and High Income (52).

Results are fairly robust across groups. In terms of real GDP, there is no clear differences across groups and no single group is radically different from the benchmark median results reported in the upper panel of Table 3. Low Income countries display a moderately lower share of variance accounted by world price shocks, 0.23, relative to the benchmark, 0.34. Likewise for Upper Middle Income countries, 0.25. However, Lower Middle Income countries display a slightly higher variance share, 0.37. And in High Income countries the share is identical to the median.

For the case of two of the remaining NIPA variables, consumption and investment, there is a positive relationship between the level of development and the share of variance accounted for world price shocks. This, however, is no longer true for the last NIPA variable considered, the trade balance share. For this case the Low and Upper Middle Income countries exhibit the highest shares, 0.24 and 0.23, and the Low and High Income countries the lowest shares, 0.16 and 0.13.

7.2 Size

In panel B of Table 3 we disaggregate the benchmark results by categorizing countries according to their size in the world economy. We do so by breaking up the distribution of GDP in current US Dollars in 2013 for all countries in the sample into five equally massed groups. We call them countries of small (27), medium-low (27), medium (28), medium-high (28), and large (26) size, where, again, numbers in parenthesis are the number of countries in each category. When doing this we loose two countries, Syria and Taiwan, due to lack of data availability for GDP in US Dollars in 2013.

Results again are fairly robust across groups. For real GDP small countries exhibit the same degree of contribution of world shocks than the median country in the entire sample, 0.34. Medium low, medium, and medium high countries display a slightly smaller contribution between 0.25 and 0.27. This increases for the group of large countries up to 0.42. This, however, may come spuriously from the fact that, for this particular group of countries, the assumption of independence between the domestic block and the commodities block may be less tenable insofar as movements in economic activity could potentially drive these prices.

For consumption and investment, results indicate that, marginally, the share of variance accounted for by world shocks is lower for small and medium-small countries relative to medium and medium-high countries. Lastly, for the trade balance share no clear differences

Table 3: Robustness Tests

Seq. Est. (corrected estimate)	**No. Countries**	**Share in total**	**Y**	**C**	**I**	**TB/Y**
Median	138*	100.0	0.34	0.22	0.21	0.16
MAD			0.20	0.15	0.15	0.15

		ROBUSTNESS				
Seq. Est			**Y**	**C**	**I**	**TB/Y**
A. Lvl. Dvpt						
- Low Income	22	15.9	0.23	0.18	0.14	0.24
- Lower Mid. Inc.	33	23.9	0.37	0.19	0.17	0.16
- Upper Mid. Inc.	31	22.5	0.25	0.21	0.22	0.23
- High Inc.	52*	37.7	0.34	0.24	0.30	0.13
B. Size**						
- Small	27	19.6	0.34	0.18	0.17	0.11
- Medium-low	27	19.6	0.25	0.11	0.16	0.16
- Medium	28*	20.3	0.29	0.23	0.20	0.15
- Medium-high	28	20.3	0.27	0.23	0.21	0.16
- Large	26	18.8	0.42	0.29	0.42	0.26
C. Oil***						
- Exporters	27*	19.6	0.36	0.22	0.22	0.28
- Importers	107	77.5	0.33	0.21	0.20	0.15
D. Net Comm. Trade***						
- Exporters	51*	37.0	0.25	0.21	0.18	0.18
- Importers	83	60.1	0.36	0.22	0.27	0.15
E. Cluster Analysis						
- Low			0.06	-0.01	0.07	0.02
No. Countries/Share			41/0.3	47/0.34	60/0.48	68/0.49
- Medium			0.33	34.00	0.32	0.29
No. Countries/Share			53/0.38	65/0.47	55/0.36	50/0.36
- High			0.61	0.59	0.62	0.54
No. Countries/Share			43/0.31	26/0.19	23/0.17	20/0.14

Table 3: Robustness tests (continued)

Seq. Est	No. Countries	Share in total	Y	C	I	TB/Y
F. Including Interes Rates						
- Seq. Estimation (median)	138*	100.0				
Corrected Estimate			0.45	0.30	0.33	0.23
MAD of Correct Estimate			0.15	0.16	0.16	0.17
G. Quarterly Dataset						
- Seq. Estimation (median)					**Y**	
Minimum No. of Quarterly Obs.			40	60	80	100
No. Countr./Share in Total (%)			86/100	77/90	64/74	38/44
Corrected Estimate			0.41	0.41	0.37	0.33
MAD of Correct Estimate			0.18	0.18	0.19	0.16
H. Terms of Trade Vs Pcom						
- Pcom (same sample as ToT)	138*	100.0				
Corrected Estimate			0.32	0.18	0.23	0.15
MAD of Correct Estimate			0.19	0.14	0.16	0.12
- ToT	138*	100.0				
Corrected Estimate			0.06	0.05	0.04	0.08
MAD of Correct Estimate			0.10	0.10	0.09	0.11
I. Alternative Filters						
- HP Filter ($\lambda = 6.25$)	138*	100.0				
Corrected Estimate			0.23	0.16	0.14	0.11
MAD of Correct Estimate			0.15	0.12	0.12	0.09
- Linear Quadratic Trend	138****	100.0				
Corrected Estimate			0.24	0.24	0.23	0.20
MAD of Correct Estimate			0.18	0.20	0.18	0.16

Note. (*) For the case of Y results exclude Oman as the sequantial VAR was not found to be stable. (**) When categorizing countries by size we drop Syria and Taiwan due to lack of data availability in 2013. (***) When categorizing countries by oil exports or net commodity exporters we drop Angola, Haiti, Myanmar, and Taiwan due to lack of information on the trade shares on commodities. (****) For the case of Y results exclude Ireland, Mongolia, Oman and Spain, and for C Ireland, Netherlands and Spain as the sequantial VAR was not found to be stable.

arise across the first four groups except that the smallest countries display a relatively lower contribution of world shocks.

7.3 Oil

In panel C of Table 3 we consider categorizing countries according to their net trade in fuel oil. We do so by computing the median country-specific net export trade of fuels. We do so by using annual information on exports and imports of fuel commodities from WDI since the 1960s. We categorize a country as oil exporter (importer) if the median net fuel export share of GDP is positive (negative). According to this criterion we identify 27 oil exporters and 107 importers.[3]

Results do not differ much between net oil exporters and importers in most variables. Only the former group displays a slightly higher role of world price shocks in terms of the variance of real GDP, consumption and investment. For the trade balance share, however, the variance share in exporters nearly doubles that of importers.

7.4 Net Commodity Trade

We also consider the joint net trade position in all three commodities that we study. We do so by computing the median country-specific net trade in fuels, metals and raw agricultural goods. We classify 51 countries as net commodity exporters and 83 as net importers.[4] Panel D of Table 3 reports the results. Unlike in the previous case when we only considered oil, we see a moderate tendency for net commodity importers to display a larger role of world shocks when accounting for the variance of real GDP, investment, and to a lesser degree, consumption. For the trade balance, however, we continue to obtain that world shocks are more relevant for exporters than importers.

7.5 Cluster Analysis

As mentioned before, while the median variance share accounted for by world shocks is large, there is also considerable dispersion across countries. We zoom in on this by arbitrarily clustering the results into three groups that can be thought of as groups with "low", "medium" and "high" variance share. This method partitions the results by minimizing the sum, over the three clusters, of the within-cluster sums of (squared Euclidean) distances with respect

[3] When categorizing countries by oil exports or net commodity exporters we drop Angola, Haiti, Myanmar, and Taiwan due to lack of information on the trade shares on commodities.

[4] Idem.

to the mean of each cluster. This allows to get an endogenous number of countries within each of the three categories defined.

The results are reported in panel E of Table 3. Across the four NIPA variables considered, the "low" variance share group accommodates between 41 to 68 countries, which represent between 30 to 49 percent of the sample. For these countries the median variance shares to from virtually nil, in consumption and the trade balance share, to about 6-7 percent for GDP and investment. The number of "medium" countries, in turn, ranges between 50 and 65, or 36 to 47 percent of the sample and the medium share varies from about a third in income, consumption and investment to 0.29 for the trade balance. Lastly, the group of high variance share, ends up being relatively smaller than the two other groups with a number of countries that ranges between 20 and 43, or 15 to 31 percent. For this group of countries the variance share ranges between 54 percent, for the trade balance share, to close to two thirds for the remaining NIPA variables.

7.6 Only World Shocks

As an extension to the benchmark results we ask the following question: What would have been the behavior of the domestic block if only shocks to the external block had been observed? To address this question, first, we back out the reduced form shocks as the fitted residuals, $\widehat{\mathbf{\Psi}}_t \equiv [\widehat{\mu}_t, \widehat{\epsilon}_t]'$, from the SVAR model. Then define

$$\widehat{\mathbf{\Psi}}_t^W = \begin{bmatrix} \widehat{\mu}_t \\ \mathbf{0} \end{bmatrix} \text{ for all } t = 1, ..., T$$

and obtain a counterfactual series for $\left\{\mathbf{Y}_t^W\right\}_{t=1}^T$ by simulating the SVAR model using only $\widehat{\mathbf{\Psi}}_t^W$ as residuals. Lastly we compare the observed process for $\mathbf{Y}_t$ against the counterfactual $\mathbf{Y}_t^W$.

The results of this counterfactual experiment are reported in Figure 2. The observed data for each of the four NIPA variables is in solid/blue while the simulated one is dotted/red. For both observed and simulated data, the information reported for each year is the median across the available countries for which the sequential estimation was made. Overall the results from the counterfactual experiment track the data pretty well, particularly for the last years of the sample that include the Great Recession. The correlation between the observed and simulated data is 0.74, 0.53, 0.59 and 0.39 for output, consumption, investment and the trade balance share. The numbers, however, increase to 0.83, 0.84, 0.77 and 0.71 when considering the period 1990 onwards.

Figure 2: Counterfactual Simulation with Only World Shocks

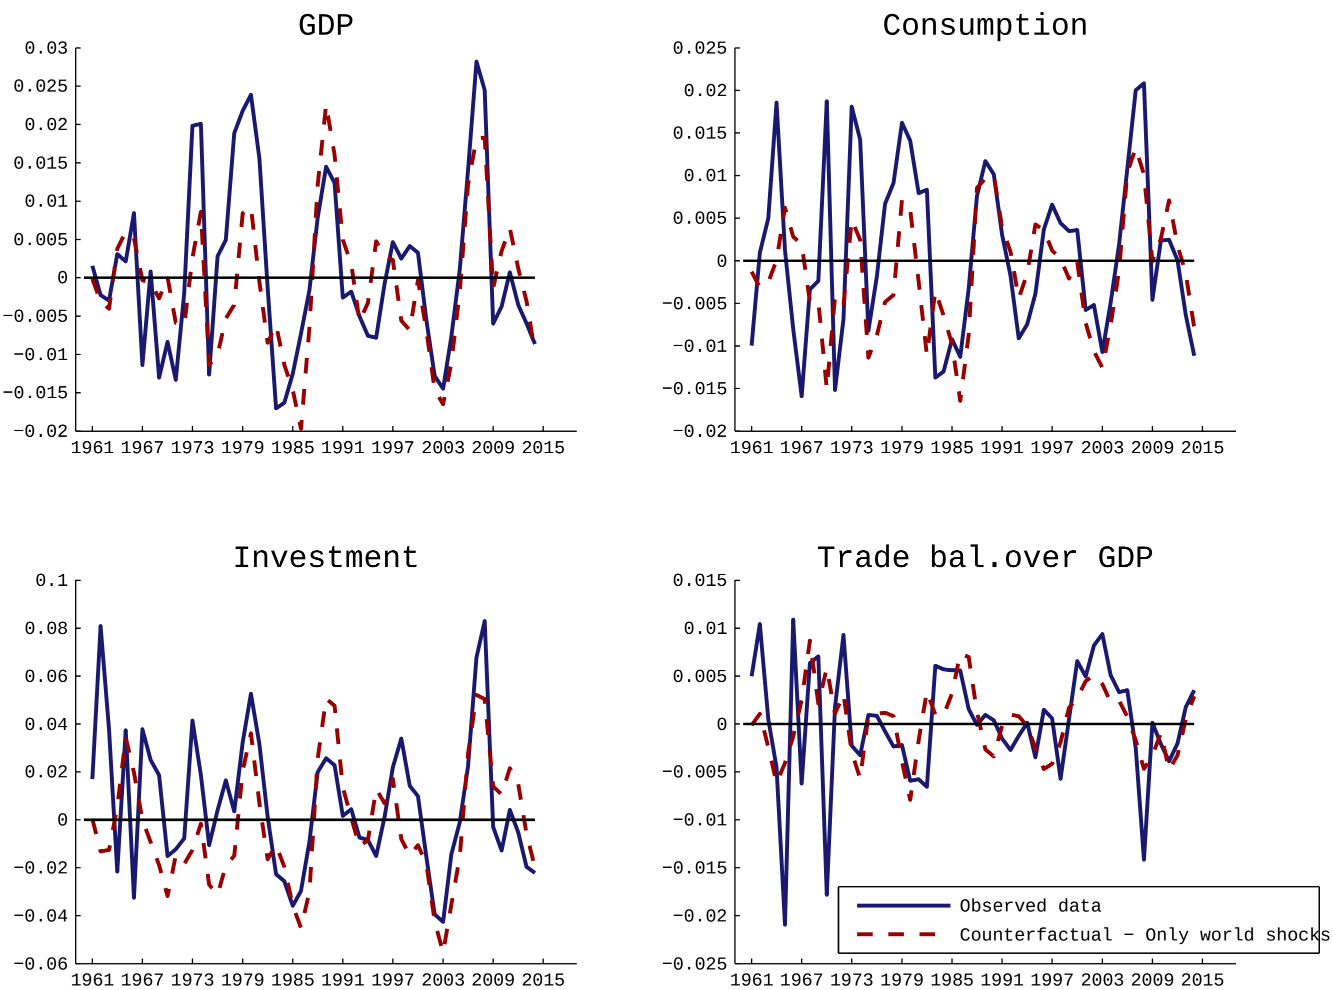

Note: Siumlated model uses only the fitted residuals from the commodities block.

7.7 World Interest Rates

As argued before, our main focus is to ascertain what fraction of business cycles in individual countries is due jointly to world shocks. Thus far we have focused on world commodity prices as the main propagation channel of these world shocks. But they need not to be the only channel. An additional natural candidate is the world interest rate. We thus now explore this by including the world interest rate as a fourth price into the world price block (1). We proxy the world interest rate with the real exante 3-month US TBills rate between 1960 and 2014 (see Section 2 for further details on the data).

Results are reported in panel F of Table 3. All variance shares accounted for the (now expanded) vector of world price shocks increase, even after properly correcting them for the aforementioned bias that comes along with adding more variables to the world price VAR. The largest increases in relative terms are recorded in investment and the trade balance share, whose (corrected) variance shares linked to world price shocks increase from 21 to 33 percent, and 16 to 23 percent, respectively. Those of consumption and GDP also increase considerably. In the former, there is an increase from 22 to 30 percent, and the latter increases almost to half from 0.34.

7.8 Quarterly Dataset

Traditionally, business cycle analysis is conducted over quarterly observations. We have nonetheless carried out our analysis on an annual panel of countries to be able to get as much cross-section and time series amplitude as possible on the four NIPA variables. We nonetheless consider here a robustness extension using quarterly data on real GDP for 86 countries, the maximum number of countries that we could collect information on with a minimum of 40 observations. The data is also unbalanced with an average time length of 116 quarterly or 29 years across countries (see Section 2 for further details on this data). The specification used is identical to the sequential estimation except that now the data in both the world price and domestic blocks are at a quarterly frequency. And, as in the preceding section, we include world interest rates in the world block.[5]

Results are reported in panel G of Table 3. We consider four alternative subsamples according to a minimum threshold of quarterly observations per country. In particular, we consider a minimum threshold of 40 observations (i.e. all 86 countries), 60 (77), 80 (64), and 100 (38). For each case the numbers are smaller, although closer in value, to those obtained with the annual dataset. For the first two larger dataset, the (corrected) variance share is 41 percent, only slightly below that in the benchmark annual case, 45 percent. The numbers

[5]We continue filtering variables with the Hodrick-Prescott filter but use a smooting parameter of 1600.

with the other two dataset, while smaller, 37 and 33 percent, still account for a contribution of about a third of the variance of real GDP.

7.9 Financialization

Some researchers have pointed to the fact that, since the early 2000s, commodity futures have become a popular asset class for portfolio investors, just like stocks and bonds. This process is sometimes referred to as "financialization" of commodity markets (see Cheng and Xiong, 2013, and references therein). A distinctive characteristic of this process is a large inflow of investment capital to commodity futures markets, generating a debate about whether this distorts commodity prices. We now explore the extent to which this process of financialization of commodity prices may have impacted in the role of world shocks when accounting for business cycles in economies around the world.

We do so by two separate but complementary experiments using the quarterly dataset that we collected on real GDP. In Experiment 1, we first estimate the commodity block (1) with three commodity prices for the entire sample 1960.Q1 to 2014.Q4, estimate the domestic block for the entire available sample in each year, and compute the share of the variance of real GDP accounted for by the shocks in the commodity block. For this particular case we focus only on the subset of 38 countries for which we have the longest time series on real GDP (with a minimum of 100 quarterly observations). In the second part of the experiment we re-estimate the commodities block in the subperiod post 2004.Q1of the so called financialization of commodity markets, and then use the new estimates together with the estimates of the domestic block and recompute the variance share accounted for world shocks.

The implicit assumption in Experiment 1 is that, even if financialization did occur and indeed affected the distribution of commodity prices, this did not imply that the way commodity prices affect real activity in economies across the world changed. However, one could also consider the possibility that the domestic economy's response to changes in prices varies with financialization. To account for this case Experiment 2 estimates the commodities and domestic block in the pre-2004 period and, separately, do the same for the post-2004 period. For each of the two cases we compute the share of variance accounted for world shocks on real GDP variance.

The quarterly estimates of A and Σ_μ in the two samples are

$$A_1 = \begin{bmatrix} 0.80 & -0.05 & 0.09 \\ 0.20 & 0.69 & 0.14 \\ 0.08 & -0.06 & 0.85 \end{bmatrix}; \quad \Sigma_\mu^1 = \begin{bmatrix} 0.0020 & 0.0005 & 0.0007 \\ 0.0005 & 0.0117 & 0.0009 \\ 0.0007 & 0.0009 & 0.0041 \end{bmatrix},$$

$$A_2 = \begin{bmatrix} 0.74 & -0.11 & 0.12 \\ 0.38 & 0.58 & 0.13 \\ -0.33 & -0.16 & 0.98 \end{bmatrix}; \quad \Sigma_{\mu}^2 = \begin{bmatrix} 0.0026 & 0.0037 & 0.0032 \\ 0.0037 & 0.0153 & 0.0103 \\ 0.0032 & 0.0103 & 0.0111 \end{bmatrix},$$

where A_1,$\Sigma_1$and A_2,$\Sigma_2$are estimated between 1960.Q1 to 2003.Q4 and 2004.Q1 to 2015.Q4, respectively.

As can be seen there are some noticeable changes relative to the benchmark estimates. The persistence of commodity prices increases for fuels and metals and, more importantly, the fitted residuals display an increase in the estimated variance and covariances.

Results of both experiments are reported in Table 4. The first two columns report the results in Experiment 1 in terms of the (bias corrected) variance shares of world shocks using the parameters estimated for the commodities block for the full sample (column1) and the counterfactual numbers using the estimates of the commodities block for the financialization period (second column). Results of the two estimates from Experiment 2, pre- and post-2004, are reported in columns 3 and 4, respectively. For both experiments we do recover a considerable increase in the share of variance accounted for by world shocks for the financialization period. In Experiment 1, the share increases from 0.26 to 0.40. In the case of Experiment 2 the increase is more radical as the statistic increases from 0.19 to 0.78. The latter is likely coming from the fact that the pre-2004 period is capturing the Great Moderation period in developed economies (which are mostly the countries considered as they tend to be those with fairly long quarterly data), while the post 2004 period is capturing the Great Recession and the wide cross-country effect that such episode had.

Table 4: Financialization

	Experiment 1		**Experiment 2**	
Minimum No. of Quarterly Obs.	100		100	
No. Countries / Share in Total (%)	38/44		38/44	
	Estimated	Counterfactual	Pre-2004	Post-2004
Noncorrected Estimate	0.30	0.40	0.25	0.79
Small-Sample Bias	0.04	0.02	0.07	0.01
Corrected Estimate	0.26	0.40	0.19	0.78
MAD of Correct Estimate	0.13	0.18	0.12	0.10

Note. Experiment 1 "estimated" uses all available information when estimating the world and domestic blocks, while the "counterfactual" uses the post 2004.Q1 period to estimate the world block but keeps the domestic block untouched. Experiment 2 splits the sample into pre and post 2004.Q1 and estimates the two blocks on these two periods.

7.10 Terms of Trade

As argued before, some of the works that have recently focused on commodity prices as business cycle drivers (Fernandez, et.al., 2015; Shousha, 2016) have argued in favor of using spot prices in international markets instead of using the ratio of export and import unit value indices, often called the terms of trade. We now compare the estimates using commodity prices as in our benchmark case with those that come from replacing the three-variable VAR in commodity prices by a simple AR(1) process for the terms of trade.

Because, for some countries in our sample, data on terms of trade may be limited to a shorter sample relative to that of NIPA variables, we must rerun our benchmark estimates in order to take into account the new (possibly shorter) balanced sample that includes NIPA variables and terms of trade. In addition to this, since the goal is to compare the role of commodity prices against that of terms of trade, in order to make the comparison consistent, we re-estimate the process of commodity prices for each country on the subsample of years for which we have data available on terms of trade.

Results are reported in panel H of Table 3. First, the new benchmark results change only marginally relative to those in Table 2. The (bias corrected) variance share of world shocks continues to lie between a third and a fifth for all NIPA variables. The share in real GDP, for instance, is now 0.32, only slightly below the one obtained before, 0.34. However, for this variable, the variance share accounted for by relative prices drops by a factor of about a fourth when terms of trade are used instead, to 0.06. This drop is also observed in the remaining NIPA variables, confirming that the relevance of world prices when trying to account for business cycles can be overshadowed when using terms of trade data. It is beyond the scope of this paper to establish the deep nature of these differences, but probably a good hypothesis to begin exploring this is that terms of trade indices may suffer from biases in unit value indices associated to discrepancies between unit value indices and price indices (see Silver, 2009).

7.11 Filters

When conducting our analysis with annual data we have detrended variables with the Hodrick-Prescott filter with a smoothing parameter of 100. We conduct a simple robustness test on our benchmark results by using two alternative filters. First, we consider also the the Hodrick-Prescott filter but reduce the smoothing parameter to 6.25 as recommended by Ravn and Uhlig (2002) when working with annual data. Second, we use a quadratic trend on all NIPA variables and commodity prices (all in logs, except for the trade balance share)
.

Results are reported in panel I of Table 3. For real GDP both alternative filters reduce the variance share accounted for by world shocks from a third to a fourth. For the remaining three NIPA variables results are not consistent across alternative filters. The modified Hodrick-Prescott filter further reduces the variance shares in all three variables relative to the benchmark, although moderately. The log-quadratic filter, instead, increases the variance share for all three and place it above 20 percent.

8 Concluding Remarks

To be included.

References

Cheng, Ing-haw and Wei Xiong. 2013. "The Financialization of Commodity Markets." Annual Review of Financial Economics, forthcoming.

Fernandez, A., A. Gonzalez, and D. Rodriguez, "Sharing a Ride on the Commodities Roller Coaster: Common Factors in Business Cycles of Emerging Economies," IDB-WP No. 640, 2015.

Kose, M. Ayhan, "Explaining business cycles in small open economies 'How much do world prices matter?'," Journal of International Economics 56, 2002, 299-327.

Mendoza, Enrique, "The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations," International Economic Review 36, February 1995, 101-137.

Ravn, Morten and Harald Uhlig (2001) "On Adjusting the HP-Filter for the Frequency of Observations," CEPR Discussion Paper 2858, June

Schmitt-Grohe, S. and M. Uribe, "How Important Are Terms Of Trade Shocks?" manuscript, NBER, 2015.

Shousha, Samer, "Macroeconomic Effects of Commodity Booms and Busts," manuscript, Columbia University, October 2015.

Silver, M., 2007. Do unit value Export, Import, and Terms of Trade Indices Represent or Misrepresent Price Indices? IMF Working Paper 07/121.