INTERNATIONAL RISK-SHARING AND COMMODITY PRICES

by

Martin Berka, Mario J. Crucini and Chih-Wei Wang

Working Paper No. 11-W21

November 2011

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

INTERNATIONAL RISK-SHARING AND COMMODITY PRICES[1]

Martin Berka, Mario J. Crucini and Chih-Wei Wang

October 19, 2011

Abstract

Cole and Obstfeld (1991) exposited a classic result where equilibrium movements in the terms of trade could make ex ante risk-sharing arrangements unnecessary: a unity elasticity of substitution across goods and production specialization. This paper extends their model to N countries and M commodities ($N > M$). Here the terms of trade provides insurance against commodity-specific shocks, not country-specific shocks. Using commodity-level production data at the national level and world commodity prices we document significant terms of trade variability and positive responses of nation-specific production to terms of trade improvements. The endogenous terms of trade insurance mechanism highlighted in CO is virtually non-existent.

1 Introduction

Financial markets provide a mechanism through which individuals and firms are able to transform the cash flows from their economic activities across time and states of nature. The consumption risk sharing literature attempts to provide an overall assessment of the outcomes of these financial market linkages by examining how closely

[1]Crucini gratefully acknowledges the financial support of National Science Foundation (SES-1030164).

consumption allocations conform to the idealized world in which complete contingent claims shield individuals entirely from idiosyncratic variation in their income. The empirical literature on this topic is broad, ranging from microeconomic approaches that study individual income and consumption dynamics using, for example, the U.S. Consumption Expenditure Survey (CEX) to macroeconomic approaches, using income and product accounts at the national or regional level. Examples of the microeconomic approach are Cochrane (1991) and Mace (1991) while examples of the macroeconomic approach are Canova and Ravn (1996) and Lewis (1996).[1]

At the national level, the empirical risk-sharing literature has been mostly focused on the large industrialized countries, due in part to the dominant fraction of world output that they produce, but also because of the relatively short time spans of reliable developing country data available at the time the literature emerged in the early 1990's. This omission has a number of implications. First, the omission is significant in the sense that the welfare implications of risk-sharing for a majority of the world's population has not been examined. This omission would also be important for studies of industrialized countries in the sense that while individual developing nations contribute very little to global capital flows, in the aggregate, the flow of capital between developing and developed countries is actually quite large and growing rapidly. Second, given the more volatile and idiosyncratic nature of the business cycles in developing countries, these nations evidently have much more to gain from risk-sharing with each other and through integration with industrialized countries than industrialized nations gain from pooling risks among themselves. Thus, the welfare implications of risk-sharing may be understated when the focus is placed on the most economically stable nations who share a large common factor in their business cycle variation.

To gain an appreciation for the heterogeneity in business cycle experiences around the world, consider **Figure 1**. It plots the standard deviation of the growth rates of consumption and GDP for 161 countries using the Penn World Tables (PWT) over the sample period 1971 to 2005.[2] The standard deviation of GDP growth ranges from an astounding 27.5% in Lebanon to a mere 1.88% in the Netherlands; the median country is Samoa (5.16%). The United States, the most scrutinized country in the business cycle literature, is hardly representative, it ranks 149th. OECD nations occupy 16 of the 20 least volatile positions in the ranking. The often-cited business cycle fact that the ratio of the standard deviation of consumption growth relative to output growth is much less than one in the United States is a feat achieved by only 11 of the 161 nations in the PWT (the line at the bottom of the figure presents this ratio by country along with the U.S. benchmark value of 0.69). The median volatility ratio is 1.18. Note, also, nations with more GDP volatility tend to have more consumption volatility: the correlation of output and consumption volatility across nations is 0.76.

The benchmark international risk-sharing model in which nations pool output endowments in a mutual fund and each nation is allocated consumption in proportion to total world output predicts that consumption volatility will be equalized across countries and furthermore, it will be equal to the variance of world GDP. In the data, however, the median consumption standard deviation is 6.63, four times the standard deviation of world GDP growth of 1.69.[3] The standard deviation of world consumption growth is even lower than that of world GDP growth (1.12) suggesting that extending the endowment version of the risk-sharing model to allow for intertemporal substitution would allow consumption variability fall below output variability by a significant margin. Interestingly, the world consumption volatility ratio is 0.66, only slightly lower than the U.S. level. Taken at face value these facts suggest considerable GDP risk at the national level and limited amounts of risk sharing. Below, we provide reduced-form estimates that one-third of national GDP variability is pooled among nations.

With a shift in focus to the full sample of nations of the world it is productive to move from a one-sector macroeconomic model emphasizing the gains from intertemporal trade to a trade model emphasizing variation in the terms of trade. The classic general equilibrium trade model of risk-sharing is Cole and Obstfeld (1991). In their two-country model, each nation receives a stochastic endowment of a perishable commodity. When the elasticity of substitution across goods is unity, the equilibrium terms of trade and relative endowments move in equal and opposite directions. Consequently, the same consumption allocations arise under financial autarky as under complete risk sharing: asset markets are redundant.

We extend the Cole and Obstfeld model in a number of directions to more accurately assess the terms of trade channel. Two crucial theoretical extensions are: i) to allow multiple producers of each primary commodity and ii) to allow for incomplete specialization. By allowing for multiple producers of each primary commodity, we show that the terms of trade insurance mechanism of the original paper remains in effect for the group of nations (or individuals) that specialize in the production of a primary commodity, but provides absolutely no insurance against idiosyncratic variation in production levels across members of the group producing the same commodity. Allowing for incomplete commodity specialization reduces supply-side risk to the extent shocks are good or sector specific, but tends to amplify the impact of nation-specific disturbances because the terms of trade effect is muted. Since our focus is on countries most vulnerable to terms of trade risk, we assume the numeraire good (a composite manufactured good) is produced by one 'country' and consider this to be an aggregate of the G-7 nations. The diversification of production across most of the goods in the consumption basket greatly reduces the output and consumption variance of the G-7 group.

Much of the novelty of the analysis lies in the empirical work, which exploits a panel of physical production data of individual primary commodities at both the national and world level. To incorporate the broadest possible picture of the role of commodity prices in the business cycle variation of nations we use a cross-section of 66 countries. The number of countries is substantially lower than the 161 nations available in the PWT which is a limitation forced by the available panel of commodity production and export data. We find a very substantial idiosyncratic component to annual production changes, which based on the theory cannot be insured by movements in the terms of trade. In fact, the results are even more stark since by measuring production variation it is possible to examine the covariance of production changes in each country with the relative price of the primary commodities that it exports. This covariance is large and positive in most cases which effectively means that supply changes actually reinforce the effect of relative prices on export earnings, opposite of the Cole and Obstfeld mechanism. Put differently, the primary commodity supply shocks that the Cole and Obstfeld theory focused upon appears to be the exception rather than the rule.

To make this a bit more concrete, consider the international coffee market. When the demand for coffee on the global market increases due to a world cyclical expansion, what do we expect will happen in individual countries that depend on coffee for a large fraction of their income and export earnings? If income elasticities of demand are unitary and the world production possibility frontier is shifting out in a neutral fashion (neutral in the sense of all dimensions of the commodity space), all commodity supplies and demands will increase in the same proportion and relative prices and relative demands will remain unchanged. More of everything is produced and consumed at exactly the same relative prices as before the expansion took place. A more realistic expectation is that asymmetric movements in prices and quantities will occur, at least in the short-run. In other words, we would expect changes in the terms of trade. This is an empirical question informed in part by the fact that the typical primary commodity relative price is procyclical, extremely volatile and persistent. Thus in a typical global expansion, the income of the coffee producer is likely to rise due to a terms of trade improvement, which may be reinforced by a positive supply response. This is not meant to completely downplay the role of commodity supply shocks. Exogenous crop failures do have significant implications for nations that concentrate their production in a few primary commodities. It is also true that the concentration of production is sufficiently high that a shock in one or a few locations could move the world price significantly in the manner described by the Cole and Obtsfeld model. Brazil, for example, produces about one-third of the world coffee supply and it together with Colombia and Vietnam, account for more than 50%. Again, an empirical question which our panel data is well-suited to address.

We find that the negative correlation between a relative commodity price and its production level needed to provide the endogenous terms of trade risk sharing in the Cole and Obstfeld model is largely inoperative in the real world. While there are periods in which a commodity relative price increases as the relative quantity of its global production falls, these periods and commodities are the exceptions to the rule. Most of the time the correlation is either ambiguous or positive. Drilling down to the level of a nation makes matters worse since a positive correlation is prevalent, consistent with a positive supply response to a terms of trade improvement. In other words, the quantity movements reinforce the relative price movements on export earnings rather than working to mitigate the effect as would be true in the Cole and Obstfeld model. Finally, the terms of trade are where most of the variance lies in the change in export earnings. This suggests either an inelastic supply curve for the primary commodity due to capacity constraints in the short-run or possibly wealth effects working to offset the substitution effects that would typically lead to an expansion of commodity supply.

Overall we find that the model predicts high consumption variance in nations that have more volatile terms of trade. Overall, the model accounts for about 15% of the cross-sectional variation in consumption volatility in the data. However, the marginal impact of variation in export revenue is too great: as we move from a country with a 10% predicted standard deviation of consumption growth to twice that level, the increase in the standard deviation of consumption growth in the actual data is only about 1.4%, not 10%. As the model assumes complete specialization in primary commodities and abstracts from non-traded goods, it is not surprising that it predicts excess consumption sensitivity to the terms of trade. The main point however, is that the terms of trade is not providing the insurance channel attributed to it by the Cole and Obstfeld benchmark model. In fact, the data seem to point to an important role for changes in world commodity demand that, first of all, move relative prices and relative quantities in the same direction, and second of all, move relative prices by far more than relative quantities, at least in the short-run. The relatively high variance of price compared to quantity seems inconsistent also with the unitary elasticity of substitution in demand. More plausibly, the elasticity of both demand and supply of primary commodities relative to manufactured goods is quite below unity. Finally, national production changes tend to be dominated by nation-specific variation, adding an important idiosyncratic component to national consumption even for producers of the same commodities.

2 The Data

The macroeconomic data on Gross Domestic Product and private consumption for 161 countries is from version 6.3 of the Penn World Tables (PWT). These data are annual and the sample period is 1970 to 2005. The commodity price and physical quantity of production data are from the Commodity Research Bureau Yearbook, 2009. The prices are annual averages of U.S. dollar spot market prices. The quantity data are production data by country and total world production. The difference between the world production estimate and the sum of national production levels is a residual referred to as 'other' in what follows. Subject to availability, the commodity price and production data starts in 1970 and ends in 2005. The commodity export data by country are collected from the United Nations Commodity Trade Statistics (UN Comtrade) Database. The sample periods are: i) 1962-2002 for the SITC classification and 1988-2009 for the HS classification.[4]

2.1 Macroeconomic Fluctuations

When output is modelled as a stochastic and perishable endowment, full risk pooling is achieved by having each nation hold shares in a global mutual fund paying dividends equal to world GDP. A natural starting for a discussion of international risk sharing, therefore, is a variance decomposition of world output.

World GDP growth is defined as a country-size weighted average of national GDP growth rates:

$$\Delta y_t = \sum_i \pi_i \Delta y_{it} \tag{1}$$

where y_{it} is the logarithm of real GDP (Layspeyres) for country i, and π_i is the fraction of world GDP produced by country i, on average, over the period 1971 to 2005. **Figure 2** presents the entire distribution of country sizes, with the largest countries at the bottom of the figure and successively smaller countries added until the total of 100%, is achieved. The extreme skewness of the world country-size distribution is readily evident: of the 161 countries in the full sample, 147 of them contribute significantly less than 1% of world GDP.

Due to the large number of countries, it is economical to use a non-standard variance decomposition:

$$var(\Delta y_t) = \sum_i \pi_i cov(\Delta y_t, \Delta y_{it}) \ . \tag{2}$$

Dividing both sides by the variance of world output growth leads to the variance decomposition:

$$1 = \sum_i \pi_i \beta_i \tag{3}$$

where $\beta_i \equiv cov(\Delta y_t, \Delta y_{it})/var(\Delta y_t)$. By construction, the country-size weighted average β is 1.

The U.S. ranks first in this variance decomposition, which is not very surprising. What is surprising is that the U.S. accounts for 35% of the variance despite producing only 23% of world output (the U.S. share of world GDP in the PWT data). The difference reflects the fact that the beta of the U.S. is 1.52. Put differently, while the U.S. may have one of the most stable growth rates in the world, it is more volatile than world GDP and has a high correlation with world GDP, thus giving it a beta far in excess of 1.

About 15% of the world's nations, 23 to be exact, have negative betas, indicating that the variance of world GDP falls as a consequence of their inclusion in the portfolio. Among this group, China and Indonesia are notable because of their economic size. At the other extreme are countries with β's far exceeding unity, oil exporters tend to fall into this category. Saudi Arabia, for example, contributes one-fifth as much to world GDP variance as the U.S., despite being about 30 times smaller economically.

The fact that the income betas are very different from unity suggests considerable scope for international risk-sharing. For evidence on the extent to which risk-sharing occurs, the consumption side of the ledger is relevant. The analogous variance decomposition of world consumption is, $1 = \sum_i \pi_i \alpha_i$, where $\alpha_i \equiv cov(\Delta c_t, \Delta c_{it})/var(\Delta c_t)$.

As is well-known, complete risk-sharing in the one-sector endowment economy implies the equality of marginal utility across nations. Moreover, under the assumption that consumption is separable from other arguments of the utility function and takes the form, $(1-\sigma)^{-1} c^{1-\sigma}$, complete international risk-sharing implies: $\Delta c_t = \Delta c_{it}$. In terms of the world consumption growth variance decomposition this means the distribution of the α's is degenerate: $\alpha_i = 1$.[5]

Figure 3 presents a scatterplot of the consumption α's against the income β's. The figure also contains three lines. The horizontal line is the relationship between the two parameters under the null hypothesis of complete risk-sharing. The 45-degree line is consistent with national GDP and consumption moving with each other, which is inconsistent with risk-sharing. The third line is a least-squares estimate of the relationship between the α's and the β' in the data. The slope of the line is 0.7, suggesting that 70% of GDP risk is not pooled among nations. In sum, the macroeconomic time series are consistent with a considerable amount of idiosyncratic variation in national GDP available to be pooled and yet the majority of the consumption risks remain specific to the nation.

2.2 Commodity Market Fluctuations

Developing countries would seem to be the ideal setting in which to investigate the Cole and Obstfeld proposition that the positive wealth effects of an increase in the physical output of a commodity are completely offset by a deterioration in the terms of trade. Developing nations are typically highly specialized in exports. Moreover the physical quantity of primary commodity production is arguably more susceptible to exogenous sources of variation than is true of manufactured goods and services produced in industrialized nations.

The properties of commodity price fluctuations have been thoroughly documented in the existing literature. The typical primary commodity price fluctuates by an order of magnitude more than other goods or services in the CPI basket. Economists continue to debate whether they are random walks or cointegrated with the price level.[6] Less is known about the evolution of the quantity of world production and we know of no study that investigates nation-specific variation in production levels. This section intends to fill these gaps in the literature before turning to their risk-sharing implications.

2.2.1 World commodity price and quantity indices

The panel data consist of spot U.S. dollar prices and physical production of 33 primary commodities. The price data are U.S. dollar spot prices from centralized exchanges. The production data are physical production by major producers and world production levels.

The analysis begins at the level of world commodity price and quantity indices, each constructed using sample average production value weights, denoted θ_j:

$$\theta_j \equiv T^{-1} \sum_t \frac{P_{jt} Y_{jt}}{\sum_j P_{jt} Y_{jt}} \tag{4}$$

where P_{jt} is the international price of commodity j in year t and Y_{jt} is world physical production of commodity j in the same year. The commodity price and quantity indices are defined as:

$$\widetilde{z}_{ct} \equiv \sum_j \theta_j z_{jt},\ z = p, y, \tag{5}$$

where lowercase variables are logarithms of their uppercase counterparts (thus the indices should be viewed as geometric averages of the micro-data with production shares as the weights).

The commodity price index is trending over time due to inflation (in U.S. dollar terms) while the commodity quantity index is trending due to the interactions of economic growth and technological innovation. The presumption is that economic growth increases the demand of primary commodities while technological innovation allows producers to economize on these inputs or develop substitutes for them over time. For this reason, we normalized by the U.S. CPI for the price index (p_t) and we normalized by world real GDP (y_t) for the quantity index. We refer to these as the relative price and relative quantity indices in what follows and denote them by: $p_{ct} = \widetilde{p}_{ct} - p_t$ and $y_{ct} = \widetilde{y}_{ct} - y_t$, respectively.

Figure 4 plots the levels of the relative commodity price and quantity indices from 1970 to 2005. For the Cole and Obtsfeld risk-sharing channel to be operating across the industrialized and developing world these two indices would need to be equal in slope and opposite in sign. In fact, they are very unequal in slope (notice different vertical scales) and equal in sign. Primary commodities are becoming relatively inexpensive and relatively scarce at the same time. The property of a secular decline in the relative price of primary commodities is called the Prebisch-Singer hypothesis and has been the subject of intense empirical study for some time. Note that the implication of these two series for the world share of expenditure on primary goods is astounding: the relative price decline is on the order of 100% while the quantity decline is about 20%, thus the 2005 share of world expenditure on primary commodities is about one-half what it was in 1970.

Figure 5 presents the same variables in growth rates to focus on the higher frequency movements in consumption and income that often are the focus of the risk-sharing and business cycle literature. The correlation of the two variables is now ambiguous (-0.08) and the relative volatility of prices is greatly amplified: to a factor of about 10. There are periods in which the correlation of relative prices and quantities is strongly negative, suggestive of a dominance of 'supply shocks.' Taking five-year moving averages of the correlations centered on the early 1970's and the late 1980's and 1990's, the correlation averages -0.6. While interesting, the standard error is very large and these periods are not typical of the last quarter century of evidence.

2.2.2 Deconstruction of world commodity price and quantity indices

It is tempting to relate the annual changes in relative commodity prices to the patterns of business cycle heterogeneity documented in Figure 1. As we move up along the business cycle volatility profile of nations in Figure 1, the commodity concentration of output and particularly exports is rising, suggesting an important role for commodity prices is elevating the volatility of GDP and consumption in small, developing nations, compared to large industrialized ones.

Since primary exports of developing countries tend to be highly concentrated in a single or small number of items, it is instructive to begin this investigation by decomposing the variance of the world commodity relative price and quantity indices into the contributions of individual commodities using the same variance decomposition method employed earlier. The variance decomposition of the commodity price index is thus: $1 = \sum_j \theta_j \gamma_{jc}$, where θ_j is once again the production value weight of commodity j and $\gamma_{jc} = cov(\Delta p_{jt}, \Delta p_{ct})/var(\Delta p_{ct})$ is the contribution of commodity j's relative price to the variance of the world relative price index before production share-weighting. Here we define $\Delta p_{jt} = \Delta \widetilde{p}_{jt} - \Delta p_t$ as the change in the logarithm of the U.S. dollar price of commodity j, $\Delta \widetilde{p}_{jt}$, relative to U.S. CPI inflation (to be consistent with our definition of Δp_{ct}). The analogous variance decomposition of the quantity index is:$1 = \sum_j \theta_j \beta_{jc}$, where $\beta_{jc} = cov(\Delta y_{jt}, \Delta y_{ct})/var(\Delta y_{ct})$.

Table 1 presents the variance decomposition of the commodity relative price index on both an equal and production-value-weighted basis. When commodities are weighted equally in the price index, commodities such as oranges and corn oil contribute far more than the average good and petroleum contributes far less. When world production value weights are used the picture changes because petroleum has a production value share of 36% compared to 3.3% under equal weighting. On a production value-weighted basis, oil now accounts for about 68% of the variance of the commodity price index changes, twice its production share. Only eight other commodities account for 1% or more of the commodity price variance: corn oil, eggs, iron, milk, rice, soybeans, sugar and wheat.

Table 2 presents the variance decomposition of the commodity relative quantity index on both an equal and production-value-weighted basis. Some parallels and some differences emerge in a comparison of price and quantity indices. Petroleum stands out much less starkly in the valued weighted commodity quantity index, accounting about one-third of the variation, which happens to match it's value share. The other difference between the price and quantity data is that unlike commodity prices, the production levels tend to move more independently of one another. There are some commodity production changes which have a negative covariance with the aggregate quantity index, such as iron, oranges, petroleum, tallow and zinc. This means there is more averaging out of quantity changes across commodities than is true of price changes which cautions against casual interpretations of relative price and quantity variation in terms of structural elasticities of supply and demand. The averaging out of quantity changes, though, is evidence of a significant role for commodity-specific supply changes that are not as evident in the relative price series. We turn, next, to nation-specific patterns of specialization and production changes.

2.2.3 National and World Commodity Supplies

Bidarkota and Crucini (2000) investigated the consumption risks faced by primary commodity-exporting developing countries under the assumption that the quantity of production was constant since they lacked production data by country. This section explores the cross-sectional and time series properties of commodity production by major world producers.

For each of our 33 primary commodities we have annual production data for major producers and total world production. The distinction between world and nation-specific variation is important for our analysis. Changes in the world supply of commodities are arguably related to changes in world demand and the equilibrium supply responses those changes evoke. Changes in commodity supply at the level of a country relative to the world are arguably related to exogenous supply shocks specific to that country.

Figure 6 plots the world market shares of coffee production for 12 major producers. Brazil and Columbia are the two dominant producers, accounting for a fairly stable 50% share of world production. By the end of the sample Vietnam has risen from being a marginal producer to contributing a share comparable to Columbia. There is considerable annual variation in Brazil's share of world production.

Figure 7 focuses on petroleum production. Here the production levels are somewhat more concentrated across the 10 major producers displayed than was true of coffee. The U.S. and Russia exhibit secular declines in their share of world production from 35% to about 20%, most of which is picked up by the 'other' category. The breakup of the Soviet Union is the main reason for the dramatic decline in the share accounted for by Russia, evident in the early 1990s. Also evident are the higher frequency movements in the production shares of Saudi Arabia and Iran. Between 1975 and 1981, the Iranian share of world production falls precipitously. Some of the early decline is taken up by increases in Saudi Arabia, later some new entrants (those at the top of the chart) account for some of the change in composition.

3 The Model

The starting point of discussion is the seminal paper by Cole and Obstfeld (1991), who developed a sharp theoretical benchmark for commodity risk sharing. Specifically, in an endowment economy setting with complete specialization and Cobb-Douglas preferences, they demonstrated that in a two-country world, the movement of a nation's terms of trade was sufficient to make financial markets redundant.

The economic mechanism works as follows. The first country specializes in X and the second in Y. In a frictionless world, with Cobb-Douglas preferences, the relative price of the two goods is $P^* \equiv \frac{P_Y}{P_X} = \frac{X}{Y}\frac{(1-\theta)}{\theta}$, where X and Y are the realizations of the endowments of the two goods, and θ is the Cobb-Douglas weight on good X. Thus, any variation in relative supply is offset one-for-one with a relative price change. When countries are in financial autarky and trade only one physical object for another the allocations in a competitive equilibrium are the same as the allocations when countries pool risks to their individual endowments. Effectively, equilibrium terms of trade adjustments provide insurance against supply shocks.

Consider the more general case: a world populated by a large developed block, the OECD, and a smaller block of developing countries who are each specialized in the production of a single primary commodity. Note that while each country may be specialized, there may be more than one specialist in each good.

The maximization problem of country i, specializing in production of commodity j is:

$$\begin{aligned} \max U(C_{it}) &= \prod_{k}^{N} C_{ikt}^{\theta_k} \\ \text{subject to} &: \sum_k P_{kt} C_{ikt} \leq P_{jt} Y_{jt} \end{aligned} \tag{6}$$

$P_{jt}Y_{jt}$ is the value the single exported commodity fetches on world markets, but could, in general, be thought of as the market value of the entire export basket of country i with appropriate definitions of the aggregates. The demand for good k, by country i, is:

$$C_{ikt} = \theta_k \frac{P_{jt}}{P_{kt}} Y_{jt} \ . \tag{7}$$

Taking a geometric weighted average across commodities gives:

$$\begin{aligned} \prod_{k}^{N} C_{ikt}^{\theta_k} &= \prod_{k}^{N} \left(\theta_k \frac{P_{jt}}{P_{kt}} Y_{jt} \right)^{\theta_k} \\ P_t C_{it} &= \Theta P_{jt} Y_{jt} \end{aligned} \tag{8}$$

The last line indicates that an appropriately defined aggregate real consumption index, C_{it}, and consumer price index, P_t, when multiplied, equal nominal income accruing to country i from selling commodity j, up to a scalar normalization. Specifically,

$P_t = \prod_k^N P_{kt}^{\theta_k}$ and $C_{it} = \prod_k^N C_{ikt}^{\theta_k}$ and $\Theta \equiv \prod_k^N \theta_k^{\theta_k}$.

3.1 Model Implications

Since aggregate consumption, so defined, maps into welfare and wealth, what we would like to know is how consumption changes with relative prices and endowments. Taking the first-difference of the logarithm of the last equation and denoting logarithms of upper-case variables as lower-case variables, we have:

$$\Delta c_{it} = (\Delta \widetilde{p}_{jt} - \Delta p_t) + \Delta y_{jt} \ . \tag{9}$$

In words: changes in aggregate consumption of country i, consist of a terms of trade effect, $\Delta \widetilde{p}_{jt} - \Delta p_t$, and an endowment effect. The terms of trade effect is the change in the value of the endowment relative to the change in the value of value of the consumption basket, at the initial endowment level. The endowment effect is the change in physical quantity produced, which at fixed prices leads to a proportional change in aggregation consumption. Thus, the volatility of consumption depends on the variance of the terms of trade and production as well as the covariance between the two. The covariance term is the key element of the Cole and Obtsfeld asset market redundancy result.

To see this, consider a pair of countries each specializing in a particular good and noting how relative prices and quantities behave in equilibrium. Noting that price level inflation is: $\Delta p_t = \theta_1 \Delta \widetilde{p}_{1t} + \theta_2 \Delta \widetilde{p}_{2t}$, the consumption growth rate of country 1 and 2 (from (9)) are given by:

$$\Delta c_{1t} = \Delta\widetilde{p}_{1t} - [\theta_1\Delta\widetilde{p}_{1t} + \theta_2\Delta\widetilde{p}_{2t}] + \Delta y_{1t} \tag{10}$$

$$\Delta c_{2t} = \Delta\widetilde{p}_{2t} - [\theta_1\Delta\widetilde{p}_{1t} + \theta_2\Delta\widetilde{p}_{2t}] + \Delta y_{2t}$$

and the consumption growth differential is:

$$\Delta c_{1t} - \Delta c_{2t} = (\Delta\widetilde{p}_{1t} - \Delta\widetilde{p}_{2t}) + (\Delta y_{1t} - \Delta y_{2t}) = 0\ . \tag{11}$$

The last equality holds by virtue of the fact that relative prices and relative quantities move in proportion and in opposite directions: the 'market value' changes cancel. Thus consumption moves one-for-one across countries as implied by risk-sharing despite the fact that no risks are pooled ex ante: the terms of trade does all the work.

The subtlety in applying this result in practice is that it actually is referring to good-specific risk, not country-specific risk. In the model with complete specialization in production the two are the same, but in practice they are not. Recall Figure 6 and 7 with the national levels of coffee and petroleum production across producers of the world. To see what happens when countries remain specialized, but there are multiple producers of the two goods (less goods than countries),[7] let the output of country i of commodity j, y_{ijt}, consist of both a commodity-specific shock, y_{jt}, and an idiosyncratic shock specific to both the commodity and the country, ϵ_{ijt},

$$y_{ijt} = y_{jt} + \epsilon_{ijt} \tag{12}$$

where $E_i(y_{ijt}) = y_{jt}$ is the commodity specific component, while ϵ_{ijt} capture i.i.d. deviations from this level across countries that produce good j.

Aggregate consumption in country i, still has the terms of trade component, but now the quantity must be indexed by the good and country to reflect the nation-specific component of the endowment:

$$c_{it} = \widetilde{p}_{jt} - p_t + y_{ijt} \ . \tag{13}$$

Consider, now, two randomly selected countries, a and b, that produce goods, 1 and 2, respectively, subject to idiosyncratic as well as nation-specific risk, the previous result becomes (after trivial substitutions):

$$\Delta c_{at} - \Delta c_{bt} = (\Delta\widetilde{p}_{1t} - \Delta\widetilde{p}_{2t}) + \Delta y_{a1t} - \Delta y_{b2t} \ . \tag{14}$$

Substituting the country-specific realizations with their assumed stochastic processes gives us:

$$\Delta c_{at} - \Delta c_{bt} = (\Delta\widetilde{p}_{1t} - \Delta\widetilde{p}_{2t}) + (\Delta y_{1t} + \Delta\epsilon_{a1t} - \Delta y_{2t} - \Delta\epsilon_{b2t}). \tag{15}$$

and using the relationship between relative prices and quantities in the Cobb-Douglas formulation, gives:

$$\Delta c_{at} - \Delta c_{bt} = \Delta\epsilon_{a1t} - \Delta\epsilon_{b2t} \ . \tag{16}$$

Remarkably, each country bears all of its idiosyncratic risks!

What has happened to the risk-sharing facet of movements in the terms of trade? Essentially, the relative price movements of one commodity to another still provide insurance, but the group that is insured is the industry that produces that good. This is easily demonstrated by averaging the consumption allocations across individuals who specialize in different sectors:

$$E_i(\Delta c_{i1t}) - E_i(\Delta c_{i2t}) = E_i(\Delta\epsilon_{i1t} - \Delta\epsilon_{i2t}) = 0\ . \tag{17}$$

In words: equilibrium terms of trade adjustments insure the producers who happen to have a change in production exactly equal to the world sectoral change. Thus, in the absence of ex ante risk-sharing, all *idiosyncratic* output deviations are completely born by the individual countries. Note also, that if all sector risks are idiosyncratic to the nation, the terms of trade will be constant and all risks will be borne by the individual producers.

The extent to which national production levels move with the world aggregate and in response to the terms of trade is an empirical question, which our data allow us to investigate in detail in the next section.

4 Consumption risk-sharing and the terms of trade

This section applies the logic of the Cole-Obtsfeld model, extended to the case of non-specialization, to conduct a variance decomposition. The variance decomposition is that of national consumption growth relative to world consumption growth with the variance allocated to terms of trade variation, world production variation and nation-specific production variation.

The starting point is the construction of the commodity terms of trade from world commodity relative prices and the national commodity export data. Toward this end, each country's terms of trade is constructed as an export-share-weighted index of changes in primary commodity relative prices, $\Delta p_{it} = \Delta\widetilde{p}_{it} - \Delta p_t$:

$$\Delta p_{it} = \sum_j \omega_{ij} \Delta p_{jt} \tag{18}$$

Recall, p_{jt} is, as before, the U.S. dollar price of commodity j in world markets relative to the U.S. CPI. Thus when Δp_{jt} is positive, commodity j's inflation exceeds that of the U.S. consumption basket (which is treated as the import deflator), indicating a relative price increase and a terms of trade improvement for exporters of commodity j. The ω_{ij} are country and commodity-specific export weights computed from our panel:

$$\omega_{ij} \equiv T^{-1} \sum_t \frac{P_{jt} Y_{ijt}}{\sum_j P_{jt} Y_{ijt}} \ . \tag{19}$$

Note that the use of common, as opposed to country-specific commodity prices on the right-hand-side of the index construction assumes that the Law-of-One-Price holds internationally for these goods and for the basket of imports.

We return, now, to the logic of the CO model. Recall the generic equation for consumption growth for country i producing commodity j is $\Delta c_{it} = (\Delta\widetilde{p}_{jt} - \Delta p_t) + \Delta y_{jt}$. The case of non-specialization involves replacing the j subscripts with i subscripts indicating the terms of trade and production value of country i, given its pattern of specialization

$$\Delta c_{it} = (\Delta \widetilde{p}_{it} - \Delta p_t) + \Delta y_{it} \tag{20}$$

Aggregating across all countries gives the deviation of consumption growth of country i relative to world consumption growth, which represents the deviation from complete risk sharing:

$$(\Delta c_{it} - \Delta c_t) = (\Delta \widetilde{p}_{it} - \Delta p_t) + (\Delta y_{it} - \Delta y_t)\ . \tag{21}$$

The goal is to decompose these deviations from complete risk-sharing, $\Delta c_{it} - \Delta c_t$, into microeconomic sources of risk based on the available data. The variance of the change in consumption of country i relative to the world (the G-7 supplier of manufactured goods) is:

$$var(\Delta c_{it} - \Delta c_t) = var[(\Delta \widetilde{p}_{it} - \Delta p_t) + (\Delta y_{it} - \Delta y_t)] \tag{22}$$

This should be viewed as the amount of commodity revenue risk to be shared after taking into account non-diversifiable movements in world consumption or income and allowing for endogenous covariances among variables to offset some of the underlying risks. One such endogenous covariance is the negative covariance of the terms of trade with aggregate world output of a commodity highlighted in the CO model.

To ensure a full variance decomposition other sources of consumption variation are ignored and the focus is on decomposing the right-hand-side which is, given our data, nominal commodity export earnings relative to nominal world GDP. Drilling down to the commodity level on the right-hand-side, we have:

$$\begin{aligned}
\Delta x_{it} &\equiv \Delta c_{it} - \Delta c_t = (\Delta\widetilde{p}_{it} - \Delta p_t) + (\Delta y_{it} - \Delta y_t) \\
&= \underbrace{\sum_j \omega_{ij}(\Delta\widetilde{p}_{jt} - \Delta p_t)} + \underbrace{\sum_j \omega_{ij}(\Delta y_{jt} - \Delta y_t)} + \underbrace{\sum_j \omega_{ij}(\Delta y_{ijt} - \Delta y_{jt})} \qquad (23) \\
&= \quad \text{terms of trade} \quad + \quad \text{world production} \quad + \quad \text{nation-specific production}
\end{aligned}$$

and the resulting variance decomposition, in terms of the β-method is:

$$1 = \sum \omega_{ij} \left(\beta_{pj} + \beta_{yj} + \beta_{yij}\right)$$

- $\beta_{pj} = cov(\Delta\widetilde{p}_{jt} - \Delta p_t, \Delta x_{it})/var(\Delta x_{it})$ captures the pure terms of trade effect, the change in value of the exported commodity holding physical production fixed at steady-state shares (ω_{ij}).

- $\beta_{yj} = cov(\Delta y_{jt} - \Delta y_t, \Delta x_{it})/var(\Delta x_{it})$ captures the common world changes in the production of commodity j, relative to world GDP.

- $\beta_{yij} = cov(\Delta y_{ijt} - \Delta y_{jt}, \Delta x_{it})/var(\Delta x_{it})$ captures the contribution of nation-specific changes in the production of commodity j, i.e., relative to world commodity production of that commodity.

Table 3 reports the results of this variance decomposition focusing on the top exported commodity of each country in our panel. We focus on the top commodity since this commodity accounts for three-fourths of the variation in our primary export variable for the median country and virtual all of the variation for half of the countries.

Variation in the terms of trade is by far the dominant contributor to export earnings, accounting for about 82% of the variance for the median country. The terms of trade of the top exported commodity alone accounts for about two-thirds of this variation. The next most important contributor to export earnings variability is the nation-specific component of the quantity of primary commodities produced, at 15% for the median country. Again, the top export accounts for the bulk of this variation. World supply variation is trivial with a median absolute value contribution of about 3.5%.

Approximately one in three countries have a negative contribution of world production, which is consistent with the terms of trade channel of Cole and Obstfeld, but the magnitude of this variation is so modest that it contributes very little to the overall variance. Nation-specific production changes are overwhelmingly positive contributors to variability. This means countries tend to increase their production when their terms of trade improves, suggestive of an endogenous supply response and thus reinforcing the effect of terms of trade variation.

Figure 8 plots the decomposition of the predicted variance of $\Delta c_{it} - \Delta c_t$ into the contribution of the terms of trade, national production and world production for each country in the cross-section. While it is obvious that the terms of trade dominates the variance decomposition in most countries, there are a substantial minority of cases in which changes in domestic production relative to world production is the dominant factor. For twenty nations variation in the physical production of the top export commodity accounts for more than 20% of the total predicted variation.

Figure 9 plots the actual standard deviation of consumption growth against the predicted variance of consumption growth, excluding outliers.[8] A linear regression produces a slope coefficient of 0.14 and an R^2 of 16%. That the two are strongly positively correlated despite the fact that many sources of consumption variation are ignored is evidence of the importance of these commodity terms of trade changes and supply responses. There are many reasons that the coefficient is less than unity: most countries in the sample produce some manufactured goods as well as non-traded goods, for example.

5 Conclusion

The heterogeneity of business cycles across countries outside the large industrialized ones is vast. The hypothesis put forward here is that a substantial fraction of this heterogeneity reflects a concentration of exports in a single or small number of primary commodities. We estimate this fraction to be about 16%.

Specialization in primary commodities poses risks that are significant due to the highly volatile and unpredictable nature of world commodity prices. These risks may be further amplified when production is also somewhat beyond the control of individuals and firms, due to weather and national policy instability or when the short-term elasticity of supply is very low.

The implications of these trade patterns for the long-run was emphasized by Prebisch (1950) and Singer (1950) who documented a secular decline in relative commodity prices and later in the immiserizing growth theory of Bhagwati (1958). Business cycle implications of the terms of trade have been studied by Kose (2002), Kouparitsas (1997) and Mendoza (1995). These studies have focused on the aggregate terms of trade and also abstracted from idiosyncratic movements in production in the cross-section of countries specializing in the same commodity.

The focus here has been to emphasize the heterogeneity across countries and the role of individual primary commodities in shaping the heterogeneity through idiosyncratic production shifts across countries by commodity. We find evidence that national production responses almost always magnify, rather than dampen the impact of commodity price movements on a nation's export earnings. In other words, while there may be a risk-sharing effect in terms of negative comovements of commodity relative supplies and commodity relative prices at the world level, this effect is swamped by a positive covariance of national production and the relative price of the export commodity which amplifies export earnings variation. It could be that the elastic supply responses of developing countries are themselves enhanced by underlying risk-sharing arrangements in accord with the theoretical work of Baxter and Crucini (1995). Further study is needed to investigate this possibility.

Our analysis is subject to a number of important caveats. In particular, the import-side of the trade balance has been broad-brushed. We assumed all countries import a common consumption basket in exchange for a small number of exported primary commodities. It will be interesting to measure the import side of the trade balance more carefully and incorporate imports of primary commodity as inputs into domestic production. The analysis above has also abstracted from dynamic considerations, including permanent income theory, capital accumulation, credit market frictions and interest rate movements. Each of these facets are extensions in progress. Much remains to be done.

Victoria University of Wellington

Vanderbilt University and NBER

Pacific Lutheran University.

References

[1] Bhagwati, Jagedish. 1958. "Immiserizing Growth: A Geometrical Note", The Review of Economic Studies, Vol 25, No. 3, June., pp. 201-205.

[2] Baxter, Marianne and Mario J. Crucini. 1995. Business Cycles and the Asset Structure of Foreign Trade. International Economic Review, 36(4), 821-854.

[3] Bidarkota, P. and Crucini, M. J. 2000. Commodity Prices and the Terms of Trade. Review of International Economics, 8, 647–666.

[4] Canova, Fabio and Morten Ravn. 1996. International Consumption Risk Sharing. International Economic Review, 37(3), 573-601.

[5] Cochrane, John. 1991. A Simple Test of Full Consumption Insurance. Journal of Political Economy 99, 957–976.

[6] Cole, Harold and Maurice Obtsfeld. 1991. Commodity trade and international risk sharing How much do financial markets matter? Journal of Monetary Economics 28, 3-24.

[7] Hess, Gregory and Eric van Wincoop (eds.). 2000. Intranational Economics, Cambridge University Press.

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

[9] Kouparitsas, Michael A. 1997. "North-South terms of trade: an empirical investigation," Working Paper Series, Macroeconomic Issues WP-97-05, Federal Reserve Bank of Chicago.

[10] Kouparitsas, Michael A. 1997. "A dynamic macroeconomic analysis of NAFTA," Economic Perspectives, Federal Reserve Bank of Chicago, January, 14-35.

[11] Lewis, Karen K. . 1996. What Can Explain the Apparent Lack of International Consumption Risk Sharing? The Journal ofPolitical Economy, 104(2), 267-297.

[12] Mace, Barbara. 1991. Full insurance in the presence of aggregate uncertainty. Journal of Political Economy, 99(5), 928-956.

[13] Mendoza, Enrique. 1995. The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations. International Economic Review, 36(1), 101-137.

[14] Prebisch, Raoul. The Economic Development of Latin America and Its Principal Problems (New York: United Nations, 1950).

[15] Singer, H. W. "The Distribution of Gains between Investing and Borrowing Countries," American Economic Review 40 (1950): 473–85.

[16] Wang, Chih-Wei. 2010. Trend-cycle Decompositions of Commodity Prices. Mimeo, Pacific Lutheran University.

Notes

[1]See Hess and van Wincoop (2000) for insightful comparisons of approaches at the regional and national level.

[2]Using the Hodrick-Prescott filter gives qualitatively similar results.

[3]World GDP and consumption growth are computed as a country-size weighted averages of national GDP and consumption growth, respectively. The country size weights used are the averages over the sample period.

[4]See the data appendix for more details about data sources and index construction.

[5]The reader will recognize that the international business cycle liteature has shown how difficult it is to reduce the consumption correlation below unity under risk pooling even when assumptions such as separable utility is relaxed and preference shocks are introduced.

[6]See Chih-Wei Wang (2011) for evidence on persistence of commodity prices relative to the CPI and the role of permanent and transitory shocks in the evolution of nominal commodity prices.

[7]It may seem odd to have the number of countries exceed the number of goods rather than the reverse, but this is more a consequence of the arbitrariness of using countries as the unit of account than the thrust of our argument. The important point is that we break the singularity of country and sector shocks.

[8]The figure and regression line do not include extreme outliers of the distribution, countries with terms of trade standard deviations in excess of 40% (Portugal, El Salvador, Uganda, Guatemala, Costa Rica, Kuwait, Ethiopia, Ecuador, Colombia, Morocco, Senegal) and countries with consumption growth standard deviations exceeding 10% (Saudi Arabia, Nigeria).

Table 1. Variance decomposition of commodity price index (levels)

	Equally weighted		Production share weighted	
	weight	var share	weight	var share
Aluminum	3.3	2.4	1.1	0.9
Apples	3.3	1.4	0.9	0.5
Butter	3.3	4.1	1.7	2.5
Cocoa	3.3	3.8	0.4	0.5
Coffee	3.3	2.0	0.3	0.2
Copper	3.3	3.5	1.4	1.8
Corn	3.3	4.2	2.1	3.1
Corn Oil	3.3	4.5	4.3	6.9
Cotton	3.3	4.0	2.1	3.0
Eggs	3.3	3.6	3.0	3.7
Hides	3.3	1.1	0.9	0.4
Iron	3.3	1.9	6.8	4.7
Lead	3.3	2.5	0.4	0.4
Milk	3.3	2.5	10.7	10.0
Nickel	---	---	---	---
Oranges	3.3	9.6	0.1	0.2
Palm oil	---	---	---	---
Peanuts	3.3	3.4	0.6	0.7
Pepper	3.3	1.1	1.2	0.3
Petroleum	3.3	1.3	36.4	24.9
Potatoes[1]	3.3	4.3	--	--
Rice	3.3	4.2	8.5	12.6
Rubber	3.3	5.0	0.5	1.0
Rye	---	---	---	---
Soy Meal	3.3	2.6	0.3	0.3
Soy Oil	3.3	3.9	0.8	1.1
Soybeans	3.3	3.5	2.3	2.9
Sugar	3.3	2.8	4.6	4.8
Tallow	3.3	2.9	0.2	0.2
Tin	3.3	3.4	0.2	0.3
Wheat	3.3	4.4	6.3	10.0
Wool	3.3	4.1	1.3	1.8
Zinc	3.3	1.9	0.7	0.5
Totals	100.0	100.0	100.0	100.0

Table 2: Variance decomposition of commodity quantity index

	Equally weighted		Production share weighted	
Commodity	weight	var share	weight	var share
Aluminum	3.4	1.8	1.0	0.5
Apples	3.4	4.2	0.9	-0.7
Butter	3.4	3.0	1.8	3.6
Cocoa	3.4	10.7	0.4	0.4
Coffee	3.4	11.3	1.5	-0.3
Copper	3.4	2.8	1.9	1.6
Corn	3.4	10.5	4.4	12.3
Corn Oil				
Cotton	3.4	8.5	2.2	4.5
Eggs	3.4	0.4	3.1	-2.1
Hides	3.4	0.3	1.0	0.1
Iron	3.4	-0.6	7.1	4.4
Lead	3.4	1.2	0.4	-0.1
Milk	3.4	4.5	11.1	19.8
Nickel	3.4	3.5	0.6	0.4
Oranges	3.4	-3.7	0.1	-0.1
Palm oil	3.4	2.7	0.5	-1.0
Peanuts	3.4	9.2	1.3	0.6
Pepper	--	--	--	--
Petroleum	3.4	-1.2	34.2	34.2
Potatoes[1]	--	--	--	--
Rice	3.4	2.2	8.9	3.8
Rubber	3.4	0.6	0.5	0.1
Rye	3.4	10.8	0.3	0.8
Soy Meal				
Soy Oil	3.4	2.8	0.7	-0.7
Soybeans	3.4	5.7	2.3	2.3
Sugar	3.4	2.3	4.8	4.8
Tallow	3.4	-2.0	0.2	-0.1
Tin	3.4	5.1	0.2	0.3
Wheat	3.4	2.8	6.5	9.3
Wool	3.4	1.5	1.4	0.9
Zinc	3.4	-0.7	0.7	0.2
Totals	100.0	100.0	100.0	100.0

Table 3. Variance Decomposition of Export Earnings

Country	Std. Dev. of Primary Exports	First-ranked commodity	Terms of trade	Nation-specific production	World production	Total contribution to variance	Other commodities
Argentina	19.9	corn	12.8	30.7	-2.4	41.1	58.9
Australia	18.6	wheat	16.8	31.7	0.3	48.8	51.2
Bangladesh	34.8	rice	95.4	5.9	-1.3	100.0	--
Belgium-Luxembourg	21.3	zinc	51.5	17.2	4.8	73.5	26.5
Bolivia	29.8	tin	58.9	26.7	11.0	96.7	3.3
Brazil	24.5	coffee	37.8	16.1	5.4	59.3	40.7
Burkina Faso	23.7	cotton	42.1	61.0	-3.1	100.0	--
Cameroon	35.0	cocoa	85.4	12.7	1.9	100.0	--
Canada	13.6	petroleum	22.4	1.3	2.3	26.0	74.0
Chile	36.2	copper	95.9	3.2	0.9	100.0	--
China	17.2	iron	49.6	3.1	5.2	58.0	42.0
Colombia	49.4	coffee	94.3	-2.6	8.0	99.7	0.3
Costa Rica	52.1	coffee	88.9	7.9	3.2	100.0	--
Côte dIvoire	36.0	cocoa	42.8	8.0	5.5	56.3	43.7
Cuba	33.4	sugar	56.6	42.3	1.2	100.0	--
Dominican Republic	34.5	cocoa	62.7	38.9	-1.6	100.0	--
Ecuador	50.9	cocoa	42.7	51.8	5.5	100.0	--
Egypt	25.4	cotton	51.7	42.9	5.2	99.9	0.1
El Salvador	65.1	coffee	71.3	23.4	5.3	100.0	--
Ethiopia	51.5	coffee	93.2	0.4	6.4	100.0	--
EU-27	18.6	sugar	50.3	5.0	-0.4	54.9	45.1
France, Monaco	24.3	iron	81.9	1.4	5.9	89.2	10.8
Germany	24.4	iron	76.0	2.1	5.6	83.8	16.2
Ghana	32.3	cocoa	79.1	15.9	5.0	100.0	--
Greece	23.8	cotton	39.7	24.0	7.4	71.1	28.9
Guatemala	52.5	coffee	92.9	0.7	6.4	100.0	--
India	20.8	cotton	28.7	-0.7	7.7	35.7	64.3
Indonesia	21.9	petroleum	56.8	2.3	1.2	60.3	39.7
Iran	30.4	petroleum	73.5	26.9	-0.4	100.0	--
Italy	24.5	iron	96.2	-3.5	5.0	97.7	2.3
Japan	28.2	iron	91.1	3.1	5.3	99.5	0.5
Kazakhstan	27.3	zinc	36.9	60.6	2.5	100.0	--
Korea, Republic	32.0	iron	64.7	32.7	2.5	99.9	0.1

Table 3. Variance Decomposition of Export Earnings (continued)

Country	Std. Dev. of Primary Exports	First-ranked commodity	Terms of trade	Nation-specific production	World production	Total contribution to variance	Other commodities
Kuwait	51.7	petroleum	23.2	76.8	0.1	100.0	--
Malaysia	24.5	rubber	67.6	13.6	4.0	85.2	14.8
Mexico	22.9	petroleum	55.2	12.2	-0.2	67.2	32.8
Morocco	47.9	oranges	90.1	6.4	3.5	100.0	--
Myanmar (Burma)	34.3	rice	96.9	4.7	-1.6	100.0	--
Netherlands	23.6	tallow	89.6	11.4	-1.0	100.0	--
New Caledonia	25.6	iron	99.7	0.0	0.0	99.7	0.3
New Zealand	18.6	milk	51.9	7.2	-3.8	55.3	44.7
Nigeria	25.5	cocoa	35.6	19.0	4.4	58.9	41.1
Norway	35.5	aluminum	92.7	2.6	4.7	100.0	--
Pakistan	27.0	rice	59.6	7.2	-1.5	65.2	34.8
Paraguay	31.7	soybeans	46.0	66.1	-12.1	100.0	--
Peru	37.1	copper	85.7	13.2	0.3	99.2	0.8
Philippines	33.4	rice	99.3	2.4	-1.8	100.0	--
Poland	36.7	copper	92.3	7.9	-0.1	100.0	--
Portugal	79.7	tin	-0.8	101.2	-0.3	100.0	--
Romania	31.4	wool	92.5	7.6	-0.1	100.0	--
Russian Federation	19.9	petroleum	75.5	10.5	0.7	86.7	13.3
Saudi Arabia	30.9	petroleum	76.8	22.1	1.0	100.0	--
Senegal	44.6	peanuts	10.0	82.7	7.4	100.0	--
South Africa	28.0	oranges	83.2	1.9	2.3	87.4	12.6
Spain	27.8	oranges	91.4	-10.5	4.8	85.7	14.3
Sri Lanka	27.5	rubber	85.5	8.8	5.7	100.0	--
Switzerland	19.0	iron	41.7	0.0	0.0	41.7	58.3
Thailand	20.7	rice	45.9	4.4	-1.0	49.3	50.7
Turkey	22.7	cotton	65.9	15.6	9.4	90.9	9.1
Uganda	53.3	coffee	86.8	6.8	6.4	100.0	--
Ukraine	23.8	sugar	75.1	23.2	-2.3	96.0	4.0
United Kingdom	30.3	iron	60.2	24.7	3.7	88.7	11.3
Uruguay	28.9	wool	92.9	6.1	1.0	100.0	--
USA and Puerto Rico	16.0	iron	18.3	5.4	2.8	26.5	73.5
Venezuela	26.5	petroleum	86.3	5.6	-0.1	91.8	8.2
Vietnam	30.2	coffee	39.5	12.3	0.4	52.1	47.9

Figure 1. International business cycles
standard deviation of output and consumption growth

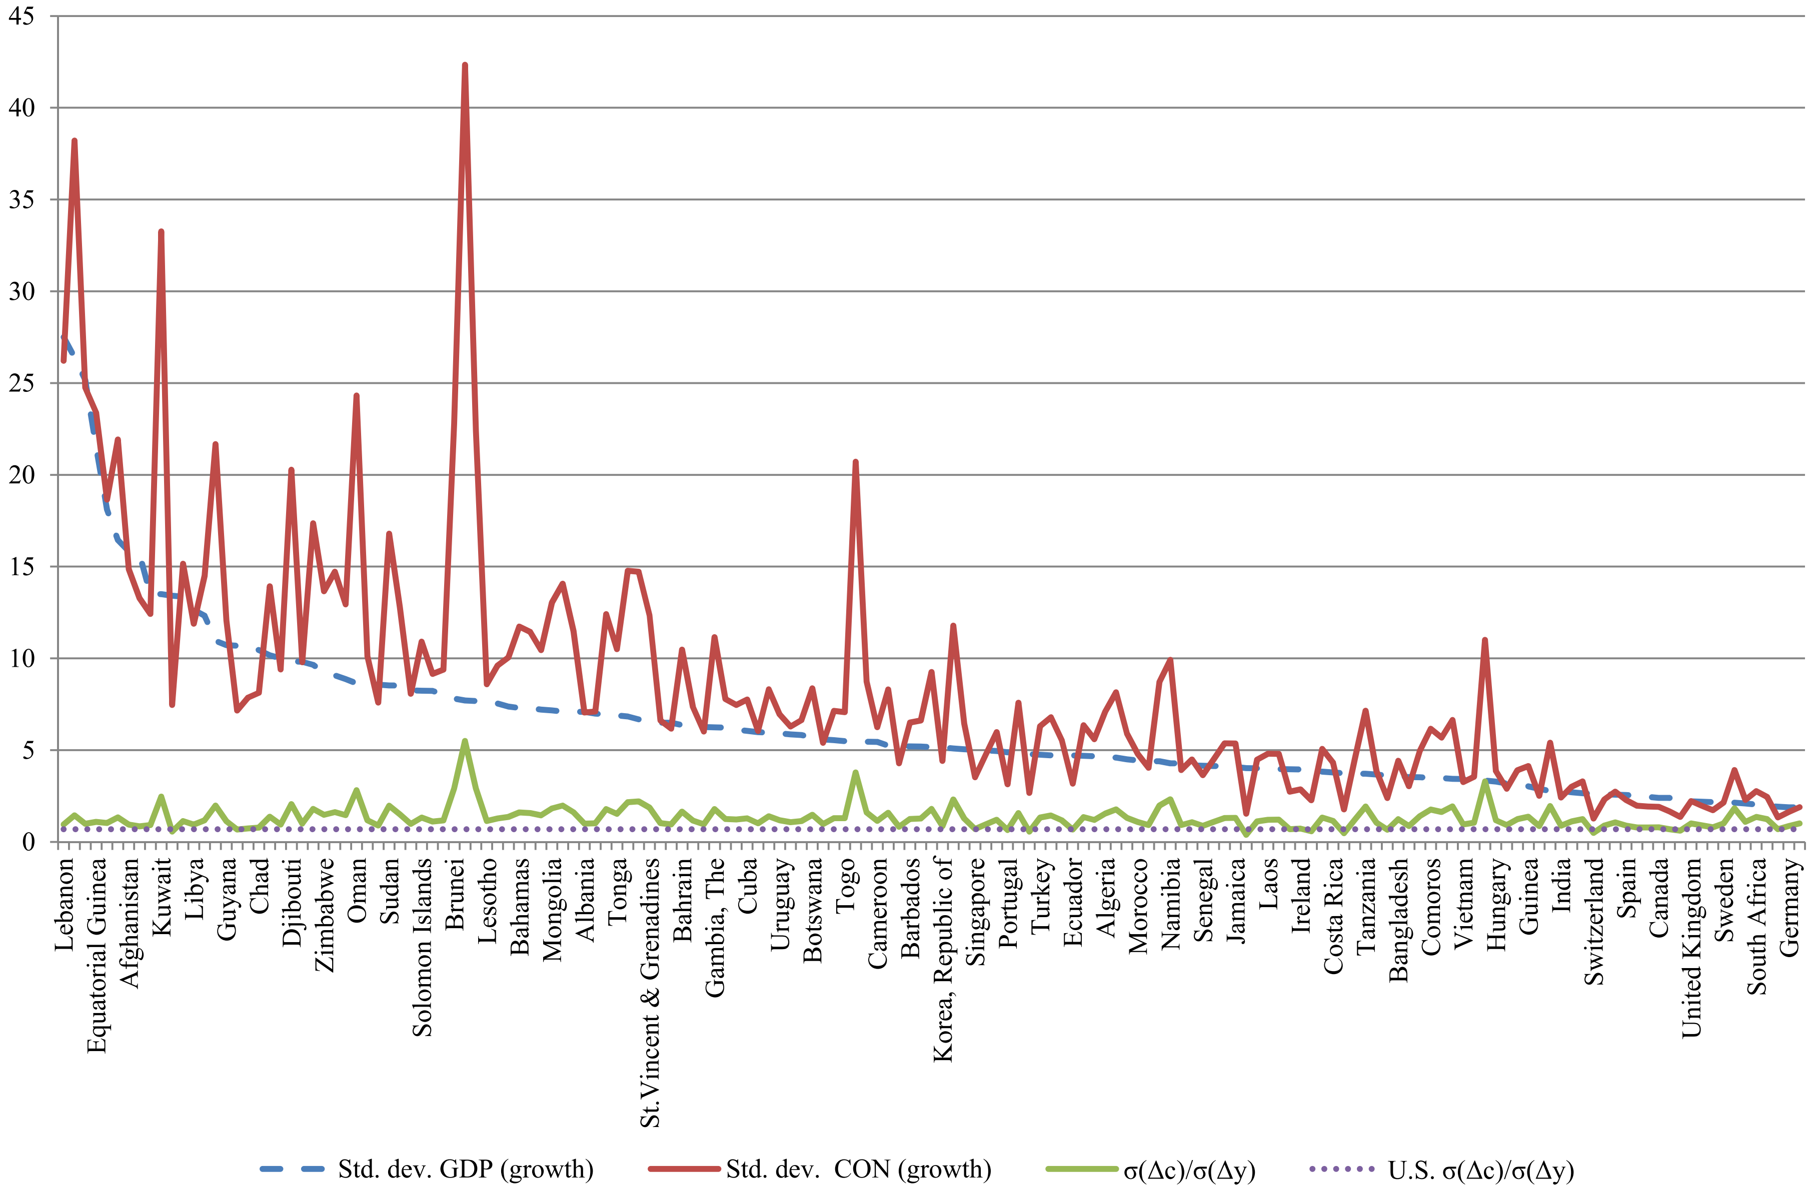

Figure 2. The evolution of relative country size

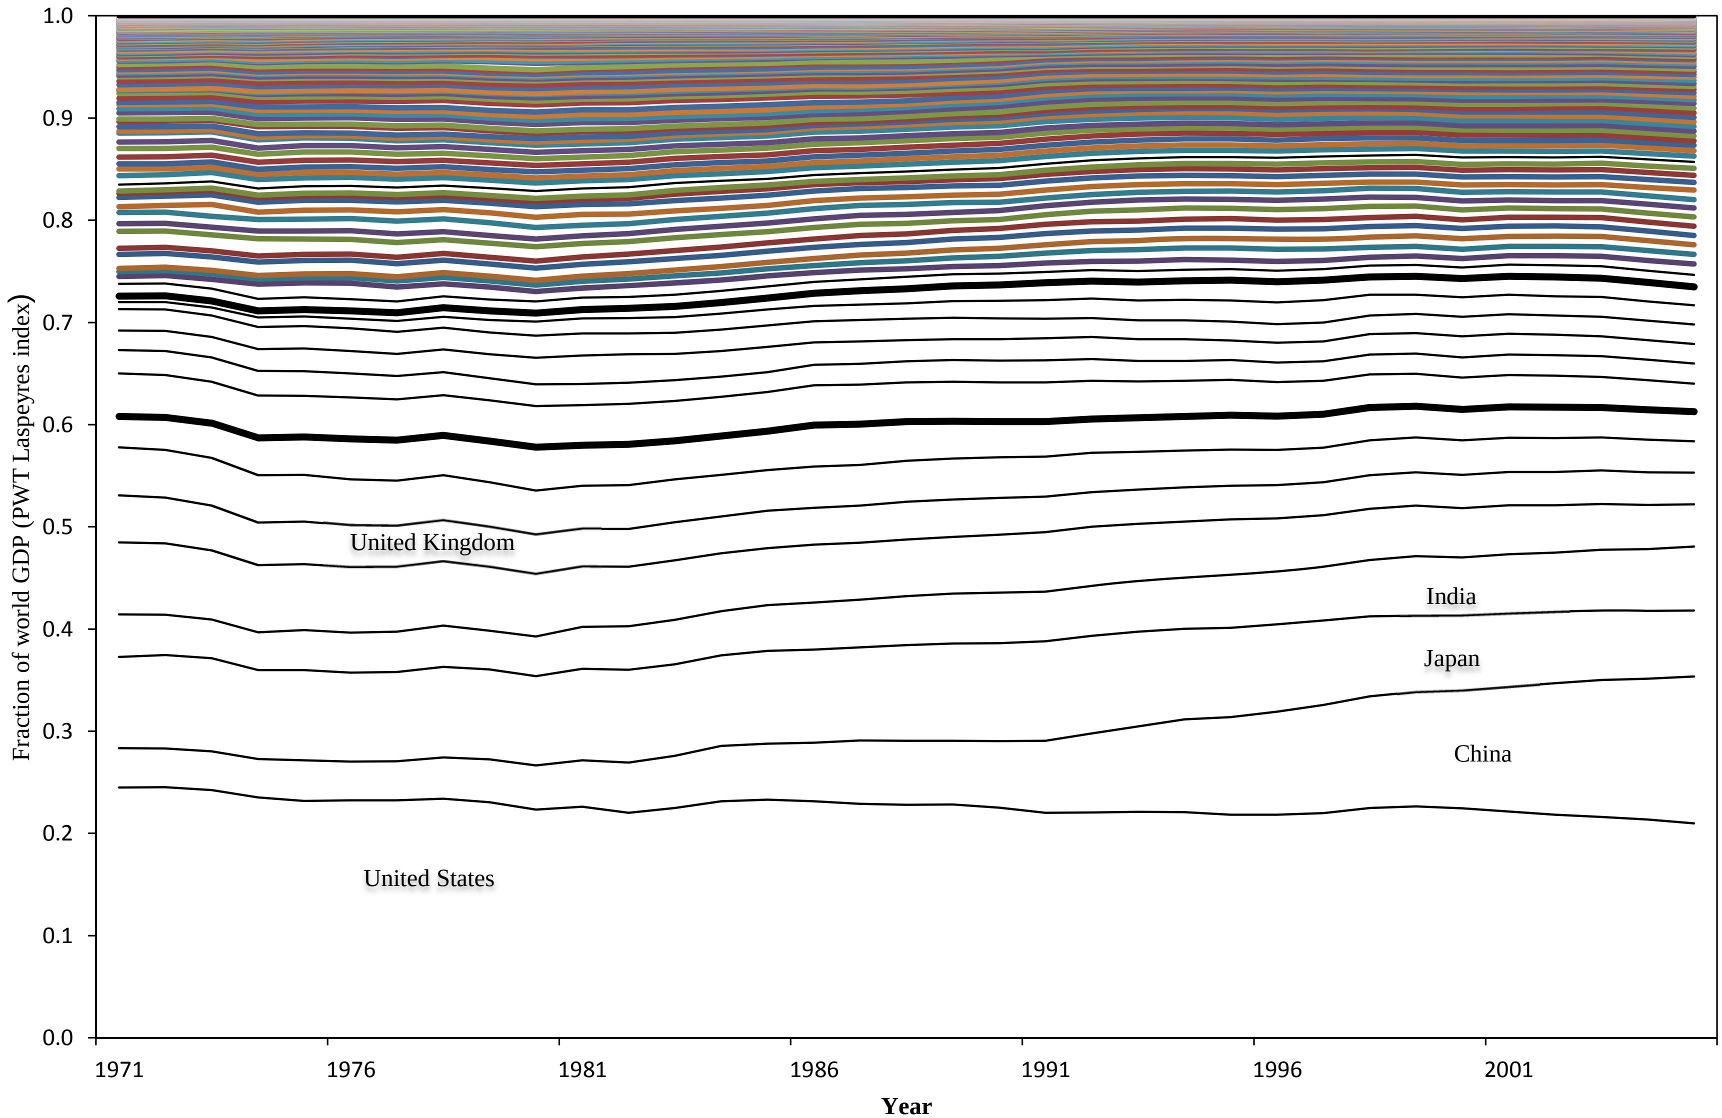

Figure 3. Risk-sharing regression

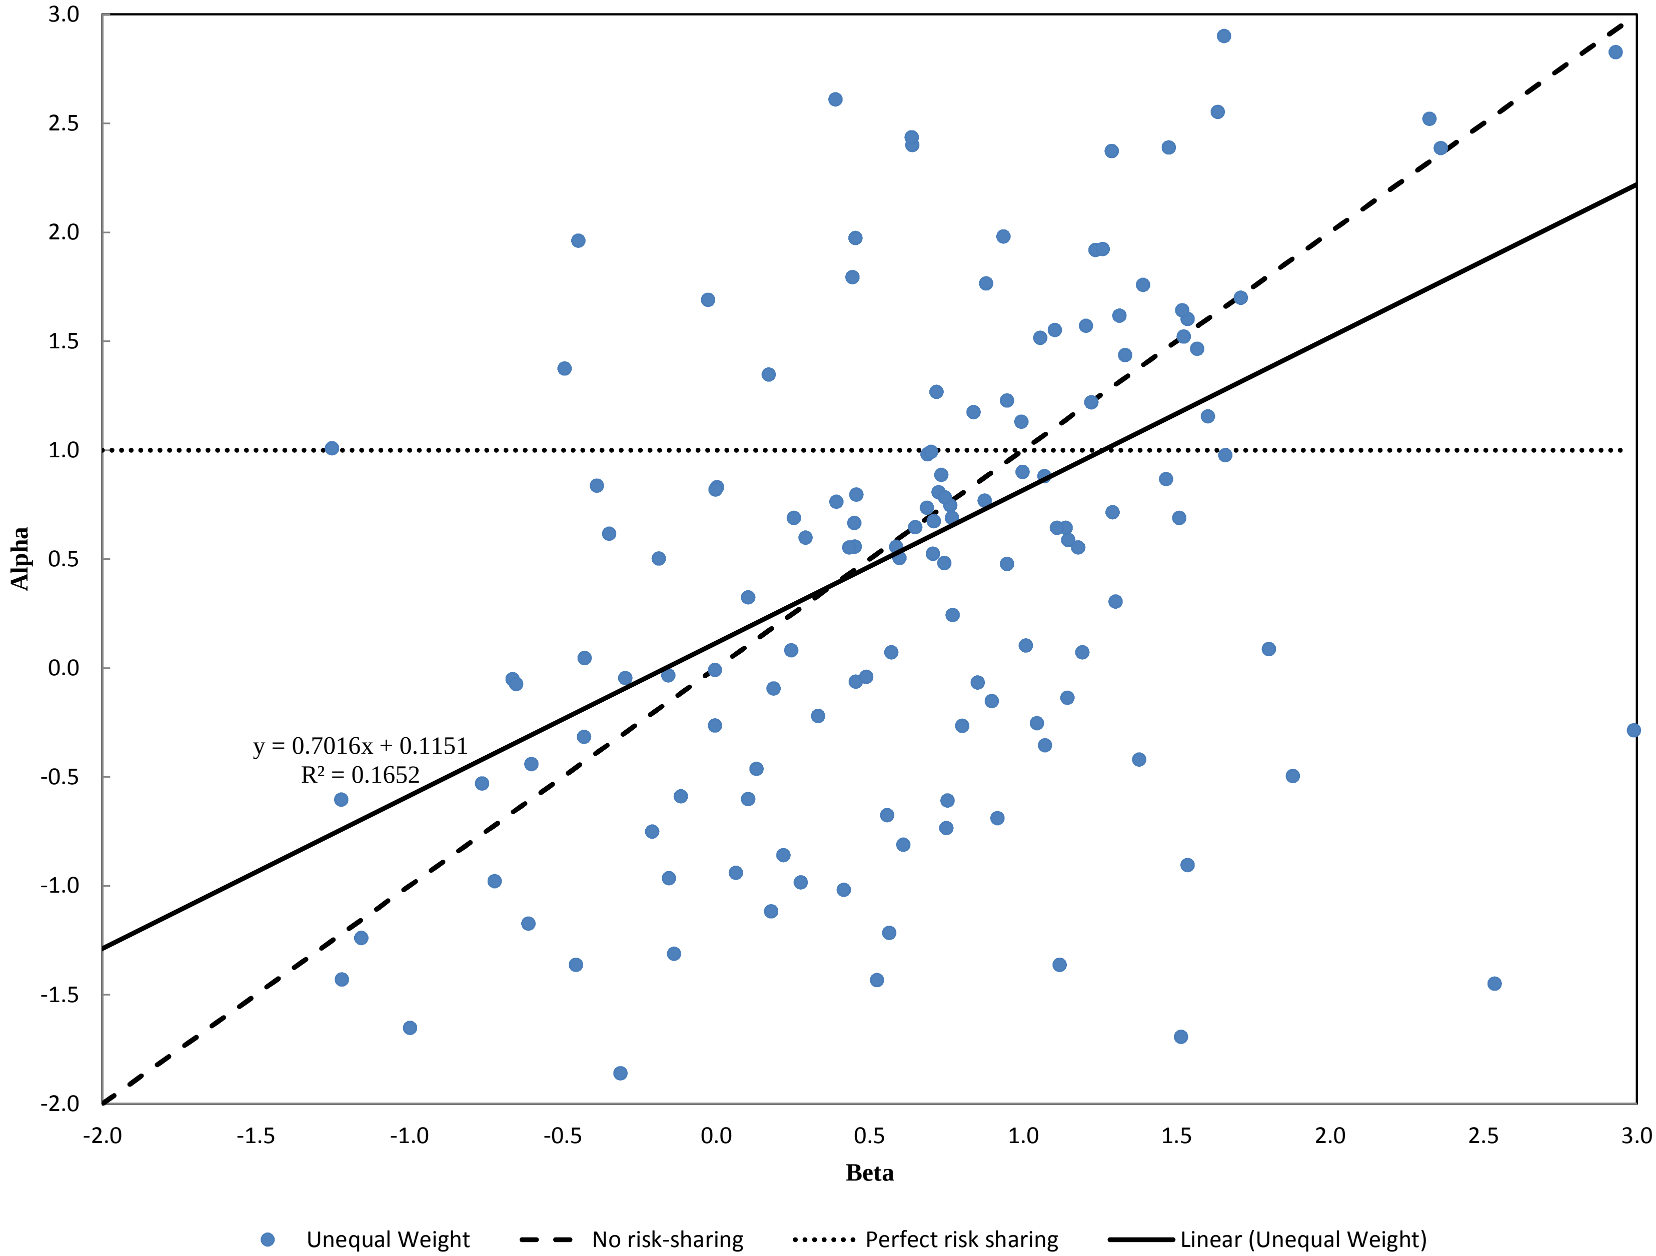

Figure 4. Commmodity Quantity and Price Trends

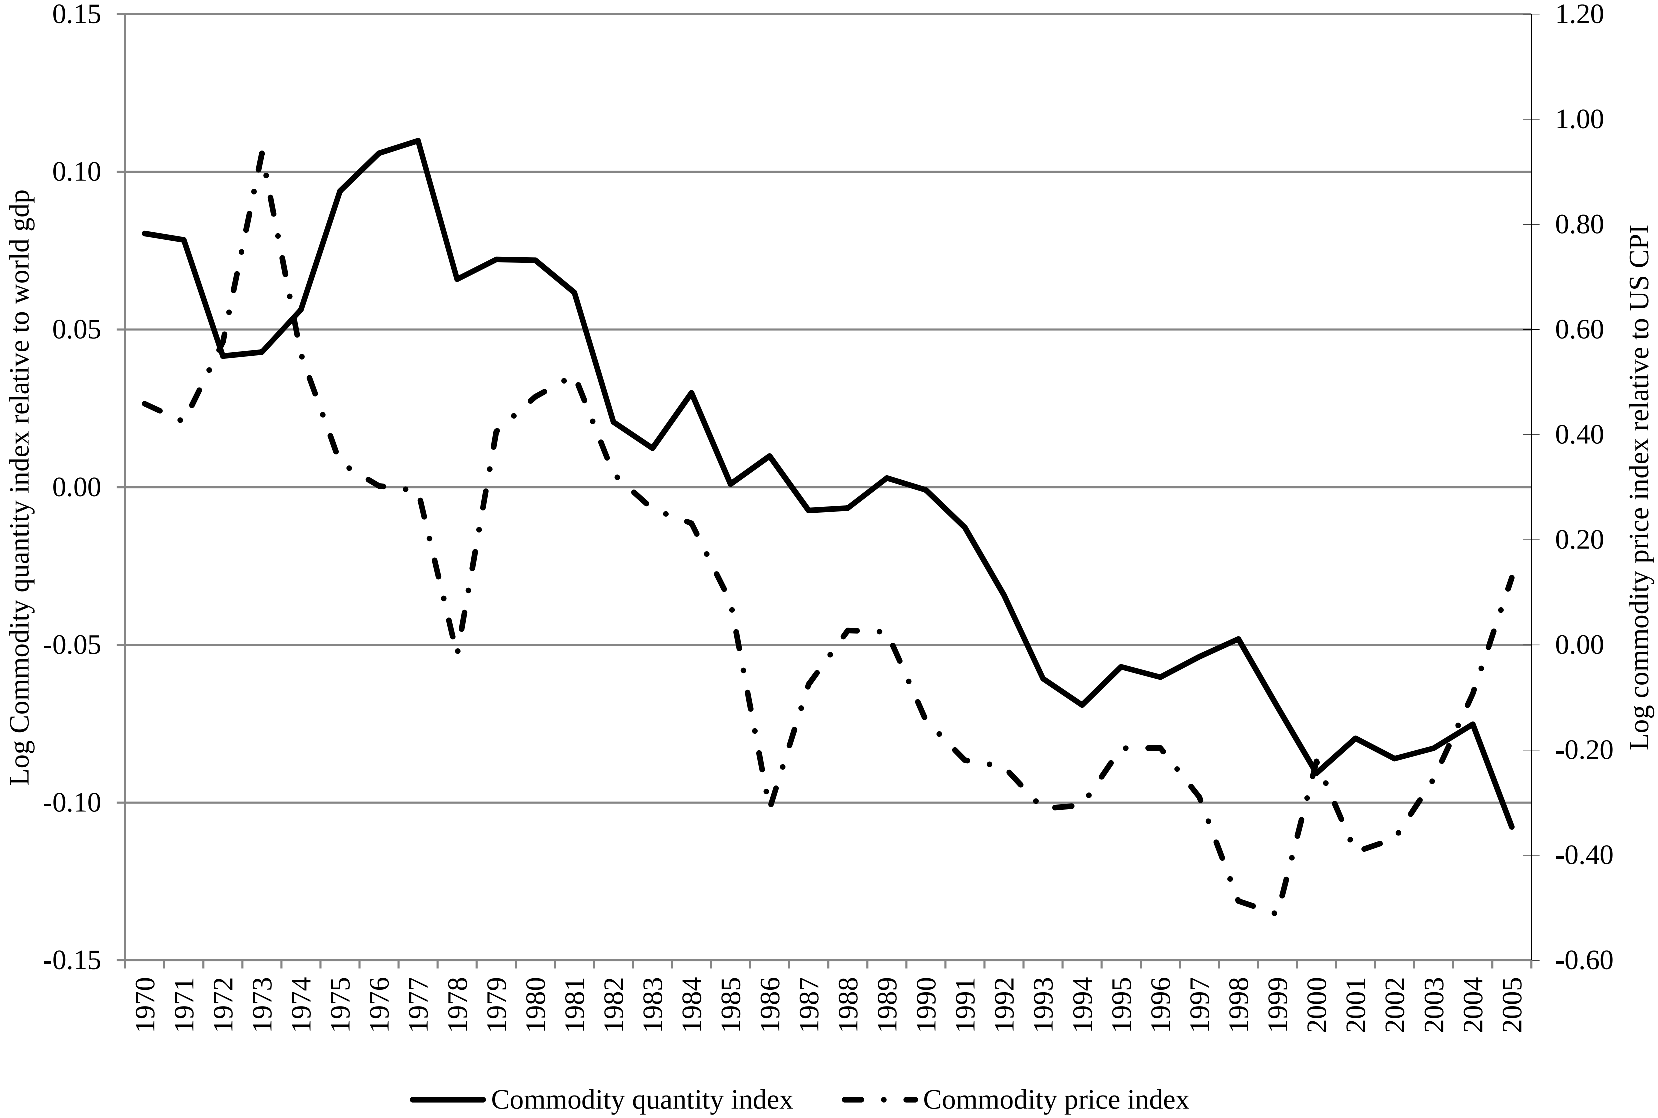

Figure 5. Commodity price and production changes

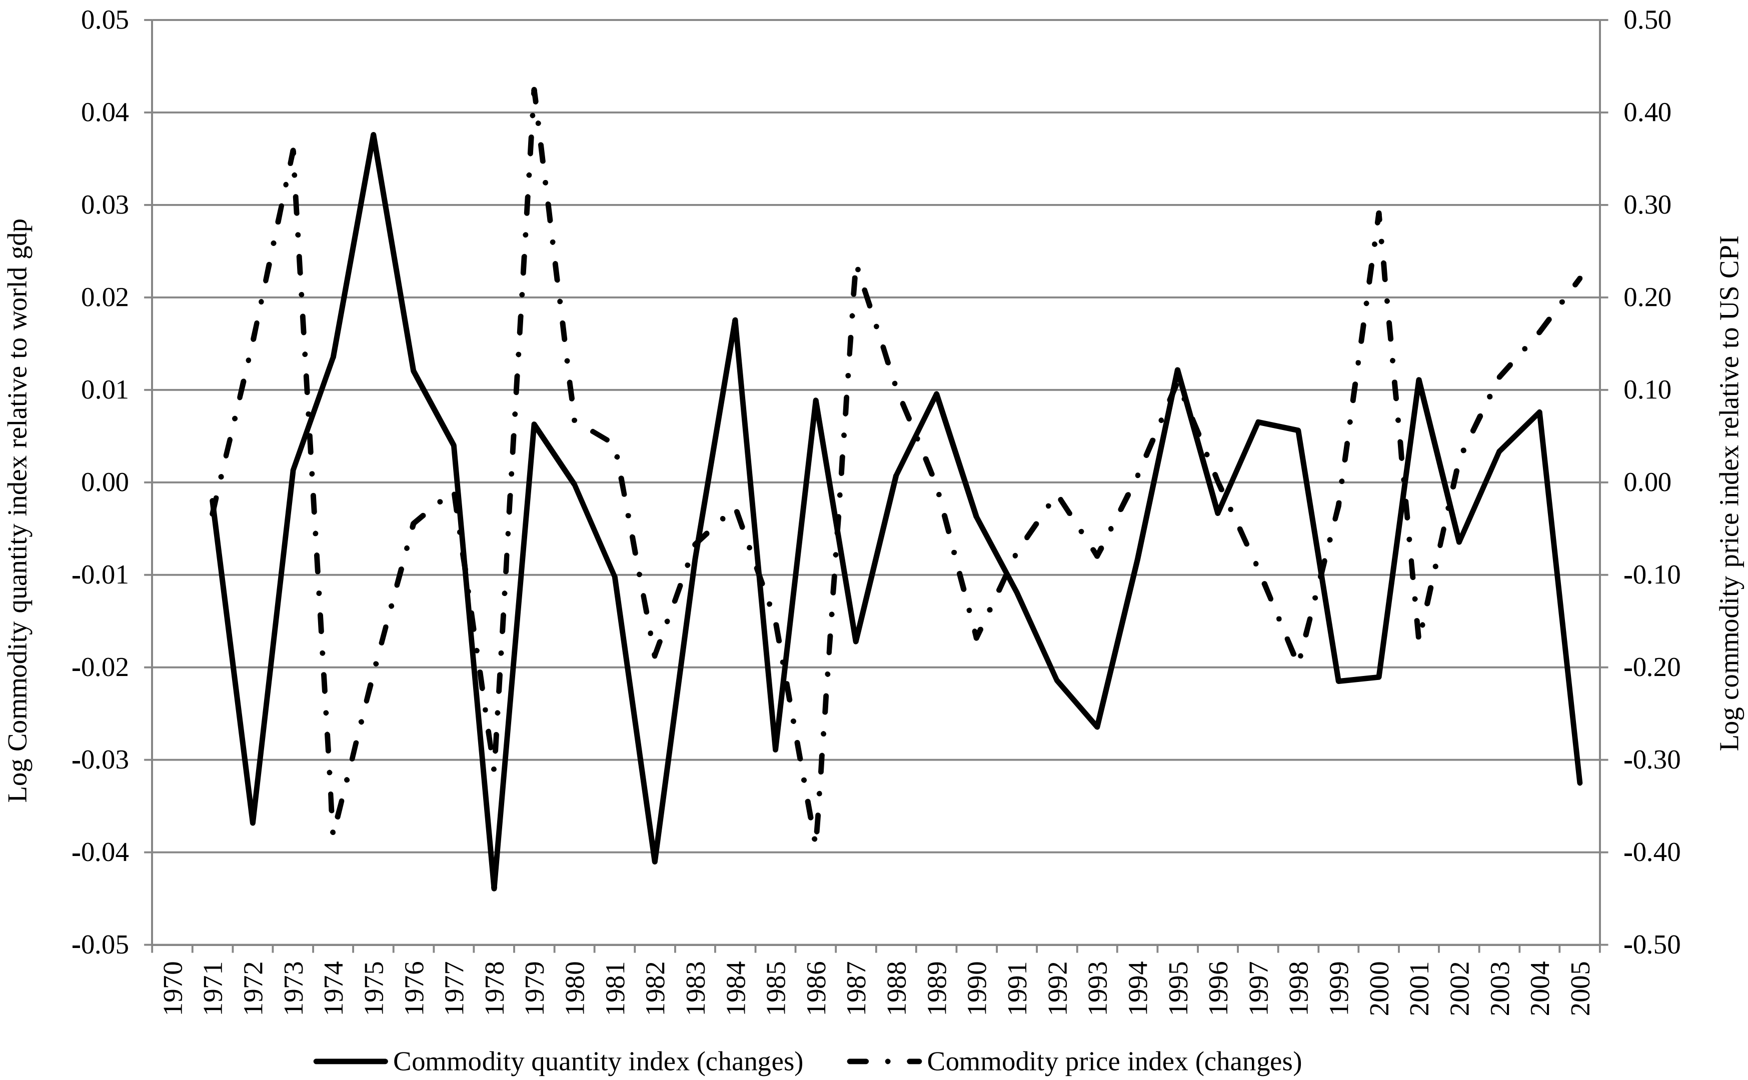

Figure 6. Coffee world production shares, by country

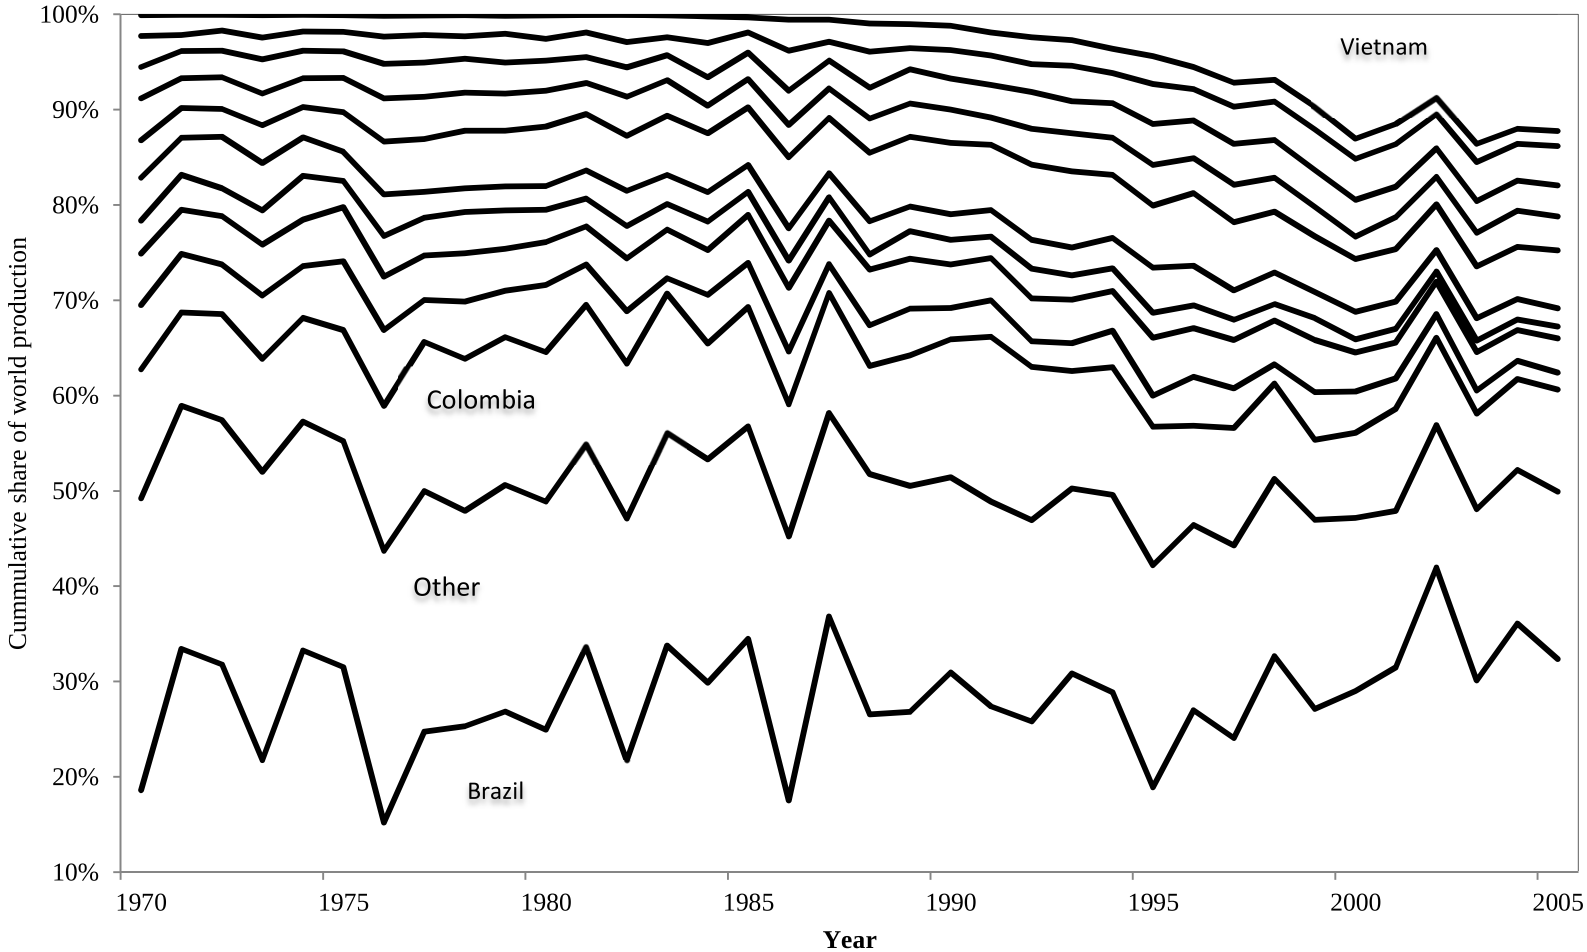

Figure 7. Petroleum world production shares, by country

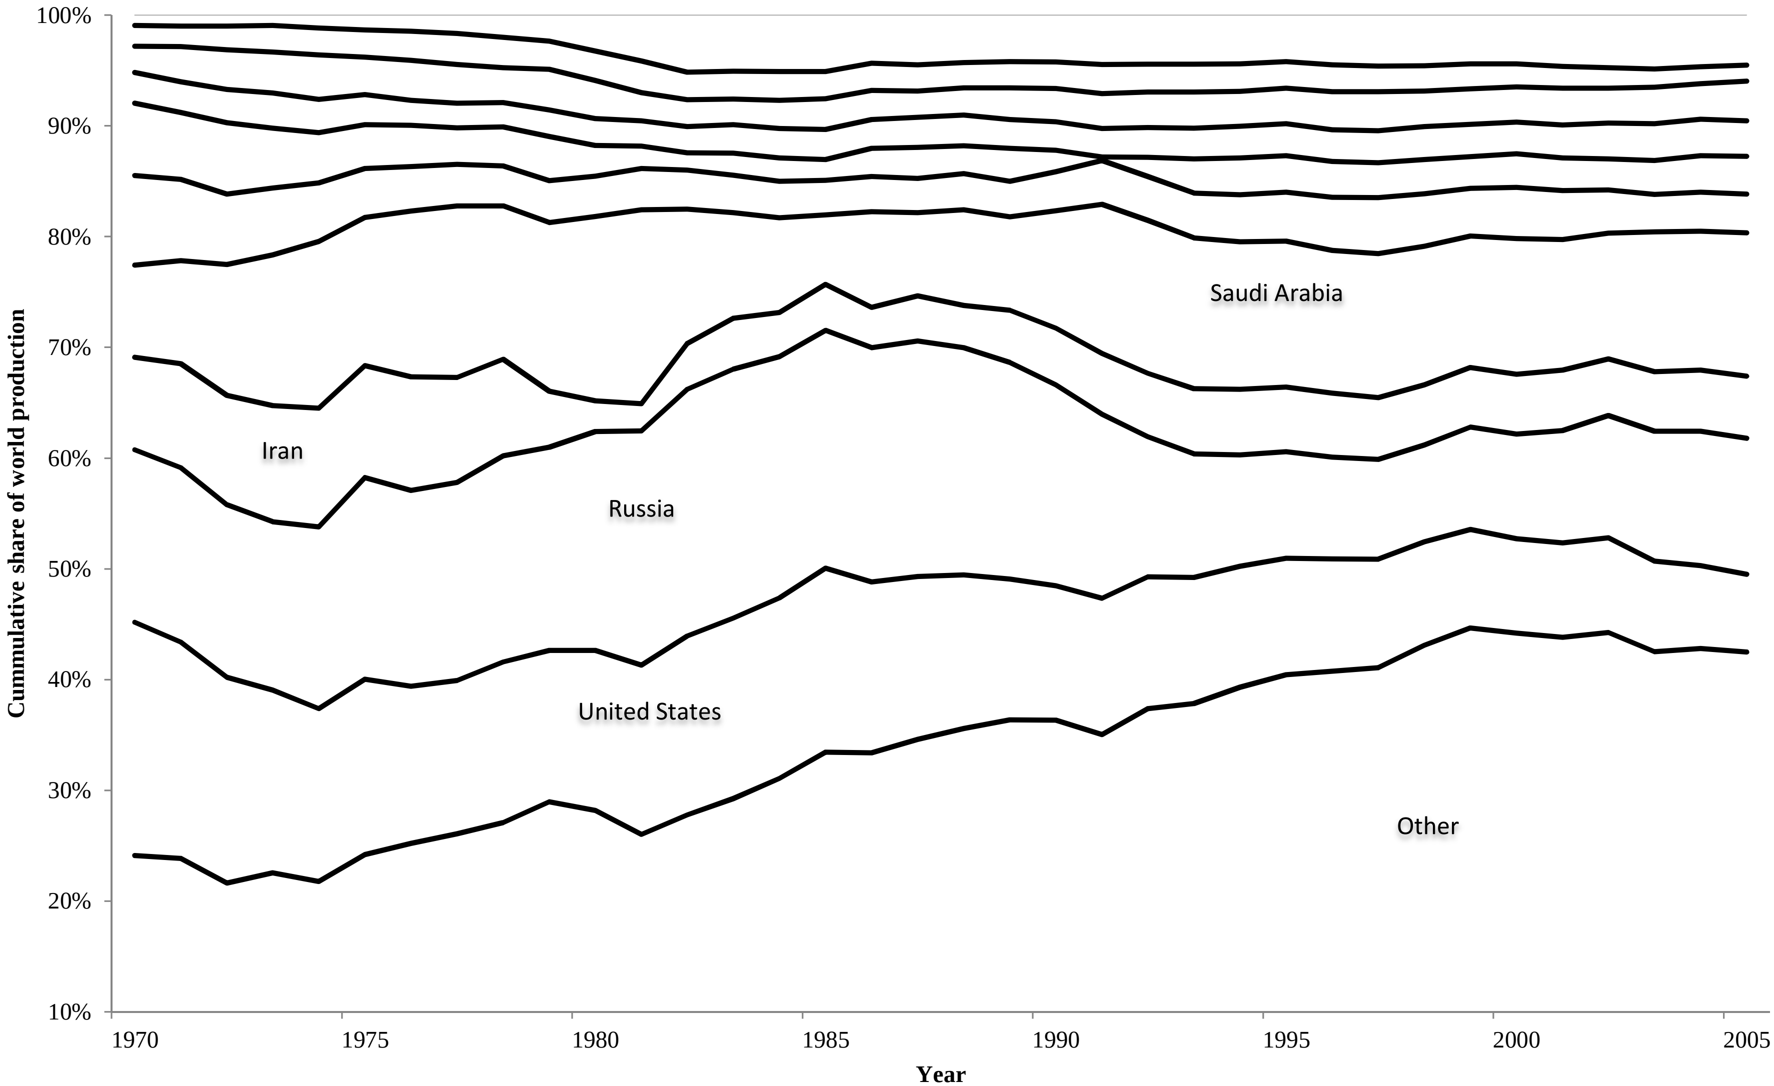

Figure 8. Sources of deviations from risk-sharing

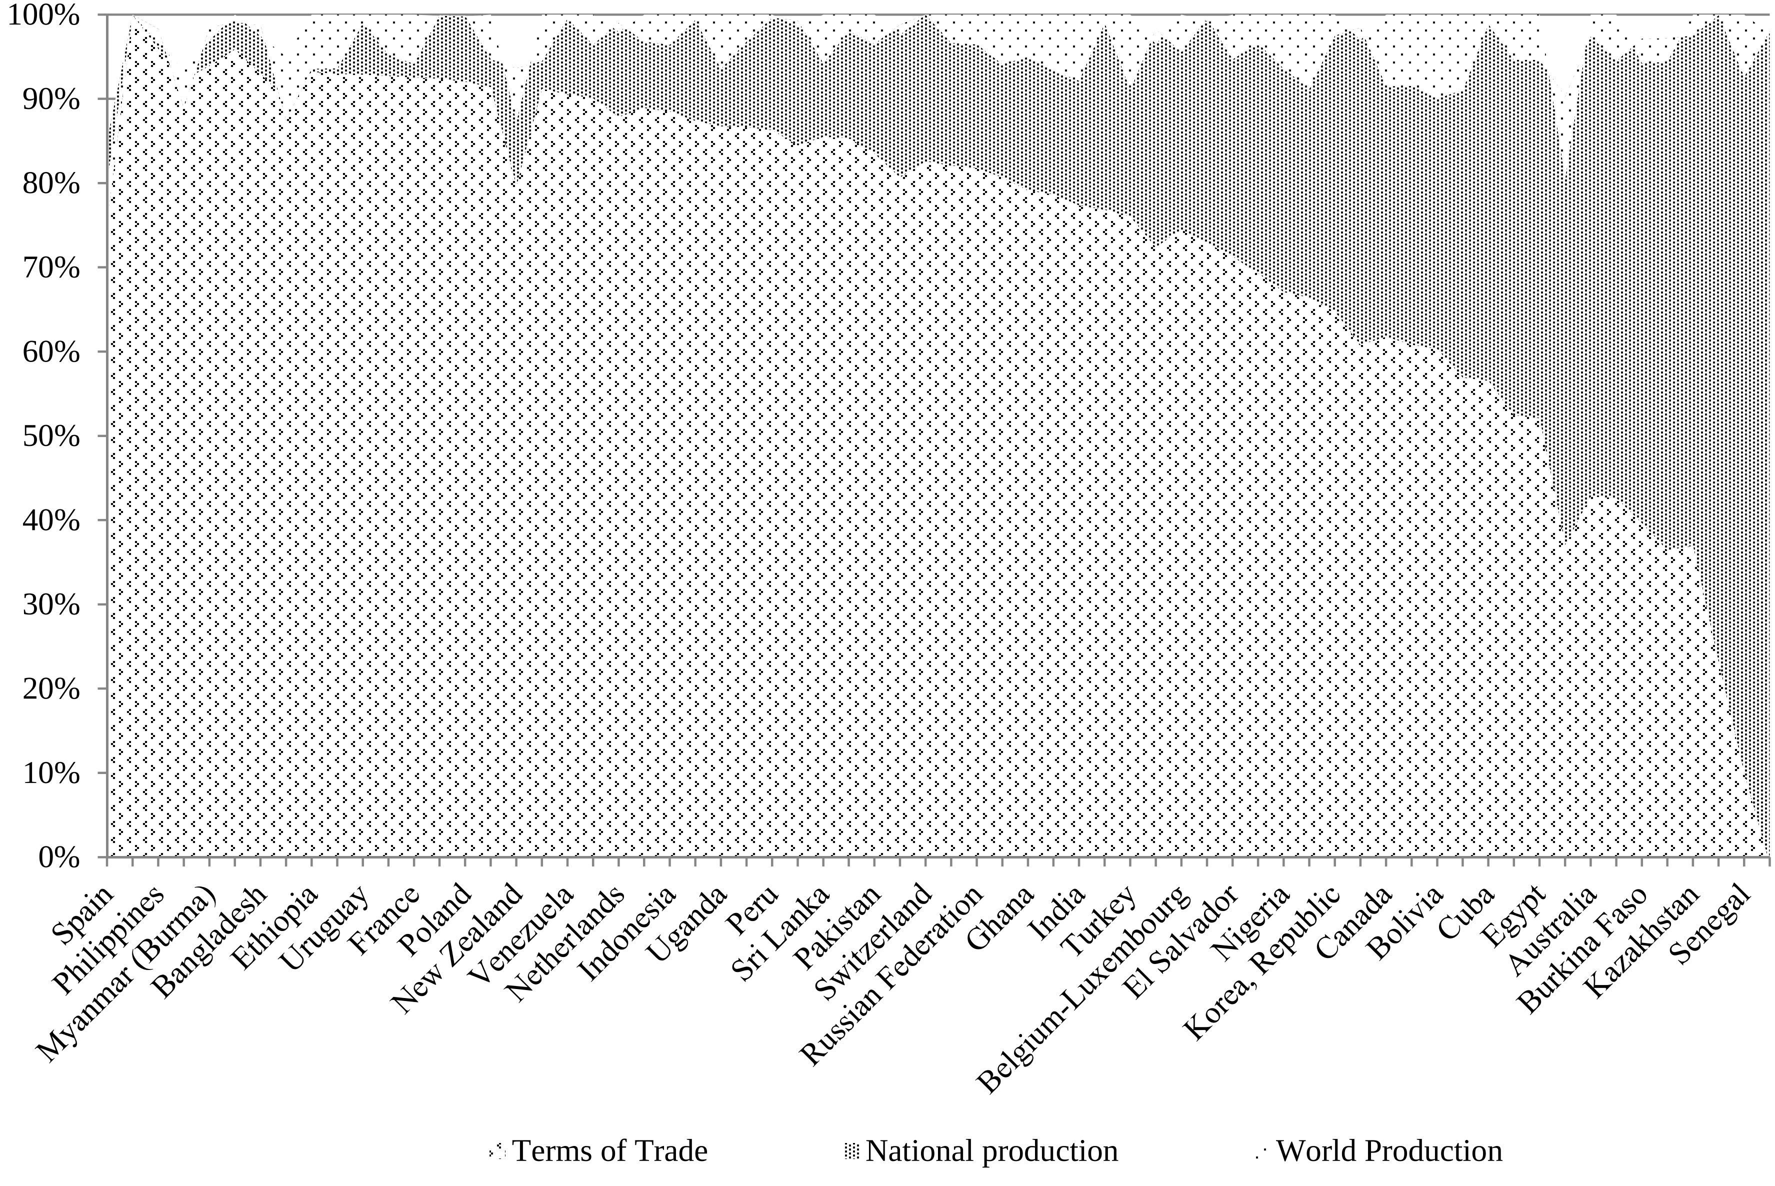

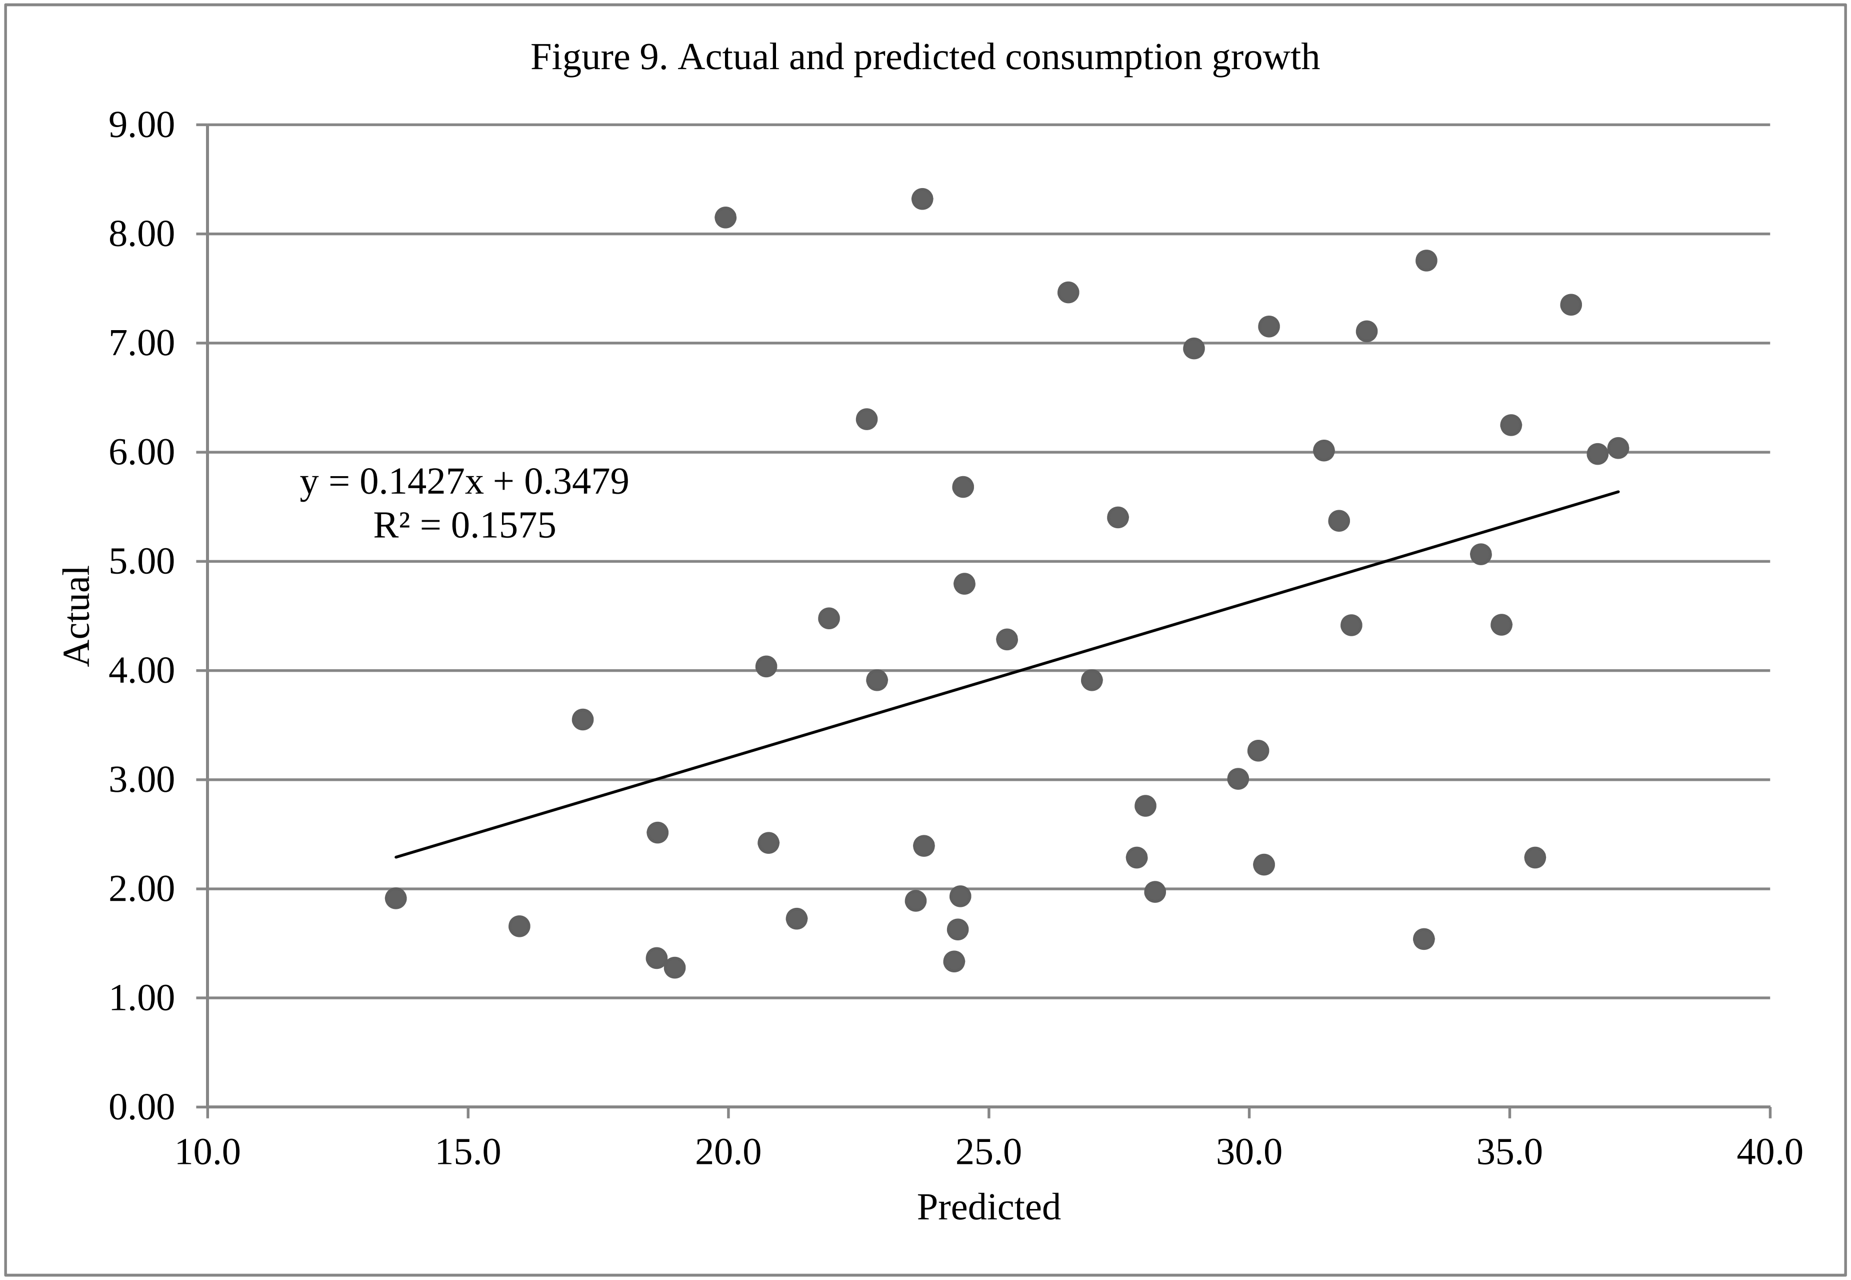

Figure 9. Actual and predicted consumption growth
y = 0.1427x + 0.3479
R² = 0.1575
Actual
Predicted
9.00
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
10.0
15.0
20.0
25.0
30.0
35.0
40.0