

This PDF is a selection from a published volume from the
National Bureau of Economic Research

Volume Title: Commodity Prices and Markets, East Asia Seminar
on Economics, Volume 20

Volume Author/Editor: Takatoshi Ito and Andrew K. Rose, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-38689-9

ISBN13: 978-0-226-38689-8

Volume URL: http://www.nber.org/books/ito_09-1

Conference Date: June 26-27, 2009

Publication Date: February 2011

Chapter Title: The Consumption Terms of Trade and Commodity
Prices

Chapter Authors: Martin Berka, Mario J. Crucini

Chapter URL: <http://www.nber.org/chapters/c11871>

Chapter pages in book: (119 - 145)

The Consumption Terms of Trade and Commodity Prices

Martin Berka and Mario J. Crucini

4.1 Introduction

Movements in a nation's terms of trade are widely viewed as important for understanding the sources of business cycle fluctuations, the dynamics of the trade balance, and economic welfare. Backus, Kehoe, and Kydland (1994) emphasize the role of productivity movements across the United States and Europe, with each assumed to specialize in a manufactured product. In their model, an increase in domestic productivity expands output at home relative to output abroad and the terms of trade deteriorates. Simply put: a large country expanding the supply of the traded good it produces must (in equilibrium) drive down the relative price of its products on world markets. The importing country's terms of trade improves, a positive spillover. Backus and Crucini (2000) add a third region to this model; a region that specializes in oil production. When the oil region cuts back production, the relative price of oil rises, improving the terms of trade of oil producers. Output falls in the United States and Europe because oil is an intermediate input in the production of manufactured goods. The business cycle implications of this model are consistent with empirical work by Hamilton (1983) showing that oil price increases are leading indicators of U.S. recessions. Mendoza

Martin Berka is senior lecturer in economics at Massey University. Mario J. Crucini is associate professor of economics at Vanderbilt University and a research associate of the National Bureau of Economic Research.

We have benefited from the comments of our discussants Mark Spiegel and Roberto S. Mariano as well as the editors, Andrew Rose and Takatoshi Ito. Mario Crucini gratefully acknowledges the financial support of the National Science Foundation (SES-0524868). Both authors are thankful to Martin Young, Head of Department, Economics and Finance at Massey University, for supporting the visit that facilitated this collaboration.

(1995) studies the terms of trade and business cycles in an extensive cross-country panel using a partial equilibrium business cycle model where terms of trade movements are exogenous. In his theoretical setting, terms of trade shocks are analogous to lotteries with the sign and magnitude of the payout dependent upon a country's pattern of specialization across an array of internationally traded goods.

Three features handicap the practical value of most theoretical models of the terms of trade, including those mentioned previously. First, the models have too few countries. In the two-country model, it must be true that at least one of the two trading partners is large enough to alter its terms of trade. In practice, countries with this amount of market power are in the minority. Even more problematic is the fact that the terms of trade of the two countries are perfectly negatively correlated with one another since one is the inverse of the other. In contrast, the terms of trade of net oil importers tend to correlate positively with each other and negatively with those of oil exporters, particularly during periods of volatile oil prices. Second, the models have too few goods. Adding countries that specialize in production means adding goods as well. This makes the aggregate terms of trade a blunt instrument for identifying the underlying sources of terms of trade movements. This is why the empirical literature tends to focus on less-aggregated measures of the terms of trade, at a minimum: energy and nonenergy components. Third, most international trade is firm-to-firm or intrafirm (e.g., a multinational purchasing parts from a foreign subsidiary) involving intermediate inputs. In contrast, consumers purchase most of the items they consume from retailers in local markets. It may be more appropriate to think of producers and consumers interacting in segmented markets that are part of a larger equilibrium process.

This chapter conducts a forensic analysis of the sources of terms of trade variation of thirty-eight countries, over the period 1990 to 2005. What makes our analysis forensic is the use of microprice data from the Economist Intelligence Unit (EIU) to parse the variance of the aggregate terms of trade into the contributions of individual goods. The microprice data in conjunction with trade shares helps us isolate the source of a nation's terms of trade variation in the space of goods. Knowledge of trade shares and economic clout of countries or regions in export markets provides indirect evidence on the national origins of terms of trade shocks (Organization of Petroleum Exporting Countries [OPEC]'s role in the oil market, for example).

The use of retail prices in our terms of trade construct is nonstandard. It is intended to distinguish between relative prices of traded goods faced by a nation's consumers from the more conventional definition using trade prices at the border. We refer to these constructs as the *consumption terms of trade* and the *production terms of trade*, respectively. Conceptually, the consumption terms of trade is the relative price that motivates shifts in consumption demand between the home export and imports, while the production terms

of trade is the relative price that influences resource allocation across the export- and import-competing sectors.

If producers and consumers face the same prices for imported and exported goods, the consumption and production terms of trade are equal and a common *terms of trade* prevails. Improvements in the terms of trade motivate domestic producers to shift resources toward the production of exports and away from imports, with consumption shifting in the opposite direction. When the terms of trade in consumption differs from the terms of trade in production, the consumption, production, and trade balance implications are altered in fundamental ways. Yet, before these implications can be understood, we need to know how much the consumption and production terms of trade differ and understand the underlying sources of those differences.

The analysis begins with the study of commodity-level retail price inflation and price level inflation for the world as a whole. This is accomplished by averaging U.S. dollar prices of individual goods and services across as many cities as available in the Economist Intelligence Unit retail price survey, the source of our consumer price data. The standard deviations of these commodity-level inflation rates range from a low of 3.7 percent to a high of 11 percent; the median is about 5 percent. Averaging these global price inflation rates across goods and services provides our world inflation estimate, which turns out to have a correlation of 0.88 with the official Organization for Economic Cooperation and Development (OECD) U.S. dollar world inflation rate. This high correlation is surprising given that the EIU sample typically comprises only one city per country, and uses a different consumption basket than official estimates. It is also reassuring in the sense that the estimate appears not to be systematically biased by these differences. Next, a variance decomposition of world inflation is performed where the contribution of each good's inflation to aggregate world inflation is estimated. Prices with relatively high variation and positive comovement with other prices in the basket will contribute more to aggregate inflation variability for the same reasons that high beta stocks contribute more than their portfolio weight to the variance of a stock price index. Individual items are found to contribute vastly different amounts to price level variability. Some of the usual suspects show up at the upper end of the distribution, such as fuels, but individual food items often display annual changes not unlike that of fuel. Qualitatively, what we see is analogous to what is already known at the national level: food and energy prices are more volatile than the typical item found in the consumption basket.

The world price series at the microeconomic level form the basis of our benchmark computations of the *consumption terms of trade at world prices*. Using microdata on trade flows, we construct import and export price indices by weighting world prices (constructed from the EIU microdata) by national import and export trade shares. The ratio of the export price index

to the import price index is the terms of trade estimate. The consumption terms of trade is somewhat less volatile than the production terms of trade in levels. However, this difference is nil in growth rates, when averages across the panel of thirty-eight countries are taken. As is true of the production terms of trade, the variability of the consumption terms of trade differs vastly across countries, from a low of about 1 percent in Australia to a high of about 10 percent in Korea (in levels). Countries with high production terms of trade variability tend to have high consumption terms of trade variability in log-levels, but this is not true of growth rates. The correlation of the two measures within a country averages 0.3 for log-levels and 0.4 for growth rates. The two measures, then, are conceptually and empirically distinct.

Decompositions of the aggregate consumption terms of trade into micro-economic sources of variation at the good level is telling. The bulk of the variability for most countries in the sample is accounted for by oil, automobiles, and medicine. The role of oil in the production terms of trade has been extensively studied in the existing literature; the evidence here suggests that this feature extends to the consumption level. The fact that one need not go beyond a few key items to account for virtually all of the terms of trade variance is a more novel finding. Focusing on oil, automobiles, and medicine accounts for much of the secular swing in the consumption terms of trade over the 1990 to 2005 period in our panel. Interestingly, oil moves in an idiosyncratic fashion relative to other world prices, helping to further distinguish its role beyond its very skewed trade shares internationally.

Movements in key world prices allow a classification of thirty of the thirty-eight countries into two groups, ten countries with U-shaped terms of trade profiles (seven of these are oil exporters), and twenty countries with inverted U-shaped patterns (all are oil importers). With few exceptions, patterns outside these two groups and subtle differences within the groups are elucidated by looking at differences in the relative importance of oil, automobiles, and medicine. Ireland, for example, experiences virtually continuous improvements in the consumption terms of trade due to the pull of medicine prices on the export side serving as an effective counterweight to the drag of oil prices on the import side.

The analysis highlights the fact that what determines the marginal contribution of a good to terms of trade variability is a country's net trade share in that good or sector. Due to the extensive volume of intraindustry trade among most industrialized countries, the bulk of terms of trade risk associated with world price movements is mostly hedged via balanced trade. Isolating the sources of variation, then, requires the variance decomposition method developed here that hinges on the use of microprice data.

The final section of the chapter constructs the *consumption terms of trade at local prices* using local retail prices to measure the import and export prices rather than world prices. The trade weights remain the same as in the benchmark case. This allows the law of one price (LOP) deviations to influ-

ence the terms of trade. As one would expect, the terms of trade using local prices is typically more volatile than when world prices are used. However, the difference is surprisingly modest because the LOP deviations that distinguish world prices from local prices tend to average out across goods. The redistribution of the attribution of terms of trade variance across goods, while notable, is also modest. Mostly, the contribution of oil falls while that of medicine rises. This shift makes intuitive sense since international LOP deviations are plausibly larger for medicine than oil. This is not to say that LOP deviations are unimportant in a more general sense. Our focus on the terms of trade means we ignore most of the service sector entirely. The prices of services are known to exhibit larger and more persistent deviations from the LOP than are the prices of traded goods. It may also be true that the LOP deviations for services fail to average out across items to the extent found among traded goods, contributing to greater volatility in the aggregate real exchange rate than the consumption terms of trade.

4.2 The Terms of Trade

Consider the *production terms of trade* constructed as a constant-share-weighted average of the logarithms of export and import prices.

$$(1) \quad q_{j,t}^p = \sum_{i=1}^M \gamma_{ij} p_{ij,t}^x - \sum_{i=1}^M \omega_{ij} p_{ij,t}^m$$

where $p_{ij,t}^x$ is the free-on-board price of good i exported from country j , at date t and $p_{ij,t}^m$ is the price inclusive of insurance and freight imported into country j . These prices are denominated in U.S. dollars throughout the chapter. The export and import shares are γ_{ij} and ω_{ij} , respectively, and are assumed to satisfy: $\sum_{i=1}^M \gamma_{ij} = \sum_{i=1}^M \omega_{ij} = 1$. The i index is used on both the import and export side to account for the fact that countries may import and export the same good, or at least goods that are difficult to distinguish given the published data. The summation should be thought of as being over the union of all goods appearing on the export and import side of the nation's income and product accounts, with many goods entering with a zero trade weight on one side of the trade balance or the other.

Our primary interest is a more novel concept, the *consumption terms of trade*. The trade weights remain the same, but prices at the border are replaced with retail prices. The benchmark is the case in which the LOP holds across retail markets in all countries, for each good, in which case the *consumption terms of trade* is given by:

$$(2) \quad q_{j,t}^c = \sum_{i=1}^M (\gamma_{ij} - \omega_{ij}) p_{i,t}$$

where $p_{i,t}$ will be computed as the average U.S. dollar price of good i , across all locations. For obvious reasons, this is defined as the *consumption terms*

of trade at world prices. Note the strong implication of the assumption of common prices in all locations. The consumption terms of trade of each country is simply a different geometric weighted average of a common vector of world prices. Put differently, the world price vector forms a common basis for determining all price indices and relative prices of interest. The key insight of this level of detail is that goods in which a country has balanced trade will not contribute to the variance of the terms of trade. For example, if a country imports oil and exports coal, commodities that produce energy demanded by consumers, energy is not going to be a large part of what determines variation in that nation's terms of trade if the country has balanced trade in energy.

In the penultimate section of the chapter, this measure is compared to the *consumption terms of trade at local prices*.

$$(3) \quad \tilde{q}_{j,t}^c = \sum_i (\gamma_{ij} - \omega_{ij}) p_{ij,t}$$

Note that the trade shares are the same as before, but the prices are now retail prices paid by final consumers in country j , not world average prices. Note that while the exports and imports of good i have the same price in the home retail market, the retail prices of these same goods may differ across countries for various reasons.

The two measures of the terms of trade may be easily contrasted, as follows:

$$(4) \quad \tilde{q}_{j,t}^c = q_{j,t}^c + \sum_i (\gamma_{ij} - \omega_{ij}) q_{ij,t}$$

where $q_{ij,t} = (p_{ij,t} - p_{i,t})$. In other words, the consumption terms of trade of country j at local prices equals the consumption terms of trade of country j at world prices, plus the net-trade-share-weighted average of the LOP deviations of country j where those deviations are computed relative to the world average price of each good, i .

4.3 The Data

The source of the retail price data is the Economist Intelligence Unit Worldwide Survey of Retail Prices. The sample period runs from 1990 until 2005 and spans 123 cities and 301 goods and services. As these data have now been quite extensively used elsewhere, our description is brief.¹ The value of these data in this application is that the basket contains the same items in all cities, which contrasts significantly with the practice of National Statistical Agencies, where the focus is on the goods typically consumed in a particular city. While one implication of this is that we may not match the

1. See, for example, Crucini and Shintani (2008), Frankel, Parsley, and Wei (2005), and Rogers (2007).

official Consumer Price Index (CPI) inflation of a city, an advantage is that we are not averaging prices of different goods across locations, which would not provide a meaningful estimate of worldwide commodity level inflation at the microeconomic level. The supplemental data we use includes very disaggregated import and export shares and official terms of trade data.

To ensure reasonably broad coverage of the consumption basket, only locations with at least 200 retail prices are included in the analysis. This restriction limits our sample to 82 of the 123 available cities, including 55 cities in 28 OECD countries and 27 non-OECD locations. These include thirteen cities in the United States, five in Australia and Germany, four in Canada, and two each in Japan, Spain, France, Italy, Switzerland, and New Zealand. The non-OECD locations include five cities in three oil-exporting countries, nine Asian and Latin American countries each, and three African countries.

Products whose prices are used in the construction of our terms of trade measure account for 21.1 percent of the expenditure found in the U.S. CPI basket. Given that the consumption share of tradables in the U.S. CPI is only about 31.8 percent, these products account for 66 percent of the tradables in the U.S. consumption basket. In terms of CPI subindices, coverage rates are: 93 percent for clothing, 72 percent for alcoholic beverages, 70 percent for food at home, 61 percent for transportation goods, 40 percent for personal care products, 24 percent for household furnishings, and 19 percent for recreation goods.

Our trade shares are drawn from the United Nations (UN) Comtrade database; specifically, the six-digit 2002 Harmonized System of import and export volumes in U.S. dollars for the year 2007. Each good in the EIU retail price survey is matched to one of these six-digit trade volumes. The sample includes members of the OECD, China, Brazil, Russia, India, several major oil exporting countries, and Asian exporters. Among these countries, our retail price data cover an average of 19 percent of imports and 18 percent of exports (we do not include re-exports or re-imports). The primary reason these numbers are not higher is that much of international trade is in intermediate goods, whereas retail purchases are mostly final goods. The import coverage ratios reach up to 30 percent for Greece, but are as low as 6 percent for Singapore, which seems due to the paucity of electronic goods in the EIU survey in the latter case. The export coverage ratios range between 77 percent for Saudi Arabia and 3.5 percent for Singapore.

4.4 World Inflation

Recall that our estimate of the *consumption terms of trade at world prices* requires estimates of world price series by good. The commodity-level inflation estimates are inflation rates, in U.S. dollars, of a particular good, i , $\Delta p_{i,t}$, averaged across all available cities in the EIU sample, indexed by j ,

$$(5) \quad \Delta p_{i,t} = N^{-1} \sum_{j=1}^N \Delta p_{ij,t}$$

The number of locations varies somewhat across goods. We restrict the sample of locations to those with at least 200 price observations.

Aggregate world inflation (again, in the units of the numeraire currency, the U.S. dollar) is the average across commodities of these inflation rates,

$$(6) \quad \Delta p_t = \frac{1}{84} \sum_{i=1}^{84} \Delta p_{i,t}$$

The use of equal weights may be justified theoretically by the zero degree homogeneity of demand functions, in which case the interpretation is that our aggregate price level is a numeraire, not a price index.

Figure 4.1 presents the inflation series for each of the eighty-four goods used to estimate world inflation while figure 4.2 presents the aggregate inflation rate, the left-hand side of the expression. The common inflation factor across goods is obvious from a visual inspection of figure 4.1, further confirmed by the fact that the median correlation of inflation at the good and aggregate level is a remarkable 0.92.

World inflation averaged 1.7 percent over the sample period (1991 to 2005). Two years exhibited significant deflation, 1997 (−6.0 percent) and 2000 (−7.6 percent), while inflation was very high in 1994 (6.9 percent), 2003 (8.6 percent), and 2005 (6.8 percent). The correlation of this inflation measure with the official estimate of OECD inflation in U.S. dollars is 0.88.

To more fully understand the role of individual prices in the evolution of the aggregate inflation rate, we use the portfolio-inspired variance decomposition used by Crucini and Landry (2009) to study the microeconomic sources of aggregate real exchange rate variation. The variance of aggregate inflation may be expressed in terms of the covariance of aggregate inflation and good-level inflation:

$$(7) \quad \text{var}(\Delta p_t) = \frac{1}{84} \sum_{i=1}^{84} \text{cov}(\Delta p_{i,t}, \Delta p_t)$$

Dividing through by the variance of aggregate inflation gives the variance decomposition:

$$(8) \quad 1 = \frac{1}{84} \sum_{i=1}^{84} \beta_i$$

The decomposition centers the distribution of the contributions of good-level inflation to aggregate inflation variability, the average beta, at unity. Thus, the average good contributes its weight to the total variance, 1/84. Goods with betas exceeding unity contribute more than the average good, while goods with betas less than unity contribute less. Since beta can be negative, a commodity may reduce aggregate inflation variability. The inter-




Fig. 4.1 World inflation rates, by commodity




Fig. 4.2 Average (across goods) world inflation

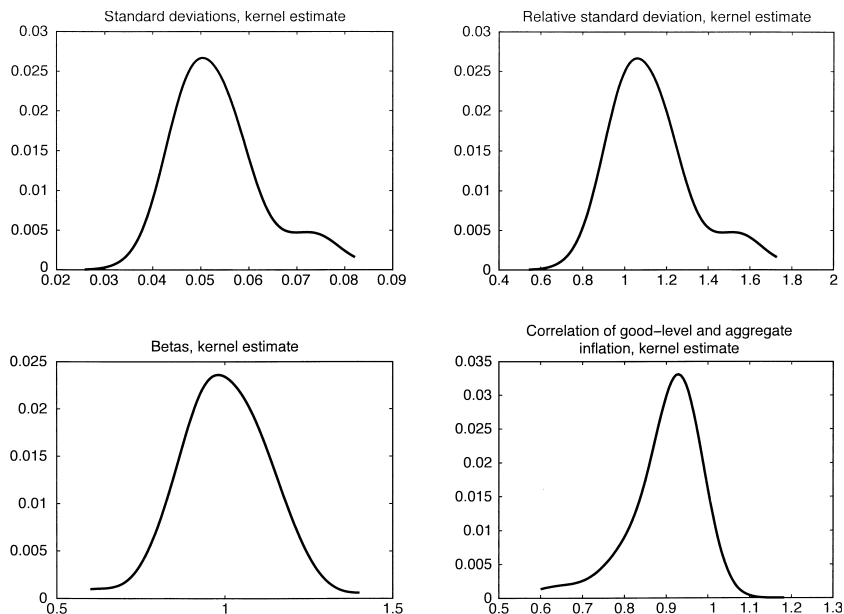


Fig. 4.3 Variance decomposition of world inflation

pretation is that adding such a price to the commodity basket will reduce the variance of the aggregate inflation rate. As it turns out, no commodity-level inflation in the microsample has a negative covariance with the world inflation level—the lowest beta is 0.56.

The betas may alternatively be expressed in terms of the standard deviation of commodity-level inflation relative to aggregate inflation multiplied by the sample correlation of the two:

$$(9) \quad \beta_i = \frac{\sigma_i}{\sigma} \text{corr}(\Delta p_{i,t}, \Delta p_t).$$

Figure 4.3 plots kernel density estimates relating to this decomposition. The upper-left chart is the kernel estimate of the standard deviation of inflation across goods; it ranges from about 3.75 percent to about 8.68 percent (see also table 4.1). There is considerable central tendency of commodity-level inflation near the level of aggregate inflation variability, 4.75 percent. The upper-right panel is the relative standard deviation, one component that influences how individual goods contribute to aggregate inflation variability. The values range from a low of 0.79 to a high of 1.83.

The lower-left panel is the estimated distribution of the betas, the distribution of the contributions of commodity-level inflation to the variance of aggregate inflation. These average to 1 by construction, but vary considerably across goods, from a low of 0.56 for oil to a high of 1.45 for lettuce. It

Table 4.1 Inflation variance decomposition

Moment	Min	Mean	Median	Max
σ_i	3.75	5.40	5.26	8.68
β_i	0.56	1.00	0.99	1.45
σ_i/σ	0.79	1.14	1.12	1.83
$\text{corr}(\Delta p_{i,t}, \Delta p_t)$	0.35	0.89	0.92	0.99

may seem surprising that oil contributes the least to world inflation variability, given the attention oil prices draw in discussions of monetary policy. The conventional wisdom that oil is among the most variable prices is valid, even at the retail level: it ranks fourth among the eighty-four commodities in our inflation construct. What sets oil apart is that it has the lowest correlation with the aggregate inflation level of any commodity in our sample, at 0.35. The median correlation of good-level inflation with aggregate inflation is 0.92. The final chart is a kernel estimate of the correlation of commodity-level inflation with the aggregate inflation rate.

4.5 The Consumption Terms of Trade

Recall that the *consumption terms of trade at world prices* is defined as:

$$(10) \quad q_{j,t}^c = \sum_{i=1}^{84} (\gamma_{ij} - \omega_{ij}) p_{i,t}$$

Figures 4.4 and 4.5 present a comparison of our estimate of the U.S. consumption terms of trade and the conventional production terms of trade, as well as the import and export price indices used in the construction of each. Because the official data is available quarterly but our retail price data is annual, we present a figure with the original official data as well as a version where we take quarterly averages to make them more comparable to our estimates. Each figure contains four charts: the left-most charts are the terms of trade, the differences between the two lines in the right-hand charts, which contain the import and export price indices.

The U.S. consumption terms of trade displays a distinctive secular swing from 1990 to 2005. This is true of ten of the twelve cases for which we have both measures of the terms of trade (not shown). The most frequent pattern, found in six of twelve cases (Finland, France, Italy, Korea, Netherlands, and the United States), is an inverted U-shape, with terms of trade improvements followed by deterioration. Four are reversed, U-shaped patterns (Australia, Canada, Denmark, and New Zealand), with the terms of trade deteriorating during the first half of the sample and then improving in the second half. Two terms of trade measures exhibit virtual continuous improvement (Switzerland and the United Kingdom). We will explore these striking similarities in the next section.

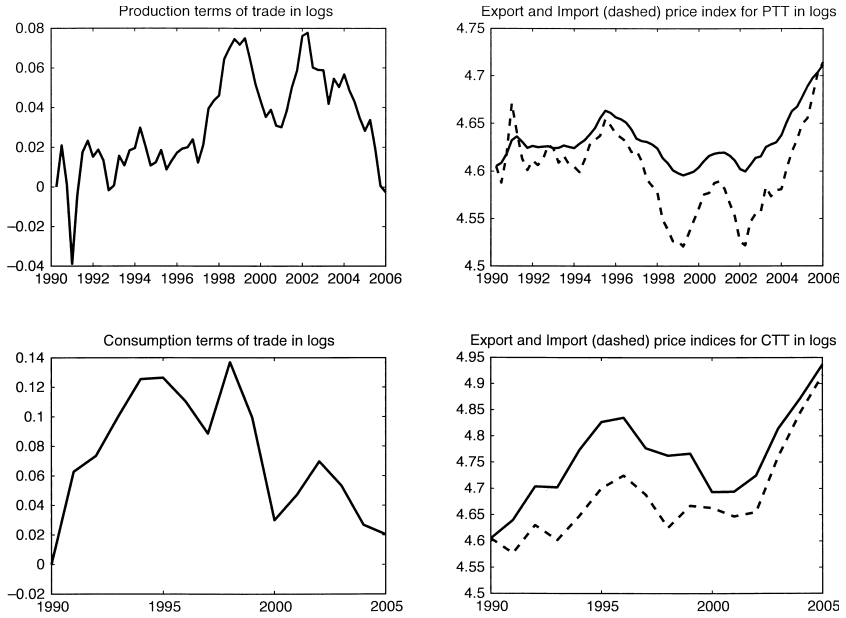


Fig. 4.4 U.S. production terms of trade (quarterly) and consumption terms of trade (annually), 1990 to 2005

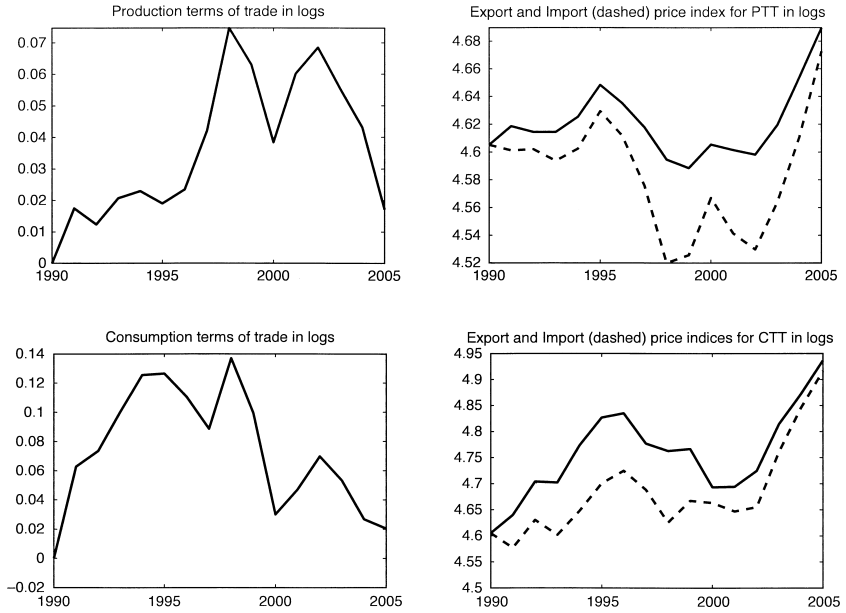


Fig. 4.5 U.S. production terms of trade and consumption terms of trade annually, 1990 to 2005

Table 4.2 Production and consumption terms of trade (world prices)

Country	Log-levels			Growth rates		
	σ_{q^p}	σ_{q^c}	$\rho_{q^p q^c}$	$\sigma_{\Delta q^p}$	$\sigma_{\Delta q^c}$	$\rho_{\Delta q^p \Delta q^c}$
Australia	9.40	1.09	0.28	4.79	0.68	0.28
Canada	3.94	1.29	0.80	2.93	0.98	0.74
Denmark	2.67	2.03	0.17	0.95	1.58	0.43
Finland	3.51	4.04	0.46	2.52	3.11	0.30
France	1.21	3.17	0.07	1.36	2.60	0.38
Italy	2.92	3.24	-0.11	2.84	2.76	0.53
Korea	16.05	9.09	0.56	3.06	6.45	0.75
Netherlands	1.72	3.00	-0.32	0.58	3.09	0.31
New Zealand	4.31	2.95	0.00	2.79	1.91	0.26
Switzerland	3.41	3.10	0.57	2.34	1.23	0.19
United Kingdom	3.02	1.30	0.77	1.78	0.62	0.02
United States	2.27	4.16	0.22	1.72	3.43	0.71
Averages	4.54	3.21	0.29	2.30	2.37	0.41

Turning to the production terms of trade, the relative price of exports to imports using prices at the border, the patterns share similarities and differences to the consumption terms of trade. The distinctive U-shapes and inverted U-shapes are largely gone. Denmark, the Netherlands, New Zealand, Switzerland, and the United Kingdom show general terms of trade improvements. Canada and Australia maintain some of their original U-shaped paths. Finland, Korea, and Italy have deteriorating terms of trade over much of the sample. The picture for France is ambiguous, due to low variability. The United States has a somewhat inverted U-shape (see figure 4.4 or 4.5), but the timing is different from what the consumption terms of trade shows.

Table 4.2 reports standard deviations of log-level and growth rates for both terms of trade measures, as well as their contemporaneous correlation. The production terms of trade tend to be more volatile than the consumption terms of trade, though this difference largely disappears in the move to growth rates, where the average standard deviation in the production terms of trade is 2.3, compared to 2.4 for the consumption terms of trade. Thus, the differences between the two is not merely less volatile prices at the retail level than at the border. The two measures move weakly together in log-levels, where the correlation is about 0.3, on average, and somewhat more strongly in growth rates, where the correlation is 0.4, on average.

Since the production and consumption terms of trade are different conceptually, it is not clear that a high correlation between them is to be expected. Moreover, the construction of the consumption terms of trade uses average international retail prices, while the production terms of trade uses prices at the customs point of entry or exit of each country. Given that

deviations from the LOP have been widely documented in the literature, this is another source of difference between the two measures. The role of LOP deviations is evaluated in the penultimate section of the chapter. The question we turn to next is: what is generating the trends and fluctuations in the consumption terms of trade? As was noted earlier, there appear to be a few common secular trends shared by certain groups of countries. It will be interesting to see if those common features are driven by trade patterns and particular properties of a few key international prices, such as the price of oil.

4.5.1 Variance Decomposition

The variance of a nation's terms of trade satisfies the following equation:

$$(11) \quad \text{var}(q_{j,t}^c) = \frac{1}{84} \sum_{i=1}^{84} (\gamma_{ij} - \omega_{ij}) \text{cov}(q_{j,t}^c, p_{i,t}).$$

Dividing both sides by the variance of the terms of trade gives the variance decomposition:

$$(12) \quad 1 = \frac{1}{84} \sum_{i=1}^{84} (\gamma_{ij} - \omega_{ij}) \beta_{i,j}.$$

The betas are effectively the coefficients from a regression of the commodity-level price on the consumption terms of trade $\beta_{i,j} = \text{cov}(q_{j,t}^c, p_{i,t}) / \text{var}(q_{j,t}^c)$.

The analysis starts by pulling back to the broadest picture, pooling all goods and locations. Figure 4.6 plots two kernel density estimates, one for the net trade share and the other for the betas. These two distributions contain all of the elements needed to decompose the variance of the terms of trade.

The net trade shares lie almost exclusively between plus and minus 5 percent. Recall that these are normalized so that for each country the import shares (and export shares) sum to unity. The net trade shares for oil are extreme outliers, averaging 40 percent across countries, while various categories of cars have an absolute net trade share that averages about 30 percent. Not surprisingly, these items will exert a disproportionately large influence on the terms of trade.

The beta distribution lies mostly between plus and minus three, with a strong central tendency toward the mean of about 0.25. Since the contribution to variance is the product of the net trade share and beta, values in the tail of the distribution are what dominate the variance decomposition of the terms of trade. To see this more clearly, it is productive to look to the details of the distribution of products of the net import share and beta: $(\gamma_{ij} - \omega_{ij}) \beta_{i,j}$.

Figure 4.7 plots the contributions to variance, the product of the net trade share and beta, for each commodity and country in our sample. Since there are eighty-four commodities and thirty-eight countries, there are 3,192 values plotted in this figure. The upper panel orders the contributions by





Fig. 4.6 Kernel estimates of net share shares (top chart) and betas (bottom chart), pooling all goods and countries





Fig. 4.7 Contributions to terms of trade variation

Notes: Top chart orders contributions by good (1 to 84), with vertical segment marking goods. Bottom chart is the same information, but with vertical segments marking countries (as labeled).

commodity, and the lower panel orders the contributions by country. In the upper panel, the vertical lines mark the variance contributions by good with points between the lines denoting a country-specific variance contribution for that good. In the lower panel, the vertical lines mark variance contribution by country with points between the lines denoting good-specific variance contributions for that country. The variance contributions organized by country will sum to unity within each interval by construction.

The clusters of extreme values in the upper chart identify commodities that contribute considerably more to terms of trade variation than the typical item. To see this more clearly, Figure 4.8 focuses on the seven most important contributors to consumption terms of trade variability. In order of ascending importance, they are: pullovers; boneless beef; luxury, compact, and large cars; medicine; and oil. The contribution from oil averages 0.6. In other words, oil alone accounts for about 60 percent of the variation in the terms of trade when we average across our thirty-eight nations. Medicine accounts for about 12 percent of terms of trade variation, the three automobile categories combined are comparable at 11 percent, while pullovers and boneless beef each account for about 2 percent. To place these numbers in perspective, the next twenty items in the ranking combine to account for the same percentage as medicine. The reader should keep in mind that the composition of influences differs across countries, which is masked by the cross-country averaging discussed here. The cross-country differences in the contribution of each commodity is visible in the variation within the commodity partitions of figure 4.8.

4.5.2 Goods Prices and the Terms of Trade

Based on the variance analysis (displayed in figure 4.8), it seems sufficient to focus on the time paths of the U.S. dollar prices of the items found to be most influential in the evolution of the aggregate consumption terms of trade: oil, automobiles, and medicine. Figure 4.9 plots the U.S. dollar prices of these five goods.

Two features of these price histories are worth emphasizing. The first is that oil prices have a large idiosyncratic component; the other four series track each other very closely. This is consistent with the earlier decomposition of the variance of world inflation. The median correlation of good-level and aggregate inflation is 0.92, while the correlation of oil inflation with world inflation is a mere 0.35.

The second striking feature of figure 4.9 is that oil prices are not much more variable than the typical commodity price in these data. This is a reflection of two facets of our analysis. First, the fact that we use prices paid by final consumers rather than prices determined in commodity exchanges such as the Chicago Board of Trade. Thus our “oil” price is a retail fuel price, not the price of a barrel of crude petroleum. The former is much less volatile than the latter at annual frequencies, in most time periods. Second, we use




Fig. 4.8 Contribution to terms of trade variance, ordered by commodity, seven most influential goods

Notes: Dashed line marks cross-country average variance contribution by good.




Fig. 4.9 Nominal, U.S. dollar, price indices of key traded commodities

microdata that highlights the fact that the retail prices do move around a great deal, and the aggregate CPI index tends to obscure this by averaging away much of the idiosyncratic variation. Thus, retail prices are much more volatile than the price level, which is more familiar to macroeconomists.

The distinctive paths of these prices along with their dominate contribution to the terms of trade variance for the median country, already documented, suggests a convenient link between the aggregate consumption terms of trade and a few key prices. In a nutshell, countries with net positive exposure to oil (net exporters) should have an inverted U-shaped terms of trade path, following the path of oil's price, while those with a net negative exposure in oil and positive exposure in medicine or automobiles should have a U-shaped pattern, following the evolution of these other prices.

To document this as clearly as possible, the terms of trade is built up in stages, beginning with the oil terms of trade, then adding medicine, then automobiles, and finally, everything else, to arrive at the aggregate consumption terms of trade. Figures 4.10 and 4.11 do precisely this for each of the thirty-eight countries in our sample. Figure 4.10 focuses on the eight oil exporters in our panel: Canada, Denmark, Mexico, Norway, Russia, Saudi Arabia, United Arab Emirates, and the United Kingdom. Figure 4.11 focuses on the thirty oil net importers.

The terms of trade for oil is, by definition, just the path of oil prices at the retail level (with 1990 = 0, the base year), scaled by the net trade share in fuel.² Thus, the dashed lines in all figures are either perfectly positively correlated with the world price of oil (for net exporters) or perfectly negatively correlated with the world price of oil (for net importers). If this was the complete picture, net exporters would experience a secular decline in their terms of trade followed by a secular rise due to oil's price movements—a U-shaped pattern. For net importers, we would see an inverted U-shaped terms of trade profile. Moreover, the dashed lines (oil terms of trade) and the solid black lines (overall terms of trade) would be the same.

While this is, of course, an oversimplification, it is the case that oil dominates the secular movements in almost every case, with net exporters of oil having U-shaped terms of trade (figure 4.10) and net importers of oil tending toward an inverted U-shaped terms of trade (figure 4.11). The pattern among net importers of oil is more complex than this stylized description, partly because their trade patterns are more complex on both the export and import side. Oil exporters, in contrast, tend to be concentrated on the export side and less concentrated on the import side. For them, complexity lies on the import side of the ledger, for the most part.

In most of the oil importing countries, the terms of trade in oil is the lower envelope of the other terms of trade constructs. What prevents the

2. The net fuel share is zero for Hong Kong, Iceland, and Luxembourg. Thus, for these three the oil terms of trade is constant, reflecting a perfect hedge.

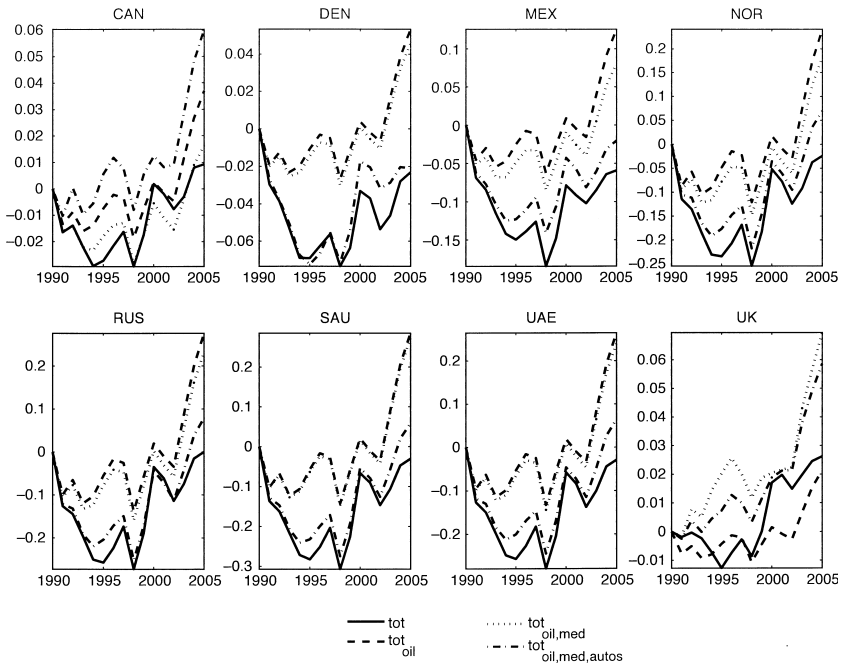
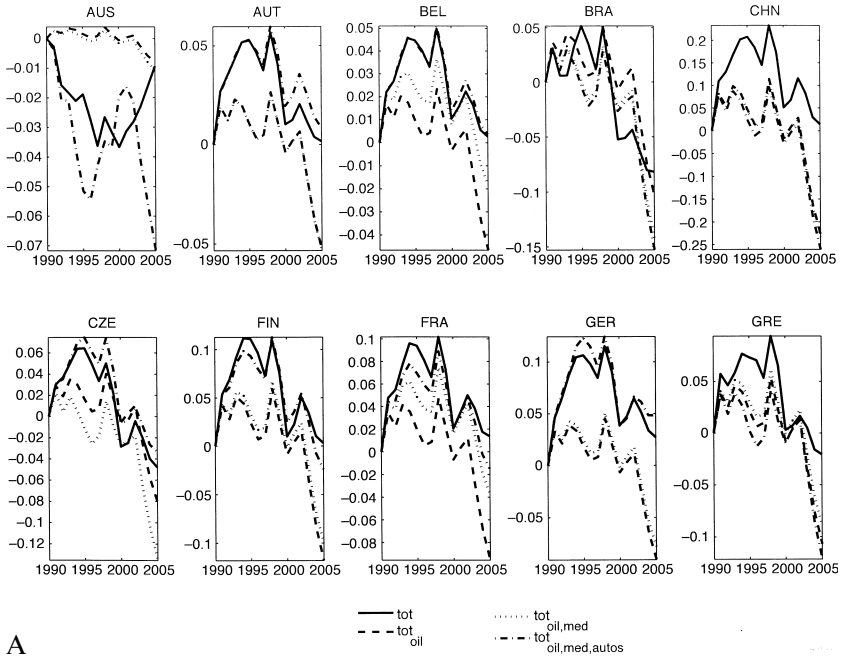


Fig. 4.10 Oil exporters: Consumption terms of trade decomposition



A

Fig. 4.11, Panel A Oil importers: Consumption terms of trade decomposition

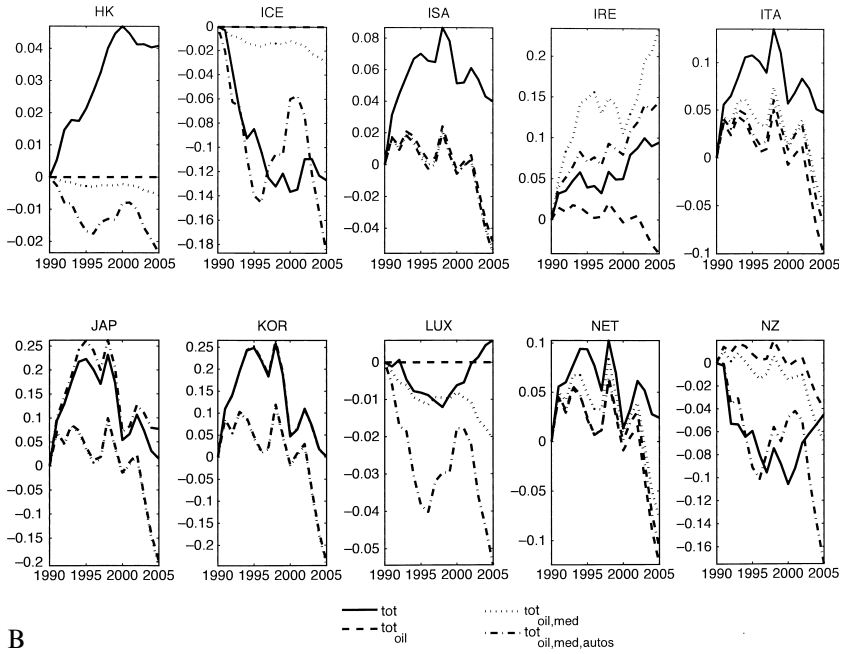


Fig. 4.11, Panel B Oil importers: Consumption terms of trade decomposition

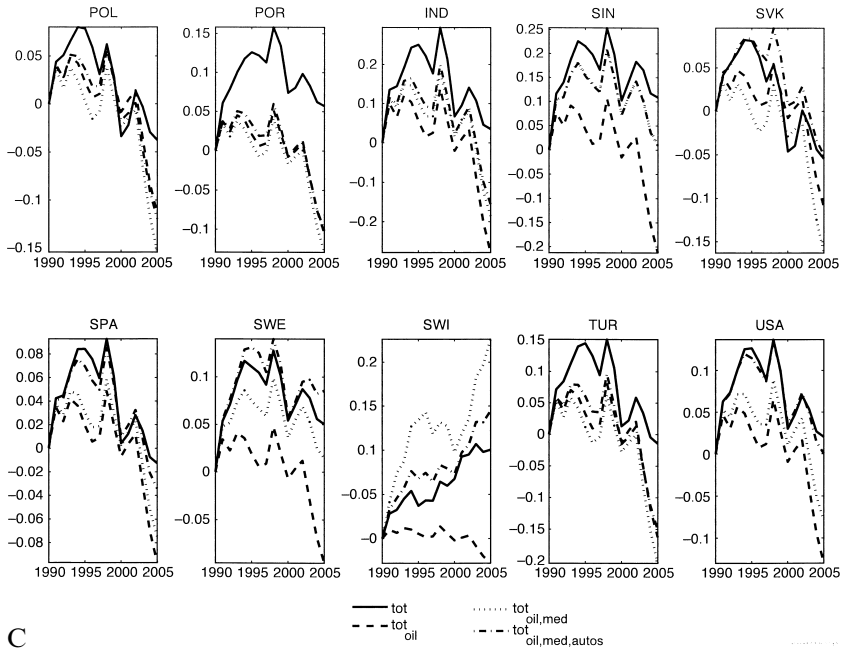


Fig. 4.11, Panel C Oil importers: Consumption terms of trade decomposition

overall terms of trade from behaving similarly is that other items are fueling improvements at the same time oil is sapping the fuel. Perhaps the clearest examples of this is Ireland (figure 4.11, panel B), where oil prices contribute to terms of trade deterioration over the last third of the sample, but this is completely swamped by the improvements in the terms of trade in medicine. Moreover, during the first third of the sample, medicine and oil prices are moving in the same direction, reinforcing the improving trend in Ireland's terms of trade. Automobiles, a net import, are a drag on Ireland's terms of trade more significant even than oil (comparing the dashed lines of panel B). However, medicine is sufficient to keep the Irish terms of trade rising on trend for almost the entire sample period.

Korea is an even clearer case in point as a major oil importer and automobile exporter. Only two lines are visible; medicine and other goods play no role (this is not to say other goods are individually unimportant, as they may average out across goods). Automobile price increases buoyed Korea's terms of trade until the last third of the sample when oil prices rose relative to automobile prices. Israel is a case where oil is a significant terms of trade drag, but automobiles and medicine are not helpful in accounting for the terms of trade. Here, exploration of the sources of variation would need to go beyond these three items.

4.5.3 The Consumption Terms of Trade at Local Prices

Up to this point, we have maintained that the LOP price holds for retail prices. However, a large literature emphasizes that LOP deviations are large and persistent, particularly so when consumer prices are the focus. The *consumption terms of trade at local prices* allows for these deviations:

$$(13) \quad \tilde{q}_{j,t}^c = \sum_i (\gamma_{ij} - \omega_{ij}) p_{ij,t}$$

Figure 4.12 presents a scatterplot of $\tilde{q}_{j,t}^c$ against $q_{j,t}^c$ pooling all cities and time periods. The asterisks are levels while the open dots are changes. The correlation of the two is 0.59 in levels and 0.45 in changes, when all locations and time periods are pooled together. As expected, the variability of the terms of trade is generally higher when local prices are used, as is evident in the increase in time series variance from 3.28 to 5.62 for the median country in levels and from 2.69 to 4.68 for growth rates.

Allowing for LOP deviations matters for some countries more than others. To see the heterogeneity across countries, table 4.3 reports the correlation of the two terms of trade measures for each country. The median correlation is 0.60, which turns out to be Finland; the range of correlations is very wide, from -0.62 for Hong Kong to 0.91 for the United States. When we examine the decomposition of the variance across goods, the reasons for the differences at the aggregate level will become clear.

Toward this end, consider how the variance of the terms of trade gets

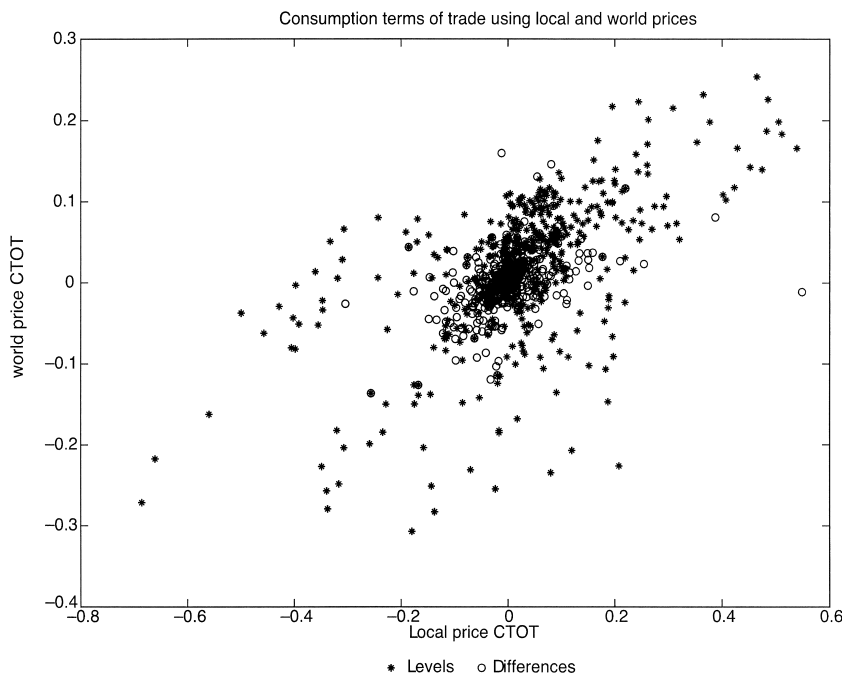


Fig. 4.12 Consumption terms of trade using local and world prices

distributed across goods based on our methodology. Taking the variance decomposition for the consumption terms of trade at local prices less the variance decomposition for the consumption terms of trade at world prices gives:

$$0 = M^{-1} \sum_{i=1}^M (\gamma_{ij} - \omega_{ij})(\tilde{\beta}_{i,j} - \beta_{i,j}),$$

where $\tilde{\beta}_{i,j} = \text{cov}(\tilde{q}_{i,t}^c, p_{ij,t})$ and $\beta_{i,j} = \text{cov}(q_{i,t}^c, p_{i,t})$ —the β 's from the variance decompositions of the two measures of the terms of trade. Since both measures use the same weights, the shift in the allocation of variance across goods boils down to a comparison of β 's across goods. Goods with $\tilde{\beta}_{i,j} > \beta_{i,j}$ will have their role in the variance decomposition elevated in the move to local prices, while those with the opposite sign will be demoted in relative importance.

The fact that relatively few goods account for the bulk of the variation of the consumption terms of trade at world prices is also true for the consumption terms of trade at local prices. Table 4.4 reports the contribution to variance of the twelve most important goods for both measures of the terms of trade. The two lists have eleven common goods. Moving from world prices to local prices causes the variance contribution of oil to fall by about

Table 4.3 Comparisons of consumption terms of trade at world and local prices

Country	Log-levels			Growth rates		
	σ_{q^c}	$\sigma_{\tilde{q}^c}$	$\rho_{q^c, \sigma_{\tilde{q}^c}}$	$\sigma_{\Delta q^c}$	$\sigma_{\Delta \tilde{q}^c}$	$\rho_{\Delta q^c, \Delta \tilde{q}^c}$
Australia	1.09	2.81	-0.23	0.68	1.47	0.31
Austria	1.98	2.49	0.29	1.42	3.41	0.39
Belgium	1.60	1.68	0.26	1.23	1.96	0.54
Brazil	4.60	15.92	0.86	2.76	8.70	0.46
Canada	1.29	1.55	0.59	0.98	1.84	0.69
Denmark	2.03	3.34	0.71	1.58	2.42	0.56
Finland	4.04	5.16	0.62	3.11	4.67	0.62
France	3.17	4.90	-0.34	2.60	4.70	0.40
Germany	3.33	2.38	0.91	2.49	1.90	0.69
Greece	3.68	8.80	0.47	2.84	6.95	0.81
Hong Kong	1.51	4.12	-0.62	0.37	1.63	0.50
Ireland	2.76	9.90	0.70	1.49	6.55	0.62
Italy	3.24	4.01	0.29	2.76	2.03	0.65
Japan	7.79	10.10	0.32	5.47	5.14	0.69
Luxembourg	0.55	2.72	-0.58	0.27	1.19	-0.35
Mexico	4.54	12.14	0.66	3.52	7.61	0.51
Netherlands	3.00	6.31	0.79	3.09	5.69	0.48
New Zealand	2.95	6.08	-0.08	1.91	4.99	-0.07
Norway	8.04	7.52	0.37	6.39	8.08	0.38
Poland	4.11	12.94	0.53	2.92	7.40	-0.14
Portugal	3.80	6.90	0.79	2.94	3.14	0.73
Saudi Arabia	9.62	31.73	0.54	7.57	21.02	0.34
Singapore	6.12	13.45	0.68	5.63	11.13	0.55
Spain	3.56	2.63	0.81	2.58	2.88	0.68
Sweden	3.22	3.66	0.63	2.63	2.82	0.75
Switzerland	3.10	3.70	-0.21	1.23	3.27	0.11
United Arab Emirates	8.72	14.16	0.88	7.02	7.72	0.65
Turkey	5.68	12.24	0.86	4.28	10.10	0.65
United Kingdom	1.30	2.71	0.65	0.62	2.25	0.15
United States	4.16	6.79	0.91	3.43	5.77	0.92
Median	3.28	5.62	0.60	2.69	4.68	0.55
Maximum	9.62	31.73	0.91	7.57	21.02	0.92
Minimum	0.55	1.55	-0.62	0.27	1.19	-0.35

10 percent, from 0.61 to 0.51, while the contribution of medicine rises by 11 percent, from 0.13 to 0.24. This is consistent with LOP deviations tending to be larger for medicine than oil.

The last part of the analysis revisits the question of the correlation between the production terms of trade and the consumption terms of trade, now allowing for LOP deviations. Table 4.5 shows that the consumption terms of trade is now consistently more variable than the production terms of trade (eight of thirteen countries). The two are less correlated than what was found using world prices, possibly due to larger deviations from the LOP at the retail stage than at the border.

Table 4.4 Microeconomic drivers of the consumption terms of trade

Good	Net trade share	At world prices		At local prices	
	$\gamma_{ij} - \omega_{ij}$	β_{ij}	var_{share}	$\widetilde{\beta}_{ij}$	\widetilde{var}_{share}
Oil	-0.16	-0.04	0.61	-0.60	0.51
Medicine	0.04	0.37	0.13	0.20	0.24
Large car	-0.01	0.30	0.06	0.26	0.04
Boneless beef	0.01	0.28	0.02	0.98	0.03
Wine	*	0.35	*	0.18	0.03
Pullovers	0.01	0.30	0.02	0.62	0.03
Footwear	*	0.29	0.01	0.43	0.02
Car compact	0.01	0.28	0.03	0.32	0.02
Fresh fish	0.01	0.21	*	0.36	0.01
Cheese	0.01	0.33	*	0.47	0.01
Creams	*	0.31	*	0.41	0.01
Pork chops	*	0.16	*	0.53	0.01
Apples	*	0.16	*	-0.22	*
Tomatoes canned	*	0.11	*	0.56	*
Pots	*	0.17	*	-0.13	*
Cigarettes	*	0.42	0.01	0.81	-0.01

Note: The cells report cross-country average values.

*Significant at less than 1 percent level.

Table 4.5 Production and consumption terms of trade (local prices), 1990–2005

Country	Log-levels			Growth rates		
	σ_{qp}	$\sigma_{\tilde{q}^c}$	ρ_{qp, \tilde{q}^c}	$\sigma_{\Delta qp}$	$\sigma_{\Delta \tilde{q}^c}$	$\rho_{\Delta qp, \Delta \tilde{q}^c}$
Australia	9.40	2.81	0.73	4.79	1.47	0.04
Canada	3.94	1.55	0.51	2.93	1.84	0.55
Denmark	2.67	3.34	-0.35	0.95	2.42	-0.04
Finland	3.51	5.16	0.17	2.52	4.67	0.21
France	1.21	4.90	0.25	1.36	4.70	0.09
Italy	2.92	4.01	-0.51	2.84	2.03	0.33
Netherlands	1.72	6.31	-0.41	0.58	5.69	-0.00
New Zealand	4.31	6.08	-0.17	2.79	4.99	-0.44
Switzerland	3.41	3.70	0.08	2.34	3.27	-0.06
United Kingdom	3.02	2.71	0.62	1.78	2.25	-0.09
United States	2.27	6.79	0.51	1.72	5.77	0.80
Averages	3.49	4.31	0.13	2.24	3.55	0.13

4.6 Discussion

Our findings regarding the importance of oil are reminiscent of the analysis of Backus and Crucini (2000), who documented the extraordinary extent to which oil dominated the variation in the terms of trade of major industrialized countries from the 1970s to the middle of the 1980s. The thrust

of their analysis was to show how business cycle comovement evolved as the importance of oil shocks relative to total factor productivity shocks changed across historical periods and affected importers differently than exporters.

Figure 4.13 displays the quarterly data used in their paper, along with the relative price of oil in U.S. terms. By the latter, we mean the U.S. dollar spot price of crude petroleum divided by the U.S. consumer price index. In the figure, the oil price is normalized so that its standard deviation matches the standard deviation of the average of the terms of trade across all the countries in the sample (the lower-right chart displays this average terms of trade variable).

The standard deviation of the average terms of trade for the group is about 5.5 percent, while the standard deviation of the relative price of oil is an astounding 74 percent. Japan has the highest terms of trade variation at 21 percent, while Switzerland's is the most stable at 3.6 percent. Differencing both the terms of trade and the relative price of oil leave the basic implication unchanged, oil plays a large role, mostly because it has enormous variation relative to the terms of trade. The ratio of standard deviation of oil to that of the terms of trade, in levels or growth rates, is about ten.

We view our preliminary findings as pointing to a broader role than was previously thought for a small set of goods to dominate a nation's terms of trade variation. Uncovering this feature of the data would have been daunting without the novel variance decomposition employed here. The fact that oil dominates in a broad cross-section is consistent with prior work on oil and the terms of trade. The notion that individual items other than oil may

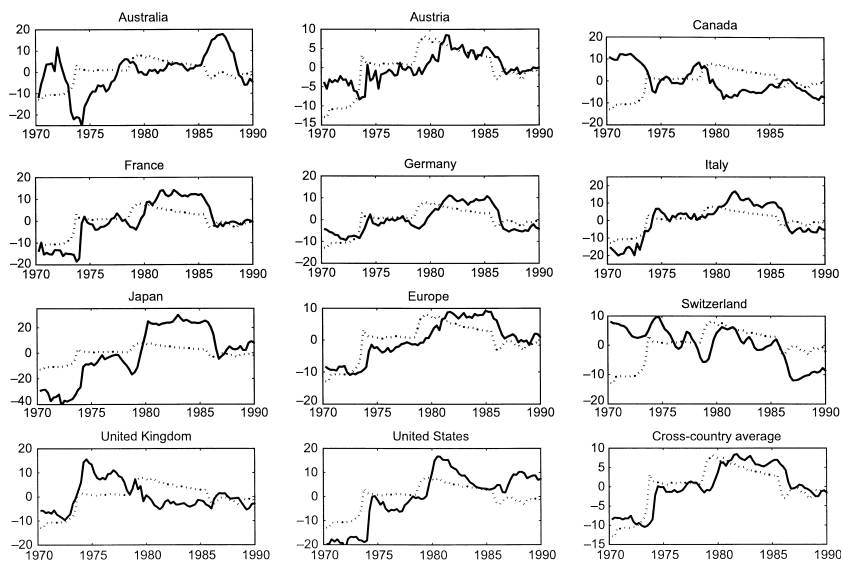


Fig. 4.13 Historical national terms of trade for major industrialized nations and the relative price of oil in U.S. terms

dominate within the cross-section of countries is novel. Moreover, it also suggests the value of organizing countries on the basis of their net export shares and a larger set of commodities than focused upon here. In focusing on the average country or group of countries, we have likely missed some important microdrivers of national terms of trade histories. It would also be interesting to consider how the influential set has evolved over time and across countries, analogous to how oil's role has been historically punctuated. Unlike oil and other commodities where comparative advantage is largely endowed, manufacturers and, increasingly, services, play a large role in trade and are likely to be more geographically footloose.

The empirical differences between the consumption and production terms of trade are compelling, though it is too early in the research program to say how they relate to the broader literature on markups and distribution costs. If the consumption terms of trade is fundamentally different than the production terms of trade, the trade balance adjustment process on the demand and supply side needs to be elaborated. The common use of one elasticity to relate prices to the trade balance condition is likely muddling consumption and production elasticities and two relative prices (the consumption and production terms of trade) rather than one.

Our results are subject to a number of important caveats. First, the Economist Intelligence Unit sample, while comprehensive, certainly does not cover the universe of consumption items and misses intermediate goods that are used by firms and not used by consumers. This, combined with the need to reconcile final goods with trade shares, leads inevitably to some errors and omissions in prices and trade weights. Second, the short sample also prevents us from backcasting our analysis before 1990, when the EIU survey was first developed. We hope to deal with some of these issues in future work, such as using the Penn World Table data to push the sample back in time. Finally, the comparisons with the official, or production terms of trade as we call it, is a crude starting point. We lack microprices at the border to conduct an analogous variance decomposition of the official terms of trade into its microeconomic determinants. Such an exercise will likely assign an important role for price variation of different brands in the terms of trade. That is, the hedging argument implicit in our use of the net trade share interacted with a single price (either the world price or the local price) becomes an imperfect hedge in the realistic cases in which the U.S. dollar inflation rates of, say, imported Mercedes and exported Infinity sedans are less than perfectly correlated. Much remains to be done.

References

- Backus, D., and M. Crucini. 2000. Oil prices and the terms of trade. *Journal of International Economics* 50 (1): 185–213.
- Backus, D., P. Kehoe, and F. Kydland. 1994. Dynamics of the trade balance and the terms of trade: The J-curve? *American Economic Review* 84 (1): 84–103.
- Crucini, M. J., and A. Landry. 2010. Accounting for real exchange rates using micro-data. Vanderbilt University. Unpublished manuscript.
- Crucini, M., and M. Shintani. 2008. Persistence in law-of-one-price deviations: Evidence from micro-data. *Journal of Monetary Economics* 55 (3): 629–44.
- Frankel, J., D. Parsley, and S.-J. Wei. 2005. Slow passthrough around the world: A new import for developing countries? NBER Working Paper no. 11199. Cambridge, MA: National Bureau of Economic Research, March.
- Hamilton, J. 1983. Oil prices and the macroeconomy since World War II. *Journal of Political Economy* 91 (2): 228–48.
- Mendoza, E. 1995. The terms of trade, the real exchange rate and economic fluctuations. *International Economic Review* 36 (1): 101–37.
- Rogers, J. 2007. Monetary union, price level convergence and inflation: How close is Europe to the USA? *Journal of Monetary Economics* 54 (3): 785–96.

Comment Roberto S. Mariano

In this chapter the authors apply a variance decomposition analysis based on microprice data to study the consumption terms of trade and the production terms of trade in various countries. The analysis starts with the study of retail price inflation and price level inflation for the world as a whole. This is done by averaging U.S. dollar prices of individual goods and services across cities covered in the EIU retail price survey. A variance decomposition of world inflation follows next—to estimate the contribution of each good to aggregate world inflation variability.

While variance decomposition analysis is a statistical tool that is well-developed and applied in the statistics and finance literature, this chapter is one of the first to utilize this methodology to undertake this study of consumption terms of trade and commodity prices. Furthermore, the level of disaggregation that the authors have chosen allows them to address detailed issues at the individual good/service level and by country. The authors adequately lay out the motivation and the basics of the approach in numerous sections of the chapter. One minor additional detail that would be useful is more information on just how the kernel density estimates of the betas and the individual good standard deviations are calculated in figure 4.3 (section 4.4).

Roberto S. Mariano is professor of economics and statistics and dean of the School of Economics at Singapore Management University, and professor emeritus of economics and statistics at the University of Pennsylvania.