Chapter 11
International Business Cycles: Theory and Evidence

David K. Backus, Patrick J. Kehoe, and Finn E. Kydland

1. Introduction

In modern developed economies, goods and assets are traded across national borders, with the result that events in one country generally have economic repercussions in others. International business cycle research focuses on the economic connections among countries and on the impact these connections have on the transmission of aggregate fluctuations. In academic studies this focus is expressed in terms of the volatility and comovements of international time series data. Examples include the volatility of fluctuations in the balance of trade, the correlation of the trade balance with output, the correlation of output and consumption across countries, and the volatility of prices of foreign and domestic goods.

We consider international business cycles from the perspective of dynamic general equilibrium theory, an approach adopted in a large and growing number of studies in international macroeconomics. In closed-economy studies, models of this kind have been able to account for a large fraction of the variability of aggregate output and for the relative variability of investment and consumption. See, for example, Prescott's (1986) review. In public finance, similar models have been used to assess the impact of fiscal policy on aggregate output, employment, and saving. Auerbach and Kotlikoff (1987) are a prominent example. In international macroeconomics, this approach has been used to account for some of the notable features of international data: the time series correlation of saving and investment rates (Baxter and Crucini 1993; Cardia 1991; and Finn 1990), the countercyclical movements of the trade balance (Backus, Kehoe, and Kydland 1994; Glick and Rogoff 1992; and Mendoza 1991), and the relation between the trade balance and the terms of trade (Backus, Kehoe, and Kydland 1994; Macklem 1993; and Smith 1993).

These efforts illustrate the insights dynamic theory has contributed to date, and is likely to contribute in the future. In our view, however, the most important aspects of this line of work for future research are those for which the theory remains significantly different from the data. These discrepancies between theory and data provide focus for future theoretical work in this area.

For this reason, we focus on two striking discrepancies between current theory and data. The first concerns the relations between business cycles across countries. In the data, correlations of output across countries are larger than analogous correlations for consumption and productivity. In theoretical economies, consumption and productivity correlations are larger than output correlations. The second discrepancy concerns the terms of trade, which we define as the relative price of imports to exports. Fluctuations in the terms of trade are much more variable in the data than in theoretical economies.

We examine cross-country comovements of aggregate quantities, including output and consumption, in the natural extension of Kydland and Prescott's (1982) closed-economy model to an international setting. In this extension, agents in the two countries produce and trade a single good. Fluctuations are driven by exogenous movements in productivity. Although the theory mimics some features of the data, the international comovements are much different from those in the data. Using parameters for the stochastic process for productivity shocks that we estimate from data for the United States and a European aggregate, we find that productivity is positively correlated across countries. In the model, however, shocks of this form give rise to output fluctuations that are less highly correlated than consumption and productivity fluctuations. The ranking of output, consumption, and productivity correlations is extremely robust: it survives large changes in a number of the model's parameters. Since these differences between theory and data are relatively insensitive to the choice of parameter values and even the model's structure, we term them collectively the consumption/output/productivity anomaly, or simply the quantity anomaly.

To examine fluctuations in relative prices, we extend the theoretical model to allow the outputs of the two countries to be imperfect substitutes. This extension allows the relative price of the two goods to differ from one. In the data, fluctuations in the terms of trade in the industrialized world have been very persistent and highly variable. These properties, and similar properties of the real exchange rate, are perhaps the most widely studied issues in international macroeconomics. We find that the model generates fluctuations in the terms of trade as persistent as they are in the data. The variability of the terms of trade, however, is generally much less in the model than in the data. We call this discrepancy the price-variability anomaly. If we lower the substitutability of foreign and domestic goods, we can increase the variability of the terms of trade, but this comes at the expense of reducing the variability of imported and exported goods far below what we see in the data.

The two anomalies concerning the behavior of international business cycles and relative prices pose a challenge for international business cycle research. With them in mind, we review a rapidly expanding body of work aimed at these and other issues and speculate on directions future work might take. Notable extensions of the theory include nontraded goods, incomplete markets, money, and imperfectly competitive firms. We argue that none of these extensions has yet to provide a persuasive resolution of the price and quantity anomalies.

2. Properties of International Business Cycles

We begin by reviewing some of the salient properties of international business cycles. These features of the data serve as a basis of comparison with theoretical economies. These properties, and others reported later, refer to moments of Hodrick-Prescott–filtered variables.[1] Our data are from the OECD's *Quarterly National Accounts* and *Main Economic Indicators* and the IMF's *International Financial Statistics*.

Table 11.1 shows a number of properties of business cycle experience since 1970 in ten developed countries and a European aggregate constructed by the OECD. We focus on volatility, measured by standard deviations; persistence, measured by autocorrelations; and comovement, measured by correlations, for a set of common macroeconomic time series. With respect to volatility, we find that while consumption has generally had about the same standard deviation, in percentage terms, as output, investment in fixed capital has been two to three times more volatile than output, and employment has been somewhat less volatile than output. There are, however, some differences across countries in the magnitudes. The standard deviation of output fluctuations ranges from a low of 0.90 percent in France to a high of 1.92 percent in the United States. We also find some differences in consumption volatility. The standard deviation of consumption, relative to that of output, is 0.75 in the United States, 1.09 in Japan, 1.14 in Austria, and 1.15 in the United Kingdom. The numbers are larger than those generally reported in studies of the United States, partly because consumption in this data set includes expenditures on consumer durables. If we exclude durables, which we can do for only five countries, the volatility ratios fall from 0.75 to 0.52 for the United States, from 0.85 to 0.59 for Canada, from 0.99 to 0.77 for France, from 0.78 to 0.61 for Italy, and from 1.15 to 0.96 for the United Kingdom. Some of these differences almost certainly reflect differences in the procedures used to construct aggregate data, but more work is needed before we can quantify the impact of disparities of measurement. In terms of persistence we see that in all countries the autocorrelation of output is high. It ranges from 0.57 for Austria to 0.90 for Switzerland.

There has been even greater variation in the volatility of employment (civilian employment from the OECD's *Main Economic Indicators*): the ratio of the standard deviation of employment to that of output ranges from 0.34 in Australia to 0.86 in Canada to 1.23 in Austria. At least some of this disparity appears to reflect international differences in labor market experience. Blackburn and Ravn (1992) and Burdett and Wright (1989) both note that fluctuations in total hours worked in the United States are largely the result of movements in employment, while in the United Kingdom changes in hours per worker are more important. We note

Table 11.1
Properties of Business Cycles in OECD Economies

	Standard Deviation (%)		Ratio of Standard Deviation to That of y					Autocorr.	Correlation with Output					
Country	y	nx	c	x	g	n	z	y	c	x	g	nx	n	z
Australia	1.45	1.23	0.66	2.78	1.28	0.34	1.00	.60	.46	.67	.15	−.01	.12	.98
Austria	1.28	1.15	1.14	2.92	0.36	1.23	0.84	.57	.65	.75	−.24	−.46	.58	.65
Canada	1.50	0.78	0.85	2.80	0.77	0.86	0.74	.79	.83	.52	−.23	−.26	.69	.84
France	0.90	0.82	0.99	2.96	0.71	0.55	0.76	.78	.61	.79	.25	−.30	.77	.96
Germany	1.51	0.79	0.90	2.93	0.81	0.61	0.83	.65	.66	.84	.26	−.11	.59	.93
italy	1.69	1.33	0.78	1.95	0.42	0.44	0.92	.85	.82	.86	.01	−.68	.42	.96
japan	1.35	0.93	1.09	2.41	0.79	0.36	0.88	.80	.80	.90	−.02	−.22	.60	.98
Switzerland	1.92	1.32	0.74	2.30	0.53	0.71	0.67	.90	.81	.82	.27	−.68	.84	.93
United Kingdom	1.61	1.19	1.15	2.29	0.69	0.68	0.88	.63	.74	.59	.05	−.19	.47	.90
United States	1.92	0.52	0.75	3.27	0.75	0.61	0.68	.86	.82	.94	.12	−.37	.88	.96
Europe	1.01	0.50	0.83	2.09	0.47	0.85	0.98	.75	.81	.89	.10	−.25	.32	.85

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables are y, real output; c, real consumption; x, real fixed investment; g real government purchases; nx, ratio of net exports to output, both at current prices; n, civilian employment; z, Solow residual, defined in text. Except for the ratio of net exports to output, statistics refer to logarithms of variables. Data are quarterly from the OECD's *Quarterly National Accounts*, except employment, which is from the OECD's *Main Economic Indicators*. The sample period is 1970:I to 1990:II.

that employment has been procyclical in all ten countries, but the magnitude of the correlation with output varies substantially across countries.

The last variable in Table 11.1 is the Solow residual, z, which we refer to as productivity. The Solow residual is defined implicitly in the Cobb-Douglas production function,

$$y = zk^{\theta}n^{1-\theta},$$

where y is real output, k is the stock of physical capital, and n is employment. This allows us to compute the Solow residual in logarithms by

$$\log z = \log y - [\theta \log k + (1-\theta) \log n].$$

We set the parameter θ equal to 0.36, as explained in the next section. Since comparable capital stock data are not available on a quarterly basis, we omit the capital part of the expression. This is probably not a serious problem, since the capital stock contributes very little to the cyclical fluctuations of output (see, e.g., Kydland and Prescott 1982, tab. IV). Productivity, by this measure, is strongly procyclical. Its volatility is generally less than that of output.

Two exceptions to this tendency for aggregate variables to move procyclically are government purchases and net exports. Government purchases are procyclical in seven countries and countercyclical in three, but the correlations are small in all cases. The ratio of net exports to output, on the other hand, has been countercyclical in all ten countries, although both its standard deviation and its correlation with output vary substantially across countries.

In Table 11.2 we report statistics with more of an international flavor. In the first column we list the correlation of output fluctuations between each country and the United States. These vary in magnitude but are all positive. The largest is 0.76 for Canada. The correlations for Japan and the major European countries lie between 0.4 and 0.7. Table 11.2 also includes correlations of consumption, investment, government purchases, employment, and Solow residuals across countries. With respect to consumption, we find that the correlations are smaller than those of output for every country, but the difference is large only for Australia. The consumption correlation between the United States and the European aggregate, for example, is 0.51, while the output correlation is 0.66. The correlations of investment, employment, and productivity are also positive in most cases. We find that Solow residuals are generally less highly correlated across countries than output. In our data the differences are generally small. Finally, the cross-country correlations of government purchases vary in sign but are generally small.

We summarize briefly. Despite some heterogeneity in international business cycle experience across the major industrialized countries over the last twenty years, most of the regularities emphasized in Kydland and Prescott's (1982) closed-economy study stand up. More interesting from our point of view are statistics that capture comovements across countries. One is of particular interest to us: the correlations of output across countries are larger than those of consumption and

Table 11.2
International Comovements in OECD Economies

	Correlation of Each Country's Variable with Same U.S. Variable					
Country	y	c	x	g	n	z
Australia	.51	−.19	.16	.23	−.18	.52
Austria	.38	.23	.46	.29	.47	.17
Canada	.76	.49	−.01	−.01	.53	.75
France	.41	.39	.22	−.20	.26	.39
Germany	.69	.49	.55	.28	.52	.65
Italy	.41	.02	.31	.09	−.01	.35
Japan	.60	.44	.56	.11	.32	.58
Switzerland	.42	.40	.38	.01	.36	.43
United Kingdom	.55	.42	.40	−.04	.69	.35
Europe	.66	.51	.53	.18	.33	.56

Notes: See Table 11.1.

productivity. The question for the next section is how these properties compare to those of a theoretical world economy.

3. A Theoretical Business Cycle Model

In our first theoretical economy, agents in two countries produce a single homogeneous good. The structure is a streamlined version of the model of Backus, Kehoe, and Kydland (1992) in which we have eliminated inventory accumulation and leisure durability, which in turn is a two-country extension of Kydland and Prescott's (1982) closed-economy real business cycle model.

In this economy each country is represented by a single agent. The preferences of the representative consumer in country i, for $i = 1, 2$, are characterized by an expected utility function of the form

$$u_i = E_0 \sum_{t=0}^{\infty} \beta^t U(c_{it}, 1 - n_{it}),$$

where c_{it} and n_{it} are consumption and employment in country i and $U(c, 1 - n) = [c^{\mu}(1 - n)^{1-\mu}]^{1-\gamma}/(1 - \gamma)$.

Production of the good takes place in each country using inputs of capital, k, and domestic labor, n, and is influenced by the technology shocks, z. Output, or

GDP, in country i is

$$y_{it} = z_{it} F(k_{it}, n_{it}),$$

where $F(k, n) = k^{\theta} n^{1-\theta}$, the same relation we used to construct Solow residuals in the last section. Since the two countries produce the same good, the world resource constraint for the good is

$$\sum_i (c_{it} + x_{it} + g_{it}) = \sum_i z_{it} F(k_{it}, n_{it}),$$

where x_{it} is the amount of the good allocated to fixed capital formation and g_{it} is government purchases, both for country i. The trade balance, or net exports, in country i is then $nx_{it} = y_{it} - (c_{it} + x_{it} + g_{it})$, the difference between goods produced and goods used.

Capital formation incorporates the time-to-build structure emphasized by Kydland and Prescott (1982). Additions to the stock of fixed capital require inputs of the produced good for J periods, or

$$k_{it+1} = (1 - \delta) k_{it} + s_{it}^{1},$$

and

$$s_{it+1}^{j} = s_{it}^{1}, \quad \text{for } j = 1, \ldots, J - 1,$$

where δ is the depreciation rate and s_{it}^{j} is the number of investment projects in country i at date t that are j periods from completion. We denote by ϕ_j, for $j = 1, \ldots, J$, the fraction of value added to an investment project in the jth period before completion. We set $\phi_j = 1/J$, so that an investment project adding one unit to the capital stock at date $t + 1$ requires expenditures of $1/J$ for the J periods prior to $t + 1$. Fixed investment at date t is

$$x_{it} = \sum_{j=1}^{J} \phi_j s_{it}^{j},$$

the sum of investment expenditures on all existing projects.

The vectors $z_t = (z_{1t}, z_{2t})$ and $g_t = (g_{1t}, g_{2t})$ are stochastic shocks to productivity and government purchases, respectively, which we model as independent bivariate autoregressions. The technology shocks follow

$$z_{t+1} = A z_t + \epsilon_{t+1}^{z},$$

where $\epsilon^{z} = (\epsilon_1^{z}, \epsilon_2^{z})$ is distributed normally and independently over time with variance V_z. The correlation between the technology shocks, z_1 and z_2, is determined by the off-diagonal elements of A and V_z. Similarly, shocks to government purchases follow

$$g_{t+1} = B g_t + \epsilon_{t+1}^{g},$$

where $\epsilon^g = (\epsilon_1^g, \epsilon_2^g)$ is distributed normally with variance V_g. Technology shocks, z, and government spending shocks, g, are independent.

It is straightforward, but notationally burdensome, to define a competitive equilibrium for this economy with complete contingent claims markets. In it consumers use these contingent claims markets to diversify country-specific risk across states of nature. By so doing, consumers end up equating the marginal utility of consumption across countries. Such an equilibrium is, of course, Pareto optimal and we can characterize the equilibrium allocations by exploiting this feature. We compute, in particular, the equilibrium associated with the optimum problem: maximize $u_1 + u_2$ subject to the technology and the resource constraint. In this optimum problem the marginal utility of consumption is also equated across countries for each state of nature, and thus country-specific risk is optimally diversified. We approximate this problem with one that has a quadratic objective function and linear constraints. Details of this procedure are described in Backus, Kehoe, and Kydland (1992, sec. II).

Quantitative properties of this theoretical economy depend to a large extent on the values of the model's parameters. Our benchmark parameter values for this economy are listed in Table 11.3. With the exception of the parameters of the shocks to productivity and government spending, they are taken from Kydland and Prescott's (1982) closed-economy study. The parameters of the technology process are based on Solow residuals for the United States and an aggregate of European countries, as described in our earlier paper (Backus, Kehoe, and Kydland 1992, sec. III). They imply that the productivity shocks are persistent and positively correlated across countries. For the time being, we set $g_t = 0$, thereby eliminating government purchases from the model.

Table 11.3
Benchmark Parameter Values

Preferences	Discount factor, $\beta = 0.99$
	Consumption share, $\mu = 0.34$
	Curvature parameter, $\gamma = 2.0$
Technology	Capital share, $\theta = 0.36$
	Depreciation rate, $\delta = 0.025$
	Time-to-build, $J = 4$
Forcing Processes	Technology shocks, $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{12} & a_{11} \end{bmatrix} = \begin{bmatrix} 0.906 & 0.088 \\ 0.088 & 0.906 \end{bmatrix}$
	$\text{var}\,\epsilon_1^z = \text{var}\,\epsilon_2^z = 0.00852^2$,
	$\text{corr}(\epsilon_1^z, \epsilon_2^z) = 0.258$
	Government spending, $g_t = 0$

Table 11.4
Business Cycles in Theoretical Economies

A. Business Cycle Properties

Economy	Standard Deviation (%) y	nx	Ratio of Standard Deviation to That of y: c	x	n	z	Autocorr. y	Correlation with Output: c	x	nx	n	z
U.S. data	1.92	0.52	.75	3.27	.61	.68	.86	.82	.94	−.37	.88	.96
Benchmark	1.50	3.77	.42	10.99	.50	.67	.62	.77	.27	.01	.93	.89
Transport cost	1.35	0.37	.47	2.91	.47	.75	.61	.81	.92	.23	.92	.98
Autarky	1.26		.54	2.65	.91	.99	.62	.90	.96		.91	.99

B. International Comovements

Economy	Correlation of Foreign and Domestic Variables: y	c	x	n	z
U.S. data	.66	.51	.53	.33	.56
Benchmark	−.21	.88	−.94	−.94	.25
Transport cost	−.05	.89	−.48	−.48	.25
Autarky	.08	.56	−.31	−.31	.25

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables are defined in notes to Table 11.1. Entries are averages over twenty simulations of length 100. The data column refers to the United States in part A and to correlation between the United States and Europe in part B.

Properties of this theoretical world economy are reported in Table 11.4. The entries are means of various statistics across 20 stochastic simulations of the economy, each for 100 periods. As with the data, the statistics refer to Hodrick-Prescott–filtered variables.

We find, first, that the variability of output in this economy is somewhat less than we see in U.S. data, but larger than that of Europe in aggregate, as well as the component European countries. The differences between theory and data in this respect are not large compared to the differences among countries. The behavior of some of the output components, however, differs substantially from the data. The variability of consumption relative to output is smaller in the model economy than in U.S. data when durables are included (0.40 versus 0.75). Since the model disregards durability, a comparison with the volatility of U.S. nondurables consumption may be more appropriate. In this case most of the discrepancy disappears (the volatility of U.S. nondurables consumption is 0.52). Investment, on the other hand, is more than three times more variable relative to output than we see in U.S. data (10.99 versus 3.28). The standard deviation of net exports is about seven times larger than in U.S. data and much larger than for any country in Table 11.1. The net exports component is essentially uncorrelated with output (with a contemporaneous correlation of 0.01), and not countercyclical as it is in the countries in Table 11.1.

We can get some intuition for these properties of the model by examining the dynamic responses pictured in Figures 11.1 and 11.2. These figures illustrate the responses in the benchmark economy to a one-time increase of one standard deviation in the home country's technology innovation, ϵ_1, starting from the steady state. In these figures, productivity is measured as a percentage of its steady-state value; the remaining variables are measured as percentages of steady-state output. Figure 11.1 shows what happens in the home country. There, the technology innovation is followed by a rise in productivity, which subsequently slowly decays. The increase in productivity is associated with increases in domestic investment, consumption, and output. The movement in investment is by far the largest, and it leads to a deficit in net exports.

In Figure 11.2, we see that the innovation to domestic productivity leads eventually, through the technology spillover, to a rise in foreign productivity. Despite this, foreign output and investment both fall initially. Roughly speaking, resources are shifted to the more productive location, the home country. This happens both with capital, as investment rises in the home country and falls abroad, and with labor, which follows a similar pattern. This tendency to "make hay where the sun shines" means that with uncorrelated productivity shocks, consumption will be positively correlated across countries while investment, employment, and output will be negatively correlated. With productivity shocks that are positively correlated, as they are in our model, all of these correlations rise, but with the benchmark parameter values none change sign. This helps to explain why the correlations of foreign and domestic output, employment, and investment

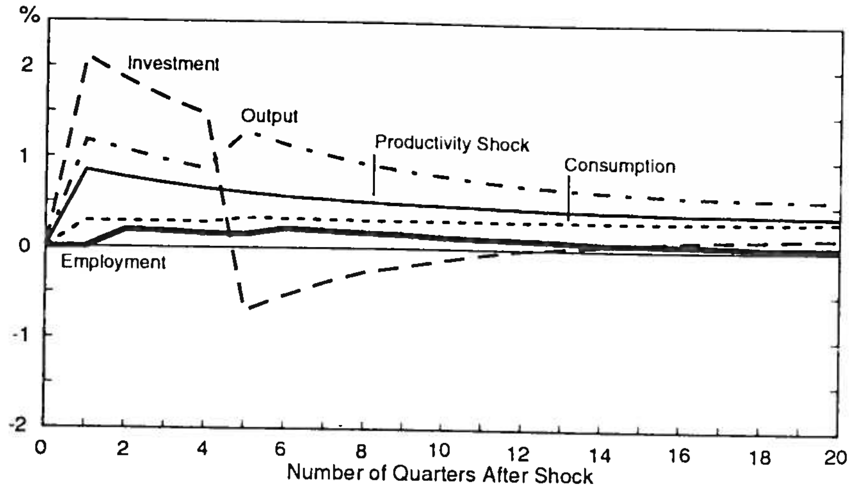

Figure 11.1 Effects of a Productivity Shock in the Benchmark Model on the Home Country

Figures 11.1 and 11.2 show percentage changes after a one-standard-deviation innovation in the home country's productivity shock. Change in productivity shock is measured as percentage of its steady-state value; change in other variables is measured as percentage of steady-state value of output.

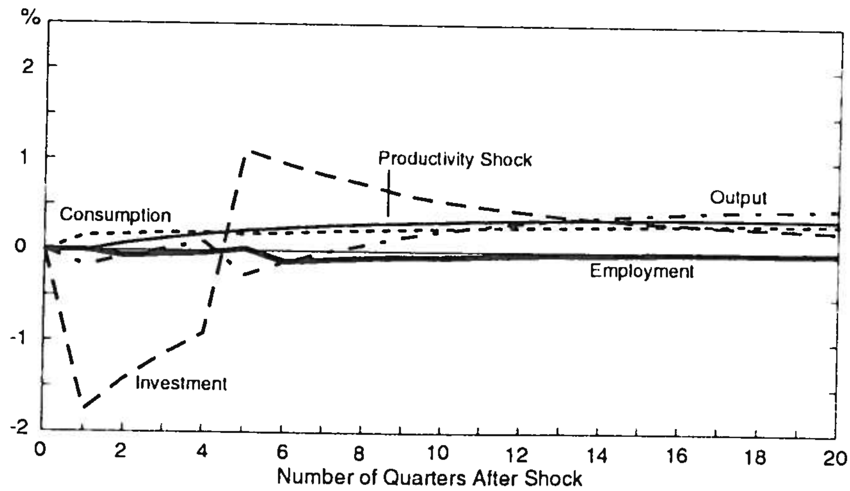

Figure 11.2 Effects of a Productivity Shock in the Benchmark Model on the Foreign Country

are negative, and why the output correlation is smaller than the productivity correlation.

The benchmark economy, then, differs from postwar international data in several respects. In the model, investment and net exports are more variable, whereas consumption is more highly correlated across countries, and output is less highly correlated. Our intuition is that the volatility of investment and net exports reflects the ability of agents in the model to shift perfectly substitutable goods costlessly between countries and to trade in complete markets for state-contingent claims. The ability to shift resources allows agents to shift capital and production effort to the country with the higher current technology shock; that movement shows up in the model as excessive variability of investment and negative correlation of output across countries. Consumers' ability to insure themselves against adverse movements in their own technology shocks suggests that the shifting of production will not be reflected in consumption plans.

We therefore investigate frictions in the physical trading process and the market structure. In the experiment labeled *transport cost*, we impose a quadratic cost on goods shipped between countries. The average cost, in equilibrium, is less than 1 percent, so that if one unit of the good is exported from country 1, more than 0.99 units arrive in country 2. As we see in Table 11.4, this cost reduces the variability of net exports substantially: the standard deviation of the ratio of net exports to output falls from 3.77 percent in the benchmark economy to 0.87 percent with transport costs. The transport cost also lowers the standard deviation of investment relative to output by a factor of almost four, from 10.99 to 2.91. Output's correlation across countries rises from −0.21 to −0.05, while consumption's correlation rises from slightly, from 0.88 to 0.89. In short, this type of friction greatly reduces the variability of net exports and investment but has little effect on the difference between the cross-country correlations of output and consumption.

In our next modification of the theory, we consider limitations on agents' ability to share risk across countries. With complete markets, we know that if preferences are additively separable between consumption and leisure, as they would be if we set $\gamma = 1$, then the ability of agents to trade in markets for contingent claims leads to a perfect correlation across countries. The nonseparability lowers this correlation, in our benchmark economy, to 0.88, which is far larger than we saw in Table 11.2. Here, we consider an extreme experiment, labeled *autarky*, in which we eliminate from the model all trade in goods and assets. The only connection between countries in this case is the correlation between technology shocks. We see in Table 11.4 that this reduces the consumption correlation to 0.56, which is only slightly larger than the correlation of 0.51 between the U.S. and Europe. Output, on the other hand, remains much less highly correlated than it is in the data. Even in this extreme experiment, the difference between theory and data is considerable. Our intuition for the large consumption correlation in the benchmark economy was that it reflected agents' ability to share risk internationally. Under autarky, risk sharing is prohibited, yet we still see a positive correlation. This correlation seems to reflect, instead, the operation of the permanent income hypothesis. The foreign agent knows that a rise in productivity in the home country will spill over to the foreign country and raise the foreign agent's own future productivity and income. In anticipation of this, the foreign agent chooses to increase consumption immediately and postpone some investment.

One way to increase the correlation between foreign and domestic output is to make the productivity shocks more highly correlated. In the benchmark economy, the correlation of productivity shocks is 0.23. If we vary the correlation of innovations we can make this correlation as large or as small as we like. In Figure 11.3 we graph the correlations of consumption, output, and productivity for different values of corr$(\epsilon_1^z, \epsilon_2^z)$. We see that as we increase the correlation of the productivity innovations, we raise the correlation of productivity shocks, as well as the correlations of consumption and output. For different values of the correlation of the productivity innovation, the model can replicate either the consumption correlation in the data or the output correlation, but not the two together. In this sense, the discrepancy between theory and data is the relative size of the consumption and output correlations, rather than either one separately. We refer to these differences between cross-country correlations as the consumption/output/productivity anomaly, or the quantity anomaly.[2]

In short, the theoretical economy generates fluctuations that differ sharply in some respects from what we see in the data. The most interesting differences, we think, concern correlations across countries. In contrast to the data, the theory generally produces output fluctuations that are less highly correlated across countries than those of consumption and productivity. We return to this issue later

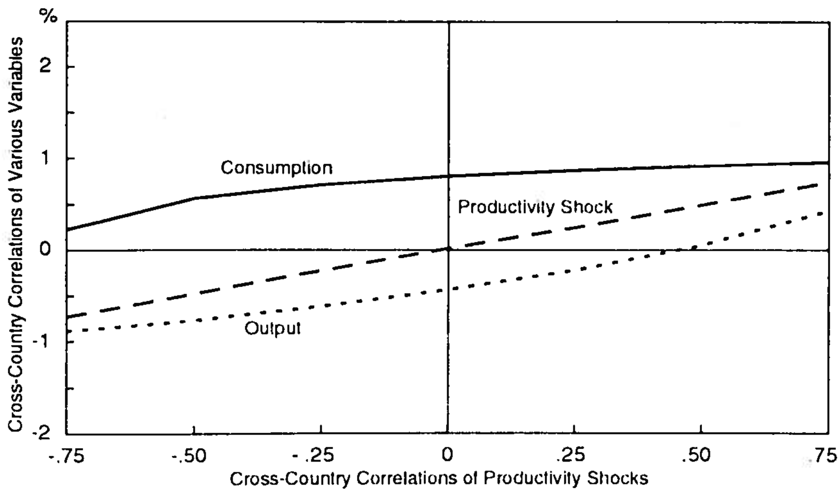

Figure 11.3 How Changes in Cross-Country Correlation of Productivity Shocks Change Cross-Country Correlations of Quantity Variables in the Benchmark Model

in the context of a theoretical economy in which foreign and domestic output are imperfect substitutes. For now we note that these properties are not unique to international economies: similar features should hold in multisector models of closed economies. The tendency for output fluctuations to be less highly correlated than productivity fluctuations, for example, should be more pronounced in a closed-economy where labor is mobile across sectors, yet we know that sectoral outputs are strongly correlated in the data. Similarly, consumption fluctuations should be strongly correlated across regions or individuals. Atkeson and Bayoumi (1991), Crucini (1992), and van Wincoop (1992b) are among those who compare related theories to data for regions within countries. Their work suggests that the one-sector methodology has also masked some interesting features of closed-economy business cycle behavior.

4. Properties of International Relative Prices

We turn now to the behavior of international relative prices, which has been one of the leading issues in international macroeconomics since the collapse of the Bretton Woods system of fixed exchange rates. The terms of trade, labeled p, is the ratio of the implicit price deflators for imports and exports—the relative price of imported goods. This definition is the inverse of the definition used by trade theorists, but corresponds to the convention applied in international macroeconomics to the real exchange rate. The deflators are from the OECD's *Quarterly National Accounts*. The exchange rate, denoted e, is the trade-weighted domestic currency price of one unit of foreign currency: the MERM effective exchange rate in the IMF's *International Financial Statistics*. As in Section 2, we measure the trade balance, labeled nx, as the ratio of net exports to output, with both measured in current prices as reported in the national income and product accounts. Real output, as before, is labeled y. Statistics for p, e, and y refer to logarithms.

We note in Table 11.5 a number of regularities in the behavior of the terms of trade. First, the terms of trade has been highly variable. The standard deviations vary somewhat, but are always greater than those of output (Table 11.1), sometimes by a factor of two or three. A second regularity is the persistence of relative price movements: the terms of trade is highly persistent, with an autocorrelation in the neighborhood of 0.8 for most countries. Finally, we find that the contemporaneous correlation between the terms of trade and net exports is negative in most countries. In France, Italy, Japan, Switzerland, and the United Kingdom the correlations are less than −0.4. The United States is the only country in our table for which these two variables have a sizable positive contemporaneous correlation.

In short, we find a number of regularities in the behavior of net exports and the terms of trade for eleven OECD countries. Prominent among them are the large standard deviations of international relative prices and the high degree of persistence of these variables.

Table 11.5
Properties of the Terms of Trade in OECD Economies

Country	SD $p(\%)$	Autocorr. p	Correlation of: (p, nx)	(p, y)
Australia	5.78	.82	−.10	−.27
Austria	1.73	.46	−.24	.04
Canada	2.99	.85	.05	−.05
France	3.52	.75	−.50	−.13
Germany	2.66	.85	−.08	−.11
Italy	3.50	.78	−.66	.38
Japan	7.24	.86	−.56	−.22
Switzerland	2.85	.88	−.61	.41
United Kingdom	3.14	.80	−.58	.09
Europe	3.68	.83	.30	−.20

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables are p, terms of trade, relative price of imports to exports; y, real output; nx, ratio of net exports to output, both at current prices. Except for the ratio of net exports to output, statistics refer to logarithms of variables. Most variables are from the OECD's *Quarterly National Accounts*. The sample period is 1970:I to 1990:II.

5. Relative Prices in a Theoretical World Economy

A theory of relative price movements of foreign and domestic goods requires, obviously, that they be different commodities. Accordingly, we modify the economy of Section 3 so that the two countries produce different, imperfectly substitutable goods. As in Section 3, the preferences of the representative agent in each country i are characterized by an expected utility function of the form

$$u_i = E_0 \sum_{t=0}^{\infty} \beta^t U(c_{it}, 1 - n_{it}),$$

where c_{it} and n_{it} are consumption and hours worked in country i and

$$U(c, 1 - n) = [c^{\mu}(1 - n)^{1-\mu}]^{1-\gamma}/(1 - \gamma).$$

The technology changes as follows. Each country specializes in the production of a single good, labeled a for country 1 and b for country 2. Each good is produced using capital, k, and labor, n, with linear homogeneous production functions of the same form. This gives rise to the resource constraints,

$$a_{1t} + a_{2t} = y_{1t} = z_{1t} F(k_{1t}, n_{1t}),$$

and

$$b_{1t} + b_{2t} = y_{2t} = z_{2t}F(k_{2t}, n_{2t}),$$

in countries 1 and 2, respectively, where $F(k, n) = k^{\theta}n^{1-\theta}$. The quantity y_{it} denotes GDP in country i, measured in units of the local good, and a_{it} and b_{it} denote uses of the two goods in country i.

Consumption, investment, and government spending in each country are composites of the foreign and domestic goods, with

$$c_{1t} + x_{1t} + g_{1t} = G(a_{1t}, b_{1t}),$$

and

$$c_{2t} + x_{2t} + g_{2t} = G(b_{2t}, a_{2t}),$$

where $G(a, b) = [\omega a^{1-\alpha} + b^{1-\alpha}]^{1/(1-\alpha)}$. The parameters α and ω are both positive, and the elasticity of substitution between foreign and domestic goods is $\sigma = 1/\alpha$. This method of treating foreign and domestic goods, widely used in computable static general equilibrium trade models, is due to Armington (1969) and the resulting function, G, is called the Armington aggregator.

We simplify the capital formation process by setting the time-to-build parameter, J, equal to one. The capital stocks then evolve according to

$$k_{it+1} = (1 - \delta)k_{it} + x_{it},$$

where δ is the depreciation rate.

To develop some intuition for this economy think of good a as aluminum and good b as bricks. Thus, country 1 specializes in making aluminum using capital and domestic labor, while country 2 specializes in making bricks using capital and domestic labor. Country 1 keeps a_1 units of aluminum for domestic use and exports the rest, namely a_2. It then imports b_1 units of bricks from country 2 and combines the bricks and aluminum to make $G(a_1, b_1)$ units of country 1 goods. One can think of G as a function that simply transforms the aluminum and bricks into a country-1–specific good, which is then used for consumption, investment, and government spending in country 1. Likewise, country 2 imports a_2 units of aluminum from country 1 and combines them with b_2 units of bricks, which it produced to make $G(b_2, a_2)$ units of country-2–specific goods. These goods are used for consumption, investment, and government spending in country 2.

As before, we compute equilibrium quantities by finding an optimal allocation. If q_{1t} and q_{2t} are the prices of the domestic and foreign goods, respectively, then the terms of trade is $p_t = q_{2t}/q_{1t}$. In equilibrium, this relative price can be computed from the marginal rate of substitution in the Armington aggregator.

$$\begin{aligned} p_t = q_{2t}/q_{1t} &= \{\partial G(a_{1t}, b_{1t})/\partial b_{1t}\}/\{\partial G(a_{1t}, b_{1t})/\partial a_{1t}\} \\ &= \omega^{-1}(a_{1t}/b_{1t})^{1-\sigma}. \end{aligned} \tag{1}$$

evaluated at equilibrium quantities. The trade balance of country 1, expressed in units of the domestic good, is

$$nx_{1t} = (a_{2t} - p_t b_{1t}).$$

Properties of this variable in Tables 11.6 through 11.8 refer to the ratio of net exports, nx_{1t}, to domestic output, y_{1t}.

With these elements and some parameter values, we can approach the behavior of the terms of trade. Relative to Table 11.3, our benchmark parameter set includes $J = 1$ and the parameters of the Armington aggregator: the elasticity of substitution, σ, which we set equal to 1.5, and the steady-state ratio of imports to GDP, which we set equal to 0.15 by choosing ω appropriately. In this benchmark version of the economy, foreign and domestic goods are better substitutes for each other than they would be with Cobb-Douglas preferences. Our choice of σ is consistent with a large number of studies, as documented by Whalley (1985, ch. 4). The import share is slightly larger than we see in the United States, Japan, or an aggregate of European countries (with intra-European trade netted out).

A number of properties of the theoretical economy with alternative parameter settings are reported in Table 11.6. Consider, first, the autocorrelation of the terms of trade. The autocorrelation for our benchmark parameter values is identical to that in U.S. data in Table 11.5: 0.83. This property is not especially surprising: the variables of the model, including the terms of trade, inherit the high degree of persistence observed in technology shocks in the data and incorporated into our technology shock process.

A second property of the model is the contemporaneous correlation between net exports and the terms of trade. Recall that in the data this correlation is generally

Table 11.6
Properties of the Terms of Trade in Theoretical Economics

Country	*SD* p(%)	*Autocorr.* p	*Correlation of:* (p, nx)	(p, y)
U.S. data	3.68	.83	.30	−.20
Benchmark	0.48	.83	−.41	.49
Two shocks (technology and government spending)	0.57	.67	−.05	.39
Large import share	0.66	.83	−.41	.55
Small elasticity	0.76	.77	−.80	.51

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables are defined in the notes to Table 11.5. Entries are averages over 20 simulations of length 100. The data column refers to the United States.

negative (see Table 11.5). In the theoretical economy, we find, for the benchmark parameter values, that the correlation is -0.41. This number is in the middle of the range observed across the countries in our sample.

Finally, consider the standard deviation of the terms of trade. With our benchmark parameter values the standard deviation is 0.48 percent, which is a factor of more than seven less than we see for the United States in Table 11.5. This large difference between the standard deviation in the model and the one in the data is our second anomaly: the terms of trade, or price-variability, anomaly.

Like the consumption/output/productivity anomaly, the price-variability anomaly is robust to reasonable changes in parameter values. We add government spending shocks in the experiment labeled *two shocks*. In this experiment, we calibrate the government spending process to U.S. data: the mean value of g in each country is 20 percent of steady-state output, $B = \text{diag}(0.95, 0.95)$, and the innovations are assigned standard deviations equal to 2 percent of mean government purchases, or 0.004. These shocks are independent across countries and of the productivity shocks, as they tend to be in international data (see Table 11.2). With these shocks added to the model, the standard deviation rises from 0.48 to 0.57, which remains far below what we see in the data. In another experiment, labeled *large share*, we raise the average share of imported goods to GDP from 0.15 to 0.25. In this case, the standard deviation of the terms of trade rises to 0.59. Nevertheless, the variability of the terms of trade in the model remains well below what we see in the data.

The variability of the terms of trade is also influenced by the elasticity of substitution between foreign and domestic goods, $\sigma = 1/\alpha$ in the Armington aggregator. In the *small elasticity* experiment we lower σ from 1.5 to 0.5; the standard deviation of the terms of trade rises from 0.48 percent in the benchmark economy to 0.76. In the theory, prices are related to quantities by the first-order condition,

$$\log p_t = -\log \omega - \sigma^{-1} \log(b_{1t}/a_{1t}), \tag{2}$$

where b_1 is imports and a_1 is output minus exports in country 1. Given a fixed amount of variability in the import ratio, b_1/a_1, we can increase the variability of p without bound by lowering the value of σ. In Figure 11.4 we see that as σ approaches zero the standard deviation of the terms of trade approaches values similar to those we see in the data. Closer inspection suggests, however, that raising the complementarity between foreign and domestic goods does not resolve the anomaly. The problem is that the variability of the import ratio in the data is not much different from that of the terms of trade. Thus choosing a small value of σ "resolves" the price variability anomaly only by making the variability of b_1/a_1 much smaller than it is in the data. Given the first-order condition, (2), it is impossible to separate the problem of insufficient variability of the price, p, from that of insufficient variability of the quantity ratio, b_1/a_1.

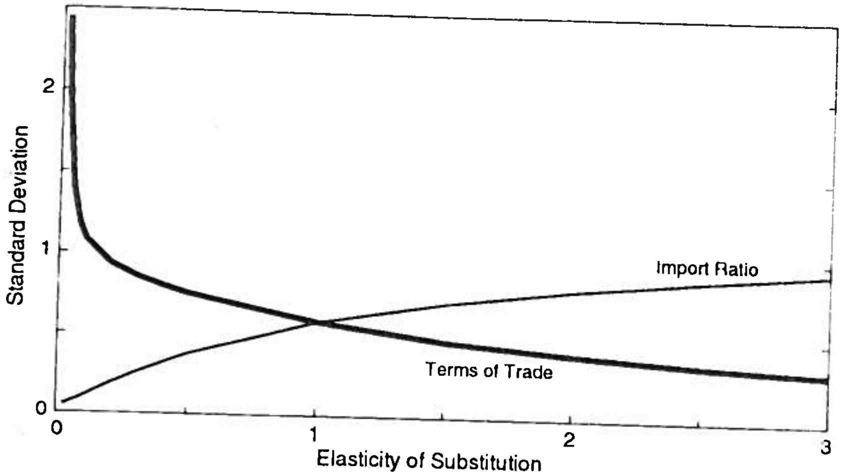

Figure 11.4 Effects of Varying the Elasticity of Substitution between Foreign and Domestic Goods on the Volatility of the Terms of Trade and the Import Ratio

Mussa (1986) adds another wrinkle to this puzzle. He argues persuasively that an important ingredient in the price variability puzzle is the sharp difference in price behavior between fixed and floating exchange rate regimes. As he shows, and we report in Table 11.7, the variability of the terms of trade has been much higher in the post–Bretton Woods period than before. By our estimates, the standard deviation of the terms of trade is higher by a factor of about three in the major countries for which we have long data series available. Mussa (1986) also notes that there has been greater price variability in other periods of floating exchange rates (for example, in Canada between 1952 and 1962), so the distinction between fixed and floating rate regimes is not simply one of time period. In our theory, and others in which there is a similar first-order condition relating prices and quantities, the standard deviation of the terms of trade is directly related to quantity variability: if the standard deviation of the import ratio doubles, then the standard deviation of the terms of trade also doubles. With this in mind, we note that while there has been greater quantity variability in most countries (Japan is an exception) in the post–Bretton Woods period, the increase has been much smaller than that for the terms of trade. The issue, then, is how to account for the sharp increase in price variability without generating a similar increase in the variability of quantities. At the very least, one must abandon the tight connection between prices and quantities implied by first-order conditions such as (1).

Finally, we return briefly to the consumption/output/productivity anomaly of Sections 2 and 3. We have noted that complementarity between foreign and domestic goods influences the variability of the terms of trade. It also influences the model's business cycle properties. As we see in Table 11.8 and Figure 11.5, the

Table 11.7
Properties of Four Economies by Subperiod

		Standard Deviation (%)						Correlation with Same U.S. Variable	
Country	*Period*	p	y	nx	im	ex	$im/(y-e)$	y	c
Canada	1955–90	2.44	1.48	0.79	5.25	5.52	3.85	.71	.52
	1955–71	1.19	1.38	0.78	4.83	2.89	4.13	.53	.59
	1972–90	3.05	1.54	0.79	5.44	4.64	4.75	.79	.48
Japan	1955–90	5.69	1.61	1.01	6.64	4.50	6.29	.20	.27
	1955–71	2.29	1.93	1.06	7.54	3.74	7.01	−.07	−.02
	1972–90	7.12	1.19	0.92	5.87	4.91	5.63	.57	.36
United Kingdom	1955–90	2.64	1.48	1.07	3.85	3.15	3.50	.46	.35
	1955–71	1.45	1.25	0.74	3.04	2.85	2.53	.15	.05
	1972–90	3.05	1.67	1.22	4.34	3.36	4.16	.57	.35
United States	1955–90	2.92	1.70	0.45	4.90	5.52	3.85		
	1955–71	1.26	1.23	0.32	3.38	5.23	3.16		
	1972–90	3.79	1.94	0.54	5.88	5.61	4.38		

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables include im, real imports; and ex, real exports. Other variables are defined in notes to Tables 11.1 and 11.5. Except for nx, statistics refer to logarithms of variables.

Table 11.8
Business Cycles and International Comovements in Theoretical Economies

A. Business Cycle Properties

	Standard Deviation (%)		Ratio of Standard Deviation to That of y					Autocorr.	Correlation with Output					
Economy	*y*	*nx*	*c*	*x*	*g*	*n*	*z*	*y*	*c*	*x*	*g*	*nx*	*n*	*z*
U.S. data	1.92	.52	.75	3.27	0.75	.61	.68	.86	.82	.94	.12	−.37	.88	.96
Benchmark	1.38	.30	.47	3.48	0.00	.42	.75	.63	.88	.93	.00	−.64	.94	.99
Two shocks	1.33	.33	.62	4.29	1.95	.42	.78	.65	.78	.89	.16	−.57	.79	.97
Large import share	1.36	.85	.47	4.70	0.00	.38	.76	.64	.86	.84	.00	−.59	.95	.99
Small elasticity	1.33	.37	.50	3.41	0.00	.37	.77	.63	.92	.93	.00	−.66	.94	.96

B. International Comovements

	Correlation of Foreign and Domestic Variables					
Economy	*y*	*c*	*x*	*g*	*n*	*z*
U.S. vs. Europe Data	.66	.51	.53	.18	.33	.56
Benchmark	.02	.77	−.58	.00	−.54	.24
Two shocks	.00	.83	−.64	−.02	−.62	.32
Large import share	.05	.83	−.76	.00	−.42	.24
Small elasticity	.10	.68	−.54	.00	−.36	.24

Notes: Statistics are based on Hodrick-Prescott–filtered data. Variables are defined in notes to Table 11.1. Entries are averages over 20 simulations of length 100. The data column refers to the United States in part A and to correlation between the United States and Europe in part B.

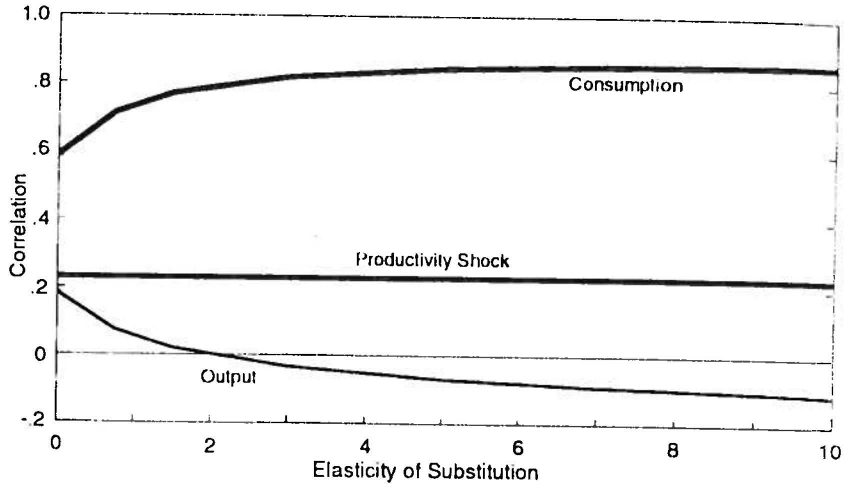

Figure 11.5 Effects of Varying the Elasticity of Substitution between Foreign and Domestic Goods on the Cross-Country Correlations of Selected Variables

correlation between consumption in the two countries of our theoretical economy falls as we reduce the elasticity of substitution between foreign and domestic goods. At the same time, the correlation between foreign and domestic output rises. Nevertheless, for values of σ above 0.025 (the smallest value we've been able to use) the consumption correlation exceeds the output correlation. The productivity correlation, of course, is not affected by our choice of σ: it equals 0.23 throughout. Thus for reasonable values of σ, there remains a substantial difference between the cross-country correlations of output, consumption, and productivity in the theory and those in the data. Imperfect substitutability between goods does not appear to resolve the consumption/output/productivity anomaly documented earlier.

In short, we find that we must add relative price variability to our list of anomalies. An interesting wrinkle to this finding is that the anomalous behavior of the relative price is closely connected, in our theory, to anomalous behavior of quantities.

6. Related Work and New Directions

We have documented two striking differences between theory and data, which we label the consumption/output/productivity and price variability anomalies. Our review of these issues has focused on our own work, but international macroeconomics has been one of the most active areas of business cycle research and includes studies that go far beyond the theoretical economies of Sections 3 and 5. Although these studies have addressed a wide range of issues, we find it useful to review them from the perspective of the two anomalies. We start by listing some of the prominent theoretical innovations, and then go on to consider their possible roles in accounting for the character of aggregate fluctuations and relative price movements.

Recent studies in international business cycle research have extended the theory in at least five directions. One of the more popular extensions has been to introduce nontraded goods. We are often reminded that haircuts and other services cannot be traded across cities, much less across countries, so this approach has some natural appeal. A second extension adds other sources of shocks including oil shocks and taste shocks. A third popular extension of the theory introduces restrictions on asset trade so that agents' ability to hedge risk is more limited than in our complete market economies. A fourth extension of the theory adds money to economies that are otherwise much like those we studied in Sections 3 and 5. A final extension introduces imperfect competition.

Adding nontraded goods does not seem to help explain either anomaly. Consider the consumption/output/productivity anomaly. Nontraded goods can, in principle, lower the cross-country consumption correlations, since the correlations between the nontraded components of consumption are not directly connected by trade in goods. They may, in addition, lower the correlation of the consumption of traded goods if the utility function is nonseparable between traded and nontraded goods consumption, as it is in Ravn (1992) and Stockman and Tesar (1991). The effect is similar to that of leisure in our models when utility is not additively separable between consumption and leisure. In both our work (Section 3) and in Ravn and Stockman and Tesar the effect of the nonseparability is quantitatively small. In Stockman and Tesar, the result of nontraded goods is that traded goods consumption, rather than total consumption, is more highly correlated across countries in the models than in the data. The anomaly, in other words, is simply pushed onto the traded component of consumption. Backus and Smith (1993) note, as well, that these models imply close connections between consumption differentials and relative prices that are not observed in aggregate data.

One of the by-products of this work has been a reconsideration of the impulses generating fluctuations. Costello and Praschnik (1992) introduce oil shocks, which increase the variability of the terms of trade in oil-importing countries and lowers the correlation of consumption across countries. They find, however, that the terms of trade for manufactured goods remains less variable in the model than we see in the data and that the cross-country correlation of manufactured goods consumption is much higher than in the data. Stockman and Tesar (1991) suggest shocks to preferences. They add a shock to the first-order condition that links consumption quantities and relative prices. This shock lowers the correlation of aggregate consumption across countries, and of consumption of traded goods alone. It has little effect, however, on the variability of the terms of trade. To date there has been no attempt to quantify such shocks, which makes it difficult to assess the effects of adding them to our models. One step in this direction might be to compute preference shocks as residuals from agents' first-order conditions, much as we compute productivity shocks as residuals from production functions.

Economies with incomplete markets would also seem to have the potential to account for low correlations of consumption across countries. With complete markets, as in the models of Sections 3 and 5, agents use asset markets to equate marginal rates of substitution across dates and states of nature. With separable preferences, this leads, as we have seen, to a perfect correlation of consumption across countries. When agents have limited ability to use international financial markets to share risk, marginal rates of substitution are not equated for all dates and states. One might guess, then, that the consumption correlation would be smaller than with complete markets. Thus Conze, Lasry, and Scheinkman (1991) show that in an economy in which agents can trade a single asset, the consumption correlation falls and the output correlation rises. Nevertheless, they still find that the consumption correlation exceeds the output correlation for most parameter values. Our autarky experiment in Section 3 makes the same point in an economy with even more limited trading opportunities. Kollmann (1990, tab. 1.1.3) studies an economy in which two agents trade a single, risk-free bond. In this economy he finds much smaller consumption correlations than with incomplete markets, but the correlation of investment across countries is sharply negative when productivity shocks are persistent, as they are in the data, and the consumption correlation remains higher than the output correlation. Baxter and Crucini (1992, tab. 4) also consider an economy in which agents trade a single risk-free bond and find that output is more highly correlated across countries than consumption, but the correlations of consumption, investment, and employment are negative. Thus these models have, to some extent, transferred the consumption/output/productivity anomaly onto other variables.

None of these models have accounted for the variability of international relative prices, like the terms of trade. The fourth extension, including money, was done with price variability explicitly in mind. Both Grilli and Roubini (1992) and Schlagenhauf and Wrase (1992) adapt Lucas's (1990c) liquidity model to the open economy. In these economies, asset markets and goods markets are separated for one period, and shocks to the stock of money have a one-period effect on interest rates, currency prices, and relative prices of goods. Thus the theory generates greater variability of relative prices than we would see in an analogous model without the segmented market structure. In its current form, however, this structure generates relative price movements with very little persistence, and thus fails to mimic this important feature of the data. The next step in this line of research is to specify a mechanism to generate the persistence we see in the data.

Another class of monetary models considers labor or goods contracts that fix wages or prices in advance. In closed-economy studies, such as Cho and Cooley (1992), this magnifies the effects of some shocks on employment and output. In open economies, one might guess that it could generate additional relative price variability, particularly if we added segmentation across national markets. This intuition has yet to be tested, but Cho and Roche (1993) and Ohanian and Stockman (1993) have made some progress in developing international business cycle models of this sort.

Imperfect competition is the final extension of the theory, and it might bear on the price variability anomaly. If imperfectly competitive firms sell their output in markets that are internationally segmented, then price discrimination might lead to greater changes in relative prices than we see with perfect competition. Studies of industries by Giovannini (1988) and Lapham (1991) show that this can lead to persistent movements in relative prices across countries, but the theory has yet to be extended to general equilibrium settings at the level of aggregation considered in Section 3 and 5. Perhaps Hornstein's (1993) or Rotemberg and Woodford's (Chapter 9, this volume) general equilibrium treatment of monopolistic competition in a closed economy could be adapted to the open economy.

All of these innovations help bring the quantitative implications of the theory closer to observed properties of international time series data. In our view, they have yet to resolve the two anomalies, but perhaps future efforts along similar lines will be more successful in this regard.

7. Final Thoughts

We have reviewed recent work on international business cycles, emphasizing two striking differences between theory and data. The first we call the consumption/output/productivity anomaly: in the data we generally find that the correlation across countries of output fluctuations is positive, and larger than the analogous consumption and productivity correlations. In theoretical economies we find, for a wide range of parameter values, that the consumption correlation exceeds the productivity and output correlations. The second anomaly concerns relative price movements: the standard deviation of the terms of trade is considerably larger in the data than in theoretical economies.

These anomalies have been met with a large and imaginative body of work in which the dynamic general equilibrium framework has been extended in ways that go well beyond the two-country versions of Kydland and Prescott (1982) that started this line of study. Our guess is that five years from now the models that have been developed will differ from this starting point in fundamental ways.

Notes

We thank Tom Cooley, John Donaldson, Tiff Macklem, Klaus Neusser, Patricia Reynolds, Julio Rotemberg, Gregor Smith, and Kei-Mu Yi for helpful comments on earlier drafts, and the National Science Foundation and the Center for Japan–US Business and Economic Studies NEC faculty fellowship program for financial support. We hope to make our data available shortly in machine-readable format.

1. See King and Rebelo (1993) and Prescott (1986) for descriptions of this filter and its relation to others.

2. Reynolds (1992) argues that our assessment of the theory is unduly pessimistic, in part because uncertainty about the parameter values makes the theory's predictions less precise. In her view, a model with multiple traded goods "is capable of replicating and explaining both the output and consumption correlations" (Reynolds 1992, abstract). Most of her point estimates, however, imply that the output correlations in her theory are smaller than the consumption correlations, and in one case the difference is significant in a statistical sense.