Consumption Risk-Sharing Across G-7 Countries

Giovanni P. Olivei

Economist, Federal Reserve Bank of Boston. Erika M. Dreyer provided outstanding research assistance.

An intensely debated issue in international economics concerns the extent to which investors exploit the gains from international trade in financial assets. Economic theory has long recognized the potential benefits of open capital markets. International financial transactions offer countries a larger pool of resources from which to finance profitable investment opportunities, thereby promoting growth and smoothing the time profile of consumption and investment. Through open capital markets, firms and households can also diversify away idiosyncratic country-specific risks. Still, despite the continuing process of financial integration and globalization, it is unclear whether such benefits are fully exploited in actual practice.

This article focuses on the extent to which investors take advantage of international diversification opportunities. At least two related pieces of evidence suggest that international diversification of risks is far from perfect. To begin with, several studies have documented that domestic investors hold too little of their portfolios in foreign assets. According to the predictions of a standard capital-asset-pricing model, U.S. investors should hold over one-half of their wealth in foreign equities (see Lewis 1995). In actual practice, this figure is less than 10 percent. Additionally, there is little evidence of international consumption risk-sharing. An optimal international portfolio allocation implies that investors perfectly pool all idiosyncratic country-specific risks. With such an allocation, a country's domestic private consumption is affected only by uninsurable global shocks. As a consequence, one should observe consumption growth rates to be highly correlated across countries, except for preference shocks and measurement errors. However, the evidence shows that domestic consumption is correlated with domestic output, implying that country-specific risk is not diversified away through international asset markets.

Of course, the observed lack of international risk-sharing could simply be a reflection of the fact that the potential welfare gains from

diversification are small. Still, a growing body of work suggests that the gains from international risk-sharing, far from being trivial, could amount to several percentage points of a country's annual domestic consumption. If this is the case, understanding the current and potential risk-allocation role of world financial markets appears to be an important task.

In this study we reexamine some of the evidence concerning the degree to which international financial markets help countries diversify away country-specific risks to achieve a mutually preferable allocation of consumption. We do so by looking at national consumption correlations across the G-7 countries. Prior studies have documented a lack of coherence in international consumption fluctuations. However, these studies do not cover the 1990s, a period that appears particularly fruitful for testing the extent of international risk-sharing. Trade and capital flows suggest that the world has become increasingly integrated in recent years. For example, some of the countries we consider (Italy and France) removed all barriers to capital account transactions only during the 1990s. Moreover, in contrast to the previous two decades, business cycle fluctuations during the 1990s in the G-7 countries have been driven, to a sizable extent, by country-specific shocks. It then seems interesting to investigate whether a trend toward increasing globalization and greater incentives to international diversification of risks have been accompanied by an increase in consumption risk-sharing.

Through open capital markets, firms and households can diversify away idiosyncratic country-specific risks.

The rest of the paper is structured as follows. Section I reviews how global consumption allocations behave in an ideal setup with perfect capital mobility and markets for insuring against all idiosyncratic risks. Section II presents evidence on international consumption correlations for the G-7 countries and shows how the correlations have evolved over time. We find that the apparent lack of consumption risk-sharing continued to persist during the 1990s and, somewhat surprisingly, that the puzzle of low international consumption correlations is probably worse than usually thought. In Section III we briefly discuss the potential benefits from international diversification and some of the issues related to improving global risk-sharing. Section IV offers concluding remarks.

I. International Consumption Risk-Sharing with Complete Markets

The empirical evidence on world capital markets' ability to diversify risks is often measured against the predictions of a highly stylized model where countries can hedge all idiosyncratic country-specific shocks. Though clearly counterfactual, the model delivers striking implications for a country's consumption behavior. Specifically, a key prediction of the model is that a country's consumption growth is a linear function of world consumption growth:

$$\log\left(\frac{C_t}{C_{t-1}}\right) = a + \log\left(\frac{C_t^W}{C_{t-1}^W}\right), \qquad (1)$$

where a is a constant, C_t denotes the country's level of per capita consumption, and C_t^W denotes the level of world per capita consumption at time t. Thus, according to equation (1), a country's consumption growth should be perfectly correlated with world consumption growth. While we refer to previous studies (for example, Obstfeld 1994a; Obstfeld and Rogoff 1996, ch. 5) for a detailed derivation of (1), here we focus on the assumptions and the intuition behind such a result.

There are three key assumptions behind equation (1). First, there must be perfect capital mobility across countries, that is, no restrictions on international capital movements and zero transaction costs. Second, asset markets must be complete, so that all idiosyncratic consumption risks are insurable. For this to be the case, agents must trade in world capital markets a set of assets sufficiently large to span all possible future contingencies. For example, a worker must be able to insure his labor income against job loss and against all other contingencies that have a positive probability of occurring. This requires the ability to foresee events with sufficient clarity to provide for them in contracts (see Obstfeld 1995). Finally, in order for international trade in assets that bear claims on a country's domestic output to be feasible, output must be tradable. If a portion of domestic output is nontradable, it can be consumed only by domestic residents. Therefore, since assets that involve claims on nontradable goods can only be traded among domestic resi-

dents, risks involving fluctuations in nontradable output must be borne entirely by domestic consumers and cannot be diversified away in international asset markets.

Equation (1) shows that domestic consumption will vary with the common component of world consumption when these three main assumptions and other auxiliary assumptions hold. Note that if equation (1) holds for each individual country, consumption growth must also be perfectly correlated across countries. The perfect correlation between domestic and world consumption growth rates in this ideal setup is explained by the ability of each country to provide mutual insurance against purely idiosyncratic country-specific risks through international trade in a complete set of securities. Consequently, all idiosyncratic domestic consumption fluctuations can be eliminated, with domestic consumption responding only to uninsurable global shocks. Each country's consumption is then internationally diversified, so that any risk associated with consumption fluctuations is due to systemic output risk, or the risk associated with world output fluctuations. Thus, consumption growth rates will be perfectly correlated across countries even when domestic output growth rates, which are affected by purely idiosyncratic country-specific shocks, do not display a high degree of coherence.

Other things equal, in the presence of increasing capital mobility one should observe an increasing coherence of consumption fluctuations across countries.

A weaker correlation between domestic consumption growth and world consumption growth than implied by equation (1) will apply when the assumptions of perfect capital mobility, market completeness, and/or full output tradability do not hold. The presence of transaction costs can reduce international trade in assets, as can informational asymmetries across countries and differential tax treatment of domestic and foreign investors. When asset markets are incomplete, some idiosyncratic country-specific risks will not be insurable, limiting risk-sharing opportunities. As previously mentioned, the presence of nontraded goods would also temper the perfect correlation outcome of equation (1).

Given that the assumptions of perfect capital mobility, market completeness, and full output tradability are obviously at odds with reality, one would not expect to find perfect coherence in the evolution of per capita consumption across countries in actual data as per equation (1). Still, other things equal, in the presence of increasing international capital mobility one should observe an increasing coherence of consumption fluctuations across countries. In a similar vein, increasing financial sophistication should, other things equal, translate into an increasing tendency toward positively correlated consumption comovements across countries (see Obstfeld 1994a). This is because greater financial sophistication, through the pricing of a larger set of contingencies, contributes to more complete asset markets.

In the next section we examine whether the implications concerning per capita consumption growth fluctuations that stem from equation (1) are supported by actual data. Specifically, we ask whether the coherence in consumption movements across highly industrialized countries has increased over the last decade. As international financial markets are becoming increasingly integrated and sophisticated, opportunities for diversifying away idiosyncratic country-specific shocks should also be increasing. To the extent that consumers are taking advantage of such opportunities, this should translate into higher coherence in international consumption fluctuations. Moreover, as the next section will illustrate, in recent years country-specific shocks appear to have played an important role in driving business-cycle fluctuations. The increasing importance of country-specific shocks might have created greater incentives for international consumption risk-sharing.

II. Output and Consumption Correlations Across G-7 Countries

Before assessing the extent of risk-sharing among G-7 countries, we briefly examine the main features of cyclical variations in output over the past 20 years. While business cycles were closely synchronized over the 1970s and 1980s, in more recent years the coherence of business fluctuations among G-7 countries has become less pronounced. This pattern suggests that while global shocks played a major role in driving

economic fluctuations in the 1970s and 1980s, country-specific shocks have been more important in the last few years (see International Monetary Fund 1998, pp. 55–58).

Figure 1 depicts output gaps, a measure of cyclical fluctuation, for the G-7 countries over the period 1970 to 1998. Panel A shows a fairly high degree of coherence across business cycle fluctuations until the early 1990s for the three largest economies (United States, Japan, and Germany). The fall in economic activity during the 1973–75 period occurred in the context of a sharp rise in oil prices. The downturn of 1979 again took place in the presence of an oil price shock, and the unrelentingly tight monetary stance in the United States thereafter eventually forced other industrialized countries into tight money in order to defend their currency values. The expansion in the 1980s occurred during a period in which several industrial countries took important steps toward liberalizing their financial systems and lifting restrictions on their capital account transactions. Additionally, the decline in oil prices in 1986 contributed to relatively subdued inflation in the presence of strong output growth.

Figure 1

G-7 Output Gaps, as a Percent of Potential GDP

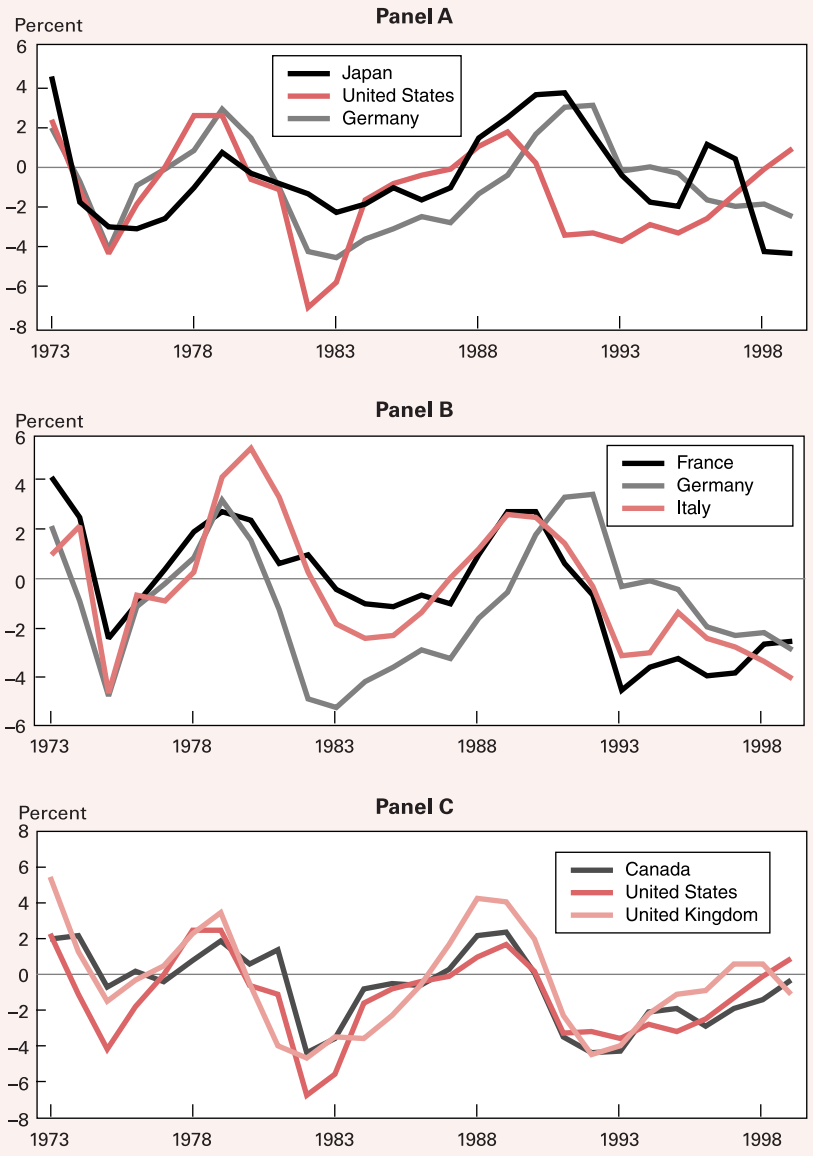

Source: International Monetary Fund, *World Economic Outlook*, September 1999.

Since the early 1990s, country-specific shocks appear to have played a greater role in cyclical fluctuations. The collapse in asset prices and the crisis in the banking sector in Japan contributed to a prolonged slump in that country, with real GDP growth only averaging about 1 percent over the period 1992 to 1999. In continental Europe (panel B of Figure 1), the tight monetary stance necessary to defend currency pegs vis-à-vis the deutsche mark following the German unification and the tight fiscal stance required to meet the Maastricht criteria marked a period of low output

growth that has only recently ended. In contrast, since the slowdown of 1990–1991, the United States has experienced a period of sustained growth in an environment characterized by low inflation and low unemployment. The U.S. experience in the 1990s thus stands in sharp contrast to developments in Japan and continental Europe. Additionally, as panel C of Figure 1 shows, activity over the same period in the United Kingdom has been more closely related to that in the United States than to activity in continental Europe. In the case of Canada, the slowdown in economic activity in the early 1990s was slightly more protracted than in the United States, but since 1993 the coherence between the Canadian and U.S. cycles has been fairly high.

The divergent paths taken by industrialized countries over the 1990s should provide a fruitful ground for testing the extent of international risk-sharing.

The divergent paths taken by industrialized countries over the 1990s should provide a fruitful ground for testing the extent of international risk-sharing. If the shocks that affect a country's level of economic activity are identical to those in all other countries, then there cannot be any welfare improvement from international risk-sharing, because it is impossible to diversify away systemic global risks. Instead, to the extent that country-specific shocks play a larger role than global shocks, the incentives to insure against output fluctuations should increase, because all the idiosyncratic country-specific risk can be diversified away when world asset markets are complete. While the potential gains from international risk-sharing will be discussed in the next section, here we note that welfare improvements can be nontrivial when country-specific output shocks are highly persistent.

We now examine the degree of international risk-sharing among the G-7 countries by assessing the evolution of output and consumption correlations over the period 1973 to 1998. As shown previously, if asset markets are complete and agents fully insure themselves against idiosyncratic risk, a country's consumption should be perfectly correlated with world consumption. In addition, the presence of country-specific shocks implies that per capita domestic output, Y, must be less than perfectly correlated with per capita world output, Y^w. Thus, a crucial prediction of the complete markets model is that, in per capita terms, a country's correlation of output growth with world output growth should be less than the country's correlation of consumption growth with world consumption growth:

$$corr(\Delta y, \Delta y^w) < corr(\Delta c, \Delta c^w),^1 \qquad (2)$$

where $\Delta x = x - x_{-1}$, and lower-case letters denote natural logarithms of the respective upper-case variables.

A closely related implication of the complete markets model is that a country's consumption growth should be more highly correlated with *world* output growth than with *domestic* output growth:

$$corr(\Delta c, \Delta y) < corr(\Delta c, \Delta y^w). \qquad (3)$$

Inequality (3) holds because, with complete insurance, domestic consumption will react only to global shocks, which, of course, are uninsurable and responsible for the fluctuations in world output. However, domestic output fluctuates in response not only to global shocks, but also in response to those country-specific shocks that domestic consumers are able to diversify away. Such a prediction is clearly in sharp contrast with the implications of a simple Keynesian consumption function in which domestic consumption closely tracks domestic output. In that case the correlation of domestic consumption with domestic output would be very high, while the correlation of domestic consumption with world consumption would be lower, especially so when country-specific shocks play a larger role than global shocks.

Several studies have previously examined output and consumption correlations in developing and developed countries (for example, Obstfeld 1994a, 1995).

[1] The definition of the correlation coefficient between two random variables X and Y is as follows:

$$corr(X, Y) = \frac{\text{cov}(X, Y)}{\sigma_X \sigma_Y},$$

where cov(X,Y) denotes the covariance between X and Y, and σ_X is the standard deviation of X. Note that it is always the case that $-1 \leq corr(X, Y) \leq 1$. If there exists a linear relationship X and Y, so that $Y = a + bX$, where $b > 0$, then there is perfect positive correlation between X and Y: $corr(X, Y) = 1$. If $b < 0$, then there is perfect negative correlation between X and Y: $corr(X, Y) = -1$.

These studies show that the correlation between domestic and world consumption growth is generally lower than the correlation between domestic and world output growth, in sharp contrast to the complete markets model prediction of relationship (2). Still, Obstfeld (1994a) notes that among the G-7 countries over the period 1951 to 1988 there has been a marked tendency for domestic consumption to become more closely correlated with world consumption. By comparing the period 1951 to 1972 with the period 1973 to 1988, Obstfeld documents a general rise in the correlations of domestic consumption growth with world consumption growth, most notably in Germany and Japan.

Both trade and capital flows suggest that the world is becoming increasingly integrated. Moreover, the overall trend toward greater financial integration has been accompanied by a tremendous increase in financial sophistication.

It is thus interesting to ask whether the pattern of increasing coherence of domestic consumption with world consumption also applies to the 1990s. After all, both trade and capital flows suggest that the world is becoming increasingly integrated. For some countries, this trend is very new. Italy and France completely removed restrictions to capital account transactions only in the 1990s. Japan and the United Kingdom did not open their capital accounts until the beginning of the 1980s.[2] Moreover, the overall trend toward greater financial integration has been accompanied by a tremendous increase in financial sophistication, with the creation of a vast array of new financial instruments. A global setting with more integrated and sophisticated financial markets should provide increased opportunities for hedging idiosyncratic shocks. To the extent that the model outlined in Section I provides a relevant, albeit stylized, description of consumer behavior, one should observe higher correlations of per capita consumption growth across countries during the most recent years. Furthermore, the increased importance of country-specific shocks in driving recent output fluctuations is likely to have raised the incentives for consumption risk-sharing.

Table 1 reports per capita output and consumption growth correlations with rest-of-world output and consumption, respectively, for the period 1973 to 1998. The frequency of the data is quarterly.[3] Since some of the countries considered constitute a nontrivial fraction of world output, we use rest-of-world instead of world aggregates in order to avoid distorting the sample correlations upward. For each country considered, the rest-of-world measure is given by a population-weighted average of the six remaining G-7 countries. Thus, for our purposes, the "world" consists of the largest industrial nations only. Note that while it is quite plausible that G-7 countries' residents engaged in risk-sharing activities with non-G-7 countries, according to the complete markets model of Section I, movements in a country's per capita consumption growth should be perfectly correlated across *all* countries trading in international capital markets. Accordingly, there is nothing restrictive in testing the hypothesis in (2) with rest-of-world output and consumption consisting of a weighted average of G-7 countries only.[4]

Table 1 provides a picture of the extent of international risk-sharing that is very similar to that reported in earlier studies (for example, Obstfeld 1994a). To begin with, the correlation of a country's consumption growth rate with world consumption growth is far below the value of unity predicted by the complete markets model with perfect capital mobility. In all cases, the hypothesis that $corr(\Delta c, \Delta c^{w}) = 1$ is rejected at standard confidence levels. Given the highly restric-

[2] Canada, Germany, and the United States removed restrictions on capital account transactions in the early 1970s.

[3] The data source is *OECD Quarterly National Accounts*, various issues. Population data are interpolated using annual values from various issues of the *International Financial Statistics*. World aggregates were computed by using 1992 PPP-adjusted exchange rates vis-à-vis the U.S. dollar.

[4] Obstfeld (1994a) notes that examining pairwise correlations between countries instead of correlations between one country and rest-of-world may be misleading because correlations between countries are likely to be downward biased. This will happen if there are significant country-specific preference shocks. Obstfeld shows that the bias can be reduced by considering a country's consumption correlation with rest-of-world consumption. In this respect, it is possible that some downward bias may remain when the world aggregate, instead of consisting of a larger set of countries, is composed of the G-7 countries only. Still, note that in 1998 the G-7 countries had a share of more than 72 percent in total OECD output (which includes all highly industrialized countries).

Table 1
Per Capita Consumption and Output Correlations Between Domestic and Rest-of-World Growth Rates: 1973–1998

Country	Canada	United States	Japan	France	Italy	United Kingdom	Germany
Correlation of Domestic and Rest-of-World Consumption Growth	.3340	.3549	.2322	.4510	.0882	.3176	.2837
Correlation of Domestic and Rest-of-World Output Growth	.4917	.3935	.2359	.4329	.2605	.4440	.3195

Source: Author's computations based on OECD *Quarterly National Accounts* data.

tive assumptions needed to obtain the prediction in equation (1), such a result is not very surprising.

Most important, with the exception of France, the sample correlations of domestic consumption with rest-of-world consumption are *lower* than the corresponding sample correlations of domestic output with world output. For Canada and Italy, the hypothesis that domestic consumption is less correlated with world consumption than domestic output is with world output, that is, $corr(\Delta c, \Delta c^w) < corr(\Delta y, \Delta y^w)$, cannot be rejected at standard confidence levels. For all other countries, including France, the correlation of domestic consumption with rest-of-world consumption is not statistically different from the corresponding correlation of domestic output with rest-of world output. Thus, there is no instance in which the coherence of domestic consumption fluctuations with world consumption is statistically larger than the coherence of domestic output fluctuations with world output. Note that not only condition (2), but also condition (3) finds no support in the data: The correlation of domestic consumption with domestic output can be shown to be statistically larger than the corresponding correlation of domestic consumption with world output.

We now investigate whether there is any evidence that the coherence of domestic consumption with world consumption has been increasing. For this purpose, we split the sample into two subperiods, 1973 to 1987 and 1988 to 1998. The split is chosen to isolate the 1990s. Panels A and B of Table 2 report per capita output and private consumption growth correlations with rest-of-world output and private consumption growth for the periods 1973 to 1987 and 1988 to 1998, respectively.

Several things are worth noting. First, the sample correlations of domestic output with rest-of-world output generally show a marked decline from the first subperiod to the next, especially for the three largest economies. Japan's output growth correlation with rest-of-world output changed from a positive value over the period 1973 to 1987 to a significantly negative value during the more recent period. This negative relationship is largely due to the combination of the prolonged Japanese recession and the simultaneous U.S. expansion since 1992. Similarly, the coherence of the U.S. business cycle with the rest of the G-7 countries has virtually disappeared over the period 1988 to 1998, reflecting the divergent developments in the United States, Japan, and continental Europe. For the same reason, Germany's correlation of output growth with the rest of the G-7 countries has also declined substantially. The evolution of per capita output growth correlations thus corroborates our previous characterization of business cycle fluctuations across highly industrialized countries as being driven in large part by global shocks during the 1970s and 1980s and, to a greater extent, by country-specific shocks in more recent years.

The correlations of domestic consumption with rest-of-world consumption also show a decline from the first subperiod to the next. It is true that in the more recent period consumption correlations have been higher than the corresponding output correlations in Germany, Japan, and the United States. However, for all G-7 countries there has been no tendency for the coherence of domestic consumption with world consumption to increase over time. The sample correlations of domestic consumption growth with rest-of-world consumption growth are lower in the later period than they were previously.

This result may seem surprising, given the increase in world financial integration and sophistication. However, it is likely that the sample correlations for consumption growth over the period 1973 to 1987

Table 2
Per Capita Consumption and Output Correlations Between Domestic and Rest-of-World Growth Rates

Panel A: 1973–1987

Country	Canada	United States	Japan	France	Italy	United Kingdom	Germany
Correlation of Domestic and Rest-of-World Consumption Growth	.3466	.4065	.2623	.4801	.0638	.3384	.3699
Correlation of Domestic and Rest-of-World Output Growth	.5237	.4744	.4232	.4303	.2850	.4357	.3881

Panel B: 1988–1998

Country	Canada	United States	Japan	France	Italy	United Kingdom	Germany
Correlation of Domestic and Rest-of-World Consumption Growth	.2579	.2253	.1718	.3637	.0610	.2409	.0276
Correlation of Domestic and Rest-of-World Output Growth	.3265	.0173	−.2017	.5021	.0727	.5014	.0043

Source: Author's computations based on OECD *Quarterly National Accounts* data.

overstated the extent of international consumption risk-sharing because of the presence of sizable global shocks. In this respect, the sample correlations for the more recent years may provide a more accurate assessment of consumption risk-sharing, since global shocks have played a relatively minor role.

Overall, the findings provide little support for much international risk diversification, even when one restricts the analysis to the period 1988 to 1998. Although in the case of Japan condition (2) is satisfied at standard confidence levels, it can be shown that condition (3) is strongly rejected. The correlation of Japan's consumption and output growth from 1988 to 1998 is 0.81, while the correlation of Japan's consumption with world output is significantly lower at 0.35. While Panel B of Table 2 shows that the sample consumption correlations for the United States and Germany are greater than the corresponding output correlations, one cannot reject the hypothesis that the correlations differ only because of sampling error. This is also true for the remaining countries with the exception of the United Kingdom, for which the output correlation is significantly greater than the corresponding consumption correlation. Furthermore, condition (3) is never met—or even approached. In all circumstances a country's consumption correlation with its own output is greater than the country's consumption correlation with world output.

In all G-7 countries, private domestic consumption is highly correlated with domestic output. In contrast to the prediction of the complete markets model with perfect capital mobility, idiosyncratic country-specific shocks play an important role in explaining consumption fluctuations. This can also be seen by regressing a country's consumption growth on rest-of-world consumption growth and domestic output growth:

$$\Delta c_t = a + b_1 \Delta c_t^W + b_2 \Delta y_t .$$

According to the model outlined in Section I, b_1 should equal 1 while b_2 should equal 0, since no country-specific component should explain domestic consumption growth. The following is the regression's result for Japan over the period 1988 to 1998, with *t*-statistics reported in parentheses:

$$\Delta c_t = -.02 + .52 \Delta c_t^W + 1.02 \Delta y_t$$
$$(-1.1) \quad (1.7) \qquad (9.0) \qquad R^2 = .66.$$

The estimates provide no compelling evidence in support of international consumption risk-sharing, and it can be shown that similar results apply to the other G-7 countries.[5]

[5] It is possible to argue that the comparison of consumption fluctuations with total output fluctuations may not be entirely appropriate for addressing the issue of the degree of global risk-

In sum, despite the increase in world financial markets' integration and sophistication, the puzzle of low international consumption risk-sharing continues to persist. The picture for the most recent period suggests that the evidence for increasing consumption risk-sharing documented by Obstfeld (1994a) over the period 1951 to 1988 is probably a consequence of the high coherence in international business cycles during the years 1973 to 1988. When the coherence in business cycles disappeared, the correlation of domestic consumption growth with world consumption growth declined. This empirical pattern should not be interpreted as evidence against greater financial integration over the most recent period, but rather as an indication that the already low international consumption correlations of the 1970s and 1980s were likely overstating the actual degree of consumption risk-sharing.

What can explain the low degree of international risk-sharing suggested by the finding of little coherence of domestic consumption with world consumption movements? One possibility is that actual trade in international assets is not being undertaken with the goal of smoothing consumption, in contrast to the presumption of the complete asset markets model with perfect capital mobility and utility-maximizing consumers that we are considering. Tesar and Werner (1995) note that the composition of the portfolio of foreign securities in five of the G-7 countries (United States, Japan, Germany, United Kingdom and Canada) reflects factors other than risk diversification. In particular, trade linkages, language, and geographical proximity appear to be more important than risk-sharing motives. Such an explanation implicitly suggests that the welfare gains from consumption risk-sharing are small, an issue discussed in the next section.

Another explanation involves informational asymmetries. Portes and Rey (1999) document an important geographical component in international asset trading. Specifically, distance appears to play a significant role in explaining cross-border equity flows, with countries farther from each other trading a smaller amount of financial assets. Portes and Rey interpret this finding as evidence of informational asymmetries that could limit the extent of international portfolio diversification and consumption risk-sharing. In this respect, Gehrig (1993) argues that domestic investors are reluctant to invest in a foreign country if investors in that country receive a more precise signal about foreign asset returns than the domestic investors.

In sum, despite the increase in world financial markets' integration and sophistication, the puzzle of low international consumption risk-sharing continues to persist.

It is also possible that the low consumption correlations simply reflect the fact that a country's output is not entirely tradable. Several studies have shown that explicitly accounting for the presence of non-traded goods in models of the kind described in Section I can lead to correlation coefficients for domestic consumption and output that are more in line with their empirical counterparts. However, in order to match actual data, these models also require the presence of large preference shocks, whose empirical relevance is inherently hard to assess. In applied work, Lewis (1996) argues that there is substantial evidence in favor of consumption risk-sharing among industrialized countries after one controls for consumption tradability and durability. Still, her findings should be interpreted cautiously in light of the potentially large measurement errors.

Despite the *embarras de richesses* for plausible explanations of the consumption correlations puzzle, there does not appear to be a wide consensus yet over the relative importance of alternative hypotheses ad-

---

sharing, since only the output remaining after investment and government consumption can be shared by private consumers (see Obstfeld 1994a). Thus, a more relevant exercise would be to regress a country's per capita consumption growth on rest-of-world consumption growth and on per capita output growth *net* of investment and government consumption:

$$\Delta c_t = a + b_1 \Delta c_t^W + b_2(\Delta y_t - \Delta i_t - \Delta g_t),$$

where i and g denote the country's (log) levels of investment and government consumption respectively. Still, the following regression's result for Japan over the period 1988 to 1998 illustrates that consumption growth remains more strongly associated with domestic than with global factors:

$$\Delta c_t = -.01 + .06\Delta c_t^W + 1.07(\Delta y_t - \Delta i_t - \Delta g_t)$$
$$(-.94)\ (.23) \qquad (11.0) \qquad R^2 = .76.$$

Similar findings hold for the other G-7 countries.

vanced in the literature. Informational asymmetries, preference shocks, nontradable consumption, incomplete markets, and transaction costs are all likely to play a role, and much of the recent research is aimed at testing their relative importance. While this is a difficult task, it is a necessary step toward a better understanding of the current and potential risk-allocation role of world capital markets.

III. Are There Gains from International Consumption Risk-Sharing?

Economic theory has long recognized the benefits of trade in international assets, but disagreement exists over their potential magnitude in actual practice, especially with regard to the gains from international consumption risk-sharing. Lucas (1987) suggests that the welfare gains from diversifying risk internationally can be quite small, on the order of a fraction of a percentage point increase in permanent aggregate consumption. If this is indeed the case, then small transaction costs could be sufficient for discouraging international consumption risk-sharing. However, caution is warranted in interpreting Lucas's estimate, since the computation appears to be sensitive to realistic changes in assumptions.

Economic theory has long recognized the benefits of trade in international assets, but disagreement exists over their potential magnitude in actual practice, especially with regard to the gains from international consumption risk-sharing.

Specifically, Lucas's assumption that a country's output fluctuations are all temporary has been questioned by several economists. If permanent shocks to a country's output sometimes occur, then the welfare gains from international risk-pooling could amount to several percentage points of permanent consumption.[6] This happens because permanent shocks can generate considerable uncertainty over a country's long-run output level. Athanasoulis, Shiller, and van Wincoop (1999) argue that over medium to long horizons, the uncertainty can be large even for highly industrialized countries, whose output fluctuations are relatively small. For a set of OECD countries, they estimate that over a 35-year horizon the probability that per capita GDP for the best-performing country rises 70 percent relative to that of the worst-performing country is more than 0.60, a significant number that would translate into sizable gains from international consumption risk-sharing.[7]

In addition, Lucas's computation relies upon the critical assumption that within a country risks have already been shared optimally, leaving only the amount of risk-sharing between countries to be determined. Thus, Lucas's estimate provides, at best, an indication of the welfare gain that international risk-sharing may yield once domestic risk-sharing opportunities have been exhausted. However, if financial markets are incomplete and it is not possible to insure, either domestically or internationally, against some contingencies, then gains from pooling insurable risks globally could be much larger. This is suggested by research showing that uninsurable risks greatly increase the price that people are willing to pay for feasible risk reduction. (See Obstfeld and Rogoff 1996, ch. 5, pp. 331–32.)

While there is no wide consensus yet on the benefits of international diversification, it seems premature to argue that the low consumption correlations found in actual data are an indication that global risk-sharing is not terribly important, as Lucas's estimate would suggest. Still, if gains from international diversification are indeed large, one is left with the puzzle of why the benefits are not fully exploited in actual practice. A possible explanation, advanced by Shiller (1993), is that international financial markets do not yet provide adequate instruments for hedging against currently nontradable income components. Shiller notes that actual opportunities for fully insuring against labor income risk are limited, not only because a worker's effort might be subject to moral hazard, but also because labor income is generally

[6] See van Wincoop (1999).

[7] Qualifications to Lucas's computation also come from work by Obstfeld (1994b), who argues that a country is likely to benefit greatly from expanded risk-sharing opportunities, because capital market integration allows the country to place a larger fraction of its wealth in high-expected-return, risky assets that significantly enhance the country's growth prospects. Far from being trivial, the gains from international risk-sharing in this case could be equivalent to an increase in permanent consumption of several percentage points.

only weakly correlated with the returns on existing financial assets (see Bottazzi, Pesenti, and van Wincoop 1996). A less than perfect correlation implies that trade in existing financial assets can provide only partial insurance against labor income fluctuations.

Shiller argues that if countries were to trade in GDP-linked securities, the possibilities for hedging against labor income fluctuations would be much greater. This is because labor income constitutes a very large fraction of a country's GDP, and thus the correlation between a country's labor income and the return on the country's GDP-linked security would be high. A country could then obtain insurance against labor income fluctuations by selling short the security linked to its own GDP and investing the proceeds globally.

While there might be a number of practical obstacles to introducing GDP-linked securities, Shiller's proposal highlights the issue that further financial innovations might be needed in order for countries to reap the benefits of international consumption risk-sharing.

IV. Conclusions

This article examines the extent of international consumption risk-sharing across G-7 countries. Output and consumption correlations indicate that per capita domestic consumption is highly correlated with per capita domestic output and, in contrast to the predictions of a model with complete asset markets and perfect capital mobility, the coherence between domestic consumption and world consumption is weak. The increasing degree of coherence between domestic and world consumption over the post-World War II period documented in other studies does not seem to hold, once the 1990s are taken into consideration. This result is surprising, since the general presumption is that the overall trend toward greater financial integration and sophistication should have led to increased consumption risk-sharing and, thus, to greater coherence in consumption fluctuations across highly industrialized countries.

We note, however, that the evidence suggesting a higher degree of international consumption risk-sharing between 1973 and 1988 than during the 1990s is probably the result of the important role that global shocks played in driving business cycle fluctuations in the former period. With the coherence in international business cycles fading over the last decade, the correlation of domestic consumption with world consumption has also declined. This is in accordance with a characterization of domestic private consumption movements as being closely related to domestic output fluctuations.

Of course, it is possible that the consumption correlations our analysis is based upon do not provide a good indication of the extent of consumption risk-sharing because, for example, nontradable consumption has not been taken into account explicitly. Nevertheless, other evidence related to foreign equity holdings by domestic residents suggests that international trade in assets is still limited: Japanese and U.S. investors, for example, hold at least 90 percent of their equity portfolios in domestic assets (Kang and Stulz 1997; Tesar and Werner 1997). Such low levels of foreign equity holdings are unlikely to provide an optimal level of international diversification. Given the continuing process of global financial integration, it is plausible that 20 years from now portfolios will be better diversified internationally, and that domestic consumption will be less responsive to idiosyncratic country-specific shocks. However, this may require the introduction of financial instruments that are currently unavailable to investors.

References

Athanasoulis, Stefano, Robert Shiller, and Eric van Wincoop. 1999. "Macro Markets and Financial Stability." *FRBNY Economic Policy Review* (April), 21–39.

Bottazzi, Laura, Paolo Pesenti, and Eric van Wincoop. 1996. "Wages, Profits, and the International Portfolio Puzzle." *European Economic Review* 40, no. 2 (February), 219–54.

Gehrig, Thomas. 1993. "An Information Based Explanation for the Domestic Bias in International Equity Investment." *Scandinavian Journal of Economics* 95, no. 1, 97–109.

International Monetary Fund. 1998. *World Economic Outlook* (May). Washington, DC.

Kang, Jun Koo, and Rene M. Stulz. 1997. "Why Is There a Home Bias? An Analysis of Foreign Portfolio Equity Ownership in Japan." *Journal of Financial Economics* 46, no. 1, 3–28.

Lewis, Karen K. 1995. "Puzzles in International Financial Markets." In Gene Grossman and Kenneth Rogoff, eds., *Handbook of International Economics,* vol. 3. Amsterdam: North Holland.

———. 1996. "What Can Explain the Apparent Lack of International Consumption Risk Sharing?" *Journal of Political Economy* 104 (April), 267–97.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles.* Oxford, UK: Basil Blackwell.

Obstfeld, Maurice. 1994a. "Are Industrial-Country Consumption Risks Globally Diversified?" In Leonardo Leiderman and Assaf Razin, eds., *Capital Mobility: The Impact on Consumption, Investment, and Growth.* Cambridge, UK: Cambridge University Press.
———. 1994b. "Risk-Taking, Global Diversification, and Growth." *The American Economic Review* 84 (December), 1310–29.
———. 1995. "International Capital Mobility in the 1990s." In Peter B. Kenen, ed., *Understanding Interdependence: The Macroeconomics of the Open Economy.* Princeton, NJ: Princeton University Press.
Obstfeld, Maurice and Kenneth Rogoff. 1996. *Foundations of International Economics.* Cambridge, MA: The MIT Press.
Pakko, Michael R. 1998. "Characterizing Cross-Country Consumption Correlations." *The Review of Economics and Statistics*, vol. 80, no. 1 (February), 169–74.
Portes, Richard and Hélène Rey. 1999. "The Determinants of Cross-Border Equity Flows." *NBER Working Paper* No. 7336. Cambridge, MA.
Shiller, Robert J. 1993. *Macro Markets: Creating Institutions to Manage Society's Largest Economic Risks.* Oxford, UK: Clarendon Press.
Tesar, Linda L., and Ingrid Werner. 1995. "Home Bias and High Turnover." *Journal of International Money and Finance* 14 (August), 467–92.
———. 1997. "The Internationalization of Securities Since the 1987 Crash." Working Paper, University of Michigan.
van Wincoop, Eric. 1999. "How Big Are Potential Welfare Gains From International Risk Sharing?" *Journal of International Economics* 47, 109–135.