

International Risk-Sharing: Globalization Is Weaker Than You Think

BY SYLVAIN LEDUC

W

ith the development of international financial markets, households should be better equipped to pool their resources so that their level of consumption varies less from year to year. Yet the extent of international risk-sharing remains surprisingly small. In this article, Sylvain Leduc digs a little further into the data to uncover why, in spite of recent trends, financial globalization remains weaker than you think.

From 1980 to 2004, world trade in goods and services increased from 36 percent to 50 percent of world GDP. As the world experienced a surge in the trade of goods and services, it also saw a substantial rise in the trade of financial assets. The share of foreign equities in U.S. investors' portfolios, for instance, increased from about 1 percent in the early 1980s to 12 percent in 2000.¹ On that dimension, the impression that we are living in a more integrated world is borne out in the data. But if we dig in a little further,

¹ See Francis Warnock's article.



Sylvain Leduc is an economist with the Board of Governors. When he wrote this article, he was a senior economist at the Philadelphia Fed.

we will find that, notwithstanding the trend toward globalization, the world's economies remain strikingly insular along many dimensions.

With the developments of international financial markets, households should be better equipped to diversify their portfolios and protect their investments against unforeseen events, which ultimately should result in more sharing of consumption risk across countries. That is, households would effectively pool their resources so that their level of consumption varies less from year to year. Yet, the extent of international risk-sharing remains surprisingly small and is one key reason that globalization is weaker than you think.

Standard macroeconomic models offer predictions regarding the extent of international risk-sharing. If consumers are diversifying internationally, we should see consumers in one country consuming more than those in another country when the price of doing so is lower than in the other

country. This relative price is the real exchange rate, that is, the exchange rate between the countries' currencies adjusted for the rate of inflation in the two countries. One reason for the lack of international risk-sharing is that, empirically, real exchange rates often move in a way that hinders the risk-sharing process. As a result, full globalization remains far away, at least along this important dimension.

INTRODUCING RISK-SHARING

At the base of the concept of risk-sharing is the idea that most people would prefer to keep a relatively stable pattern of consumption instead of a highly variable one. The challenge is to achieve this smooth consumption pattern even though income may vary a lot from year to year. For instance, many workers are, at times, temporarily laid off because of a slowdown in their particular line of business. Or people may have to temporarily quit their jobs for health reasons. Depending on the frequency of such events, incomes can vary quite a bit in any given year.

If households do not save or borrow, their level of consumption will follow their variable level of income. For instance, imagine a simple economy composed of two households, the Greens and the Verdis, that have fluctuating incomes from year to year.² Suppose we look at how much money these households made over the last two years and we find that the Greens had an after-tax income of \$10,000 in year 1 and \$30,000 in year 2. For

² See also Keith Sill's *Business Review* article for a discussion of risk-sharing.

simplicity, imagine that the opposite is true for the Verdis: in year 1, the Verdi household took home \$30,000, while it earned \$10,000 in year 2.

First, to keep the argument simple, assume that both households use their income to consume the same basket of goods and that they pay the same price for one unit of those goods, \$1. This is an important assumption that I will relax in the next section. If the households do not save or borrow, their level of consumption will follow their level of income. That is, in year 1 the Greens consume 10,000 units of goods and the Verdis 30,000 units of goods, and vice versa in year 2.

How could the Greens and the Verdis achieve a relatively more stable consumption pattern? It could be simply achieved if we let the households pool their income each year and divide the total equally between them. Both households could therefore keep a constant consumption level of 20,000 units of goods per year. Notice that, in this example, one implication is that risk-sharing equalizes consumption across the two households. That is, by pooling their resources, households are able to “share” the risks of their fluctuating incomes and therefore eliminate or “insure” against their consumption risk.

However, it might be quite difficult to find another household that will agree to pool its income with yours. In practice, this risk-sharing process is instead carried out through financial markets. For instance, households can save by buying stocks of firms or government bonds when their income is unexpectedly high, or they can buy goods with credit when their income is unexpectedly low and repay their debt in more prosperous times. Through borrowing and lending in financial markets, households can smooth out the bumps in their income streams and

achieve a more stable consumption pattern. As long as households keep a well-diversified investment portfolio, they are better equipped to smooth out their consumption risk. Indeed, one of the tenets of modern finance is that households should hold a well-diversified investment portfolio so that the portfolio’s overall risk is less subject to the vagaries of one particular sector or one particular stock.

In the above example, note that I did not mention the country of residence of the two households. In fact, the argument does not depend on the households’ locations. As long as household incomes do not move in the same direction — up or down — at the same time, there is scope for sharing consumption risk, be it within or between nations. Since world economies are not always in sync, and some countries fall into recession while others continue to expand, household incomes in different countries do not

both the Greens and the Verdis to be temporarily laid off at the same time. In this case, there is no scope for mutually beneficial trade by which to insure against consumption risk.

Global risks will necessarily trigger movements in consumption. But every household’s consumption will be moving in the *same* way. Therefore, in a world in which households can use financial markets to insure against all possible *idiosyncratic* risks to their income and in which households consume the same basket of goods and pay the same price for those goods, theory predicts that consumption should move in the same direction across countries.

INTERNATIONAL RISK-SHARING AND RELATIVE PRICES

Obviously, this prediction is derived under relatively strong conditions. For instance, it is unlikely that households consume the same basket of goods and services. There

As long as household incomes do not move in the same direction — up or down — at the same time, there is scope for sharing consumption risk, be it within or between nations.

always move together. So there is potential for sharing consumption risk across countries.

However, households cannot insure against every type of risk. For instance, global risk (as opposed to idiosyncratic risk) is not insurable, since it affects everyone in the same manner, at the same time.³ In terms of our previous example, global risk could include a recession that leads

is also ample evidence that different consumers do not pay the same price for the same goods, especially when these consumers live in different countries (see *Where You Are Affects How Much You Pay*). Once we relax those assumptions, we obtain a more general prediction about sharing consumption risk. In this case, efficient risk-sharing dictates that the household facing the lower relative price consume more.

To see that, let’s look again at our previous example. Suppose that the Greens’ and the Verdis’ income patterns in year 1 and year 2 continue to

³ Contrary to global risk, which affects everybody in the economy, idiosyncratic risks affect only particular individuals.

Where You Are Affects How Much You Pay

I

n the early 1980s, total trade in goods accounted for 36 percent of world GDP; 23 years later, that ratio surged to 50 percent. The fall in trade barriers, initiated after World War II under the General Accord on Tariffs and Trade (GATT), in large part triggered the rise in the trade of goods. As more goods are traded, you might expect the prices of these goods in different parts of the world to converge. That is, what economists called the *law of one price* would hold: A product would sell for the same price (expressed in the same units of currency) in different locations, absent natural or government-imposed trade barriers.

Imagine that you can freely trade cars between the U.S. and Canada and you notice that a Ford Explorer sells for \$5,000 more in Montreal than in Detroit, once you convert the price of a Ford Explorer from Canadian dollars into U.S. dollars using the exchange rate. A profitable business opportunity, called arbitrage, would be to buy Ford Explorers in Detroit at the cheaper price and sell them in Montreal for a profit of \$5,000. As long as prices (expressed in a common currency) of Ford Explorers differ between these two markets, there is an opportunity for arbitraging the price difference. Obviously, it is not costless to trade goods, since businesses have to pay transportation costs, tariffs, or the costs associated with different regulations in different locations. The presence of these costs will allow prices to differ across locations. However, as long as goods can be freely traded, prices of goods should be equalized across countries. In this case, prices would obey the law of one price.^a

You can arbitrage price differentials not only in markets in different countries but also in markets located in the same country.^b Arbitrage opportunities should tend to equalize prices in

different locations. However, it appears that price differentials are much larger across countries than across locations in a given country. For instance, in a widely cited article, economists Charles Engel and John Rogers documented that prices vary much more between Toronto and New York, say, than between Detroit and New York. This implies that price differentials across countries are not solely the result of transportation costs, since the distance between Toronto and New York is about the same as that between Detroit and New York. Rather, there seems to be something special about crossing borders.

Prices can indeed differ widely across countries.^c Mario Crucini, Chris Telmer, and Mario Zachariadis documented the price differentials for selected traded goods in different European countries. They found that price differentials are indeed large, once prices are converted in common currency units. For instance, they found that Austrians pay twice the amount Belgians pay for one pound of long-grain rice. Washing detergent is twice as expensive in Greece as it is in Germany. And two pounds of coffee is 40 percent cheaper in France than in Italy.

Moreover, it appears that deviations from the law of one price are fairly stable through time. In a National Bureau of Economic Research paper, economists Kenneth Froot, Michael Kim, and Kenneth Rogoff showed that for many commodities (for instance, barley, butter, and silver), the deviations from the law of one price are not just a property of modern economies; they were present as far back as the 13th century.

In a nutshell, the law of one price fails dramatically, and this failure provides another example that globalization is weaker than you think.

^a When the law of one price holds for every good in the economy, exchange rates will be determined according to what economists call purchasing power parity, or PPP. PPP states that nominal exchange rates should move to offset differences in inflation across countries, leaving real exchange rates constant over time. Notice that this simple approach to exchange-rate determination cannot explain the high volatility of real exchange rates.

^b See Leonard Nakamura's *Business Review* article for a discussion of the failure of the law of one price across U.S. retailers and its impact on the measurement of inflation.

^c Kenneth Rogoff's article provides a survey of the large empirical literature documenting the failure of the law of one price.

be the same as before: the Greens have an after-tax income of \$10,000 in year 1 and \$30,000 in year 2. Further suppose that the opposite is true for the Verdis. However, let's now assume that the two households do not pay the same price for the goods. Suppose that in year 1, the Greens continue to pay \$1, but the Verdis now must spend \$2 to obtain the same goods and that the

reverse is true in year 2.

If the households do not pool their resources, the Greens will consume 10,000 units of goods the first year and 15,000 units in the following year, since it must then pay \$2 for the goods. For the same reasons, the Verdis' consumption will fluctuate between 15,000 and 10,000 units between year one and year two. In

this case, the household that faces the cheaper price does not consume more. For instance, even though the Greens pay half the price as the Verdis in year 1, they consume 5,000 fewer units.

By pooling their income (\$40,000 in each year) and dividing the total equally between them (\$20,000 per household in each year), the Greens and the Verdis can take advantage of

the price differentials and achieve a more efficient consumption pattern. In year 1, the Greens would consume twice as much as the Verdis (20,000 versus 10,000 units of goods), since it must pay half the price the Verdis pay for the same goods (\$1 versus \$2). Since, in the second year, the Verdis face a lower price than the Greens (\$2 versus \$1), they will consume more (20,000 versus 10,000 units).

Note that when households face different prices, efficient risk-sharing does not state that consumption should move together across households. Rather, efficient risk-sharing dictates that the household facing the lower relative price should consume more. Intuitively, this criterion makes sense, since the world economy should channel more consumption to places where it is relatively cheap to consume.⁴

Once again, it is immaterial whether these two households live in the same country. The only difference is that when households live in different countries, the relative price of goods has a particular name: the real exchange rate.

⁴ Another way to think about optimal risk-sharing is to think in terms of costs and benefits. Optimal risk-sharing occurs when the benefit of transferring one extra dollar from the Verdis to the Greens (or vice versa) equals the cost. As long as the marginal benefit of the transfer exceeds the marginal cost, it is beneficial to transfer resources from the Greens to the Verdis. For instance, in year 1 the benefit of transferring one extra dollar from the Verdis to the Greens is that the Greens now consume one more unit. However, such a transfer has a cost. To transfer one extra dollar to the Greens, the Verdis have to lower their consumption by half a unit, since the Verdis pay twice as much as the Greens for the same basket of goods. Therefore, the cost of the transfer is the relative price, 2, times 0.5 units of consumption, which is 1 unit of consumption. Therefore, optimal risk-sharing occurs because the marginal benefit of transferring one extra dollar from the Verdis to the Greens exactly equals the marginal cost.

INTERNATIONAL RELATIVE PRICES: REAL EXCHANGE RATES

People usually think about nominal exchange rates, which denote the price of one currency in terms of another. For instance, in the first quarter of 2003, one British pound was worth 1.60 U.S. dollars. One year later, the British pound traded for 1.84 U.S. dollars. Therefore, the U.S. dollar lost 15 percent of its value against the British pound over that year.⁵

The real exchange rate, on the other hand, is the nominal exchange rate multiplied by the ratio of price levels in the two countries, as measured, for instance, by the consumer price index.⁶ A change in the real exchange rate, therefore, represents a change in the relative price of two countries' goods, controlling for inflation.

For instance, in the first quarter of 2004, the consumer price index in the United States was 121.4, and the consumer price index in the U.K. was 179.2, implying a real exchange rate of 2.36: the nominal exchange rate of 1.60 U.S. dollar per British pound times the ratio of U.K. to U.S. price indices. By the first quarter of 2003, however, the U.S. consumer price index had risen to 123.4, while the

⁵ Throughout this article I will denote the exchange rate in foreign currency units, i.e., how many U.S. dollars one unit of foreign currency (in the above example, a British pound) is worth. In this case, an upward movement in the exchange rate implies a depreciation of the U.S. dollar.

⁶ The consumer price index, or CPI, measures the cost of living for a typical urban family. The index shows how the price of a typical basket of goods changes from year to year. So the real exchange rate between the U.K. and the U.S. equals the number of dollars per British pound times the ratio of prices in the U.K. relative to that in the U.S.: $\left(\frac{\text{dollar price level in the UK}}{\text{pound price level in the US}} \right)$. Again, notice that a rise in the real exchange rate implies a depreciation of the U.S. dollar in real terms.


U.K.'s had increased to 183.8; thus, the real exchange rate rose to 2.74. So the real exchange rate increased 16.1 percent from the first quarter of 2003 to the first quarter of 2004. In other words, while \$1 would buy 15 percent fewer pounds in the first quarter of 2004 compared with one year earlier, \$1 of U.S. goods could be traded for 16.1 percent fewer British goods in the first quarter of 2004 than in the first quarter of the previous year.

The real exchange rate is the nominal exchange rate multiplied by the ratio of price levels in the two countries, as measured by the consumer price index.

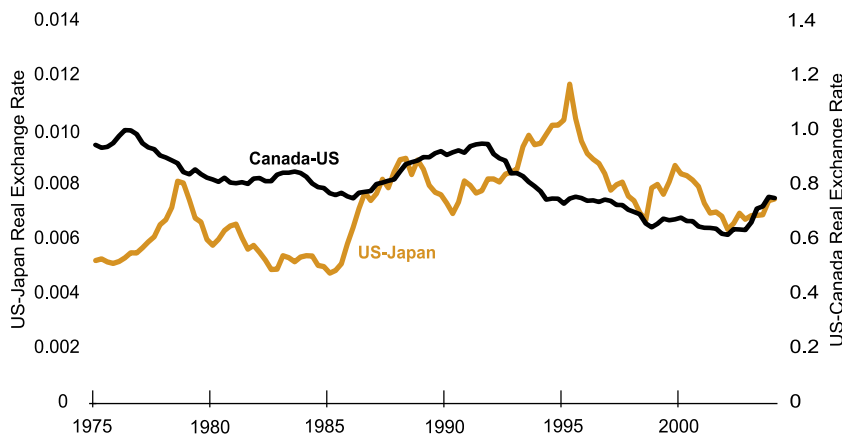
The variations in the U.S.-U.K. real exchange rate between 2003 and 2004 are not unusual. In fact, the real exchange rate has been varying widely over time (Figure 1). Moreover, other currencies, such as the Canadian dollar or the Japanese yen, have experienced similarly large fluctuations (Figure 2). The reasons for those large swings in real exchange rates have intrigued and puzzled international economists for quite a while.

What underlies the large fluctuations in real exchange rates? John Rogers and Michael Jenkins found that the source of movements in real exchange rates is the failure of the law of one price (see *Where You Are Affects What You Pay*).⁷ In fact, they found that 81 percent of the movements in real exchange rates

⁷ Under the law of one price, a good should sell for the same price in different locations, once the prices of the good are expressed in the same currency units and if there are no transport or trade-related costs.

FIGURE 1**U.S.-U.K. Real Exchange Rate**

The real exchange rate is constructed using CPI indices in the U.S. and the U.K. The exchange rates are number of U.S. dollars per unit of British pound.

FIGURE 2**U.S.-Japan and U.S.-Canada Real Exchange Rates**

The real exchange rates were constructed using CPI indices in Canada, Japan, and the U.S. The exchange rates are number of U.S. dollars per unit of Canadian dollar or Japanese yen.

occur because traded goods do not sell for the same price in different countries, once those prices are expressed in common currency units. Using a different methodology, Charles Engel showed that over 95 percent of the variations in real exchange rates are the result of deviations from the law of one price.

As we saw in the previous section, when households do not face the same price for the same goods, risk-sharing has to be modified to take into account the movements in relative prices. For households located in different countries, efficient risk-sharing dictates that consumption should be higher in the country where the rela-

tive price of consumption (that is, the real exchange rate) is lower. In other words, when the U.S. experiences a fall in the price of its consumption basket relative to that in Europe (a *depreciation* of its real exchange rate), it should also be consuming more. However, this does not appear to be the case.

THE LACK OF INTERNATIONAL RISK-SHARING

A simple way to look at the extent of consumption risk-sharing is to look at the correlation between the real exchange rate and the ratio of consumption between different countries. Here we focus on this correlation for the U.S. vis-à-vis other OECD countries (Table).⁸ The correlation captures how these two variables move over time. For instance, a positive correlation implies that when the real exchange rate increases (a depreciation of the U.S. dollar in real terms),⁹ consumption in the U.S. should rise relative to that in the foreign country. (I will call *relative consumption* the movement in U.S. consumption vis-à-vis that of the foreign country.) On the other hand, if the real exchange rate rises as relative consumption falls, the correlation would be negative.

Under efficient risk-sharing, consumption should be higher when its relative price is lower. This implies that the correlation between relative consumption and the real exchange rate should be positive.¹⁰ When the

⁸ The Organization for Economic Cooperation and Development (OECD) is a group of 30 countries that share a commitment to democratic government and the market economy.

⁹ Remember that the exchange rates are U.S. dollars per unit of foreign currency, so that an increase in the real exchange rate implies a fall in the relative value of the dollar in real terms.

¹⁰ It can be shown that, under certain conditions, the correlation between the real exchange rate and relative consumption should be exactly one.

TABLE

Correlations Between Real Exchange Rates and Relative Consumption*

Country	Correlation with U.S.
Australia	-0.01
Austria	-0.35
Belgium	-0.12
Canada	-0.41
Denmark	-0.16
E.U.	-0.30
Finland	-0.27
France	-0.18
Germany	-0.27
Italy	-0.26
Japan	0.09
South Korea	-0.73
Mexico	-0.73
Netherlands	-0.41
New Zealand	-0.25
Portugal	-0.56
Sweden	-0.52
Spain	-0.60
Switzerland	0.16
Turkey	-0.31
U.K.	-0.47

* Consumption and real exchange rate data are annual series from the OECD Main Economic Indicators data set, from 1973 to 2001.

real exchange rate increases, which implies a fall in the relative value of the dollar in real terms, consumption in the U.S. should be higher than it is abroad. The correlations reported in the table demonstrate that there is little consumption risk-sharing among the OECD countries. In fact, all of the correlations are negative, which means that consumption is higher in

the country in which the relative price of consumption is higher — the exact opposite of what efficient sharing of consumption risk predicts. Therefore, sharing of consumption risk across the different countries of the world remains small, even though over the last several decades the world has become seemingly much more integrated.

What underlies the lack of international consumption risk-sharing across countries? One reason is obviously that investors fail to hold a well-diversified portfolio. Indeed, a large literature has documented the puzzling fact that most investors hold a disproportionate share of assets of their country of residence in their portfolio, yet another sign that globalization is weaker than you think. In other words, U.S. investors hold mostly U.S. assets, while French investors' portfolios are mainly composed of French assets. For instance, Francis Warnock, an economist at the Federal Reserve Board, reports that, in 2000, the share of foreign equities in U.S. investors' equity portfolios was about 12 percent, a substantial increase from the 1 percent share in the early 1980s. Yet, U.S. investors remain far from being well diversified: Warnock estimates that, in 2000, a well-diversified U.S. portfolio would have roughly 50 percent in foreign equities. As a result, U.S. investors are exposed to specific risk originating in the U.S., for instance, a recession in the U.S. economy. To the extent that country-specific risks are not perfectly positively correlated across countries, investors could lower the risk of their portfolios by holding stocks of different countries' firms. Trying to understand why investors do not do so remains a very active area of research. Yet, even given that investors' portfolios are not well diversified, it remains puzzling that a country's consumption is higher when its exchange rate is high relative to that of other countries.

REAL EXCHANGE RATES AND RISK-SHARING

We have seen that real exchange rates exhibit large fluctuations, sometimes gaining 10 percent to 20 percent in value in a couple of years, followed by equivalent or larger losses in value. In fact, like any other prices in the economy, real exchange rates react to changes in demand and supply conditions, which can be affected by a variety of fundamental factors such as monetary and fiscal policy or technological innovations. In a recent paper, Giancarlo Corsetti, Luca Dedola, and

Under efficient risk-sharing, consumption should be higher when its relative price is lower.

I documented one reason behind the lack of risk-sharing: Real exchange rates often move in a way that hinders risk-sharing in response to technological changes (Table).

Theory predicts that as a country becomes more productive because of an improvement in technology, it should produce and consume more goods relative to other countries, and it should also experience a depreciation of its real exchange rate, i.e., the price of its goods (in real terms) relative to that in the other country should fall. With an improvement in technology, a country can produce more goods for a given level of inputs, such as the number of workers or machines in the economy. As the supply of goods increases, prices fall. Remember that the real exchange rate is the relative price of goods across countries. As the prices of the goods a country produces fall, the real exchange rate, in general, depreciates.¹¹ Moreover, as a country becomes more productive, it also becomes richer, and its level of consumption should therefore rise

relative to the level of consumption in the rest of the world. Notice, once again, theory predicts that following a technological improvement, a country's consumption should be higher when its real exchange rate is lower. But are these predictions consistent with the data?

To verify whether improvements in technology affect economies as theory predicts, we conducted an analysis based on an empirical model, a simple vector autoregression (VAR). A VAR is a system of linear equations that link different variables together over time. For instance, a VAR with two variables — let's say the real exchange rate and consumption — would also have two equations. One equation would try to explain the movements in the real exchange rate; the other would try to explain the movements in consumption. To do so, both equations would use previous values of the real exchange rate and consumption.

Our VAR included five variables: labor productivity, real GDP, real consumption, net exports, and the real exchange rate.¹² We used a rise in U.S. labor productivity vis-à-vis the

¹¹ Note that a productivity increase can theoretically raise the real exchange rate if the productivity improvement is concentrated in the traded-goods sector and countries produce very similar traded goods. However, models in which countries specialize in the production of a particular array of traded goods generally predict a depreciation of the real exchange rate following a technological improvement.

¹² All of our variables are in growth rates. For labor productivity, real GDP, and real consumption, we take the difference between the growth rate of these variables in the U.S. and in the rest of the OECD countries. Our measure of labor productivity is that of the manufacturing sector.

rest of the OECD countries as a proxy for technological improvement in the U.S.¹³ Using our model, we estimated the effect that a sudden increase in the rate of U.S. technological progress would have on the U.S. and foreign economies. We did that by determining the impact that the change in labor productivity would have on the other variables in our statistical model.


We can chart the responses of the variables in our model to a one-time, unanticipated increase in the growth rate of labor productivity (Figure 3). The dotted line represents the estimated response of the variable to the sudden change in labor productivity; the grey area around the dotted line tells us how much confidence we can place in this estimate. In particular, when the entire area is above zero or below zero, we can say with a 90 percent level of confidence that the estimated response of, say, the real exchange rate to the unanticipated jump in productivity is significantly different from zero — that is, the unanticipated jump has an impact on the variable.

For instance, following the jump in labor productivity, the growth rate of output in the U.S. increases relative to the rest of the OECD countries. The rise in productivity is also accompanied by a rise in relative real GDP and consumption growth. These effects are the standard ones predicted by theory. However, contrary

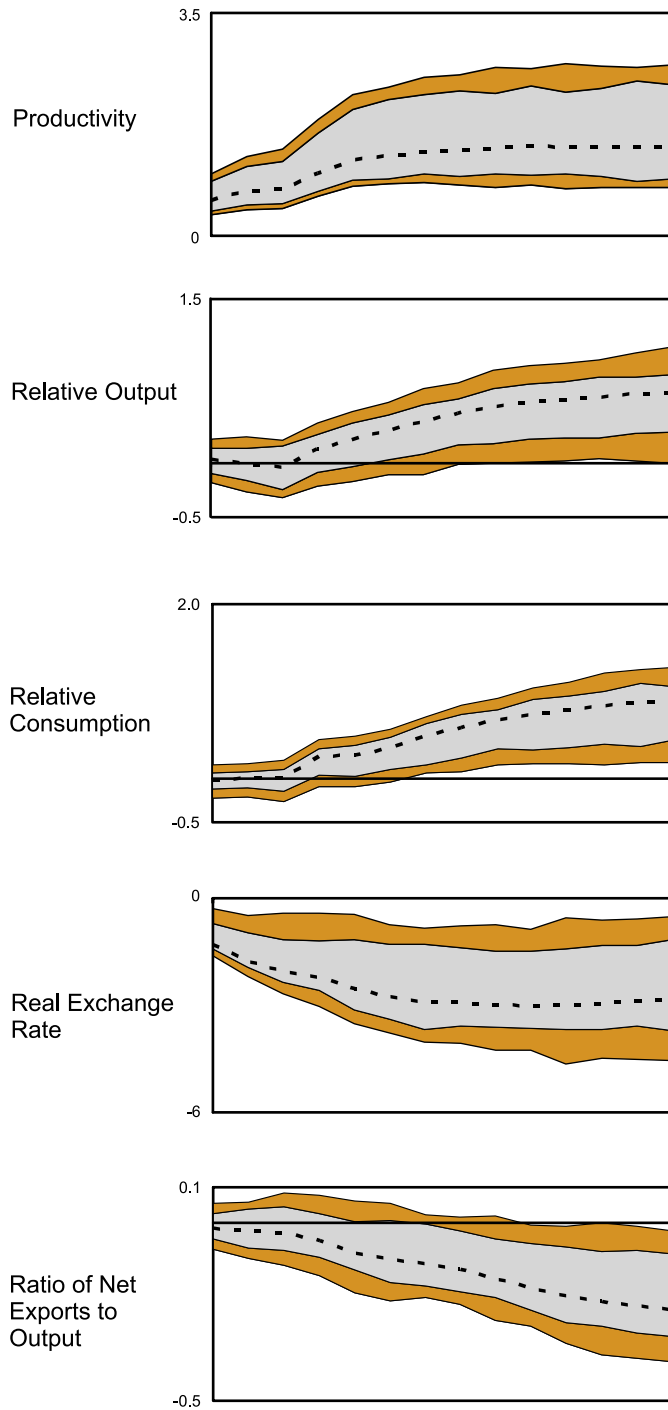
¹³ We also looked at the sensitivity of our results when we substituted total factor productivity for labor productivity: Our results are robust to this change. See my working paper with Giancarlo Corsetti and Luca Dedola for more details.

to what theory predicts, the U.S. real exchange rate *appreciates* following an improvement in productivity (that is, the real exchange rate falls), which implies, once again, that consumption is higher when its price is higher. The appreciation of the real exchange rate hinders risk-sharing. As the real exchange rate appreciates, foreign countries can consume fewer imported products, a situation that makes it more difficult for the foreign country to sustain its level of consumption. This is reflected in the fact that net exports of U.S. goods fall following an increase in labor productivity.¹⁴

SUMMARY

Notwithstanding the emergence of globalization over the last couple of decades, economies remain, to some extent, strikingly insular. Indeed, theory predicts that as the world becomes more integrated, consumption should be higher in countries where the relative price of consumption, the real exchange rate, is lower. In fact, we observe the exact opposite in the data: Consumption is higher in countries where the relative price of consumption is higher! One reason for this puzzling fact is that real exchange rates often move in a way that hinders the risk-sharing process in response to technological changes, accentuating the benefits to winners and the losses to losers. 

¹⁴ In our working paper, Corsetti, Dedola, and I detail the theoretical reasons underlying an appreciation of the real exchange rate and the terms of trade following an increase in the productivity of the traded-goods sector.

FIGURE 3**Impulse Responses to a Technology Shock in the U.S.**

The charts describe the responses from a five-variable VAR, using quarterly data. The variables are labor productivity, the real exchange rate (constructed using CPI indices), relative consumption (i.e., domestic minus foreign consumption), relative output, and net exports. All series are in percent.

REFERENCES

- Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc. "International Risk-Sharing and the Transmission of Productivity Shocks," Federal Reserve Bank of Philadelphia Working Paper 03-19, 2003.
- Crucini, Mario J., Chris I. Telmer, and Mario Zachariadis. "Understanding European Real Exchange Rates," Vanderbilt University Working Paper 01-W20 (2001).
- Engel, Charles. "Accounting for U.S. Real Exchange Rate Changes," *Journal of Political Economy* 107 (1999), pp. 507-538.
- Engel, Charles, and John H. Rogers. "How Wide Is the Border?" *American Economic Review* 86 (1996), pp. 1112-1125.
- Froot, Kenneth A., Michael Kim, and Kenneth Rogoff. "The Law of One Price Over 700 Years," National Bureau of Economic Research Working Paper 5132 (1995).
- Nakamura, Leonard. "The Retail Revolution and Food-Price Mismeasurement," Federal Reserve Bank of Philadelphia *Business Review*, May/June 1998.
- Sill, Keith. "The Gains from International Risk-Sharing," Federal Reserve Bank of Philadelphia *Business Review*, Third Quarter 2001, pp. 23-32.
- Rogers, John H., and Michael Jenkins. "Haircuts or Hysteresis? Sources of Movements in Real Exchange Rates," *Journal of International Economics* 38 (1995), pp. 339-60.
- Rogoff, Kenneth. "The Purchasing Power Parity Puzzle," *Journal of Economic Literature* 34 (1996), pp. 647-68.
- Warnock, Francis E. "Home Bias and High Turnover Reconsidered," *Journal of International Money and Finance* 21, pp. 795-805.