Why Are Goods So Cheap in Some Countries?*

BY GEORGE ALESSANDRIA AND JOSEPH KABOSKI

L ooking around the world, we observe substantial differences across countries in prices for most goods. These price differences also tend to be positively correlated with income differences, so that citizens of high-income countries tend to pay more for the same goods than citizens in low-income countries. In this article, George Alessandria and Joseph Kaboski summarize some of the evidence related to the big price differences across countries for a broad set of goods. They then discuss the relationship between prices and income levels and some possible explanations for that relationship.

Lovers of Big Macs will find China to be a true paradise and Switzerland quite the opposite, since the money spent to buy one Big Mac in Switzerland will get you almost four Big Macs in China.[1] These big international price differences haven't led Swiss Big Mac lovers to move to Beijing. In fact,

despite a much higher price, based on annual income data for 2005, the average Swiss citizen earned enough to eat eight times as many Big Macs as the average Chinese citizen.[2]

These differences in prices and purchasing power extend beyond just Switzerland and China and Big Macs. In fact, when we look across the world, we find substantial differences across countries in prices for a broad range of goods. These price differences also tend to be positively correlated with income differences so that citizens of high-income countries tend to pay more for the same goods than citizens of low-income countries.

In this article, we will summarize some of the evidence of the big price differences across countries for a broad set of goods. We will then discuss the relationship between prices and income levels. Finally, we'll discuss some possible explanations for this relationship.

MAKING INTERNATIONAL PRICE COMPARISONS
Comparing prices across countries can be difficult because prices are typically quoted in different currencies. For instance, to compare the yuan price of a Big Mac in China with the franc price in Switzerland, we need to use the nominal exchange rate between the yuan and the Swiss franc to convert the prices into a common currency. Movements in the nominal exchange rate[3] over time can thus lead Swiss Big Macs to become relatively more or less expensive compared with Big Macs in China. We will ignore the short- to medium-run fluctuations

[1] Big Mac™ is a registered trademark of the McDonald's Corporation.

George Alessandria is a senior economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www. philadelphiafed.org/econ/br/.

[2] Based on 2005 data on gross national income taken from the World Development Indicators: China $1,700 and Switzerland $54,930 (U.S. dollars).

[3] The nominal exchange rate is the value of one country's currency in terms of another country's currency.

Joseph Kaboski is an assistant professor in the Department of Economics at the Ohio State University.

related to exchange rates and instead focus on long-run differences in prices across countries.

Even though we've already seen otherwise, a natural expectation is that the price of a Big Mac should be the same everywhere; after all, it is the same good.[4] This idea is known as the law of one price (LOP). More formally, the LOP states that once prices are converted to a common currency, the same good should sell for the same price everywhere, provided there are no barriers to trade and markets are competitive.

The basic idea behind the LOP is that if prices differ across locations, firms can make some profits by buying in the low-price place and selling in the high-price place. This activity, which is called arbitrage, will continue until prices are similar in the two locations.

While the LOP is described as a "law," it does not hold for all goods. Gold and Big Macs provide evidence of its respective successes and failures as a description of world prices. The prices of Big Macs across countries reported in Table 1 provide a clear example of its failure. When converted into U.S. dollars, Big Macs sell for up to 65 percent more than in the U.S. and down to 57 percent less than in the U.S. On the other hand, from Table 2, which reports the price of one troy ounce of gold quoted on the same day at nearly the same moment on different exchanges throughout the world, we see that the LOP seems to hold, since the price of gold ranges in a 3 percent band around the price in the U.S.

One important reason the LOP does not hold is that there are barri-

TABLE 1

Some Big Mac Prices

Country	Big Mac Price		Actual Exchange Rate 1 USD=
	in Local Currency	in U.S. Dollars	
Switzerland	SFr6.30	5.12	1.23
U.S.	$3.10	3.10	1.00
China	Yuan10.50	1.31	8.00

Based on Big Mac prices and exchange rates as of March 25, 2006. The Big Mac index is published periodically by *The Economist*. Go to http://www.economist.com/markets/bigmac/ to find more information about Big Mac prices and exchange rates across many countries.

TABLE 2

Gold Prices Around the World

Exchange	Time (Eastern Standard Time)	USD/Troy Ounce
United States	10:28	$625.01
Australia	10:28	$625.00
Brazil	9:53	$617.71
Switzerland	10:28	$625.51
India	6:07	$634.89
United Kingdom	10:21	$622.75
Luxembourg	5:58	$624.50
Hong Kong	0:51	$623.40

Prices were downloaded from Bloomberg on November 3, 2006.

ers that make international trade, and thus arbitrage, costly.[5] These barriers can be man-made, such as tariffs, taxes, or trade restrictions, or physical, such as distance, which incurs shipping costs. The costs of these barriers differ quite a bit across goods. For instance, shipping costs primarily depend on the distance, weight, and mode of transportation. For goods such as gold,

[4] There are some minor differences in the size and condiments across countries. The biggest difference is in India, where Hindu and Islamic Sharia dietary laws prohibit eating beef, and so the Big Mac is made with two all-chicken patties.

[5] For a detailed breakdown of the costs to trading goods across countries, see the *Business Review* article by Edith Ostapik and Kei-Mu Yi.

which have a high value to weight ratio, shipping costs are fairly minor. For Big Macs, which, based on U.S. prices, are 1/1400 as valuable per ounce as gold and don't travel particularly well, shipping costs are relatively large.[6] However, even though it's expensive to ship a Big Mac, Big Mac prices might be the same in different countries if the inputs to producing it are very easy to trade. This is essentially true for the beef and special sauce, but it's not true for the workers who fry it up or the building in which it is consumed. For some goods, such as buildings or haircuts, the shipping costs are so high that they are almost never traded. Economists call these goods nontraded goods.

Another reason prices may differ across countries is that the competitive environments may differ. For instance, in some countries, there may not be many close substitutes for a Big Mac, and so Big Macs might be relatively expensive. However, in countries with lots of low-cost alternatives, Big Macs might cost relatively less. Or it might be the case that people in some countries are just willing to pay more for certain goods. Firms take advantage of these differences in willingness to pay for certain goods by charging different prices across countries. Charging different people different prices for the same good is known as price discrimination, and it is a common practice in many industries.[7] To make this strategy effective, firms make arbitrage

difficult by changing their product slightly across countries. For instance, film studios embed region codes on their DVDs so that they work only on DVD players in particular parts of the world.[8] Similarly, makers of cameras, electronics, and cars often won't honor warranties of products purchased in a different country.

A Broader Test: Comparing the Price of a Basket of Goods. As we have already discussed, prices of individual goods may not be equated across countries for many reasons. We would like to know if these deviations from the LOP are systematic. One way to do this is to see if these individual price differences wash out when we buy a broad basket of goods. But what basket should we compare? In the U.S. the consumer price index measures the price of the basket of goods the typical U.S. consumer purchases. Similarly, many countries measure the price of a basket of goods that their consumers purchase.

There are two problems with comparing these price indexes across countries. First, they are indexes, so their level is not meaningful, and therefore, we can talk only about how prices change over time relative to one another. Second, countries do not sample the same basket of goods, so we are comparing the prices of different baskets of goods, making price

comparisons meaningless. Fortunately, there is a way around these problems.

The International Comparison Program (ICP) and the Penn World Tables (PWT) collect data that allow us to compare prices and income across countries. The ICP is a series of statistical surveys that collect prices on a representative sample of approximately 3000 goods and services. These surveys are conducted in many countries and are very careful to sample the price of very similar goods. Surveys are large projects involving each country's

> Charging different people different prices for the same good is known as price discrimination, and it is a common practice in many industries.

national statistical agency and are coordinated by the World Bank and the Organization for Economic Cooperation and Development (OECD).[9] The last survey took place from 1993 to 1996 and involved 117 countries, and it provides a useful starting point for analyzing prices and income across countries.[10]

Measuring Prices and Income. Based on the prices collected in each country, it is possible to come up with a world price for each good as a weighted average of all the prices in the world. For each country, real

---

[6] A Big Mac weighs 7.5 oz. (www.mcdonalds.com/app_controller.nutrition.index1.html) and 1 oz. equals 0.9 troy ounces. So a Big Mac weighs 6.75 troy ounces. Based on a U.S. price of $3.10, a Big Mac costs $0.46 per troy ounce compared with gold, which costs $625.01 per troy ounce.

[7] For example, by allowing children to fly for half price, airlines are engaging in price discrimination.

[8] This also occurs with video games and video consoles.

[9] The World Bank is an international organization that provides financial and technical assistance to developing countries. The OECD is a group of 30 countries committed to democracy and market economies, and this organization collects and publishes a range of economic and social statistics.

[10] For a brief overview of the ICP, refer to the article by Sultan Ahmed. A new survey is underway with almost 150 countries.

income is then calculated as the value of the goods purchased at world prices. Because each country's income is measured using the same prices, these measures of income are directly comparable across countries. The value of each country's purchases is then calculated at its own prices; this is a measure of its income at local prices. The ratio of income at local prices to income at world prices is a measure of a country's price level relative to world prices. In this way the Penn World Tables construct a measure of the price level and purchasing power (real income in each country). The procedures for measuring income and prices across countries are quite similar to how the Bureau of Economic Analysis measures income and prices in the U.S. over time.

Figure 1 presents a scatter plot of the relative price of the common basket of goods (on the y-axis) against the relative income of each country (on the x-axis). These data are from the 1996 Penn World Tables, and each point is relative to the U.S. and measured in logarithms, which means that the slope approximates the percentage change in the price level for a given percentage change in per capita GDP. There are obviously substantial differences in price level and income per capita. Turkmenistan has the lowest prices (-2.18, or 11 percent of the U.S. level), while Switzerland has the highest prices (0.53, or 170 percent of the U.S. price level). Tanzania has the lowest income per capita (-4.12, or 1.6 percent of the U.S. level), and Luxembourg has the highest income per capita (0.18, or 120 percent of the U.S. level).

From Figure 1 we see that there is a positive relationship between prices and income. As we saw with Big Macs and Switzerland and China, the countries with the highest income also pay the highest prices for a broad range

of goods. A measure of the strength of this relation can be found by estimating how much relative prices increase with relative income. The results of this estimate are reported in the lower right corner of Figure 1. We find that a doubling of income per capita is associated with a 43 percent increase in the price level.

The differences in price levels and income per capita are quite persistent over time. For instance, of the 32 countries with price levels one-half of those in the U.S. in 1996 for which we also have data on price levels in 1985, 26 also had price levels less than half of those in the U.S. in 1985.

EXPLAINING THE PRICE-INCOME RELATIONSHIP

Economists tend to attribute the price-income relationship seen in Figure 1 to either differences in the prices of tradables (those goods that are either traded frequently or easy to trade) or differences in the prices of nontradables, those goods that are both costly and infrequently traded across countries.[11] We will discuss an explanation for the price-income relationship based on deviations from the LOP in tradables. (An alternative, complementary explanation based on deviations from the LOP in nontradables is presented in *Another Theory to Explain the High Prices in High-Income Countries.*)

Examining the role of prices for tradables for the relationship seen in Figure 1 requires a measure of the price of tradable goods. Fortunately, the ICP contains prices for over 3000 goods, so we can compare the price of a basket of those goods that are traded frequently across countries. Examples of the types of goods classified as

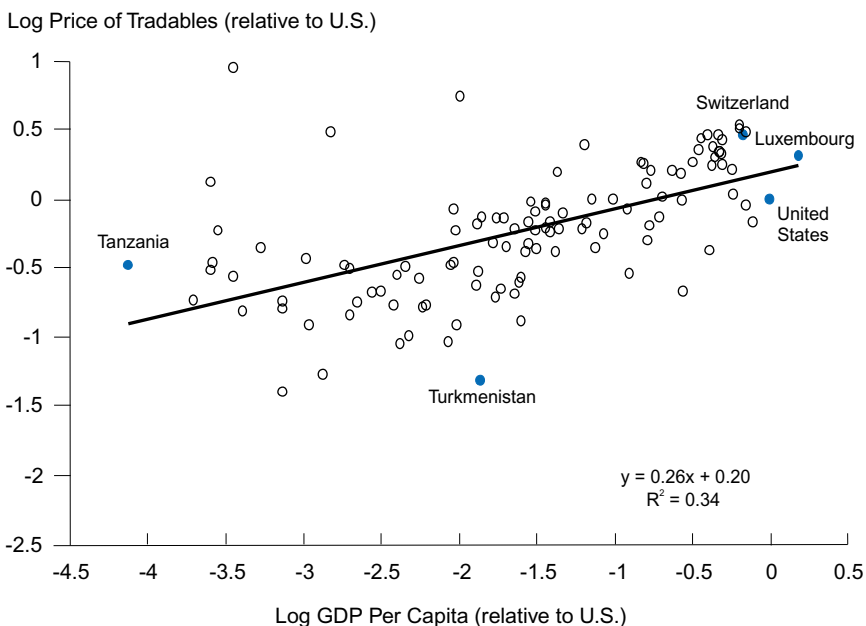tradable are machinery and equipment, tobacco, alcohol, and personal transportation equipment.

Figure 2 shows the relationship between the price of a basket of tradable goods and income per capita. Similar to what we saw in Figure 1, there is a positive relationship between the price of this tradable basket and income. In the lower right-hand corner of Figure 2, we estimate that a doubling of income per capita is associated with a 26 percent increase in the price of tradable goods. Comparing our measures of price differences of tradable goods to the measure of price differences for all goods, we find that differences in prices for tradables account for about 60 percent of the aggregate price-income relationship.[12]

One possible explanation for the positive relationship between prices for tradables and income is that prices for tradable goods include some nontradable inputs, which are cheaper in low-income countries. For instance, the price of a car includes the cost of transporting it from the factory to the dealership as well as the costs the car dealership incurs in selling the car. The costs in getting products to the consumer, essentially retail and wholesale distribution, are mostly nontradable and contribute to price differences across countries. If retail and wholesale distribution services are cheaper in low-income countries,[13] that

---

[11] Much of this section is based on our working paper.

[12] In our paper, we show that for certain commonly used price indexes, the contribution of differences in the prices of tradables to the relationship between price levels and income can be measured by comparing the coefficient from the regression of prices for tradables on income to the coefficient from the regression of price levels on income.

[13] Obviously, wholesale and retail distribution also includes some tradable inputs, such as trucks, airplanes, and fuel, which would tend to make their prices similar across countries.

FIGURE 1

Prices and GDP Per Capita

Log Price Level (relative to U.S.)

$$y = 0.43x - 0.02$$
$$R^2 = 0.49$$

Log GDP Per Capita (relative to U.S.)
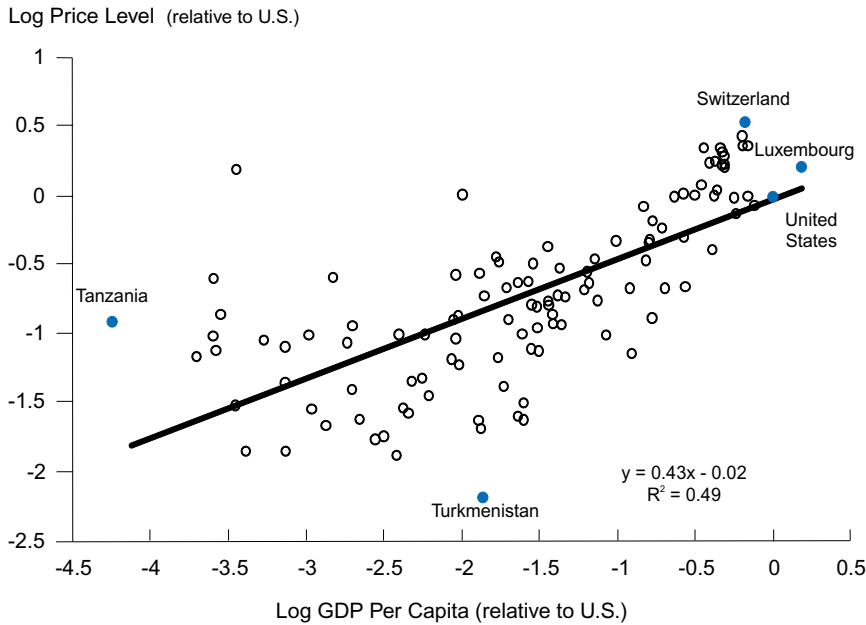
FIGURE 2

Price of Tradables and GDP Per Capita

Log Price of Tradables (relative to U.S.)

$$y = 0.26x + 0.20$$
$$R^2 = 0.34$$

Log GDP Per Capita (relative to U.S.)

may explain why prices for tradables are lower in low-income countries.

To isolate the source of differences in prices for tradables, we must compare the price of goods before these retail and wholesale distribution services are added. One way of doing this is to measure the price of goods as they leave the U.S. and are being shipped to different destination markets.

Measuring U.S. Export Prices at the Border. Destination-specific export prices can be constructed using data collected from shippers' export declaration forms. These are forms filed with Customs for every shipment of goods that leaves the U.S.[14] For each good, there are data, by destination country, on the average price of all shipments in each year from 1989 to 2000. These prices are measured at the U.S. border or the shipping dock before any taxes or nontradable services are added. Goods are classified according to the Harmonized Commodity Description and Coding System (HS). This is a system of names and numbers for classifying traded products.[15] The data cover 10,471 goods.

We focus on shipments to OECD countries plus some low-income countries for which we also have wage data. The complete list of countries

---

[14] These forms aren't necessary for shipments with values below $2000. These small-value shipments account for a very small share of U.S. exports.

[15] The HS system is an international classification system based on broad six-digit categories. Many countries classify traded goods in more detail. For example, the U.S. defines products using 10-digit HS codes. Export codes (which the U.S. calls Schedule B) are administered by the U.S. Census Bureau. In this system, 10-digit goods can be incorporated into nine-digit goods, nine digits into eight digits, and so on. The U.S. Census Bureau's website offers this example: Concentrated frozen apple juice is assigned a 10-digit number, but this product can be included in the broader six-digit category described as apple juice, which, in turn, can be incorporated into the broader four-digit category, fruit juices and vegetable juices, and so on.

A number of explanations of the aggregate price-income relationship attribute it to deviations from the law of one price (LOP) in nontraded goods. Recall that nontraded goods are those goods that have high international shipping costs and thus are infrequently traded across countries; such goods include haircuts, restaurant meals, housing, and medical services.

To get an idea of how the price of nontraded goods differs with income, we plot in the Figure the relative price of nontraded goods to traded goods against real income per capita. By looking at how the ratio of nontraded to traded prices differs with income, we can isolate anything that affects nontradables separately from tradables. As we saw with the relationship between aggregate price levels and incomes, we find that the ratio of nontradable prices to tradable prices also rises with income. In fact, a doubling of income is associated with a 34 percent increase in the relative price of nontradables.

There are a variety of competing explanations of this observation. The most common is known as the Balassa-Samuelson theory.[a] It contains two main elements. First, the theory assumes that the LOP holds in tradables.

_____

[a] For alternative explanations of this observation, see the work by William Baumol and William Bowen; Irving Kravis and Robert Lipsey; and Jagdish Bhagwati.

Second, it assumes that across countries, there are much larger differences in the productivity of workers producing tradable goods than nontradable goods. Since the LOP holds for tradable goods, the cost of producing tradables must be the same everywhere. This means that international wage differences are determined by differences in labor productivity in traded goods and are quite large. With large wage differences across countries and relatively small differences in labor productivity in nontradables, prices for nontradables will differ substantially across countries and will be higher in high-wage/high-income countries.

A simple two-country, two-goods example might help to explain how the theory works. Suppose the two countries, call them Richland and Poorland, can make cars, which can be freely traded, and haircuts, which are impossible to trade. The table below describes the productivity of workers in each country. Starting with case 1, we see that in Poorland, one worker can produce either one car or one haircut per day, while the typical worker in Richland is more productive and can produce either four cars or two haircuts per day. To keep things simple, suppose that workers in both countries get paid in dollars and that the daily wage in Poorland is $1.

Given that a worker in Poorland earns $1 per day and can produce one car per day, the price of a car must be $1 everywhere, since cars can be freely traded. Now, since Richland workers can produce four cars a day, they will earn $4 per day. With these wages, the price of haircuts

TABLE

A Two-Country Example of the Balassa-Samuelson Model

	# of units produced per worker			Prices		Price Level	Real Wage
Case 1	Cars	Haircuts	Wages	Cars	Haircuts		
Poorland	1	1	1	1	1	2	0.50
Richland	4	2	4	1	2	3	1.33
Case 2							
Poorland	1	0.5	1	1	2	3	0.33
Richland	4	2	4	1	2	3	1.33

will be $1 in Poorland, since a worker earning $1 can give one per day, while in Richland a haircut will cost $2, since it takes a worker earning $4 a half a day.

To see how prices vary with real income, we must define the bundle of consumption goods. Let's suppose that the typical basket of goods is composed of one car and one haircut. Given the prices for individual goods, this basket will cost $2 in Poorland and $3 in Richland. We can use these prices to get a measure of real wages in each country as the wage divided by the price. So notice that real income is 50 cents in Poorland and $1.33 in Richland. Clearly, then, the higher price country, Richland, also has a higher real income, as in the data.

To see how prices and income depend on productivity in each sector, let's look at case 2 in the table. In this case, workers in Poorland are one-quarter as productive as workers in Richland for both goods. A Poorland worker still gets a daily wage of $1 and produces a car a day, so the price of a car is $1 and the price is $1 everywhere, since cars are freely traded. The price of a haircut will be $2, since it now takes two days to produce a haircut. In this case, the price level in Poorland will rise to $3, and the real wage will fall to 33 cents, while it is the same as case 1 in Richland. The price level is now the same across the two countries, and there is no positive relationship between prices and income. Thus, to get a positive relationship between prices and income, it is necessary for low-income countries to be relatively productive in producing nontradables compared to high-income countries.

There are three reasons to question the Balassa-Samuelson theory as a complete explanation of the aggregate price-income relationship. First, as we have seen, there are large deviations from the LOP in tradable goods. Second, the Balassa-Samuelson theory requires relatively large differences in the efficiency of producing trad-

able goods compared to nontradable goods across rich and poor countries. While we don't have good measures of these productivity differences across countries, we do have good measures from the U.S. Using data on sectoral labor productivity growth in the U.S. from a paper by Dale Jorgenson and Kevin Stiroh, we find that labor productivity in the nontradables sector has grown by about two-thirds as much as labor productivity in the tradables sector. Finally, for nontradable goods to explain the aggregate price-income relationship, the nontradables price-income relationship must be much stronger than the aggregate price-income relationship.[b] Comparing Figure 1 in the text and the figure in this box, we see that this is not the case.

---

[b] Notice from our first case that while the price of haircuts is twice as high in Richland, the price level is only 50 percent higher. The weaker relationship between aggregate prices and income is due to nontradables' accounting for only a part of the final basket of goods. In their paper, Alan Stockman and Linda Tesar measure the size of the tradables sector in OECD economies and find that it accounts for about one-half of the economy. This implies that the relationship between prices for nontradables and income per capita needs to be twice the relationship between the full basket of goods and income to explain the data.

FIGURE

Relative Price of Nontradables and GDP Per Capita
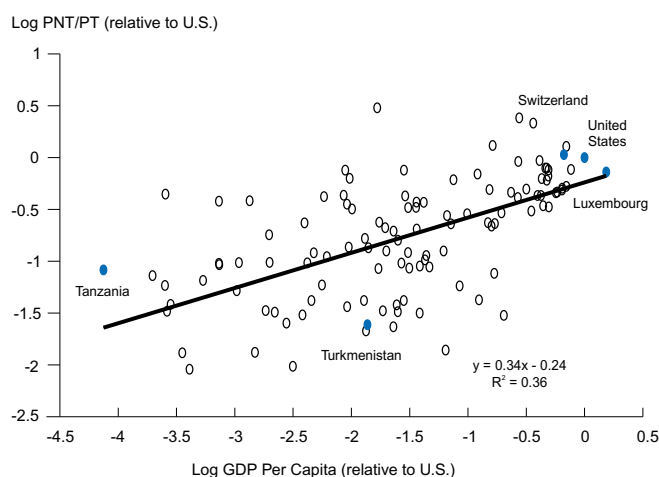
can be found at the bottom of Table 3. Overall, there are almost 1.2 million observations, where an observation is a particular good sold to a country in a particular year, accounting for about 75 percent of the value of U.S. trade in goods over the period.

We can use these data to ask whether, on average, goods being shipped to markets with relatively high income tend to be sold for relatively high prices (a description of the empirical specification can be found in the footnotes to Table 3).[16] The results of our analysis using these data on export prices confirm what we found using retail prices for tradables from the Penn World Tables: Prices for tradables increase as income per capita increases. Moreover, in export prices, this effect is about two-thirds as strong as that for retail prices for tradables in the Penn World Tables. This finding suggests that differences in the factory prices of tradables account for about 40 percent of the differences in retail price levels across countries, while wholesale and retail margins account for about 20 percent.[17]

For the most part, then, the evidence points to retail prices for tradables being higher in high-income countries because exporters sell these goods at higher prices in these countries.

Digging a Little Deeper into Export Prices. Even though we perform the analysis using data that have been broken down into subcategories, one might suspect that the price differences uncovered may be related to differences in the quality of the products being sold. For instance, it could be that a 10-digit category contains different quality goods, say, a high-quality 11-digit good and a low-quality 11-digit good, and that high-income countries purchase relatively more of the high-quality good. While this idea can't be directly tested for goods classified at the 10- and 11-digit levels, we can see if this is happening at broader levels of classification. For instance, we can compare the price-income relationship on 10-digit goods to the same goods classified at the nine-digit level. If rich countries purchase relatively more of the high-quality, more expensive 10-digit goods, we should find that the relationship between prices and income is stronger at the nine-digit level. We actually find the opposite and conclude that quality differences do not explain the differences in export prices by destination.

To get at the source of international price differences, we next examine the association between export prices and the real wage in the destination market. Not surprisingly, since high-income countries also tend to have high wages, we find a strong positive relationship with wages in destination countries (see the column labeled "Wages only"). However, wages and income per capita are not perfectly correlated, since there are differences in labor force participation, hours worked, capital income, and taxes across countries. When we examine the independent effect of income and

wages on export prices, in the last two columns of Table 3, we find that wages explain export prices by destination. This leads us to conclude that high prices are associated with high wages.[18]

Since this is a big data set, we can dig a little deeper. We next examine the relationship between export prices and destination characteristics for different types of goods. This analysis can be found in the bottom seven rows of Table 3. We find that export prices increase more with wages for

> For the most part, the evidence points to retail prices for tradables being higher in high-income countries because exporters sell these goods at higher prices in these countries.

consumption goods than for capital goods, industrial supplies, autos, and a range of other products. We also find that the price of medicinal products tends to be most affected by the wage and income in the destination market. Finally, notice that when we control for wages and income per capita in the final two columns, for each type of good we find that wages are always positively associated with prices, while income per capita may have a negative or positive association with prices.

The analysis of export prices tells us three things. First, high prices for tradables are largely due to exporters' charging high prices as goods leave the country. Second, export prices are more strongly related to wages in the destination market than income per capita. Third, this effect is stronger for consumer goods than industrial supplies or capital goods.

[16] Here we are looking at prices for individual goods rather than baskets of goods. Since not all goods are exported to all countries, we cannot construct a representative basket as in the previous analysis.

[17] Recall that differences in the prices of tradables account for 60 percent of the difference in price levels. Since differences in export prices account for two-thirds of the differences in prices for tradables, we can conclude that differences in export prices account for 40 percent of the difference in price levels.

[18] These results hold even after controlling for a wide range of factors such as trade costs, the share of intra-firm trade, and the level of intellectual property protection.

TABLE 3

Export Prices and GDP Per Capita and Wages*
(t-statistics in parenthesis)

	Number of Obs.	Fraction of Total Value Exported	GDP Per Capita Only	Wages Only	Both Together GDP Per Capita	Both Together Wages
All Goods**	1,177,803	0.751	0.170 (64.3)	0.162 (86.2)	0.012 (3.25)	0.156 (57.5)
Consumption Goods	228,074	0.085	0.236 (39.6)	0.218 (51.9)	0.036 (4.3)	0.200 (33.8)
Food/Feed/Beverages	109,646	0.078	0.156 (31.7)	0.091 (26.5)	0.128 (18.2)	0.027 (5.6)
Capital Goods	322,105	0.248	0.087 (14.5)	0.146 (33.5)	-0.126 (14.4)	0.213 (33.4)
Industrial Supplies	484,661	0.247	0.201 (52.9)	0.168 (63.4)	0.063 (11.4)	0.136 (34.9)
Autos	25,694	0.082	0.158 (8.84)	0.113 (8.9)	0.090 (3.5)	0.066 (3.6)
Agricultural Goods	61,991	0.044	0.140 (21.1)	0.077 (16.7)	0.128 (13.0)	0.012 (1.7)
Medicine	15,859	0.014	0.187 (6.4)	0.282 (13.2)	-0.201 (4.7)	0.390 (12.4)

* Income per capita and wages are measured in real terms using price deflators in the Penn World Tables.
The table reports the relationship between export prices and the characteristics of the export destination from a regression of export prices on the characteristics of the export destination. The regression takes the form: $p_{ijt} = a_{it} + b_1 \ast y_{jt} + b_2 \ast w_{jt} + e_{ijt}$, where p_{ijt} measures the logarithm of the price of good i sold to country j at time t. In country j at time t, income per capita is measured as y_{jt}, and the hourly manufacturing wage is measured as w_{jt}. The term e_{ijt} accounts for errors. The term a_{it} is a dummy variable that accounts for good-specific attributes, such as marginal cost. To explore the relationship between destination prices and just income per capita (or wages), we can run the regression constraining $b_2=0$ ($b_1=0$). The equation can be estimated by ordinary least squares. We construct White robust standard errors that allow for heteroskedasticity in e_{ijt} and also allow for country-year clustering.

** Countries include Australia, Austria, Belgium-Luxembourg, Brazil, Canada, Denmark, Finland, France, Germany, Greece, Hong Kong, Ireland, Israel, Italy, Japan, South Korea, Mexico, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sri Lanka, Sweden, Switzerland, and the total U.S. exports of these goods.

Explaining the Export Price-Income Relationship. The export data confirm that exporters ship goods at lower prices to low-income locations. As we have already discussed, this type of price discrimination by destination market is possible only if trade barriers make it difficult for other firms to arbitrage these destination-specific prices away. Given that we see these price differences, it must be the case that arbitrage is limited, so prices are determined by either differences in competitive environments or consumers' tastes for particular goods, or a combination of the two.

In a recent paper, we develop a theory of price discrimination that can generate a positive relationship between export prices and wages. It builds on our second and third findings from studying export prices: Wages seem to matter most for export prices, and export prices increase with

wages more for consumer goods. The key idea of the theory is that searching to find goods at a really low price doesn't pay. We mean this literally. Since searching takes time away from working, the higher one's wage, the more costly it is to search repeatedly to find a good at a lower price. For this reason, high-wage individuals will be willing to accept higher prices than low-wage individuals and will pay more, on average.

To be a little clearer, our theory assumes that consumers do not know where to buy goods at the lowest price and must spend some time searching for goods. This is a theory of a cost that limits arbitrage, the time it takes to search, and is consistent with everyone's experience of finding the same good selling for different prices in different stores. It is also consistent with consumers' trade-off of paying a higher price at a local store to save time rather than traveling to a store farther away that sells goods at lower prices. As individuals search, they find goods at some price and must decide whether to accept a store's price or continue to search. Because search takes time away from work, consumers consider the forgone labor income of continuing to search, so the consumer's wage determines which prices the consumer will accept. Individuals with higher wages have a higher opportunity cost of time and therefore are willing to accept higher prices rather than search repeatedly.

Firms, knowing consumers' purchasing behavior, will charge higher prices in markets where it is more costly for the average consumer to search repeatedly. This implies that prices are higher in high-wage locations. Now, as long as the time it takes to shop is not so different between high- and low-income countries, low-income countries will have a comparative advantage in search, so prices will

be lower in low-income countries. This is a natural extension of the Balassa-Samuelson mechanism described in the box on page 6.

The theory developed here also tells us something about the source of income differences across countries. In this model, countries with more productive workers will earn higher wages and be willing to pay higher prices for all goods, both tradables and nontradables. In contrast, in the Balassa-Samuelson theory, for prices to rise with wages, high-wage countries must be relatively more productive at producing tradables than nontradables. Thus, the Balassa-Samuelson theory requires that cross-country productivity differences in tradables be much larger than productivity differences in nontradables. The Balassa-Samuelson theory suggests that countries mainly become richer by becoming better at producing tradables, while our theory suggests a more balanced approach to growth in which workers in a country become better at producing everything.

Evidence on Shopping Time, Prices, and Income Per Capita. The theory we have described also implies a relationship between wages, shopping time per purchase, and prices. There is some evidence that these variables

are related based on time-use surveys, which are studies in which respondents are asked to track their every activity in small time increments over the course of a day or week. Examples of activities tracked are sleeping, eating, working, commuting to work, shopping, traveling to shopping, and listening to the radio.

Two recent papers use time-use survey data to confirm a positive relationship between wages and prices paid and a negative relationship between wages and time spent shopping predicted by our theory. Using time-use data from the U.S., economists Mark Aguiar and Erik Hurst find that when people retire, and the opportunity cost of their time declines, they spend more time shopping per purchase and tend to pay less per unit purchased. Likewise, using time-use data from Argentina, David McKenzie and Ernesto Schargrodsky find that higher income individuals spend less time shopping per purchase and pay higher prices, on average (Table 4). In fact, shopping time per expenditure of people in the lowest income quartile is about 80 percent higher than that of people in the highest income quartile. Moreover, after the economic crisis in 2001, which lowered all Argentineans'

TABLE 4

Shopping Frequency Per Real Expenditure by Income in Argentina

		Households by Income Quartile			
	All	Lowest	2nd	3rd	Highest
Pre-Crisis	0.24	0.29	0.26	0.21	0.16
Post-Crisis	0.28*	0.35*	0.31*	0.26*	0.20*

* Statistically different from 2001 mean at 1 percent level.
Shopping frequency measures time spent shopping.
From Table 3: McKenzie and Schargrodsky (2005)

real income, shopping time increased by about 25 percent per expenditure across all income levels.

We can also compare results of time-use surveys in different countries to get an idea of how shopping time differs by income per capita. Figure 3 presents a scatter plot of time spent per purchase against income per capita based on data collected from countries that participated in the European Harmonized Time Use Survey. The data show that shopping time per purchase tends to fall with income per capita, so that in low-income countries people tend to search more intensively than in high-income countries. As we have already seen, prices and wages tend to rise with income per capita. Thus, both the within-country evidence and the cross-country evidence are consistent with the model we have described.

SUMMARY

There are large differences in prices across countries that are related to income per capita. On average, the cost of a basket of goods tends to be relatively high in high-income countries. These price differences exist both for goods that are easily and frequently traded and those goods that are not traded. Moreover, these price differences show up at the dock, so that export prices to high-income countries tend to be higher than export prices to low-income countries.

Understanding the determinants of the price-income relationship sheds light on the source of the large differ-

FIGURE 3

Time Spent Shopping Per Purchase

Log Shopping Time Per Expenditure (relative to UK)

$y = 0.80x - 0.11$
$R^2 = 0.90$

Log Real GDP Per Capita (relative to UK)
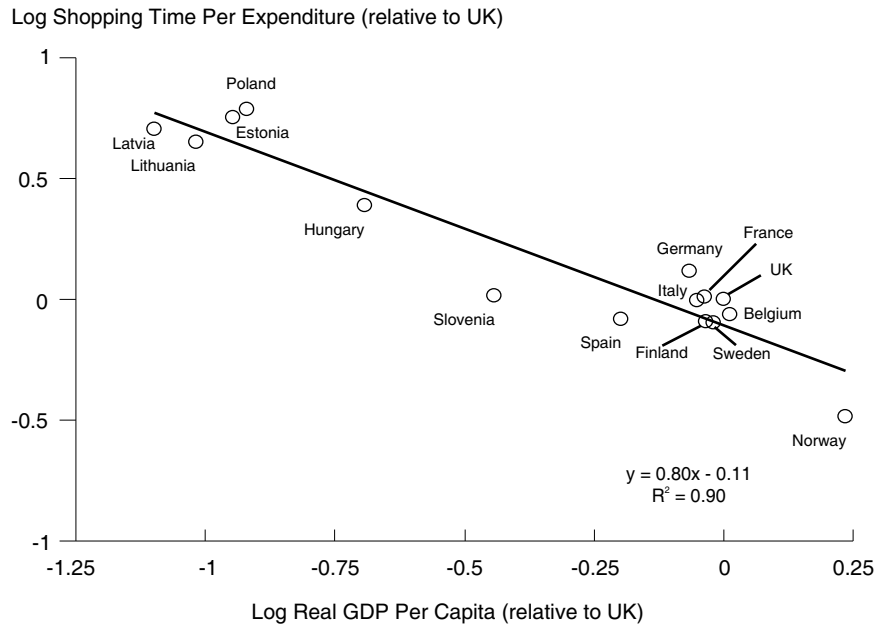
Data: European Harmonized Time Use Survey and Eurostat (2002)

ences in income and well-being across countries. Traditional models of these price differences have focused on differences in prices of nontradable goods and thus attributed income differences largely to differences in productivity in the tradables sector. The evidence presented here that price differences are quite large in the tradables sector as well suggests a more balanced view of productivity differences across sectors and countries. The large price differences in tradable goods suggest that policymakers should target improving efficiency across the entire economy and not just in the tradables sector.

The discussion has purposely avoided nominal exchange rates. However, a good theory of price levels across countries is also useful as a long-run theory of nominal exchange rates. It provides a natural benchmark for determining whether a currency is overvalued or undervalued. For countries that actively manage their exchange rate, this may be a useful guide in determining an appropriate target level. BR

REFERENCES

Aguiar, Mark, and Erik Hurst. "Lifecycle Prices and Production," mimeo, University of Chicago (2005).

Ahmed, Sultan. "Historical Overview of the International Comparison Program," available at http://siteresources.worldbank.org/ICPINT/Resources/historical_overview-Icp_SA.doc.

Alessandria, George, and Joseph Kaboski. "Violating Purchasing Power Parity," Philadelphia Federal Reserve Bank Working Paper 04-19 (2004).

Balassa, Bela. "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, 72 (1964), pp. 244-67.

Baumol, William, and William Bowen. *Performing Arts: The Economic Dilemma*. New York: The Twentieth Century Fund (1966).

Bhagwati, Jagdish. "Why Are Services Cheaper in Poor Countries?," *Economic Journal*, 94 (1984), pp. 279-86.

Jorgenson, Dale W., and Kevin J. Stiroh. "U.S. Economic Growth at the Industry Level," *American Economic Review*, 90 (December 2000), pp. 161-67.

Kravis, Irving, and Robert Lipsey. *Toward an Explanation of National Price Levels*. Princeton Studies in International Finance, No 52, Princeton University (1983).

McKenzie, David, and Ernesto Schargrodsky. "Buying Less, but Shopping More: Changes in Consumption Patterns During a Crisis," BREAD Working Paper 092 (February 2005).

Ostapik, Edith, and Kei-Mu Yi. "International Trade: Why We Don't Have More of It," Federal Reserve Bank of Philadelphia *Business Review* (Third Quarter 2007).

Samuelson, Paul A. "Theoretical Notes on Trade Problems," *Review of Economics and Statistics*, 46 (1964), pp. 145-54.

Stockman, Alan, and Linda Tesar. "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements," *American Economic Review*, 85(March 1995), pp. 168-85.

Liquidity Crises*

BY RONEL ELUL

Financial markets have experienced several episodes of "liquidity crises" over the past 20 years. One prominent example is the collapse of the Long Term Capital Management hedge fund in 1998. The recent market disruption brought about by the downturn in subprime mortgages also shares many features with liquidity crises. What is liquidity? Why does it sometimes seem that the market's supply of it is insufficient? Can anything be done about it? In this article, Ronel Elul outlines some theories of market liquidity provision, how it breaks down in times of crisis, and some possible government responses.

Over the past 20 years, financial markets have experienced several episodes of "liquidity crises." Among these are the 1998 collapse of the Long Term Capital Management hedge fund and the disruption in financial markets that began in the summer of 2007, sparked by the downturn in subprime mortgage markets.

In many of these cases, the market's supply of liquidity seemed to be insufficient, and moreover, liquidity does not always appear to be allocated to those who need it most. Lack of liquidity also sometimes forces "fire sales," actions that, in turn, push down asset prices, thus making liquidity problems worse. Economists have sought to understand the nature of market liquidity provision, how it breaks down in times of crisis, and possible government responses.[1]

ANATOMY OF A LIQUIDITY CRISIS

What Is Liquidity? One author has pointed out that "liquidity, like pornography, is easily recognized but not so easily defined."[2] For understanding liquidity crises, however, it may be useful to think of liquidity as the ease of selling an asset at its "true," or fundamental, value. This fundamental value may be defined as the present value of the asset's future cash flows. Alternatively, liquidity can be viewed as the extent to which it is possible for the holder of an asset to borrow against these future cash flows.

The Collapse of LTCM. The events of the summer and fall of 1998 provide an illustration of many of the main features of liquidity crises. These events revolve around the collapse of the Long Term Capital Management (LTCM) hedge fund.[3]

During the summer of 1998, LTCM took large losses on many of its trades; these losses were intensified when Salomon Smith Barney's arbitrage group, which had positions very similar to LTCM's, was broken up and its positions liquidated. But LTCM's position became much more precarious on August 17, 1998, when the Russian government devalued the ruble and declared a moratorium on repaying 281 billion rubles ($13.5 billion) of its Treasury debt. The fact that the IMF had allowed a major economy to default shocked the markets.[4]

---

[1] I use the term "government intervention" broadly. In principle, this might include fiscal policy or central bank monetary policy. In this paper, I will focus on monetary policy.

---

*The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

[2] See the book by Maureen O'Hara.

[3] Much of this account is drawn from Roger Lowenstein's book.

[4] In addition, a further surprise occurred when Russian banks and securities firms exercised *force majeure* clauses and refused to honor the derivatives contracts they had sold to foreign customers. These clauses, which are common in many contracts, are intended to free a party from liability when an extraordinary event prevents him from fulfilling his obligation.

Ronel Elul is a senior economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/econ/br/.

Business Review Q2 2008 13

LTCM had indeed invested in Russian bonds and lost money following this default. However, the resulting *flight to quality* had an even bigger effect on the value of LTCM's portfolio. Investors who had become nervous as a result of these events pulled out of risky assets and rushed to assets considered safe. For example, the yield on the 30-year U.S. Treasury bond (a safe security) fell to its lowest level up to that time. Many of LTCM's strategies had involved betting that the spread between safe and risky assets would actually decline; thus, the flight to quality caused it to lose substantially more. Finally, in addition to a flight to quality in security markets, there was a broad-based drying up of liquidity as banks chose to preserve their liquidity and cut back on lending.[5]

As a result of declines in prices on the risky assets in its portfolio, LTCM breached collateral agreements with its lenders and was forced to sell assets to meet these margin calls.[6] These asset sales had ramifications for other markets and other hedge funds. Mark Mitchell, Lasse Pedersen, and Todd Pulvino recount an example: "When LTCM incurred large losses on macroeconomic bets, the firm was forced to liquidate large convertible bond positions.[7] These sales led to depressed valuations of convertible bonds despite

the fact there was little change in overall fundamentals.[8] As a result, other hedge funds incurred large losses and were also forced to sell their convertible bond holdings." The authors show further that prices of convertible bonds fell far below their "fair" value, as calculated by mathematical models.[9]

Because of concerns that the forced liquidation of LTCM's huge

> The events of the summer and fall of 1998 provide an illustration of many of the main features of liquidity crises.

portfolio would cause further upheaval in financial markets, the Federal Reserve helped coordinate a private-sector bailout of the fund in September 1998.[10] The Fed also cut its fed funds rate target by 75 basis points during the fall of 1998, in part because of concerns that financial market turmoil

might spill over to the real side of the economy.

From this account we can identify several key features of liquidity crises. The apparent trigger for the crisis was an unexpected event that called long-standing models into question. Lenders responded by cutting back on providing liquidity. The effect of the crisis was to push prices below their fundamental, or fair, value. More precisely, the prices of risky assets fell, while those of assets perceived to be safe rose; that is, there was a flight to quality. There was *commonality of illiquidity* — problems spilled over from one market to another. A *liquidity spiral* was created: These falling prices caused margin requirements to be breached, thus leading to asset sales, which then led to further drops in prices and thus to further losses, and so on. Government intervention played a role in resolving the crisis.

The Current Financial Market Turmoil. Many of these features are also present in the disruption in financial markets that began in the summer of 2007, sparked by the downturn in subprime mortgage markets. Not surprisingly, the sharp increase in default rates on mortgages called into question models of subprime mortgage credit quality (as well as lenders' underwriting standards). There was also a flight to quality — for example, the premium paid by high-quality (AAA-rated) corporate borrowers over U.S. Treasury bonds nearly doubled in the summer of 2007 (Figure 1). In this case, market participants suddenly demanded much more compensation to bear even a small amount of risk. The cutback in the provision of private-sector liquidity was even more dramatic than in the case of LTCM. This may be seen most strikingly in the interbank market; the London InterBank Offered Rate (LIBOR) that banks charge one another in the London interbank market

---

[5] See, for example, the September 1998 Federal Reserve Senior Loan Officer Opinion Survey: http://www.federalreserve.gov/boarddocs/snloansurvey/199810/default.htm.

[6] Like most hedge funds, LTCM had borrowed heavily to finance its portfolio and borrowing allowed it to generate higher returns per dollar of outside investment. However, LTCM's lenders required that the value of these assets, which served as collateral to secure its loans, stay above a certain minimum *margin requirement*. When the prices of its assets fell, these collateral agreements were breached, and lenders issued *margin calls*. This required LTCM to come up with additional cash or securities in order to avoid the forced liquidation of its portfolio.

[7] A convertible bond is a type of bond that can be converted into shares of stock in the issuing company, usually at some pre-announced ratio. Hedge funds are significant traders of convertible bonds, as part of a popular strategy known as convertible arbitrage.

[8] In this example, prices are below their fair value because of binding collateral constraints. However, another reason that prices may fall in a crisis is that one side of the market has more information than the other, and thus, asset sales may be interpreted as negative information about fundamentals. For a similar model motivated by the 1987 stock market crash, see the paper by Gerard Gennotte and Hayne Leland.
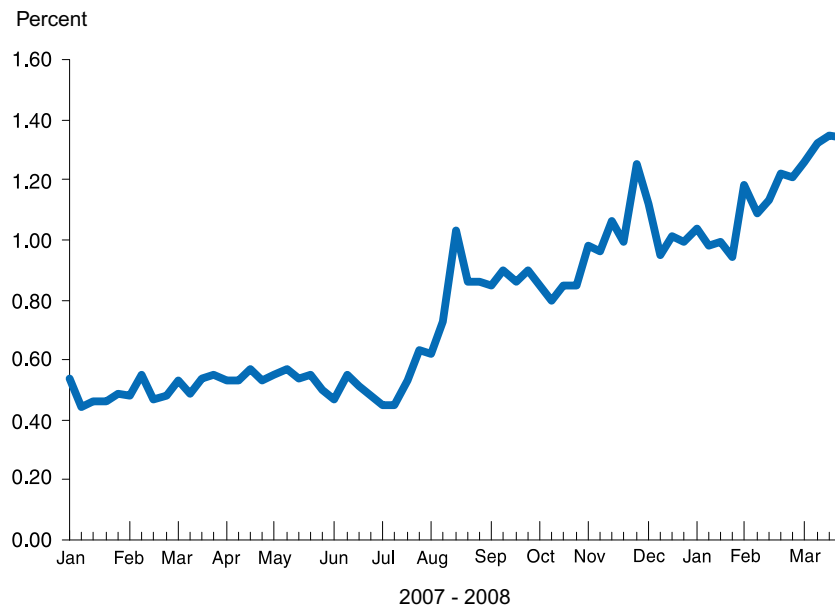
[9] The fair value of the convertible bond is calculated by using an option valuation model; these models are extensions of the well-known Black-Scholes pricing formula.

[10] Although it will not be discussed here, the economic rationale behind the Fed's coordinating role is also of interest; see the paper by Stephen Morris and Hyun Shin.

FIGURE 1

Five-Year AAA Corporate Bond Rate —
Five-Year Treasury Bond Rate

Percent, Jan Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan Feb Mar, 2007 - 2008

shot up relative to the baseline U.S. Treasury bill rate as banks sought to conserve their scarce liquidity. The problems were particularly pronounced in term (that is, not overnight) interbank markets (Figure 2). These events were widely understood in the popular press as reflecting *liquidity hoarding.* Finally, the Federal Reserve and other central banks intervened in several different ways. (See *Federal Reserve Responses to Recent Problems in Interbank Markets.*)

PRIVATE MARKETS MAY PROVIDE TOO LITTLE LIQUIDITY

One central feature of these episodes is the inadequacy of the private market's provision of liquidity. In studying this issue, Bengt Holmstrom and Jean Tirole explore various means by which firms may obtain liquidity and show that the private market may

not always be able to provide adequate liquidity on its own. They then consider possible government responses.

They consider firms that have long-lived projects, for example, manufacturing plants that can produce a good for several years before becoming obsolete. These firms may experience a "liquidity shock," a sudden need for funds to keep the project going. This could be due, for example, to an unanticipated, temporary shortfall in sales, so that internal funds that were previously used to keep projects going are no longer available. But if these funds are not available, the firms' assets must be liquidated immediately, at a loss. How can the firms obtain enough liquidity to continue their projects?

These projects are still profitable, so one might think that a firm that has been hit by a shock could simply borrow against its future project returns. But lenders may be unwilling to offer

sufficient funds to the firm because the greater the firm's required debt payments, the smaller the firm's own share of the returns from the project. This means that the firm has less incentive to exert enough effort to ensure that the project succeeds.

To guard against this risk of illiquidity, a firm might hold cash or other safe assets, such as Treasury securities, that can be sold in case it experiences a shock. Since these assets are safe, the firm can always sell them to raise funds. But Holmstrom and Tirole also show that this is not generally an ideal way to allocate scarce liquidity because lucky firms that do not experience a shock will be left with assets they do not need, while unlucky firms have no way to gain access to those assets.

What is needed instead is some way for firms to obtain insurance against unexpected liquidity needs. This can be facilitated through a financial intermediary that can offer lines of credit to firms, which they draw upon only if they experience a shock. In effect, the financial intermediary takes stakes in all of the firms' future returns and lends only to those firms that have been hit by a shock. When liquidity shocks are *idiosyncratic* — that is, the shock hits only a few firms at once — Holmstrom and Tirole show that this is indeed the best way to provide liquidity to the private sector.

However, in a liquidity crisis, in which the liquidity shock is an *aggregate* one (that is, it hits many firms at once — for example, a recession that hits all firms' sales), the private market is not able to meet each firm's liquidity needs. The reason is that firms' aggregate demand for liquidity will exceed the private sector's ability to meet this need. In this case, there is scope for the government to provide liquidity in times of crisis. The government is able to commit to providing liquidity when

FIGURE 2

Three-Month LIBOR — Three-Month T-Bill Rate

Percent

2007 - 2008
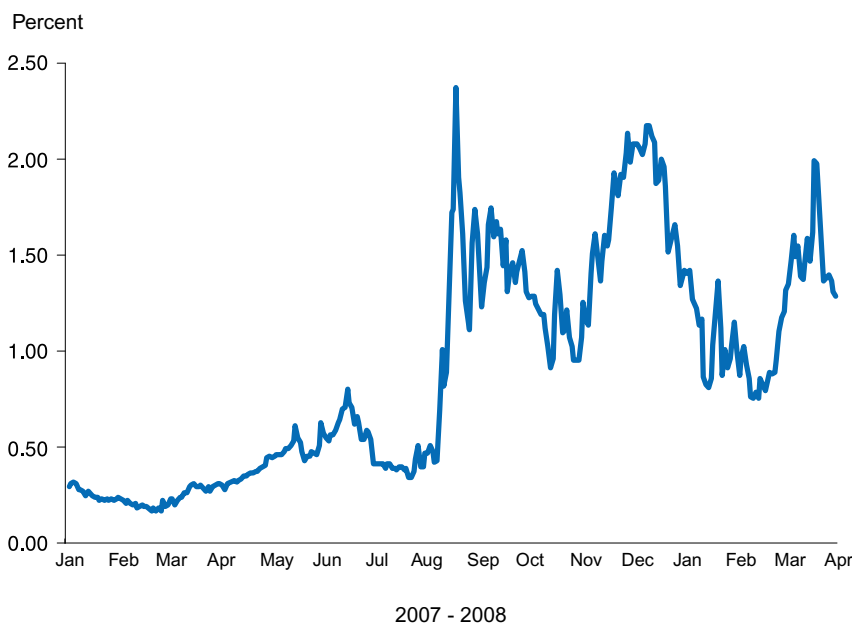
the private market can't, either by taxing consumers or by printing money.

Holmstrom and Tirole suggest that this intervention may take many forms. For example, it could take the form of government securities that pay off only in the event of a particular aggregate liquidity shock. Sundaresan and Wang document this in connection with the run-up to Y2K. They show that privately supplied liquidity dried up as the millennium approached. In response, the Federal Reserve intervened by issuing options on the fed funds rate.[11] Alternatively,

Holmstrom and Tirole suggest that monetary policy could serve this role by easing financing conditions in times of crisis.

VERY UNLIKELY CONTINGENCIES CAN AFFECT THE AVAILABILITY OF LIQUIDITY

While Holmstrom and Tirole focus on the lack of sufficient liquidity in the private sector as a rationale for government intervention, another feature of some liquidity crises is that what liquidity is available is not efficiently allocated. That is, liquidity is not allocated to those who need it the most. The reason is that the liquidity crisis may make market participants overly concerned about extremely unlikely risks and lead them to hoard liquidity so as to insure against these risks.

Ricardo Caballero and Arvind

Krishnamurthy study this phenomenon and show how government intervention may be useful in resolving it. They focus on liquidity crises that begin with unexpected events that call widely held beliefs and models into question. We have seen that this may characterize both the 1998 LTCM collapse as well as the recent disruption in financial markets that began with the downturn in subprime mortgage markets. Having scrapped old models, but without well-articulated new models to take their place, investors may tie up so much capital in response to concerns about extreme — but unlikely — events that they are unwilling to provide financing to meet more moderate — but likelier — liquidity needs.

Consider the example of corporations that deposit funds in a bank and, in return, have access to lines of credit that they can draw on should they experience a liquidity shock. In Caballero and Krishnamurthy's model, a sudden liquidity shock hits some firms in the economy and generates a need for borrowing. But those firms not affected by this first shock grow concerned that they may be hit by a second shock, even though this second shock is very unlikely. The unaffected firms react by preemptively drawing down their own lines of credit.[12] That is, they *hoard* liquidity. The result is that there is much less available for those firms that actually need liquidity because they have been hit by the first shock.

Reports in the popular press during the recent financial market disruption frequently refer to liquidity hoarding motivated by uncertainty. For example, in explaining elevated inter-

---

[11] These options, which were sold to Treasury bond dealers, each gave the holder the right to borrow $50 million from the Fed at a pre-specified interest rate, on a specific date between December 23, 1999, and January 12, 2000. The Fed also responded in other ways, for example, by creating a "century date change special liquidity facility" for banks.

[12] Firms may act preemptively because they are concerned that when the second shock hits, their credit quality will deteriorate so much that they will violate the covenants in their lines of credit and, thus, will be unable to borrow any further.

est rates in the interbank market, the *Wall Street Journal* quoted one banker as saying that "[banks and investors] are still fearful of each other and everybody is worried about counterparty risk and so people are hoarding their balance sheets."[13] This article also suggested that government intervention might reassure market participants and so reduce the impetus to hoard liquidity. We will see that, in Caballero and Krishnamurthy's model, government intervention can play such a role.

But why would banks hoard capital in response to an unforeseen shock? Caballero and Krishnamurthy assume that market participants are *uncertainty averse*. That is, when evaluating outcomes about which they are *uncertain*, they use the most pessimistic probability assessments. In particular, each participant overweights the probability that he will be among those hit by the second shock. (See *Uncertainty Aversion*.) This creates a desire to hoard liquidity against this unlikely shock.

Caballero and Krishnamurthy then discuss how government intervention might remedy this market failure. Their prescription is for the government to act as a lender of last resort. More precisely, by committing to provide liquidity in the event that the second shock occurs, the government thereby frees the private market to insure itself against the first, more likely, shock. Indeed, Caballero and Krishnamurthy quote former Fed Chairman Alan Greenspan to this effect: "'... [p]olicy practitioners operating under a risk-management paradigm may, at times, be led to undertake actions intended to provide insurance against especially adverse outcomes.'"

LIQUIDITY AND ASSET PRICES

We have seen that one way for firms to generate liquidity in times of need is to sell assets. But the level of liquidity can affect the value of these assets. This can then result in a "spiral," in which falling liquidity reduces asset values, which, in turn, leads to lower liquidity, and so on.[14] We have already discussed one example of this: the convertible bond market during the 1998 collapse of LTCM. These forced sales, Mitchell, Pedersen, and

Falling liquidity reduces asset values, which, in turn, leads to lower liquidity.

Pulvino argue, were the result of binding capital constraints.

Markus Brunnermeier and Lasse Pedersen develop a model that explains these spirals, along with many other features of liquidity crises. They focus on a particular aspect of liquidity: the need for immediacy. In their model a *customer* may arrive with an immediate need to sell an asset today, but no buyer may be available. So there is a need for temporary liquidity to bridge this gap (what they term *market liquidity*). This need for immediacy is provided by *speculators* (for example, securities dealers or hedge funds). The speculators serve a valuable economic role: They buy the asset today and then sell it at some later date when a buyer arrives. The speculators require funds in order to operate, and they obtain these funds from *financiers*, for

example, banks. So they themselves also have a need for liquidity. Brunnermeier and Pedersen call this *funding liquidity.*

In normal times, financiers themselves have adequate funding liquidity; therefore, they are able to provide market liquidity to their customers and thus assets are priced "fairly." That is, the price of an asset accurately reflects its expected future cash flows. However, when funding liquidity is scarce, there will also be insufficient market

liquidity, and asset prices will need to fall below this fair value to induce speculators to buy.

But why would funding liquidity be scarce? One reason is that speculators may have incurred losses on their other activities (as LTCM did). In addition, falling prices can themselves negatively affect speculators' funding liquidity. The reason is that speculators are limited in how much they can borrow by a collateral constraint.[15] That

---

[13] See the article by Greg Ip and Joellen Perry.

[14] The feedback between asset values and financing conditions has also been explored by macroeconomists seeking to explain the depth and persistence of economic downturns. An early and influential example is Irving Fisher's "debt deflation" theory of the Great Depression.

[15] Brunnermeier and Pedersen model this as a maximal value-at-risk (VaR) for the speculators. Banks commonly use value-at-risk to measure market risk, both for themselves and for their counterparties. Indeed, the Basel II Accord — an international agreement regarding how much capital banks need to put aside to guard against financial and operational risks — encourages the use of VaR to determine the amount of regulatory capital a bank must hold against its market risk. In the Basel II framework, VaR is calculated using a 10-day horizon, at a 1 percent probability level. So if a bank's market risk model predicts that there is only a 1 percent chance that the value of its portfolio will decline by *more than* $1 million in the next 10 days, its VaR is $1 million. VaR thus depends critically on the volatility of the value of the assets. In Brunnermeier and Pedersen's model this goes up in a crisis because price declines make the world appear to be more volatile.

As of March 2008, the Federal Reserve has responded in several ways to the liquidity problems associated with the recent disruptions in financial markets. These interventions provide examples of the policy instruments available to central banks.

• *Discount Window* – The Fed took two broad classes of actions to ease disruptions in financial markets by making it less costly for depository institutions to borrow directly from the Fed through the discount window. The discount window offers some advantages over private markets during episodes of tight credit. First, the Fed accepts a wider variety of collateral than do bank lenders (particularly during periods of financial market turmoil). In addition, by lending directly to depository institutions, the Fed can supplement the interbank market at times when it is not functioning well.[a] However, depository institutions are often reluctant to borrow directly from the Fed because of the perceived stigma it carries.

One step the Fed took was to narrow the spread between the discount rate (which is the rate that depository institutions must pay to borrow directly from the Fed's primary credit facility) and the federal funds rate (the rate at which banks borrow and lend among themselves, for one day at a time, on an unsecured basis). It also extended the terms of discount window loans; before the summer of 2007 they were overnight or very short-term loans. The Fed did this in two stages: On September 18, 2007, the Fed reduced the spread from 100 basis points above the target fed funds rate to 50 basis points and extended the maturity of discount window loans to up to 30 days. On March 16, 2008, it lowered the spread further, to 25 basis points, and also extended the maximum maturity of discount window loans to 90 days.

• *Term Auction Facility (TAF)* – This was a new policy tool announced on December 12, 2007. The Fed undertook to make 28-day loans directly to depository institutions at rates determined through competitive auctions. From these institutions' perspective, the TAF has several potential advantages over the discount window. One is that borrowing from the TAF may carry less of a stigma for a depository institution than accessing the discount window. In addition, depository institutions were able to place bids below the discount rate, so that they had the possibility of receiving funding at lower rates. While the TAF was new for the Fed, the European Central Bank regularly uses a similar tool.

• *Primary Dealer Credit Facility (PDCF)* – On March 16, 2008, the Fed announced the PDCF, which is a new, temporary discount window facility. The PDCF provides overnight funding to primary dealers[b] at the discount rate, in exchange for a specified range of collateral, including investment-grade mortgage-backed securities and asset-backed securities. This facility is intended to improve the primary dealers' ability to provide financing to participants in securitization markets and promote the orderly functioning of financial markets more generally.

• *28-Day Single-Tranche Repurchase Agreements* – On March 7, 2008, the Federal Reserve announced that it would initiate a series of term repurchase agreements[c] that are expected to cumulate to $100 billion. There

---

[a] Recall that there was a particularly large spread between LIBOR and Treasury rates, suggesting that there were indeed problems in the interbank market.

[b] Primary government securities dealers (*primary dealers*) - of which there are 20 - are banks or securities broker-dealers who may trade directly with the Fed. They are active participants in the Fed's open market operations as well as in U.S. Treasury auctions. Several of these are investment banks, and many others are subsidiaries of commercial banks. The current list of primary dealers may be found at http://www.newyorkfed.org/markets/pridealers_current.html.

[c] A repurchase agreement (or "repo") is a collateralized borrowing agreement structured as a sale of the collateral (in this case by dealers to the Fed), along with an agreement to buy it back at a higher price in the future (in this case, 28 days later). This higher price implicitly determines an interest rate, known as the "repo rate."

[d] Although these 28-day single tranche repos differ from the ones typically conducted by the Fed, in the past the Fed has occasionally conducted either 28-day, or single-tranche, repos.

are two main differences between these agreements[d] and typical Fed repurchase agreements. First, they are 28-day repos; typically the term is shorter. In addition, they are "single-tranche": dealers may submit any of the following types of collateral — Treasuries, agency debt, and agency mortgage-backed securities — and pay the same repo rate regardless of its type. By contrast, the repo rate typically differs by the type of collateral, with those pledging Treasuries paying the lowest rate and those pledging mortgage-backed securities the highest. Since under the new program market participants face the same rate regardless of the collateral, they have an incentive to submit only mortgage-backed securities.[e]

• *Term Securities Lending Facility (TSLF)* – On March 11, 2008, the Fed announced an expansion of its securities lending program. Under the new program, the Term Securities Lending Facility, the Federal Reserve will lend up to $200 billion of Treasury securities to primary dealers for a term of 28 days (rather than overnight, as in the existing program) by a pledge of other securities, including federal agency debt and both agency and AAA-rated nonagency mortgage-backed securities. The TSLF is intended to promote liquidity in the financing markets for Treasury and other collateral and thus to foster the orderly functioning of financial markets more generally.

• *Cooperation with Other Central Banks* – Other central banks also undertook to increase their liquidity provision through similar means. In addition, the Fed entered into "reciprocal currency arrangements," in which it lent dollars to the European Central Bank and the Swiss National Bank, which, in turn, offered dollar loans to their member banks. This was the first time since September 11, 2001, that the Fed had entered into such arrangements with central banks in Europe.

This period was also characterized by a slowing economy and by concerns that continued financial market turmoil could slow the real economy further. In response, the Fed reduced its target for the fed funds rate. Between September 2007 and March 2008, the Fed cuts its target from 5.25 percent to 2.25 percent

---

[e] That this was in fact the case can be seen from the New York Fed's announcements of the results of its single-tranche repos on March 7, 11, 18, and 25, 2008.

---

is, the speculators must post securities to back any loan they take. (See *Margins and Liquidity*.) This collateral constraint becomes tighter during a crisis for two reasons. First, the value of the collateral is lower in the crisis, because asset prices have fallen. In addition, falling prices make the world appear more volatile, which also leads the bank to tighten its collateral constraint. This results in a liquidity spiral, as described above.

In addition to liquidity spirals, Brunnermeier and Pedersen are able to explain other characteristic features of liquidity crises. The first is the flight to quality: a rush to buy safe assets, which is reflected in a relative appreciation in their price. This is an outcome of Brunnermeier and Pedersen's model because when liquidity is scarce, market participants prefer to conserve it by investing in less risky assets. This causes the price of riskier assets to fall more than those of safe ones (and hence causes their yield to rise relative to that of safe assets). Another feature that Brunnermeier and Pedersen are able to explain is the commonality of liquidity across securi- ties and markets, that is, the tendency for liquidity problems to spread from one market to another (as was the case for the convertible bond market in 1998, discussed earlier). This occurs because speculators provide liquidity in many markets simultaneously, and so a deterioration in their financing position (even if initially caused by a shock to a single market) will affect all of these markets.

Brunnermeier and Pedersen also discuss the possible regulatory responses to a crisis. First, they argue that if the regulator "knows" that this

R icardo Caballero and Arvind Krish-namurthy use uncertainty aversion in their model to explain the hoarding of liquidity. There is evidence that decisions are indeed characterized by aversion to uncertainty; that is, individuals seek to avoid ambiguous situations in which probabilities are not known. A classic example is the Ellsberg paradox.

In the Ellsberg paradox, there is an urn containing 90 balls: 30 red balls and 60 black and yellow balls in unknown proportion. A ball is drawn, and an individual is then asked whether he prefers to bet that this ball is red (gamble A) or black (gamble B). Most people choose gamble A. That is, they prefer to bet on a red ball, which they know will occur one-third of the time, against the chance of a black ball, whose probability lies somewhere between 0 and two-thirds. Note that gamble B reflects *uncertainty*, in that the precise probability that a ball is black is unknown.

Then the ball is replaced and another is drawn. The (same) individual is now asked whether he prefers to bet that the ball is either red or yellow (gamble C) or black or yellow (D). Most people choose gamble D in this case. That is, they prefer to bet that a black or yellow ball is drawn, which they know will happen two-thirds of the time, against the chance of a red or yellow ball (the latter probability lies somewhere between one-third and 1).

But there is a certain inconsistency in these two choices because the second gamble is really equivalent to the first, with the addition of the possibility of yellow balls to each option. So if someone prefers gamble A over B, he should actually prefer gamble C over D. For example, if one believes there are more red balls than black balls (and hence prefers gamble A over B), one also believes that there are more red or yellow balls than black or yellow balls, and so should prefer gamble C over D.

Ellsberg explained these seemingly contradictory choices as reflecting individuals' dislike for uncertainty, that is, for unknown probabilities. They prefer A over B because they know that one-third of the balls are red; conversely, they prefer D over B because they know that two-thirds of the balls are either black or yellow.

Since this evidence is inconsistent with the canonical economic model of expected utility maximization,* it has led to the development of alternative models of decision-making under uncertainty. One of these, developed by Itzhak Gilboa and David Schmeidler, models individuals faced with unknown probabilities as pessimistic – that is, as making decisions under the "worst case" probabilities. This is the approach used by Caballero and Krishnamurthy.

---

* In particular, the "sure-thing principle."

---

is, in fact, a temporary liquidity shock, he should try to convince banks to lend to speculators. If the regulator is correct, banks' profits in times of crisis will actually be higher. However, attributing superior information to the regulator is, of course, a very strong assumption. In addition, Brunnermeier and Pedersen suggest that the regulator can help stabilize prices by providing liquidity directly to the speculators (or to the banks that finance them).

CONCLUSION

Financial intermediaries serve to allocate the private market's wealth so as to meet firms' and investors' liquidity needs. But "liquidity crises," in which the market fails to function properly, are a recurrent feature of financial markets.

During these episodes the demand for liquidity may be so great that it cannot be met by the private sector alone. This creates a role for govern-ment intervention; the government can provide liquidity in these circumstances through its ability to raise funds by taxing consumers.

In addition, these episodes may also be characterized by a misallocation of private liquidity. Market participants may hoard liquidity because they become concerned about extremely unlikely events. In this case there may be a further role for the government in insuring against these extreme events.

Margins and Liquidity

In Markus Brunnermeier and Lasse Pedersen's model, margin requirements play a key role. They may exacerbate liquidity crises and thus facilitate a liquidity spiral. Likewise, Mark Mitchell, Pedersen, and Todd Pulvino document the role that capital constraints have played in many real-life liquidity crises (for example, in the convertible bond market in the 1998 LTCM crisis). Given that they may exacerbate crises, why do margin requirements and capital constraints exist? And are they optimally determined by private markets?

On the most basic level, requiring borrowers to post margin, or collateral, facilitates lending by increasing the likelihood that lenders will be repaid.[a] Viral Acharya and S. Viswanathan show how this can help to provide liquidity. Consider firms facing liquidity shocks. If they are able to borrow enough, they can meet their liquidity needs without selling any assets. However, firms may be limited in what they can borrow because lenders are concerned that the firm might divert funds to risky projects. In general, borrowers have a tendency to prefer riskier projects than their lenders because lenders bear the brunt of any losses when the firm defaults. By pledging their assets as collateral, firms are able to reassure lenders that they will indeed invest efficiently; so firms are able to borrow more and thus meet larger liquidity needs.

While collateral does facilitate borrowing, if the liquidity shock is large, not all firms will be able to borrow enough to survive, and some will need to be liquidated early. In cases where there are many more sellers than buyers, this will, in turn, lower asset prices. These lower prices then lead to tighter collateral constraints and thus to a liquidity spiral, as described earlier.

Although Acharya and Viswanathan show that the ability to post collateral is valuable to society, does it follow that the private market sets collateral requirements optimally? In a different model, John Geanakoplos and Felix Kubler present an example suggesting that this is not always the case.[b] They show that margins may be too loose in good times and too tight in crises because lenders do not take into account the effect that the margins they set have on market prices. So there may also be a role for government regulation of margin requirements (such a view was also recently expressed by former Treasury Secretary Robert Rubin). Indeed, Geanakoplos and Kubler suggest that margins may be too loose in good times, since lenders do not realize that their lending increases the risk of a future crisis; conversely, margins may be too tight in crises, since increasing lending would raise prices and thereby ease the crisis. The extent to which this intuition can be generalized is unclear, and further research in this area would be valuable.

_____

[a] For more on the economic role of collateral, see the article by Yaron Leitner.

[b] See also the presentation by John Geanakoplos at the Philadelphia Fed Policy Forum.

This frees up the private markets' liquidity and allows it to be used more effectively.

Another feature of liquidity crises is the interaction between liquidity and asset prices. A lack of liquidity can lower asset prices below their fair, or fundamental, value. Since liquidity is often obtained by using these assets as collateral for loans, this in turn can lead to lower liquidity provision. The outcome is a "liquidity spiral," in which the resulting illiquidity can lower asset prices even further, and so on. Further research is needed on this role of collateral in providing liquidity and, in particular, on the question of whether government intervention can improve on the private market's use of collateral. ⓑ

REFERENCES

Acharya, Viral V., and S. Viswanathan. "Moral Hazard, Collateral and Liquidity," manuscript, December 2007.

Brunnermeier, Markus K., and Lasse Heje Pedersen. "Market Liquidity and Funding Liquidity," NBER Working Paper 12939 (2007); *Review of Financial Studies* (forthcoming).

Caballero, Ricardo J., and Arvind Krishnamurthy. "Collective Risk Management in a Flight to Quality Episode," J*ournal of Finance* (forthcoming).

Fisher, Irving. "The Debt-Deflation Theory of Great Depressions," *Econometrica 1* (1933), pp. 337–57.

Gilboa, Itzhak, and David Schmeidler. "Maxmin Expected Utility with Non-unique Priors," *Journal of Mathematical Economics*, 18 (1989), pp. 141-53.

Geanakoplos, John, and Felix Kubler. "Leverage, Incomplete Markets and Pareto Improving Regulation," mimeo, Yale University (November 2005).

Geanakoplos, John. "The Leverage Cycle," presentation at the Philadelphia Fed Policy Forum, November 30, 2007. http://www.philadelphiafed.org/econ/conf/forum2007/presentations/Geanakoplos_leverage_cycle_6_11-20-2007.pdf

Gennotte, Gerard, and Hayne Leland. "Market Liquidity, Hedging, and Crashes," *American Economic Review*, 80:5 (December 1990), pp. 999-1021.

Holmstrom, Bengt, and Jean Tirole. "Private and Public Supply of Liquidity," *Journal of Political Economy*, 106:1 (February 1998), pp. 1-40.

Ip, Greg, and Joellen Perry. "Central Banks Launch Effort to Free Up Credit," W*all Street Journal*, December 13, 2007.

Leitner, Yaron. "Using Collateral to Secure Loans," Federal Reserve Bank of Philadelphia *Business Review* (Second Quarter 2006).

Lowenstein, Roger. *When Genius Failed: The Rise and Fall of Long-Term Capital Management.* New York: Random House, 2000.

O'Hara, Maureen. *Market Microstructure Theory.* Blackwell, 1995.

Morris, Stephen, and Hyun Shin. "Coordination Risk and the Price of Debt," *European Economic Review*, 48 (2004), pp. 133–53.

Mitchell, Mark, Lasse Pedersen, and Todd Pulvino. "Slow Moving Capital," *American Economic Review*, Papers and Proceedings, 97:2 (2007), pp. 215-20.

Sundaresan, Suresh, and Zhenyu Wang. "Y2K Options and the Liquidity Premium in Treasury Bond Markets," Federal Reserve Bank of New York, Staff Report No. 206.

The Evolution of the World Income Distribution*

BY KEITH SILL

There is tremendous disparity in the levels of individuals' incomes across countries. However, this disparity in per capita income has not always existed. In this article, Keith Sill investigates some facts about the evolution of per capita income across countries and reviews a simple model that broadly captures the observed evolution of the world income distribution since 1800. He then discusses what predictions can be made about future cross-country distributions of income and some policy prescriptions that follow from our understanding of the past and our predictions about the future.

There is tremendous disparity in the levels of individuals' incomes across countries. Those fortunate enough to live in the richest countries have an average income that is about 30 times greater than the average income of residents of the world's poorest countries. Such a large disparity in income across countries implies large differences in living standards

and well-being. A significant share of the world's population has a living standard well below that of the average U.S. citizen. Indeed, inhabitants of the world's poorest countries face daily hardships and deprivations that are so foreign to the citizens of rich countries as to be hard to believe.

However, this large difference in per capita income across countries has not always existed. It wasn't until the early 19th century that countries began to experience significantly different growth rates in income as some countries were quicker to begin the process of industrialization. Conse-

quently, before the late 1800s, there was relatively little income disparity across countries, at least by today's standards. But it doesn't take long for small differences in income growth rates to lead to wide divergence in per capita income levels. From the late 1800s until about the 1960s, there was a steady and rapid increase in inequality. Since then, the cross-country dispersion in per capita income has become somewhat more stable, while, at the same time, world poverty has been decreasing as countries with large populations, like China and India, begin to industrialize.

We'll investigate some facts about the evolution of per capita income across countries and review a simple model that broadly captures the observed evolution of the world income distribution since 1800. Given our analysis of what happened in the past, we'll discuss what predictions can be made about future cross-country distributions of income. We'll also discuss some policy prescriptions that follow from our understanding of the past and our predictions about the future.

EVOLUTION OF COUNTRY PER CAPITA INCOMES BEFORE 1800

Before 1800 and the onset of the Industrial Revolution, the distribution of world income looked very different than it does today. While cross-country data on incomes and population prior to 1800 are incomplete and challenging to piece together, the available information suggests that there was little, if any, growth in per capita incomes in any of the world's economies. Before the Industrial Revolu-

Keith Sill is an economic advisor and economist in the Research Department of the Philadelphia Fed. This article is available free of charge at www.philadelphiafed.org/econ/br/.

tion, economies were agricultural, and living standards were similar across countries and over time. People born before 1800 could expect to be about as well off as their parents, grandparents, and great-grandparents. In addition, they could expect their children to be about as well off as they were. Moving to a different country wouldn't have improved living standards much either — the agricultural technology across countries was about the same. This stands in stark contrast to today's world, in which living standards have increased rather consistently over time (at least in the developed countries) and vary greatly between poor and rich countries.

We will measure the standard of living, or economic well-being, of the typical resident of a country using real gross domestic product (GDP) per capita, which is real GDP divided by a measure of the population. Real GDP is all of the goods and services produced domestically by residents of a country. Higher real GDP means that a country produces more goods and services for its residents to consume and invest in. By itself, real GDP is not a particularly good measure of how rich a country is because a country with a large population is likely to produce more than a country with a small population. When we divide a country's real GDP by its population, we get a measure of goods and services produced per person: Rich countries will produce more per person than poor countries.

However, real GDP per capita is not an all-inclusive measure of a country's well-being. Factors that affect well-being include leisure time, income sharing within households, environmental quality, and health. These factors may be imperfectly correlated with output per person, and there is some evidence from survey data that the correlation between output per

capita and happiness is weak across OECD countries.[1] Despite these potential problems, we will treat real GDP per capita as a useful summary measure of well-being for purposes of cross-country comparison. After all, it seems implausible to argue that in some broad sense Africa's poorest residents are as well-off as the residents of the U.S.

Unfortunately, official statistics on GDP for most countries start after World War II. So how do we measure the world income distribution far back

in history? Recent work by Angus Maddison, the eminent economic historian, pieces together various bits of evidence to develop measures of real GDP per capita for several regions of the world going back to 1 AD. Going that far back in time means that there is considerable uncertainty about particular measures of living standards, since the recorded evidence on how people lived is sparse. However, over that time span, virtually all societies were traditional agricultural societies, and agricultural technology seems not to have varied greatly across countries. Furthermore, we could hypothesize

that average living standards before 1800 were probably similar to the living standards of today's poorest countries, which are agricultural societies that do not have much capital or technology to work with.

Figure 1 plots some of the per capita income data from Maddison's study for several regions of the world from 1 AD to 1820.[2] The figure shows that even the fastest growing regions of the world, which are denoted Western Europe and Western Offshoots (Australia, Canada, New Zealand, and the
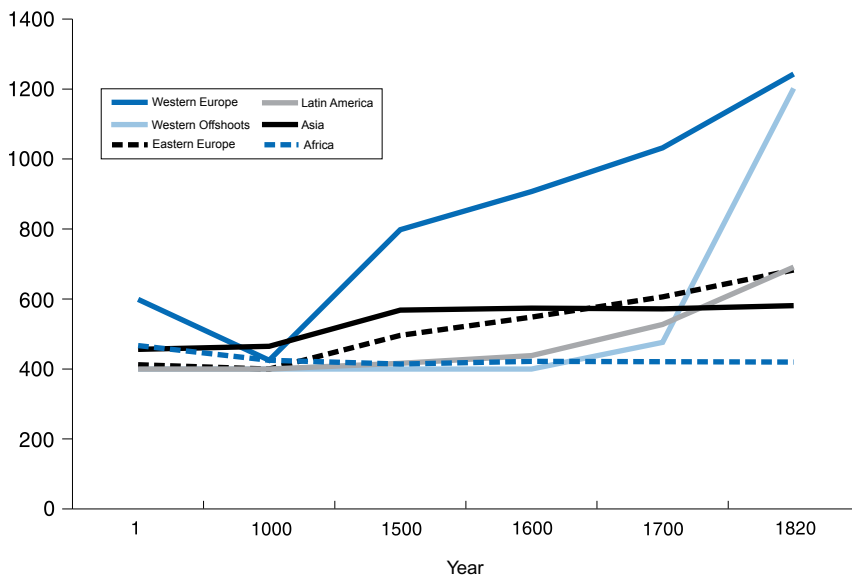
United States), had per capita incomes that increased only by a factor of two to three over a span of 1800 years. This amounts to minuscule growth of only about 0.04 percent per year. By 1820, the richest region (Western Europe) had per capita income about three times that of the poorest region (Africa). But this is nothing like the 30-fold difference we see today between the richest and poorest countries. The story of economic growth before 1800 appears to be one of stagnation in living standards.

The near-zero growth of per capita incomes between 1 AD and 1820 does not mean that there was no technological progress during that time. Productivity-improving inventions such as the stirrup, the heavy plow, and the three-field system of crop rotation were being adopted. However, societies responded to technological advance by increas-

> We could hypothesize that average living standards before 1800 were probably similar to the living standards of today's poorest countries, which are agricultural societies that do not have much capital or technology to work with.

[1] The Organization for Economic Cooperation and Development is a group of 30 countries that share a commitment to democratic government and the market economy. The working paper by Romina Boarini, Asa Johansson, and Marco Mira D'Ercole provides an overview of the literature on wealth and happiness.

FIGURE 1

Per Capita Income 1 AD to 1820 AD

Per Capita Income $1990

Legend:
- Western Europe
- Western Offshoots
- Eastern Europe
- Latin America
- Asia
- Africa

Source: Maddison (2006) at http://www.ggdc.net/maddison

ing their populations rather than by accumulating more capital per worker and thus increasing output per worker. In other words, population grew at the same rate as output, so that output per worker, or living standards, stayed nearly constant.[3]

Why was there no growth in per capita income in the pre-1800 period? Before 1800, land was a very important factor of production, since economies were largely agrarian. However, the quantity of land available for production was to some extent fixed. If the return to adding a worker to a plot of land diminishes as more workers are added, output per worker declines with population growth. However, technological advance is an offsetting factor that makes workers more productive. If these two opposing forces of population growth and technological advance approximately balance each other, output per capita remains roughly constant over time, even though technological progress occurs.[4] It is not a coincidence that

these forces roughly cancel each other. Thomas Malthus famously argued that the stability of living standards at a low level of subsistence resulted from the tendency of population to grow whenever living standards rose above a subsistence level. The benefits afforded by steadily improving technology were absorbed in supporting a steadily rising population.

POST-1800 INCOME DISTRIBUTION: THE INDUSTRIAL REVOLUTION AND THE ERA OF MODERN ECONOMIC GROWTH

While the pre-1800 era is marked by relatively little, if any, growth in per capita incomes across countries, the post-1800 era is marked by a dramatic surge in growth for some countries but not others. We see this in Figure 2, which plots income per capita for the same regions as in Figure 1 but includes the period 1820 to 2000. Some regions, such as the West, experience sustained increases in their per capita incomes. Other regions, such as Africa, show virtually no increase in per capita income over this time span. A consequence of this disparity in growth is that the levels of income per capita vary greatly across regions. Indeed, by 2000, the countries labeled Western Offshoots (Canada, the U.S., Australia, and New Zealand) had per capita incomes that averaged about 15 times higher than the per capita income of Africa. The U.S. has a per capita income that is about 30 times higher than that of the poorest countries in Africa.

The central question of economic development is why some countries make the transition to modern growth while others stagnate. Many factors appear to be at work, and the particular set of circumstances, policies, and institutional structure that coincides with the transition to modern growth

FIGURE 2

Per Capita Income 1 AD to 2000 AD

Per Capita Income $1990

Source: Maddison (2006) at http://www.ggdc.net/maddison
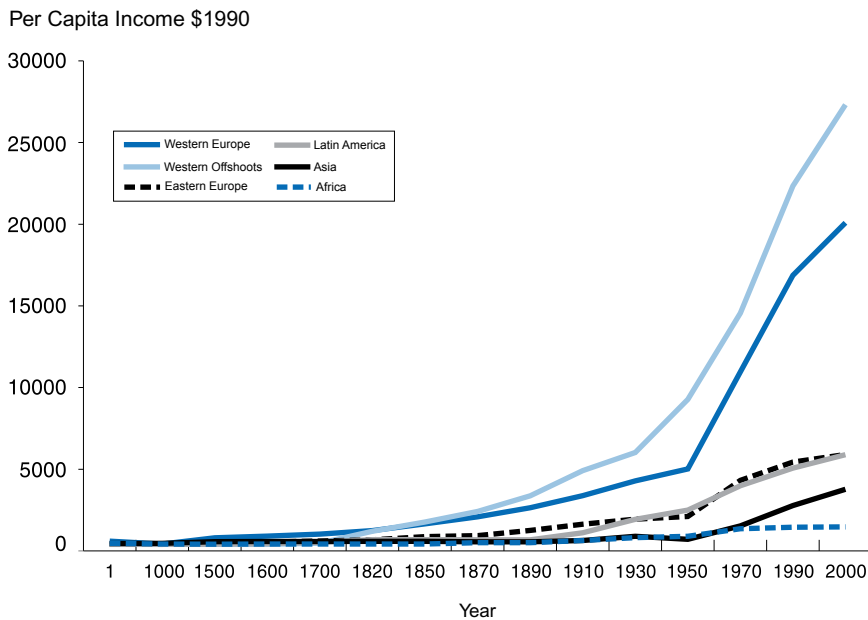
varies from country to country. Developing well-defined property rights seems to be an important factor as does developing human capital so that people can implement and take advantage of new technologies.

What we see happening as countries make the transition from stagnation to economic growth is that countries industrialize, technological change allows for increasing productivity, and capital accumulation begins. That is, residents invest in buildings and machines that help produce more future output. Countries that enter this modern growth phase shift from primarily agricultural output to industrial output that is not as land intensive. That is, countries undergo an industrial revolution. A second key observation is that fertility changes, so that improvements in technology

no longer result in matching increases in population. Parents have fewer offspring and invest more resources in their care and development. This demographic transition is a defining feature of the transition to sustained growth for economies.[5]

Great Britain was the first country to enter the industrialization process. The exact date of the Industrial Revolution's beginning is difficult to determine, but many historians place it in the latter half of the 1700s. From Great Britain, it spread to Western Europe and the United States. The onset of the Industrial Revolution led to divergence in per capita incomes across countries, since not all countries began

[5] The article by Aubhik Khan discusses the relationship between the demographic transition and industrialization in more detail.

the industrialization process at the same time. Consequently, the world income distribution started to become more unequal as some countries started to grow (at different dates) and others that had not started to industrialize continued to stagnate.

Income inequality across countries largely increased from about 1820 until the latter half of the 20th century as more and more countries began to industrialize. Since 1960, there are good cross-country data that can be used to examine the change in the world income distribution more carefully. Alan Heston, Robert Summers, and Bettina Aten have developed these data in the Penn World Tables. In these tables, these researchers provide national income accounts converted to international prices for 188 countries from 1950 to 2004, assuming purchasing power parity holds.[6] Their data set allows a direct comparison of income across countries because incomes are measured in common units.

Figure 3 plots the distribution of income per capita across countries using the Penn World Table Mark 6.2.[7] The distribution is approximated by a histogram, which shows what proportion of countries fall into a particular category of per capita income levels. In Figure 3, income levels are measured relative to the U.S., so the scale runs from zero to one.[8] Histograms are

[6] Purchasing power parity is a relative price concept that measures the number of units of country A's currency that are needed in country B to purchase the same quantity of a good or service as one unit of country A's currency will purchase in country A. See the OECD's glossary of statistical terms at http://stats.oecd.org/glossary/index.htm.

[7] The data are available at http://pwt.econ.upenn.edu/.

[8] Some countries have higher per capita income than the U.S. We placed all countries with per capita incomes greater than or equal to the U.S. in the same data category as the U.S.

plotted for three years: 1970, 1985, and 2000. All countries that have data available for each of the three years are included in the charts. Perhaps the most striking observation from the charts is that so many countries have per capita incomes far below that of the U.S. In fact, about 75 percent of countries in the 2000 sample have per capita incomes that are less than half that of the U.S. Furthermore, the distribution has been fairly stable over the 30-year span.

Although there wasn't a large change in the number of poor countries between 1970 and 2000, that doesn't mean that worldwide poverty has not decreased. First, the histograms in Figure 3 assume that each resident of a country gets the same share of real GDP. However, we know that there is income inequality within countries and that this inequality can change through time. Any such change in within-country inequality is not captured by the histograms in Figure 3. Nevertheless, the difference in living standards across countries tends to exceed the within-country differences by most measures. For example, in 1988 the ratio of the per capita income of a country in the richest 10 percent of countries to that of a country in the poorest 10 percent was 20. But for the U.S., the ratio of the permanent income of an individual in the richest 10 percent to that of an individual in the poorest 10 percent was less than four.[9] Second, the histograms in Figure 3 give each country equal weight, despite the fact that some countries have much larger populations than others. So if China, with its massive population, jumps to a new income category in the figure,

[9] Permanent income is a measure of a wage earner's long-term income that ignores short-term fluctuations in earnings.

FIGURE 3

## World Income Distribution Per Capita Relative to U.S.
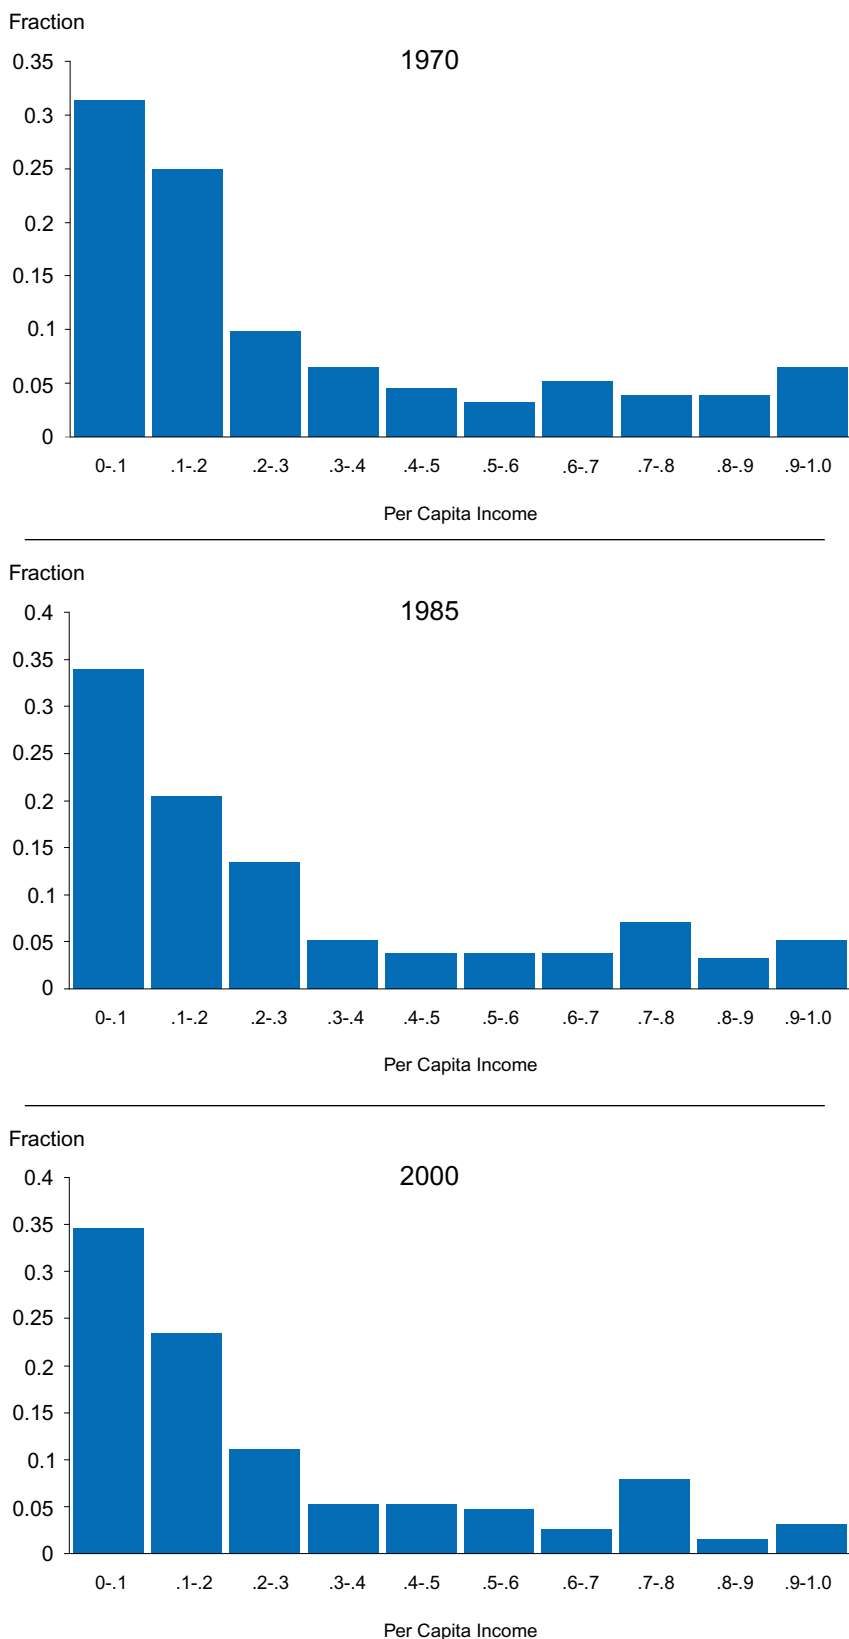
it has the same effect as if Bermuda jumped income categories.[10]

We can partially correct our histograms for country size by multiplying income per capita in each country by that country's share of world population (note, though, that this still does not correct for within-country inequality). The weighted country distributions are shown in Figure 4, and now the picture looks somewhat different. We see that in 1970, ignoring within-country inequality, nearly 60 percent of the world population had a per capita income that was 15 percent or less than that of the U.S. By 2000 there had been a significant shift, in that the 15 percent of U.S. income or less category shrank at the expense of the 15 to 40 percent category. Thus, many fewer people were in the lowest income category compared to 1970. What's happening is that high-population countries like China and India are starting to experience sustained increases in per capita incomes. This evidence suggests that world poverty has been decreasing even though the distribution of country per capita incomes has remained relatively stable.

A SIMPLE MODEL OF THE WORLD INCOME DISTRIBUTION

The evidence presented so far suggests three phases for the dynamics of the world income distribution. The first phase, the pre-1800 period, is characterized by economic stagnation because per capita incomes were not growing. The second phase, from 1800 to the latter half of the 20th century, is characterized by increasing inequality as industrialization diffused through the world's economies. The third phase, from about 1960 to now, was one in which the world income

---

[10] See the article by Xavier Sala-i-Martin for a thorough discussion of this issue.

FIGURE 4

Population-Weighted World Income Distribution Per Capita Relative to U.S.
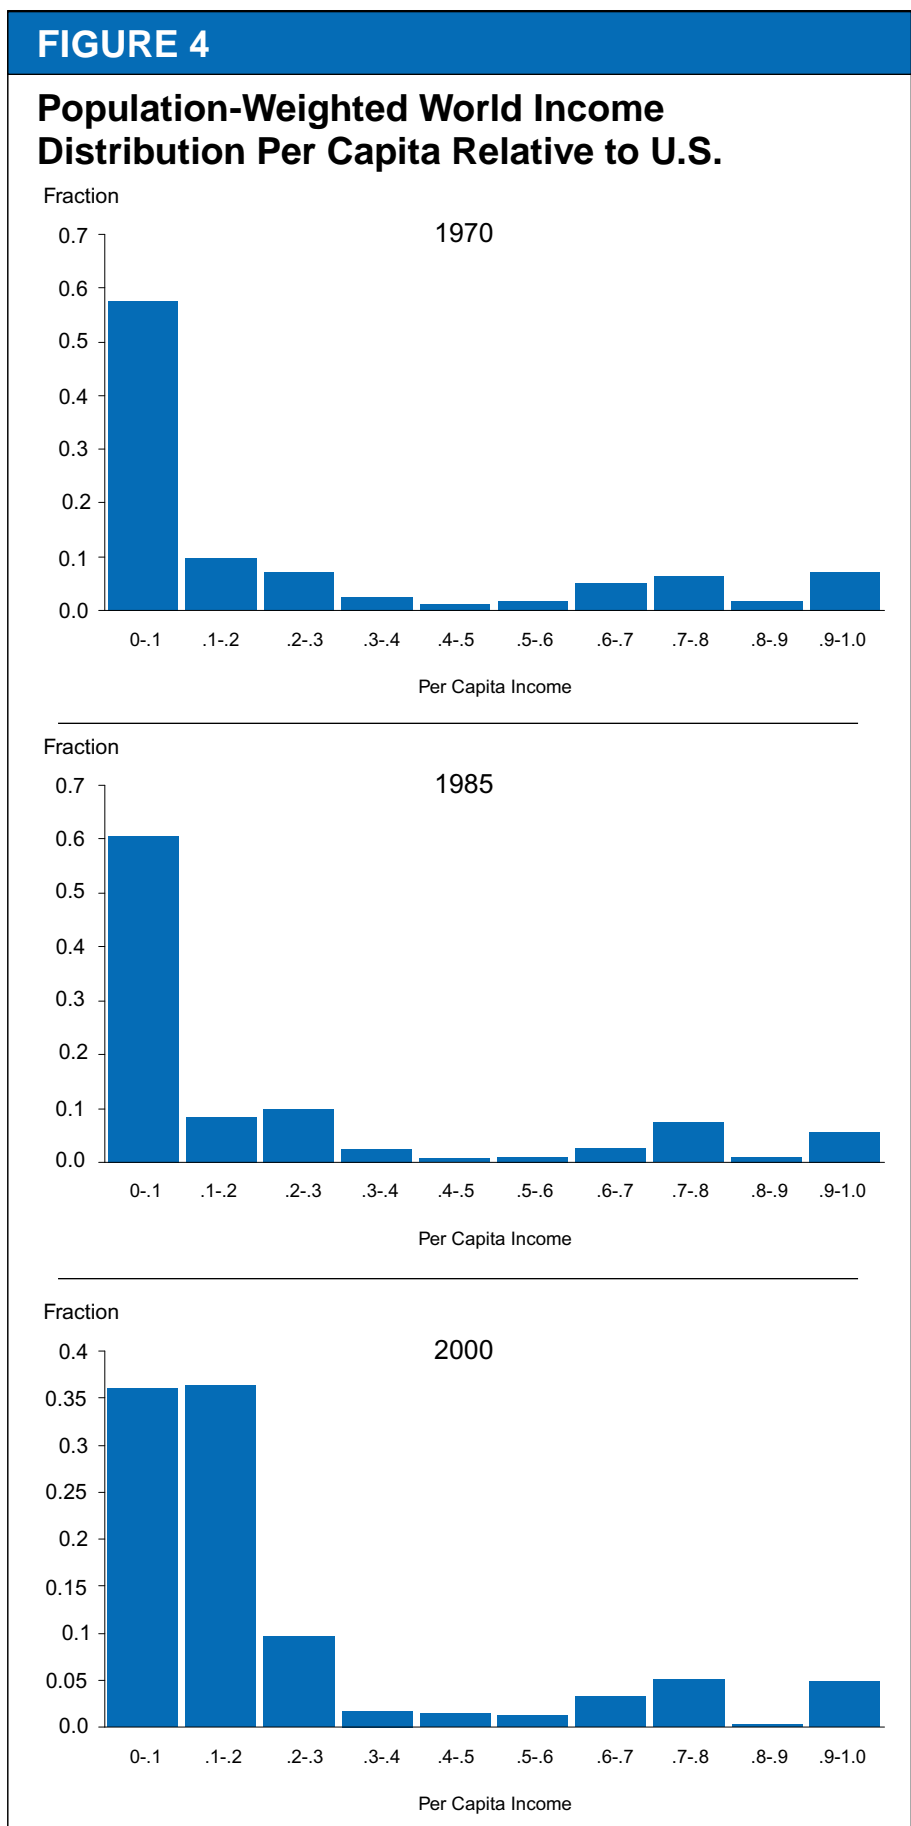
distribution was somewhat more stable.

How should we think about these phases and the recent stability of the distribution? Is the recent period a pause before a return to a further increase in inequality? Is it the shape of things to come, in that things have settled down and we should expect the future to look much like the recent past? Or is the most recent phase merely a transition period that will usher in a period of declining income inequality across countries?

The last conjecture is nicely advocated in a recent paper by Nobel laureate Robert Lucas titled "Some Macroeconomics for the 21st Century." We are going to explore Lucas's reasoning further and, as a byproduct, see that the conjecture about a "pause before a return to increasing inequality" and the conjecture about the "shape of things to come" are somewhat at odds with the long-run world-growth mechanism. However, none of the views can really be proven or disproven based on the existing set of facts, since the truth will be revealed only as the future unfolds.

Lucas uses a simple mechanical model to articulate the view that the last 40 years represent a transition period for the world income distribution. The transition can be thought of as moving from the pre-industrial era to the post-industrial era. The post-industrial era is one in which industrialization has largely diffused throughout the world's economies and ever fewer countries remain stagnant. The model is mechanical in the sense that while it is based in part on an economic interpretation of the facts, there are no choices, policies, or meaningful actions at the model's core.

The model has several key elements. Consider first the concept of the *technology frontier*. The frontier represents the best, most recent technology that countries can use for transforming labor and capital inputs into output. These new technologies represent things like information technology, genetics, advances in medical care, improved organizational methods, and a host of other things that make labor and capital more productive. If we examine the data on per capita incomes for countries over the past 200 years, we could reasonably argue that, on average, the richest countries, those at the technology frontier, increase their real incomes at a pace of about 2 percent per year. If we take this as giving us some information about the pace of technological advance, we can say that the technology frontier grows about 2 percent per year. This represents the maximum long-run growth a country can achieve. The model takes this as given and does not describe an economic mechanism for why the technology frontier grows at 2 percent or how to make it grow faster.

For simplicity, assume that all countries have the same population and that the population does not grow.[11] Also, think of these countries as starting out at some time before the Industrial Revolution, so that there is no growth in real income per capita. The model begins at a time when the world consists of a bunch of poor, stagnant economies that have the same incomes and the same populations. How does growth occur and industrialization diffuse? Lucas uses a racetrack metaphor to describe how the model works.

Imagine all countries lined up in a row like horses at the starting gate of a racetrack. But instead of all gates opening at the same time, they open randomly. When the gate opens, the corresponding country starts to grow. Thus, some economies start to grow at time 1, others at time 2, others at time 3, etc. In any year after the starting period (taken to be 1800), the world economy is composed of countries that have already been growing, those that just started growing, and those still waiting to start growing. The Lucas model does not posit any economic reasons for a country to make the transition to modern growth and so does not offer policy advice to kick start economies into the growth phase of development. Rather, it models the random process by which gates are opened using observations on world income growth since 1800.[12]

The model makes two other key assumptions based on observations of historical growth rates. First, when a country begins to grow, it does not necessarily grow at 2 percent, the rate of growth of the technology frontier. Rather, it grows at 2 percent plus a growth rate determined by the income gap between itself and the richest country. The later a country starts to grow, the larger is the income gap between itself and the leader (which we will take to be the U.S.) and the faster it grows. As a country closes the income gap with the leader, its growth rate begins to slow toward 2 percent. In the long run, a country that is growing does so at 2 percent, the rate of growth of the technology frontier.[13] This modeling behavior is based on observations from countries such as South Korea, Japan, and China, which experienced very high growth rates as they began the industrialization

---

[11] The model assumes that all countries are the same size and so makes the same prediction for inequality whether or not countries are weighted by population. In the model, country size does not matter for growth or development or for the inequality consequences of development of large-population economies.

[12] Thus, the model does not explain why the Industrial Revolution started in England and then quickly spread to continental Europe and the United States.

[13] See Lucas's 2000 paper for a detailed description of the model and how it is calibrated.

process. This can happen because these later entrants to the growth club do not have to reinvent all of the advanced technologies they see in countries like the U.S. To some extent they can import advanced technologies and try to copy existing best practices and methods. This allows them to industrialize more rapidly than if they had to develop all the new technologies themselves.
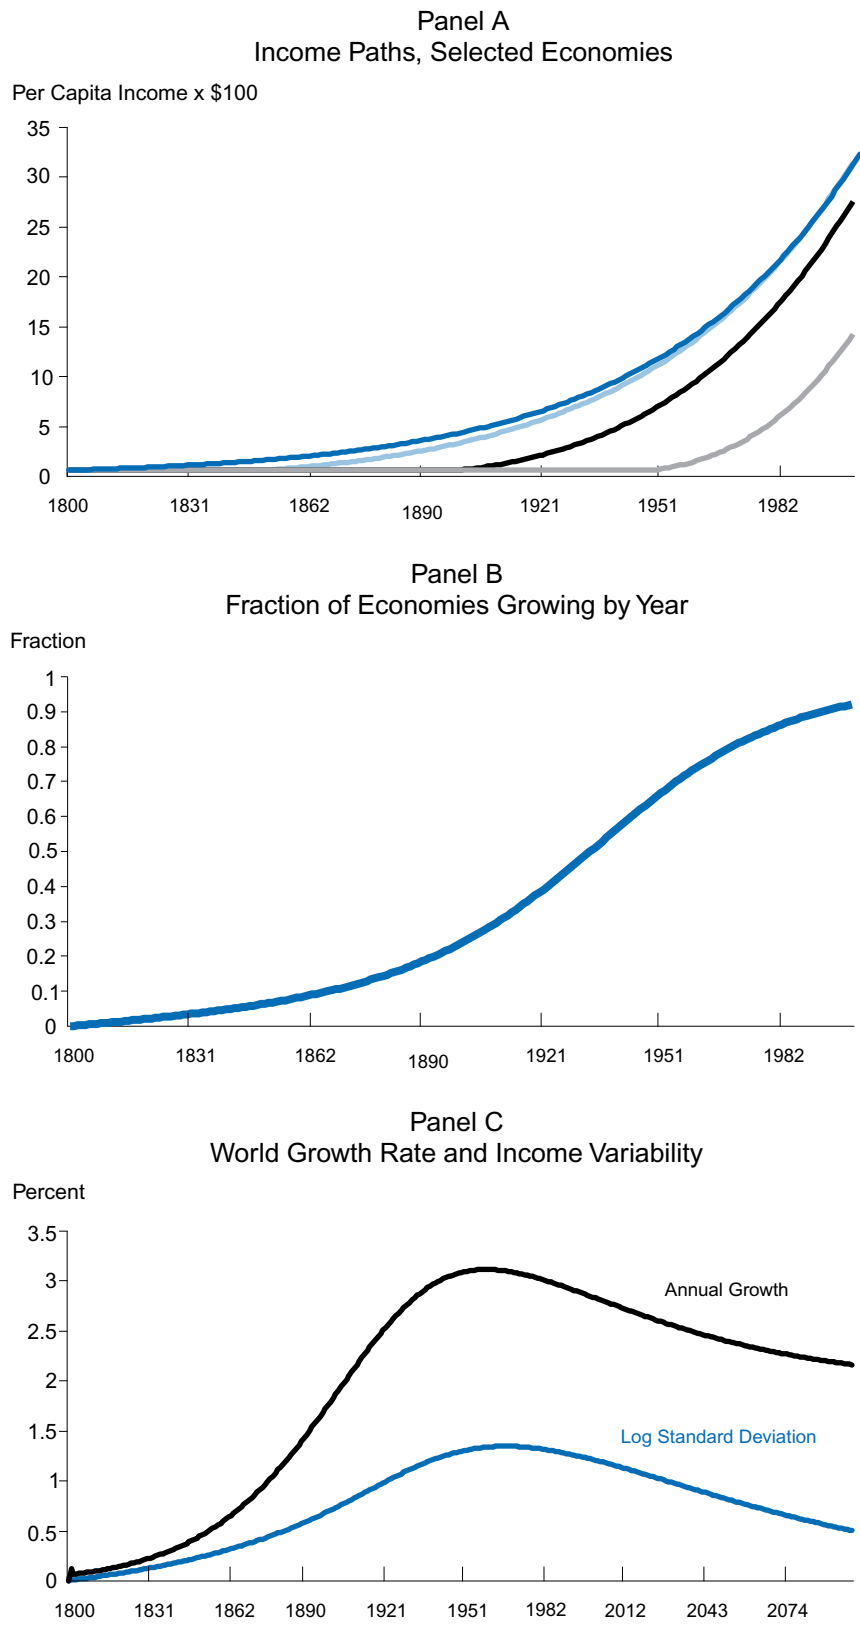
Another key assumption is that the probability that a stagnant country begins to grow in any given period is positively related to the level of income in the rest of the world. The more countries that are growing in the rest of the world, the higher the probability that a stagnant country will start to grow. This is a model of spillovers: The more technology the world acquires — the more prevalent or diffuse it is — the easier it is for a stagnant country to begin taking advantage of it and start the industrialization process.

Figure 5 shows some results from simulating Lucas's basic model, reproducing some figures from his article. Panel A shows how the model behaves for a sample of four countries that start growing at different times: 1800, 1850, 1900, and 1950. The first country that begins to grow (shown in blue) – the leader country – grows at a rate of 2 percent. The countries that begin to grow later (shown in light blue, black, and grey) do so at progressively faster rates initially because their income gap with the leader is wider, the longer they remain in stagnation. Over time, the income per capita in these countries just about catches up with the leader's per capita income. Panel B plots the model's prediction about the fraction of countries that have entered the modern growth phase for the period 1800 to 2000.[14] By 2000,

[14] The model calibration is the same as in Lucas's 2000 paper.

FIGURE 5

Results from Simulating Lucas Model*

Panel A
Income Paths, Selected Economies

Panel B
Fraction of Economies Growing by Year

Panel C
World Growth Rate and Income Variability

* Figures recreated by author based on Lucas (2000)

about 90 percent of the world's economies have entered the modern growth phase. Panel C shows the implied world annual average growth rate and a measure of cross-country inequality. The annual average growth rate is simply the average rate of growth of all the economies in the model. The income inequality measure plotted is the cross-country standard deviation of incomes (in logs).

The key figure is Panel C. It shows that the growth of world average annual income peaks at a rate a bit above 3 percent around the 1970s and then begins to decline. Average annual growth can exceed 2 percent (the growth rate of the technology frontier) because countries that start to grow later than 1800 get to grow faster than the leader country in order to close the income gap. The model is set up such that long-run annual growth is 2 percent for all growing countries, so eventually world average annual growth will stabilize at 2 percent when all countries are growing at the same rate as the leader country. The income inequality measure also peaks in the 1970s. It starts at zero, when no country has yet started to grow and all have the same per capita income. It then rises until the 1970s, after which it declines; eventually it will return to zero when all countries are growing at a rate of 2 percent. Thus, the calibrated model predicts that the maximum dispersion in per capita income across countries occurred around 1970, and inequality will subsequently decline as more and more countries begin to experience modern growth and fewer are left in stagnation.

We see then how the model helps us think about the plausibility of the "pause before a return to further inequality" story and the "shape of things to come" story mentioned earlier. The basic force in the model is that industrialization diffuses through-

out the world economy and eventually all countries enter the club of growing economies. What happened in the past guides the model's predictions about what will happen in the future. Eventually all countries get in on the act, and incomes become less and less unequal across countries. In this sense, the stability of the income distribution over the past 40 years or so shouldn't be extrapolated into the future to

Each country is unique and the particular combination of factors that will push an economy over the threshold to sustained growth varies from country to country.

predict permanent large income differences across countries.

Of course, the basic model is very simple and lacks many features that we expect to influence the growth paths of individual economies and the transition of stagnant economies to growth economies. For example, no business cycles are built into the model: Once countries start to grow, they continue to grow without experiencing recessions. There are no disasters like wars that could have long-lasting effects on growth paths. In addition, the model does not incorporate the demographic transition we spoke of earlier or capital flows across countries. Despite these omissions, the model is basically in accord with the evidence on cross-country income growth since 1800. Importantly, it provides a plausible guide for thinking about the future of the world income distribution, one that is optimistic about the prospects. It may take a long time, but eventually

all countries move from pre-industrial to post-industrial.

POLICY PRESCRIPTIONS

The Lucas model assumes that poor countries will eventually develop the environment necessary for sustained growth in per capita income to begin. It does not offer prescriptions for policymakers on how to make that happen. What can policymakers in poor countries do to push their countries into the sustained-growth phase? This question has generated enormous interest on the part of economists and policymakers. Perhaps the best that can be said is that each country is unique and the particular combination of factors that will push an economy over the threshold to sustained growth varies from country to country. However, rich countries do appear to share some common factors that suggest directions for policymakers in poor countries.

Recent work by Stephen Parente and Edward Prescott examines the question of why some countries are richer than others. They come to the view that cross-country differences in incomes can be traced to the differential knowledge that societies apply to the production of goods and services. It's not that societies differ fundamentally in the knowledge available to them; that knowledge is largely available to all countries by observing the methods and practices of advanced countries, by trading with advanced countries, or by licensing advanced technologies. Rather, Parente and Prescott conclude that poor countries do not fully exploit the existing stock of usable knowledge because poor countries implement policies that constrain work practices and hinder firms' ability to implement more advanced production methods. These barriers are often put in place to protect the vested interests of entrenched groups.

Parente and Prescott reach this conclusion in part because they find that differences in savings rates across countries are unable to account for differences in international incomes. This is so even when they define savings broadly to include intangible capital and human capital. Intangible capital includes things like expenditures allocated to developing and launching new products and expenditures on increasing the efficiency of existing practices. Human capital includes the acquired knowledge that individuals obtain from education and on-the-job training. In Parente and Prescott's reading, differences in savings rates across countries can account for only a small portion of the international differences in incomes we observe.

In a famous 1988 paper, Robert Lucas laid out a theoretical case for the role of human capital in the analysis of growth and development. Subsequent research, including that of Parente and Prescott, failed to find a large role for disparities in human capital in accounting for income differences. However, recent work by Rodolfo Manuelli and Ananth Seshadri and by Andres Erosa, Tatyana Koreshkova, and Diego Restuccia re-evaluates the role that differences in human capital across countries play in accounting for income differences. Manuelli and Seshadri argue that when one measures human capital correctly, paying careful attention to differences in the quality of human capital across countries, it turns out that differences in human capital play a large role in accounting for income disparities. Indeed, in some versions of their model, the differences in the stock of usable knowledge across countries that was emphasized by Parente and Prescott play little, if any, role in accounting for differences in income across countries; the difference is largely accounted for by accumulated factors: labor and capital broadly defined to include human capital. In Manuelli and Seshadri's model, policymakers would do well to find ways to facilitate human capital accumulation among the residents of poor countries.

Erosa, Koreshkova, and Restuccia similarly emphasize the role of human capital in accounting for cross-country income differences. But their model differs from that of Manuelli and Seshadri in important details, such as allowing wage growth over people's lifetimes to be influenced by investment in physical capital and by technological progress. They argue that technology is strongly amplified by the stock of human capital. Their finding of significant differences in the average level of human capital across countries then goes a long way toward explaining cross-country differences in per capita incomes.

The somewhat opposing findings of Parente and Prescott vs. those of Manuelli and Seshadri and Erosa, Koreshkova, and Restuccia highlight some of the difficulties researchers face in studying cross-country income differences. The data on growth and development experiences of countries do not point to an obvious, unique recipe for making poor countries rich. The development process is the outcome of a complex mixture of policies and institutions. Securing property rights and creating a level playing field are likely to help individuals and firms acquire the capital and technology required to push an economy onto a sustainable growth path. 🅑🅡

REFERENCES

Boarini, Romina, Asa Johansson, and Marco Mira d'Ercole. "Alternative Measures of Well-Being," OECD Working Paper 33 (2006).

Boserup, Esther. *Population and Technological Change: A Study of Long-Term Trends.* Chicago: University of Chicago Press, 1981.

Eros, Andres, Tatyana Koreshkova, and Diego Restuccia, "How Important Is Human Capital? A Quantitative Theory Assessment of World Income Inequality," manuscript, 2007.

Hansen, Gary D., and Edward C. Prescott. "Malthus to Solow," *American Economic Review*, 92:4 (2002), pp. 1205-17.

Heston, Alan, Robert Summers, and Bettina Aten. Penn World Table Version 6.2, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania (September 2006).

Khan, Auhbik. "The Industrial Revolution and the Demographic Transition," Federal Reserve Bank of Philadelphia *Business Review* (First Quarter 2008).

Lucas, Robert E. Jr. "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 22 (1988), pp. 3-42.

Lucas, Robert E. Jr. "Some Macroeconomics for the 21st Century," *Journal of Economic Perspectives* (Winter 2000), pp. 159-68.

Lucas, R.E., Jr. "The Industrial Revolution: Past and Future," Federal Reserve Bank of Minneapolis Annual Report (2003).

Maddison, Angus. *Contours of the World Economy, 1-2030 AD: Essays in Macroeconomic History.* Oxford: Oxford University Press, 2007.

Manuelli, Rodolfo, and Ananth Seshadri. "Human Capital and the Wealth of Nations," Working Paper, University of Wisconsin-Madison.

Parente, Stephen L., and Edward C. Prescott. *Barriers to Riches.* Cambridge, MA: MIT Press, 2002.

Sala-i-Martin, Xavier. "The Disturbing 'Rise' in Global Income Inequality," NBER Working Paper 8904 (2002).