NBER
Macroeconomics
Annual 2000

Editors
Ben S. Bernanke and
Kenneth Rogoff

Maurice Obstfeld and Kenneth Rogoff
UNIVERSITY OF CALIFORNIA, BERKELEY AND NBER;
AND HARVARD UNIVERSITY AND NBER

The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?

1. Introduction

International macroeconomics is a field replete with truly perplexing puzzles, and we generally have five to ten (or more) alternative answers to each of them. These answers are typically very clever but far from thoroughly convincing, and so the puzzles remain. Why do people seem to have such a strong preference for consumption of their home goods (the home-bias-in-trade puzzle)? Why do observed OECD current-account imbalances tend to be so small relative to saving and investment when measured over any sustained period (the Feldstein–Horioka puzzle)? Why do home investors overwhelmingly prefer to hold home equity assets (the home-bias portfolio puzzle)? Why isn't consumption more highly correlated across OECD countries (the consumption correlations puzzle)? How is it possible that the half-life of real exchange-rate innovations can be three to four years (the purchasing-power-parity puzzle)? Why are exchange rates so volatile and so apparently disconnected from fundamentals [the exchange-rate disconnect puzzle, of which the Meese–Rogoff (1983) forecasting puzzle and the Baxter–Stockman (1989) neutrality-of-exchange-rate-regime puzzle are manifestations]?

What we attempt to do in this paper is to offer a unified basis for

understanding *all* of these puzzles, in which the key friction is a (significant but plausible) level of international trade costs in *goods* markets. These trade costs may include transport costs but also tariffs, nontariff barriers, and possibly other broader factors that impede trade.

We do not pretend to be the first to make this connection. In a fundamental contribution to the literature on international trade and finance, Samuelson (1954) argued that the existence of an international transfer problem depends critically on whether there is a home bias in consumption, and he showed explicitly how a home bias could be derived from transport costs.[1] In subsequent research, however, Samuelson's straightforward approach has generally been abandoned in favor of a more stylized paradigm based on breaking up a country's products into two dichotomous categories, traded and nontraded goods.[2] The analysis of the present paper suggests that for many purposes, this dichotomous grouping is far less helpful than the natural alternative of simply introducing trade costs.

Especially in our treatment of capital-market anomalies, the approach in this paper differs from the one that is conventionally taken in the literature. Typically, an author chooses from a menu of plausible capital-market imperfections the one best suited to explain a particular puzzle. We do not deny the importance of a variety of imperfections peculiar to international asset markets. Our goal here, however, is to show how far one can go in elucidating major empirical riddles without appealing to intrinsically international capital-market imperfections. Remarkably, we find that once one allows for trade costs in *goods* markets, many of the main empirical objections to the canonical models of international macroeconomics disappear. Our approach, which is based on a very simple stylized model, seems to be particularly successful in resolving the realside quantity puzzles. To explain adequately the various pricing puzzles, we would need to develop a much richer framework featuring imperfect

1. Obstfeld and Rogoff (1996, Chapter 4) embed trade costs in a version of the Dornbusch–Fischer–Samuelson (1977) Ricardian model and show that Samuelson's transfer-problem analysis can be extended to a modern dynamic setting. See Krugman (1991) on the relevance of the transfer problem to contemporary debates in international macroeconomics.
2. A notable exception is Backus, Kehoe, and Kydland (1992), who find that their approximate method for incorporating small trade costs does not resolve the consumption correlations puzzle in a calibrated one-good global real-business-cycle model. Another is Dumas (1992), who looks at a dynamic, stochastic, one-good open economy model with transport costs and explores a number of issues, including the forward exchange-rate premium. His work is theoretical and qualitative, however, and he does not calibrate his model's empirical implications for the various puzzles we look at here. Also, our main points in this paper really require an extension to the multigood case. In a more recent contribution, Ravn and Mazzenga (1999) examine further the business-cycle implications of transport costs in a variant of the Backus–Kehoe–Kydland model.

competition plus sticky prices and/or wages, as in the extensive recent literature on the "new open-economy macroeconomics." Although we do not present such a model here, we do demonstrate why trade costs must constitute an essential element, implicitly if not explicitly.

The first puzzle we address is the home-bias-in-trade puzzle (McCallum, 1995), which, as we have already noted, is closely related to the classic transfer problem. Following Wei (1998) and Evans (1999), we discuss how empirically plausible trade costs, combined with fairly standard estimates of elasticities of substitution across imports and exports, can explain much of the puzzle.

Having established the trade friction at the core of our analysis, we next turn to one of the most robust and intractable puzzles in international finance, the Feldstein–Horioka puzzle. We show that trade costs can create a wedge between the effective real interest rates faced by borrowers and lenders. In our model, the effect is highly nonlinear, manifesting itself strongly only when current-account imbalances become very large. We argue that it is precisely such incipient real-interest-rate effects that keep observed current-account imbalances within a modest range. Though we rely primarily on the theoretical force of the argument, we do demonstrate empirically that current-account-deficit countries tend to have higher real interest rates, as our model predicts.

Next, we show that the same approach can simply and elegantly explain the widely discussed home bias in equity holdings (or, more generally, in overall asset positions). The following section covers the consumption correlations puzzle, which is closely related to the preceding three, so it is not surprising our same approach again applies. We also briefly address other related puzzles.

We largely ignore nominal rigidities in our discussion of the first four puzzles, because our main argument does not depend upon having them. But as we turn to our last two puzzles—the purchasing-power-parity puzzle and the exchange-rate disconnect puzzle—we obviously must think about adding other ingredients, in the form of imperfect competition and some degree of price or wage rigidity. We nevertheless argue that trade costs in output markets must be an essential ingredient in resolving these puzzles as well. The final section concludes and also evaluates our results in the light of long-term trends in world transport costs.

2. The Puzzle of Home Bias in Trade (Puzzle 1)

The starting point for all the puzzles we examine is the growing evidence that international *goods* markets appear to be far more segmented than is

commonly supposed. Perhaps the most dramatic suggestion of segmentation stems from the work of McCallum (1995). Using a simple Tinbergen (1962) gravity model of trade that controls for distance, trading-partner sizes, and a small number of other factors, McCallum found that trade among individual Canadian provinces was *twenty* times greater than trade between individual Canadian provinces and individual U.S. states, a surprisingly high differential. It is true that the subsequent literature has both tempered McCallum's estimates and challenged their interpretation. McCallum's calculations were based on the year 1988, still at the dawn of the U.S.–Canada free-trade agreement, and before trade patterns had time to adjust fully. Using data for 1993–1996, Helliwell (1998) found that the unexplained home bias had fallen to a factor of 12, which remains a surprisingly large number. Though intracountry trade data are available only for Canada and the United States, Wei (1998) and Evans (1999) use indirect methods to test home bias for other OECD country pairs. Wei suggests that the average bias may be as low as 2.5, while Evans finds values intermediate between Wei's and Helliwell's.[3] Van Wincoop (2000) argues that even though McCallum controls for state and province size in his gravity equation, his trade-diversion measure gives an exaggerated impression of home bias in global trade because it calculates the bias from the perspective of the small country, Canada, rather than from the perspective of the large country, the United States.[4] Overall, a balanced interpretation of the literature is that countries do exhibit a considerable degree of home bias in trade, but the bias is not as extreme as McCallum's original estimates suggested.

But if there is still a significant degree of home bias in international trade, how can we explain it? Clearly, international trade does involve added border costs such as tariffs, nontariff barriers, and exchange-rate risk (and it is also possible that domestic transportation costs are lower due to greater coordination problems in constructing international transportation networks). Do these border costs need to be implausibly large to generate observed home bias, even in the more modest range of Wei's

3. Wei (1998) tries to estimate home bias indirectly by assuming that the amount a country imports from itself is the difference between total production and total exports. However, Wei's 2.5 home bias estimate could be downward biased due to his exclusion of the service sector. Evans (1999) uses data on selected industries for a number of OECD countries.

4. Van Wincoop (2000) shows that McCallum's measure of trade bias must be carefully interpreted to ascertain the negative border effect on U.S.–Canada trade. Because Canada's economy is so small relative to that of the United States, a moderate percentage diversion of U.S.–Canada trade into intra-Canada trade amounts to a spectacular percentage increase in intra-Canada trade. Using U.S. interstate trade data, van Wincoop estimates that the U.S.–Canada border reduces trade between the two countries by at most 30%.

and Evans' estimates? Not necessarily, since what really matters is the interaction between border costs and the elasticity of substitution between home and foreign goods. As this is a recurring theme in our discussion of the various quantity puzzles, it is helpful to take up a simple illustrative example.

2.1 A MODEL OF THE INTERACTION BETWEEN TRADE COSTS AND THE PRICE ELASTICITY OF DEMAND

Here, we show how costs of international trade can dramatically skew domestic consumption in favor of home-produced goods.

Consider an extremely simple two-country endowment economy, in which the utility function of the representative home consumer is given by

$$C = \left(C_H^{(\theta-1)/\theta} + C_F^{(\theta-1)/\theta} \right)^{\theta/(\theta-1)}, \tag{1}$$

where C_H is home consumption of the home-produced good and C_F is home consumption of the foreign-produced good. Foreign agents are assumed to have identical utility functions in C_H^* and C_F^*. Home agents are endowed with Y_H per capita of the home good, and foreign agents are endowed with Y_F. We assume *iceberg* shipping costs τ, so that for every unit of home (foreign) good shipped abroad, only a fraction $1 - \tau$ arrives at the foreign (home) shore. Let P_H (P_F) be the home price of the home (foreign) good, and P_H^* (P_F^*) the corresponding foreign prices, with all prices measured in terms of a common world monetary unit. (Since we are in a flexible-price world here, it is not important whether the two countries share a common currency.) Then, if markets are competitive, arbitrage implies that

$$P_F = P_F^*/(1 - \tau), \tag{2}$$

$$P_H = (1 - \tau)P_H^*. \tag{3}$$

Thus, if $p \equiv P_F/P_H$, and $p^* \equiv P_F^*/P_H^*$,

$$p^* = p(1 - \tau)^2. \tag{4}$$

(We will maintain the assumption of competitive markets through the first four sections, though our main points would still apply in an imperfectly competitive setting, as in our discussion of puzzles 5 and 6.)

From the first-order conditions for utility maximization by home and foreign agents, we have

$$\frac{C_H}{C_F} = p^\theta, \qquad \frac{C_H^*}{C_F^*} = (p^*)^\theta. \tag{5}$$

Combining (4) and (5) implies

$$\frac{C_H}{C_F} = (1 - \tau)^{-2\theta} \frac{C_H^*}{C_F^*}. \tag{6}$$

For illustrative purposes, consider the easy symmetric case in which $Y_H = Y_F$. Under that assumption, $C_H/C_F = C_F^*/C_H^*$ and equation (6) reduces to

$$\frac{C_H}{C_F} = \frac{C_F^*}{C_H^*} = (1 - \tau)^{-\theta} = p^\theta.$$

This equation shows that the ratio of home (foreign) expenditure on imports relative to home (foreign) goods is

$$\frac{C_H}{pC_F} = \frac{p^*C_F^*}{C_H^*} = (1 - \tau)^{1-\theta}.$$

Thus, for example, if there were no trade costs ($\tau = 0$), then $pC_F/C_H = 1$. If $\tau = 0.25$ (a large number just for goods actually traded but conservative when applied to all of GNP) and $\theta = 6$, then $C_H/pC_F = 4.2$. This ratio is consistent with those we observe for many OECD countries, and the degree of home bias can easily be made larger by raising τ, raising θ, or assuming that the home country is a small one trading with many like-sized foreign partners.

2.2 THE NONLINEAR RELATIONSHIP BETWEEN TRADE COSTS AND HOME BIAS IN TRADE

The higher trade costs (the closer τ is to 1), the greater the impact of a 1% reduction in τ on home bias:

$$\frac{d \log(C_H/pC_F)}{d \log \tau} = \frac{\tau}{1 - \tau}(\theta - 1).$$

For our baseline case of $\tau = 0.25$ and $\theta = 6$, the elasticity of home bias with respect to trade costs is $\tau(\theta - 1)/(1 - \tau) = 1.67$.

Obviously, this example is wildly oversimplified. It implicitly assumes a common substitution elasticity across any individual pair of home and foreign goods, and similarly lumps all goods together as having common trade costs. It ignores the potential importance of substitution between domestic and foreign inputs in production. Nevertheless, it neatly illustrates how a high elasticity of substitution can explain a large observed home trade bias even with low trade costs. What then are plausible values for the parameters τ and θ?

2.3 EMPIRICAL ESTIMATES OF θ

Though there is a range of estimates for θ, recent trade studies typically find values for the elasticity of import demand with respect to price (relative to the overall domestic consumption basket) in the neighborhood of 5 to 6. Examples include Trefler and Lai (1999), who present panel estimates over 1972–1992 for a panel of 28 industries in 36 countries; their preferred estimate is 5.3. That average number reflects estimated disaggregated substitution elasticities as high as 21.4 (for industrial chemicals) and 18.9 (for electrical machinery and electronics) but as low as 1.2 (for printing and publishing). Harrigan (1993) looks at three-digit 1983 SITC data for 13 OECD countries representing 90% of OECD output and finds elasticities in the range of 5 to 12.

Recognizing that much of trade involves imperfectly competitive industries, one can attempt to infer the value of θ by looking at markups of price over marginal cost. Using that approach, Cheung, Chinn, and Fujii (1999) look at two-digit industry level data for a range of OECD countries, and impute elasticities typically in the range of 3.5 to 4. Hummels (1999a) tries to disentangle the effects of trade elasticities from those of substitution elasticities within a cross-section framework. Using linear least squares, he comes up with an average markup of 22%, translating into a θ of 5.6, although other of his estimates of θ are higher. Finally, in their classic article on the demand for automobiles—including both domestic and foreign makes—Berry, Levinsohn, and Pakes (1995) find price elasticities of demand between 3.1 and 6.4.[5]

Of course, these studies refer to goods actually traded. As Hummels emphasizes, one would expect that elasticities of substitution would be higher on average for goods that are not traded. In this case, an estimate

5. Studies of monopoly markups in domestic sales, while not necessarily directly applicable here, also yield similar estimates for θ. For example, Rotemberg and Woodford (1992) find a markup for the United States of around 20%, corresponding to $\theta = 6$. In subsequent discussion, Rotemberg and Woodford (1995) argue that there is great uncertainty about actual markups in U.S. industry, but favor estimates in the range of 20% to 40%, that is, θ between about 3.5 and 6.

of θ = 20, as Wei proposes, does not seem so wild-eyed. Brown and Stern (1989) use θ = 15 for their policy experiments.

2.4 EMPIRICAL ESTIMATES OF TRADE COSTS τ

There is far less consensus about the size of trade costs, which include (among other things), tariffs, nontariff barriers, and transport costs. For 1993, average tariffs, on a domestic-production-weighted basis, were 4.9% for the United States, 7.7% for the European Union, 3.5% for Japan, and 8.9% for Canada.[6] As for nontariff barriers, official statistics only give information on their existence, not their effectiveness, which must be estimated using an economic model. Anderson and Neary (1998) use a simple computable general equilibrium model to estimate uniform tariff equivalents for nontariff barriers for a broad range of countries, and typically find estimates on the same order of magnitude as for tariffs, larger of course for some countries (such as Japan) than for others (such as the United States). This result is also consistent with the trade-equation estimates of Lee and Swagel (1997).[7]

Differential international transportation costs are also an important potential element of τ.[8] If one looks across all commodities on a value-weighted basis, freight and insurance charges for U.S. imports averaged 3.6% in 1995, and 3.3% in 1996 and 1997.[9] But these numbers considerably understate average costs in international shipping. First, As Hummels (1999a) shows, average costs are much higher for many other countries (the United States has a vast coastline, and sea shipping tends to be much cheaper than shipping by land). Second, these numbers do not include other considerable costs of international shipping, including preparing the paperwork (bills of lading) needed to clear international customs, and the costs of delays either in transit or at port of entry.

Just as empirical measures of the elasticity of substitution between home and foreign goods may be biased downwards, it is also likely that simple estimates of average transport costs grossly understate average τ across all goods in the economy (due to substitution effects). Table 1 is drawn from Hummels (1999a), who based his estimates on highly disag-

---

6. See OECD (1996, Table 1.1, row 9).
7. Harrigan (1993), however, finds nontariff barriers insignificant compared to tariffs and transportation costs.
8. Recall that Helliwell's and McCallum's estimates use distance in an attempt to control for transport costs. Geographical distance is an imperfect measure of these, however.
9. The authors are grateful to Robert Feenstra for compiling these numbers based on *U.S. Imports of Merchandise*, U.S. Census Bureau. The estimates give shipping and freight charges as a percentage of total value of imports excluding these charges. Importantly, these numbers do not include any inland shipping at point of departure or port of arrival.

Table 1   COMMODITY DISTRIBUTION OF FREIGHT RATES (UNITED STATES, 1994)[a]

Commodity	Average freight rate	
	Trade-weighted	unweighted
All goods	3.8	
Food and live animals	8.2	14.1
Beverages and tobacco	6.9	14.4
Crude materials	8.2	15.1
Mineral fuels, lubricants	6.6	15.7
Animal and veg. oils, fats	7.1	10.6
Chemicals and related products	4.5	9.0
Manuf. goods (by material)	5.3	10.3
Machinery and transp. equip.	2.0	5.7
Misc. manufactures	4.7	8.3
All other goods, NES	1.0	2.5

[a]Source: Hummels (1999a), compiled from U.S. Census Bureau, *U.S. Imports of Merchandise*. Freight costs include shipping and insurance as a percentage of total FAS value. Calculations are based on 10-digit Harmonized System level data (over 15,000 categories of goods). Unweighted shipping costs are based on all individual goods imported.

gregated 10-digit data. We see from the table that shipping costs for many categories of goods are quite a bit larger than the average (trade-weighted) shipping costs—and this table excludes goods that are not traded at all.

What other factors might be included in τ? In a provocative paper, Rose (2000) uses a gravity model to argue that countries with currency unions trade two to three times as much with each other as countries with separate currencies. Certainly, currency conversion costs and exchange-rate uncertainty can add to trade costs. While exchange-rate variability can have direct negative effects on capital flows, any direct negative effect on trade flows will result in an additional, indirect source of capital-market imperfection according to our analysis. A similar point can be made for various informational costs of international trade; see Rauch (1999) and Portes and Rey (1999) for discussion and some empirical evidence.[10] Differences in legal and payments systems may also add to τ.

Last, but not least, it is important to emphasize that our analysis has assumed no home bias in preferences. Suppose we replace the representative home agent's utility function (1) with

10. Anderson and Marcouiller (1999) argue that corruption and imperfect contract enforcement are major factors in disrupting trade, especially in developing countries. Since our main focus is on industrialized countries, we will not consider these categories of trade cost further here.

$$U = (C_H^{(\theta-1)/\theta} + \omega C_F^{(\theta-1)/\theta})^{\theta/(\theta-1)} \tag{7}$$

and the representative foreign agent's utility function with

$$U^* = (\omega C_H^{*(\theta-1)/\theta} + C_F^{*(\theta-1)/\theta})^{\theta/(\theta-1)}.$$

One can easily show that the effects of home bias in preferences ($\omega < 1$) can be isomorphic to the effects of trade costs τ. Helpman (1999) argues that once one controls for income, there is no clear evidence of home bias in preferences. Indeed, it is more illuminating to derive trade biases from other frictions. Nevertheless, it is important to recognize that a home bias in demand for goods can work similarly to trade costs, at least for the trade and portfolio-bias puzzles.

2.5 OTHER REAL-TRADE PUZZLES

Though international-finance puzzles are our main focus, we note that trade costs can explain a number of real-trade puzzles as well. For example, Trefler's (1995) favored explanation of the "missing trade" puzzle combines Hicks-neutral productivity differences across countries with a home bias in consumption (which, per our discussion above, may be induced by transport costs).[11] Deardorff (1998) points out that with transport costs, the standard conditions for factor-price equalization in a Heckscher–Ohlin world break down, also implying greater specialization. Since factor-price equalization fails miserably empirically, this implies another important puzzle that can be at least partially resolved by transport costs. Anderson (1979), Deardorff (1998), and others have shown that transport costs can help explain the surprising empirical robustness of the gravity equation of trade flows. Not only do trade costs help to resolve a number of puzzles in the data, they also seem to be important in determining economic performance. Radelet and Sachs (1998) argue that countries that have high shipping costs due to adverse geography (for example, high mountains or limited port access) grow much more slowly than countries with natural transport advantages. Finally, we note that evidence on international price differentials seems quite consistent with the high degree of market segmentation evinced on the quantities side; we shall refer to this work later in discussing puzzles 5 and 6.

Armed with a simple understanding of how plausible trade costs together with high elasticities of substitution in consumption can explain

---

11. "Missing trade" is how Trefler describes the puzzle that the imputed factor content of trade does not seem to reflect comparative advantage.

substantial home biases in trade, we are ready to explore the linkages to other macro puzzles.

3. The Feldstein–Horioka Puzzle (Puzzle 2)

There has been no shortage of explanations for the famous saving–investment puzzle of Feldstein and Horioka (1980), with numerous articles on the topic having been published in most of the leading journals. The problem is that none of the explanations advanced to date (including our own attempts) has been terribly convincing. Most explanations tend to be clever but empirically inadequate and, more troublesome still, tend to fix one puzzle at the expense of creating others. The fact that the Feldstein–Horioka regularity does not seem to characterize intranational regional data suggests that factors intrinsic to trade between different nations are at work.[12]

3.1 STILL A PUZZLE

What Feldstein and Horioka actually demonstrated, of course, is that across OECD countries, long-period averages of national saving rates are highly correlated with similar averages of domestic investment rates. Indeed, in the original data sample examined by Feldstein and Horioka, covering 1960 through the mid-1970s, cross-section regressions of investment on saving yielded slope coefficients near unity. True, this Feldstein–Horioka coefficient—which the original work interpreted as measuring the effect of the saving rate on the investment rate, or a "savings retention" measure—has fallen over time. As Table 2 illustrates, however, it still remains large and significant. The table gives simple cross-country regressions of investment (relative to GDP) against national saving (relative to GDP), taking eight-year averages for the most recent period, 1990–1997. For the OECD countries, the coefficient (0.60) is a good deal smaller than the 0.89 found in Feldstein and Horioka's original work, but it is still larger than one might expect in a world of fully integrated capital markets where global savings should flow to the regions with the highest rates of return. The coefficient falls further once one includes countries outside the OECD (particularly poor countries), although the extended results must be viewed with extreme caution given the poor quality of national income and product data for most non-OECD countries. (The data underlying the regressions in Table 2 are reported in Table 7 in the appendix, which also describes how the countries in the sample were chosen.)

12. See Obstfeld (1995), Obstfeld and Rogoff (1996), and Coakley, Kulasi, and Smith (1998) for recent surveys.

Table 2 FELDSTEIN–HORIOKA REGRESSIONS, $I/Y = \alpha + \beta\,NS/Y + \epsilon$, 1990–1997[a]

	No. of obs.	α	β	R^2
All countries[b]	56	0.15 (0.02)	0.41 (0.08)	0.33
Countries with GNP/cap. > 1000	48	0.13 (0.02)	0.48 (0.09)	0.39
Countries with GNP/cap. > 2000	41	0.07 (0.02)	0.70 (0.09)	0.62
OECD countries[c]	24	0.08 (0.02)	0.60 (0.09)	0.68

[a]OLS regressions. Standard errors in parentheses.
[b]Israel is excluded from all regressions in this table. If Israel is added to the samples of size (56, 48, 41), the estimates of β are (0.39, 0.45, 0.63).
[c]If one adds Korea to the OECD sample, the estimate for β rises to 0.76. Korea is included in the larger samples.

The Feldstein–Horioka puzzle is durable because the core regression simply summarizes in a compact way the fact that OECD current accounts tend to be surprisingly small relative to total saving and investment, especially when one averages over any sustained period. For developing countries, notably the many that have repeatedly had trouble servicing debts, it is perhaps not so surprising that creditors prevent them from running up large sustained deficits. But it is hard to appeal to sovereign-default risk for OECD countries, especially when one considers that gross international flows of financial assets are much bigger than *net* international flows. Indeed, for OECD countries, asset price comparisons suggest a high degree of integration; arbitrage in similar nominally risk-free assets appears to be nearly perfect. We leave it to the reader to look at other sources (for example, Obstfeld and Rogoff, 1996, Chapter 3) for assessments of previous attempts to explain the Feldstein–Horioka conundrum.

A fair summary of the literature is that there are at least five or six leading explanations (and ten or so close seconds). All are unconvincing empirically—some because they are based on very special assumptions about the nature of the exogenous shocks (e.g., Obstfeld, 1986, or Mendoza, 1991), others because they raise collateral empirical contradictions. For example, in the asymmetric information model of Gordon and Bovenberg (1996), a "lemons" problem is invoked to explain why foreigners finance so little domestic investment, yet departures from covered

interest parity must also be assumed if there is to be any foreign equity inflow at all. Explanations that try to maintain the assumption of perfect capital mobility often have the strong implication that one should also observe high saving–investment correlations across states or regions within a given country. But the partial evidence available on saving and investment by subnational regions simply does not produce the Feldstein–Horioka regularity; see, for example, Helliwell (1998, Chapter 4). We are going to propose here an entirely new explanation, based on transaction costs for international trade in goods. An especially attractive feature of our approach is that it seems to help resolve other puzzles rather than exacerbating them.

It is important to emphasize that whereas our model includes trade costs for goods, it is consistent with free and costless trade in securities. Thus, it is perfectly consistent with the observation that gross international flows of securities are substantial even though net flows are small. Our account is also notable both for endogenizing the price and interest effects of trade impediments, and for showing how moderate transport costs could generate empirically significant international differences in real interest rates despite full asset-market integration.

3.2 TRANSPORT COSTS CAN INDUCE A NONLINEAR RELATIONSHIP BETWEEN THE CURRENT ACCOUNT AND THE REAL INTEREST RATE

The basic intuition of why transport costs can temper current-account imbalances can be illustrated in a standard two-period, two-good, small-country endowment model. It would not be difficult to endogenize the world real interest rate, or to incorporate uncertainty (as in the next section), but neither generalization is essential here. We will later discuss investment to confirm that the basic argument we make still goes through.

The model below is entirely standard except that we will again allow for Samuelsonian "iceberg" costs in trade, so that τ percent of any good is lost in transit. The utility function of a representative home resident is

$$u(C_1) + \delta u(C_2),$$

where total real consumption C depends on consumption of the home and foreign goods, C_H and C_F, with constant elasticity of substitution θ as in equation (1). The small country is endowed only with good H, with $Y_{H,1}$ in period 1, and $Y_{H,2}$ in period 2. Good F must always be imported. (Endowing the country with both goods would not overturn our argument.) The home country is small in the sense that its actions have no

effect on the world prices P_H^* and P_F^*, which are constant across the two periods in terms of a world unit of account (money). Nor can it affect the foreign real interest rate r^* (which equals the foreign nominal interest rate assuming there is zero foreign inflation). Because of iceberg transit costs τ in shipping either good, however, home consumption patterns can affect home relative prices and the home real interest rate.

Though we shall give a formal analysis below, the basic argument is simple. Suppose, for example, that the country's endowment pattern and rate of time preference δ are such that in the first period, net exports of good H are negative (in which case intertemporal solvency dictates they must be positive in period 2). Then, as we shall shortly confirm, the relative price of good H will be higher in period 1 than in period 2. There will be expected deflation, and the home real interest rate will be above r^*. The situation is reversed when the country is initially running a sufficiently large current-account surplus, so that its real consumption-based lending rate must lie below r^*. As we demonstrate formally below, this effect can be quite dramatic, assuming realistic values for trade costs τ and the elasticity of substitution θ (values similar to those needed to resolve the home-bias-in-trade puzzle).[13]

3.3 BUDGET CONSTRAINTS AND TRANSPORT COSTS

A formal analysis requires one to think carefully about the budget constraints facing the representative agent. In general, the first-period budget constraint can be written as

$$P_{H,1}Y_{H,1} + D = P_{H,1}C_{H,1} + P_{F,1}C_{F,1} = P_1C_1,$$

where P_H (P_F) is the home-soil price of good H (F) in terms of the world currency unit, and D is borrowing from abroad in world currency units. The overall home price level, in terms of world currency units, is

$$P = (P_H^{1-\theta} + P_F^{1-\theta})^{1/(1-\theta)}. \tag{8}$$

Therefore, given a total real consumption level C in any period, the consumptions of the two individual home and foreign goods are

$$C_H = \left(\frac{P_H}{P}\right)^{-\theta} C, \qquad C_F = \left(\frac{P_F}{P}\right)^{-\theta} C. \tag{9}$$

---

13. Dumas (1992) likewise shows how international real interest-rate differentials can arise in a model with transport costs, though, as we have already noted, his one-good model is very different and he does not explore the implications for the Feldstein–Horioka puzzle.

Similarly, the second-period budget constraint, measured in world currency units, is

$$P_{H,2}Y_{H,2} - (1 + r^*)D = P_{H,2}C_{H,2} + P_{F,2}C_{F,2} = P_2C_2.$$

Combining the period budget constraints gives the consolidated intertemporal budget constraint as

$$P_1C_1 + \frac{P_2C_2}{1 + r^*} = P_{H,1}Y_{H,1} + \frac{P_{H,2}Y_{H,2}}{1 + r^*},$$

or, in terms of the domestic real interest rate $1 + r = (1 + r^*)P_1/P_2$, as

$$C_1 + \frac{C_2}{1 + r} = \frac{P_{H,1}Y_{H,1}}{P_1} + \left(\frac{1}{1 + r} \right) \frac{P_{H,2}Y_{H,2}}{P_2}. \tag{10}$$

3.4 INFLATION

Since good F is always imported, its home price, $P_F = P_F^*/(1 - \tau)$, must be higher than the foreign price in both periods, per equation (2). By the same logic, when good H is exported— as it must be in at least one of the two periods—its home price $P_H = P_H^*(1 - \tau)$ must be lower than the foreign price, per equation (3). However, if total domestic spending is high enough relative to income in any given period, it is possible that good H is *imported* rather than exported ($C_H > Y_H$), in which case its home price $P_H = P_H^*/(1 - \tau)$ must be *higher* than the foreign price. As we shall see, there are also important intermediate cases where $C_H = Y_H$ in one period, in which case P_H will turn out to lie between $P_H^*(1 - \tau)$ and $P_H^*/(1 - \tau)$, despite the fact that no goods roundtrip.

3.5 A GRAPHICAL ANALYSIS OF THE LINK BETWEEN REAL INTEREST RATES AND CURRENT ACCOUNTS

The link between the effective real interest rate faced by the home country and its first-period borrowing decision is illustrated in Figure 1, which plots total real consumption in period 1, C_1, against the domestic real interest rate, $1 + r$. (Note that the period 1 current account deficit is simply $Y_1 - C_1$). The resulting graph is a step function that can be divided up into five segments; it shows the schedule of effective real interest rates faced by the country as a function of its borrowing–lending decision.

In the first segment C_1 is so low, and the period 1 current-account surplus so high, that in period 2 the country will consume an amount

Figure 1 DOMESTIC SPENDING AND THE DOMESTIC REAL INTEREST
RATE IN A TWO-GOOD MODEL WITH TRADE COSTS
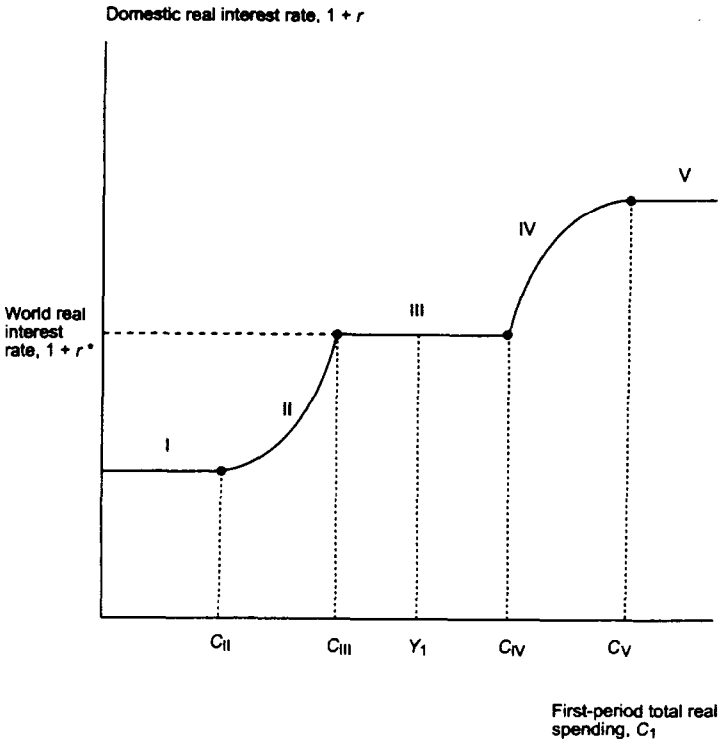
$C_{H,2} > Y_{H,2}$. Since in period 2 the home good must be imported, while in
period 1 it is exported, we have

$$1 + r = \frac{(1 + r^*) (P_{H,1}^{1-\theta} + P_F^{1-\theta})^{1/(1-\theta)}}{(P_{H,2}^{1-\theta} + P_F^{1-\theta})^{1/(1-\theta)}}$$

$$= \frac{(1 + r^*) \{[P_H^*(1 - \tau)]^{1-\theta} + P_F^{1-\theta}\}^{1/(1-\theta)}}{\{[P_H^*/(1 - \tau)]^{1-\theta} + P_F^{1-\theta}\}^{1/(1-\theta)}} < 1 + r^*$$

in segment I. If the country contemplates being a big lender, it will face
an effective real interest rate significantly *below* the world real interest
rate.

Segment II starts when period 1 consumption first reaches the level C_{II}
such that $C_{H,2} = Y_{H,2}$. In this region, $P_{H,2}$ is determined by equation (8) and
the first relation in equation (9), with $C_{H,2} = Y_{H,2}$ (so there is no round-

tripping). Period 2 consumption of the home good remains constant at $Y_{H,2}$ as long as $P_{H,2}$ remains strictly between $P_H^*(1 - \tau)$ and $P_H^*/(1 - \tau)$, but equation (9) implies that $P_{H,2}$ falls as C_1 rises and C_2 falls, until $P_{H,2}$ reaches $P_H^*(1 - \tau)$. Accordingly, the real interest rate *rises* over segment II.[14] At the point $C_1 = C_{III}$, Segment III begins as the home country becomes a period 2 exporter of its endowment good. On this stretch, $1 + r = 1 + r^*$. Because here, $C_H < Y_H$ in both periods, the overall price level is constant over time. In region III, the country is running a sufficiently small current-account surplus or deficit that there is never any reversal of the pattern of trade in either good. It is precisely in this region that trade costs have no effect on the real interest rate.

At $C_1 = C_{IV}$, however, $C_{H,1}$ reaches $Y_{H,1}$, and the real interest rate begins to rise once more. In segment IV, $C_{H,1}$ remains stuck at $Y_{H,1}$ as C_1 rises, pushing $P_{H,1}$ up with it until $P_{H,1}$ reaches $P_H^*/(1 - \tau)$. As $P_{H,1}$ rises along segment IV, with $P_{H,2}$ constant at $P_H^*(1 - \tau)$, the real interest rate rises. At C_V, however, where $P_{H,1}$ first reaches $P_H^*/(1 - \tau)$, the country becomes a period 1 importer of its own endowment good, and the real interest rate stabilizes (along segment V) at the level

$$1 + r = \frac{(1 + r^*) \{[P_H^*/(1 - \tau)]^{1-\theta} + P_F^{1-\theta}\}^{1/(1-\theta)}}{\{[P_H^*(1 - \tau)]^{1-\theta} + P_F^{1-\theta}\}^{1/(1-\theta)}} > 1 + r^*.$$

The range of possible real interest rates produced by this simple example can encompass a wide distribution. For example, with $r^* = 0.05$, $\tau = 0.1$, $\theta = 6$, and $P_H^* = P_F^* = 1$, we find that the highest possible real interest rate is 20% (15% above the world level) while the lowest is −8% (13% below the world level). The interplay between the commodity transport costs τ and the substitution elasticity θ is similar to what we saw in the preceding section. As θ rises, the maximum and minimum real domestic interest rates move apart—with higher substitutability, the price-level impacts of changes in P_H are more pronounced. In the limiting case as $\theta \rightarrow \infty$ the two goods are asymptotically perfect substitutes, in which case the country's effective real borrowing rate will be 30%, and its lending rate, −15%!

Of course, the range of real domestic interest rates encompassed by Figure 1 is far greater than what we usually observe in practice, especially for OECD countries. But this simply reflects the fact that incipient

14. The increasing portions of the schedule have the shapes we show for θ-values that are plausibly high.

real interest differentials put a sharp check on a country's incentives to run large current-account deficits or surpluses.[15]

3.6 EXTENSIONS AND ALTERNATIVE FORMULATIONS

The preceding account of the effect of domestic spending on real interest rates is overly stylized, but a number of obvious extensions can add to realism without diluting the main message.

3.6.1 A Continuum of Goods and Transport Costs Assume, for example, that countries are endowed with multiple goods in various proportions and that these goods display a distribution of transport costs. Then, as domestic spending rises, progressively more types of goods are imported from abroad, leading to a steadily rising real domestic interest rate. In this more realistic setup, the relationship between expenditure and the home real interest rate will still resemble a version of Figure 1, but with very many small steps—to the naked eye, a smoothly upward-sloping curve. With a rich enough range of goods, transport costs, and elasticities of substitution, even small current-account deficits may produce trade reversals in a small number of goods, thereby resulting in an interest-rate effect. We conjecture, though it remains to be proved, that one would obtain a similar nonlinearity to that depicted in Figure 1, with small current-account imbalances having relatively little effect on interest differentials.

3.6.2 Long-Term Borrowing and Lending An obvious question is how the results here might be tempered in a model with many periods so that there are opportunities for long-term borrowing and lending. For example, if a country ran a big current-account deficit in the initial periods, it could repay slowly over many periods. Though a more careful analysis is required than we can provide here, it seems unlikely that this consideration would overturn our basic point; there would still be a big price swing between the big deficit periods and surplus periods—which is precisely why a country would seek to avoid such swings. We note also that in a richer model with a continuum of goods, the current account would not necessarily have to swing between deficit and surplus to induce real-interest-rate effects. In general, the range and type of goods

15. The suggestion that idiosyncratic real-interest-rate developments might help explain the Feldstein–Horioka puzzle can be found in earlier work, for example, Frankel (1986). However, to the extent that real-interest-rate effects have been touched upon in the literature, no one has taken the idea very seriously, since earlier models could not give any reason why the real interest rate might be so important quantitatively. Nor could they really explain the durability of the Feldstein–Horioka relationship across different time periods and regimes.

being imported and/or exported will vary more or less continuously in the level of trade-balance deficit (or surplus). Thus, the real-interest-rate effect will arise along any path where there are big trade-balance swings, either over any short period, or cumulatively over any long period. This would be true even in a setting with growth in which countries could, in principle, run perpetual deficits and surpluses.

3.6.3 Investment   How is the preceding analysis affected by introducing investment? In the case where the country desires to be a large net borrower (segments IV and V), the real-interest-rate effect will be tempered to the extent the country can cut back on investment instead of borrowing from abroad. But that very mechanism dictates that reductions in national saving will be accompanied by reductions in domestic investment. In segments I and II, the country could channel some of its higher savings into higher investment, again tempering the fall in the effective real interest rate but creating the positive Feldstein–Horioka correlation between increases in saving and increases in investment.[16]

3.6.4 Deriving Similar Results in a More Conventional Setup with Traded and Nontraded Goods   The reader may well ask whether we needed such an extravagant formulation to make the basic point that the consumption-based real interest rate can be linked to the current account. Couldn't we have made the same point in the context of a standard Salter–Swan model having two classes of goods, one with infinite trade costs and the other with zero trade costs (as discussed, for example, in Chapter 4 of Obstfeld and Rogoff, 1996)? Indeed, for a pure endowment case, the standard traded–nontraded model does produce a graph very much like Figure 1. Holding endowments of both goods flat, if the country chooses to run a large deficit in period 1, the price of nontraded goods will be high in that period, and low in the following period. This implies a consumption-based real interest rate above the world interest rate, just as in segments IV and V of Figure 1, and the effect can similarly be nonlinear. We prefer our formulation largely because it is much easier to think concretely about trade costs than about the arbitrary dividing line between traded and nontraded goods. Perhaps the ideal model would be a richer one incorporating a range of transport costs in which the degree of tradability is endogenous and some goods are consistently produced exclusively for the home market.

16. Though their focus is on the short-run time-series properties of the data rather than on the Feldstein–Horioka regularity, Backus, Kehoe, and Kydland (1992) do note that a small trade cost can sharply reduce the variability of net exports in their simulations.

3.6.5 Monopoly Pricing and Sticky Prices Our analysis assumes that prices are flexible and set in competitive markets. Introducing realistic features such as price rigidity and monopoly pricing, as in our discussion of puzzles 5 and 6, would enrich the model without overturning the main points. Also, the most troubling manifestations of the Feldstein–Horioka puzzle are at medium-term horizons of five to fifteen years, when price flexibility is much greater and firms' ability to preserve monopoly power is less.

3.7 EMPIRICS

The model does contain one simple prediction that can easily be checked. Countries running current-account surpluses should have lower real interest rates than countries running deficits.

This connection is illustrated in the panel regression results reported in Table 3. Specification 1 regresses the domestic real interest rate, defined as the average three-month nominal interest rate in a given year less lagged annual inflation, on the ratio of the current-account surplus to GDP. Specification 2 forms real interest rates by using December average nominal interest rates in year *t* less year *t* inflation, in an attempt to

Table 3   REAL INTEREST RATES AND THE CURRENT ACCOUNT, 1975–1998

	Coefficient on CA/GDP	Significance	ρ	R^2
Specification 1				
OLS	−36.9	0.00	0.65	0.05
Country fixed effects	−46.3	0.00	0.65	0.08
Country fixed effects, time dummies	−32.3	0.00	0.55	0.50
Specification 2				
OLS	−17.9	0.00	0.58	0.02
Country fixed effects	−19.4	0.00	0.58	0.05
Country fixed effects, time dummies	−18.9	0.01	0.54	0.32

The dependent variable is the annualized three-month nominal interest rate less lagged annual inflation CPI rate (specification 1) or less the contemporaneous inflation rate (specification 2). The sample uses annual data and covers the years 1975–1998 and all OECD countries except Iceland, Korea, Mexico, and Turkey. Current accounts (as a percentage of GDP) are reported by the OECD. We use three-month interest rates, usually a Treasury bill rate, but an interbank rate if no government rate is available. These data come from *International Financial Statistics* and the OECD. CPI inflation rates are based on IFS data. For the specification 2 regressions, four countries did not report monthly interest-rate data until after the start of our sample. The countries, with their starting dates in parentheses, are Spain (1977), Greece (1980), Portugal (1985), and Finland (1987).

capture that agents can incorporate contemporaneous information into forming inflation expectations. In both specifications we employ an autoregressive correction. We report estimates for simple ordinary least squares (OLS), a model with country fixed effects, and a model with fixed effects and time dummies (the latter to capture global influences on national real interest rates).

The results show highly significant negative correlations between the current-account surplus and the real domestic interest rate, as our model suggests. However, the two specifications differ somewhat in their numerical predictions, with specification 1 giving an effect that is substantially larger than that given by specification 2. Taking the regressions with country fixed effects and time dummies as likely to be most reliable, we see that a 1% of GDP rise in an OECD country's current-account surplus is associated with roughly a 20- to 30-basis-point decline in its real interest rate.[17]

## 4. The Puzzle of Home Bias in Equity Portfolios (Puzzle 3)

Despite the rapid growth of international capital markets toward the close of the twentieth century and a much expanded world market for equities, stock-market investors maintain a puzzling preference for home assets. When they first highlighted the extent of the home-bias portfolio puzzle at the end of the 1980s, French and Poterba (1991) observed that Americans held roughly 94% of their equity wealth in the U.S. stock market whereas the Japanese held roughly 98% of their equity wealth at home.[18] Figure 2, drawn from Tesar and Werner (1998), suggests that the home equity bias is muted for smaller countries and has shown some tendency to decline over time—by the mid-1990s about 10% of U.S. equity wealth was invested abroad. Standard models of optimal international portfolio diversification imply, however, that equity investors still have not diversified internationally nearly as much as they should, and so the puzzle remains.[19]

17. We experimented with a number of other specifications, expected inflation proxies, and time periods, almost always finding results similar to those reported in Table 3. Gordon and Bovenberg (1996) also establish a relationship between current accounts and real interest rates for OECD countries, but their test and their specification are motivated by a model that is very different than ours.

18. See also Golub (1990), who compared gross international asset flows with gross domestic asset creation for OECD countries.

19. University of California investment policies illustrate the extent and persistence of home bias even for large, sophisticated investors. On April 20, 2000, the U.C. regents announced a revision in investment guidelines for the university's retirement and endowment funds. The overall target portfolio share for equities remained at 65%, but the recommended target share for non-U.S. equities, previously *zero*, was raised

Figure 2 HOME BIAS IN EQUITY PORTFOLIOS: 1987–1996
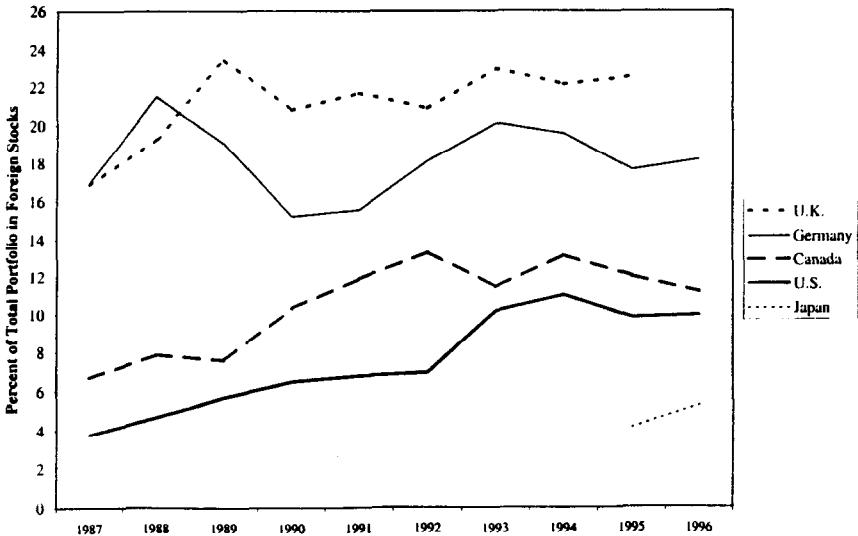
From Tesar and Werner (1998)

Potential explanations range from nontraded factors such as human capital (which may worsen or reduce the puzzle; see Baxter and Jermann, 1997) to nontraded consumption goods to asymmetries of information to data inadequacy. Yet it is fair to say that none of the available stories alone has provided a quantitatively satisfactory account of the observed home bias; see Lewis (1999) for an up-to-date and thorough survey.

To set the stage for our discussion of trade costs, it is worthwhile briefly reviewing what is perhaps the leading explanation, which is based on the classic Salter–Swan traded–nontraded-goods dichotomy we have already mentioned. While these two types of goods lie at polar extremes in terms of their tradability, equity claims on *either* type of industry can be frictionlessly traded. Thus, even though cement is prohibitively costly to transport, there is nothing to stop foreign investors from buying shares in the domestic cement industry. Earnings, of course, must be redeemed in traded goods, since nontradables cannot be shipped to foreign equity holders by assumption. The key result one gets out of this framework is that, for the baseline case of separable

to 7%. The positive target position in foreign equities, meant to "reduce risk and broaden portfolio diversification while maintaining or improving investment performance," represents a substantial advance. It still falls far short, however, of the optimal foreign equity share that simple models of international diversification would predict.

preferences (across the two types of good), investors hold a globally diversified portfolio of traded-goods industries. But nontraded-goods industries are held entirely domestically. The intuition is that, since payments can only be made in traded goods and utility is separable, there is no way to enhance risk sharing in tradables by linking the allocation of tradables consumption to returns in nontraded-goods industries. (That intuition has to be modified for the case of nonseparable preferences, but it is still a useful reference point.[20]) Thus, if nontraded goods constitute, say, 50% of total output (a popular rule of thumb based on the fact that for many OECD countries, services, construction, and transport constitute roughly 50% of GDP; see Stockman and Tesar, 1995), then agents will (loosely speaking) hold more than half their equity in home assets.

While elegant, this explanation still is not entirely satisfactory. First, although it goes some way toward explaining home bias, it falls short of explaining the 80% to 90% domestic equity shares we actually observe (Figure 2). Second, the sharp dichotomy between traded and nontraded goods is a contrived one, since in reality transport costs differ across goods, and a particular good may or may not enter trade under different market conditions. For most goods, tradability is not absolute and tradedness is endogenous.[21]

Here we will take an approach based on intuition similar to that in the preceding discussion. We explore just how far can one get in explaining the home portfolio bias by explicitly introducing trade costs, rather than splitting goods into two arbitrary and dichotomous categories. What we will show is that, with a plausible elasticity of substitution across goods and reasonable-sized costs for trading them, our model can produce a very high and realistic level of home portfolio bias.

4.1 A SIMPLE MODEL

We now add uncertainty to the two-country general equilibrium version of our model, with each country having a random endowment of its distinct perishable consumption good, along the lines of Lucas (1982) or Cole and Obstfeld (1991). To keep notation simple, we again abstract from dynamics and consider a one-period portfolio problem. We assume a completely symmetric joint distribution for the national outputs (Y_H, Y_F).

A home or foreign individual chooses state-contingent consumptions C_H and C_F of the home and foreign goods in order to maximize

---

20. Baxter, Jermann, and King (1998) develop some results for the case of nonseparable preferences.
21. For a more thorough discussion, see Obstfeld and Rogoff (1996, Chapter 5).

$$EU = E\left\{ \frac{1}{1-\rho}\left[\left(C_H^{(\theta-1)/\theta} + C_F^{(\theta-1)/\theta} \right)^{\theta/(\theta-1)} \right]^{1-\rho} \right\} = E\frac{C^{1-\rho}}{1-\rho}. \tag{11}$$

Above, C is the index of total real consumption [per equation (1)], θ is consumers' elasticity of substitution between the two goods, and ρ is the coefficient of relative risk aversion.

There is free and costless international trade in a complete set of state-contingent Arrow–Debreu securities. (Imagine again that the securities' payoffs are made in a costlessly tradable international monetary unit of account.) We continue to assume that there are iceberg costs of trade, such that only a fraction $1 - \tau$ of a unit of good shipped abroad reaches its destination, so that under competitive markets $P_F = P_F^*/(1 - \tau)$ and $P_H = (1 - \tau)P_H^*$, per equations (2) and (3).[22]

Because, in addition, the countries are symmetric, free trade in Arrow–Debreu securities yields an allocation in which

$$\frac{1}{P_H}\frac{\partial U}{\partial C_H} = \frac{1}{P_H^*}\frac{\partial U^*}{\partial C_H^*}$$

and

$$\frac{1}{P_F}\frac{\partial U}{\partial C_F} = \frac{1}{P_F^*}\frac{\partial U^*}{\partial C_F^*},$$

for every state of nature, or

$$C_H^{-1/\theta} C^{1/\theta-\rho} = (1 - \tau)C_H^{*-1/\theta} C^{*1/\theta-\rho} \tag{12}$$

and

$$(1 - \tau)C_F^{-1/\theta} C^{1/\theta-\rho} = C_F^{*-1/\theta} C^{*1/\theta-\rho}. \tag{13}$$

Together these conditions imply the ex post consumption efficiency condition

$$\left(\frac{P_F}{P_H} \right)^{\theta} = \frac{C_H}{C_F} = (1 - \tau)^{-2\theta}\frac{C_H^*}{C_F^*} = (1 - \tau)^{-2\theta}\left(\frac{P_F^*}{P_H^*} \right)^{\theta}.$$

The model is closed by the output-market clearing conditions:

22. Just as in our discussion of the trade-bias puzzle, one could obtain similar results on home bias in equity holdings if trade costs were zero but there existed a home bias in preferences along the lines of equation (7).

$$C_H^* = (1 - \tau)(Y_H - C_H),$$
$$C_F = (1 - \tau)(Y_F - C_F^*).$$

Four of the preceding five equations are independent and yield solutions for the consumption levels C_H, C_F, C_H^*, and C_F^*.

4.2 INTERPRETING THE MODEL

It may puzzle some readers that we focus on the Arrow–Debreu allocation when in fact we are interested in relating our analysis to observed trade in the narrower class of equity-type assets that one observes in the real world. One rationale, perhaps, is that we do not want our theoretical home bias results to be driven by ad hoc assumptions about the kinds of securities that can be traded, especially since many assets like debt and direct foreign investment have complex optionlike qualities that may be difficult to summarize in a simple model. A second, more pragmatic, rationale is that the equilibrium for the complete-markets case is relatively simple to compute. A final rationale, as we shall see, is that for realistic parameters, trade in equities alone can come quite close to attaining the complete-markets consumption allocation, so that the home bias evident under complete markets is a good guide to the home bias in an equities-only model.

4.3 EVALUATING THE HOME BIAS

It is helpful to begin by analyzing the special case $\rho = 1/\theta$, in which the Arrow–Debreu conditions (12) and (13) simplify enormously. One can also show that the Arrow–Debreu allocation is then identical to the one in which people can trade only straight equity shares. Given our assumption of symmetry, the equilibrium portfolio shares are

$$\chi_H = \frac{1}{1 + (1 - \tau)^{\theta-1}} Y_H,$$

$$\chi_H^* = \frac{(1 - \tau)^{\theta-1}}{1 + (1 - \tau)^{\theta-1}} Y_H,$$

$$\chi_F = \frac{(1 - \tau)^{\theta-1}}{1 + (1 - \tau)^{\theta-1}} Y_F^*,$$

$$\chi_F^* = \frac{1}{1 + (1 - \tau)^{\theta-1}} Y_F^*,$$

where χ_H (χ_F) denotes the home agent's share of total equity in the home (foreign) industry, and χ_H^* (χ_F^*) denotes the foreigner's optimal equity shares. Note that if we were to translate these equity positions into *consumption* shares, we would find c_F and c_H^* lower than χ_F and χ_H^* by a factor of $1 - \tau$, reflecting the trade costs. Of course, in the absence of trade costs all the portfolio (and consumption) shares would equal 0.5, reflecting full diversification (under symmetry).

For $\theta = 6$ and trade costs of $\tau = 0.25$ (again, a seemingly reasonable number when applied to all of output, especially compared to the usual assumption that fully half of output is nontraded), one obtains $\chi_H = 0.81$, $\chi_H^* = 0.19$. Since share prices will be equal due to symmetry, this implies a home equity share of 81%. If $\theta = 10$, then the home portfolio share of home equities is 72% even with trade costs of just 10%. (As in the case of home bias in trade, there is significant nonlinearity: the elasticity of foreign shareholdings with respect to trade costs is very high when τ is near 1, but falls as trade costs fall.) The preceding calculations constrain the value of ρ to equal $1/\theta$, but, as we shall now demonstrate numerically, the results turn out to be remarkably insensitive to this assumption, given realistic levels of output uncertainty.

If we relax our restriction $\rho = 1/\theta$, the exact conditions needed to implement the Arrow–Debreu allocation through equity trade alone are broken. Trade costs create an international wedge between marginal rates of substitution such that standard stock-market spanning theorems no longer apply.[23] Nevertheless, one can still gain a good deal of insight into home bias by computing the *state-contingent* consumptions of the two goods in the Arrow–Debreu efficient allocation.

We can reduce the dimensionality of our numerical simulations by noting that, in equilibrium, the ratios of consumption to output, $c_H \equiv C_H/Y_H$, $c_F \equiv C_F/Y_F$, $c_H^* \equiv C_H^*/Y_H$, and $c_F^* \equiv C_F^*/Y_F$, depend only on the output ratio $y_H \equiv Y_H/Y_F$. Table 4 illustrates how the consumption ratios c_H and c_F differ both across states of nature and across a number of settings of the parameters τ, θ, and ρ.[24] (The values of c_H^* and c_F^* are apparent from the assumed symmetry of the model.) Notice that the home country's output shares decline across states of nature as its relative endowment rises. That pattern compensates the foreign country for the greater share of transport costs it must pay in states of nature such that home output is relatively high, and it is naturally more pronounced the higher the risk aversion parameter ρ.

For the cases in which $\theta = 6$ and τ is 10% or 20%, the table documents

23. See Obstfeld and Rogoff (1996, Section 5.3).
24. In a "baseline" model with trade frictions in which individuals nonetheless consume the proceeds from fully diversified portfolios, we would have $c_H = \frac{1}{2}$ and $c_F = (1 - \tau)/2$.

Table 4 PORTFOLIO POSITIONS IN HOME AND FOREIGN GOODS FOR
STATE OF NATURE $y_H \equiv Y_H/Y_F$

Parameter settings			Portfolio shares c_H, c_F				
τ	θ	ρ	$y_H = 0.8$	0.9	1.0	1.1	1.2
0.1	2	2	0.53, 0.43	0.53, 0.43	0.53, 0.43	0.53, 0.43	0.52, 0.42
0.1	3	2	0.56, 0.41	0.55, 0.40	0.55, 0.40	0.55, 0.40	0.55, 0.40
0.1	5	2	0.61, 0.37	0.61, 0.36	0.60, 0.36	0.60, 0.35	0.60, 0.35
0.1	6	$1/\theta$	0.63, 0.33	0.63, 0.33	0.63, 0.33	0.63, 0.33	0.63, 0.33
0.1	6	2	0.64, 0.35	0.63, 0.34	0.63, 0.33	0.62, 0.33	0.62, 0.32
0.1	6	5	0.64, 0.35	0.64, 0.34	0.63, 0.33	0.62, 0.33	0.62, 0.32
0.2	2	2	0.56, 0.36	0.56, 0.36	0.56, 0.36	0.55, 0.35	0.55, 0.35
0.2	6	$1/\theta$	0.75, 0.20	0.75, 0.20	0.75, 0.20	0.75, 0.20	0.75, 0.20
0.2	6	2	0.78, 0.22	0.76, 0.21	0.75, 0.20	0.74, 0.19	0.73, 0.18
0.2	6	5	0.78, 0.22	0.77, 0.21	0.75, 0.20	0.74, 0.18	0.73, 0.18
0.3	6	2	0.89, 0.13	0.87, 0.11	0.86, 0.10	0.84, 0.09	0.83, 0.08
0.3	8	2	0.95, 0.08	0.94, 0.07	0.92, 0.05	0.91, 0.04	0.89, 0.04

how insensitive the portfolio shares are even to large changes in ρ. Because the results turn out to be fairly insensitive to ρ, we find that our earlier calculations are indeed little affected by relaxing the assumption $\rho\theta = 1$. The low sensitivity to ρ over the range of relative output outcomes in Table 4 is consistent with the conjecture by Cole and Obstfeld (1991) that, for moderate uncertainty, the gains from global risk sharing may be so low as to be mostly offset by costs of trade. Here, the equilibrium with a rich variety of assets is not so different from the one in which individuals can hold only equity. Another conclusion we can draw from these numbers is that trade costs have to be quite large before there is a substantial discrepancy between the Arrow–Debreu consumption allocation and the one that trade in equities alone would produce. Even for trade costs of 30%, an equity allocation that gave each country the same consumption share in *every* state of nature as it would have in the Arrow–Debreu equilibrium only when the realization was $y_H = 1$ would not entail a large departure from efficiency. As a result, even when $\rho\theta \neq 1$, the home bias evident in the complete-markets example is quite close to what a model of pure equity trade would imply.[25]

25. Backus, Kehoe, and Kydland (1992) report some relevant experiments with their calibrated two-country, complete-markets version of the Brock–Mirman stochastic growth model. It is true that they do not focus on the equity home-bias puzzle and that they allow for only a single consumption good, effectively making the elasticity of substitution between national outputs infinite. However, the fact that they find that moderate transportation costs produce an allocation close to that with full autarky is quite in accord with our results here.

4.4 CAVEATS

We do not believe that trade costs in goods markets are necessarily the whole story in explaining observed portfolio biases, and we certainly expect that the kinds of information asymmetries and legal restrictions emphasized in earlier work also play a role. These frictions can be viewed as trade costs in a broader sense, as we have noted, and they can affect portfolios through the trade-cost channel that we have emphasized in this paper. Nevertheless, it is remarkable that our simple model based on the trade-cost channel alone matches up so well to the data. As we have noted, our explanation not only has the merit of (extreme) simplicity, but is also more convincing because the same basic approach seems to help explain such a diverse range of puzzles. Finally, we note that our results are consistent with recent empirical work by Portes and Rey (1999). They find that international trade in both equities and goods is surprisingly well explained by an enhanced gravity model in which informational distance proxies supplement the standard set of geographical explanatory variables.[26] These results are certainly in accord with our model's prediction that equity biases in large measure reflect goods-market biases.[27]

A caveat to our findings is that transaction costs, and the resulting home bias, would be reduced somewhat in a fully dynamic model. Investors could then reinvest dividends abroad rather than repatriating them immediately. As is true for a tax-deferred asset, they could earn dividends on wealth that would otherwise be burned up as shipping costs. The question deserves further research. Dumas and Uppal (2000) develop a dynamic two-country growth model with shipping costs, but their focus is on welfare rather than on the home-bias puzzle. (They also assume $\theta = \infty$ throughout by positing a single consumption good.) Our guess is that trade costs will remain an important determinant of home bias even in a realistic dynamic setting.

We have used a complete-markets model to illustrate how trade costs can generate a home equity bias. By taking that modeling approach, however, we certainly do not intend to endorse an empirical view that

---

26. Portes and Rey (1999) report that their information variables are quite significant in explaining goods-market trade, even after controlling for geographical distance.
27. One consideration that dovetails nicely with our explanation is illustrated in the model of Martin and Rey (1999), which provides the closest antecedent to our approach. In Martin and Rey's (endogenously) incomplete-markets setup, the main driving force behind home bias is that owners of home firms retain a disproportionate share of their equity in order to extract a higher monopoly price for remaining shares from other agents. Martin and Rey focus on transaction costs in asset rather than in goods markets, in the tradition of Aiyagari and Gertler (1991). They posit an asymmetry between transaction costs for home and foreign agents, and this cost also affects share values. It does not interact with θ, however, so the effects are much smaller than here.

real-world asset markets are complete or nearly complete, either domestically or internationally. The complete-markets assumption is not essential, and our arguments would go through in a fully articulated incomplete-markets model, for example, one in which households have unequal access to equity markets, so that only some hold equity (Mankiw and Zeldes, 1991). The home-equity-bias puzzle has a strong empirical basis that is independent of any narrow theoretical framework. The consumption correlations puzzle, which we turn to next, encompasses a broader notion of market completeness, but its exact formulation is also more model-specific.

5. The International Consumption Correlations Puzzle (Puzzle 4)

If one believes that both domestic and international capital markets are well approximated by an Arrow–Debreu complete-markets framework, then it is a puzzle that international consumption growth correlations are not much higher than they appear to be. In an Arrow–Debreu world, country-specific output risks should be significantly pooled, and therefore domestic per capita consumption growth should not depend too heavily on country-specific income shocks. Of course, in some sense, the consumption correlations puzzle is almost a corollary of the Feldstein–Horioka and home-equity-bias puzzles. Given that the most transparent market means of consumption smoothing—debt and equity trade—are far less operative across borders than within them, it should not come as any great surprise that international consumption correlations are low. However, there are many reasons for thinking about consumption correlations independently. One is that we have only very imperfect measures of international trade in equity and debt, and another is that there may be other market channels, such as direct investment, for pooling risk.

The international consumption correlations puzzle has spawned a variety of subpuzzles. Backus, Kehoe, and Kydland (1992) highlight the fact that international output growth rates are actually more highly correlated than consumption growth rates. Backus and Smith (1993) note that in a world with traded and nontraded goods, efficient risk sharing calls for giving higher rates of consumption growth to countries that experience relative drops in the real price of consumption. (Very loosely speaking, the United States and Canada should write contracts that imply big transfers to Canada in states of nature where the Canadian dollar is very weak so that Canadians can exploit bargain Canadian prices, and vice versa when the Canadian dollar is high.)

As we shall see, most consumption correlations puzzles tend to be

quite model-specific (depending on factors like the completeness of markets and the exact form of the utility function), so they are not quite as obviously puzzles about the real world in the same way that, say, the equity-home-bias puzzle is. One does not have to believe that the world is Arrow–Debreu to think it a puzzle that agents do not take more advantage of international diversification opportunities. Nevertheless, consumption correlation puzzles play a very important role in assessing alternative general equilibrium models, and, at a more fundamental level, we can ask why consumption risk pooling tends to be higher across regions within a country's boundaries than across national boundaries.

5.1 THE PUZZLE OF LOW INTERNATIONAL CONSUMPTION CORRELATIONS

Consider a single-good world with time-separable preferences in which all agents have identical period utility functions of the form $u(C) = C^{1-\rho}/(1-\rho)$. Then, if there are no trade costs, trade in a complete set of Arrow–Debreu securities would imply that home and foreign consumption growth rates are equalized:

$$\frac{C_{t+1}}{C_t} = \frac{C_{t+1}^*}{C_t^*}, \tag{14}$$

regardless of relative shocks to home and foreign outputs. (See Obstfeld and Rogoff, 1996, Chapter 5.) This is hardly what one observes in practice, as Table 5, which gives consumption growth-rate correlations based on Penn World Table data from the Group of Seven industrial countries, illustrates. The strong prediction of equation (14) is relaxed somewhat in models where utility depends nonseparably on both consumption and leisure. However, in this case, the benchmark frictionless world economy model of Backus, Kehoe, and Kydland (1992) still predicts a cross-country consumption correlation of almost 0.9, far above the correlations we see in the table.

Since, as we have already noted, the low-consumption-correlation puzzle is virtually a corollary of the previous two puzzles we have studied, the reader will hardly be surprised when we note that introducing trade costs works just as well in explaining it. Indeed, our model of the equity-home-bias puzzle can easily generate correlations of the sort seen in Table 5.[28]

---

28. Lewis (1999) points out that when a significant share of output is absolutely non-tradable, international consumption correlations will be sharply reduced. However,

Table 5   CORRELATIONS IN PER CAPITA PRIVATE CONSUMPTION
GROWTH, 1973–1992

	France	*Germany*	*Italy*	*Japan*	*U.K.*	*U.S.*
Canada	0.25	0.31	0.44	0.05	0.40	0.64
France		0.52	0.27	0.68	0.43	0.51
Germany			0.27	0.40	0.33	0.51
Italy				0.21	0.30	0.13
Japan					0.59	0.50
U.K.						0.65

Source: Penn World Table. Correlations of log differences in per capita real consumption. Simple average of correlation coefficients is 0.40.

5.2 THE BACKUS–SMITH PUZZLE

Backus and Smith (1993) derive a generalization of equation (14) that holds when trade is costly and, as a consequence, national price levels for the consumption baskets entering $u(C)$ generally differ. Let P denote the home price level and P^* the foreign price level, with both price levels measured in the same numeraire currency. As in the last section, currency and securities can be traded without transport costs even though goods are costly to trade. Then complete markets in state contingent assets ensure that growth rates in the marginal utility of *currency*—the medium in which state-contingent insurance payments are made—are equalized across countries. If the utility-of-consumption function exhibits constant relative risk aversion and is independent of leisure, as in equation (11), that equality implies

$$\frac{C_{t+1}^{-\rho}/P_{t+1}}{C_t^{-\rho}/P_t} = \frac{C_{t+1}^{*-\rho}/P_{t+1}^*}{C_t^{*-\rho}/P_t^*}. \tag{15}$$

This generalizes equation (14) in that $P = P^*$ absent international trade frictions.

Given the high volatility of real exchange rates under floating together with the low volatility of consumption, it is perhaps not surprising that Backus and Smith's empirical work forcefully rejects the optimal risk-sharing condition (15). In fact, the empirical rejection of condition (15) is

---

Stockman and Tesar (1995) observe that, insofar as the data can be trusted, international consumption correlations for apparently tradable goods are not appreciably higher than those for goods generally classified as nontradable. Their finding supports the view that the dichotomous distinction between tradables and nontradables is overdrawn, and simultaneously suggests that there are substantial impediments to international risk sharing in traded goods.

even more devastating, since even very high values of ρ cannot reconcile that condition with the data. One possible explanation is that their assumption that preferences are separable in consumption and leisure is too strong, so that one needs to look instead at a generalized version of (15). In our view, however, incompleteness of asset markets is the major reason why condition (15) fails so miserably in practice. Indeed, given the volatility of exchange rates, the size of transfers required for (15) to hold would require a level of risk sharing even greater than we observe in domestic markets.

The alert reader will note that a version of the Backus–Smith condition will hold in a dynamic extension of our earlier model of the home-equity-bias puzzle. That model implicitly assumed flexible nominal prices, and would not produce nearly the level of real-exchange-rate volatility one sees in the data. We do not take this as damning, since for us the complete-markets assumption was only a useful device for calibration, and not a conviction. Trade costs would play essentially the same role in a world with, say, trade in debt and equities but not a complete set of Arrow–Debreu securities. Indeed, in the context of this paper, the really interesting issue is not why international consumption correlations are difficult to replicate in a complete-markets model, but the extent to which consumption risk sharing is less prevalent across distinct countries than within countries.

5.3 INCOMPLETENESS OF DOMESTIC VS. INTERNATIONAL MARKETS

Certainly, empirical studies based on domestic micro data reject resoundingly the proposition that markets are complete. For example, Attanasio and Davis (1996) find that consumption risk sharing is strikingly incomplete within the United States, and for reasons that apparently are unrelated to asymmetric information. The question the present paper raises is whether risk sharing is even more impaired internationally than domestically due to costs of specifically international trade. Our discussion of home equity bias, which does not rely fundamentally on a complete-markets assumption, suggests that this should be the case, since regional equity bias seems to be far less than the strong national home bias that we see in international data. Backus and Smith's theoretical proposition points in the same direction.

A growing body of empirical evidence supports the prediction that financial markets are less effective in promoting risk sharing among countries than among regions within a country. A full review of this literature would take us too far afield, but we can mention briefly a few relevant papers. Atkeson and Bayoumi (1993), in one of the first empirical studies in this area, find that regional financial transfers within the United States

are much larger than those among the major industrial countries. A comparison of the variance-decomposition results of Asdrubali, Sørensen, and Yosha (1996) on the United States with those of Sørensen and Yosha (1998) on the OECD suggests that financial markets play a much bigger role in consumption smoothing among U.S. states than is the case among industrial countries. Crucini (1999), using an alternative method, concludes that Canadian provinces pool risks more effectively than U.S. regions, and that either country shows more internal risk pooling than does the sample of industrial countries. Bayoumi and Klein (1997) find that Canadian provinces display more financial integration with each other than with the outside world.[29]

So there indeed is a puzzle as to why *intra*national consumption risk sharing is more efficient than *inter*national risk sharing, but it can be resolved in the same manner as we have resolved the home-bias and Feldstein–Horioka puzzles.

5.4 THE RELATIVE CORRELATIONS OF INTERNATIONAL CONSUMPTION AND OUTPUT GROWTH RATES

Backus, Kehoe and Kydland (1992) emphasize the puzzle that empirical consumption correlations are actually lower than output correlations. That pattern holds in the Penn World Table data analyzed here: the average international correlation in per capita real GDP growth rates is 0.53 over 1973–1992, while the corresponding average consumption correlation is only 0.40.

Our model, on its own, does not offer a new rationalization of their finding. However, we do not consider this to be a fundamental problem, since the existence of international risk sharing need not generate higher correlation among consumptions than outputs across countries. The reason is that only the output remaining after investment and government consumption can be shared by private consumers. Thus, a more appropriate comparison to assess the degree of global risk sharing is that between international consumption correlations and correlations in growth rates of output *net* of investment and government consumption $(Y - I - G)$. Table 6 reports these correlations for the same sample period and data set used to construct Table 5. The average international correlation in the growth of $Y - I - G$ is 0.17, far below the average correlation 0.40 of international consumption growth rates. For six of the 21 country pairs that ranking is reversed, but in most of these cases the discrepancy is not significant.

So in fact, the puzzle concerning the relative variability of output and

---

29. Obstfeld (1995) adds a number of caveats to some of this literature.

Table 6   CORRELATIONS IN PER CAPITA $Y - I - G$ GROWTH, 1973–1992

	France	Germany	Italy	Japan	U.K.	U.S.
Canada	0.17	0.19	0.36	−0.18	0.50	0.66
France		0.13	0.34	0.20	0.02	0.11
Germany			0.19	−0.19	0.13	0.18
Italy				−0.31	0.33	0.46
Japan					−0.25	−0.22
U.K.						0.73

Source: Penn World Table. Correlations of log differences in per capita real GDP net of investment and government consumption. Simple average of correlation coefficients is 0.17.

consumption is not necessarily incompatible with a high level of international asset market integration. Indeed, using a dynamic new open-economy macroeconomic model, Chari, Kehoe, and McGrattan (1998) are able to produce realistic cross-country correlations of output as well as of consumption.[30] The main additional assumptions that lie behind their results include sticky nominal prices and, implicitly, transport costs high enough to result in segmented national output markets. Both transport costs and nominal rigidities are central to the resolution of the fifth and sixth puzzles, to which we now turn.

6. The Purchasing-Power-Parity Puzzle (Puzzle 5) and the Exchange-Rate Disconnect Puzzle (Puzzle 6)

Our last two puzzles differ from the preceding ones in being fundamentally about the real effects of a nominal variable—the exchange rate, which is the relative price of currencies. Here, also in contrast to the preceding four puzzles, the difficulty seems to lie primarily in explaining short- to medium-term phenomena rather than phenomena that persist over very long periods. (The Feldstein–Horioka puzzle, for example, is typically framed using decade-average data). Finally, the last two puzzles can be viewed as pricing puzzles, because they refer to price behavior, including the dynamic covariation between prices and other macroeconomic variables.

Any realistic attempt to address these pricing puzzles formally would require a much more elaborate framework than the one we have used thus far, incorporating, among other things, elements of monopoly and sticky

30. In the Chari–Kehoe–McGrattan sticky-price model, highly correlated national monetary shocks can make national outputs covary more closely than national consumptions. Highly correlated monetary shocks, however, also tend to reduce real-exchange-rate variability counterfactually in the model. We suspect that an extended version of the model could handle the latter problem.

nominal prices for goods and/or labor. In fact, there is already a great deal of exciting research along these lines now taking place [see, for example, the recent survey by Lane (2001) on the new open-economy macroeconomics]. Unfortunately, we do not have nearly enough space remaining here to present a fully articulated model. Nevertheless, we will try to make clear why trade costs are as essential to resolving the pricing puzzles as they are to resolving puzzles 1 through 4, which are quantity puzzles.

The first pricing puzzle we take up is the purchasing-power-parity (PPP) puzzle (Rogoff, 1996), which highlights just how weak the connection is between exchange rates and national price levels. It is based on the observation that in hundreds of studies, using widely varying techniques and data sets, researchers have repeatedly found very long half-lives—on the order of 3 to 4 years—for shocks to real (CPI) exchange rates. As we shall explain, half-lives of this magnitude are hard to understand if financial-market disturbances with only transitory real effects are very important in explaining short-run volatility.

Our term for the second pricing puzzle is the *exchange-rate disconnect puzzle*, a name that alludes broadly to the exceedingly weak relationship (except, perhaps, in the longer run) between the exchange rate and virtually *any* macroeconomic aggregates. It manifests itself in a variety of ways. For example, Meese and Rogoff (1983) showed that standard macroeconomic exchange-rate models, even with the aid of ex post data on the fundamentals, forecast exchange rates at short to medium horizons no better than a naive random walk. Baxter and Stockman (1989) argued that transitions to floating-exchange-rate regimes lead to sharp increases in nominal- and real-exchange-rate variability with no corresponding changes in the distributions of fundamental macroeconomic variables.[31] (The PPP puzzle is really just an example, albeit a very important one, of the broader exchange-rate disconnect puzzle.)

A critical difference between the (relatively short-term) pricing puzzles and the (longer-term) quantity puzzles is that we can no longer appeal to high elasticities of substitution to lever up the effects of modest-sized trade costs. (At the very least, the connection is no longer as simple and direct.) If there are only modest obstacles to short-term price arbitrage across borders, there can be only modest short-term price differentials. In fact, at the consumer level, arbitrage costs *are* likely to be rather large, and, after all, most goods embody very large nontraded content once they reach consumers at the retail level. But one cannot make this argument for wholesale importers who trade in bulk, so here

---

31. Flood and Rose (1995) extend Baxter and Stockman's results and arrive at similar conclusions.

we need a more nuanced discussion. As we shall see, importer-level prices do appear to exhibit somewhat less anomalous behavior than do consumer-level prices.

6.1 THE PPP PUZZLE

Let Q be the real exchange rate between two countries, and consider the regression equation

$$\log Q_t = \alpha + \eta t + \gamma \log Q_{t-1} + \epsilon_t,$$

where ϵ_t is a random disturbance. The real exchange rate, Q, is defined as $\mathscr{E}P^*/P$ using overall CPI data for price levels, where the nominal exchange rate $\mathscr{E}$ is the price of foreign currency in terms of home currency. (In deference to conventional usage, we now switch notation and use P to denote the domestic price level measured in home currency and P^* the foreign price level measured in foreign currency.)

Using monthly 1973–1995 data for Canada, France, Germany, Japan, and the United States, and constructing all 10 possible real exchange rates in this sample, we find values of γ ranging from 0.99 (U.S.–Canada, implying a half-life of 69 months) to 0.97 (Germany–Japan, implying a half-life of 21 months). The mean half-life across these real exchange rates is around 39 months, or $3\frac{1}{4}$ years.[32]

Such long half-lives would not necessarily be a puzzle but for the remarkable volatility of real and nominal exchange rates, volatility that seems hard to explain without assigning a major role to monetary and financial shocks. If monetary and financial shocks are the predominant source of volatility, however, it is hard to imagine what source of nominal rigidity could be so persistent as to explain the prolongation of real-exchange-rate deviations. This is the PPP puzzle.

6.2 THE PPP PUZZLE FOR TRADABLES VERSUS NONTRADABLES

One might think that the slow mean reversion just documented applies primarily to goods with extremely high international trade costs, whereas, at least for goods that are heavily traded, mean reversion in relative international consumer prices might be more rapid. That is not the case, however, as documented most strikingly by Engel (1999).

If we are willing to set our qualms aside temporarily and adopt a conventional dichotomy of traded versus nontraded consumer goods, we can use Figure 3 to illustrate the empirical significance of the distinc-

---

32. Data on end-of-month nominal exchange rates and on consumer price indexes come from *International Financial Statistics.*

tion for real-exchange-rate dynamics. The figure is based on monthly 1962–1995 data from Engel (1999, Section I) for the United States, France, Germany, and Japan. The overall real exchange rate $Q = \mathcal{E}P^*/P$ is compared with relative price indexes for tradables and nontradables, $\mathcal{E}P_T^*/P_T$ and $\mathcal{E}P_N^*/P_N$, where we adopt Engel's disaggregation of OECD sectoral CPI data into tradable and nontradable subindexes.[33] Each panel of the figure plots the correlations of percentage changes between pairs of relative prices, where the number of months over which the data are differenced is measured on the horizontal axis.

Consistent with Engel's results, the data reveal no significant difference between short-term and long-term correlations, indicating extremely slow mean reversion in shocks to the relative prices of tradables. Interestingly, it seems to make rather little difference whether we use tradables or nontradables prices to compute real exchange rates: all the price ratios are highly correlated with each other even out to horizons of five years. Engel's results focused on the U.S. real exchange rate against various trading partners, but as one can see from the figure, the results are (almost) as striking for a pairing of Germany and Japan. Other non-U.S. pairings that we have examined look similar.

We have argued that the traded–nontraded-goods distinction is much too finely drawn—at the retail level, many "traded" goods already embody very large nontraded components, and the dividing line is arbitrary and likely endogenous. It is nevertheless surprising just how little difference there is between the measures of real exchange rates in Figure 3. These findings probably cannot be ascribed merely to price aggregation problems, since many researchers report similar sluggish responses even for relatively disaggregated data on consumer goods that are commonly perceived as highly tradable. (See, for example, Isard, 1977; Giovannini, 1988; and Engel and Rogers, 1996.) The results certainly seem to suggest that even over the medium term, the consumer prices of supposedly tradable goods are nearly as insulated from the forces of international arbitrage as are the consumer prices of nontradables.

6.3 ADJUSTMENT IS FASTER AT THE PRODUCER LEVEL

It is important to emphasize that there seems to be considerably more adjustment of prices to exchange-rate changes at the importer level than at the consumer level. In their excellent survey of the empirical literature on exchange rates and international prices, Goldberg and Knetter (1997) conclude that the passthrough of exchange rates to relative international prices is about 50% after one year, much faster than what we have just

---

33. See Appendix A of Engel (1999). Figure 3 looks much the same if attention is restricted to data from the floating-exchange-rate period, 1973–1995.

Figure 3

U.S. and Germany

U.S. and France

| —▲— Q and EPt*/Pt | —●— Q and EPn*/Pn | —■— EPt*/Pt and EPn*/Pn |
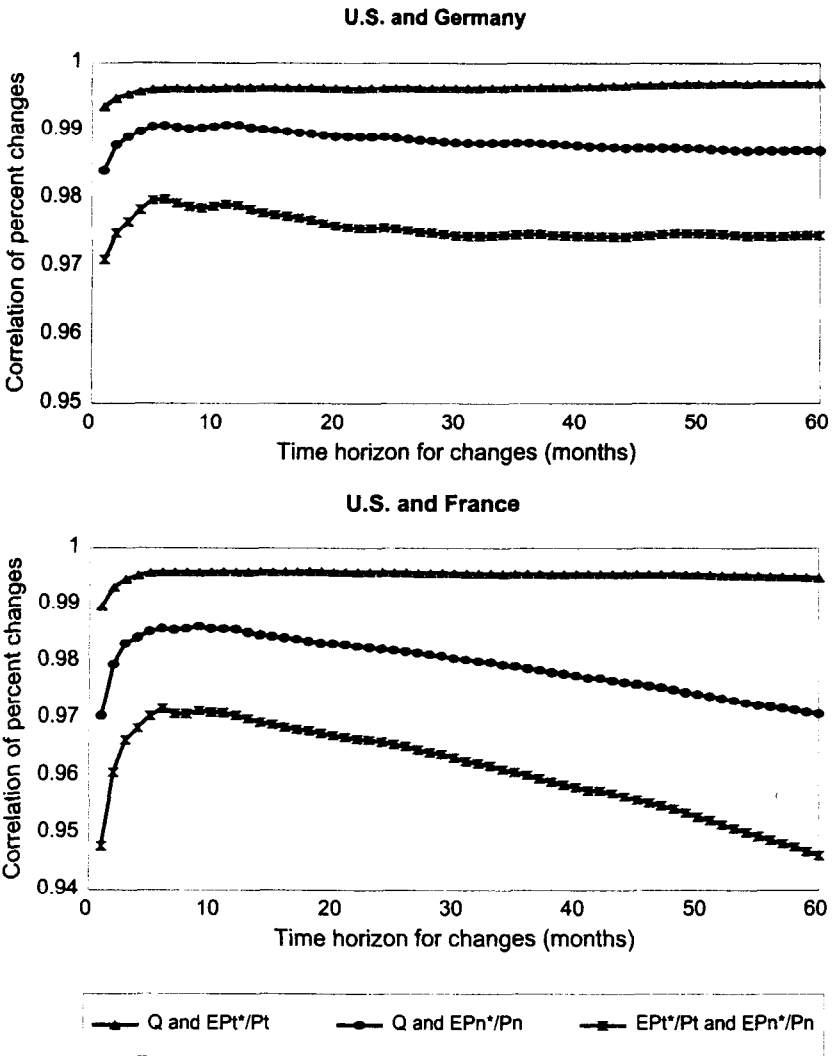
seen in the consumer-price data. Thus, relatively large elasticities in international trade between exporters and importers can be consistent with exceedingly sluggish adjustment in the relative *consumer* prices of tradables.

Obstfeld and Rogoff (2000) observe that if this were not the case—if prices paid by importers moved as sluggishly as prices paid by con-

Figure 3 continued

U.S. and Japan

Germany and Japan


```
—▲— Q and EPt*/Pt       —●— Q and EPn*/Pn       —■— EPt*/Pt and EPn*/Pn
```
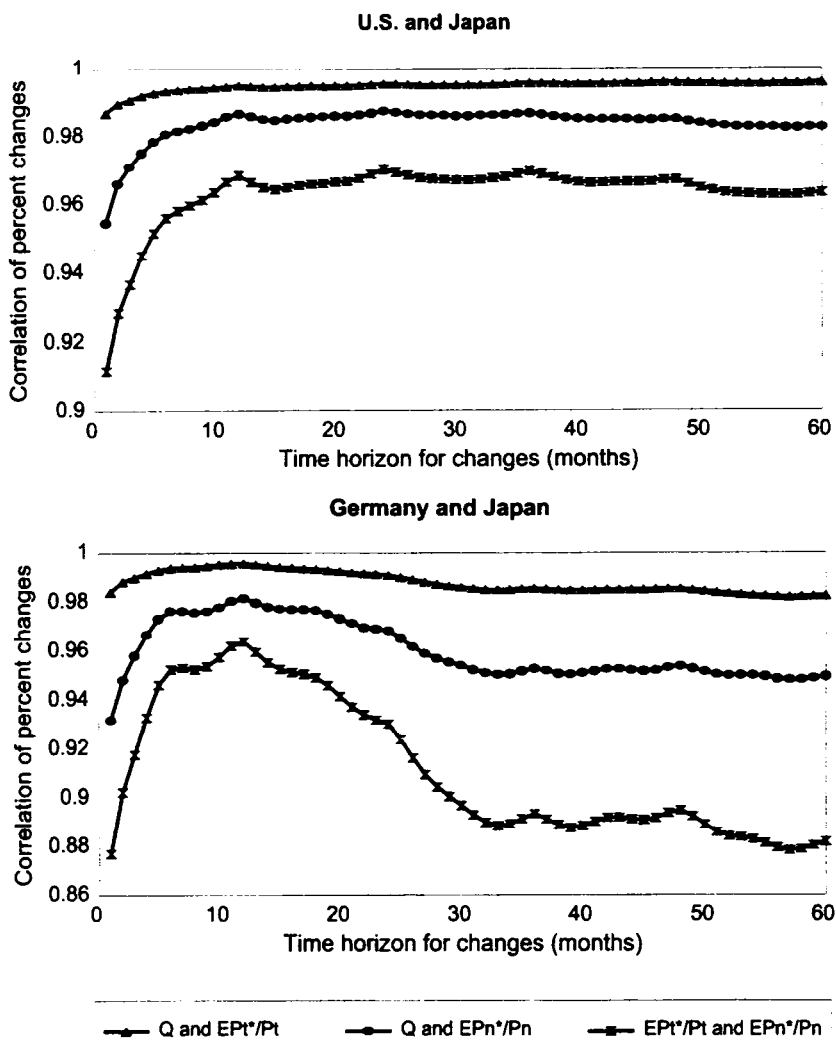
sumers—a country's terms of trade would actually improve, rather than worsen, after a depreciation of the exchange rate. For example, if the dollar depreciates against the pound and all prices are sticky, the dollar price paid by Americans for British goods remains fixed whereas the price paid by British citizens for American goods *rises* when translated into dollars. They find that this does not seem to be the case empirically, and instead

find significant support for the conventional view—that exchange-rate depreciation worsens the term of trade of the depreciating country.

6.4 TRADE COSTS AND PRICING TO MARKET

Whereas the home bias in trade could, in principle, be explained simply by a home bias in preferences, the failure of markets to arbitrage international price differentials for seemingly identical goods cannot. The most popular explanation of persistent international price differentials argues that most goods are supplied monopolistically, and that (by assumption) monopoly producers have very broad scope to *price to market* by charging different prices in home and foreign markets (see, for example, Dornbusch, 1987; Krugman, 1987; Betts and Devereux, 1996; Bergin and Feenstra, 2000; or Devereux and Engel, 2000). Goldberg and Knetter (1997) survey a large body of supportive empirical evidence.

This explanation of international price differences for very similar or identical goods is appealing, but incomplete. What is to prevent consumers from arbitraging between home and foreign prices? Any explanation—and the pricing-to-market literature offers many; see Dornbusch (1987)—has to be consistent with the tenuous connection between exchange rates and the relative prices for virtually any type of consumer good. Rationales for pricing to market that might make sense for big-ticket items such as cars (the steering wheels on American and Japanese cars are on opposite sides, dealers can refuse warranty service for vehicles purchased abroad, etc.) are not very appealing when applied to, say, basic clothing items.

In our view, trade costs simply must play a central role in any explanation of international price differentials. However, to make sense of the price data, we must refine our earlier discussion of trade costs to distinguish between bulk wholesale and individual consumer trade costs. We must also think carefully about the ability of producers to control international distribution chains at the wholesale level. Otherwise—if the only wedge between home and foreign markets were moderate trade costs—one would only observe moderate price differentials.

6.5 WHOLESALE BULK VS. RETAIL INDIVIDUAL TRANSSHIPPING COSTS

At the consumer level, it is likely that for many goods, trading costs are in fact quite large, and far, far larger than trading costs faced by bulk wholesale shippers. (Individual consumers cannot profitably arbitrage even large differences in Coca-Cola prices across countries, but bulk wholesalers can.) The real question is what prevents international price arbitrage at the wholesale level. One answer is that in many cases, a firm

can establish legal rights to control distribution of its product in different countries. Exclusive national marketing licenses are extremely common. For example, to protect its ability to price-discriminate across home and foreign markets, the Coca-Cola company sued a couple of small American wholesalers who, during the late 1990s, were trying to arbitrage the difference between Coca-Cola's $11.50-per-case wholesale price in Japan (as of January 2000) and its wholesale $5.50-per-case price in the United States—a differential far in excess of bulk shipping costs.[34] True, for small firms, the costs of establishing sole country distribution rights, and even more the legal costs of enforcing such rights, are likely to be prohibitive. Such firms also are likely to deal only with a very small number of bulk wholesalers, however, so it is still quite possible that they can price-discriminate, either by exploiting long-term relationships with their downstream wholesalers or even by taking over more portions of their wholesale distribution network.

6.6 PRICING TO MARKET AND THE PPP PUZZLE

To explain the data adequately, one must flesh out many details that we are omitting here. Very simple models of the kind we used in the first four sections are simply not adequate. For example, it is well known that with constant elasticities of demand, a monopolist may charge different prices in different countries, but exchange-rate changes will not cause fluctuations in relative prices charged [see Dornbusch (1987) and Marston (1990) for partial equilibrium models, and Betts and Devereux (1996), Obstfeld and Rogoff (1996, Chapter 10), and Hau (2000a) for general equilibrium models]. The nature of price rigidities is also quite important; Devereux and Engel (2000) emphasize that to make sense of the consumer-price data, one must think of final consumer goods prices as being sticky largely in domestic-currency terms for both domestically produced goods and importables.

Once one allows for pricing to market, however, it does become possible to develop models that can generate large price differentials exhibiting considerable persistence. Leading examples of such models are in Bergin and Feenstra (2001) and Chari, Kehoe, and McGrattan (1998), both of which develop new open-economy macroeconomic models with rich price dynamics. These authors do not explicitly base their models on trade costs—they do not try to rationalize the existence of pricing to market, but just assume it—so trade costs are only implicit. An example of a model with explicit trade costs is given by Dumas (1992), who observes that moderate trade costs can generate real-exchange-rate per-

---

34. See Constance L. Hays, "In Japan, What Price Coca-Cola?" *New York Times*, January 26, 2000, p. C1.

sistence even in a competitive world of fully flexible prices. However, the Dumas model cannot simultaneously generate anywhere near the volatility and persistence needed to match the data. Monopoly and nominal rigidities appear to be essential elements of any resolution of the PPP puzzle.[35]

Finally, one should note that in the presence of trade costs, econometric estimates of the half-life of real-exchange-rate movements may be exaggerated. Price differentials dissipate very slowly within transaction-cost bands, but more quickly outside them, and proper econometric estimation should take these nonlinearities into account (see Michael, Nobay, and Peel, 1997; Obstfeld and Taylor, 1997; and Taylor, 2001).[36]

6.7 THE EXCHANGE-RATE DISCONNECT PUZZLE

The same reasoning we have applied to thinking about the PPP puzzle can be applied to a much broader range of puzzles, all relating to the remarkably weak short-term feedback links between the exchange rate and the rest of the economy. We term this broader class of puzzles the *exchange-rate disconnect puzzle*. In a sense, the PPP puzzle is simply a very important special example of this broader class of phenomena. Of course, one may well ask why the exchange-rate disconnect puzzle should be any different from the *stock-price disconnect puzzle*, that is, the fact that stock markets seem to gyrate wildly without having any sizable contemporaneous effects on the real economy. We ourselves (Obstfeld and Rogoff, 1996, Chapter 9) have argued that to understand exchange-rate volatility, one ultimately needs a broader model that explains the high volatility we seem to observe in all asset markets. While we still maintain that view, it is also true that the links between the exchange rate and the real economy are much more direct than for stock prices. In most economies, the exchange rate is the single most important relative price, one that potentially feeds back immediately into a large range of transactions. Because the potential links are so direct, it is surprising indeed that they are not stronger.

Though much work remains to be done, it appears to us that a framework such as the one we have outlined earlier in this section (under puzzle 5) holds great potential for explaining the other disconnect puzzles as well. For example, exchange rates are remarkably volatile relative to any model we have of underlying fundamentals such as interest rates,

---

35. Working in a competitive flexible-price model with transport costs, Ravn and Mazzenga (1999) are also unable to rationalize both the real-exchange-rate volatility and the real-exchange-rate persistence in the data. Ohanian and Stockman (1997) develop an exploratory theoretical model of trade costs in a flexible-price monetary model.

36. Rogoff (1996) posits trading-cost bands as an essential element of any explanation of the PPP puzzle.

outputs, and money supplies, and no model seems to be very good at explaining exchange rates even ex post. The traditional thinking is that even though a broad range of goods is nontraded, there is always a broad range of goods that *are* traded, and these tie down the exchange rate. But a recurring theme here is that markets for most "traded" goods are not fully integrated, and segmentation due to various trade costs can be quite pervasive. In fact, the spectrum of goods subject to low trade costs may be very narrow.

In the type of model we described earlier in this section, a financial-market shock that moves the exchange rate may have little economic effect even over a fairly long horizon. With pervasive pricing to market at the retail level, consumers will be largely insulated from exchange-rate effects until these have had the time to feed through to wholesale import prices and, from there, to retailers. The magnitude of the PPP puzzle suggests how long that process might take.

Thus, interacting with the segmentation caused by trade costs, nominal price rigidities can produce a disconnect in which the exchange rate responds wildly to shocks. With the prices of most goods preset in local currency and real variables such as aggregate consumption largely insulated from exchange rates in the short run, exchange-rate adjustments have minimal short-run economic effects and therefore must be huge to clear financial markets. Only gradually will the responses of importers and exporters feed through to the retail level—and the adjustments might well be too slow to be picked up in the kinds of tests performed by Baxter and Stockman (1989). High volatility and the exchange-rate disconnect therefore both result from a combination of trade costs (costs that are especially high for consumers), monopoly, and pricing to market in local currency. A full model would incorporate those factors, while also modeling fully the dynamics of price adjustment through retail distribution networks, as well as other channels through which exchange rates might affect the real economy.[37]

We do not have space to explore the many implications that this intriguing class of models suggests. Can heightened exchange-rate volatility due to transport costs act to further segment markets internationally, with a resulting multiplier effect on volatility?[38] What are the welfare

37. Engel (1996) proposes that if all consumer prices are preset in local currency and firms fully hedge currency risks, exchange-rate changes will have no real effects and therefore exchange rates will be indeterminate. Hau (2000b) develops a new open-economy macroeconomic model in which exchange-rate volatility is decreasing in the degree of openness to international trade.

38. The theoretical work of Bacchetta and van Wincoop (1998) and Obstfeld and Rogoff (1998, 2000) and the empirical work of Obstfeld and Taylor (1997) and Rose (2000) suggest that currency volatility may itself act as a barrier to international trade.

costs of the exchange-rate disconnect? But the general approach strikes us as a very promising and realistic way to think about a host of exchange-rate volatility puzzles.

7. Conclusions

The need for research on the effects of trade costs in standard models of international finance seems compelling to us. We find that introducing plausible proportional (iceberg) trade costs into the most standard international macroeconomics models substantially resolves many of the core empirical puzzles in the field, including especially the (seemingly intractable) Feldstein–Horioka puzzle, the home-bias-in-equities puzzle, the home-bias-in-trade puzzle, and the low-consumption-correlations puzzle. We cannot claim the same degree of success in elucidating pricing puzzles as in the case of quantity puzzles, at least not with the kind of very simple models we have featured here. To tackle the PPP puzzle and the exchange-rate disconnect puzzle properly, a much richer framework featuring imperfect competition and wage-price rigidities is needed (therefore one in which, at a very fundamental level, neither domestic nor international markets are perfect). It is also necessary to build in a distinction between retail and wholesale pricing to account for the sharply different behavior of terms-of-trade indexes vs. consumer price indexes in response to exchange-rate changes (see Obstfeld and Rogoff, 2000, and Tille, 2000). We have argued, however, that introducing trade costs (implicitly or explicitly) must be an essential ingredient in resolving the international pricing puzzles as well. Richer models might consider fixed costs of trade as well as the proportional costs on which we have focused here.[39]

Although we take an eclectic perspective on the degree of completeness of international capital markets, our analysis does not rely on the assumption that their performance is intrinsically inferior to that of domestic capital markets (at least not in analyzing data for OECD countries). Our focus, instead, is on the distinctive ramifications for asset-market performance of the imperfect integration of *goods* markets. One attractive feature of our approach is that it is entirely consistent with the observation that gross flows in international capital markets are much larger than the small net flows.

An obvious potential criticism of our central theme is that transport technology has been steadily improving over the past half century, and tariffs have fallen dramatically, especially among the OECD countries. Has the home bias in trade and equities lessened, and are the consump-

39. O'Connell and Wei (1997) give an example of a theoretical model of price arbitrage involving fixed as well as variable costs.

tion-correlations and Feldstein–Horioka puzzles less acute than they were half a century ago? The short answer is that trade, capital movements, and equity flows all *have* expanded sharply since 1950, so the major quantity puzzles *are* less acute. For example, the ratio of total trade (the sum of imports and exports) to GDP has roughly doubled across the OECD between 1950 and 1995; for the United States, it has risen from 9% in 1950 to 24% in 1995.[40] (This calculation may significantly understate the true growth rate, since a large fraction of trade is in manufactures, the relative price of which has been falling over time.) And, as we have already seen, OECD savings–investment correlations have fallen significantly (from 0.89 for 1960–1974 to 0.60 for 1990–1997), while holdings of foreign equity have risen sharply (for the United States, from a 4% share in 1987 to a 10% share in 1996). At the same time, while transport technology has steadily improved, labor costs have risen sharply, so there is actually some debate about whether net transport costs have fallen. Hummels (1999b) argues that, until recently, the overall effect has been relatively small, with shipping costs falling sharply for bulk commodities but actually rising for manufactures, which account for over 70% of OECD trade. Greenspan (1989), on the other hand, emphasizes that trade is getting lighter, as many of the goods and services being traded today are highly knowledge-intensive. Overall, the data for the past half century certainly do not provide any prima facie case against our approach.

It would be interesting to look at time spans beyond just the past fifty years, so that trend declines in trade costs become more pronounced. Williamson (2000) calculates that transport costs for internationally traded goods fell by 1.5% per annum in real terms from 1850 to 1913, with the rate slowing down substantially over 1913–1950. Although prewar data are much thinner than postwar, and although there are many other factors to control for (large fluctuations in tariff rates, decolonization, wars, changes in the international monetary regime, etc.), this would nevertheless be a useful exercise. Cross-sectional empirical work is also needed.

Finally, a small apology to readers who were expected us also to address the forward-premium puzzle. We simply have not yet tackled this particular pricing puzzle, which we regard as much more of a pure finance question than a macroeconomic puzzle (and hence this paper's title). We note, however, that Dumas (1992) has produced a model of the forward premium in which trade costs do pull in the right direction, so getting a trade-cost model with the right quantitative effects may indeed be possible.

40. The only outliers are Australia and Japan, with trade ratios that remained roughly constant between 1950 and 1995 at 40% and 19%, respectively. See Baldwin and Martin (1999) or World Bank, *World Development Report, 1995.*

Appendix

Table 7 presents saving and investment rates by country for 1990–1997.

Table 7   SAVING AND INVESTMENT RATES, 1990–1997

Country	NS/Y^a	I/Y^b	$OECD^c$
Switzerland	0.29	0.23	1
Japan	0.33	0.30	1
Norway	0.27	0.23	1
Singapore	0.50	0.36	
Denmark	0.17	0.15	1
Iceland	0.16	0.17	1
United States	0.15	0.17	1
Germany	0.21	0.22	1
Austria	0.23	0.24	1
Belgium	0.22	0.18	1
Sweden	0.15	0.16	1
France	0.20	0.19	1
Netherlands	0.25	0.21	1
Finland	0.18	0.18	1
United Kingdom	0.14	0.15	1
Australia	0.17	0.22	1
Italy	0.19	0.19	1
Canada	0.16	0.18	1
Ireland	0.21	0.19	1
Countries with GNP/cap.^{d} > 18,000 (ave.)	0.22	0.21	
New Zealand	0.16	0.19	1
Israel	0.07	0.24	
Spain	0.20	0.22	1
Greece	0.15	0.17	1
Korea	0.35	0.37	
Portugal	0.22	0.22	1
Countries with GNP/cap. 5000–18,000 (ave.)	0.19	0.24	
Saudi Arabia	0.28^e	0.21	
Uruguay	0.12	0.13	
Chile	0.21	0.25	
Malaysia	0.33	0.39	
Trinidad and Tobago	0.18^f	0.16	
Mauritius	0.24	0.29	
Mexico	0.19	0.23	1
Venezuela	0.22	0.17	
Turkey	0.20	0.21	1
Panama	0.23	0.25	

Table 7 continued

Country	NS/Y[a]	I/Y[b]	OECD[c]
Thailand	0.34[f]	0.41	
Costa Rica	0.21	0.27	
Iran, I.R. of	0.26	0.27	
Colombia	0.18	0.21	
Namibia	0.15	0.21	
Tunisia	0.17	0.27	
Paraguay	0.12	0.23	
Countries with GNP/cap. 2000–5000 (ave.)	0.21	0.24	
El Salvador	0.01	0.17	
Dominican Republic	0.13	0.23	
Ecuador	0.16	0.20	
Jordan	0.01	0.32	
Guatemala	0.07[g]	0.15	
Morocco	0.18	0.22	
Philippines	0.17	0.23	
Sri Lanka	0.14	0.25	
Zimbabwe	0.14[f]	0.21	
Honduras	0.17	0.30	
Pakistan	0.16	0.19	
Zambia	0.10	0.24	
Kenya	0.13	0.20	
Burkina Faso	0.07	0.24	
Malawi	0.01	0.18	
Countries with GNP/cap. < 2000 (ave.)	0.11	0.22	
All countries (average)	0.19	0.22	

[a]*NS/Y*: gross national saving/gross domestic product, averaged over 1990–1997. For OECD countries, data on *NS* and *Y* are from the OECD database. For non-OECD countries, *NS* was constructed, from *International Financial Statistics* (IMF), as follows: *NS* = GNP − private consumption -- government consumption. Our measure of *NS* for non-OECD countries does not exactly match the theoretical definition. The main difference is that it does not take account of the balance-of-payments component "net current transfers from abroad." Most of the countries that report data to the IMF and are not in the sample were excluded for one of four reasons: (1) IFS has data only for GDP and not GNP; (2) there are no IFS data on inventory investment; (3) there is a significant statistical discrepancy either between GDP and its components (more than 3%), or between GNP and the sum of GDP and net factor income/payments from abroad (more than 2%); (4) population is under 1 million.
[b]*I/Y*: investment/GDP, average over 1990–1997. Investment is the sum of gross fixed capital formation and increase (decrease) in inventory stocks. Sources are as in note *a*.
[c]The OECD sample of countries includes those that were members in 1995.
[d]GNP per capita measured in U.S. dollars, for 1997.
[e]No data for 1996 and 1997.
[f]No data for 1997.
[g]No data for 1991.

REFERENCES

Aiyagari, R., and M. Gertler. (1991). Asset returns with transactions costs and uninsured individual risk. *Journal of Monetary Economics* 27:311–331.

Anderson, J. (1979). Theoretical foundations of the gravity equation. *American Economic Review* 69(March):106–116.

————, and D. Marcouiller. (1999). Trade, insecurity, and home bias: An empirical investigation. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7000.

————, and P. Neary. (1998). The mercantilist index of trade policy. Boston College. Mimeo.

Asdrubali, P., B. Sørensen, and O. Yosha. (1996). Channels of interstate risk sharing: United States 1963–90. *Quarterly Journal of Economics* 111(November):-1081–1110.

Atkeson, A., and T. Bayoumi. (1993). Do private capital markets insure regional risk? Evidence from the United States and Europe. *Open Economies Review* 4:303–324.

Attanasio, O., and S. Davis. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104(December):1227–1262.

Backus, D., P. Kehoe, and F. Kydland. (1992). International real business cycles. *Journal of Political Economy* 100(August):745–775.

————, and G. Smith. (1993). Consumption and real exchange rates in dynamic economies with non-traded goods. *Journal of International Economics* 35(November):297–316.

Bacchetta, P., and E. van Wincoop. (1998). Does exchange rate stability increase trade and capital flows? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6704.

Baldwin, R., and P. Martin. (1999). Two waves of globalization: Superficial similarities, fundamental differences. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6904.

Baxter, M., and U. Jermann. (1997). The international diversification puzzle is worse than you think. *American Economic Review* 87(March):170–191.

————, ————, and R. King. (1998). Nontraded goods, nontraded factors, and international non-diversification. *Journal of International Economics* 46:211–229.

————, and A. Stockman. (1989). Business cycles and the exchange rate regime: Some international evidence. *Journal of Monetary Economics* 23(May):377–400.

Bayoumi, T., and M. Klein. (1997). A provincial view of economic integration. *International Monetary Fund Staff Papers* 44(December):534–556.

Bergin, P., and R. Feenstra. (2001). Pricing-to-market, staggered contracts, and real exchange rate persistence. *Journal of International Economics*, forthcoming.

Berry, S., J. Levinsohn, and A. Pakes. (1995). Automobile prices in market equilibrium. *Econometrica* 63(July):841–890.

Betts, C., and M. Devereux. (1996). The exchange rate in a model of pricing to market. *European Economic Review* 40(April):1007–1021.

Brown, D., and R. Stern. (1989). Computable general equilibrium estimates of the gains from U.S.–Canadian trade liberalization. In *Economic Aspects of Regional Trading Arrangements,* D. Greenaway, T. Hyclak, and R. Thornton (eds.). New York: New York University Press.

Chari, V., P. Kehoe, and E. McGrattan. (1998). Monetary shocks and real

exchange rates in sticky price models of international business cycles. Federal Reserve Bank of Minneapolis. Research Department Staff Report No. 223.

Cheung, Y., M. Chinn, and E. Fujii (1999). Market structure and the persistence of sectoral real exchange rates. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7408.

Coakley, J., F. Kulasi, and R. Smith. (1998). The Feldstein–Horioka puzzle and capital mobility: A review. *International Journal of Finance and Economics* 3:169–188.

Cole, H., and M. Obstfeld. (1991). Commodity trade and international risk sharing: How much do financial markets matter? *Journal of Monetary Economics* 28(August):3–24.

Crucini, M. (1999). On international and national dimensions of risk sharing. *Review of Economics and Statistics* 81(February):73–84.

Deardorff, A. (1998). Determinants of bilateral trade: Does gravity work in a neoclassical world? In *The Regionalization of the World Economy*, J. Frankel (ed.). Chicago: University of Chicago Press.

Devereux, M., and C. Engel. (2000). Monetary policy in the open economy revisited: Price setting and exchange rate flexibility. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7665.

Dornbusch, R. (1987). Exchange rates and prices. *American Economic Review* 77(March):93–106.

———, S. Fischer, and P. Samuelson. (1977). Comparative advantage, trade, and payments in a Ricardian model with a continuum of goods. *American Economic Review* 67(December):823–839.

Dumas, B. (1992). Dynamic equilibrium and the real exchange rate in a spatially separated world. *Review of Financial Studies* 5:153–180.

———, and R. Uppal. (2000). Global diversification, growth and welfare with imperfectly integrated markets for goods. *Review of Financial Studies*, forthcoming.

Engel, C. (1996). A model of foreign exchange rate indetermination. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 5766.

———. (1999). Accounting for U.S. real exchange rate changes. *Journal of Political Economy* 107(June):507–538.

———, and J. Rogers. (1996). How wide is the border? *American Economic Review* 86(December):1112–1125.

Evans, C. (1999). *Do national borders matter?* Doctoral Dissertation, Harvard University.

Feldstein, M., and C. Horioka. (1980). Domestic savings and international capital flows. *Economic Journal* 90(June):314–329.

Flood, R., and A. Rose. (1995). Fixing the exchange rate regime: A virtual quest for fundamentals. *Journal of Monetary Economics* 36(August):3–37.

Frankel, J. (1986). International capital mobility and crowding-out in the U.S. economy: Imperfect integration of financial markets or of goods markets? In *How Open Is the U.S. Economy?* R. W. Hafer (ed.). Lexington, MA: Lexington Books.

French, K., and J. Poterba. (1991). Investor diversification and international equity markets. *American Economic Review* 81(May):222–226.

Giovannini, A. (1988). Exchange rates and traded goods prices. *Journal of International Economics* 24(February):45–68.

Goldberg, P., and M. Knetter. (1997). Goods prices and exchange rates: What have we learned? *Journal of Economic Literature* 35(September):1243–1272.

Golub, S. (1990). International capital mobility: Net versus gross stocks and flows. *Journal of International Money and Finance* 9:424–439.

Gordon, R., and L. Bovenberg. (1996). Why is capital so immobile internationally? Possible explanations and implications for capital income taxation. *American Economic Review* 86(December):1057–1075.

Greenspan, A. (1989). The economic value of ideas: Looking to the next century. *Japan Society of New York Newsletter* 36 (July).

Harrigan, J. (1993). OECD imports and trade barriers in 1983. *Journal of International Economics* 35:91–111.

Hau, H. (2000a). Real exchange rate volatility and economic openness: Theory and evidence. Centre for Economic Policy Research. Working Paper 2356.

———. (2000b). Exchange rate determination: The role of factor rigidities and nontradables. *Journal of International Economics* 50(April):421–427.

Helliwell, J. (1996). Do national borders matter for Quebec's trade? *Canadian Journal of Economics* 29(August):507–522.

———. (1998). *How Much Do National Borders Matter?* Washington: Brookings Institution.

Helpman, E. (1999). The structure of foreign trade. *Journal of Economic Perspectives* 13(Spring):121–144.

Hummels, D. (1999a). Toward a geography of trade costs. University of Chicago, Graduate School of Business. Mimeo.

———. (1999b). Have international transportation costs declined? University of Chicago, Graduate School of Business. Mimeo.

Isard, P. (1977). How far can we push the law of one price? *American Economic Review* 67(December):942–948.

Krugman, P. (1987). Pricing to market when the exchange rate changes. In *Real-Financial Linkages among Open Economies*, S. Arndt and J. Richardson (eds.). Cambridge, MA: The MIT Press.

———. (1991). *Has the Adjustment Process Worked?* Washington: Institute for International Economics.

Lane, P. (2001). The new open economy macroeconomics: A survey. *Journal of International Economics*, forthcoming.

Lee, J., and P. Swagel. (1997). Trade barriers and trade flows across countries and across industries. *Review of Economics and Statistics* 79:372–382.

Lewis, K. (1999). Trying to explain the home bias in equities and consumption. *Journal of Economic Literature* 37(June):571–608.

Lucas, R. E., Jr. (1982). Interest rates and currency prices in a two-country world. *Journal of Monetary Economics* 10(November):335–360.

Mankiw, N., and S. Zeldes. (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics* 29(March):97–112.

Marston, R. (1990). Pricing to market in Japanese manufacturing. *Journal of International Economics* 29(December):217–236.

Martin, P., and H. Rey. (1999). Financial supermarkets: Size matters for asset trade. Centre for Economic Policy Research. Working Paper 2232.

McCallum, J. (1995). National borders matter: Canada–U.S. regional trade patterns. *American Economic Review* 85(June):615–623.

Meese, R., and K. Rogoff. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics* 14(February):3–24.

Mendoza, E. (1991). Real business cycles in a small open economy. *American Economic Review* 81(September):797–818.

Michael, P., A. Nobay, and D. Peel. (1997). Transaction costs and nonlinear adjustments in real exchange rates: An empirical investigation. *Journal of Political Economy* 105(August):862–879.

Obstfeld, M. (1986). Capital mobility in the world economy: Theory and measurement. *Carnegie-Rochester Conference Series on Public Policy* 24(Spring):55–103.

———. (1995). International capital mobility in the 1990s. In *Understanding Interdependence: The Macroeconomics of the Open Economy*, P. B. Kenen (ed.). Princeton, NJ: Princeton University Press.

———, and K. Rogoff. (1996). *Foundations of International Macroeconomics*. Cambridge, MA: The MIT Press.

———, and ———. (1998). Risk and exchange rates. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6694.

———, and ———. (2000). New directions for stochastic open economy models. *Journal of International Economics* 50(February):117–153.

———, and A. Taylor. (1997). Nonlinear aspects of goods-market arbitrage and adjustment: Heckscher's commodity points revisited. *Journal of the Japanese and International Economies* 11(December):441–479.

O'Connell, P., and S. Wei. (1997). "The bigger they are, the harder they fall": How price differences across U.S. cities are arbitraged. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6089.

Ohanian, L., and A. Stockman. (1997). Arbitrage costs and exchange rates. University of Rochester, Mimeo.

Organization for Economic Cooperation and Development. (1996). *Indicators of Tariff and Nontariff Barriers*. Paris: Organization for Economic Cooperation and Development.

Portes, R., and H. Rey. (1999). The determinants of cross-border equity flows: The geography of information. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7336.

Radelet, S., and J. Sachs. (1998). Shipping costs, manufactured exports, and economic growth. Harvard Institute for Economic Development. Mimeo.

Rauch, J. (1999). Networks versus markets in international trade. *Journal of International Economics* 48(June):7–35.

Ravn, M., and E. Mazzenga. (1999). Frictions in international trade and relative price movements. London Business School. Mimeo.

Rogoff, K. (1996). The purchasing power parity puzzle. *Journal of Economic Literature* 34(June):647–668.

Rose, A. (2000). One money, one market: Estimating the effect of common currencies on trade. *Economic Policy* 30(April):7–45.

Rotemberg, J., and M. Woodford. (1992). Oligopolistic pricing and the effects of aggregate demand on economic activity. *Journal of Political Economy* 100 (December):1153–1207.

———, and ———. (1995). Dynamic general equilibrium models with imperfectly competitive output markets. In *Frontiers of Business Cycle Research*, T. F. Cooley (ed.). Princeton, NJ: Princeton University Press.

Samuelson, P. (1954). The transfer problem and transport costs: Analysis of effects of trade impediments. *Economic Journal* 64(June):264–289.

Sørensen, B., and O. Yosha. (1998). International risk sharing and European monetary unification. *Journal of International Economics* 45:211–238.

Stockman, A., and L. Tesar. (1995). Tastes and technology in a two-country model of the business cycle: Explaining international comovements. *American Economic Review* 85(March):168–185.

Taylor, A. (2001). Potential pitfalls for the purchasing-power-parity puzzle? Sampling and specification biases in mean-reversion tests of the law of one price. *Econometrica,* forthcoming.

Tesar, L., and I. Werner. (1998). The internationalization of securities markets since the 1987 crash. In *Brookings-Wharton Papers on Financial Services,* R. Litan and A. Santomero (eds.). Washington: The Brookings Institution.

Tille, C. (2000). "Beggar-thy-neighbor" or "beggar-thyself"? The income effect of exchange rate fluctuations. Federal Reserve Bank of New York. Mimeo.

Tinbergen, J. (1962). An analysis of world trade flows. In *Shaping the World Economy,* J. Tinbergen (ed.). New York: Twentieth Century Fund.

Trefler, D. (1995). The case of the missing trade and other mysteries. *American Economic Review* 85(December):1029–1046.

———. and H. Lai (1999). The gains from trade: Standard errors with the CES monopolistic competition model. University of Toronto. Mimeo.

van Wincoop, E. (2000). Borders and trade. Federal Reserve Bank of New York. Mimeo.

Wei, S. (1998). How reluctant are nations in global integration? Harvard University. Mimeo.

Williamson, J. (2000). Land, labor, and globalization in the preindustrial third world. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7784.