Real Exchange Rates, Saving and Growth:
Is there a Link?

Peter J. Montiel

Williams College

Luis Servén

The World Bank

Abstract

The view that policies directed at the real exchange rate can have an important effect on economic growth has been gaining adherents in recent years. Unlike the traditional "misalignment" view that temporary departures of the real exchange rate from its equilibrium level harm growth by distorting a key relative price in the economy, the recent literature stresses the growth effects of the *equilibrium* real exchange rate itself, with the claim being that a depreciated equilibrium real exchange rate promotes economic growth. While there is no consensus on the precise channels through which this effect is generated, an increasingly common view in policy circles points to saving as the channel of transmission, with the claim that a depreciated real exchange rate raises the domestic saving rate -- which in turn stimulates growth by increasing the rate of capital accumulation. This paper offers a preliminary exploration of this claim. Drawing from standard analytical models, stylized facts on saving and real exchange rates, and existing empirical research on saving determinants, the paper assesses the link between the real exchange rate and saving. Overall, the conclusion is that saving is unlikely to provide the mechanism through which the real exchange rate affects growth.

1

1 Introduction

The view that policies directed at the real exchange rate can have an important effect on economic growth has been gaining adherents in recent years (Hausmann and Rodrik 2003, Rodrik 2006, Prasad et. al. 2006, Eichengreen 2007, Rodrik 2007). This "new view" on the links between the real exchange rate and growth goes beyond the traditional "disequilibrium" view that persistent real exchange rate misalignment harms growth by distorting a key relative price in the economy, or that real exchange rate volatility is harmful for growth because it tends to obscure an important macroeconomic relative price signal. The traditional view was primarily concerned with the macroeconomic impact of temporary departures of the real exchange rate from its equilibrium level. In contrast, the new view focuses on highly persistent departures – extending over as long as several decades (see the examples in Rodrik 2007). We interpret this new view as stressing the growth effects of the *equilibrium* real exchange rate itself, with the claim being that a depreciated equilibrium real exchange rate promotes economic growth. However, the literature addressing this issue is in its infancy, and there is no consensus on the precise channels through which this effect is generated. There are two dominant views (Levy-Yeyati and Sturzenegger 2007a).

The first, articulated in both policy and academic circles, emphasizes the effect of real exchange rate depreciation on the composition of domestic production. Specifically, a depreciated real exchange rate shifts production from nontraded to traded goods. The link from output composition to growth is through economy-wide productivity improvements generated by the production of some types of traded goods (e.g., exported manufactured goods) through mechanisms such as technology and skill transfers associated with "learning by doing" that is external to the firm. If such mechanisms are important, a change in the composition of output in favor of goods that create external effects if this type would increase the economy's growth rate without necessarily affecting the rate of capital accumulation. We can refer to this as the "TFP growth" channel. The second, which seems to be increasingly common in policy circles, links a depreciated real exchange rate to growth through effects on the domestic saving rate. The claim here is that a depreciated real exchange rate tends to increase the domestic saving rate, and a higher saving rate stimulates growth by increasing the rate of capital accumulation.[1] We can refer to this as the "capital accumulation" channel. Both channels are potentially relevant, but the TFP growth channel, though it has recently drawn increased interest, is of older vintage (it provided one interpretation of Dutch disease in the literature of the 1980s, for instance), and has consequently received much more attention than the capital accumulation channel. Neither the internal consistency nor the empirical plausibility of the latter have yet been carefully explored.

---

[1] A depreciated real exchange rate could also encourage investment by raising profitability in the tradable sector, if the latter is more capital-intensive than the nontradable sector (e.g., as in Kouri 1978). However, this mechanism – recently stressed by Levy-Yeyati and Sturzenegger (2007b) -- is conceptually distinct from the saving channel that is the focus of this paper.

The view that a capital accumulation channel may be of potential empirical importance is motivated by the observation that, both currently as well as in the recent past, several countries that have achieved sustained high growth rates have simultaneously exhibited depreciated real exchange rates and high domestic saving rates. China is the most prominent current example, but over previous decades the association between depreciated real exchange rates and high saving rates has also been observed in high-performing Asian economies such as Korea, Malaysia, and Thailand, as well as in Chile. The opposite constellation of outcomes -- an appreciated real exchange rate, low domestic saving, and low growth -- has also been apparent in highly visible cases, for example pre-1995 Mexico. This raises the questions of whether there is a systematic macroeconomic relationship between the real exchange rate and the saving rate, and whether that relationship is the key to these countries' growth performance.

It is important to note that there are two conceptual links in the causal chain that underlies the capital accumulation channel: from the real exchange rate to the saving rate, and from the saving rate to growth. The second of these links is a familiar one, and has long underpinned mainstream views of the growth process even though, as we shall note below, it is far from uncontroversial, both theoretically as well as empirically. The first link is a more novel one. Analytically, it is based on the notion that if the real exchange rate is adopted as a policy target, and if a depreciation in the real exchange rate is to improve a country's current account balance, then it must increase the country's saving rate relative to its rate of investment. The capital accumulation view implicitly assumes that the increase in the domestic saving rate is devoted only in part to an accumulation of claims on the rest of the world (a current account improvement). The rest of it results in an increase in domestic investment that increases the country's growth rate. Conditions for the latter to be true are of interest in their own right, but more intriguing -- because less familiar --- is the relationship between the real exchange rate and the saving rate. That relationship is the focus of this paper.

Some observers have recently argued that such a relationship exists, but have not necessarily interpreted it as reflecting causation from a depreciated real exchange rate to a high saving rate to a high growth rate driven by capital accumulation. Bernanke (2005), for example, argues that, in order both to wean themselves from foreign financing of domestic investment as well as to amass a stock of liquid claims on the rest of the world as self-insurance against a repeat of the 1997-98 Asian crisis, several important developing countries have explicitly sought to achieve high domestic saving rates through means other than exchange rate policy. However, because a high saving rate tends to depress domestic demand, these countries have been able to sustain internal balance only by maintaining a depreciated real exchange rate. In Bernanke's view, therefore, causation runs from a high saving rate to a depreciated real exchange rate. If this view is correct, an empirical correlation between saving rates and real exchange rates therefore cannot be interpreted as evidence of a capital accumulation channel at work.

In contrast to the direction of causation emphasized by Bernanke, other an-alysts have devised models in which an explicit policy of pursuing a depreciated real exchange drives a high domestic saving rate. For example, Dooley, Folkerts-Landau and Garber (DFG, 2004) argue that a depreciated real exchange rate has

long been central to the export-driven development strategy pursued by several Asian countries, including China. In their view, high saving rates in these countries are at least in part the result of the pursuit of such an exchange rate policy. The link between a depreciated real exchange rate and the saving rate arises because a depreciated real exchange rate tends to shift aggregate demand away from traded to nontraded goods, requiring an increase in the real interest rate to maintain internal balance. The high real interest rate restrains aggregate demand in part by increasing the domestic saving rate. Thus, from this perspective causation runs from the real exchange rate through the real interest rate to the saving rate.

Levy-Yeyati and Sturzenegger (LYS, 2007b) also claim that a more depreciated real exchange rate results in higher saving, but through a different channel: a more depreciated exchange rate is associated with lower real wages, inducing firms to invest more and to increase their saving to finance the additional investment, thereby raising overall saving.

While all of these models describe links between the real exchange rate and the saving rate, they do not provide the analytical basis for a causal chain from a depreciated real exchange rate to a higher saving rate and from a higher saving rate to a higher rate of capital accumulation[2]. As already mentioned, in Bernanke's framework causation runs from the saving rate to the real exchange rate, rather than vice versa. In the DFG framework, the direction of causation runs in the right direction -- from the real exchange rate to the saving rate -- but the next link in the capital accumulation channel --from higher saving to more rapid growth --would not tend to emerge in that framework, because the same high real interest rate that induces a higher domestic saving rate would also tend to discourage domestic investment, making this a better story about the current account outcome (which is what it was intended to be) than about domestic capital accumulation. Finally, even if a more depreciated real exchange rate induces higher investment and thus higher corporate saving as in the LYS framework, causation also runs in the wrong direction in this framework: from the real exchange rate directly to investment and from the latter to saving. Thus, exchange rate policy affects investment directly, rather than through saving.

Accordingly, both the existence of a link between the real exchange rate and the saving rate, as well as the interpretation of that link if it exists, are open to question. This paper is a preliminary exploration of both of these issues. It seeks to answer the following questions:

a. Is the association between the real exchange rate and the saving rate in the episodes described above *sui generis*, or is there a broader correlation between real exchange rates and saving rates in the international experience?

b. Do standard open-economy models offer an analytical basis for a causal relationship between exchange rate policy -- and specifically a depreciated real exchange rate -- and the national saving rate?

c. If an analytical basis for such a link exists, is there any empirical evidence for it? Specifically, what role --if any -- does exchange rate policy play in the literature on cross-country differences in saving rates? In particular, to what

---

[2] To be fair, none of these models was designed for this purpose.

extent have explanations of saving behavior in high-saving countries identified the behavior of the real exchange rate as a key determinant of the saving rate?

The structure of the paper is as follows. In the next section we examine the empirical evidence on the correlation between the real exchange rate and saving in the international experience, to determine whether this correlation has arisen only in special cases or is a broader feature of that experience. Section 3 explores the effects on the saving rate of an exchange rate policy designed to sustain a depreciated value of the real exchange rate in the context of a standard intertemporal open-economy model. Section 4 provides a brief review of the empirical literature on the determinants of cross-country differences in saving rates, as well as of that explaining saving behavior in the high saving-high growth countries, to establish what - - i f any -- role has been accorded to the behavior of the real exchange rate in this literature. The final section summarizes the implications of our findings for the capital accumulation channel.

## 2 Saving and the Real Exchange Rate: International Evidence

Aside from the episodic evidence mentioned earlier, is there more systematic evidence linking real exchange rates and saving rates? To determine whether the association between the real exchange rate and saving rate is a characteristic of the international experience, in this section we examine the correlation between the saving rate and the real exchange rate in a broad cross section of countries. The real exchange rate, defined as the relative price of traded goods in terms of nontraded goods, is measured by the (log) GDP deflator from the Penn World Tables, which captures each country's national price level in dollar terms (relative to the U.S.). If the law of one price holds, the dollar price of traded goods will be equalized across countries, but the dollar price of nontraded goods will be lower where nontraded goods are cheaper -- i.e., where the real exchange rate is more depreciated. Hence the dollar price level, which is the weighted average of the dollar prices of traded and nontraded goods, should be lower where the real exchange rate is more depreciated. In turn, the saving rate is measured as the gross national saving rate.[3]

Using these measures and international panel data covering 94 countries over 1975-2005, Figure 1 plots the national saving rate against the (log) real exchange rate. In this context, a higher value of the latter represents a real appreciation. Since our interest is in the relationship between the equilibrium real exchange rate and the saving rate, lower-frequency observations of both variables are of greater interest than annual data. Accordingly, the figure plots decade averages

---

[3] Gross national saving is defined as the difference between gross national disposable in-come (GNDI for short, equal to GNP plus net unrequited transfers from abroad) and total consumption expenditure. The saving rate therefore is the ratio of gross national saving to GNDI. See Schmidt-Hebbel and Servén (2000) for an overview of this and other measures of aggregate saving.

(top panel) as well as 30-year averages (bottom panel). The results are very similar in both cases: a higher saving rate is strongly associated with a more *appreciated* real exchange rate -- the opposite of the view held in the recent literature cited earlier. Each figure shows a regression line drawn through the scatter plot. The slope of this line is positive and highly significant in both cases.[4] Furthermore, this same pattern emerges in other exercises (not shown to save space) that consider countries grouped by geographic region. All regions considered show a positive correlation between saving and the real exchange rate, at both time frequencies analyzed.[5]

Figure 1.   Cross-Country Association between the Real Exchange Rate and Saving Rate.
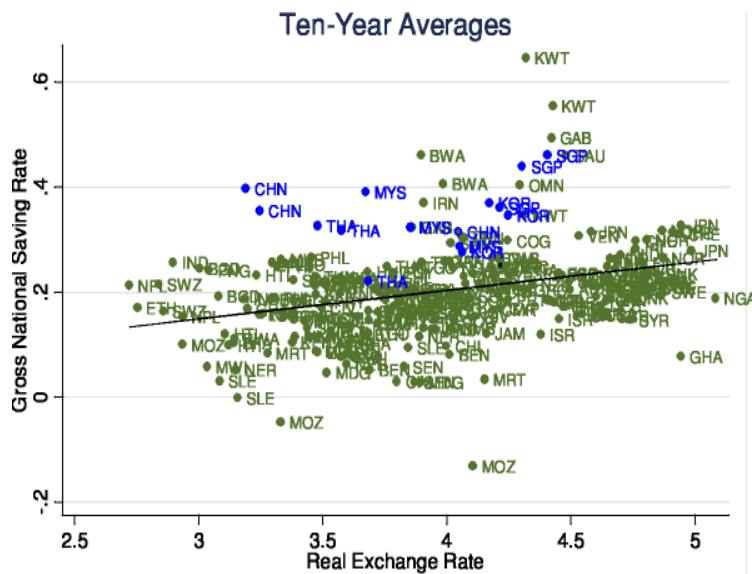
4 Robust standard errors are used to mitigate the potential impact of outliers.

[5] However, the correlation is not always significant, due mainly to the varying sample sizes by region.

6

Thirty-Year Averages
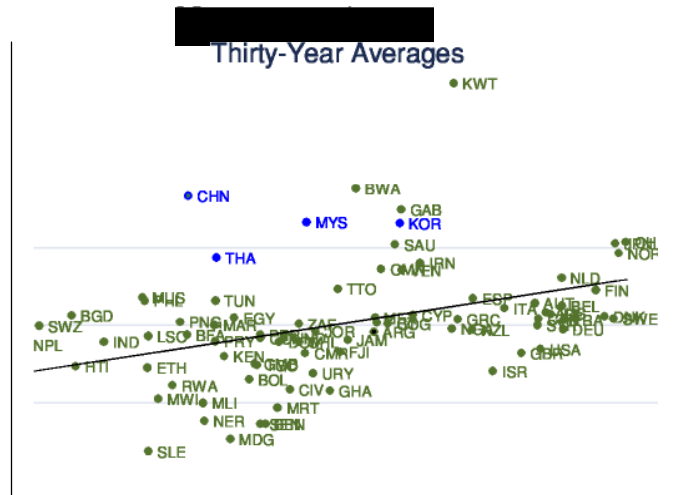
Looking at the 30-year averages, it is interesting to note that China (labeled CHN) is a substantial outlier, with a much higher saving rate than would be predicted from its real exchange rate on the basis of the international experience. The same is true for the high-growth Asian economies of Korea, Malaysia, Singapore and Thailand. However, among these high-saving economies, only China, Thailand and Malaysia appear to have sustained relatively depreciated real exchange rates over this 30-year period. Korea and Singapore have operated with real exchange rates that have been relatively appreciated compared to the average in the sample.

Simple correlations of the type just presented are useful as a corrective to the temptation to draw generalizations from a small set of highly visible episodes, but they can only provide limited evidence about the empirical relevance of the postulated link between the real exchange rate and the saving rate. The link hypothesized in the capital accumulation channel is between a real exchange rate that is depreciated by policy relative to what it would otherwise have been and a saving rate that is high relative to what it would otherwise have been. Simple correlations do not satisfy the *ceteris paribus* conditions on which this link is predicated. Because these unconditional correlations do not hold constant other real exchange rate and saving determinants, the postulated relationship between the real exchange rate and the saving rate might be hidden in these correlations. One clear candidate for such a determinant is the level of per capita income, which has been found to exert a positive impact both on the level of the real exchange rate (through the Balassa-Samuelson effect) as well as on saving rates (e.g., Loayza, Schmidt-Hebbel and Servén 2000).

To control for these effects, it is useful to examine the conditional correlation between saving and the real exchange rate, after removing the influence of per capita income from both variables. Figure 2 plots the residuals from regressing saving rates and real exchange rates on (log) per capita income, with the latter

7

taken from the Penn World Tables (we refer to these as orthogonalized residuals). As before, the figure shows ten-year and long-term averages. In contrast with the previous figures, the scatters now reveal a negative slope: more depreciated real exchange rates are associated with higher national saving rates. The slope of the regression line between the two variables is very small, however. It is statistically significant using the 10-year averages, but not in the longer-term data.

Figure 2.   Cross-Country Association between the Real Exchange Rate and Saving Rate: Orthogonalized Residuals.

Thirty-Year Averages
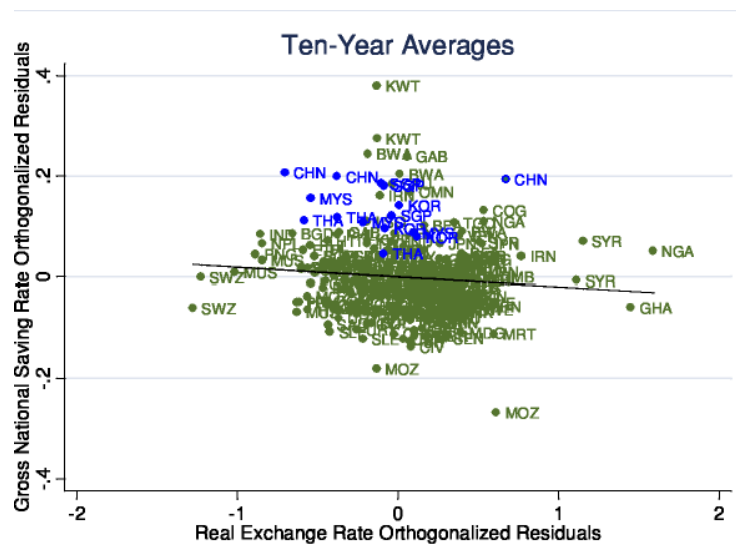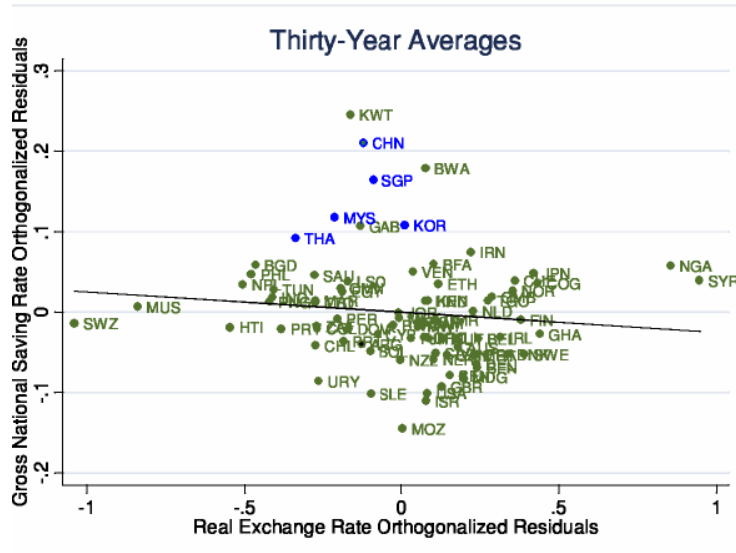
Looking at the bottom panel in the figure, it remains the case that the high-growth Asian economies of China, Korea, Malaysia, Singapore and Thailand have unusually high saving rates. All of these countries have higher saving rates than would be predicted from their real exchange rates on the basis of the international experience, even after removing the effects of income levels on both variables.

To summarize, there is only modest evidence that more depreciated real exchange rates are associated with higher saving rates, as postulated in the capital accumulation channel. Indeed, contrary to this view, the unconditional correlation between saving and the level of the real exchange rate is strongly positive, though this appears to at least partially reflect the positive association of both variables with the level of per capita income. Controlling for this factor changes the picture, revealing a negative correlation between saving rates and the real exchange rate. However, the correlation is very small and statistically significant only at the higher 10-year frequency. It is not statistically significant at the 30-year horizon. This latter fact makes it problematic to interpret the negative correlation as reflecting the effects of a sustained real exchange rate depreciation on the saving rate, as would be expected from an entrenched development strategy. Overall, then, the international experience does not appear to provide much empirical support for a growth strategy predicated on the effects of a depreciated equilibrium real exchange rate on the saving rate, and the effects of the latter on the rates of capital accumulation and growth. [4]

However, the crude nature of this empirical exercise leaves its results open to different interpretations. One such interpretation is that the hypothesized negative link between the real exchange rate and the saving rate emerges only

---

[4] To be clear, this is not to say that the level of the real exchange rate is of no consequence for long-term growth. What we are questioning here, in the light of the international evidence, is the role of the saving rate as the key channel of transmission from the real exchange rate to growth.

9

when the real exchange rate is explicitly used as an instrument of development policy. In the absence of such a systematic policy, the real exchange rate and saving rate are both endogenous variables, and the correlation between them reflects the influence of a variety of exogenous variables whose behavior differs across countries and over time. This being the case, we should perhaps not expect to observe a systematic correlation in the international data. However, if the hypothesized link is valid, the subset of countries that have adopted the real exchange rate as a tool of development policy should exhibit both more depreciated real exchange rates and higher saving rates than would be predicted from their income levels. In other words, such countries should be found in the Northwest quadrant in the bottom panel of Figure 2. Since this is indeed the case for most of the high-growth Asian economies that have been said to adopt this policy, the evidence may not be inconsistent with the real exchange rate-saving link emphasized in the capital accumulation channel after all. This justifies probing deeper in the sections that follow. Specifically, we ask whether the proposed link is analytically consistent with standard open-economy macroeconomic models, and whether research on the determinants of saving rates --including in the high-growth Asian economies themselves -- has uncovered evidence of a causal link from exchange rate policy to the saving rate.

3 Saving and the Real Exchange Rate: Analytics

As argued in the introduction, the analytical underpinnings of a channel of transmission from exchange rates to growth through the saving rate remain to be established. In this section we examine these analytical underpinnings by considering whether modern open-economy macroeconomic models would tend to predict the existence of such a link. Specifically, we analyze the effect of a depreciation of the equilibrium real exchange rate on the saving rate in the context of a representative-agent, intertemporal-optimizing model of the type that has become the standard framework of analysis in the "new" open-economy macroeconomics. The distinguishing feature of the model we use is the fact that the real exchange rate is a policy target. The main conclusion from this type of model is that the saving rate will increase in response to a depreciation of the equilibrium real exchange rate only if the real exchange rate depreciation generates a transitory increase in real income. As we will argue in this section, it is possible to construct such a mechanism, but it is by no means easy to do so, and the empirical relevance of the particular mechanism that we describe is questionable at best.

3.1 The model

We use the simplest possible model to investigate this issue analytically. The framework is taken from Montiel (2007), where the technical details are described. To conduct this analysis, we need to consider an economy with two domestic production sectors --producing traded and nontraded goods -- so that

10

we can define the real exchange rate (the relative price of traded goods in terms of nontraded goods), and distinguish it from the terms of trade (the price of importables in terms of exportables). Since our interest is only in household saving behavior, we will neglect investment by assuming that each sector possesses a fixed and sector-specific capital stock. We assume that a given stock of homogeneous labor can move freely between sectors, allowing sectoral output levels to change. The model contains a representative household that supplies a fixed amount of labor to the economy, and chooses a path of consumption over time that it allocates between traded and nontraded goods so as to maximize discounted utility over an infinite horizon, subject to an intertemporal budget constraint that depends on its initial financial wealth and expected flow of future income.

Below we use the model to explore the effects of real exchange rate targets achieved through suitable nominal exchange rate policy. To properly consider the latter, the model has to describe a monetary economy, so our representative household will also have to decide how to allocate its financial wealth at each instant between money and other assets (which we consider to be internationally-traded interest-bearing "bonds"). We therefore also need to explain why such a household would hold money at all, as opposed to interest-paying assets. There are various ways to motivate the holding of money in models of this type, including inserting money directly in the household's utility function or imposing a cash-in advance transactions technology. We adopt a generalized and more flexible version of the latter, reflecting the view that money is valued because it reduces the transactions costs associated with consumption, but allowing households to alter the ratio of money holdings to consumption expenditures in response to changes in the opportunity cost of holding money. As we will discuss further below, transactions costs may be incurred in the form of either of the two types of goods in the model. The model does not really require a government sector (other than a central bank to formulate exchange rate policy), so we assume that any revenues generated by seignorage are simply rebated to the household in lump-sum fashion.

As a final preliminary matter, we need to describe the mechanism that allows the economy to reach a steady-state equilibrium. The difficulty here is that the representative household will seek to adopt a time-varying path of consumption as long as its rate of time preference differs from the market real interest rate, because only by doing so would it be able to equalize the marginal utility of consuming at different points in time. If the rate of time preference and market real interest rate happen to be equal to each other, the household chooses a flat consumption profile, a situation which does not lend itself to analyzing the question we are interested in in this section, because in this case satisfying the intertemporal budget constraint would imply that the household never saves. Thus, we need to adopt an analytical framework that allows for at least temporary differences between the household rate of time preference and the real interest rate. As pointed out by Schmidt-Grohe and Uribe (2003), there are essentially three ways to "close" small open-economy models while retaining this property. We choose the simplest of these, and one that has some empirical support. Specifically, we take the household's rate of time preference to be exogenous, but allow the market real interest rate faced by the representative domestic household to respond endogenously, by modeling the risk premium faced by the

11

domestic economy as a decreasing function of the economy's net international investment position (its net financial claims on the rest of the world). A steady state is reached in this framework when, given the world risk-free interest rate, the domestic economy accumulates enough net claims on the rest of the world to make this risk premium small or large enough to close the gap between the domestic rate of time preference and the world risk-free interest rate. We are interested in the behavior of the saving rate during the transition path to this steady state.

The issue that we seek to investigate can be described as follows. Suppose that the domestic central bank simply targets some nominal variable --say the nominal exchange rate -- and keeps it at a fixed value, allowing the economy to reach a steady-state equilibrium in which the real exchange rate adjusts endogenously to whatever equilibrium value is implied by the economy's "real" fundamentals. We can label this the "baseline" scenario. Now consider what would happen if instead of targeting the nominal exchange rate the central bank adopts a specific *real* exchange rate target that differs from that which prevails under the baseline scenario, and specifically that is more depreciated than that prevailing in the baseline scenario. It can do so using the nominal exchange rate instrument at its disposal by devaluing the nominal exchange rate, and then continually depreciating it by an amount equal to the contemporaneous difference between the domestic inflation rate and the inflation rate of the country's trading partners. This policy can be described as "real exchange rate targeting." The questions are: what does the economy's new steady state look like under this policy, and -- more importantly for present purposes -- how does the saving rate behave during the transition from the original steady state to the new one?[5]
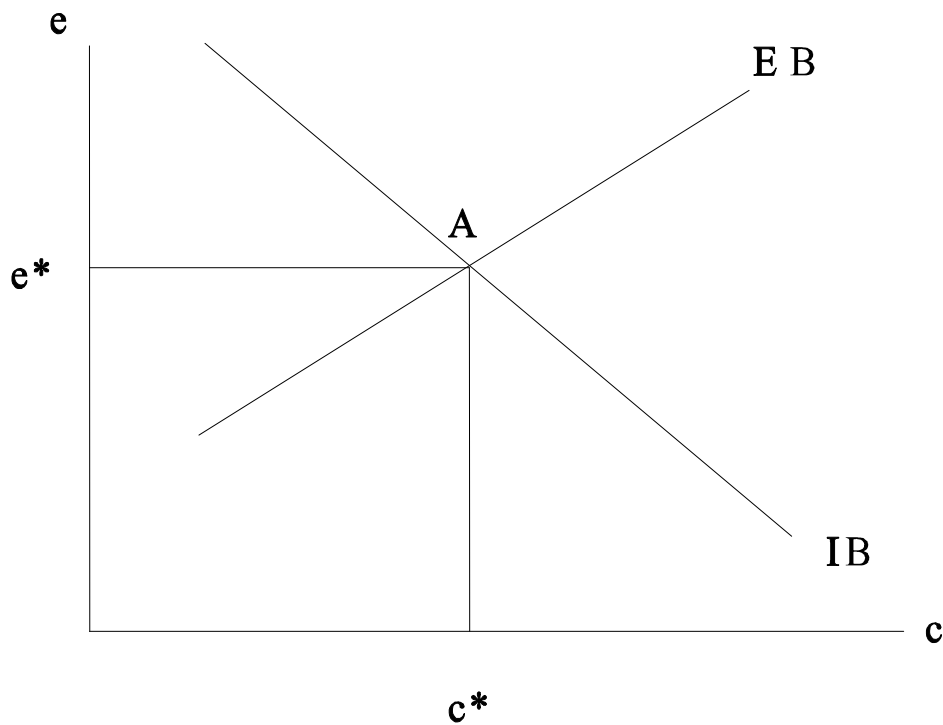
3.2 The baseline

A reasonable point of departure is to examine what the economy's equilibrium looks like in the baseline scenario -- i.e., in the absence of an explicit real exchange rate policy. In the steady state, household consumption must be constant and equal to household income, so the saving rate is zero -- the domestic economy is neither accumulating nor decumulating claims on the rest of the world. As mentioned above, the economy's steady-state net claims on the rest of the world must be such as to generate a risk premium in international financial markets that just bridges the gap between the domestic rate of time preference and the world real interest rate. Thus, the domestic economy's net claims on the rest of the world will be small (and its risk premium correspondingly high) when it has a relatively high rate of time preference or when the world interest rate is low. Both of those conditions are associated with more international borrowing by the domestic economy.

The economy's steady state must satisfy two conditions: an external balance condition, under which the country's trade balance (the excess of its production of

---

[5] In this model, steady-state saving is constant and, with no capital depreciation, equal to zero. Along the adjustment path, however, saving can show persistent departures from its steady-state level in response to disturbances. Such persistent swings in the saving rate are the focus of our analysis,

traded goods over its consumption of such goods) must be just equal to its (positive or negative) net income on the financial claims that it holds on the rest of the world, and an internal balance condition, which requires the market for nontraded goods to be in equilibrium (since such goods are only produced and consumed domestically, domestic output of nontraded goods must be equal to domestic demand for them in equilibrium). Two endogenous variables adjust to ensure that these conditions are met in the baseline scenario: the real exchange rate and the level of private consumption expenditure. A depreciation of the real exchange rate increases output of traded goods and decreases output of nontraded goods, as labor moves from the nontraded to the traded goods industry. At the same time, for a given level of private consumption expenditure, the depreciation switches household demand from traded to nontraded goods. Thus, a depreciation of the real exchange rate increases the excess supply of traded goods and reduces the excess supply of nontraded goods. On the other hand, an increase in private consumption expenditure increases excess demand for both types of goods. It follows that to maintain internal balance (equilibrium in the market for nontraded goods) in the steady state, an increase in private consumption expenditure must be associated with an appreciation of the equilibrium real exchange rate, but to sustain external balance, an increase in private consumption expenditure must be associated with a depreciation of the equilibrium real exchange rate. These relationships are illustrated in the form of the positively-sloped external balance locus EB and negatively-sloped internal balance locus IB in Figure 3. The equilibrium values of the real exchange rate and the equilibrium level of household consumption are determined at the intersection of these two locii at the point A, which therefore determines the economy's steady-state equilibrium real exchange rate and steady-state level of private consumption expenditure.

13

Figure 3. Steady-State Equilibrium with Nominal Exchange Rate Targeting

The equilibrium level of the domestic inflation rate plays a role in determining the equilibrium values of these two variables in the baseline scenario. Since the real exchange rate is constant in the steady state, the steady state inflation rate must be equal to the rate of depreciation of the nominal exchange rate, which is determined exogenously by the central bank, plus the world rate of inflation. If the nominal exchange rate is fixed, for example, the domestic inflation rate will be equal to the world inflation rate in steady-state equilibrium. This steady-state inflation rate is added to the steady-state real interest rate (recall that this is equal to the domestic rate of time preference) to determine the economy's steady-state nominal interest rate. In turn, this determines the real demand for money held by households per unit of consumption, and the transactions costs that households must incur to consume. The reason that this affects the economy's steady-state equilibrium values of the real exchange rate and private consumption is that these transactions costs affect the effective supplies of traded and nontraded goods in the economy. If these costs are borne in the form of nontraded goods, for example, then a higher steady-state inflation rate will reduce the effective supply of such goods, shifting the internal balance locus downward, implying a more appreciated equilibrium real exchange rate and lower level of private consumption. If they are borne in the form of traded goods, then it is the supply of

14

traded goods that will be reduced. In this case, the external balance locus is shifted upward, implying a more depreciated equilibrium real exchange rate and once again, a lower level of private consumption.
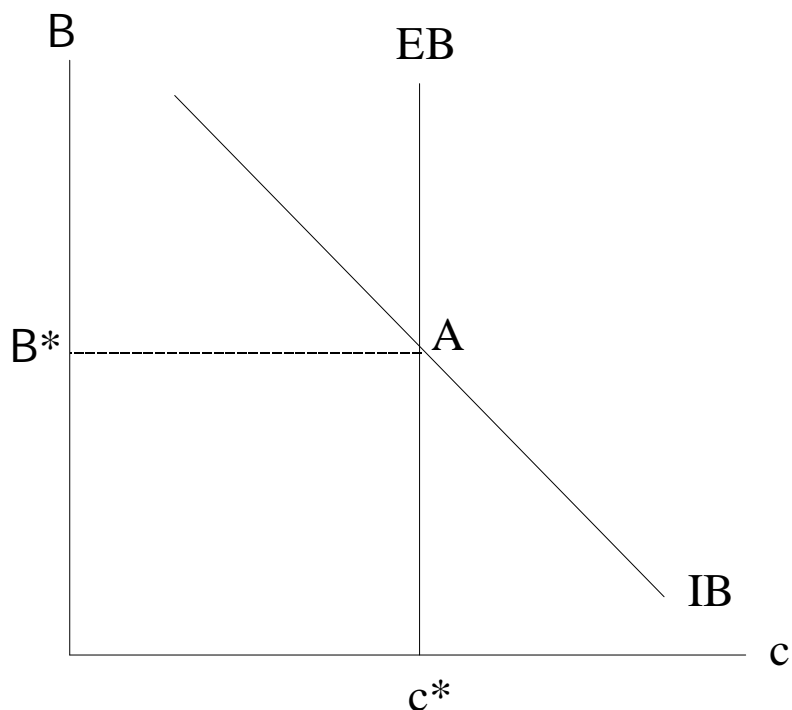
3.3 Real exchange rate targeting

To examine the effect of policies directed at the real exchange rate on the saving rate, consider next what happens to the economy if the central bank decides to adopt a real exchange rate target, and proceeds to target a more depreciated real exchange rate than that corresponding to the baseline steady state just described.

This policy implies that, since the real exchange rate is set exogenously by the central bank, it can no longer play the equilibrating function in domestic markets described above. If the real exchange rate cannot adjust to allow the economy to simultaneously achieve internal and external balance, then something else must do so. Notice that under real exchange rate targeting, the central bank is no longer fixing the value of the nominal exchange rate. Thus, the steady-state inflation rate is no longer pinned down by central bank policy. Formally, the steady-state inflation rate becomes the endogenous variable that adjusts to allow the economy to simultaneously achieve internal and external balance.

To see how this works, suppose that transactions costs are incurred in the form of nontraded goods only, and consider how the inflation rate and level of private consumption expenditure jointly interact to ensure internal and external balance under real exchange rate targeting. The analysis is provided in Figure 4. Since the real exchange rate is now an exogenous policy variable and the inflation rate is the adjusting endogenous variable, the steady-state inflation rate replaces the real exchange rate on the vertical axis. Since all transactions costs are borne in the form of nontraded goods, the inflation rate has no effect on the external balance locus, which is therefore vertical at the level of private consumption expenditure required to maintain external balance (c*). However, the internal balance locus has a negative slope: an increase in private consumption expenditure creates an excess demand for nontraded goods, which requires a reduction in the domestic inflation rate to reduce transactions costs and thus free up the resources required to meet the additional demand for non-traded goods. The economy's steady state inflation rate and level of private consumption expenditure are determined at the intersection of these two locii.

Figure 4.   Steady-State Equilibrium under Real Exchange Rate Targeting

Suppose that initially the central bank sets the real exchange rate target at a level equal to the equilibrium real exchange rate from Figure 3. In that case, the domestic inflation rate must be equal to the world inflation rate as before, since the new regime just mimics the baseline scenario. But now suppose that the central bank adopts a more depreciated real exchange rate target. Notice that this change in the real exchange rate target has no effect on the domestic rate of time preference or on the world risk-free real interest rate, so the economy's steady-state level of net financial claims on the rest of the world is unaffected. However, the policy disturbs both the external and internal balance locii, and so will alter the steady-state values of domestic inflation and private consumption expenditure. In this case, the external balance locus will shift to the right, as in Figure 5, since with a more depreciated real exchange rate, which increases production of traded goods and switches private consumption away from traded goods, a higher level of total private consumption expenditure is compatible with external balance. At the same time, the internal balance locus shifts to the left, since the reduced supply of nontraded goods induced by the real depreciation calls for a reduced level of total private consumption expenditure to maintain equilibrium in the market for nontraded goods. The upshot is that the steady-state inflation rate falls and the steady-state level of

16

private expenditure increases. The reason that households are able to sustain a permanently higher level of private consumption expenditure in this case is that the reduction in the domestic inflation rate effectively increases the supply of nontraded goods available to the economy (by reducing transactions costs) and thus acts like a favorable permanent productivity shock. In other words, households are able to sustain higher consumption levels because they have a higher income level.

Figure 5. Steady-State Effects of a Depreciated Real Exchange Rate Target
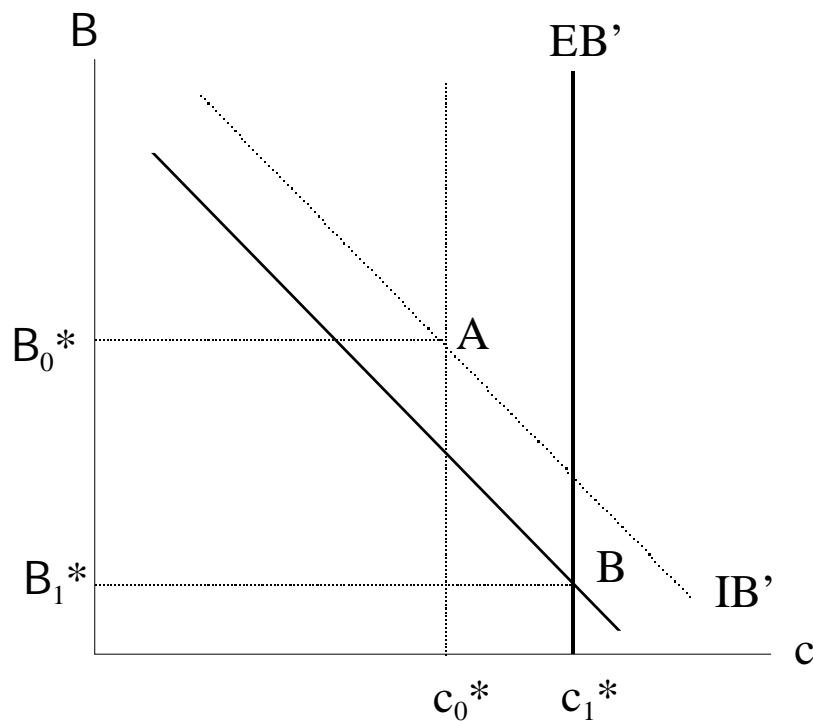
The important point, however, is that this happens as soon as the more depreciated real exchange rate target is adopted. The inflation rate falls immediately, the effective supply of nontraded goods rises immediately, and the level of private consumption also rises immediately. Thus, household consumption expenditure immediately rises to meet the new higher level of permanent income. Consequently, this policy does not create a gap between household income and household consumption expenditure -- in other words, there is no effect on the national saving rate. This is an important result, because it highlights the difficulty of producing a real exchange rate effect on the saving rate in a standard intertemporal optimizing model. To achieve such an effect what is required is that the real depreciation create an incentive for households to temporarily depress

17

their consumption relative to their income levels. Since an unanticipated, permanent depreciation in the central bank's real exchange rate target does not provide such an incentive, it has no effect on the saving rate.

Is there any way to create such an effect in the context of this model? The answer is yes. Consider what happens if the pursuit of the more depreciated real exchange rate target is (and is correctly expected to be) transitory, rather than permanent. The result is now quite different. Again, the steady-state value of the economy's net international investment position is unchanged. But now, while the more depreciated real exchange rate target is in place, output of the traded goods sector expands and that of the nontraded goods sector contracts. Given the more depreciated real exchange rate, private consumption demand also changes its composition, with households demanding more non-traded goods and fewer traded goods at any given level of private consumption expenditure. Thus, this situation creates an incipient excess demand for non-traded goods. To clear the nontraded goods market, the inflation rate must fall to reduce transactions costs and thereby increase the supply of nontraded goods. This reduction in the inflation rate acts like a positive productivity shock, because the reduction in transactions costs that it induces effectively increases the total supply of goods available for consumption in the economy. But this effective increase in the economy's productivity arising from a reduction in the inflation rate is only temporary: it only lasts as long as the more depreciated real exchange rate target is kept in place. Because this increase in real income is only temporary, consumption-smoothing motives prevent the representative household from consuming all its income gains while they are in place. The household therefore increases its consumption level, but by *less* than the increase in its real income -- in other words, it is induced to save in order to be able to sustain a higher level of consumption after the real exchange rate target is abandoned and the inflation rate and associated transactions costs return to their initial levels.

It may be worth noting that, while the mechanism just described indeed generates a positive effect of a more depreciated real exchange rate target on the household saving rate, it falls well short of establishing the analytical link between a permanently more depreciated real exchange rate and a higher household saving rate that would be required to explain low-frequency correlations between the real exchange rate and the saving rate in cross-country data. The reasons are two. First, the emergence of this association explicitly requires a transitory depreciation in the central bank's real exchange rate target, rather than a permanent one. However, this may be acceptable, since the duration of the real exchange rate depreciation would coincide with that of the positive effect on the saving rate. But second --and more importantly --it can be shown that the longer the more depreciated real exchange rate target is expected to be sustained, the *weaker* will be the observed positive effect on the household saving rate, even while the more depreciated real exchange rate is in place. The reason is that the longer the positive productivity effect is expected to last, the higher the value of consumption the household can maintain without violating its intertemporal budget constraint. Thus, even if this mechanism is operative, sustained equilibrium real exchange rate depreciations of the type we have been considering would at best tend to have weak effects on the household saving rate.

18

3.4 Implications

The model we have just analyzed contains a very specific mechanism linking the level of the real exchange rate to the saving rate --i.e., through the equilibrium rate of inflation and its effects on the transactions costs associated with consumption. But this particular mechanism serves as an illustration of a more general point. Essentially what is required for the pursuit of a more depreciated real exchange rate target to generate an increase in the saving rate in the popular open-economy framework just described is that the real exchange rate depreciation generates a front-loaded increase in real income -- i.e., an increase that is greater in the short run than in the long run. The increase in real income needs to be front-loaded so that consumption smoothing will cause households to save some of their additional real income rather than simply consume it all, as they do in the model when the more depreciated real exchange rate is adopted as a permanent target.

One way to achieve this analytically is as we have done it in this section -- i.e., by building into the model a mechanism through which the timing of the real income rise coincides with the period of real exchange rate depreciation, and by making the real exchange rate depreciation itself temporary. Though the inflation-cum-transactions cost mechanism incorporated in the model of this section produces this result, the empirical relevance of such a mechanism is questionable, and it is hard to think of other mechanisms that would have a similar effect. An alternative way to achieve the same result would be for a permanent real depreciation to result in a short-run productivity improvement greater than its long-run counterpart. However, it is hard to think of mechanisms that would have such an effect. Certainly the types of effects on TFP expected under the TFP channel discussed above, in which the rate of growth of domestic TFP depends on the share of traded goods production in GDP, would be unlikely to behave in this way. The reason is that the gradual accretion of economy-wide productivity improvements caused by an increase in the level of production of traded goods would accumulate over time, resulting in larger long-run than short-run increases in productivity. This would be particularly so in the (realistic) scenario in which the relocation of productive inputs is subject to adjustment costs. If productivity levels are expected to increase over time, however, forward-looking optimizing households would tend to increase consumption in anticipation of higher future incomes, thus reducing, rather than raising, the saving rate on impact.

The upshot is that, while it is possible to generate positive effects on the saving rate of a depreciation in the equilibrium real exchange rate in a mainstream open-economy model, it is by no means easy to do so, and it is unclear whether the types of effects that can be generated in this framework could be empirically important.

19

4 Interpreting High Saving Rates around the World

While theory is thus not inconsistent with the proposition that a more depreciated real exchange rate could have a positive effect on household saving, the conditions for this effect to emerge in standard open-economy models appear to be stringent. Nevertheless, as long as such a link is admissible theoretically, it is at least possible that it has been important empirically. This section examines whether such a link, featuring causation from the real exchange rate to the saving rate, has indeed been identified as empirically important in determining cross-country differences in saving rates.

The factors behind these differences have been the focus of a large literature, which we can only touch upon. The issue we explore here is what role exchange rate policy, and in particular a depreciated real exchange rate, has played in existing explanations of cross-country differences in saving rates. We consider first the broad international evidence, and then focus more specifically on explanations for high saving in East Asia, the region of the world where the real exchange rate-saving link has received the most attention.

4.1 Determinants of international differences in saving rates

As the empirical literature has noted, there are significant regional differences in saving rates, with East Asia and Sub-Saharan Africa at the extremes. Countries in the former region have tended to save in excess of 30 percent of GDP on average, while those in the latter have typically featured saving rates in the range of 10-15 percent of GDP. Industrial countries are in an intermediate position, generally saving in the vicinity of 20 percent of GDP. As we will discuss below, explanations of cross-country differences in saving rates have drawn on a variety of different saving determinants suggested by theory, but two have played particularly prominent roles:[6] the level and growth rate of income per capita, both of which tend to be positively related to saving rates. The role of the level of income per capita may arise because of the pressures of subsistence needs at low levels of income, such as would emerge under a Stone-Geary specification of utility. This interpretation is supported by the fact that the positive effect of income levels on private saving rates tends to taper off at higher levels of income (see Masson, Bayoumi, and Samiei 1998, as well as Loayza, Schmidt-Hebbel and Serven 2000a,b). In turn, the effect of growth on the saving rate may be due to a growth-associated shifting of income from low-saving older cohorts to high-saving young cohorts within countries, or to the influence of habit formation on consumption, which causes changes in consumption to lag those in income (Deaton 1999 offers a comprehensive discussion). In any case, almost all studies of international differences in saving rates include these variables, and they are often found to have statistically significant explanatory power.

---

[6] See, for example, the overview by Schmidt-Hebbel, Serven, and Solimano (1996).

Aside from the basic roles of income per capita and growth rates, the large cross-country differences in saving rates has caused researchers to examine the roles of a large number of other explanatory factors suggested by theory. These include structural features of economies such as demographic composition, the presence of mandatory saving schemes, uncertainty in the macroeconomic environment, financial depth and the importance of borrowing constraints; short-run policy variables such as fiscal policy (because of potential Ricardian equivalence effects), interest rates, and tax incentives; and exogenous shocks such as inflows of foreign aid and fluctuations in the terms of trade. Results can be summarized briefly: increases in the young-age and old-age dependency ratios both tend to lower private saving rates, in line with life-cycle theories of consumption. Research on countries with particularly high saving rates, such as Singapore, or that have secured a dramatic increase in saving over a short period of time, such as Chile, has emphasized the role of mandatory saving schemes such as Singapore's Central Provident Fund (e.g., World Bank 1994), finding that such schemes have had a significant effect on the domestic saving rate. Uncertainty in the macroeconomic environment, arising either from the structure of production (e.g., agricultural income tends to be more uncertain than non-agricultural income) or from poor domestic macroeconomic management, has been found to increase saving for precautionary reasons (although uncertainty in the domestic macroeconomic environment may cause such saving to be invested abroad). Financial structure seems to matter, but in ambiguous ways: in particular, borrowing constraints have been found to exert a strong positive impact on household saving (Jappelli and Pagano 1994), but the effects of financial development (which is often associated with a relaxation in borrowing constraints) are less clear cut (Bandiera et al 2000). On the other hand, the evidence suggests that interest rates and tax incentives have only weak effects at best on household saving (Edwards 1995), and that Ricardian equivalence typically does not hold, so that private saving offsets to public saving are small at best. The effects of aid on national saving are typically not clear-cut (Obstfeld 1999); in turn, terms of trade windfalls raise saving (Loayza, Schmidt-Hebbel and Serven 2000a, Agenor and Aizenman 2004).

As is clear from this overview, a substantial amount of ingenuity has been expended on attempting to explain international differences in saving rates, and the theory has been mined extensively to identify potential explanatory variables. However, despite the large and diverse set of explanatory factors that has been considered in these studies, the empirical literature concerned with international differences in saving rates has not even considered, much less provided evidence for, a possible role for the level of the real exchange rate in influencing national saving rates.[7]. On the whole, exchange rate policy has not been considered a potentially important determinant of national saving rates in th empirical literature.

---

[7] A few papers have looked for direct effects of real exchange rate misalignment on saving, without much success. One exception is Ghura and Greenes (1993), who conclude that overvaluation discouraged saving in Sub-Saharan Africa.

4.2 Saving in East Asia

As indicated above, countries in the East Asian region have recently stood out for their relatively high saving rates. Although there is a large body of literature that has identified high saving rates and relatively depreciated real exchange rates as two of the characteristics of the high-growth economies in East Asia (World Bank 1993), these two variables have not typically been linked to each other causally. This is so despite the fact that the interpretation of high saving rates in East Asia, like so much else associated with the growth experience of these countries, has been controversial. As with the broader literature on international differences in saving rates, explanations for high saving in East Asia have focused on other factors.

One clue to explaining high saving rates in East Asia is that saving rates in the region have not always been high. Private saving rates in East Asia trended up continuously during the region's high growth period from the mid-70s to the mid-1990s that preceded the Asian crisis, starting at about 15 percent and rising to about 30 percent of GDP. Though China's (national) saving rate has been high for some time, averaging 29 percent of GDP during 1970-77, before the country's change in economic policy launched it on its high-growth path, it also increased over time, reaching 41 percent of GDP in 1993-94, just before the exchange rate system was unified. There is little variation within the region in the pattern of saving rates over time. Overall, saving rates increased sharply in all countries except the Philippines during the two decades from the mid-1970s to the mid-1990s. Since public saving increased as well over the same period, the increase in national saving was even more dramatic than that in private saving alone.[8] Whereas national saving rates in Southeast Asian and Latin America were similar during the mid-70s, by 1995 the regional gap was about 20 percent. The implication of this change in saving rates over time is that high saving rates in the region do not simply reflect factors that could be captured by a regional fixed effect. Thus, although cultural factors are often cited to explain Asian saving rates, the increase in saving rates in the region over time does not support the view that high saving was driven by a cultural predisposition, but rather was induced by phenomena that varied over time. The interpretation of high saving rates in the region has indeed been linked to a variety of such factors.

An orthodox approach to explaining high saving rates in the region is typified by the World Bank's 1993 "miracle" study. This study cites an early demographic transition, macro stability, and policies toward the financial sector as the key policy contributions to the encouragement of higher private saving. An early demographic transition tended to increase the relative size of the high-saving working age cohort and decrease dependency ratios (see also Bloom and Williamson 1997). Macro stability, in turn, was judged to contribute to higher saving rates both directly and indirectly. The direct effect operates through higher public saving rates (see also Edwards 1995), while the key indirect contribution is through the achievement of low and stable inflation. Because high inflation tends

---

[8] The surge in saving in Thailand in the 1980s, for example, was largely driven by the public sector.

22

to be volatile inflation, it has the effect of making real interest rates both negative as well as unpredictable, tending to increase the risk associated with financial saving and thus discouraging such saving. Policies toward the financial sector included the achievement of financial sector stability, resisting severe financial repression, and making formal saving vehicles more convenient to small savers. The governments of high-performing Asian countries were credited with achieving the security of banks by protecting them from competition (mainly through restrictions on entry, except in Hong Kong, Indonesia, and Singapore), and safer banks encourage saving by reducing the risk associated with financial saving. In contrast with other literature of the time, the elimination of financial repression is not cited in the Bank study as a reason for higher saving rates, because repression was retained for most of the high growth period in the high-performing Asian economies. However, these countries practiced a relatively mild form of financial repression, which was consistent with positive real interest rates and financial deepening. In addition, a reduction in the degree of financial repression practiced in the region is cited as possibly having contributed to an increase in saving rates over time. Finally, policies to promote the accessibility of the formal financial sector during the high-growth period included bank branching to rural areas and the adoption of postal saving schemes. Postal saving systems in Japan, Korea, Malaysia, Singapore and Taiwan, for example, were designed to attract small savers (by offering greater security and lower transaction costs than the informal financial sector). However, much of this reasoning seems more relevant for savers' decisions regarding the allocation of their assets between the domestic financial system and other alternatives, than for the aggregate saving flow.

Other explanations focus on less orthodox factors. Some authors, for example, have emphasized the roles of restrictions on consumer credit and forced saving schemes in explaining high Asian saving rates. Restrictions on consumer credit in the context of financial repression, coupled with taxes on luxury goods, are perceived to have been important in encouraging household saving. The financial systems in countries such as Japan, Korea and Taiwan channeled resources into infrastructure and productive business uses, and away from housing and consumer credit which, according to some authors (see Patrick 1994), may have resulted in higher saving rates and more rapid growth, but at some welfare cost. In turn, the implementation of such provident funds in some countries (Japan, Singapore and Malaysia) may also have made a large contribution to saving rates.[9] In addition, other observers (Birdsall and Sabot 1998) have argued that the adoption of a "shared growth" strategy, in which growth occurred in the context of a relatively egalitarian income distribution, may have encouraged saving by moving the poor away from subsistence and thus allowing them to save.

All of these mechanisms, orthodox or otherwise, are consistent with high saving making a key contribution to growth by facilitating capital accumulation. A final set of interpretations, however, sees saving not as a cause of capital

---

[9] Whether such funds encourage total saving depends on the extent to which they force low savers to save more that they would have otherwise chosen to do. Faruqee and Husain (1998) conclude that provident funds did help raise aggregate saving in Singapore and Malaysia.

accumulation, but rather as its result. Several authors have argued, for example, that the saving transitions in East Asia reflect the effects of investment incentives on investment rates, and that higher investment rates resulted in higher saving rates, either directly or through higher growth rates. Singh (1995), for example, argues that high saving rates in the successful Asian economies did not reflect particularly high levels of household saving, but rather of corporate saving. These were the results of large corporate profits plus inducements and incentives to retain earnings. Large corporate profits were achieved by restrictions on domestic competition (in Japan and Korea) and import protection (there and elsewhere). Government provided fiscal and other incentives to retain these -- i.e., not to distribute them as dividends -- for example by taxing dividends and not retained earnings or capital gains.[10] Rodrik (2000) similarly emphasizes the driving role of investment incentives, seeing these as resulting in larger investment and growth rates in Korea, Singapore and Taiwan, with higher saving rates resulting from the increase in growth rates.

Careful empirical studies do not greatly restrict the range of possible determinants of Asian saving. The panel-data econometric evidence in Dayal-Gulati and Thimann (1997), for example, attributes high saving rates among the high-growth Asian economies to macro stability, public saving, mandatory saving schemes, and financial deepening (presumably reflecting positive real interest rates and secure banks). By contrast, in studying China's household saving behavior, Kraay (2000) emphasizes the role of an entirely different set of variables. He finds that household saving was affected by above-subsistence income levels and expectations of future growth, but demographic factors and income uncertainty do not help explain household saving in China. Other authors have suggested instead that China's high saving rate is largely a result of the low coverage of its pension and health systems, together with restrictions on family size.

In short, there is no consensus at present on the extent to which increases in saving rates in East Asia, and their mobilization through the financial system, were due to growth, demographics, reduced inflation, higher real interest rates, the spread of financial institution offices, or other variables. What seems to be true, however, is that students of this issue have not considered exchange rate policy among the myriad of possible explanations for high saving rates in Asia.

5 Summary and Conclusions

The bottom line is that the real exchange rate-to-saving link -- or what we have labeled the capital accumulation channel -- is conceptually and empirically weak. While there are clearly examples of countries in which high saving rates, depreciated real exchange rates, and high growth have been observed at the same time, as well as some in which the reverse of each of these outcomes has also materialized simultaneously, there is no clear-cut association between more depreciated

---

[10] As in the discussion of Levy-Yeyati and Sturzenegger (2007b), this argument would need to establish that households did not pierce the corporate veil and reduce their own saving as corporate saving increased. Otherwise, private saving would be unaffected by whether saving is done by corporations or by households. Singh does not address this issue, leaving the matter of high private saving still to be explained.

real exchange rates and higher saving rates in the international experience, even after controlling for levels of GDP per capita. Moreover, even in those countries where high saving rates and relatively depreciated real exchange rates have been observed simultaneously, Bernanke's (2005) view that causation runs from high saving rates to a more depreciated real exchange rate is a plausible explanation.

Perhaps a more serious difficulty is that there is as yet no obviously plausible analytical basis for the capital accumulation channel itself. While a well-known model by Dooley, Folkerts-Landau and Garber (2004) generates a link from a depreciated real exchange rate to a higher saving rate, that model predicts that the higher saving rate will result in larger current account surpluses rather than more domestic capital accumulation. Moreover, while it is possible to specify standard modern open-economy models that generate the desired causation from a more depreciated real exchange rate to a higher household saving rate, the requirement for this outcome to emerge in the context of such models -- i.e., that the exchange rate policy results in a front-loaded productivity boost -- appears to be quite stringent. We have described one mechanism that would indeed generate this result, operating through inflation-induced transactions costs, but the empirical relevance of this mechanism remains to be established, and in any case the result emerges only when the exchange rate policy is transitory.

Finally, the literature on cross-country differences in saving rates has not identified exchange rate policy in general, or a depreciated real exchange rate in particular, as an important factor in explaining such differences. Even in the high-saving, high-growth countries that have provided the main impetus for professional interest in the capital-accumulation channel, explanations for high saving rates have tended to focus on factors such as demographics, financial-sector policies, mandated saving schemes, and fiscal policies, rather than exchange rate policies.

Though we have not focused on it in this paper, the second link in the capital accumulation channel, from higher saving rates to higher growth rates through higher investment rates, appears to be at least as problematic as the first. There are two major difficulties with this link. The first one is empirical: as indicated above, theory is consistent with reverse causation from growth to saving, so a positive empirical correlation between saving and growth rates is subject to more than one interpretation (Deaton 1999). The possibility of bidirectional causality between saving and growth has led several authors to perform various tests of Granger causality with international data, often with the result that growth typically precedes saving, rather than vice versa.[11] Although Granger tests can hardly be

---

[11] The classic reference is Carroll and Weil (1994). Baharumshah, Thanoon and Rashid (2003) reach similar conclusions using data for East Asian countries. Attanasio, Picci and Scorcu (2000) find that the patterns of Granger causality between saving and growth in a large cross-country sample are not robust to minor changes in specification and econometric technique. Rodrik (2000) found in an international sample of saving transitions that increases in saving rates tended to be followed by only temporary increases in growth rates, while increases in growth rates tended to be associated with permanent increases in saving rates, consistent with the East Asian experience summarized in the last section. He interpreted this as suggesting that saving is insignificant as a causal factor for long-run growth, and tends to be an outcome -- rather than a determinant -- of high growth rates. More recently, Mohan (2006) derived a similar result for a

viewed as conclusive, the upshot is that time series studies provide little evidence that the positive cross-country correlation between saving and growth rates reflects causation from the former to the latter. The second problem is analytical: the link between saving and domestic investment may be quite loose in financially open economies, in which higher saving rates may simply be associated with a larger surplus on the current account of the balance of payments. Though Feldstein and Horioka (1983) documented a strong correlation between saving and domestic investment even in economies that appear to be highly open financially, this correlation has weakened over time as capital mobility has increased, and there is a now a host of potential explanations for the saving-investment correlation that do not rely on causation running from saving rates to investment rates (see Coakley, Fulasi, and Smith 1998).

In short, there is as yet little analytical or empirical support for the view that an exchange rate policy geared to the maintenance of a depreciated real exchange rate promotes increased domestic saving, and through more saving, larger rates of domestic capital accumulation and growth. If exchange rate policy indeed has a significant impact on economic growth in developing countries, it is likely to do so through what this paper has labeled the TFP channel, rather than through the capital accumulation channel.

6 References

Agenor, Richard and Joshua Aizenman (2004): "Savings and the terms of trade under borrowing constraints", *Journal of International Economics* 63, 32 1-340.

Attanasio, Orazio, Lucio Picci and Antonello Scorcu (2000): "Saving, growth and investment: a macroeconomic analysis using a panel of countries", *Review of Economics and Statistics* 82, 182-211.

Baharumshah, Ahmad, Marwan Thanoon and Salim Rashid (2003): "Saving dynamics in the Asian countries*", Journal of Asian Economics* 13, 827-845.

Bandiera, Oriana, Gerard Caprio, Patrick Honohan and Fabio Schiantarelli (2000): "Does financial reform raise or reduce saving?", *Review of Economics and Statistics* 82.

Bernanke, Ben S. (2005), "The global saving glut and the US current account deficit," remarks at the Sandridge Lecture, Virginia Association of Economics, Richmond, Virginia (March).

Birdsall, Nancy, Thomas Pinckney and Richard Sabot (1998): "Why low inequality spurs growth: savings and investment by the poor", in A. Solimano (ed.), *Social Inequality: Values, Growth, and the State*. Michigan University Press.

Coakley, Jerry, Farida Fulasi, and Ron Smith (1998), "The Feldstein-Horioka puzzle and capital mobility: a review," *International Journal of Finance and Business* 3, pp. 169-188.

Deaton, Angus (1999), "Saving and Growth", in Klaus Schmidt-Hebbel and Luis Serven (eds.), *The Economics of Saving and Growth: Theory, Evidence*

large sample of countries at various income levels.

and Implications for Policy, Cambridge University Press.

De Gregorio, Jose, Alberto Giovannini and Holger Wolf (1994): "International evidence on tradables and nontradables inflation", *European Economic Review* 38, 1225-44.

Dhayal-Gulati, Anuradha, and Christian Thimman (1997): "Private saving in Southeast Asia and Latin America compared: searching for policy lessons", in J. Hicklin et al (eds.), *Macroeconomic Issues Facing ASEAN Countries*. International Monetary Fund.

Edwards, Sebastian (1995), "Why are saving rates so different across countries? An international comparative analysis," NBER Working Paper 5097.

Eichengreen Barry (2007): "The real exchange rate and economic growth", unpublished manuscript.

Faruqee, Hamid and Aasim Husain (1998): "Saving trends in Southeast Asia: a cross-country analysis", *Asian Economics Journal* 12, 195-217.

Frankel, J. (2005): "On the renminbi: the choice between adjustment under a fixed exchange rate and adjustment under a flexible rate," NBER Working Paper 11274.

Ghura, Dhaneshwar, and Thomas Grennes (1993): "The real exchange rate and macroeconomic performance in Sub-Saharan Africa", *Journal of Development Economics* 42, 155-174.

Hausmann, Ricardo and Dani Rodrik (2003), "Economic development as self-discovery," John F. Kennedy School, Harvard University, mimeo.

Hausmann, Ricardo, Lant Pritchett, and Dani Rodrik (2004), "Growth accelerations," mimeo, John F. Kennedy School, Harvard University.

Jappelli, Tullio and Marco Pagano (1994), "Saving, growth, and liquidity constraints," *Quarterly Journal of Economics*, Vol. 109, pp. 83-109.

Johnson, S., J. Ostry and A. Subramaniam (2007): "The prospects for sustained growth in Africa," IMF Working Paper 07/52.

Kraay, Aart (2000), "Household saving in China," *Journal of Development Economics*, Vol. 14, No. 3, pp.

Kouri, Pentti (1978): "Profitability and growth in a small open economy", Cowles Foundation Discussion Paper 490.

Lee, J. and M. Tang (2007): "Does productivity growth appreciate the real exchange rate?," *Review of International Economics* 15, 164-187.

Levy-Yeyati, Eduardo, and Federico Sturzenegger (2007a): "The effect of monetary and exchange rate policies on development", forthcoming in *Handbook of Development Economics.*

Levy-Yeyati, Eduardo, and Federico Sturzenegger (2007b). "Fear of floating in reverse: exchange rate Policy in the 2000s", paper presented at the LAMES-LACEA Annual Meetings.

Loayza, Norman, Klaus Schmidt-Hebbel and Luis Serven (2000a): "What drives private saving across the world?", *Review of Economics and Statistics*

Loayza, Norman, Klaus Schmidt-Hebbel and Luis Serven (2000b), "Saving in developing countries: an overview," *World Bank Economic Review*, Vol. 14, No. 3, pp. 393-414.

Montiel, Peter (2007), "The real exchange rate and the saving rate in an intertemporal model of the current account," Williams College, mimeo.

Obstfeld, Maurice (1999): "Foreign resource inflows, saving and growth", in Klaus Schmidt-Hebbel and Luis Serven (eds.), *The Economics of Saving and Growth: Theory, Evidence and Implications for Policy*, Cambridge University Press.

Prasad, Eswar, Raghuram G. Rajan and Arvind Subramanian (2006). "Patterns of international capital flows and their implications for economic development," Proceedings, Federal Reserve Bank of Kansas City, pages 119-158.

Razin, Assaf and Susan Collins (1999), "Real exchange rate misalignments and growth", in A. Razin and E. Sadka (eds.), *The Economics of Globalization: Policy Perspectives from Public Economics*. Cambridge University Press.

Rodrik, Dani (2000), "Saving Transitions," *World Bank Economic Review*, Vol. 14, No.1, pp. 481-507.

Rodrik, Dani (2006), "Industrial development: stylized facts and policies," John F. Kennedy School of Government, mimeo (August).

Rodrik, Dani (2007), "The Real Exchange Rate and Economic Growth: Theory and Evidence," John F. Kennedy School of Government, mimeo.

Schmitt-Grohe, Stephanie and Martin Uribe (2003), "Closing small open economy models," *Journal of International Economics* 61, pp. 163-185,

Schmidt-Hebbel, Klaus, and Luis Serven (1999): " Saving in the world: the stylized facts", in Klaus Schmidt-Hebbel and Luis Serven (eds.), *The Economics of Saving and Growth: Theory, Evidence and Implications for Policy*, Cambridge University Press.

Schmidt-Hebbel, Klaus, Luis Serven, and Andres Solimano (1996), "Saving and Investment: Paradigms, Puzzles, Policies," *World Bank Research Observer*, Vol. 11, No. 1, 87-117.

World Bank (1993), *The East Asian Miracle: Economic Growth and Public Policy* (New York: Oxford University Press).

World Bank (1994), *Averting the Old Age crisis: Policies to Protect the Old and Promote Growth* (New York: Oxford University Press).