Capital flows and exchange rates

By Andrew Bailey of the Bank's International Economic Analysis Division, and Stephen Millard and Simon Wells of the Bank's Monetary Instruments and Markets Division.

This article focuses on the possible role of capital flows in explaining exchange rate movements. Some commentators have suggested that a substantial increase in capital flows into the United States could have accounted for the recent appreciation of the US dollar. This could imply that capital inflows have increased in response to a rise in the rate of return on capital, which in turn has reflected the structural increase in US productivity seen in recent years. We find evidence to suggest that this may explain part of the recent dollar appreciation, but unsurprisingly it does not provide a full explanation.

Introduction

Much attention has been paid recently to the possible links between flows of capital into the United States, US equity prices and the US dollar.(1) It is argued that a structural improvement in US productivity(2) increased the rate of return on capital, which led to a substantial increase in capital flows into the United States. These flows, in turn, caused the dollar to appreciate in both real and nominal terms. This possible explanation for the dollar's appreciation has been quite widely discussed in market commentaries,(3) in the May 2001 issue of the IMF's *World Economic Outlook*, and in the December 2000 issue of the Bank of England's *Financial Stability Review*.

The aim of this article is to show how, focusing on the recent performance of the US dollar, it might be possible to tie together movements in the real and nominal exchange rates, the real interest rate, equity prices and the current/capital account balance. These relationships are complex, and their form at any time depends on the shock hitting the economy. Here we examine the effects of a permanent shock to the level of future productivity.(4) We also show how the effects of a productivity shock that is concentrated in a 'tradable sector' differ from the effects of a shock that affects the whole economy.(5) Finally, we examine the US evidence to see how well it fits this theoretical analysis.

The effects of a productivity shock on capital flows ...

A productivity shock that raises expected future output in the home country will tend to lead to capital inflows for two reasons.

First, if consumers in the home country expect to be richer in the future, they would want to borrow from abroad—ie draw an inflow of capital—to increase their consumption today. (We assume that consumers are sufficiently forward-looking to wish to smooth their consumption over the present and future time periods.) In other words, they would increase their current consumption to reflect the expectation that their permanent income has increased. The annex on pages 316–18 sets out a simple one-good model that demonstrates how this may work.

(1) By 'capital flows' we mean, specifically, flows of funds from savers to investors, either directly or through intermediaries. (Note that this definition is much broader than simply flows of financial assets.)
(2) See Berry, S and England, D (2001), 'Has there been a structural improvement in US productivity?', *Bank of England Quarterly Bulletin*, Summer, pages 203–09 (henceforth Berry and England (2001)), which presents evidence of such a structural change.
(3) Some examples of recent brokers' reports that discuss this are Lehman Brothers, *Global foreign exchange strategies* (1/3/2001), Morgan Stanley Dean Witter, *Currency strategy and economics* (22/11/2000), Goldman Sachs, *European weekly analyst* (17/11/2000), and Deutsche Bank, *Global FX outlook and strategy* (1/12/2000).
(4) Such a shock would have a temporary effect on productivity growth while the economy was adjusting to its new higher level of productivity.
(5) We can think of 'non-tradable' goods and services as those whose transportation costs are so large relative to their costs of production that it would never be profitable to produce them in an economy and then sell them abroad. Goods and services that can profitably be sold abroad are called 'tradable'.

Second, the expected increase in future productivity would raise expected future profits. This, in turn, would lead to an increase in equity prices and, other things equal, would encourage investment. Residents of the home country would want to take advantage of current investment opportunities that enhance future output (ie stimulate an increase in productivity) but without forgoing current consumption. So the increase in investment demand that is not financed by current domestic savings would be financed by inflows of capital. And inflows of foreign direct investment and foreign equity investment are particularly likely to increase as overseas investors also take advantage of the higher rates of return to capital in the home country.

... and the real exchange rate(1)

The simple model in the annex shows that a productivity shock can generate large capital inflows even when the real exchange rate is constant. However, the real world is clearly more complicated than assumed in the model. In particular, we need to extend our analysis to a world in which the real exchange rate can vary.

A useful starting-point is the idea of uncovered interest parity (UIP), which suggests that the expected change in the real exchange rate over any period should equal the difference between the domestic real interest rate and that of the rest of the world. Within this framework, a shock to productivity that raised the future level of productivity, and hence temporarily raised productivity growth, would lead to an increase in the domestic real interest rate relative to the world real interest rate; that, in turn, would prompt a jump appreciation of the real exchange rate. As productivity growth returned to trend, bringing the domestic real interest rate back into line with the world real interest rate, the real exchange rate would depreciate back to its equilibrium value. One caveat to this is that within this model a more or less continual series of productivity shocks would lead to a longer-run appreciation of the real exchange rate.

But what happens to the equilibrium real exchange rate? This depends importantly on whether the productivity shock is concentrated in the tradable or the non-tradable sector. As is shown in the annex, a productivity shock that affects both sectors equally is likely to lead to a depreciation of the equilibrium real exchange rate. This happens because such a shock implies an increase in the relative supply of domestic goods and services; given this, their relative price must fall.(2)

On the other hand, a productivity shock concentrated in the tradable sector is likely to lead to an appreciation of the equilibrium real exchange rate. This is commonly known as the Balassa-Samuelson effect and occurs because product market arbitrage between countries equilibrates prices for the tradable goods and services at the same time as labour market arbitrage within economies means that wages are equalised at the margin between the tradable and non-tradable sectors. If one country has an increase in the productivity of its tradable sector, other things equal, real wages will increase in both the tradable and non-tradable sectors. Because there has been no productivity change in the non-tradable sector this leads to a rise in the price of non-tradable goods and services relative to tradables in the home economy, and an appreciation of the real exchange rate. By extension, if the economy undergoes a series of productivity shocks (ie residents keep being surprised by the level of productivity growth), the Balassa-Samuelson effect would indicate an ongoing appreciation of the real exchange rate in the face of such productivity shocks.

We use the simple two-good, two-period, two-country model in the annex to illustrate this effect. We find that a productivity shock that raises expected future output in both the tradable and non-tradable sectors by 5% leads to an initial appreciation of the real exchange rate of 2.16%. It then depreciates to an equilibrium level 1.84% below its starting value. On the other hand, a productivity shock that raises the expected future output of the tradable sector by 5% while having no effect on the non-tradable sector leads to an immediate appreciation of 2.16% with no further change in the real exchange rate. Chart 1 plots these stylised responses.

Although the numbers should be seen as purely illustrative, the key point is that the persistence of the real exchange rate appreciation depends crucially on the effect of the shock on the productivity of the tradable sector relative to the non-tradable sector.

(1) Throughout this article we define the real exchange rate to be the price of a representative consumption bundle in the home country relative to that of a representative consumption bundle in the rest of the world.
(2) We are assuming that the relative price of tradable goods and services in different countries cannot change; it is the fall in the relative price of non-tradable goods and services that drives the depreciation of the real exchange rate.

Chart 1
Effect of a productivity shock on the real exchange rate

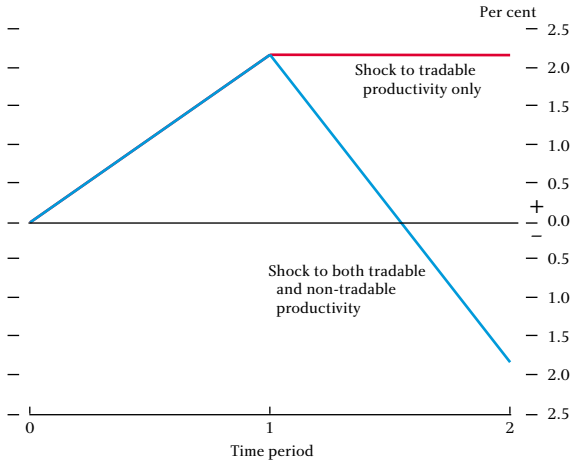

One strong caveat to this analysis is that we have assumed that the relative price of tradables cannot change. In practice, the real exchange rate will depend not only on the relative price of tradable versus non-tradable goods and services in the domestic economy, but will also depend on the terms of trade between domestic and foreign tradables. To the extent that a productivity shock in the tradable sector only affects domestic producers of tradables, the appreciation of the exchange rate that arises from the Balassa-Samuelson effect will be muted. This will happen because such a shock implies an increase in the relative supply of domestic tradables; given this, their relative price (the terms of trade) must fall.

What about the nominal exchange rate?

So far we have only considered the real exchange rate. To examine what will happen to the nominal exchange rate, we need to make an assumption about monetary policy in the home country and the rest of the world. We assume that the rest of the world has the same inflation target as the home monetary authority and that monetary authorities in the rest of the world do not react to the domestic productivity shock.

Suppose that the home economy experiences a shock that raises output relative to trend. If the monetary authority accommodated the shock, it would lower interest rates, raising nominal demand by the same amount as the rise in output to leave inflation on target. In this case, the nominal exchange rate would move one-for-one with the real exchange rate. Alternatively, suppose that the monetary authority did not accommodate the shock. In this case, inflation would fall below target and the nominal exchange rate would rise relative to the real exchange rate.

How well does this tie in with the US story?

Notwithstanding recent revisions, the United States saw a quite sharp pick-up in labour productivity growth from around 1995. This was associated with rapid technological advances in the information and communications technology (ICT) sector and a pick-up in total factor productivity growth in the United States, as ICT became more widely used. Berry and England suggest that at least some of this pick-up represented a structural improvement in US productivity. In this section, we examine how capital flows and the real and nominal exchange rates have responded to this productivity shock in light of the theory presented above. Of course, we need to bear in mind that if the increase in productivity proves to be a purely cyclical phenomenon then all the effects described are likely to be reversed over the economic cycle.

Chart 2 shows the rapid widening of the US current account deficit over the past ten years. For any current account position there must be equal and offsetting capital flows, and so the widening US current account deficit was, by definition, financed by large inflows of capital from abroad. Put another way, and as described earlier, the productivity shock in the United States has appeared to lead to large capital inflows as US residents have borrowed against expected future income.

Chart 2
US current account balance

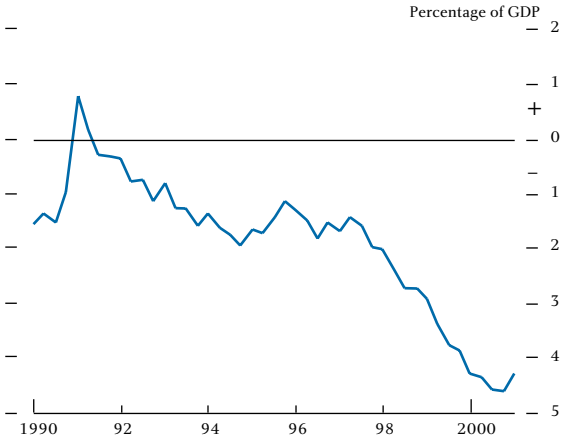

Chart 3 shows net capital inflows to the United States, illustrating the strong increase in both equity and foreign direct investment (FDI) flows since 1995. The IMF has suggested that these net equity capital inflows can help to explain recent movements in the euro-dollar

Chart 3
Net capital inflows to the United States

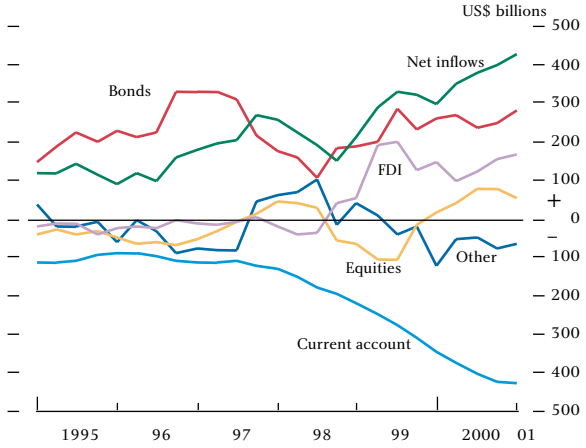

exchange rate, although long-term interest rate differentials also have an explanatory role.(1)

Chart 4 shows movements in the real effective dollar exchange rate (deflated using consumer price indices) over the past ten years. Over the past five years, the real effective exchange rate has appreciated by 31%. This appreciation is large and our model suggests that, to explain the appreciation as a result of the productivity shock alone, the shock would have to be extremely large. However, it is possible that the increase in US productivity is at least partly responsible for the dollar's appreciation in real terms. In addition, for the appreciation to have persisted, economic agents would have had to take time to realise that some elements of the productivity improvements appeared more likely to represent a structural shock to productivity rather than being cyclical. This seems a reasonable characterisation of a consensus view of US developments, ie that part of the productivity gains has over time come to be regarded as more structural in nature.

Chart 4
US real effective exchange rate

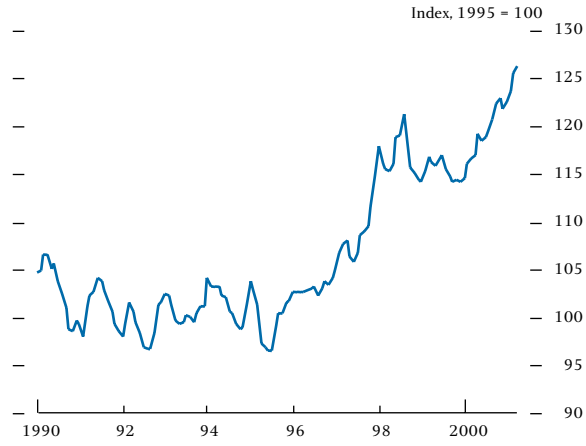

The evidence described in Berry and England (2001) suggests that the improvements in productivity have been most marked in the ICT sector. And imports of ICT goods (computers, accessories and telecommunications equipment) have accounted for more than 40% of total capital expenditure on ICT in the United States since 1995, suggesting that ICT goods appear to be fairly characterised as tradable. Further evidence to support the Balassa-Samuelson argument is provided by Chart 5, which illustrates that the relative price of ICT goods has fallen dramatically at a time when the real exchange rate for the dollar has risen.

Chart 5
Dollar real effective exchange rate (REER) and relative price of ICT goods

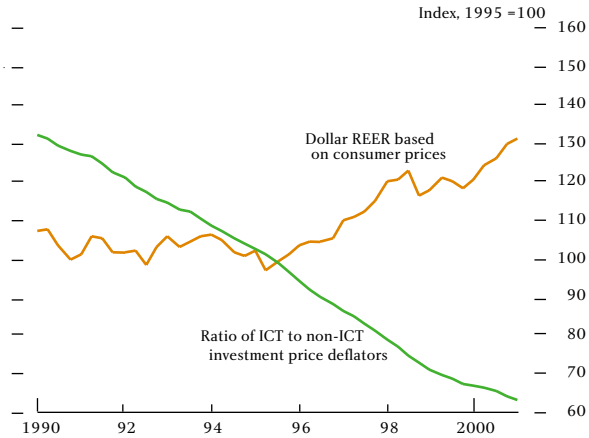

Source: Thomson Financial Datastream.

Chart 5 suggests that since 1995 the fall in the relative price of ICT goods is of the same order as the rise in the real exchange rate. This may be seen as evidence for the productivity shock in the ICT sector being able to explain the appreciation of the real exchange rate. However, the fall in the price of goods relative to the price of services—a broader proxy for the relative price of tradable versus non-tradable goods—is much more muted. We show this in Chart 6 by setting the services/goods CPI ratio in the United States alongside the rise in the real exchange rate.

Turning to other possible explanations, Chart 7 shows, furthermore, that movements in the terms of trade cannot account for the appreciation of the US real effective exchange rate. Indeed, since the end of 1998, the terms of trade have fallen, which by itself would imply a real depreciation.

One possible explanation for the strength of the dollar over this period, which does not rely on the productivity

(1) See, for example, chapter 2 of the IMF's May 2001 *World Economic Outlook*.

Chart 6
Dollar real effective exchange rate (REER) and price of services relative to goods

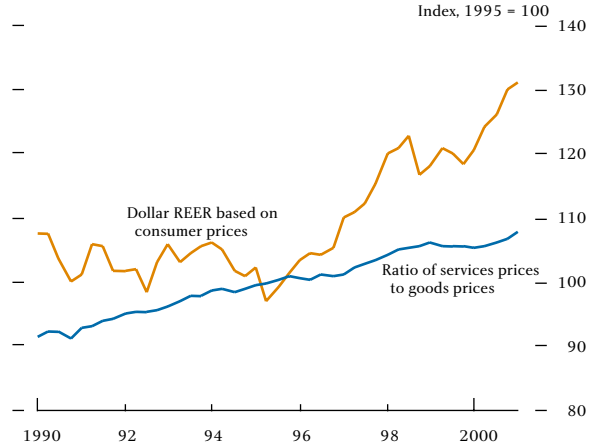

Chart 7
Dollar real effective exchange rate (REER) and terms of trade

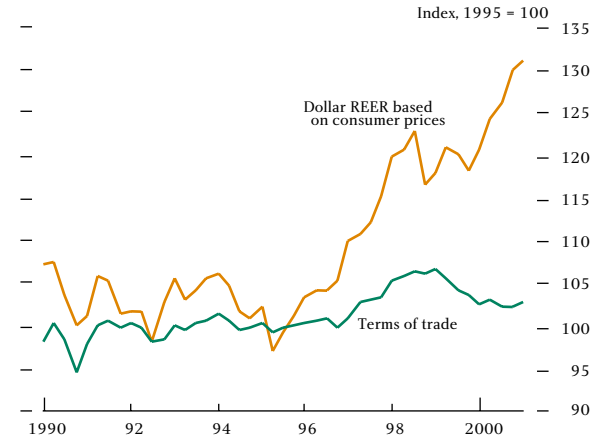

shock, is that it is a purely cyclical phenomenon. A positive shock to domestic demand, such as a large increase in investment or fall in saving, would tend to raise the real interest rate relative to the rest of the world. As a result of the UIP reasoning discussed above, this would lead to a real exchange rate appreciation. This appreciation would be reversed as demand fell back to more sustainable levels.

Finally, Chart 8 shows the nominal exchange rate index for the dollar. Since the start of 1999, the dollar has appreciated by slightly more in nominal terms (37%) than in real terms (31%). This suggests that policy has, by and large, accommodated the real exchange rate increase. (If policy had been too tight—ie interest rates were kept higher than was necessary to keep inflation on track—the nominal exchange rate would have appreciated by considerably more than the real exchange rate, as happened in the mid-1980s.) Further evidence of this is provided by the fact that US inflation has tended to be higher than inflation in the United Kingdom and the euro area since about 1995. If policy had not been accommodating, US inflation would have fallen relative to inflation in the rest of the world.

Chart 8
Dollar nominal effective exchange rate

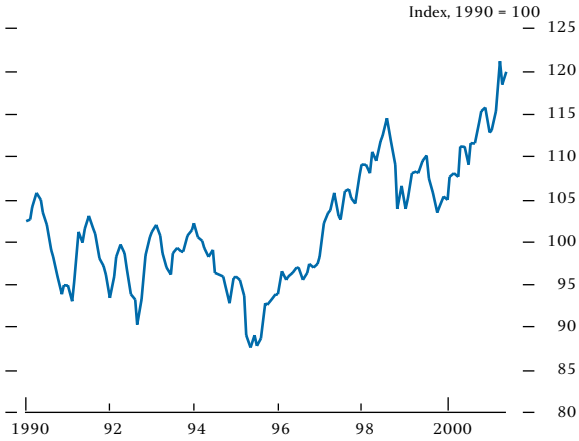

Conclusions

In the context of assessing how economic models can cast light on the recent appreciation of the dollar, we have sought to tie together movements in real and nominal exchange rates, the real interest rate, equity prices and the current/capital account balance. These relationships are likely to be sensitive to the particular shocks assumed to be affecting an economy. We focus on the effects of a permanent shock to the level of future productivity (a temporary shock to the growth rate of productivity) in the United States, and how these effects might be sensitive to whether the shock is concentrated in tradable goods production, particularly the ICT sector.

It is relatively easy to show that such a shock would lead to larger capital inflows, higher expected future profits, an increase in equity prices and a higher level of investment. Moreover, this is widely accepted as representative of developments over the past five years or so in the United States. It is more complicated to make the link to a sustained appreciation of the exchange rate. A conventional UIP argument would suggest that a productivity shock will lead to an immediate appreciation of the real exchange rate, followed by a depreciation back to its equilibrium level as productivity growth returns to its trend rate. If we assume that it takes time for economic agents to realise that the increase in productivity is structural rather than purely cyclical, we can extend the period

over which the real exchange rate appreciates before it eventually depreciates. If we assume that the productivity shock is spread evenly across the tradable and non-tradable sectors, then the equilibrium real exchange rate will—other things equal—depreciate, because such a shock implies an increase in the relative supply of domestic goods and a fall in their relative price.

But this story is sensitive to the assumption that the productivity shock is spread across tradable and non-tradable sectors. Other work suggests that the US productivity shock of recent years has reflected a strong ICT contribution, and that ICT is highly traded. Using a simple model we show that a productivity shock concentrated in tradable goods can lead to an appreciation of the equilibrium real exchange rate. But if we relax the assumption that the relative price of tradables in different countries cannot change, the appreciation will be more muted.

A key question is how far such a model helps to cast light on the appreciation of the dollar. Unsurprisingly, the evidence suggests that the model will not provide a full explanation of the recent appreciation of the dollar, but along the way there is an interesting point to the exercise. Put simply if, to use an extreme assumption, we suppose that all tradable goods in the United States are ICT, then we can explain the large appreciation of the real dollar exchange rate because of the large relative fall in ICT prices. However, if we use relative goods/services prices to proxy the ratio of tradable goods prices to non-tradable goods prices, then it becomes harder to explain within this framework the scale of the appreciation of the dollar. Moreover, the fall in the US terms of trade since the end of 1998 is consistent with a rise in productivity in the US tradable sector relative to that in the rest of the world. That fall in the relative price of US tradables might imply a more muted appreciation of the real exchange rate as a result of a productivity shock.

Capital flows: technical annex

This annex develops a model to describe more formally the effect on capital flows and the real exchange rate of an anticipated future productivity shock. In the simplest version of the model, there is only one good.(1) This version of the model is used to show that a productivity shock will generate capital flows; this happens despite the fact that the real exchange rate never moves (purchasing power parity holds at all times). In the more developed version of the model, each economy is endowed with two goods: one that is tradable with the other country and one that is not tradable. In this case, we find that a shock that raises the productivity of both sectors in one country will always lead to a long-run real exchange rate depreciation. If, on the other hand, the shock affects only the tradable sector the real exchange rate will appreciate in both the short and the long run. Provided the productivity shock has some effect on the tradable sector, we will still observe capital flows towards the country experiencing the shock.

The one-good case

In this model, the representative domestic consumer maximises his utility function (we assume a simple log utility function) subject to a two-period budget constraint in which he is allowed to borrow from/lend to foreign consumers at the world real interest rate (determined endogenously). Mathematically, we can write this problem as:

$$\text{Maximise } 1n\left(c_1^D\right)+\beta \mathrm{E}_1 1n\left(c_2^D\right) \tag{1}$$

$$\text{subject to } A = y_1^D - c_1^D \tag{2}$$

$$\text{and } \mathrm{E}_1 c_2^D = (1+r)A + \mathrm{E}_1 y_2^D \tag{3}$$

where c^D is domestic consumption, y^D is domestic output (exogenously given),(2) β is the consumer's discount factor, E is the expectations operator and A is net external assets (loans to foreigners) held at the end of period 1. This will equal capital inflows in period 1 and will result in capital outflows of $A(1+r)$ in period 2 as the loans are repaid with interest.

The first-order conditions for this problem imply:

$$\frac{1}{c_1^D} = \beta(1+r)\mathrm{E}_1\left(\frac{1}{c_2^D}\right) \tag{4}$$

$$c_1^D + \mathrm{E}\frac{c_2^D}{1+r} = y_1^D + \mathrm{E}_1\frac{y_2^D}{1+r} \tag{5}$$

Similarly, the first-order conditions for the (symmetric) foreigner's problem are:

$$\frac{1}{c_1^F} = \beta(1+r)\mathrm{E}_1\left(\frac{1}{c_2^F}\right) \tag{6}$$

$$c_1^F + \mathrm{E}_1\frac{c_2^F}{1+r} = y_1^F + \mathrm{E}_1\frac{y_2^F}{1+r} \tag{7}$$

And, for the world as a whole, we have the two budget constraints:

$$c_1^D + c_1^F = y_1^D + y_1^F \tag{8}$$

$$c_2^D + c_2^F = y_2^D + y_2^F \tag{9}$$

For ease of exposition, we assume perfect foresight. Assume that home consumers see a productivity shock that will raise output in period 2. In particular, we suppose that output is 5% higher in period 2 than originally expected in period 1. If we suppose that the two countries are of equal size, the effect of the shock will be to raise world output in period 2 by 2.5%. In this model, the growth rate of world output equals the growth rate of world consumption. Since both countries face the same real interest rate, domestic consumption growth and foreign consumption growth will be the same. Given 2.5% higher world consumption growth, this implies domestic and foreign consumption growth 2.5% higher than previously expected. From the Euler equations we can also see that the world real interest rate rises by $0.025/\beta$:

$$\frac{c_2^D}{c_1^D} = \frac{c_2^F}{c_1^F} = \beta(1+r+0.025/\beta) = (1+g+0.025) \tag{10}$$

where g is the growth rate of output originally expected in each country.

With assumptions about the initial world growth rate and real interest rate, we can use this simplistic model to measure the capital inflows to the domestic economy, in the face of the productivity shock. Suppose the quarterly growth rate is 0.6% and the quarterly real interest rate is 1% (equivalent to annual rates of 2.4% and 4% respectively). Given these values, the first-order

(1) For ease of exposition the term 'goods' in this annex encompasses goods and services.
(2) We abstract from production in this model; hence there is no capital and no investment.

conditions imply an annual value for β of 0.985. We can rearrange equation **(5)** to show that when domestic consumers anticipate a 5% increase in period 2 output, they raise period 1 consumption by 1.2%. Substituting this into equation **(2)** shows that this is funded by capital inflows equivalent to 1.2% of period 1 domestic output.

The desire of domestic consumers to smooth consumption by borrowing in period 1 drives up the real interest rate to 6.5% per annum. So in this simple (constant elasticity of intertemporal substitution) case, a large anticipated rise in domestic period 2 output induces a modest rise in consumption in period 1. This is because a large rise in the world real interest rate is required to tempt foreigners into forgoing consumption in period 1.

A two-good model

In this section, we look at what happens when we consider an economy that produces tradable and non-tradable goods. We can think of non-tradable goods as goods and services whose transportation costs are so large relative to their costs of production that it would never be profitable to produce them in an economy and then sell them abroad. The inclusion of such goods enables us to define a real exchange rate that is not constant (as was the case in the one-good world).

Consider a domestic consumer. His problem is again to maximise utility subject to the budget constraints he faces in each period. He is allowed to borrow/lend tradable goods from/to foreign consumers. He cannot borrow or lend non-tradable goods. Mathematically, we can write his problem as:

$$\text{Maximise } \gamma \ln\left(c_{T,1}^{D}\right)+(1-\gamma)\ln\left(c_{N,1}^{D}\right) + \beta \mathrm{E}_1\left(\gamma \ln\left(c_{T,2}^{D}\right)+(1-\gamma)\ln\left(c_{N,2}^{D}\right)\right) \quad (11)$$

$$\text{subject to } A = y_{T,1}^{D} + p_1^{D} y_{N,1}^{D} - c_{T,1}^{D} - p_1^{D} c_{N,1}^{D} \quad (12)$$

$$\text{and } \mathrm{E}_1\left(c_{T,2}^{D} + p_2^{D} c_{N,2}^{D}\right) = (1+r)A + \mathrm{E}_1\left(y_{T,2}^{D} + p_2^{D} y_{N,2}^{D}\right) \quad (13)$$

where c_T^D is domestic consumption of tradable goods, c_N^D is domestic consumption of non-tradable goods, y_T^D is domestic output of tradable goods, y_N^D is domestic output of non-tradable goods and p is the relative price of non-tradable goods expressed in terms of tradable goods. The particular utility function assumed implies that the share of consumer spending that goes on tradable goods is constant and equal to γ.

The first-order conditions for this problem imply:

$$\frac{1}{c_{T,1}^{D}} = \beta(1+r)\mathrm{E}_t\left(\frac{1}{c_{T,2}^{D}}\right) \quad (14)$$

$$\frac{\gamma}{1-\gamma} = \frac{c_{T,1}^{D}}{p_1^{D} c_{N,1}^{D}} = \frac{c_{T,2}^{D}}{p_2^{D} c_{N,2}^{D}} \quad (15)$$

In equilibrium, non-tradable output will equal non-tradable consumption. Using this fact and combining together the two budget constraints (equations **(12)** and **(13)**) gives:

$$c_{T,1}^{D} + \mathrm{E}\frac{c_{T,2}^{D}}{1+r} = y_{T,1}^{D} + \mathrm{E}_1\frac{y_{T,2}^{D}}{1+r} \quad (16)$$

We can use the consumer's utility function to define a consumption-based aggregate price index, P. The index is defined as the P that minimises total consumption expenditure, pc_N+c_T, subject to the utility function defined above. In this case, we can show that, if we use tradable goods as the numeraire, it will be given by:

$$P = \frac{p^{1-\gamma}}{\gamma^{\gamma}(1-\gamma)^{1-\gamma}} \quad (17)$$

Furthermore, we can define the real exchange rate between domestic and foreign consumption bundles as simply the ratio of the two aggregate price indices:

$$e = \frac{P^{D}}{P^{F}} = \left(\frac{p^{D}}{p^{F}}\right)^{1-\gamma} \quad (18)$$

where a rise in e signals a real exchange rate appreciation.

The first-order conditions for the (symmetric) foreigner's problem will include:

$$\frac{\gamma}{1-\gamma} = \frac{c_{T,1}^{F}}{p_1^{F} c_{N,1}^{F}} = \frac{c_{T,2}^{F}}{p_2^{F} c_{N,2}^{F}} \quad (19)$$

Log-linearising equations **(15)**, **(18)** and **(19)** and noting that, in equilibrium, non-tradable consumption equals non-tradable output gives:

$$0 = \hat{c}_T^{D} - \hat{y}_N^{D} - \hat{p}^{D} = \hat{c}_T^{F} - \hat{y}_N^{F} - \hat{p}^{F} \quad (20)$$

$$\hat{e} = (1-\gamma)\left(\hat{p}^{D} - \hat{p}^{F}\right) \quad (21)$$

where ^s denote small percentage changes. We can combine these to give an expression for the real exchange rate:

$$\hat{e} = (1-\gamma)\left(\left(\hat{c}_T^D - \hat{c}_T^F\right) - \left(\hat{y}_N^D - \hat{y}_N^F\right)\right) \qquad \textbf{(22)}$$

We showed in the one-good case how to calculate the effect on domestic consumption of a productivity shock that raised future output in the tradable sector. In particular, given our parameterisations, a productivity shock that raised next period's domestic output of tradable goods by 5% led to an increase in the own rate of interest on tradable goods from 4% to 6.5% and a rise in first-period tradable-goods consumption of 1.2%. As this had to be financed by a capital inflow from abroad, we can immediately note that foreign first-period consumption of tradable goods would drop by 1.2%. If we assume a value for γ of 0.1 (reasonable for the United States), equation **(22)** suggests that the real exchange rate would appreciate by 2.16% in the period of the shock.

Moving into the following period, we can note that both the home and foreign countries experience tradable goods consumption growth 2.5% higher than previously. This implies that domestic consumption of tradable goods in period 2 will be 3.7% higher than its original trend and foreign consumption of tradable goods will be 1.3% higher than its original trend. The difference between the two remains the same and hence equation **(22)** implies that the real exchange rate stays 2.16% above its original level. This response of the real exchange rate to tradable-sector productivity shocks is known as the Balassa-Samuelson effect.

Notice that equation **(22)** also has strong predictions about the effect of a productivity shock in the non-tradable sector. In particular, a 1% current-period shock to output in the domestic non-tradable sector will lead to a 0.8% depreciation of the real exchange rate in the current period, with the future path of the exchange rate depending on whether or not this shock is temporary or permanent. The intuition for this result is exactly as stated in the main text: such a shock implies an increase in the relative supply of domestic goods and, given this, their relative price must fall. A perfectly anticipated shock to future output in the non-tradable sector will lead to no change in the current exchange rate and a depreciation of the future exchange rate. Because shocks to the non-tradable sector do not raise the domestic country's wealth in terms of tradable goods relative to the rest of the world, there will be no capital flows in response to such a shock. All that changes is the relative price of non-tradable goods and, by implication, the real exchange rate. This is a direct result of the form assumed for the utility function and may not hold in more general cases.

Putting these results together enables us to consider the effects of a productivity shock to both sectors that raised expected future output. Suppose that the home country is hit with a productivity shock that raises output in both sectors by 5% in period 2 and this is perfectly anticipated in period 1. The above analysis suggests that the real exchange rate would appreciate by 2.16% in period 1 before depreciating by 4% in period 2, finishing 1.84% below its original level. The intuition here is that, since the non-tradable sector represents 90% of the economy, the effects of the shock in the non-tradable sector are going to dominate those in the tradable sector. In the short run, the desire of domestic residents to smooth consumption will lead to a rise in their real interest rate relative to the rest of the world and hence a jump appreciation of the real exchange rate. In the long run, the productivity shock will lead to an increase in the relative supply of domestic goods and, given this, their relative price (the real exchange rate) must fall.