RESERVE BANK
OF NEW ZEALAND

DP2006/04

Other stabilisation objectives within an inflation targeting regime: Some stochastic simulation experiments

Tim Hampton, David Hargreaves

and James Twaddle

May 2006

JEL classification: E52, E58, F47

Discussion Paper Series

DP2006/04

Other stabilisation objectives within an inflation targeting regime: Some stochastic simulation experiments[1]

Tim Hampton, David Hargreaves and James Twaddle[2]

Abstract

We use the Reserve Bank of New Zealand's macroeconomic model (FPS) to look at the feasibility of using monetary policy to reduce variability in output, the exchange rate and interest rates while maintaining an inflation target. Our experiment suggests that policy could be altered to increase the stability of interest rates, the exchange rate, inflation, or output, relative to the base case reaction function in FPS, but such a policy would incur some cost in terms of the variability of the other variables. In particular, we find that greater exchange rate stability would have relatively large costs in terms of the stability of all three other variables, primarily because monetary policy that leans too dramatically against exchange rate disturbances can create significant real economy variability. Relative to West (2003), we find larger costs of operating monetary policy to achieve exchange rate stabilisation. We attribute this finding to the relatively inertial inflation expectation process in FPS.

---

[1] The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of New Zealand. We thank David Archer and Christie Smith for comments. All errors and omissions are the responsibility of the authors.

[2] Address: Economics Department, Reserve Bank of New Zealand; PO Box 2498, 2 The Terrace, Wellington, New Zealand. Tel: 64-4-471 3665; fax: 64-4-473 1209; *email address:* David.Hargreaves@rbnz.govt.nz.

1 Introduction

The Reserve Bank of New Zealand's primary objective is achieving and maintaining stability in the general level of prices. What price stability means in practice is articulated in the Policy Targets Agreement (PTA), an agreement between the Reserve Bank Governor and the Minister of Finance that guides the Bank's operation of monetary policy. The centrepiece of all PTAs to date has been a target for inflation, albeit with further considerations to guide the Bank when responding to cyclical variability in the economy and inflation. However, there could be a number of ways to operate monetary policy consistent with medium term price stability. To take an extreme example, a PTA could direct the Bank to maintain a fixed level of the exchange rate against a country that is maintaining price stability. Since real equilibrium exchange rates remain fairly constant over long periods, the fixed nominal exchange rate would be the "anchor" for the price level of the economy and fluctuations in external competitiveness would occur though cyclical fluctuation in prices rather than the exchange rate.

That PTAs to date have focused on a numerical target range for the rate of inflation in the Consumer Price Index (CPI) does not mean that the Bank is meant to be unconcerned about fluctuations in the exchange rate, or other key macroeconomic variables. While having a domestic inflation target limits the ability of monetary policy to stabilise the real exchange rate, recent PTAs have included a directive (clause 4b) that the Bank should act to prevent unnecessary fluctuations in the exchange rate and output and interest rates. In other words, to the extent that it is possible to do so without compromising price stability, the Bank is expected to act to lessen the scale of fluctuations in the exchange rate.

One example where there is very little conflict between the various objectives of monetary policy is a so-called "portfolio" exchange rate shock. If financial markets suddenly consider New Zealand a safer place to invest than they had previously thought and a flood of capital begins to push up the value of the New Zealand dollar, we would typically see a weakening in the real economy (as exports fell and imports increased). The weak real economy and cheaper imports would both tend to lead to a fall in the inflation rate. To a degree, most economic models suggest that these effects could be lessened through a temporary reduction in the New Zealand interest rate.

However, circumstances of conflicting objectives are also likely to be reasonably common. When there is very strong domestic demand in New Zealand and inflation is rising, it will often be necessary to raise interest rates to stop a persistent inflationary problem from developing. But the higher interest rate is also likely to push up the real exchange rate. These conflicts mean that there are likely to be trade-offs between the variability of inflation, interest rates, output, and the exchange rate.

In this paper, we use the Bank's primary model of the New Zealand economy (the Forecasting and Policy System or FPS) to consider a wide variety of ways of operating monetary policy. How FPS and our experiment work are described in section 3.

The Reserve Bank uses FPS to generate projections of the outlook for the New Zealand economy. These projections contain a stylised description of how monetary policy might respond to the ongoing events they describe. The stylised monetary response is determined by the policy reaction function in FPS.[3] The primary objective of this paper is to consider whether the current FPS policy reaction function is appropriate, given the Bank's PTA commitment to avoid unnecessary instability. We look at this issue in detail in section 4. A particular focus of our research is whether FPS suggests there are ways of running monetary policy that would produce greater exchange rate stability without significantly compromising other objectives. Our answer is *no*: our experiment suggests there is little scope to lean against exchange rate variability without generating a significant degree of additional variability in other variables, including the interest rate, echoing the results of West (2003).

One of the greatest uncertainties about the structure of the economy concerns the behaviour of the exchange rate, and it seems likely that the way the exchange rate responds to monetary policy is a factor in our results. Consequently, in section 5, we look at how our conclusions change if we describe the exchange rate's behaviour in different ways. We find our conclusions to be fairly robust to several changes in exchange rate behaviour

---

[3] The Bank's published projections and actual policy settings may deviate from the policy reactions produced by the baseline policy reaction function, reflecting the fact that FPS and the reaction function consider only a subset of the factors that underlie actual policy deliberations. See Hampton (2002) for further discussion of the role of the FPS in the Bank's formation of interest rate projections.

in the FPS model, but identify some areas for further work.

2 Previous literature

The methodology we adopt has been used before to assess alternative policy reaction functions in large scale macroeconomic models, in both Canada (Black, Macklem and Rose 1997) and New Zealand (for example, Conway *et al* 1999). These studies tend to demonstrate a trade-off between inflation, output and interest rate variability. Policy reaction functions that attempt to promptly and aggressively stabilise inflation tend to demonstrate significantly higher output and interest rate variability than reaction functions that respond more gradually. This mirrors conclusions drawn by researchers that have worked with smaller models that pay greater attention to micro-foundations (eg Svensson 2000, who advocated a "flexible" approach to inflation targeting based on his results).

Relative to this literature, the main extension in our analysis is the consideration of exchange rate stabilisation, as well as the stabilisation of the other three variables of common interest. The extent to which the exchange rate can be stabilised by a central bank that is operating monetary policy primarily to support an inflation target has been dealt with relatively infrequently in the literature. Two notable pieces of work that we are aware of are summarised here.

Leitemo and Söderström (2005) and West (2003) both work with small open economy extensions of the New Keynesian model. Their models have forward-looking IS and Phillips curves, an uncovered interest parity (UIP) exchange rate specification, and a generalised Taylor rule that allows for an exchange rate term as an additional element in the reaction function.

Leitemo and Söderström concentrate on whether an exchange rate term improves the Taylor rule's ability to stabilise the economy. The method they use is to first calculate optimal Taylor rule coefficients for a classic inflation/output gap reaction function in their model, and then calculate the optimal coefficients for a reaction function that includes the deviation from equilibrium or the change in the exchange rate. The reaction function with the exchange rate term is always going to do at least as well as the classic Taylor rule because the coefficient on the extra term can always be set to zero to produce identical results to the classic Taylor rule. However, Leitemo and Söderström find that the gains from including an exchange rate term are small. They test this conclusion further by varying the characteristics of the model, notably the forward lookingness of exchange rate expectations, the extent of exchange rate pass-through to the CPI, and the persistence of shocks in the exchange rate equation. In general, they still find only minor gains from incorporating an explicit response to the exchange rate in the reaction function.

Finally, Leitemo and Söderström examine the consequences of model uncertainty by testing the ability of an optimal reaction function from their base case model to stabilise the economy under alternative assumptions. They find that reaction functions that include the change in the exchange rate (either real or nominal) tend to be less robust to model uncertainty than classic Taylor rules or reaction functions that react to the level of the exchange rate.

West (2003) is less concerned with model robustness and more concerned with obtaining a realistic empirical framework for his simulation experiment. Working with a similar base model, West estimates a rich structure of cross- and auto-correlations for the disturbance terms that permit the model to map to observed data on the New Zealand economy fairly well. Working with a Taylor rule and adding an explicit response to the level of the exchange rate, West finds that there are significant costs to exchange rate stabilisation. Specifically, West suggests that "reducing exchange rate variability by about 25 percent would require increasing output variability by about 10 to 15 percent, inflation variability by about 0 to 15 percent and interest rate variability by about 15 to 40 percent."

Moreover, West considers his results a lower bound, for reasons that appear to apply equally to our study. Specifically, both we and West assume the monetary authority can perfectly diagnose exchange rate disequilibria, correctly predict the impact of monetary policy on those disequilibria, and respond to them instantly. As West describes, none of these things are likely to hold in reality, further complicating any policy designed to stabilise the exchange rate.

In this paper, we test West's conclusion that there are significant costs to exchange rate stabilisation. Similar to West, we investigate how macroeconomic volatilities change in response to changes to the monetary policy reaction function. Compared to West, we work with a larger model of the New Zealand economy, calibrated to reflect the Reserve Bank's understanding of the business cycle. In addition, the model incorporates an inflation forecast-based policy reaction function, rather than a backward-looking Taylor rule.

3 FPS and Stochastic Simulations

FPS is documented more fully elsewhere, and described here only very briefly.[4] The FPS model describes the interaction of five economic agents: households, firms, a foreign sector, the fiscal authority and the monetary authority. The model has a two-tiered structure. The first tier is an underlying steady-state structure that determines the long-run equilibrium to which the economy converges. The second tier is the dynamic adjustment structure that traces out how the economy converges towards that long-run equilibrium. In experiments like this one, which deal with the behaviour of monetary policy over the business cycle, it is the dynamic structure that drives the results.

The dynamic structure is designed to capture the persistent macroeconomic cycles that are a feature of the New Zealand and other countries' economies; simple rational models of dynamic behaviour often do not exhibit this degree of persistence. In FPS, this persistence comes from a number of sources. First, many dynamic equations in FPS include lag relationships that mimic the empirical properties of the data rather than being formally micro-founded. For example, the peak effect of the real exchange rate on exports takes around six quarters. Inflation expectations are partly modelled as a forward looking rational process, but also respond very gradually to lagged inflation out-turns. Second, many equations are overlaid with costs of adjustment using a polynomial adjustment cost framework (see Tinsley 1993). Third, costs of adjustment are also proxied in several cases by gradual stock-flow adjustment: for example, any gap between the actual and desired capital stock is closed only gradually by investment.

Our basic tool in the calculation of the results of different sorts of monetary policy behaviour is the stochastic simulation. In the FPS context, these are

---

4 See Black, Cassino, Drew, Hansen, Hunt, Rose and Scott (1997) for a more complete description of the FPS core model.

described in some detail in Drew and Hunt (1998a). Essentially, the model is struck with a series of randomly generated shocks to key variables over a 25 year simulation. These shocks have inflationary implications, and lead to monetary policy responses. The success of those policy responses in stabilising variables like inflation, output, the exchange rate, and interest rates are tracked. The process is repeated 100 times with a different randomly generated series of shocks each time, so we can build up a picture of how each reaction function performs in stabilising the economy in a wide variety of situations. The same set of shocks is applied to each reaction function to ensure that any differences in economic outcomes reflect only the different reaction functions.

How can these experiments help inform monetary policy in the light of the Bank's commitments under section 4 of the PTA? One way to operationalise the avoidance of "unnecessary instability" is to ensure that there appears to be no alternative monetary policy reaction function that would reduce the variability of interest rates, the exchange rate, inflation or output without increasing the variability of one of the other variables. A policy that has a better alternative in this sense is inefficient.

Suppose there were three ways of operating monetary policy available to a central bank (A, B and C). Suppose further that research determined that, on average over a long period, those policies would lead to different amounts of inflation and output variability, as depicted in figure 1. If the sole criterion for choosing between these three ways of running policy involved minimising the cyclical variability of output and inflation, policy B is "inefficient" in the sense that it produces a greater amount of cyclical variability in output and inflation than policy A. Whether we prefer policy A or policy C depends on decisions about the trade-off between inflation and output variability.

Figure 1
Inflation and output variance of 3 stylised reaction functions
Standard deviations

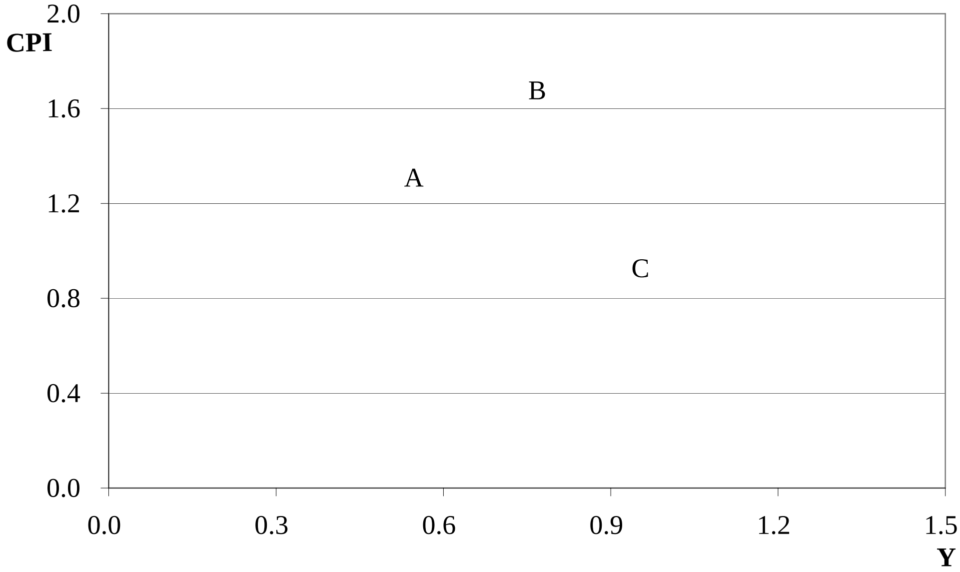

Details of the experiments

Past stochastic simulations using FPS have used shocks based on a method that induces cross- and auto-correlation in the shocks so that the shocks mimic the impulse responses of a vector auto regression model estimated on New Zealand data. Our method in this paper differs in a couple of important respects:

- Rather than estimate using the levels of official data, we look at historical disequilibrium or "gap" data from the Reserve Bank's projection databases. As described in (for example) Basdevant and Hargreaves (2003), forecasting analysis at the Bank generally makes recourse to filtering methods to extract the perceived disequilibrium component of a variable. Because we are trying to simulate business cycles, it seems appropriate to use these cyclical series.

- We also follow more closely the way that FPS is used for forecasting by the Bank by explicitly allowing for autocorrelation in the shocks, but not cross-correlation.[5] That is, we estimate the persistence of the disequilibria in each of the variables we are shocking and give the shock terms sufficient autocorrelation to make the responses of the variable after a shock consistent with the data.

- When we run a simulation with the FPS base case policy reaction function, the variability (or second moments) of the simulated data is reasonably close to the historical variability of the same variables. This gives us some confidence that the shocks we are generating are a reasonable representation of the disturbances to the economy seen through history.[6]

Table 1
Root mean squared deviations from short-run equilibria

	Output	Exchange rate	Interest rate	Annual CPI inflation	Annual domestic inflation
Simulated	1.4	9.1	1.7	1.1	1.2
Historical (1991-2003)	1.6	11.0	1.8	0.7	0.9

The variables used are from the Bank's model database (March 2004). The exchange rate is a real TWI measure, the interest rate is the nominal 90 day rate, and domestic inflation is a synthetic model variable (effectively the fitted values from a historical Phillips curve). The exchange rate is expressed as percentage deviations from the equilibrium values in the model database.

In departing from the base case reaction function, we consider reaction functions of the form shown in equation (1).

---

5 See Drew and Hunt (1998b) for a description of how FPS is used to create projections. In general, starting point shocks are allowed to run-off with autoregressive decay, with the speed of decay informed by looking at the implied shocks obtained from inverting the model over history.

6 An alternative explanation of this result is that the shocks that we are using and the parameters assumed in the model are both wrong, but in an offsetting way.

$$i_t = i^* + \alpha \left[\sum_{i=j-1}^{j+1} \theta \left(\pi_{t+i}^e - \pi^T \right) + \phi_1 (z_t - z^*) + \phi_2 (z_{t+3}^e - z^*) \right] + (1-\alpha)(i_{t-1} - i^*) \quad (1)$$

where i_t is the nominal policy interest rate, i^* is the neutral nominal interest rate, π_t^e is expected inflation, π^T is the target inflation rate, z_t is the real exchange rate, z^* is the equilibrium exchange rate, and z^e_{t+i} is the expectation of z_{t+i} formed at time t. In these equations, the exchange rate is measured as the price of foreign currency in units of domestic currency; a rise in z_t is therefore a depreciation.

We search over the following parameters

- θ: the weight on future deviations of inflation from target – we consider 0.5, 1.7, 5.0 and 14.0; base case = 1.7.
- j : the average horizon at which inflation deviations are targeted – we consider 3, 7, and 11; base case = 7.
- ϕ_1: the weight on contemporaneous deviations of the exchange rate from equilibrium – we consider 0 and 0.2; base case = 0.
- ϕ_2: the weight on expected deviations of the exchange rate from equilibrium in 3 quarters – we consider 0, 0.2 and 0.7; base case = 0.
- α: the extent to which interest rates are adjusted towards 'desired' levels each quarter – we consider 0.2, 0.66 and 0.9; base case = 0.2.

This represents a search over 216 possible policy reaction functions.[7] How the results change across these different reaction functions is discussed in the next section.

Because the exchange rate equation is of crucial importance to the analysis in this paper, we discuss it in more detail. The crucial details of real

[7] Working numerically, this is a large computational problem, made more feasible by continuing advances in computer speed: the search enumerated above took around 100 hours of computer time on a 1.7 GHz Pentium 4.

exchange rate determination in FPS are contained in equations (2) and (3).[8]

$$z_t = \eta_1 \cdot z_{t-1} + \eta_2 \cdot z_t^e \cdot \frac{\left(1 + r_t^* + rp_t\right)}{\left(1 + r_t\right)} + \left(1 - \eta_1 - \eta_2\right) \cdot z^* + \varepsilon_t \qquad (2)$$

where:

$$z_t^e = \mu_1 \cdot z_{t+1} + \mu_2 \cdot z_{t-1} + \left(1 - \mu_1 - \mu_2\right) \cdot z^* \qquad (3)$$

and the variables are defined as above. In addition, r_t is the real interest rate, r_t^* is the foreign real interest rate, rp_t is the risk premium (constant in our experiments), and ε_t is a shock term.

Equation (2) is for the real exchange rate, which is a function of its own lag, expectations for the real exchange rate, domestic real interest rates relative to foreign real interest rates (allowing for a risk premium), and its long-run fundamental value. Expectations for the real exchange rate, shown in equation (3), are a function of exchange rate lags and leads, as well as its fundamental long-run value. The nominal exchange rate (not shown here) is solved for in FPS using the ratio of foreign and domestic price levels.

This exchange rate model is similar to the UIP theory pioneered by Dornbusch (1976). The essential point of that theory is that exchange rates should move to equalise the risk-adjusted returns on assets in two economies. This can be most easily illustrated with an example. Suppose the NZD/USD exchange rate and NZ and US interest rates are at "normal" levels, when the US Federal Reserve unexpectedly announces a 100 basis point reduction in interest rates that markets expect to last for three years, so that the yield on all US securities of less than three years in duration falls by 100 basis points. A US investor who wants to invest for three years can now earn 1 percent per annum more by investing in New Zealand. Abstracting

8 Since this analysis was completed during 2004, there have been a number of changes to FPS. Most importantly for our results here, the exchange rate equation in FPS has been modified to include a response to export prices. The effects of this change on our results would be complicated, but we speculate that it would tend to increase the costs of an explicit response to the exchange rate because higher terms of trade would lead to an exchange rate appreciation and also a boost to output and inflationary pressures.

from risk, uncovered interest parity suggests that the New Zealand dollar should appreciate until this extra interest rate return is counterbalanced by an equal expected loss from the currency (of 1 percent per annum). In this example, UIP implies that on the day of the announcement, the NZD should appreciate against the USD by 3 percent, before gradually falling back to normal levels.

The difference between the "pure UIP" specification and the FPS specification described above is in the inclusion of lagged and equilibrium terms in FPS, both in the formation of exchange rate expectations and in the exchange rate equation itself. These lags and equilibrium terms attenuate and slow the response of the exchange rate to interest rates, although it remains UIP-like. A full UIP response would be embedded in the equations if we set η_1=0, η_2=1, μ_1=1, μ_2=0. The actual FPS base case calibration is η_1=0.28, η_2=0.62, μ_1=0.4, μ_2=0.6.

To illustrate this point, we switch from the FPS exchange rate equation to something very close to a pure UIP equation (we keep a very small weight on equilibrium in the exchange rate equation for model-stability purposes by setting η_2=0.95). The impact of this change is very briefly illustrated in figure 2, with a permanent change to the inflation target used to motivate a shift up in interest rates and illustrate the impact of monetary policy on the exchange rate. An approximately 50 basis point increase in short-term real interest rates (first panel) leads to a gradual 0.8 percent exchange rate appreciation according to the standard model (dotted line, panel 2), and a somewhat faster and larger appreciation assuming the alternative UIP model (solid line, panel 2).

Figure 2
Impact of monetary policy shock (permanent 1% reduction in inflation target) assuming different exchange rate equations

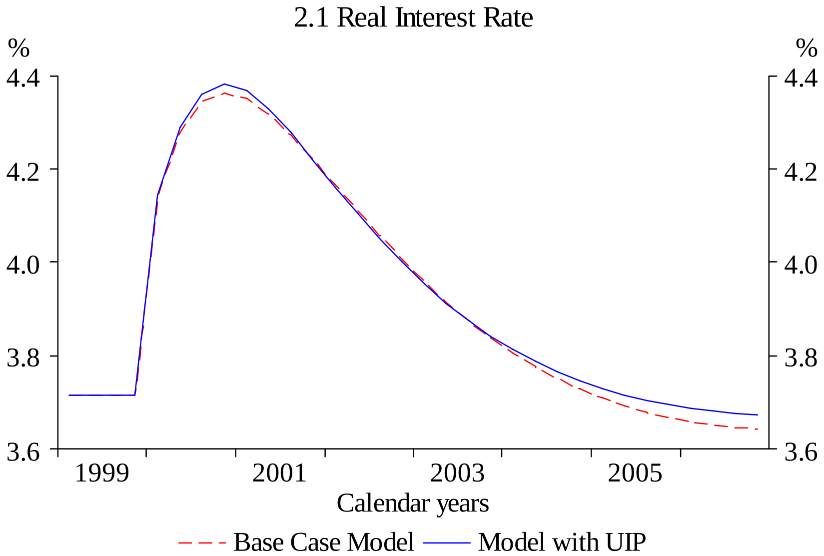

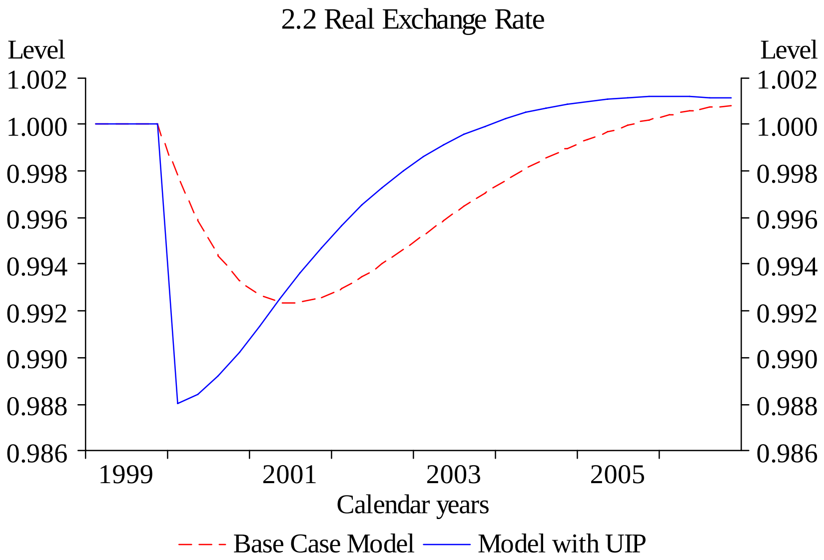

This exercise has also illustrated that even if pure UIP held in practice, then even relatively large monetary policy moves would result in a relatively small move in the exchange rate unless the policy moves were expected to be very prolonged. Hence, while UIP is the dominant theory of exchange rate determination, there seems to remain a lot of exchange rate variability unexplained by UIP.[9]

The large scale of unexplained exchange rate variability means that the shock terms added to the exchange rate equation in stochastic or analytic studies of policy reaction functions are substantial. This applies both to our study, and the work by West and Leitemo and Söderström. The need for large shocks tends to mean that the exchange rate is presumed to be primarily a source of macroeconomic variability, rather than a shock absorber. We consider this further in section 5, where we examine the consequences of removing the stochastic shock from the exchange rate equation, so that the exchange rate functions solely as a shock absorber.

4 Results

The inflation/output trade-off is a common result in quantitative and analytical exercises that consider different monetary policy reaction functions. The trade-off is depicted in figure 3, which shows the average variability in inflation and output produced by the different reaction functions we consider. As discussed earlier, we are likely to find that some of these outcomes are dominated by other outcomes where inflation variability or output variability is reduced without increasing variability in the other variable. By eliminating dominated outcomes, we end up with an efficient frontier – shown here by the black dashed curve.

---

9 At least without appealing to a time-varying risk premium.

Figure 3
Inflation and output variability under alternative reaction functions
Standard deviations

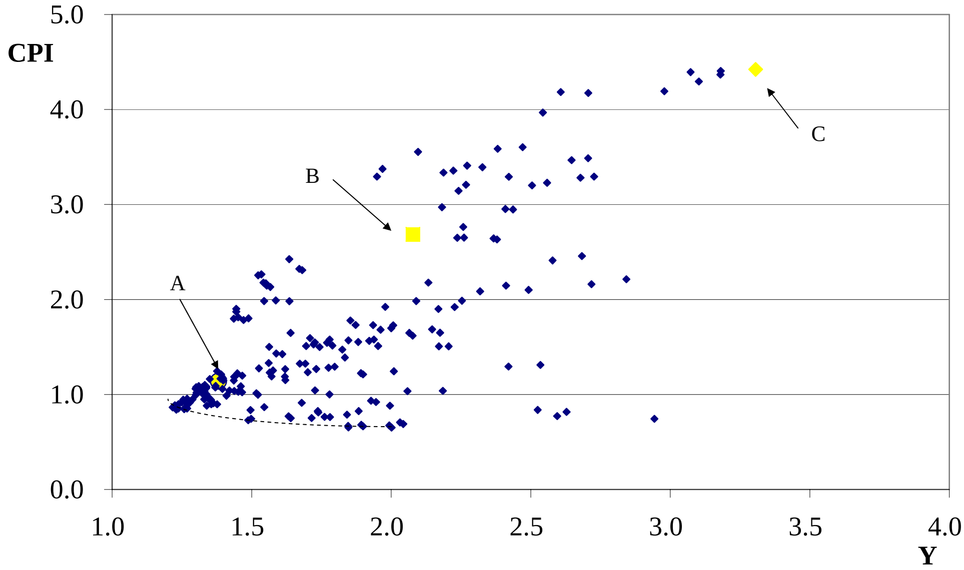

Figure 3 shows that there are a large number of inefficient reaction functions in inflation/output space. Included amongst the inefficient reaction functions is a circular yellow marker, marked A, which denotes the base case reaction function (markers B and C are discussed below). The inefficiency of the base case reaction function in this dimension suggests that, if all the Bank was concerned about was inflation and output variability, it would not be appropriate to use the base case reaction function. However, if the Bank is also concerned about not creating excessive variability in the exchange rate or interest rates, it is possible that the base case reaction function is efficient. We investigate this further below.

Figure 4 documents the combinations of exchange rate and interest rate variability achieved by these different reaction functions. It shows us that there is a significant trade-off between exchange rate and interest rate variability in the experiment. The exchange rate is assumed to be innately volatile, and the monetary authority needs to respond aggressively to exchange rate shocks if it wants to stabilise the exchange rate. Specifically, we can consider a move from the base case reaction function to the rule marked B, which puts explicit weight on targeting the exchange rate. This reduces the average “deviation from equilibrium” of the real exchange rate from around 9.0 to 7.0 (roughly 25 percent) but implies an increase in real interest rate variability of around 75 percent (from 1.8 to 3.2), and output variability of around 60 percent (from around 1.3 to 2.1).[10]

Figure 4
Exchange rate and interest rate variability under alternative reaction functions
Standard deviations

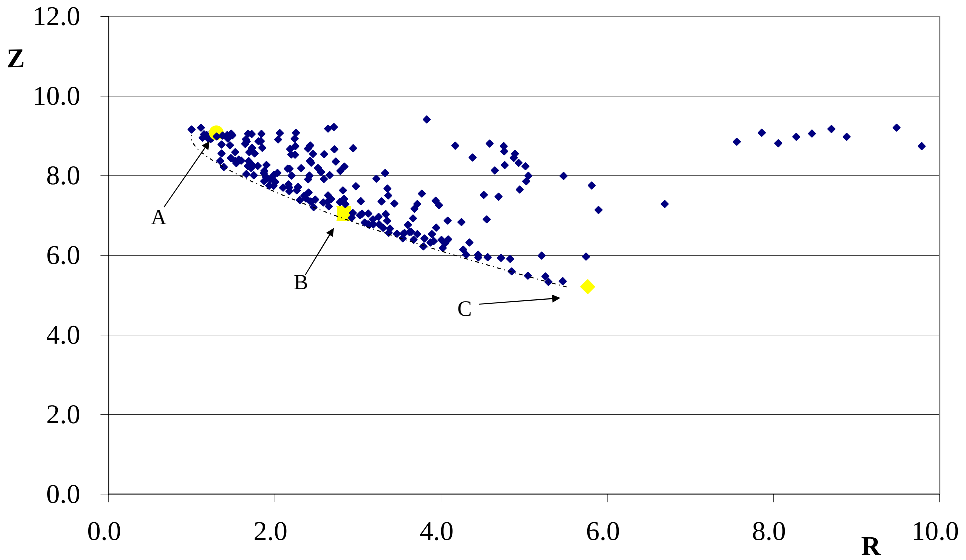

These costs of exchange rate stabilisation are qualitatively similar to West’s, but of a larger magnitude. We look at the reasons for the larger magnitude further in the next section. Before that, we are interested in examining whether the base case reaction function used in FPS is efficient in this experiment, or whether there are other policy reaction functions that appear to give better results.

---

10 To see this, look at the policy reaction function represented by point C in figures 3 and 4. Reaction functions that are efficient in exchange rate/interest rate space are often quite inefficient in inflation/output space.

An instructive way to examine efficiency across more than 2 dimensions is to look again at figure 3 and observe that there are a number of policy reaction functions that produce lower variances of inflation and output than the base case reaction function. If some of these reaction functions produce lower variance in exchange rates and interest rates *as well*, then the base case reaction function is inefficient. We test this in figure 5, by showing the real interest rate and exchange rate variances attainable by the base case reaction function (marked A) *and all reaction functions that dominate the base case reaction function in inflation/output space*.

Figure 5
Exchange rate and Real interest rate variability under reaction functions that dominate the base case in output/inflation space
Standard deviations

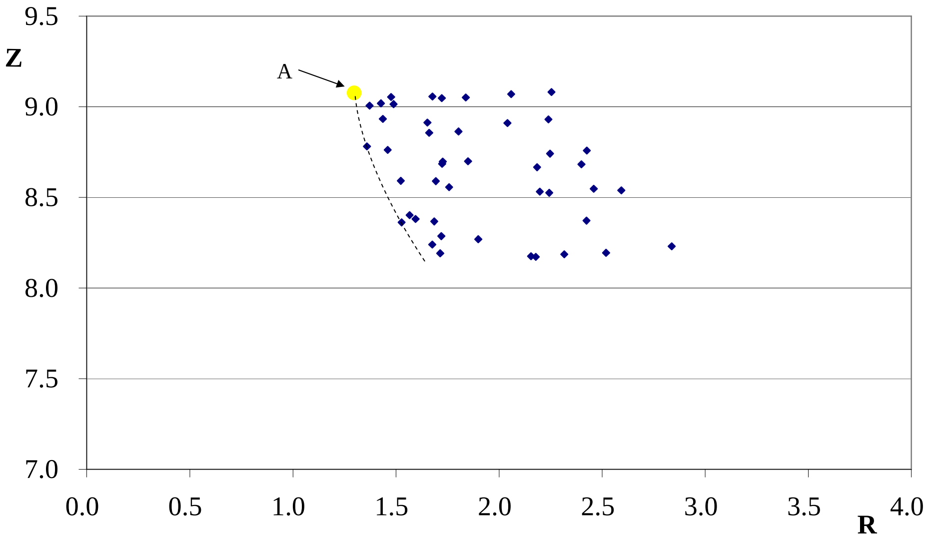

As figure 5 shows, the base case reaction function is efficient in the sense that no reaction function simultaneously produces lower inflation, output, and real interest rate variability in this experiment[11] (the base case policy

[11] Of course, this efficiency result is for a search over a limited number of rules and a particular set of shocks. More generally, we do not intent to imply that the base case rule is necessarily the literally best way of running monetary policy in all

reaction function has the lowest real interest rate variability of all the reaction functions that dominate it in inflation/output space). Many of the reaction functions that produce better inflation/output variability have significantly more real interest rate variability.

There are a number of reaction functions that appear to stabilise the exchange rate slightly (by up to 10 percent – note that the scale does not begin at zero) better than the base case reaction function, at the cost of only a relatively small increase in real interest rate variability (around a 15 percent increase). Examining these points reveals that they are reaction functions with a very high weight on inflation forecast deviations from target (5 or 14) and a long horizon for targeting inflation (10 to 12 quarters). In practice, what these reaction functions are doing is always ensuring forecast inflation returns to target within 3 years but, as far as possible within that framework, also leaning explicitly against the exchange rate. Remembering that we have restricted our attention here to reaction functions that also stabilise inflation and output rather well, these reaction functions appear to offer a legitimate alternative to the FPS base case rule.

However, studies of model robustness (notably Levin, Wieland and Williams 2003) note potential dangers of inflation forecast-based reaction functions with long forecast horizons. A reaction function that targets inflation forecast deviations from target 10 to 12 quarters ahead will not respond to a shock if the model suggests the impact on inflation will only last two years. Of course, it is not always easy to predict the persistence of inflation disturbances in advance. For example, a sharp rise in import prices may only have a temporary effect on inflation, or it may have a prolonged effect via inflationary expectations, and it is difficult to predict in advance which outcome will predominate. Reaction functions that target inflation at shorter horizons (eg responding to any inflation expected to last for six quarters or more, as in the FPS base case reaction function) are less dependent on the model assumptions about the persistence of inflation. This risk represents a reason for caution about long-horizon forecast based policy reaction functions.

---

circumstances – see Hampton (2002) for more discussion of how we use the reaction function in practice.

5 Why are significant reductions in exchange rate variability so costly in our experiment?

The previous section documented that there are high interest rate volatility costs from stabilising the exchange rate. When interest rate reaction functions seek to reduce exchange rate variability, they also significantly increase the variability of domestic spending and inflation.

Our results are qualitatively similar to those of West (2003), but our results suggest the costs of exchange rate stabilisation would exceed those posited by West. It is thus interesting to compare our experiment to West's and consider what essential characteristics of our model and experiment make exchange rate stabilisation relatively more costly.

West works with a four equation model that is consistent with recent New Keynesian theory (for example, see Gali and Monacelli 2002). The model is micro-founded, with explicit assumptions about the nature of price setting creating temporary impediments to price adjustment. The rational expectations microfoundations mean that both the IS curve and Phillips curves are fully forward looking.

In contrast, FPS is a much larger model. While its core can be derived from rational expectations microfoundations, its dynamic equations exhibit a considerable degree of inertia that is not explicitly derived from microfoundations. For example, inflation expectations are backward looking to a significant degree.

If a central bank responds directly to an exchange rate depreciation by raising interest rates, the higher interest rates will create a negative impulse to inflation. Once the exchange rate returns to equilibrium, the monetary authority eases interest rates in order to bring inflation back to target. The difference between the models is that bringing inflation back to target will be considerably more difficult in a model where inflation expectations are backward looking and have become entrenched below the inflation target.

To illustrate the importance of the inertia in inflation expectations, we run a simple experiment in FPS in the base case (Figure 6) and in a recalibrated model where we make inflation expectations fully forward looking or rational (Figure 7). We shock each model with a 5 percent exchange rate depreciation (in 2000Q1), and illustrate the impact of that depreciation assuming two different reaction functions in each model: the base case reaction function and one which responds directly to the exchange rate as well as forecast inflation.

Figure 6
Impact of direct exchange rate targeting after an exchange rate depreciation: FPS base model

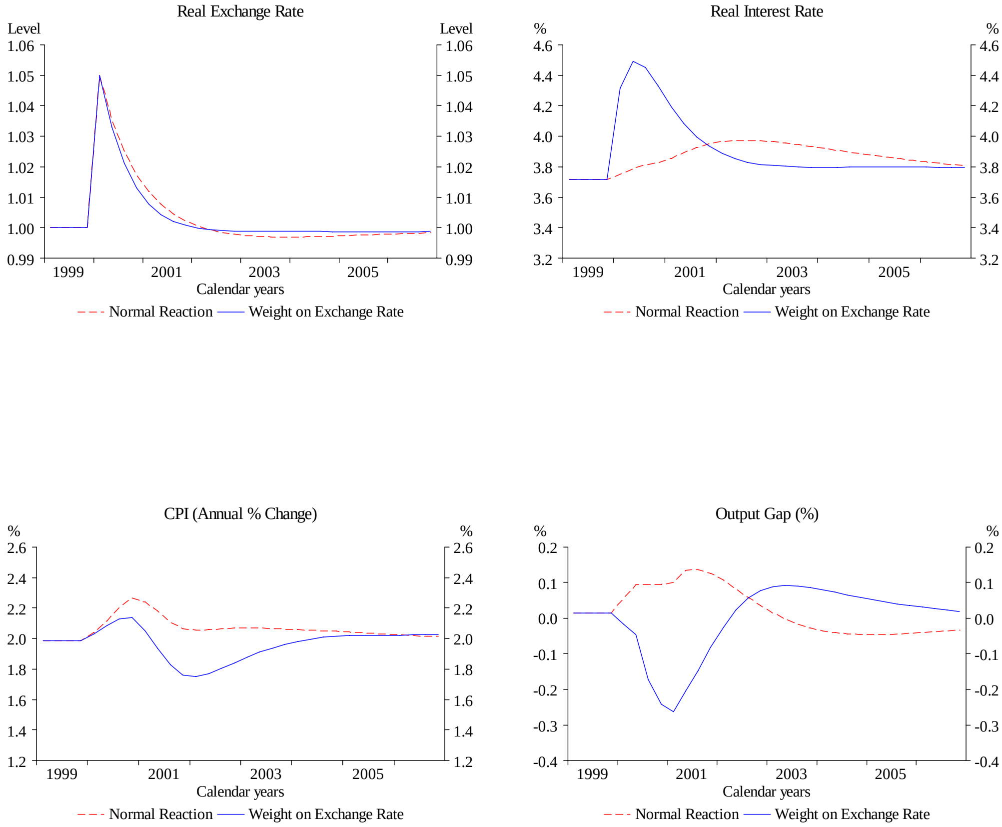

Figure 7
Impact of direct exchange rate targeting after an exchange rate depreciation: Model with rational expectations

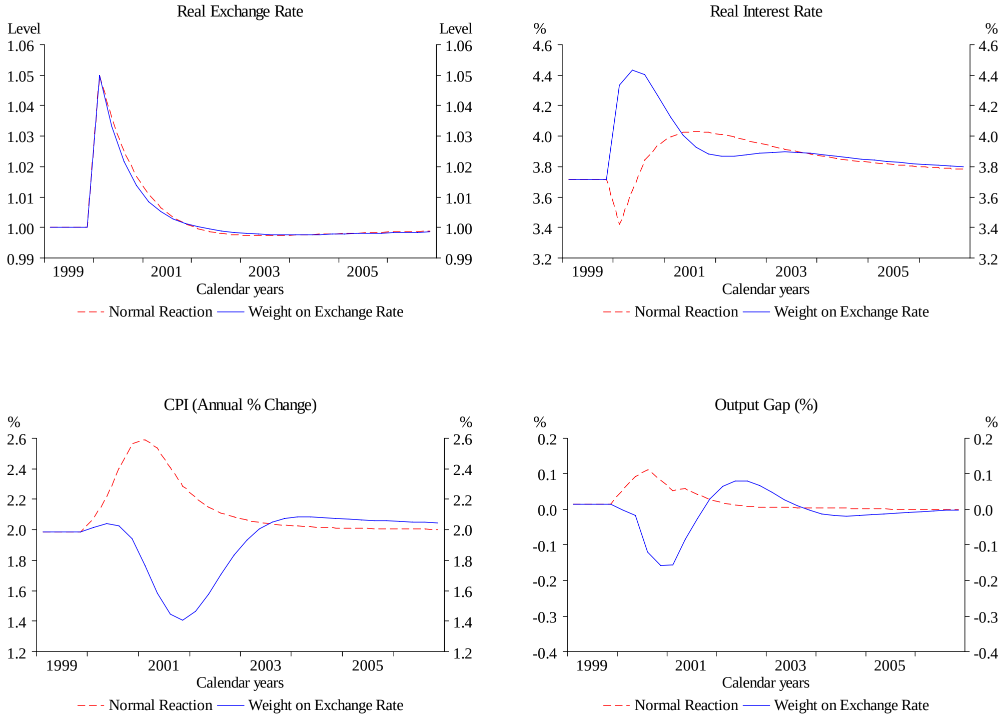

Imposing rational expectations appears to change this simulation in two key ways. First, the ability of monetary policy to respond to the exchange rate is somewhat constrained by the unanchored nature of expectations: a monetary policy tightening in the rational expectations model causes inflation to fall faster than in the base model, so policy cannot be tightened as much in response to this exchange rate appreciation. But second, after the exchange rate shock has been leaned against, it is also much easier in the rational expectations model to return inflation to target. To see this, note how the secondary cycle in the output gap when using the exchange rate targeting reaction function is positive for about two years in the rational expectations model, but more like five years in the base case model. In other words, departing from inflation control to lean against the exchange rate creates a short-term inflation problem in the rational expectations model, but that short term inflation movement is more easily reversed. Overall, it seems plausible that more forward looking inflation expectations makes it easier to trade exchange rate variability off against interest rate and output variability.

An additional possible cause for the easier trade-off in West is the fact that he uses a purer form of UIP, while we use a calibrated equation for the exchange rate that delays and attenuates the impact of UIP. However, when we rerun our experiments in a model with a pure uncovered interest parity equation (as described in section 3), we obtain results that are very similar to our previous ones, and actually suggest a sharper trade-off between exchange rate and interest rate variability. We illustrate this point by plotting a stylised representation of the frontier from the first (standard) experiment on the chart of results from the pure UIP experiment (figure 8).

Figure 8
Exchange rate and interest rate trade-offs assuming pure UIP
Standard deviations

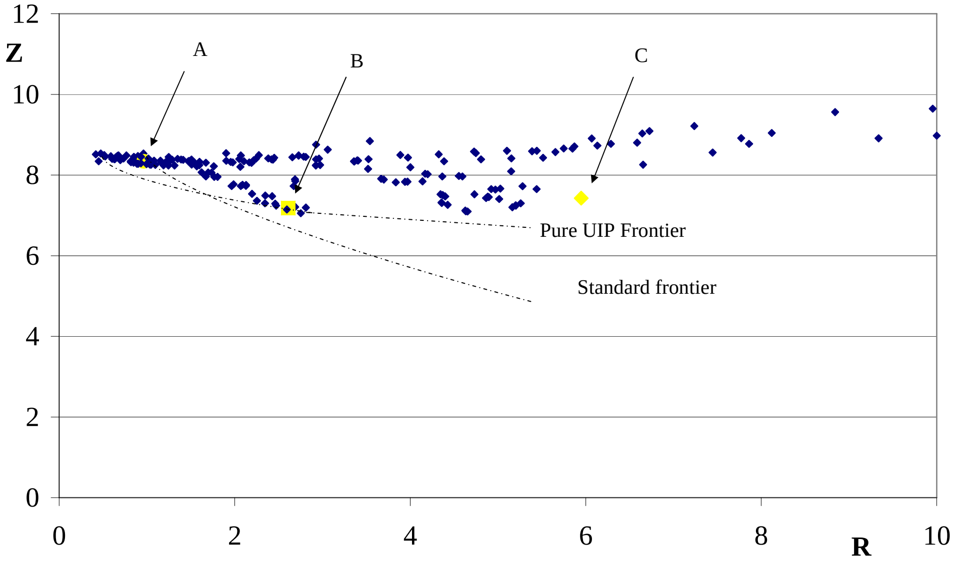

The reason why stabilising the exchange rate seems harder if we assume pure UIP relates to the way we obtain the stochastic shocks in our experiments. If we presume UIP, we have removed all intrinsic inertia in the exchange rate, so that the impact of stochastic disturbances is transient. But we design our stochastic disturbances so that the behaviour of the exchange rate is representative of the business cycle. Since the exchange rate is not intrinsically persistent in the pure UIP case, greater persistence is assigned to the shock term, making it harder for monetary policy to control the exchange rate.

In a world with a pure UIP exchange rate specification and no exchange rate disturbances, all exchange rate variability would be caused by monetary policy. Thus, there would not be a trade-off between exchange rate variability and interest rate variability. To illustrate this, and look at the implications for output and inflation, we repeat the simulations above with all non-interest rate-driven exchange rate disturbances removed. As expected, exchange rate variability becomes an increasing function of interest rate variability (Figure 9).

Figure 9
The exchange rate/interest rate frontier with no exchange rate shocks
Standard deviations

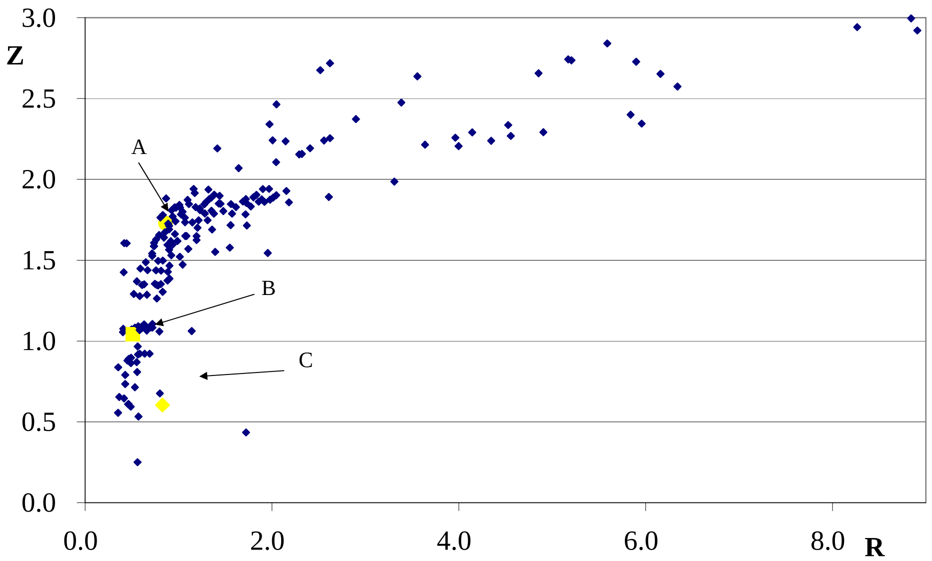

In other words, if we believed that exchange rate movements were fully determined by expected future interest rate differentials, stabilising the exchange rate would involve attenuating interest rate responses to domestic economy shocks. Such a policy approach would lead to significant additional output and inflation variability, so there would still be a trade-off between volatility in those variables and the exchange rate.

However, our empirically-driven disturbance terms support the result that much of the variability in the exchange rate is unrelated to monetary policy. This is shown by the fact that the standard deviation (from equilibrium) of the exchange rate is as little as 1 percent in our pure UIP simulation – which compares with historical estimates of around 11 percent.

Throughout all of this work, we have assumed that monetary policy is seeking to stabilise the real exchange rate. A final experiment tests how our conclusions change if monetary policy instead seeks to stabilise the nominal exchange rate. This experiment can of course be thought of as stabilising the real exchange rate plus the price level (measured as a deviation from foreign prices). We test if targeting nominal exchange rates would help stabilise the exchange rate without letting inflation get out of control.

In brief, worrying about the nominal exchange rate instead of the real exchange rate makes little difference to our conclusions. The price level varies much less than the real exchange rate in our experiments, which means that introducing concern for the nominal exchange rate is rather similar to real exchange rate targeting.

Such a result presents a potential disconnect with some available empirical evidence that suggests that economies that fix their exchange rate do not exhibit increases in real economy variability as large as our results suggest.[12] In our view, this potential disconnect points to a Lucas-critique concern with our work. It is possible that inflation expectations would begin to be formed differently in an economy with a very stable nominal exchange rate and a lot of transient variation in inflation. During a period of real exchange rate appreciation, the economy might experience significant temporary inflation, but the steady nominal exchange rate would likely anchor inflation expectations at lower levels.

These sorts of considerations led Black, Macklem and Rose (1997) to impose a (somewhat ad hoc) change to the structure of inflation expectations in a stochastic simulation experiment depicting a price level target. Without resorting to explicit modelling of how exchange rate stability would affect

[12] For example, see Ghosh et al (1997).

the formation of expectations, we think it is reasonable to say that possible changes in expectation formation mean we over-predict the variance costs of a fixed exchange rate. It is less clear whether there would be significant changes in expectation formation in a regime where the exchange rate was still floating and variable, but was more aggressively "leaned against" by monetary policy. This is an interesting area for further work, but may require a model with deeper microfoundations.[13]

Conclusion

The most general contribution of this paper is to consider whether the base case reaction function in FPS is efficient. We find that the base case reaction function is efficient, in the sense that no reaction function we examine can simultaneously produce lower real interest rate, exchange rate, inflation, and output variability.

A more specific concern in this paper was to assess the potential for, and trade-offs involved in, running monetary policy in a manner that stabilised the exchange rate better than the base case reaction function currently incorporated in FPS. We found that achieving significantly greater exchange rate stability with monetary policy would be a costly exercise. The costs that we estimated are significantly higher than those in a similar, recent study (West 2003). Further analysis suggests the additional costs in FPS are a consequence of the intrinsic inertia within FPS (when compared to smaller, directly micro-founded models of the economy). In particular, more backward-looking inflation expectations mean that leaning against the exchange rate can entrench inflation at a level some distance from the target. These represent an additional cautionary point about active exchange rate stabilisation in an inflation targeting regime, over and above those described by West.

---

13 Models with fully forward-looking inflation expectations but substantial nominal rigidities are beginning to be used to examine the relationship between optimal monetary policy in one country and the policy of the trading partner. For example, see Coletti, Lalonde, Laxton, Muir and Pesenti (2006). These complex DSGE models would presumably have greater endogenous changes in the behaviour of inflation expectations when a monetary policy rule that targeted exchange rate stability was imposed on one country.

We have investigated the robustness of our results to some simple changes in the way exchange rates are determined. We found that switching to a pure UIP specification for the exchange rate did not greatly alter the trade-offs involved, unless we were prepared to assume there were no disturbances to the UIP relationship, which would be inconsistent with New Zealand's historical experience. Even if we assumed there were no exchange rate disturbances, there was still a significant trade-off between output and exchange rate variability associated with stabilising the economy after an output shock.

There are many other sensitivity tests we could conduct for our results. One example is alternative exchange rate specifications which are based on UIP but non-rational (eg adaptive). And the analysis of alternative inflation expectation processes in section 5 shows that the exchange rate/output variability trade-off depends on other model facets besides exchange rate determination. It would probably be easiest to examine alternative features in a smaller model, along the lines of the work done by Leitemo and Söderström (2003) and Stephens (2006).

We have also noted that our analysis assumes that greater exchange rate stability does not have significant effects on the deep structure of the economy. That is, the Lucas Critique applies to our results, and an additional possible strand for future work is to consider models with explicit learning or deeper microfoundations affecting the expectations formation processes. If greater exchange rate stability reduces the variability of inflation expectations, or reduces the portfolio shocks that drive exchange rate stability, the trade-off between exchange rate stability and stability of output and inflation could become less severe.

References

Basdevant, O and D Hargreaves (2003), "Modelling structural change: The case of New Zealand," *Reserve Bank of New Zealand Discussion Paper,* DP2003/03.

Black, R, T Macklem and D Rose (1997), "On policy rules for price stability" in *Price Stability, Inflation Targets and Monetary Policy*, Bank of Canada conference volume.

Black, R, V Cassino, A Drew, E Hansen, B Hunt, D Rose and A Scott

(1997), "The Forecasting and Policy System: The core model," *Reserve Bank of New Zealand Research Paper*, 43.

Conway, P, A Drew, B Hunt, and A Scott (1999), "Exchange rate effects and inflation targeting in a small open economy: a stochastic analysis using FPS," *Reserve Bank of New Zealand Discussion Paper*, G99/4.

Coletti, D, R Lalonde, D Laxton and P Pesenti (2006), "Inflation-targeting and price-level-path targeting in a DSGE model," *Mimeo*.

Dornbusch, R (1976), "Expectations and exchange rate dynamics," *Journal of Political Economy*, 84, 1161-76.

Drew, A and B Hunt (1998a), "The Forecasting and Policy System: Stochastic simulations of the core model," *Reserve Bank of New Zealand Discussion Paper*, G98/6.

Drew, A and B Hunt (1998b), "The Forecasting and Policy System: Preparing economic projections", *Reserve Bank of New Zealand Discussion Paper,* G98/7.

Galí, J and T Monacelli (2002), "Monetary policy and exchange rate volatility in a small open economy," *NBER Working Paper*, 8905.

Ghosh, A , A-M Gulde, J Ostry, and H Wolf (1995), "Does the nominal exchange rate regime matter?" *International Monetary Fund Working Paper*, 95/121.

Hampton, T (2002), "The role of the Reserve Bank's macro model in the formation of interest rate projections," *Reserve Bank of New Zealand Bulletin*, 65(2), 5-11.

Leitemo, K and U Söderström (2005), "Simple monetary policy rules and exchange rate uncertainty," *Journal of International Money and Finance,* 24, 481-507.

Levin, A., V Wieland and J Williams (2003), "Performance of forecast-based monetary policy rules under model uncertainty," *American Economic Review*, 93(3), 622-645.

Stephens, D (2006), "Should monetary policy attempt to reduce exchange rate volatility in New Zealand?" *Reserve Bank of New Zealand Discussion paper*, DP2006/05.

Svensson, L (2000), "Open-economy inflation targeting," *Journal of International Economics*, 50, 155-183.

Tinsley, P (1993), "Fitting both data and theories: Polynomial adjustment costs and error-correction decision rules," *Finance and Economics Discussion Series*, 93-21, Board of Governors of the Federal Reserve System.

West, K (2003), "Monetary policy and the volatility of real exchange rates in New Zealand," *Reserve Bank of New Zealand Discussion Paper*, DP2003/09.