New Zealand's Risk Premium and Its Role in Macroeconomic Imbalances*

Craig Burnside†

June 16, 2011

Abstract

It is sometimes argued that interest rates in New Zealand are high, and that, therefore, investment is low, and growth is slow. I study the joint behavior of interest rates and the exchange rate in New Zealand in the inflation targeting period. I find that high interest rates in New Zealand, through much of this period appear to have represented compensation for the risk of rare and extreme events of the type observed during the recent global financial crisis. These factors also explain Australia's risk premium. I argue, therefore, that interest rates are likely not the central channel through which investment is low and growth is slow. I argue, instead, that New Zealands main source of vulnerability is its large net external debt position. I argue that policies aimed at increasing saving and removing distortions in favor of housing investment would reduce this vulnerability.

*This is a first draft of a paper written for the New Zealand Treasury, Reserve Bank of New Zealand, and Victoria University of Wellington policy forum "New Zealand's Macroeconomic Imbalances: Causes and Remedies," 23 & 24 June 2011, Wellington, NZ. Without implicating any of them, I would like to thank Christie Smith and John McDermott for comments, and the many economists and policy analysts that shared their thoughts on the New Zealand economy with me during my visit to the RBNZ in January and February 2011.

†Duke University and NBER. Email: burnside@econ.duke.edu.

New Zealand's policy makers have, for a considerable time, been concerned with New Zealand's low growth rate (and level) of per capita income relative to that of otherwise seemingly comparable countries. A striking illustration of the issue is provided by Figure 1, which shows the average growth rate of per capita real GDP in a sample of twenty-five OECD countries plotted against the income level of those countries in 1960. As is readily seen in the graph, New Zealand had the slowest growth in this group of countries, at a rate of 1.4% per annum. But a little caution is warranted here. Standard neoclassical growth theory predicts that the richest countries should grow the slowest. Since New Zealand was the fourth richest country in 1960, its lower growth rate is not at all surprising from the perspective of growth theory. It is surprising, from a theoretical perspective, that New Zealand's growth has been so slow that is has transitioned from being fourth richest to being ranked eighteenth out of twenty-five countries in 2009. It has been leap-frogged by 14 of the 21 countries that used to lie below it.

One factor in New Zealand's slower growth appears to be its lower rate of investment in physical capital. It has consistently, since the 1960s, ranked about 20th in terms of investment as a share of GDP. Although investment rates are not perfectly correlated with growth rates, within the OECD group of countries they are strongly correlated in the cross-section. So any explanation of New Zealand's low rate of investment constitutes a potential explanation of its slow growth since the 1960s.

A plausible explanation for New Zealand's low investment rate is that those wishing to invest in plant and equipment face high costs of borrowing. High costs of borrowing, furthermore, could be partially responsible for lower income levels through a working capital channel. To the extent that firms need to borrow to finance investment and current operations, there is both a growth effect and a level effect from high interest rates. In this paper I investigate the plausibility of this interest rate channel.

New Zealand also faces a substantial negative external position, and, for a considerable time, has run a substantial current account deficit. The interest rate channel has also been cited as a cyclical phenomenon that contributes to these imbalances on the external account. Some policy analysts, and academic economists tell the following story. High interest rates attract foreign capital that is chasing yield (the so called "carry trade"). These foreign capital inflows heat up the domestic economy and create inflationary pressure. The RBNZ responds to this pressure by raising rates. These interest rate increases attract more foreign capital. There is a further heating up of the economy, etc. In this view high interest rates, and an aggressive stance against inflation, in a sense, destabilize the economy, and exacerbate the business cycle.

I first measure the extent to which interest rates in New Zealand differ from those in other industrialized economies. The gap is substantial. I then ask whether New Zealand's high interest rates imply genuinely higher borrowing costs than those faced by firms and households in other economies. To do so, I investigate whether uncovered interest parity (UIP) holds between the New Zealand dollar (NZD) and other major currencies. UIP implies that the interest differential between two currencies is explained by the fact that the high interest rate currency is expected to depreciate relative to the low interest rate currency. If UIP were to hold, high interest rates in New Zealand would not imply higher borrowing costs. However, I find that UIP cannot explain the typically positive and large interest rate differentials between the NZD and the US dollar and Japanese yen, at least not for the full sample period studied in this paper 1985–2010. It comes close, however, to explaining the interest differential with the Australian dollar.

I also find that simple risk-based explanations of the failure of UIP do not work. By a simple risk-based explanation I mean that a model like the CAPM for stocks would explain why NZD rates are higher. This would be the case, for example, if the ex-post returns to NZD-denominated riskless securities measured in foreign currency terms, say US dollars, happened to be highly correlated with US share prices or the US business cycle. I find that this is not the case.

In the end, I argue that New Zealand's risk premium, vis-a-vis the US dollar and the Japanese yen, is largely due to the influence of extreme events, such as the recent global financial crisis. I outline a simple model in which these rare, but extreme, events affect the pricing of nominally riskless NZD securities. Consequently, in periods of relative tranquility, such as the pre-crisis inflation-targeting period (1990–2007), sharp apparent deviations from UIP are observed. But these apparent deviations from UIP reflect the fact that investors are being compensated for risks that are rarely realized. This leads me to be skeptical that interest rates are the channel through which investment demand in New Zealand is modest.

I also study the behavior of the risk premium over time. I find that it is largely driven by global factors, and these factors also drive the behavior of the risk premium in the Australian dollar. I find that innovations to the RBNZ's policy rate, if anything, seem to be negatively correlated with market expectations of the change in the value of the NZD. This leads me to be skeptical that monetary policy somehow fuels a speculative dynamic, even if speculation is part of the process by which exchange rates are determined.

I conclude by returning to some central questions about New Zealand's macroeconomic imbalances. If interest rates are not central to these imbalances then what is? And what should policy makers do? I suggest that while investment and growth are legitimate concerns, the central problem that New Zealand faces is its external vulnerability. Even here there is some good news: New Zealand is naturally hedged against some of this risk because its external borrowing is substantially in domestic currency rather than foreign currency. This puts it in a better position that the countries facing debt crises in Europe. Additionally, New Zealand's external debt position does not stem from a weak fiscal position. Nonetheless, the trend in New Zealand's net external position does not strike me as sustainable. Consequently, my own view is that policies aimed at increasing domestic saving, and otherwise reducing the vulnerability of the external position seem advisable. I conclude by arguing that historical property price dynamics have likely played an important role in explaining household's consumption and borrowing behavior. I argue in favor of a tax policy framework with less distortions in favor of housing investment.

1 Interest Rates and Risk Premia in New Zealand

I first review some facts about interest rates in New Zealand compared to the rest of the world. Figure 2 illustrates the interest differential between one-month money market rates in New Zealand and the U.S., Japan, and Australia during the inflation-targeting period. As the graph indicates, New Zealand's interest rates have always exceeded those in Japan, almost always (89% of the time) exceeded those in the U.S., and, more often than not (72% of the time), exceeded those in Australia. The average interest differentials were 5.6% versus Japan, 2.8% versus the U.S., and 0.7% versus Australia. Figure 1 also shows substantial variation over time in New Zealand's interest differentials. During the period leading up to 1998 interest rates were quite stable vis-a-vis the USD, experienced a substantial rise vis-a-vis the JPY, and gradually crept upward vis-a-vis the Australian dollar (AUD). The movements versus the U.S. and Japan in this period largely reflect the rate cuts that took place during the U.S. recession of 1991 and Japan's slump, that saw rates decline from 8% in 1991 to close to 0% by 1996. In 1998, the Reserve Bank of New Zealand cut rates substantially in the face of the recession that accompanied the Asian crisis, at a time when rates in the other countries did not decline much. After 1998 and until the recent financial crisis, rates crept upward against all three of the other currencies, with additional variation in the USD-NZD interest differential seemingly driven by volatility in U.S. rates: the substantial decline in rates in the U.S. during and after the 2001 recession, the rise in U.S. rates up to mid-2006, and the subsequent cuts that occurred after August 2007 in the lead-up to the U.S. recession and financial crisis.

1.1 Borrowing Costs and Uncovered Interest Parity

The fact that interest rates are different in two currencies does not, of course, mean that the economic cost of borrowing in the two associated countries is different. Even if borrowers only take out loans denominated in their own currency, the relative borrowing costs in the two countries can only be compared if they are measured in common currency units.

To see this, consider a New Zealand based firm that chooses to finance an investment by borrowing New Zealand dollars (NZD). The borrowing cost to the firm, of course, is i^{NZ} the interest rate on loans denominated in NZD. If the firm chooses to take out a foreign currency loan, however, its borrowing cost is i^{FOR}, the interest rate for loans denominated in foreign currency units (FCUs), minus the percentage change in the value of the NZD during the life of the loan. To summarize, the two costs of borrowing, measured in NZD are

$$i^{\text{NZ}} \qquad \text{and} \qquad i^{\text{FOR}} - \%\ \text{change in value of NZD}.$$

The two costs of borrowing can also both be measured in FCUs, in which case the comparison is between

$$i^{\text{NZ}} + \%\ \text{change in value of NZD} \qquad \text{and} \qquad i^{\text{FOR}}.$$

Investors in debt securities and foreign exchange make similar comparisons. A foreign investor considering investing in NZD-denominated nominally-riskless securities knows that his ex-post return, measured in FCUs, is i^{NZ} plus the percentage change in the value of NZD. His return on a FCU-denominated nominally-riskless security is simply i^{FOR}. If investors are risk neutral, in that they care only about the expected returns on different assets, economic theory predicts that arbitrage in asset markets will equate expected returns on all assets, implying

$$i^{\text{NZ}} + \text{expected } \% \text{ change in value of NZD} = i^{\text{FOR}} \tag{*}$$

This is the uncovered interest parity (UIP) condition, and it implies that expected borrowing costs in New Zealand are actually equivalent to foreign borrowing costs.[1] If UIP held, the data in Figure 2 would misleadingly give the impression that borrowing costs in New Zealand are typically much higher than those in Japan. If UIP holds, it also means that higher borrowing costs are not a plausible explanation of New Zealand's low rate of investment.

A logical next step, therefore, is to test whether UIP holds. The typical way this is done in the international finance literature is to rearrange (*) as

$$\text{expected } \% \text{ change in value of NZD} = i^{\text{FOR}} - i^{\text{NZ}} \tag{**}$$

This equation is useful in two respects. First, it offers us another way of interpreting the UIP condition. If the interest rate in New Zealand is consistently higher than the foreign interest rate, it means that this reflects investors expectation that the NZD will depreciate. Equation (**) also suggests that if we look at actual, rather than expected, changes in the value of the NZD they should be equal to the interest differential, $i^{\text{FOR}} - i^{\text{NZ}}$, plus unexpected forecast errors. Tests for UIP do exactly this, and involve running regressions of actual changes in the value of the NZD on the interest differential. UIP predicts that the constant in this regression should be zero, and the slope coefficient (attached to the interest differential) should be one.

There is a vast literature that tests UIP in the way I have just described. This literature overwhelmingly concludes that UIP is an empirical failure (see, for example, Fama 1984, and the surveys by Hodrick 1987, Lewis 1995, and Engel 1996). Nonetheless, here I test UIP for New Zealand for two reasons. First, since the focus here is specifically on New Zealand, I want to be as informative as possible about UIP for the NZD vis-a-vis a relevant set of foreign currencies. Second, it turns out that these tests are highly informative about the story I tell in the rest of the paper.

Table 1 provides results of running the regression

$$\Delta s = \beta_0 + \beta_1 \left(i^{\text{FOR}} - i^{\text{NZ}} \right) + \text{error}$$

where Δs is the percentage change in the foreign currency value of the NZD, β_0 is the constant (which UIP predicts should be 0) and β_1 is the slope coefficient (which UIP

[1]Here I am abstracting from default premiums that might be faced by an individual firm borrowing to finance an investment. I am assuming, for the sake of simplicity that these default premia are zero, or more precisely, that they would not be different based on the currency denomination of the loan.

predicts should be 1). I run this regression for the three different sample periods, and for three different choices of the foreign currency: the U.S. dollar (USD), the Japanese yen (JPY) and the Australian dollar (AUD).

Table 1 shows that, in the case of New Zealand, it is difficult to draw sharp conclusions about UIP from running the regression, (11). Although estimates of β_1 are consistently less than 1 across all sample periods and for all three currencies, they are sometimes not statistically significantly different from 1.[2] For example, consider the full sample period, from 1985, when the NZD was floated, to the end of 2010. For this sample period, we can reject UIP against the USD and the AUD, on the basis that $\hat{\beta}_1$ is negative and statistically very significantly different from 1. On the other hand, we cannot reject UIP against the JPY on the basis of the regression results, at least not at the 5% level of significance.

There is a good argument, however, for excluding the early floating rate period from the sample. Interest rate differentials between New Zealand and other currencies were much more volatile during this period than during the rest of the sample, and a few observations in this period are big outliers.[3] If we consider only the period of inflation targeting (1990–2010), the results are strikingly different. For all three currencies the point estimates of the slope coefficient, β_1, are very close to zero, but they are, in all cases, statistically insignificantly different from 1. So if we consider the inflation-targeting period as a whole, the evidence, at least from regression (11), against UIP is very weak.

It is interesting, however, to delve a little deeper. Suppose we also exclude the period of the recent global financial crisis (GFC, 2008–2010), and consider only the period from 1990 to 2007. In this period, for the USD and JPY, the evidence against UIP seems to be very strong. Now point estimates of the slope coefficient, β_1, are sharply negative, and statistically well below 1.[4] For the AUD, however, it remains true that the estimated slope is close to 0 and statistically insignificantly different from 1. This sharp contrast between the period leading up to the GFC, and the full sample plays a key role in may later discussion.

1.2 Currency Trading Profits

While the regressions in the previous section are informative about UIP, they only provide one test based on the predictability of exchange rates one period ahead. A different way of testing UIP involves looking at the profitability of currency trading.

To see this, consider a trader or investor who holds a long position in the New Zealand dollar (NZD), and finances his position by borrowing a foreign currency. Abstracting from transactions costs, if the trader borrows one unit of the foreign cur-

[2]Frankel and Chinn (1993) and Chinn and Meridith (2004) report that the UIP puzzle seems to diminish at very long investment horizons such as 5 or 10 years. In contrast, Burnside et al. (2011b) run UIP regressions for investment horizons up to one year, and find no tendency for the UIP puzzle to diminish as the investment horizon increases at these relatively short horizons.

[3]This is reflected in the small standard errors associated with the point estimates of β_1.

[4]This is the typical result in the literature. Contrary to what UIP predicts, it is often the case that the higher interest rate currency appreciates, rather than depreciating.

rency, the trader's profit is

$$\text{profit} = i^{\text{NZ}} - i^{\text{FOR}} + \% \text{ change in value of NZD}$$

If, in a particular period, the trader makes profits (losses) it means that borrowing in NZD was, *ex-post*, more (less) costly than borrowing in the foreign currency.

The UIP condition implies that the trader's expected profit is zero. This should be true in terms of predicting the profits one period ahead, but it should also be true about average profits, computed over a sufficiently long period of time. Carry traders, that is traders who borrow one currency to lend in another, presumably operate under the assumption that UIP does not hold. In the pre-GFC period, the proverbial Belgian dentists and Japanese housewives that invested in NZD-denominated securities presumably did so because they believed the positive interest differential between, say, NZD and Japanese yen did not reflect a likely depreciation of the NZD against the yen. We know that these proverbial carry traders took a hit with the onset of the GFC, but how have carry traders in the NZD fared over longer periods of time?

Consider Table 2. It shows the profitability of carry trades in the NZD when the USD, JPY and AUD are the foreign currencies. The results are, broadly speaking, consistent with what we would expect given the results from the regressions in the previous section. In all cases, carry trades profits were positive, on average. The results are weakest in the full inflation targeting sample (1990–2010), with mean payoffs and Sharpe ratios being insignificantly different from zero at the 5% level for all three currencies (although significant at the 10% level for the USD).[5] In the 1990–2007 subsample, mean payoffs and Sharpe ratios are statistically significant for the USD and JPY, but are insignificantly different from zero for the AUD.

1.3 Traditional Risk-Based Explanations

The UIP condition should only hold, and carry trade profits should only be zero, on average, if investors in the market are risk neutral. Therefore, a natural explanation for the failure, however weak, of the UIP condition, the profitability of carry trades in the NZD, and for the large interest differentials between NZD, USD and JPY, is that the NZD is exposed to risk that investors care about. When traders are risk averse, the UIP condition, (**) is replaced by

$$\text{expected } \% \text{ change in value of NZD} - \text{risk premium} = i^{\text{FOR}} - i^{\text{NZ}} \qquad \text{(R)}$$

The risk premium is positive if the NZD has a tendency to depreciate at times when investors are particularly sensitive to losses. This might, for example, be the case if the NZD tended to depreciate during global recessions or periods of financial turmoil.

Why consider the possibility that a standard risk premium is responsible for New Zealand's high interest rate? The answer is simple. A high interest rate might, indeed, be an explanation of New Zealand's low rate of investment, but if the this high

[5]The Sharpe ratio is the average payoff of the strategy divided by its standard deviation, and is a standard indicator of the profitability of a trading strategy.

interest rate reflects a risk premium, as traditionally defined, then that low rate of investment may, indeed, be "optimal". Put differently, the marginal product of capital in New Zealand may (for some yet to be described reason) need to remain high, to compensate investors in New Zealand for risk.

Conditional on somehow measuring it, a risk premium clearly has the potential to rationalize the results of the regressions in Table 1, since its existence implies that UIP fails. It also has the potential to rationalize the profitability of the carry trade in NZD, because it implies that these profits are simply compensation for investors' willingness to take on risk. And it is directly obvious, from (R), that it has the potential to explain the high relative level of interest rates in New Zealand.

While risk is a perfectly coherent and plausible explanation for the returns to carry trades, it fails in practice. Villanueva (2007), Burnside et al. (2011a), Burnside (2010) and Burnside, Eichenbaum and Rebelo (2011a) have extensively documented that the profits associated with portfolios of carry trades are approximately uncorrelated with conventional measures of risk, such as fluctuations in measures of U.S. consumption growth, measures of stock market performance, etc. This is true despite the recent experience during the global economic downturn. Their work confirms earlier results relevant to consumption-based measures of risk, for example Bekaert and Hodrick (1992) and Backus, Foresi and Telmer (2001). To the extent that some of these risk measures are correlated with carry trade profits, the degree of correlation is much too small to explain the returns to the carry trade. Burnside, Eichenbaum and Rebelo (2011a) and Menkhoff et al. (2011b) further document that unconventional measures of risk, such as those proposed by Lustig, Roussanov and Verdelhan (2009) and Menkhoff et al. (2011a), while capable of explaining carry trade returns, do not explain the profitability of momentum-based trading strategies in foreign exchange.

Rather than focus on portfolios of carry trades, suppose we consider carry trade returns that are specific to the New Zealand dollar. In Tables 3 and 4 I present results suggesting that, consistent with the broader literature, NZD-specific carry trade profits are uncorrelated with conventional measures of risk. Table 3 presents results of these regressions for carry trades in the NZD-USD pair. Table 4 presents equivalent results for carry trades in the NZD-JPY pair. The candidate risk factors are the U.S. stock market excess return, U.S. output growth, U.S. investment growth, and U.S. consumption growth. As is clear from the tables, when the GFC is excluded from the sample, there is no significant correlation between the carry trade profits and the risk factors, and sometimes the correlations have counterintuitive sign (for example, when US output growth is low, carry trade profits increase). When the GFC is included in the sample, there is some indication that the currency returns for the USD and JPY trades are exposed to stock market risk, but the estimated correlation is quantitatively small.[6]

Overall, simple risk-based explanations of the returns to carry trades in the NZD do not work. To the extent that these trades are exposed to observed market risk,

[6]Carry trades in the NZD have similar average returns to investments in stocks. The "beta" reported in Table 3 for the CAPM case would need to be close to 1 (i.e. six times larger) for the CAPM risk model to rationalize the profitability of these trades.

the amount of exposure is small, on average, and stems mainly from the GFC period. This highlights the importance of extreme events in explaining New Zealand's risk premium.

2 The Role of Extreme Events

In this section, I discuss how the possible occurrence of extreme events affect interest rates in New Zealand. Some of the discussion is technical ...

2.1 UIP and Extreme Events

As we saw in the previous section, the evidence against UIP is somewhat weak when we consider the full inflation targeting period. On the other hand, a researcher testing UIP prior to the GFC would have concluded that there was strong evidence against it. To understand why, consider Figure 2. It illustrates the 6-month moving average of carry trade profits against the USD and JPY over the full sample. The losses incurred at the height of the financial crisis (the second half of 2008) are noticeable in both graphs. Profits were sharply lower in this period against both the USD and JPY. In the case of the JPY, the impact on the overall sample mean is sufficient to make it statistically insignificant in the 1990–2010 period, even though the mean is still large and positive. In the case of the USD, losses during the peak of the crisis are offset by other gains in the 2008–10 period, but, as a consequence of the volatility experienced during the GFC, the sampling uncertainty associated with the mean payoff rises, and is sufficient to make the mean statistically insignificant.

If we consider the inflation targeting period (1990–2010), with just August through December 2008 excluded, carry profits against the USD and JPY are large and significantly positive: the annualized mean payoffs are 6.4 and 7.1%, respectively, with standard errors of 2.6 and 3.0%.[7] So the marginal significance of these profits over the full sample can be attributed almost entirely to a few months worth of observations in 2008. The mean profits against the AUD remain positive but small and statistically insignificant: 2.0% with a standard error of 1.6%.

Consistent with the story I tell below, this shows the dramatic effect that single extreme events have, not only on carry trade profits, but also on our inference about UIP. These results are also highly suggestive that extreme events play a crucial role in determining the NZD risk premium versus the USD and JPY. They also suggest that UIP vis-a-vis the AUD is a reasonable approximation, which, at first glance, suggests that the risk premium is not driven by idiosyncratic features of the New Zealand economy, at least not ones that are distinct from those of Australia.

[7] The annualized Sharpe ratios associated with these trades are also large and statistically significant: 0.60 and 0.52, respectively, with standard errors of 0.25 and 0.22.

2.2 Characterizing Extreme Events

Returning to the idea that extreme events play a crucial role in understanding the determination of New Zealand's risk premium, I introduce a simple framework that specifically highlights their role. In the technical appendix I formalize the ideas presented in this section with a mathematical model. Here, I provide only an intuitive discussion based on that model. Some of the discussion is technical, as needed.

The basic idea is that at any point in time, the economic outcomes that we observe can be characterized as either *normal* or *extreme*. Good and bad outcomes can occur in either case, but extreme events are distinct from normal ones in that they are (i) rare and (ii) the outcomes that occur in these events are distinct in some important way from normal events. These two assumptions have two important implications. The first assumption makes it plausible that a large sample of data might contain no observations corresponding to the extreme event, or, more generally, that extreme events could be under-represented in the data. The first implication and the second assumption, together, imply that sample averages of observed data may be quite different that the true averages associated with the underlying probability distributions.

My notion of extreme events is related to, but not as restrictive as, the notion of a *peso problem*. The original idea behind a peso problem is that during a period where a fixed exchange rate regime is sustained, one observes no changes in the exchange rate. However, one might well observe an interest differential between the relevant pair of currencies, due to traders' concern that the exchange rate regime might be abandoned. In any sample where the fixed exchange rate regime is not abandoned, the state of the world where it is abandoned is under-represented, and UIP will appear to fail, even if it holds. The distinction between peso problems and extreme events is simply that the former was introduced as an explanation of the failure of UIP under the assumption or risk neutrality.[8]

If we think of the GFC as an extreme event, the ways in which outcomes were different in this particular case were that, apparently, the risk tolerance of international investors declined sharply, and profits in the NZD carry trade were somewhat lower, on average, than during normal periods. Suppose that events such as the GFC are more likely to occur, or, at the very least, investors believe them to be more likely to occur, than the frequency with which have occurred in the recent past. In this case, the high cost of borrowing in New Zealand can be rationalized. If the expected change in the value of the NZD in an extreme event is negative, and if the risk premium attached to these losses is large, the left hand side of equation (R) can be negative, even if the extreme events have a small probability.

Is there evidence that investors' concerns about extreme events are "priced" into debts denominated in NZDs? Currency options provide some evidence that this is,

[8]My discussion of extreme events may also bring to mind the concept of a black swan. However, this concept is very different than what I have in mind. Typically, black swans are events that, before their occurrence are treated as outside the realm of possibility by the relevant set of individuals who experienced them. So though they share the property of being rare, and the property of having a large effects, they are, effectively ignored by individuals. In that respect, the two ideas are very different, as will become clear below.

indeed, the case. Figure 3 is derived from data on currency options for the NZD against the USD. Put options are the right to sell the NZD at a specific pre-set strike price measured in USD, so they provide insurance against downward movements in the value of the NZD. Call options are the right to buy the NZD at a specific pre-set strike price measured in USD, so they provide insurance against upward movements in the value of the NZD. The data shown in Figure 3, essentially, show that insurance against downward movements in the value of NZD is more expensive than insurance against equivalent upward movements in the value of the NZD. This is because the puts are more expensive than the equivalent calls.

To understand Figure 3 completely requires a slight digression into the way option prices are quoted. Rather than quoting the outright price for insurance in a one USD position in NZD, option prices are quoted in terms of "implied volatility". Not surprisingly, if the quoted implied volatility is larger, the outright price of an option is also larger. But prices of currency options also depend on how far the pre-set strike price is from the current forward exchange rate. Intuitively, the further away from the forward rate the strike price is, the cheaper the option. The implied volatility provides a way of quoting the price that factors out the second aspect of the price, under the assumption that the underlying variation in the exchange rate is normally distributed. Putting it differently, if the value of the exchange rate were normally distributed the implied volatility would be the same for all options. In practice, however, implied volatilities vary with the strike price of the option and this sheds light on how the underlying distribution of the exchange rate differs from a normal distribution.

Figure 3 plots the time series of differences between the implied volatilities of 25 and 10-delta options and the at-the-money-forward options between 1997 and 2007.[9] The jargon used here is not important. What is important is that the 25-delta put and call options have strike prices that are symmetrically below and above the current forward rate. Notice that the 25-delta put consistently has an implied volatility that exceeds the implied volatility of the 25-delta call (the blue line is almost always the red line). Similarly, the 10-delta put consistently has an implied volatility that exceeds the implied volatility of the 10-delta call (the green line is almost always the cyan line). This means that investors were more concerned about downward movements in the value of the NZD than they were about upward movements during this period. This is consistent with the extreme events story.

Figure 3 also suggests that investors are especially concerned with large downward movements in the value of the NZD. This can be seen in the fact that the 10-delta put, which has the lowest strike price, consistently has a higher implied volatility than the 25-delta put, which has the next lowest strike price (the green line is always above the blue line). Investors are concerned about large movements in the other direction as well, as evidenced by the fact that the 10-delta call, the option with the highest strike price, tends to have a higher implied volatility than the 25-delta call, the option with the next highest strike price (the cyan line usually lies above the

[9] The currency options data on which Figure 3 is based are only reliably available beginning in 1997. I cut the sample off in 2007 because this is arguably the end of the "normal" period and because the increased volatility that occurred during the GFC makes it hard to distinguish pre-crisis variation in the graph.

red line). But here the effect is much weaker.

In the technical appendix I describe in more detail how currency options data can be used to make more precise inference about the nature of extreme events. The model assumes that in each period investors believe that an extreme event could occur with some very small probability. In this case the NZD is expected to depreciate by some amount. On the other hand, if a normal event occurs (which is much more likely), and if the interest differential is positive, the NZD is expected to appreciate by some amount. Investors are assumed to be risk neutral with respect to normal events, but are assumed to be risk averse in extreme events. The model is calibrated to the currency options data in Figure 3, and takes into account the observed profitability of carry trades in the 1997–2007 period. The data suggest that extreme events, on average, consist of a 5% depreciation of the NZD within one month. If investors were risk neutral with respect to this low probability outcome, then extreme events would not be able to rationalize the typical size of the interest differential in the 1997–2007 period, which was about 2.6% on an annual basis. Consequently the data and model, together, imply that either (i) investors aversion to risk rises *very sharply* during the extreme event or (ii) investors assign a *much* higher probability to extreme events than the frequency with which they occur in the data. Theoretically it is impossible to distinguish between these two possibilities. However, intuitively (i) seems the more likely explanation since, otherwise, we should have observed more GFCs in sample.

Turning back to the question of what determines the risk premium in New Zealand, the conclusion one is tempted to draw from options prices is that investors are concerned about the realization of an extreme global event that would raise their aversion to risk dramatically above its normal level. This gets priced into the interest rate on nominally riskless NZD-denominated assets because the NZD is expected to depreciate if this global event occurs. But the size of the expected depreciation is relatively modest. The remaining question is whether the expected depreciation of the NZD is a New Zealand specific risk, or a more widespread risk shared by other countries. I address this question in the next section.

2.3 Australasian Risk or Idiosyncratic Risk?

In this section, I argue that New Zealand's risk premium is mainly an issue of Australasian risk, or at least, does not appear to be driven by New Zealand specific factors that are distinct from those driving the Australian risk premium. There are a number of pieces of evidence that point in this direction. The first of these is that in the inflation targeting period (1990–present), uncovered interest parity between the AUD and NZD appears to hold. We saw evidence of this in the UIP regressions reported in Table 1, both in the full 1990–2010 sample, as well as in pre-crisis (1990–2007) sample. Carry trade profits in the 1990–2010 period were also small, and statistically insignificant. In the 1990–2007 period, carry trade profits were positive, and significant at the 10% level, but I would argue that average profits in this period were small enough that they might easily be discounted on the basis of transactions costs.

Secondly, the average interest differential between the two currencies has been quite small during the inflation targeting period: 0.8% in the pre-crisis period and

0.7% in the full sample.

Third, the AUD-USD and NZD-USD interest differentials are highly correlated with one another, the correlation being 0.73 over the period 1990–2007, and 0.71 over the period 1990–2010.

Fourth, during the pre-crisis period the AUD-USD and AUD-JPY interest differentials are close to being as good predictors of the log changes in the NZD-USD and NZD-JPY exchange rates as are the NZD-USD and NZD-JPY interest differentials. If we rerun the UIP regressions of Table 1 with the AUD interest differentials in place of the NZD interest differentials, the coefficients are similar in magnitude and significance and the R^2s of the regressions are similar. Additionally, the sign of the AUD-USD interest differential is a better predictor of the sign of the change in the NZD-USD exchange rate than is the NZD-USD interest differential.

3 Time-Variation in the Risk Premium

In this section I consider time variation in the risk premium. That is, my main purpose of the section is to attempt to identify the ups and downs in the size of the risk premium. I discuss the theoretical and policy implications of this variation in Section 4. A discussion of variation in the risk premium necessarily involves building forecasting models, so, at times, this section, like the last one, uses technical jargon from statistics and options theory, although I try to explain the jargon as needed. A non-technical summary is provided at the end of the section.

3.1 Identifying the Risk Premium

As equation (R) indicates, if we want to identify the risk premium we need a forecasting model for the value of the NZD. Meese and Rogoff (1983) famously argued that the forecasting performance of various models of the exchange rate was poor relative to simply assuming that the exchange rate is a martingale, that is, that changes in the exchange rate are unforecastable. If this statement is literally correct, of course, then the risk premium, given in equation (R) is simply the interest differential, $i^{\mathrm{NZ}} - i^{\mathrm{FOR}}$. In this case, we could identify the variation in New Zealand's risk premium versus other economies simply by looking at Figure 1. More generally, if the regressions reported in Table 1 correctly identify the best forecasting model of the exchange rate, then the risk premium is not equal to the interest differential, but does move one-to-one with it. We could simply look to Figure 1, rescaled, in order to see the variation in the risk premium.

Our discussion in the previous section, however, suggested that extreme events play a role in determining the risk premium. We also saw hints that currency options data shed light on investors' expectations about the exchange rate. Indeed, in this section, I show that changes in the exchange rate are forecastable by a wider set of variables than the interest differential alone. In fact, I find that options prices contain information that is useful in forecasting the change in the exchange rate in sub-samples in which, arguably, there were no extreme events. I use information contained, not only in options on the NZD, but, also, in options data for other currencies. Unfortunately, these data are only available over a limited sample period (December 1998 to the present).

I use data on implied volatilities for the at-the-money-forward (ATMF), 25-delta put and call, and 10-delta put and call options for ten major currencies over the full sample period. The currencies are the Australian dollar, Canadian dollar, Swiss franc, Danish krone, euro, British pound, Japanese yen, Norwegian krone, New Zealand dollar, and Swedish krona, all of which are priced against the US dollar. I use three time series extracted from these data for each country:

- The implied volatility of ATMF options, as a general indicator of volatility. For currency i, I denote this variable as v_{it}. The ATMF options have a strike price equal to the forward exchange rate, i.e. the strike price that is most centrally located of the five options. Consequently it seems reasonable to use the implied volatility of these options as an indicator of the variance in the exchange rate anticipated by investors.

- An indicator of skewness in the distribution of the change of the exchange rate, given by the average of the implied volatilities of the two call options minus the average of the implied volatilities of the two put options. For currency i, I denote this variable as sk_{it}. Described in non-technical terms this variable captures the difference between the price of insurance against upward movement in the value of the NZD and downward movement. To the extent that investors are more concerned with downward movement, this variable will be negative.

- An indicator of kurtosis in the distribution of the change of the exchange rate, given by the average of the implied volatilities of the two 10-delta options, minus the implied volatility of the ATMF option. For currency i, I denote this variable as ku_{it}. Described in non-technical terms this variable takes on a larger value if investors become more concerned about tail risk in both directions.

I exclude the NZD from the data set, and use the remaining nine currencies to conduct a principal components analysis of the three time series. This analysis indicates a great deal of comovement across currencies.[10] I define global currency volatility, skewness (against the USD), and kurtosis factors using the first principal components of each group of series, and denote these v_t, sk_t and ku_t. I scale these factors so that they represent a weighted average of the underlying currency level series. They are illustrated in Figure 4 along with the country level series for New Zealand.

Figure 4 shows that the New Zealand based factors are highly correlated with the global factors despite the fact that the NZD options series were not included in

[10] In the case of the volatility indicator, I find that the first principal component has an eigenvalue more than 10 times larger than the second principal component. For the skewness indicator, I found that the first principal component has an eigenvalue more than twice as large as the second principal component, which, in turn, is six times as large as the third principal component. For the kurtosis indicator, I found that the first principal component has an eigenvalue more than 20 times larger than the second principal component.

the principal components analysis. This is true, not only during the financial crisis period, but also during the period of relative tranquility leading up to it.

It turns out that some of the options-based measures are useful in forecasting the change in the value of the NZD one month ahead. Here I let s_t denote the logarithm of the NZD-USD exchange rate in month t. Table 5 presents results for regressions of the change in the value of the exchange rate, $\Delta s_{t+1} = s_{t+1} - s_t$, on various sets of regressors in the period Dec. 1998–Dec. 2007, where time is measured in months, but the data are observed daily. The first regression is the traditional UIP regression, with Δs_{t+1} regressed on a constant and the interest differential, $i_t^{\mathrm{US}} - i_t^{\mathrm{NZ}}$:

$$\Delta s_{t+1} = \beta_0 + \beta_i \left(i_t^{\mathrm{US}} - i_t^{\mathrm{NZ}}\right) + \epsilon_{t+1}. \tag{1}$$

Consistent with the results in Table 1, the estimated slope coefficient in this regression, $\hat{\beta}_i$, is negative, and is highly statistically significant. Once again, we have the usual UIP puzzle, that when the interest differential between NZD and USD is positive (or strictly speaking, above average), the NZD is actually expected to appreciate (or strictly speaking, has an above average rate of appreciation). The R^2 of the regression, however, is quite modest (0.032).

The second regression projects Δs_{t+1} on a constant, and the three global measures, v_t, sk_t, ku_t:

$$\Delta s_{t+1} = \beta_0 + \beta_v v_t + \beta_s sk_t + \beta_k ku_t + \epsilon_{t+1}. \tag{2}$$

The estimated coefficients on volatility and kurtosis are statistically significant. The signs of the estimated coefficients on those variables ($\hat{\beta}_v < 0$, $\hat{\beta}_k > 0$) indicate that in times of higher global currency volatility, and lower global kurtosis the NZD is expected to perform worse than average. The first result is somewhat counterintuitive in that we might expect that higher expected returns in normal times would act as compensation for the risks imposed by volatility. Interestingly, however, this regression has more explanatory power than the usual UIP regression, the R^2 being 0.035.

The third regression projects Δs_{t+1} on a constant, the three global measures, v_t, sk_t, ku_t, and the part of the New Zealand specific measures, v_t^{NZ}, sk_t^{NZ}, ku_t^{NZ}, that is orthogonal to v_t, sk_t, and ku_t. I denote the latter variables $\hat{v}_t^{NZ}$, $\widehat{sk}_t^{NZ}$, and $\widehat{ku}_t^{NZ}$, so that the regression is

$$\Delta s_{t+1} = \beta_0 + \beta_v^G v_t + \beta_s^G sk_t + \beta_k^G ku_t + \beta_v^{NZ} \hat{v}_t^{NZ} + \beta_s^{NZ} \widehat{sk}_t^{NZ} + \beta_k^{NZ} \widehat{ku}_t^{NZ} + \epsilon_{t+1}. \tag{3}$$

The estimated coefficients on the global variables, and their statistical significance, remain roughly the same. The coefficient on New Zealand specific skewness, β_s^{NZ}, is significant, and, intuitively, negative. The R^2 of the regression is 0.070.

The last regression adds the interest differential to (3):

$$\Delta s_{t+1} = \beta_0 + \beta_i \left(i_t^{\mathrm{US}} - i_t^{\mathrm{NZ}}\right) + \beta_v^G v_t + \beta_s^G sk_t + \beta_k^G ku_t + \beta_v^{NZ} \hat{v}_t^{NZ} + \beta_s^{NZ} \widehat{sk}_t^{NZ} + \beta_k^{NZ} \widehat{ku}_t^{NZ} + \epsilon_{t+1}. \tag{4}$$

In this regression, as in the earlier ones, the coefficients on the interest differential, global volatility, global kurtosis and New Zealand specific skewness are statistically significant. Only the global volatility measure has a counterintuitive (negative) sign, indicating, surprisingly, that the NZD is a hedge against persistent changes in global currency volatility. The R^2 of this regression is quite impressive (0.098) even if there is a large number of explanatory variables. Figure 5 illustrates the improved in-sample forecasts generated by (4) relative to (1). As is easily seen in the graph, the forecast generated by (4) tracks the overall fluctuations of the exchange rate much better than the model based purely on the interest differential, (1).

Do these regressions provide insight into fluctuations in the risk premium? They do, as long as we are willing to make some assumptions along the lines of the model outlined in the technical appendix, and if we are willing to assume that the period Dec. 1998–Dec. 2007 is representative of *normal* states of the world. If the latter assumption is valid then the fitted values from the regressions, say, in the case of regression (4),

$$\Delta\hat{s}_{t+1} = \hat{\beta}_0 + \hat{\beta}_i\left(i_t^{\mathrm{US}} - i_t^{\mathrm{NZ}}\right) + \hat{\beta}_v^G v_t + \hat{\beta}_s^G sk_t + \hat{\beta}_k^G ku_t + \hat{\beta}_v^{NZ}\hat{v}_t^{NZ} + \hat{\beta}_s^{NZ}\widehat{sk}_t^{NZ} + \hat{\beta}_k^{NZ}\widehat{ku}_t^{NZ},$$

are estimates of the expected change in the exchange rate conditional on the next period being a normal event. Furthermore, as the appendix shows, the estimated risk premium is proportional to $\Delta\hat{s}_{t+1} + i_t^{\mathrm{US}} - i_t^{\mathrm{NZ}}$.

3.2 Performance of the Model During the Crisis

Perhaps not surprisingly, the model does quite poorly if it is fit over the full-sample (the second-last column of Table 5). New Zealand-specific skewness remains statistically significant, and global kurtosis remains marginally significant, but the other variables lose their statistical significance. This result isn't that surprising because the model is aimed at identifying the risk premium in normal times. When a crisis period is entered, a large forecast error will be generated because the change in the exchange rate is drawn from the extreme distribution not the normal distribution. Unfortunately, there simply isn't enough data available from the crisis period to estimate an appropriate Markov-switching model of the risk premium through the full sample. It is interesting to note, however, that the model fit over the period 1998–2007 does quite well at tracking observed changes in the exchange rate in the post-crisis period (late 2009–present). This can be seen in Figure 6 which compares the in-sample fit of the full-sample model to the out-of-sample fit of the 1998–2007 model in the crisis period.

3.3 Australasian Risk Once More

In this section, I argue that while two New Zealand-specific variables (the interest differential and the options-based skewness measure) are useful in forecasting the exchange rate during normal times, and therefore seem to play a role in determining the risk premium, Australia-specific variables are just as useful. This is demonstrated in Table 5 which re-runs regression (4) with Australia-specific variables replacing the New Zealand-specific variables. The signs and statistical significance of the various estimated coefficients remains roughly unchanged, and the R^2 of the regression is actually marginally larger. There are only small differences between the forecasts generated by this model and the forecasts generated by the New Zealand-specific regression.

3.4 Summary

In this section, I built a forecasting model of the NZD-USD exchange rate. I used this model, which used information extracted from currency options data, to identify variation in the risk premium on short-term NZD securities. The model demonstrates modest, but statistically significant, improvement in forecasting performance relative to a model based solely on the size of the interest differential. The model confirms two conclusions I drew in Section 2. First, variation in New Zealand's risk premium appears to be driven by Australasian factors rather than New Zealand specific factors. The evidence for this comes from the fact that when New Zealand-specific variables are replaced by Australian-specific variables, the forecasting ability of the model does not deteriorate. Second, there is evidence that extreme events have important effects on the risk premium. Currency options data suggest that investors are concerned about downside risk when holding long positions in the NZD, and that their degree of risk tolerance varies over time. Variation in indicators of these concerns helps predict the exchange rate, and therefore, appears to affect the risk premium.

4 Macro Policy and the Risk Premium

4.1 Monetary Policy

To this point I have said little about the actions of the Reserve Bank of New Zealand and how they may interact with the behavior of the risk premium. Obviously the Reserve Bank has an important role to play, since the Official Cash Rate (OCR), the interest rate set by the Reserve Bank to meet its inflation target, is closely linked to short term interest rates. For the purposes of this section, I take the one-month interest rate and the policy rate to be one and the same thing, so that, taking the foreign interest rate as given, and independent of developments in the New Zealand economy, changes in the policy rate directly determine the interest differential.

If uncovered interest parity held, in its simplest form, central banks could manipulate short term interest rates to their hearts content, but in so doing their policy maneuverings would simply change the rates of expected depreciation across currencies. The most naive example of this can be seen in a frictionless monetary model, where to achieve a higher interest rate, a central bank would need to print money at a faster rate, thus inducing a higher rate of expected depreciation of the local currency. In this type of framework, because of a lack of frictions, the central bank's actions would have little or no impact on the real side of the economy, including on capital flows, and the real exchange rate. The models have capital flows in them, but these are determined entirely on the real side of the model, and are driven by the demand of the local economy for consumption and investment, relative to the domestic supply of goods, with the supply of funds being elastic at the world interest rate.

4.1.1 One View

One interpretation of the models I have outlined in the previous section is that they are risk-adjusted versions of the same frictionless story. If the RBNZ raises short term interest rates, say because of domestic inflationary pressure, this does not imply higher risk-adjusted expected returns for investors in NZD-denominated assets who are financing these with foreign currency. Symmetrically, it does not imply higher risk-adjusted borrowing costs for New Zealand based borrowers, at least not under the assumption that they have the same risk profile as foreign investors (an important caveat). This is because equation (R) always holds. Assuming that foreign variables are independent of the actions of the RBNZ, any action on the part of the RBNZ to, say, raise rates, must either decrease the rate of expected appreciation of the NZD (or increase its expected rate of depreciation), or it must increase the risk premium. The latter would come about through an increased tendency of the NZD to depreciate if times are "bad" for the investor.

In the risk-based model, it remains true that capital flows are determined by the demand for consumption and investment relative to the local supply of goods, and the fact that there is an elastic supply of funds at the global interest rate. Since the risk-adjusted rate of return to NZD-denominated speculative investments is always zero, these play no role in the determination of capital flows.

Put differently, conventional asset pricing theory says it doesn't matter if firms contemplating real investments in New Zealand finance these investments with low interest rate loans in foreign currency, or high interest rate loans in NZD. Adjusted for risk, these loans are equally costly. The NZD risk can be borne by the foreign lender, in which case the New Zealand based firm is hedged by borrowing in NZD. Or the firm can borrow in foreign currency, in which case it bears the foreign exchange risk. Conventional theory would also argue that any *monetary* policy aimed at generally lower NZD interest rates would be irrelevant, because it would not affect risk-adjusted cost of borrowing.

4.1.2 An Alternative View

An alternative view of small open economies is that the central bank's policy decisions play an important role in determining capital flows because of frictions in goods markets, labor markets and financial markets. One such view is articulated, at least informally, by Moutot and Vitale (2009). They imagine the following scenario. Suppose a central bank in an economically significant country maintains low interest rates, which leads to credit creation, and an expanding pool of liquidity. In this scenario, investors may be attracted to a higher interest rate, higher expected return economy. Moutot and Vitale describe the subsequent outcome as follows

> To the extent that the recipient country is a credible inflation fighter, its bilateral exchange rate will tend to appreciate because the capital inflows will increase its asset prices. As a consequence, wealth effects will increase consumption growth while inflation remains subdued, at least for some time. In such circumstances, investors in the originating country will see their expectation of higher returns validated and may be encouraged to continue pouring capital into the high-yielding assets.

Moutot and Vitale argue that this story describes well development in the pre-crisis period in the carry trade between the JPY and the NZD and AUD.

A more formal, but similar, articulation of the alternate view is provided by Plantin and Shin (2011). In Plantin and Shin's model there is a feedback loop between interest rates and capital flows. High interest rates (presumably relative to those of a large economy like Japan) lead to capital inflows. These capital inflows tend to overheat the local economy leading to inflationary pressure. The inflationary pressure leads to interest rate increases by an inflation targeting central bank. This leads to more capital inflows, and so on. Carry trade dynamics are drawn out in Plantin and Shin's model because traders rebalance their portfolios infrequently, according to a reduced form Baumol-Tobin type of rule. Because the local currency ends up being overvalued in this scenario, there is downside (imported inflation) risk that goes against continuation of the carry trade dynamics. On the other hand, monetary policy keeps them going. The result is a potential for what Plantin and Shin describe as "bubbly" equilibria, in which the exchange rate may appreciate for a while (go up the stairs) and then revert sharply towards its fundamental value (go down the elevator). The tighter the financial market friction, and the stronger the central bank's response to inflationary pressure, the more likely are realized paths with occasional sharp depreciations of the recipient economy's currency.

4.1.3 Which View is More Relevant?

Distinguishing the two stories, one in which there are no frictions, and the other in which frictions are "everything", is important from a policy perspective, but doing so using data is non-trivial, and neither model provides a fully articulated macroeconomic framework. For example, in both cases, increases in a central bank's policy rate will alter the dynamics of the future exchange rate, potentially leading to short term appreciation, while conditionally skewing the exchange rate more negatively. In the extreme risk model, an increase in rates can raise the average rate of appreciation in normal times, as long as there is a larger risk-adjusted increase in the average rate of depreciation in extreme times. In Plantin and Shin's model the short-term appreciation takes place gradually due to the financial friction, and this leads to more fragility in the carry trade dynamics, with increased likelihood of sharp reversion towards the long-run fundamental value of the exchange rate.

It is beyond the scope of this paper to develop a formal test that can distinguish between the two stories. Some facts in the data are, however, suggestive. First, first-differences of the OCR (and quasi-differences with persistence parameters greater than 0.84) are actually negatively correlated with fluctuations in the expected rate of appreciation generated from the forecasting model outlined in Section 3 [specifically, regression model (4)]. This is not what we would expect given the Plantin and Shin story. There, positive innovations in the OCR would presumably be expected to lead to a gradual upward adjustment in the value of the NZD, which would be anticipated by investors and would show up in our measure of expected appreciation. Second, as Table 6 indicates, the realized skewness of changes in the value of the NZD during the pre-crisis period was quite modest. In Plantin and Shin's story, trips down the elevator happen, not only during global crises, but also during relatively stable periods of time, with potentially idiosyncratic drops in the values of some currencies.

Somewhat in favor of Plantin and Shin's story, however, is the more general cross-country evidence presented in the May 2008 working paper version of Burnside et al. (2011a). In a sample of 31 years and up to 20 currencies, they show that the payoffs to carry trades display increasingly large negative realized skewness as the absolute value of the interest differential between currencies increases.

4.1.4 Would a Different Approach to Monetary Policy be Advisable?

To fully answer this question, we need a fully articulated model, in which the question of optimal monetary policy can be addressed. To my knowledge an off-the-shelf macroeconomic model that accounts for rare and extreme events is not currently available for this purpose.

Most existing work on small open economies suggests that there is some scope for alternatives to the RBNZ's approach to monetary policy. For example, several articles in the literature have suggested that the appropriate target for the central bank in a small open economy is to target either the inflation rate of domestically-produced goods (e.g., Gali and Monicelli, 2005), or, in some circumstances, the inflation rate of nontraded goods, as opposed to the overall consumer price index (e.g., Devereux, Lane and Xu, 2006). In some of this work, the appropriate target is not dependent on which shocks appear in the model, so the fact that rare event shocks aren't in the model wouldn't change the implications of the theory for the appropriate target. Were the RBNZ to have focused more narrowly on alternative domestic price indices, however, its interest rate policy would have been more aggressive (not less so) during the period between 2001 and 2006 when NZD rates rose while USD rates fell and remained low. A less aggressive approach, even if it were to have eliminated some speculative dynamics, would not, according to standard theory have been desirable.

I am skeptical that monetary policy can play a big role in solving the low investment problem through the cost of borrowing channel. From my perspective, even if asset market frictions are important, as in, for example, the Plantin and Shin view, the importance of these frictions is only centrally relevant when deciding whether the RBNZ targets the right measure of inflation and whether it does so with the right amount of aggression. In other words, we can legitimately ask whether the RBNZ's actions are appropriately stabilizing, or whether they add new self-fulfilling "shocks" to the economy.

4.2 Other Policy Instruments

Overall, the evidence I have presented suggests that while New Zealand's interest rates have been higher than those in other countries, this may be rationalizable in terms of New Zealand's exposure to extreme events such as the GFC. I have also suggested, in the previous section, that there is no glaring case for the RBNZ to adopt an alternative policy stance in order to affect the risk premium.

What then for New Zealand? Does this mean that the level of investment in New Zealand is optimal, and that there is no risk stemming from the country's overall external position? Does this mean the low growth New Zealand has experienced is also optimal? Not at all. My analysis, instead, leads to the conclusion that the risk premium is not the first place we should be looking for answers.

Coming back to the overall macroeconomic picture some good news should be discussed first. The good news comes in two parts. In my introduction I stated that New Zealand had experienced the slowest growth in a group of OECD countries, and had slipped 18 places in ranking by per capita income within that group between 1960 and 2009. What I did not mention was that the most of the damage, if you like, occurred prior to 1990. New Zealand's ranking by income has actually increased by one place since 1990, and its growth (ranked 9th out of 25) has been above average for the OECD. Comparisons with Australia are perhaps inevitable, but it should be kept in mind that Australia has been one of the top performers.

Second, New Zealand is in the fortunate position that a very substantial portion of its external debt is denominated in domestic currency. This means that the country is naturally hedged against the risk that I have argued is the source of its large risk premium.

The bad news, in the end, lies in the sheer size of the external position, and, perhaps, in the speed of New Zealand's growth, given its location in the pacific rim. It simply isn't sustainable for a country with a large net debt position to continue to run large current account deficits and see its debt stock relative to GDP grow by 8% per decade. The basic calculus of the external position states that the sustainable size of the current account deficit is given by $(r - g) \times nfa$, where r is the real financing cost faced by the economy, g is the growth rate of real GDP and nfa is the country's net external asset position. The only case for a country being able to sustain a current account deficit and a large negative asset position at the same time is if its real financing cost is exceeded by its growth rate. While I have argued that New Zealand's interest rates are, in a sense, not radically out of line with the rest of the world, even once risk corrections are taken into account, it would be hard to make the case that $g > r$ given the country's modest growth rate.

Consequently, policies aimed at reducing New Zealand's vulnerability to an external financing crisis are well advised. Policies aimed at increasing domestic savings are one way to go, and New Zealand has obviously taken several steps in this direction in the last decade, in terms of both government saving as well as private sector saving. This would be particularly advised if, in order to address growth concerns, the government were to undertake policies to stimulate higher levels of investment spending.

Additionally, one has to wonder why New Zealand is borrowing so much from abroad. Partly this is a function of the globalization of financial markets. New Zealand's debts to foreigners have increased by roughly 20 percent of GDP since 2000, but its claims on foreigners have also increased by over 10 percent of GDP. This continued a trend begun in the 1990s. But why is net borrowing increasing?

Standard models imply that strong anticipated productivity growth (more generally, strong anticipated income growth) would lead to a negative current account balance. Investment demand would be strong because high levels of forecasted productivity growth would raise the perceived marginal product of capital. Consumption demand would be strong if anticipated income levels were substantially above current income levels. Given strong investment and consumption demand, financing would be sought from abroad. But in New Zealand's case, the first factor seems to be absent. Investment demand is not particular strong. So we must turn to consumption demand. Why is it so strong? If productivity growth isn't driving up the level of investment, where are the income increases anticipated by households going to come from?

The only obvious thing I can point to is the housing sector, where prices rose substantially in the past decade, and, while moderating, remain substantially elevated. While fundamental reasons for this can be pointed to, fundamentals have a hard time explaining why prices remain high relative to rents. If households continue to believe that house prices will rise, and this is where a substantial portion of New Zealanders' savings lie, then consumption demand will remain strong. My own research on housing prices and expectations (Burnside, Eichenbaum, and Rebelo, 2011b) suggests that it is difficult to know whether a boom in house prices is sustainable. Some booms end in busts. Others end in prices reaching plateaus that are sustained. This doesn't mean policy makers should stand by and wait for a property price crash to happen (or not). Policy makers who are concerned with the economy's vulnerability should implement policies that tend to insure the economy, even if the insurance does not eventually pay off. Towards this end, as far as property is concerned, policies that remove distortions favoring residential investment over other forms of investment seem advisable.

5 Conclusion

In this paper I have provided evidence that New Zealand's interest rates are substantially higher than those in the US and Japan, largely because of extreme event risk. Australia is also exposed to this risk, and interest rates in Australia and New Zealand are sufficiently similar that uncovered interest parity between the two countries is a reasonable approximation. This casts doubt on the notion that high interest rates are the channel through which New Zealand has a low level of investment and a low level of growth. I also find evidence suggesting that RBNZ interest rate policy does not seem to drive changes in investor expectations about the value of the NZD.

I conclude that the central issue facing New Zealand is its substantial negative external position. Even if this position is largely naturally hedged through its currency denomination, it does not seem consistent with long run sustainability. I argue, therefore, in favor of a balanced approach in which the New Zealand government continues policies aimed at increasing domestic saving, while moving away from policies that favor residential over other forms of investment.

REFERENCES

Backus, David K., Silvio Foresi, and Christopher I. Telmer (2001) "Affine Term Structure Models and the Forward Premium Anomaly," *Journal of Finance 5*6, 279–304.

Bekaert, Geert, and Robert J. Hodrick (1992) "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets," *Journal of Finance* 47, 467–509.

Burnside, Craig (2011) "Carry Trades and Risk," forthcoming in *The Handbook of Exchange Rates*, J. James, I.W. Marsh and L. Sarno, eds. Hoboken, NJ: Wiley.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo (2011a) "Do Peso Problems Explain the Returns to the Carry Trade?" *Review of Financial.Studies* 24, 853–91.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2011a) "Carry Trade and Momentum in Currency Markets," submitted, *Annual Reviews of Financial Economics*, doi: 10.1146/annurev-financial-102710-144913.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2011b) "Understanding Booms and Busts in Housing Markets," NBER Working Paper 16734.

Burnside, Craig, Bing Han, David Hirshleifer and Tracy Y. Wang (2011b) "Investor Overconfidence and the Forward Premium Puzzle," *Review of Economic Studies* 78, 523–58.

Chinn, Menzie and Guy Meredith (2004) "Monetary Policy and Long-Horizon Uncovered Interest Rate Parity," *IMF Staff Papers* 51, 409–430.

Devereux, Michael B., Philip R. Lane and Juanyi Xu (2006) "Exchange Rates and Monetary Policy in Emerging Market Economies," *Economic Journal* 116: 478–506.

Engel, Charles (1996) "The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence," *Journal of Empirical Finance* 3, 123–192.

Fama, Eugene (1984) "Forward and Spot Exchange Rates," *Journal of Monetary Economics* 14, 319–338.

Frankel, Jeffrey and Menzie Chinn (1993) "Exchange Rate Expectations and the Risk Premium: Tests for a Cross-Section of 17 Currencies," *Review of International Economics* 1, 136–144.

Galí, Jordi and Tommaso Monacelli (2005) "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 72: 707-34.

Hodrick, Robert J. (1987) *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Chur, Switzerland: Harwood Academic Publishers.

Jurek, Jakub W. (2008) “Crash-Neutral Currency Carry Trades,” SSRN Paper 1262934.

Lewis, Karen (1995) “Puzzles in International Financial Markets,” in *Handbook of International Economics*, G. Grossman and K. Rogoff, eds., 1913–1972. Amsterdam: Elsevier.

Lustig, Hanno, Nick Roussanov, and Adrien Verdelhan (2009) “Common Risk Factors in Currency Markets,” SSRN Paper 1139447.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf (2011a) “Carry Trades and Global Foreign Exchange Volatility,” Forthcoming, *Journal of Finance.*

Meese, Richard A. and Kenneth Rogoff (1983) Empirical Exchange Rate Models of the Seventies : Do They Fit Out of Sample? *Journal of International Economics* 14, 3-24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf (2011b) “Currency Momentum Strategies,” mimeo, Cass Business School, City University, London.

Moutot, Philippe and Giovanni Vitale (2009) “Monetary Policy Strategy in a Global Environment,” European Central Bank Occasional Paper Series No. 106.

Nakamura, Emi, Jon Steinsson, Robert J. Barro, and Jose F. Ursúa (2010) “Crises and Recoveries in an Empirical Model of Consumption Disasters,” National Bureau of Economic Research Working Paper No. 15920.

Plantin, Guillaume, and Hyun Song Shin (2011) “Carry Trades, Monetary Policy and Speculative Dynamics.” Mimeo, Princeton University.

Villanueva, O. Miguel (2007) “Forecasting Currency Excess Returns: Can the Forward Bias be Exploited?” *Journal of Financial and Quantitative Analysis* 42, 963–90.

TABLE 1: **UIP Regressions for the New Zealand Dollar Against Major Currencies**

	1985–2010			1990–2010			1990–2007		
Foreign currency	$\beta_0 \times 100$	β_1	R^2	$\beta_0 \times 100$	β_1	R^2	$\beta_0 \times 100$	β_1	R^2
USD	0.21	-0.99	0.012	-0.09	-0.09	0.000	0.52	-2.81	0.021
	(0.26)	(0.39)		(0.50)	(2.12)		(0.37)	(1.34)	
JPY	0.04	0.27	0.001	0.17	-0.10	0.000	1.59	-3.30	0.026
	(0.36)	(0.47)		(0.92)	(2.11)		(0.53)	(1.10)	
AUD	0.07	-1.36	0.026	-0.01	0.16	0.000	-0.10	0.39	0.001
	(0.12)	(0.40)		(0.11)	(0.81)		(0.13)	(0.90)	

Note: The table reports estimates of the equation $\Delta s_{t+1} = \beta_0 + \beta_1 (f_t - s_t) + \epsilon_{t+1}$ where s_t and f_t are the logarithms of the spot and forward exchange rates measured as FCUs per NZD. The forward premium, $f_t - s_t$, is approximately equal to the interest rate differential $i_t^* - i_t$, where i_t^* is the foreign interest rate and i_t is the NZ interest rate. The data are monthly. Heteroskedasticity consistent standard errors are in parentheses.

TABLE 2: **Carry Trade Profits for the New Zealand Dollar Against Major Currencies**

	1985–2010			1990–2010			1990–2007		
Foreign currency	Mean	SD	SR	Mean	SD	SR	Mean	SD	SR
USD	7.7	12.4	0.62	5.4	11.3	0.48	5.3	9.4	0.56
	(2.8)	(1.1)	(0.23)	(2.9)	(0.9)	(0.27)	(2.6)	(0.7)	(0.28)
JPY	6.0	15.0	0.40	5.4	14.3	0.37	6.7	12.3	0.54
	(3.0)	(1.1)	(0.21)	(3.4)	(1.3)	(0.25)	(3.1)	(0.7)	(0.26)
AUD	4.3	9.7	0.44	2.3	7.3	0.32	3.0	6.8	0.44
	(2.0)	(0.9)	(0.20)	(1.5)	(0.4)	(0.21)	(1.6)	(0.4)	(0.24)

Note: Trades are executed monthly. SD indicates standard deviation. SR indicates the Sharpe Ratio, the ratio of the mean and standard deviation. Heteroskedasticity consistent standard errors are in parentheses.

TABLE 3: **Risk Exposures of NZD-USD Carry Trades**

	Jan. 1985–Jun. 2008			Jan 1985–Dec. 2010		
Factor	a	β	R^2	a	β	R^2
Mkt-Rf (CAPM)	0.006	0.03	0.001	0.005	0.16*	0.043
	(0.002)	(0.06)		(0.002)	(0.06)	
U.S. GDP growth	0.034	-0.58	0.003	0.025	1.07	0.011
	(0.008)	(1.30)		(0.008)	(1.28)	
U.S. investment growth	0.032	-0.10	0.003	0.030	0.06	0.001
	(0.006)	(0.19)		(0.006)	(0.18)	
U.S. consumption growth	0.027	0.96	0.003	0.022	1.90	0.015
	(0.008)	(1.72)		(0.010)	(1.93)	

Note: The table reports estimates of the equation $z_t^C = a + f_t'\beta + \epsilon_{t+1}$, where z_t is either the monthly excess return, or the quarterly excess return, of a carry-trade position that is long (short) the NZD and short (long) the USD if the New Zealand interest rate is higher than the U.S. interest rate, and f_t is a risk factor. The CAPM factor is the excess return on the value-weighted US stock market (Mkt-Rf). The other factors are real per capita log growth rates at the quarterly frequency. Heteroskedasticity consistent standard errors are in parentheses. Slope coefficients that are statistically significant at the 5 percent level are indicated by an asterisk (*).

TABLE 4: **Risk Exposures of NZD-JPY Carry Trades**

	Jan. 1985–Jun. 2008			Jan 1985–Dec. 2010		
Factor	a	β	R^2	a	β	R^2
Mkt-Rf (CAPM)	0.005	0.12	0.017	0.004	0.24*	0.066
	(0.002)	(0.07)		(0.002)	(0.07)	
U.S. GDP growth	0.033	-0.58	0.002	0.020	1.53	0.015
	(0.010)	(1.49)		(0.011)	(1.87)	
U.S. investment growth	0.031	-0.36	0.021	0.026	-0.08	0.001
	(0.007)	(0.25)		(0.008)	(0.30)	
U.S. consumption growth	0.024	1.17	0.003	0.016	2.56	0.017
	(0.012)	(1.85)		(0.012)	(2.13)	

Note: The table reports estimates of the equation $z_t^C = a + f_t'\beta + \epsilon_{t+1}$, where z_t is either the monthly excess return, or the quarterly excess return, of a carry-trade position that is long (short) the NZD and short (long) the JPY if the New Zealand interest rate is higher than the Japanese interest rate, and f_t is a risk factor. The CAPM factor is the excess return on the value-weighted US stock market (Mkt-Rf). The other factors are real per capita log growth rates at the quarterly frequency. Heteroskedasticity consistent standard errors are in parentheses. Slope coefficients that are statistically significant at the 5 percent level are indicated by an asterisk (*).

TABLE 5: **Forecasting the Change in the NZD-USD Exchange Rate**

	Sample Period					
	(1998–2007)					(1998–2011)
Regressor	(a)	(b)	(c)	(d)	(e)	(f)
N.Z. interest differential	-3.98*			-5.11*		0.05
	(1.77)			(2.05)		(2.06)
Global volatility		-0.49*	-0.51*	-0.43*	-0.71*	-0.26
		(0.25)	(0.26)	(0.18)	(0.32)	(0.19)
Global skewness		0.26	0.03	-0.25	-0.44	0.17
		(0.39)	(0.41)	(0.40)	(0.44)	(0.26)
Global kurtosis		2.74*	2.86*	4.75*	3.85*	2.99†
		(1.29)	(1.45)	(1.55)	(1.65)	(1.71)
N.Z. volatility			0.01	0.18		0.21
			(0.21)	(0.20)		(0.16)
N.Z. skewness			-1.26*	-0.99†		-1.25*
			(0.57)	(0.51)		(0.47)
N.Z. kurtosis			-1.73	-0.98		0.07
			(2.16)	(2.03)		(1.50)
Aus. interest differential					-5.33†	
					(2.74)	
Australian volatility					-0.14	
					(0.26)	
Australian skewness					-1.29*	
					(0.58)	
Australian kurtosis					-0.56	
					(2.97)	
R^2	0.032	0.035	0.070	0.098	0.100	0.056

Note: The data are daily, from December 10 1998 to December 31 2007 (or February 28 2011). If the regressors are measured on date t, the one-month change in the exchange rate (the lefthand side variable) is measured as Δs_{t+d} where day $t+d$ is the business day closest to, but less than or equal to, 30 days ahead. Newey-West heteroskedasticity consistent standard errors are in parentheses, using a lag window of 40 business days.

TABLE 6: **Skewness of New Zealand Dollar Exchange Rate Changes**

Foreign currency	1985–2010	1990–2010	1990–2007
USD	-0.18	-0.15	-0.25
	(0.27)	(0.40)	(0.16)
JPY	-0.54	-0.48	-0.28
	(0.23)	(0.33)	(0.19)
AUD	0.25	0.02	0.01
	(0.48)	(0.16)	(0.20)

Note: The table reports the skewness of Δs_{t+1} where s_t is the logarithm of the spot exchange rates measured as FCUs per NZD. The data are observed monthly. Heteroskedasticity consistent standard errors are in parentheses.

FIGURE 1

Growth and Income Levels in the OECD Countries

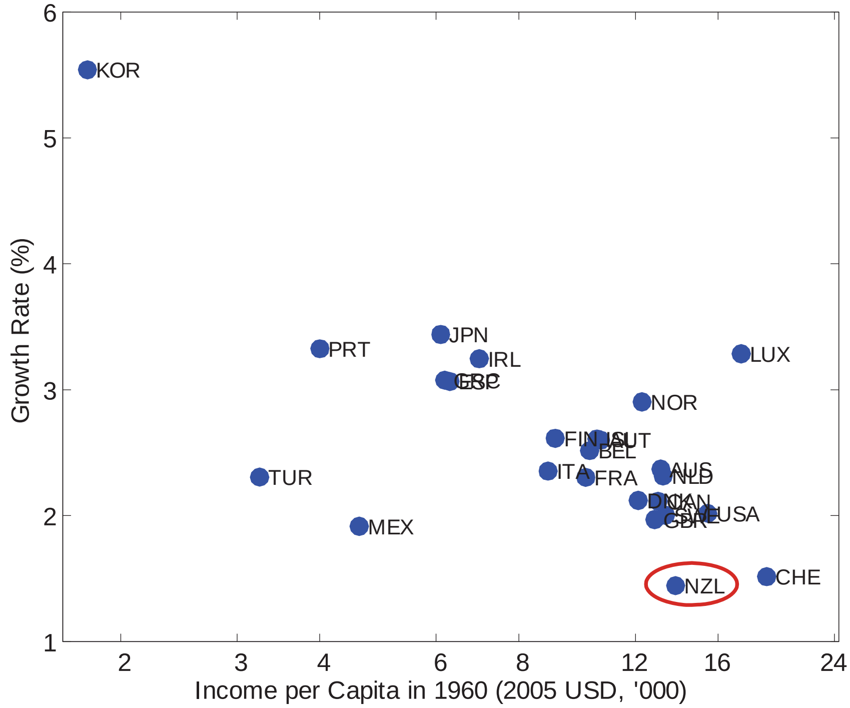

Source: Penn World Table.

FIGURE 2

One-Month Interest Rate Differentials Between New Zealand Dollars and Major Currencies

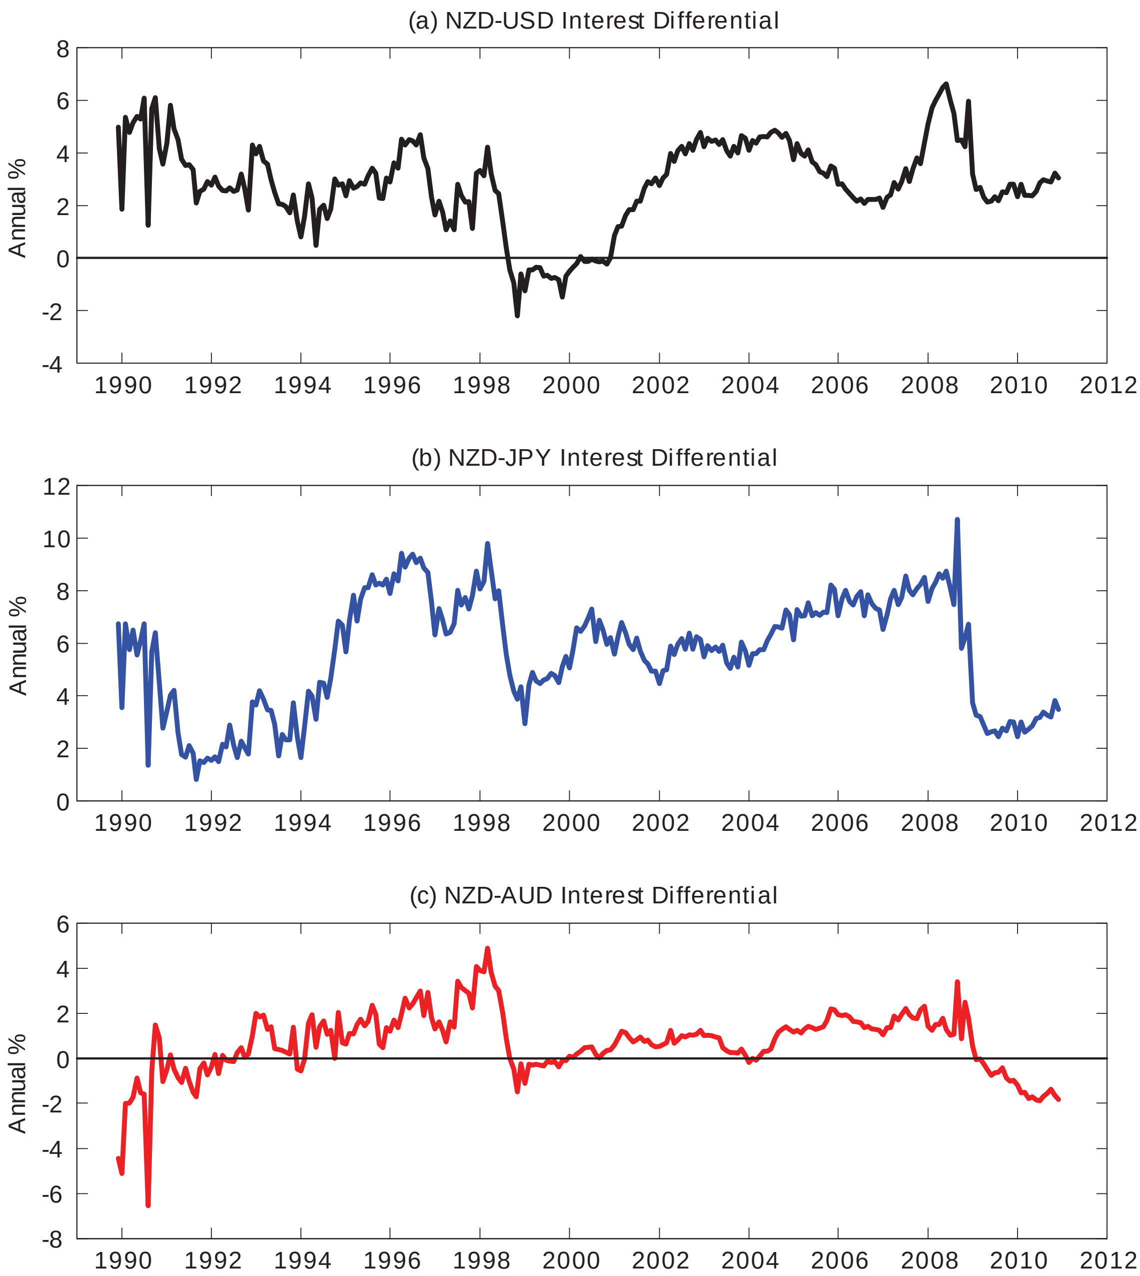

Note: End of month values, December 1989 to December 2010.

FIGURE 3

Carry Trade Profits Between New Zealand Dollars, U.S. Dollars and Japanese Yen

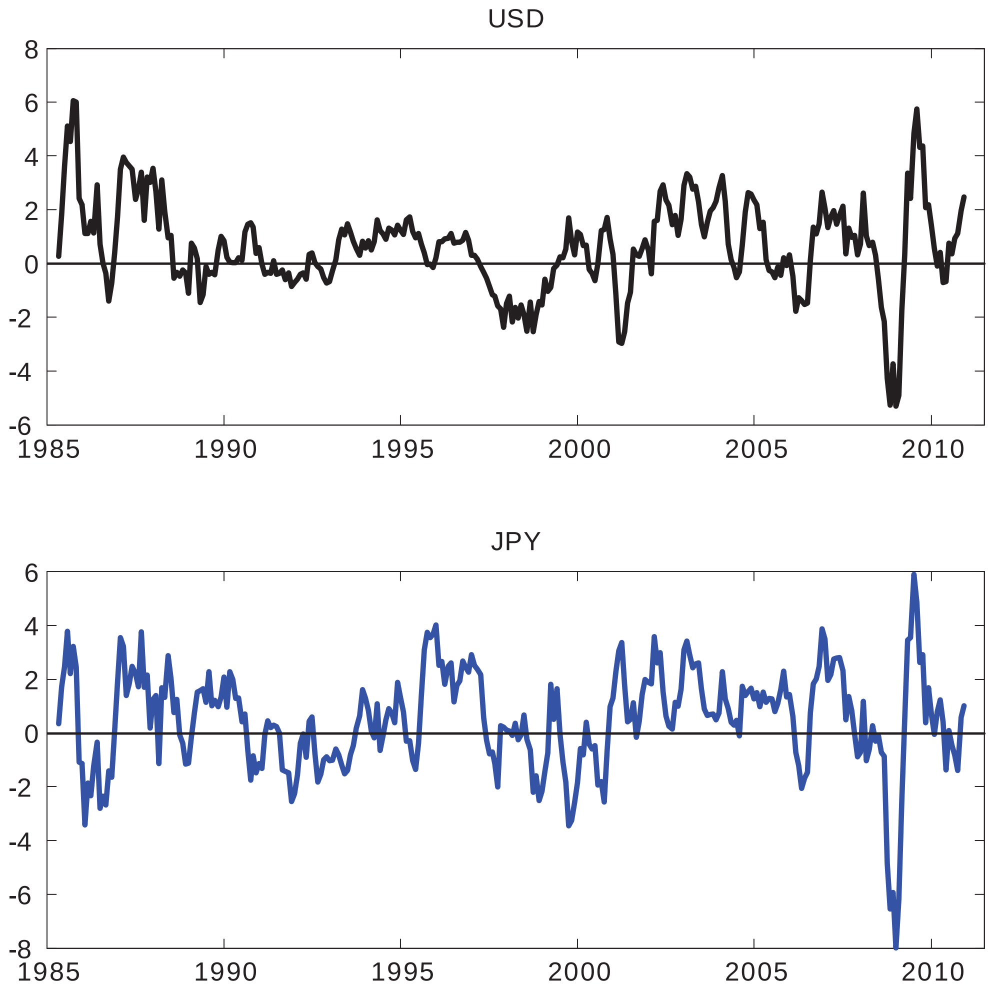

Note: The sample period is January 1995 to December 2010. The graphs indicate the average percentage return in a 6-month rolling window.

FIGURE 4

Relative Implied Volatilities of NZD-USD Options

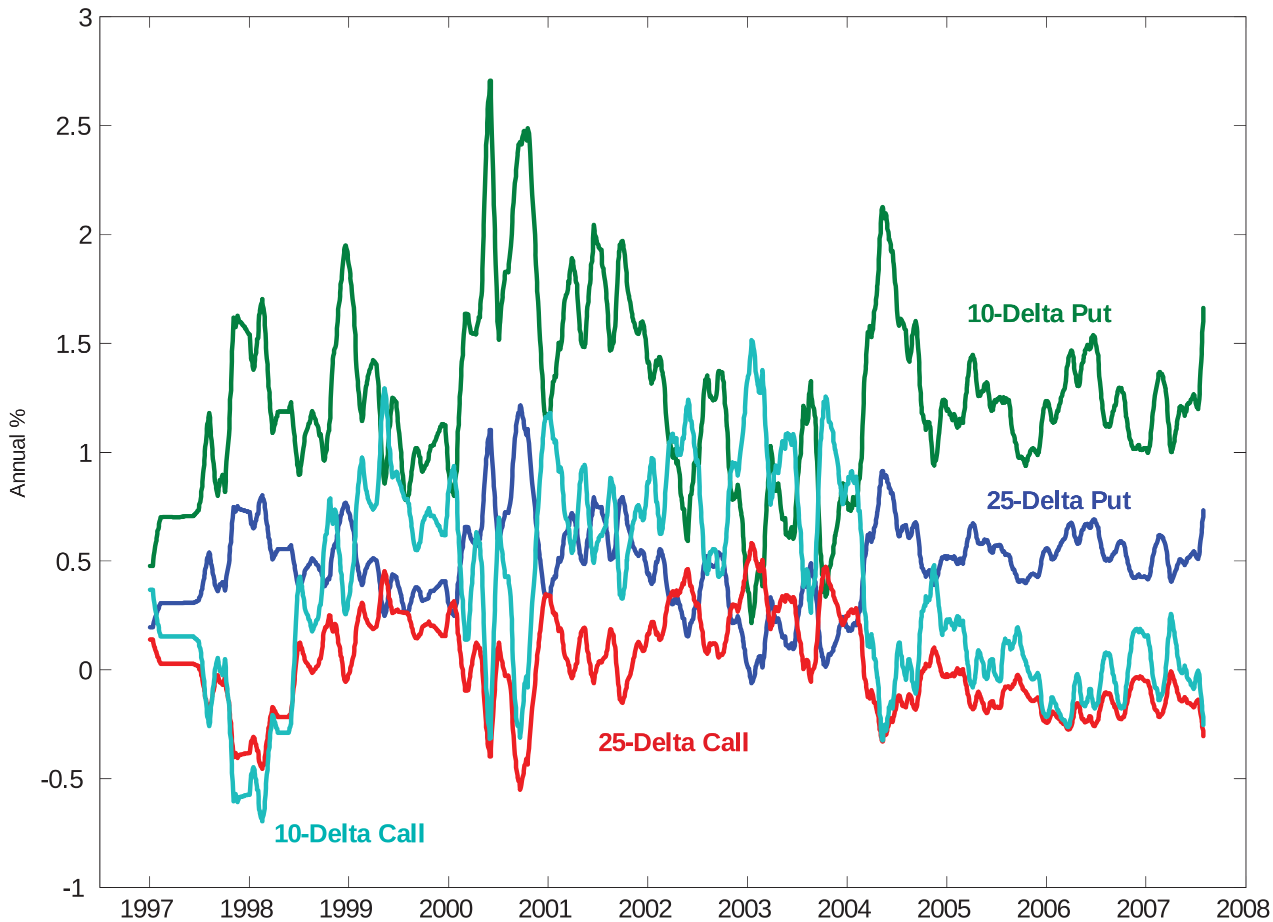

Note: January 1997 to July 2007. The figure plots the time series of differences between the implied volatilities of the 25 and 10-delta options put and call options and at-the-money (forward) NZD-USD options. The data are smoothed with a tent-shaped filter over a 21 business day window.

FIGURE 5

Global and New Zealand-Specific Measures of Currency Volatility, Skewness and Kurtosis

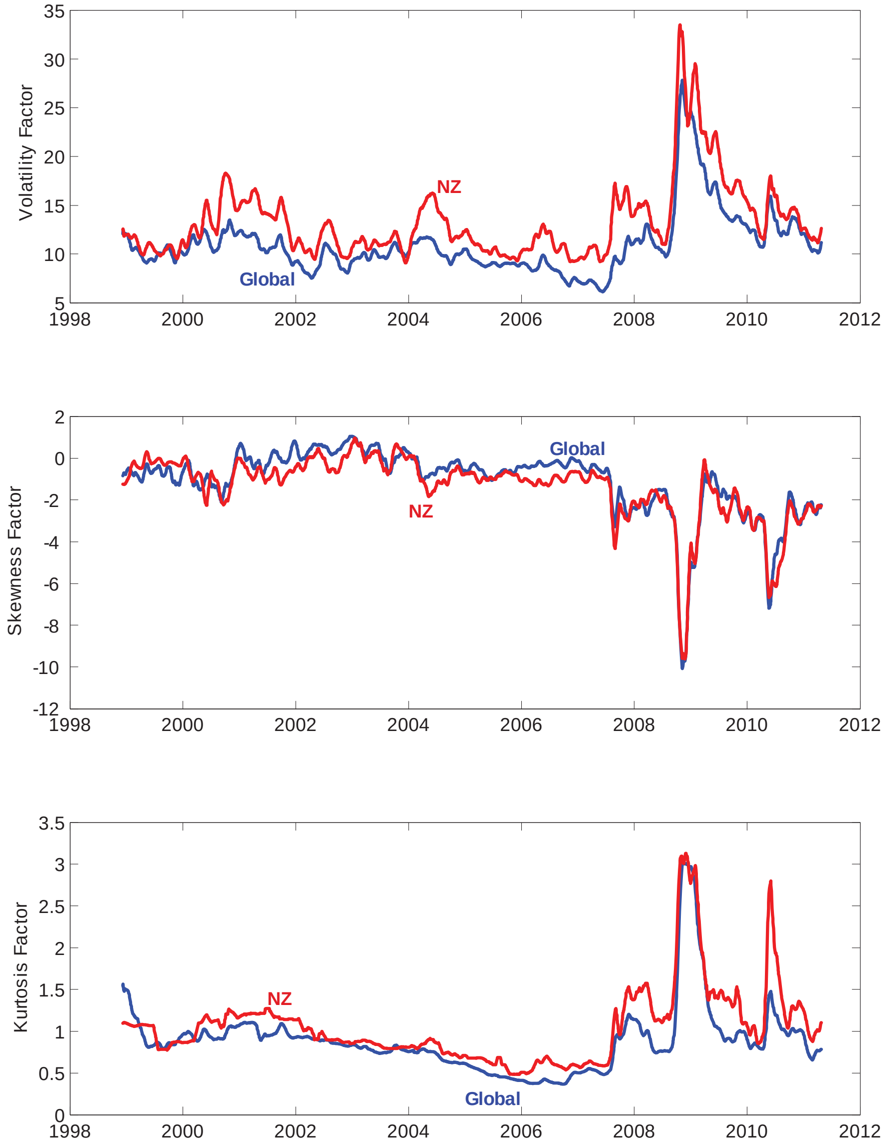

Note: December 1998 to April 2011. The data are smoothed with a tent-shaped filter over a 21 business day window.

FIGURE 6

Exchange Rate Forecasts from the Regression Models

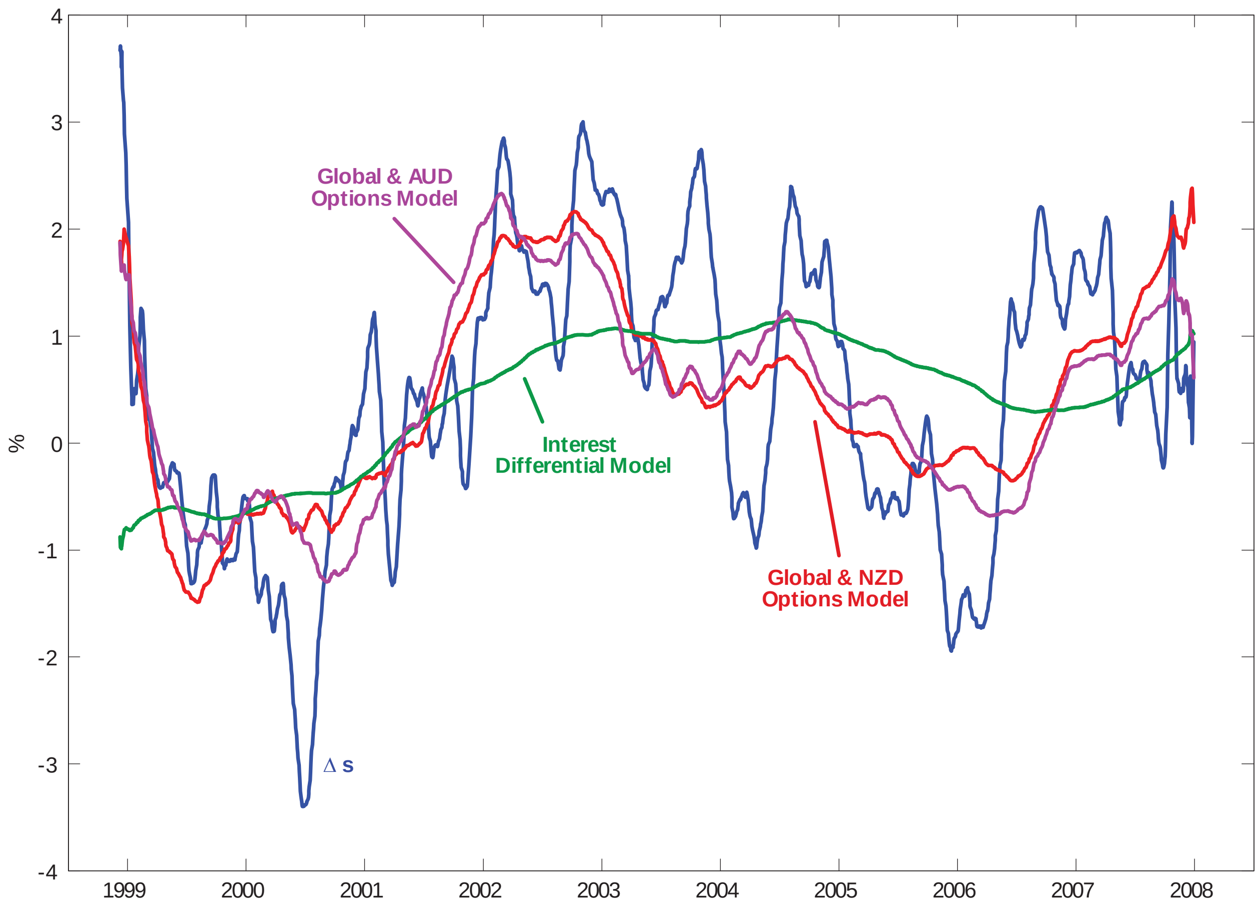

Note: December 1998 to December 2007. The data are smoothed with a tent-shaped filter over a six month period. The blue line is the change (appreciation) of the value of the NZD in USD terms over a month-long period. The green line is the forecast generated, in-sample, using the regression (1) which includes only the interest differential as a right-hand side variable. The red line is generated, in-sample, using the regression (4), which includes the interest differential, the global currency options series, and the NZ-specific currency options series as regressors. The magenta line is generated, in-sample, using a similar regression to (4), in which New Zealand-specific variables are replaced with Australia-specific variables.

FIGURE 7

Exchange Rate Forecasts During the Crisis Period

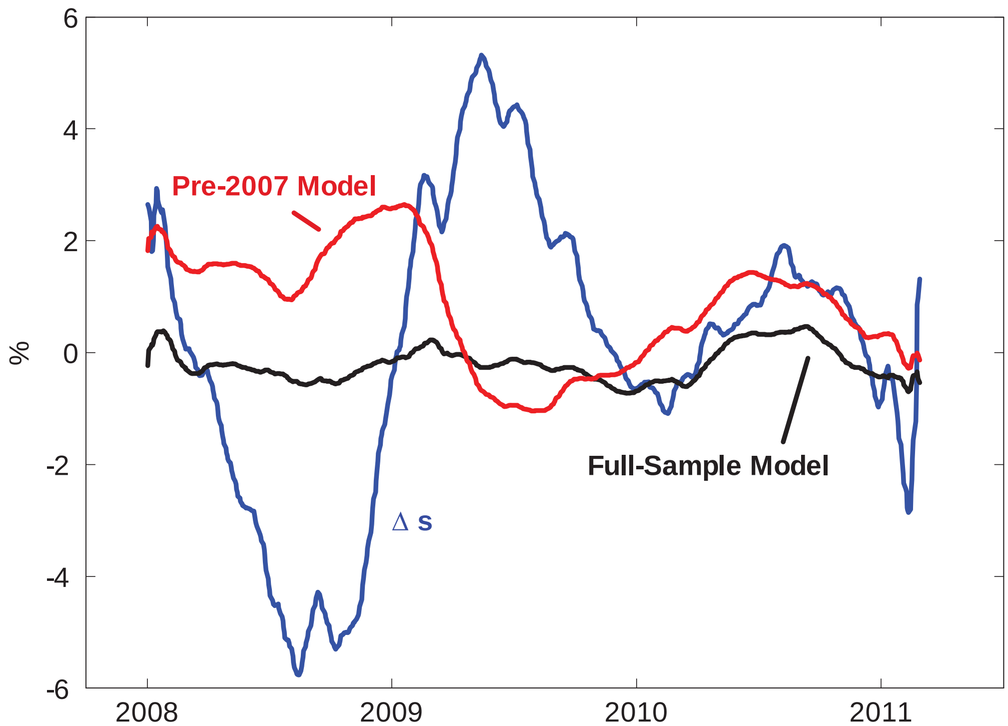

Note: January 2008 to February 2011. The data are smoothed with a tent-shaped filter over a six month period. The blue line is the change (appreciation) of the value of the NZD in USD terms over a month-long period. The red line is generated, out-of-sample, using the regression (4), which includes the interest differential, the global currency options series, and the NZ-specific currency options series as regressors, fit over the period 1998–2007. The black line is generated, in-sample, using the same regression fit over the period 1998–2010.

6 Technical Appendix

6.1 Uncovered Interest Parity

Consider a trader or investor who holds a long position in the New Zealand dollar (NZD), and finances his position by borrowing a foreign currency. Abstracting from transactions costs, if the trader borrows one unit of the foreign currency, the trader's profit is

$$z_{t+1} = (1 + i_t)\frac{S_{t+1}}{S_t} - (1 + i_t^*), \tag{5}$$

where i_t is the interest rate in New Zealand, i_t^* is the interest rate in a foreign country, and S_t is the exchange rate measured as foreign currency units (FCUs) per NZD. I measure time in months. Standard asset pricing theory implies that the payoff to any zero cost investment, such as the one above, should satisfy $E_t(M_{t+1}z_{t+1}) = 0$ where M_t is the relevant stochastic discount factor (SDF) for the marginal investor, and E_t is the expectations operator given time t information. If investors are risk neutral, $M_t = M$ for all t, and $E_t z_{t+1} = 0$. Consequently we end up with the uncovered interest parity (UIP) condition:

$$(1 + i_t)\frac{E_t S_{t+1}}{S_t} - (1 + i_t^*) = 0. \tag{6}$$

For expositional purposes I adopt the following first order approximation for the trader's profit:

$$z_{t+1} = s_{t+1} - s_t + i_t - i_t^*, \tag{7}$$

where $s_t = \ln S_t$.[11] The standard textbook form of the UIP condition comes from assuming that $E_t z_{t+1} = 0$ for this approximation, in which case

$$E_t s_{t+1} - s_t = i_t^* - i_t. \tag{8}$$

A typical test of UIP is to regress $\Delta s_{t+1} \equiv s_{t+1} - s_t$ on the interest differential $i_t^* - i_t$:

$$\Delta s_{t+1} = \beta_0 + \beta_1 (i_t^* - i_t) + \epsilon_{t+1}. \tag{9}$$

If equation (8) holds, the estimated constant and slope coefficients in this regression should be indistinguishable, respectively, from 0 and 1. The typical point estimate for the slope coefficient, β_1, however, is negative, when pairs of industrialized country currencies are used in the regression. However, the results are often found to be sample dependent, and certainly depend quite importantly on which currency is adopted as the baseline.

In running the regressions that appear in the paper, I exploit covered interest parity, which implies that

$$(1 + i_t)\frac{F_t}{S_t} = (1 + i_t^*), \tag{10}$$

[11]Although I use log approximate equations in the text, I use exact expressions for z_t and exact moment restrictions in the empirical work.

where F_t is the one-month forward exchange rate measured in FCUs per NZD. This means that, up to a first order approximation, $i_t - i_t^* = s_t - f_t$, where f_t is the logarithm of the forward rate. So I run the regression as:

$$\Delta s_{t+1} = \beta_0 + \beta_1 (f_t - s_t) + \epsilon_{t+1}. \tag{11}$$

6.2 Currency Trading Profits

Consider a trader who chooses to always be long the NZD, and finances his position by borrowing the foreign currency. The trader's profit is given by (5), and, the log approximate expression for his profit is given by (7). If UIP holds, then, $E_t z_{t+1} = 0$ and, therefore, by the law of iterated expectations, $E(z) = 0$. That is, the unconditional mean of the trader's payoff is zero. This implies that in finite samples, observed profits to always being long the NZD should usually be statistically close to 0.

On the other hand, suppose, for the sake of argument, that UIP does not hold and that $E_t \Delta s_{t+1} = \gamma(i_t^* - i_t)$ with $\gamma < 1$. This assumption is arguably consistent with the results in Table 1. In this case, $E_t z_{t+1} = (1-\gamma)(i_t - i_t^*)$ and $E(z) = (1-\gamma) E(i - i^*)$. So if UIP fails in this particular way, and if the NZD pays a higher rate of interest than the foreign currency, on average, the trader earns positive profits, on average.

Alternatively, suppose the trader engages in the carry trade, where he goes long the NZD if the NZD interest rate is higher than the foreign interest rate, but he goes short the NZD if the NZD interest rate is lower. In this case, the trader's profit is

$$z_{t+1}^C = \text{sign}(i_t - i_t^*) z_{t+1}. \tag{12}$$

If UIP holds, then $E_t z_{t+1}^C = \text{sign}(i_t - i_t^*) E_t z_{t+1} = 0$ and, therefore, $E(z^C) = 0$. So, once again, the trader earns a zero mean payoff in large samples. If, on the other hand, UIP fails because $E_t \Delta s_{t+1} = \gamma(i_t^* - i_t)$ with $\gamma < 1$, then $E_t z_{t+1}^C = (1-\gamma) |i_t - i_t^*|$ and $E(z^C) = (1-\gamma) E|i - i^*|$. In this case, the trader earns positive profits, on average, as long as the NZD and foreign interest rates occasionally differ from one another.

While positive profits were earned in the carry trades summarized in Table 2, they are not of the same magnitude as predicted by the assumption that $E_t s_{t+1} - s_t = \gamma(i_t^* - i_t)$. For example, in the 1990–2007 subsample, $\hat{\gamma} = -2.8$ for the USD. The mean absolute interest rate differential between NZD and USD over this period was 2.9%. If it were the case that $E_t s_{t+1} - s_t = \gamma(i_t^* - i_t)$ then we would expect annual returns of roughly 10.9% in the carry trades between the two currencies, as opposed to the observed 5.3% return observed in the data. The actual returns are more consistent with γ being roughly -0.8 for the USD over the 1990–2007 period. I use this result in the modeling that underlies Section 2.

6.3 Traditional Risk-Based Explanations

The UIP condition should only hold, and carry trade profits should only be zero, on average, if investors in foreign currency are risk neutral. When traders are risk averse, the trader's SDF, M_t, varies over time, and, $E_t M_{t+1} z_{t+1} = 0$ is the relevant theoretical pricing equation for zero cost investments. It follows that

$$E_t\left(z_{t+1}\right) = -\frac{\operatorname{cov}_t\left(M_{t+1}, z_{t+1}\right)}{E_t\left(M_{t+1}\right)} \text{ and } E(z) = -\frac{\operatorname{cov}\left(M, z\right)}{E\left(M\right)} \tag{13}$$

Consider, for example, a long position in NZD, whose payoff is given by equation (7). Given that $z_{t+1} = \Delta s_{t+1} + i_t - i_t^*$, we can rewrite the first part of equation (13) as

$$rp_t \equiv E_t \Delta s_{t+1} + i_t - i_t^* = -\frac{\operatorname{cov}_t\left(M_{t+1}, s_{t+1}\right)}{E_t\left(M_{t+1}\right)}. \tag{14}$$

If the risk premium, rp_t, is positive at date t it means that the value of the NZD conditionally covaries negatively with the investor's SDF. That is, if times turn out to be worse than expected for the investor (M_{t+1} is above $E_t M_{t+1}$), the NZD will likely be worth less than expected ($s_{t+1} < E_t s_{t+1}$).

Similarly, if carry trading in NZD is exposed to risk, this means that M and z^C covary negatively. That is, when times are bad for the investor (M is high), the carry trade, on average, does poorly (z^C is below its mean).

6.4 An Extreme Events Model

Suppose that at any date the state of the world corresponds to one of two possible types: *normal* and *extreme*. I let p_t denote the probability of the state of the world being *extreme* if the state of the world, today, is *normal*. If the state of the world tomorrow is *normal* the stochastic discount factor, M_{t+1}, and the change in the exchange rate, Δs_{t+1} are drawn from the probability distribution $\mathcal{F}_t^N$. If the state of the world tomorrow is *extreme,* M_{t+1} and Δs_{t+1} are drawn from the probability distribution $\mathcal{F}_t^X$. The distributions have time subscripts to indicate that they are potentially time varying and conditioned on time t information.

As before, standard asset pricing theory implies that $E_t(M_{t+1} z_{t+1}) = 0$. In this simple model this means that

$$\left[(1 - p_t) E_t^N (M_{t+1} z_{t+1}) + p_t E_t^X (M_{t+1} z_{t+1})\right] = 0, \tag{15}$$

where E_t^N and E_t^X represent conditional expectations taken with respect to $\mathcal{F}_t^N$ and $\mathcal{F}_t^X$. In order to derive useful implications from equation (15), I make the following assumptions about the distributions $\mathcal{F}_t^N$ and $\mathcal{F}_t^X$:

- First, I assume that the SDF is a time invariant constant, M^N, in all normal states, and $M^N = 1$. The second part of the assumption is inconsequential, since the overall scale of the SDF is irrelevant to the pricing condition for zero cost investments. The first part of the assumption implies that $E_t^N(M_{t+1}) = M^N$, and that there is no covariance between the SDF and the currency payoff in normal states. This implication of my assumption is broadly consistent with the fact that conventional risk factors appear to be uncorrelated with currency returns in the non-financial crisis part of the sample. So we have $E_t^N(M_{t+1} z_{t+1}) = E_t^N(z_{t+1})$.

- Second, I assume that $E_t^N \Delta s_{t+1} = \gamma(i_t^* - i_t)$ with $\gamma < 1$. This assumption is broadly consistent with the estimated UIP regressions, when the financial crisis is excluded from the sample.
- Third, I assume, for simplicity, that the SDF is a time invariant constant, M^X, in all extreme states. This implies that $E_t^X(M_{t+1}\Delta s_{t+1}) = M^X E_t^X(\Delta s_{t+1})$.

The three assumptions, together with (15), imply that

$$\left[(1-p_t)(1-\gamma) + p_t M^X\right](i_t - i_t^*) = -p_t M^X E_t^X(\Delta s_{t+1}). \tag{16}$$

Notice that if $p_t = 0$ then $i_t = i_t^*$. So, for this simple model to explain a positive interest differential between NZD and foreign currency, the probability of the extreme state must be strictly positive. Furthermore, because the term in square brackets on the left hand side of (16) is positive, $E_t^X(\Delta s_{t+1})$ must be negative. If the probability of the extreme event is small, then $M^X E_t^X(\Delta s_{t+1})$ must be a large negative number, either because:

- M^X is very large, because the extreme event is very *bad* from the perspective of the investor (say, because the investor has a very high marginal utility of consumption in the extreme state), or because
- $E_t^X(\Delta s_{t+1})$ is very negative, because the extreme event is characterized by a large depreciation of the NZD.

In this simple model, the risk premium during normal times is given by

$$E_t^N z_{t+1} = E_t^N \Delta s_{t+1} + i_t - i_t^* = (1-\gamma)(i_t - i_t^*)$$

so it moves in proportion to the interest differential. The interest differential, in turn, is determined by equation (16), which can be rearranged as:

$$i_t - i_t^* = -\frac{p_t M^X}{(1-p_t)(1-\gamma) + p_t M^X} E_t^X(\Delta s_{t+1}). \tag{17}$$

Since $\gamma < 1$ and $E_t^X(\Delta s_{t+1}) < 0$, the level of the risk premium is increasing in p_t (the probability of an extreme state occurring), $-E_t^X(\Delta s_{t+1})$ (the size of the expected depreciation in the extreme state), and M^X, the value of the SDF in the extreme state. Changes in the risk premium, over time, reflect changes in p_t and $E_t^X(\Delta s_{t+1})$.[12]

An important aspect of the model is that it implies that investors anticipate a type of conditional skewness in currency returns that may not be observed in a sample containing no extreme events. There are two senses in which this is true. The first is conventional skewness, in that returns in the extreme states are below the conditional mean (since $E_t^X \Delta s_{t+1} < 0$ and $E_t^N \Delta s_{t+1} > 0$). The second is skewness in the risk

[12] For simplicity I rule out another possible explanation for time variation in the risk premium: that the expected value of M^X varies over time. Such variation is not separately identifiable from variation in p_t so I abstract from it here.

neutral distribution of returns, due to the fact that investors may put more weight on extreme states; i.e. it may be the case that $M^X > 1$.[13] In either case, this conditional skewness will be priced into NZD-denominated assets.

The statements about currency options in the main text are based on standard pricing theory and practice for currency options. The implied volatilities (IVs) are the standard deviation coefficients that render the outright prices of the options consistent with the Garman-Kohlhagen pricing formula (which is the currency option equivalent of the Black-Scholes formula for option prices). This formula makes the theoretical assumption that the log change in the exchange rate is Gaussian and centered at the forward premium. If the assumption of the Garman-Kohlhagen formula was correct, then the IVs of the five sets of options prices that underlie Figure 3 would be the same and the four lines in the graph would all be zero. If the risk neutral distribution was fat-tailed but symmetric around the forward premium, then we might expect the IVs of the 10-delta options to be larger than the IVs of the 25-delta options, which would, in turn, be larger than the IV of an ATMF option. In this case all four lines would lie above zero, with the green line and the cyan line being higher than the blue and red lines. If the risk neutral distribution was skewed downward (i.e. in the direction of a depreciation of the NZD), but was not fat tailed, we would expect the IVs of the call options to be smaller than the IV of the ATMF option, while the IVs of the put options would be larger than the IV of the ATMF option. In this case, we would expect the red and cyan lines to lie below zero, while the green and blue lines would lie above zero. Figure 3 suggests that the risk neutral distribution is both skewed downward and fat tailed, consistent with results in Jurek (2008).

With additional assumptions, currency options data can be used to directly shed light on the small model described above. I assume that conditional on drawing a normal state tomorrow, $\Delta s_{t+1} \sim N(\gamma(i_t^* - i_t), \sigma_t^2)$, where σ_t^2 is the potentially time-varying conditional variance of Δs_{t+1}. I assume that conditional on drawing an extreme state tomorrow, $\Delta s_{t+1} \sim N(\mu_t, \psi_t^2)$, with $\mu_t < 0$. With this assumption, equation (16) can be written as:

$$\left[(1-p_t) + p_t M^X\right](i_t - i_t^*) = \gamma(1-p_t)(i_t - i_t^*) - p_t M^X \mu_t. \tag{18}$$

Dividing through by $(1-p_t) + p_t M^X$, (18) can be written as

$$i_t - i_t^* = \gamma(1-\omega_t)(i_t - i_t^*) - \omega_t \mu_t, \tag{19}$$

where

$$\omega_t = \frac{p_t M^X}{(1-p_t) + p_t M^X} \tag{20}$$

is the *risk neutral probability* of the rare event.

The assumption that the log change in the exchange rate is a mixture of normal distributions also allows us to price currency options. Take, for example, a put option with strike price K_t, which has price C_t^P measured in USD per NZD. Suppose an

[13] The fact that $E_t^N \Delta s_{t+1} > E_t^X \Delta s_{t+1}$ implies a certain amount of skewness in the overall objective conditional distribution of Δs_{t+1}. If $M^X > 1$, the risk neutral distribution will be more skewed than the objective one.

investor buys $(1+i_t^*)/F_t$ units of this put option, say, to hedge a currency position that is short 1 USD and long in NZD. The cost to the investor is $c_t^P = (1+i_t^*)C_t^P/F_t$, which is observable given the currency options data on which Figure 3 is based.

In the theoretical model, what should be the price of this option? The gross payoff of the investor's position in put options is $\max\{(K_t - S_{t+1})(1+i_t^*)/F_t, 0\}$, which, in log approximation terms, means that the net payoff to the option is

$$\begin{aligned} z_{t+1}^P &= \max\{k_t - \Delta s_{t+1}, 0\} - c_t^P \\ &= k_t - \min\{\Delta s_{t+1}, k_t\} - c_t^P \end{aligned} \tag{21}$$

where $k_t = \ln(K_t/S_t)$. Since the claim to the net payoff of the option is zero cost, $E_t M_{t+1} z_{t+1}^P = 0$, which, given the small model, implies:

$$(1-\omega_t)E_t^N\left(k_t - \min\{\Delta s_{t+1}, k_t\} - c_t^P\right) + \omega_t E_t^X\left(k_t - \min\{\Delta s_{t+1}, k_t\} - c_t^P\right) = 0, \tag{22}$$

or

$$k_t = -(1-\omega_t)E_t^N \min\{\Delta s_{t+1}, k_t\} - \omega_t E_t^X\{\Delta s_{t+1}, k_t\} + c_t^P. \tag{23}$$

Given our assumptions about the distribution of Δs_{t+1} we have

$$E_t^N \min\{\Delta s_{t+1}, k_t\} = -\gamma(i_t - i_t^*)\Phi\left(\hat{k}_t^N\right) + k_t\left[1 - \Phi(\hat{k}_t^N)\right] - \frac{\sigma_t}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\left(\hat{k}_t^N\right)^2\right] \tag{24}$$

and

$$E_t^X \min\{\Delta s_{t+1}, k_t\} = \mu_t\Phi\left(\hat{k}_t^X\right) + k_t\left[1 - \Phi(\hat{k}_t^X)\right] - \frac{\psi_t}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\left(\hat{k}_t^X\right)^2\right] \tag{25}$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution function, and

$$\hat{k}_t^N \equiv \frac{k_t + \gamma(i_t - i_t^*)}{\sigma_t} \qquad \hat{k}_t^X \equiv \frac{k_t - \mu_t}{\sigma_t}. \tag{26}$$

Equation (19) and equation (23), contain observable data, $i_t - i_t^*$, k_t, and c_t^P, as well as four unobservables, ω_t, σ_t, μ_t, and ψ_t. But with options data at different strike prices, we can repeat equation (23) without adding unobservables, while only adding new observables (the strike prices and option prices for the additional options). Our options data set, in fact, allows us to repeat (23) three times, for put options that are ATMF, 25-delta and 10-delta. Hence, we effectively have four equations in four unknowns.

While it is possible to solve these equations in every time period, I opt, instead, to measure the average values of the observables that enter the four equations over the available pre-crisis sample for the options data (1997–2007). The average value of $i_t - i_t^*$ is 0.0022, or 2.6% on an annual basis. The average values of k_t and c_t^P are -0.0014 and 0.0139 (for ATMF options), -0.0251 and 0.0054 (for 25-delta options) and -0.0486 and 0.0018 (for 10-delta options). I set $\gamma = -0.8$, which is the value of γ consistent with average profits to the carry trade between the USD and NZD in the 1990–2007 period. The implied values of the parameters are $\omega_t = 0.074$, $\sigma_t = 0.031$, $\mu_t = -0.052$

and $\psi_t = 0.035$. The resulting risk-neutral distribution of Δs is illustrated in Figure A1.

The implied values of the parameters imply that $E_t^N \Delta s_{t+1} = 0.0018$ and $E_t^X \Delta s_{t+1} = \mu_t = -0.052$, so that $E_t^X \Delta s_{t+1}$ is less than two standard deviations (measured either as σ_t or ψ_t) away from $E_t^N \Delta s_{t+1}$. This suggests that the extreme event is not particularly extreme in terms of the implied depreciation of the NZD.

Instead, what is worthy of note is that the risk neutral probability of the extreme event is $\omega_t = 0.074$. This is statistically close to the estimates provided in Burnside, Eichenbaum and Rebelo (2011) for a wide set of currencies. Could this plausibly be the true probability of the event, p_t? It seems unlikely, for two reasons. The first is that between 1976 and 2010 (a period spanning 420 months), there is really only one six-month episode (the recent financial crisis) which can be deemed extreme. This suggests a value of p_t more on the order of 0.014, which is much smaller than ω_t. Second, Nakamura et al. (2010) estimate the probability of consumption disasters (which motivate the notion of extreme events) to be 1.7% at the annual frequency, which corresponds to a monthly frequency of less than 0.2%. Consequently, the main reason extreme events affect asset prices so much is that the implied value of the SDF in the extreme state, M^X, is large relative to the SDF in normal states, 1. In our case, it would appear that M^X is 5 or larger, depending upon one's assumptions about the probability of an extreme event.

6.5 Time-Variation in the Risk Premium

If we maintain two assumptions from the previous section, that the SDF in all normal states is $M^N = 1$, and in all extreme states is $M^X > 1$, then (15) implies

$$(1 - p_t)E_t^N(\Delta s_{t+1} + i_t - i_t^*) + p_t M^X E_t^X(\Delta s_{t+1} + i_t - i_t^*) = 0. \tag{27}$$

Hence, the risk premium, $rp_t \equiv E_t(\Delta s_{t+1} + i_t - i_t^*)$, is given by

$$rp_t = (1 - p_t)E_t^N(\Delta s_{t+1} + i_t - i_t^*) + p_t E_t^X(\Delta s_{t+1} + i_t - i_t^*). \tag{28}$$

If the theoretical pricing equation, (27), is used to substitute out $E_t^X \Delta s_{t+1}$ from (28) then

$$rp_t = (1 - p_t)\frac{M^X - 1}{M^X} E_t^N(\Delta s_{t+1} + i_t - i_t^*).$$

The regressions presented in Section 3, by isolating time series variation in $E_t^N \Delta s_{t+1}$, identify fluctuations in the risk premium, as long as the implied variation in $E_t^N \Delta s_{t+1}$ is not swamped by variation in the objective probability of an extreme state.

FIGURE A1

Probability Distributions for the Exchange Rate in the Simple Model

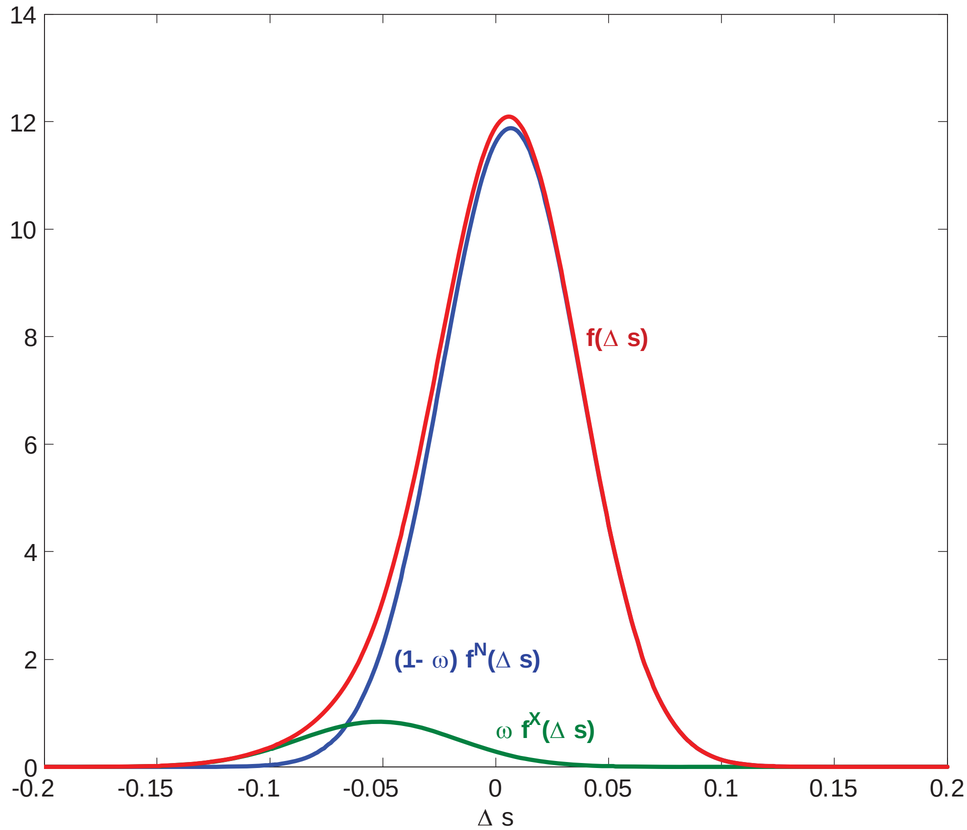

Note: The graph illustrates the contribution of the two probability distributions, $\mathcal{F}^N$ and $\mathcal{F}^X$ to the overall risk-neutral distribution of the change in the exchange rate, Δs. The green line is $\omega f^X(\Delta s)$, the distribution across extreme states times the risk-neutral probability of an extreme state occurring. The blue line is $(1-\omega)f^N(\Delta s)$, the distribution across normal states times the risk-neutral probability of a normal state occurring. The red line is the sum of the blue and green lines and is the overall distribution. It has notable negative skewness.