Three essays on exchange-rate misalignment

Mark Joy

Submitted in fulfillment
of the requirements for the Degree
of Doctor of Philosophy

Department of Economics
Faculty of Law, Business and Social Science
University of Glasgow
June 2010

Abstract

Theories of exchange-rate determination have generated a vast theoretical and empirical literature. This thesis adds to that body of literature by asking three questions. (i) How do policymakers respond to exchange-rate misalignment? (ii) How does misalignment affect the decisions of financial-market participants? (iii) What do exchange-rate dynamics reveal about the choices of investors in the face of currency risk? These three questions are tackled with studies that offer broad and tractable conclusions and contribute to furthering the current field of research.

Contents

List of Figures

List of Tables

Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Signature ....

Printed name ....

Chapter 1

Overview

This chapter describes the motivating forces behind the three studies that make up this thesis and offers an overview of the methodologies, findings and conclusions.

1.1 Introduction

The desire to understand what governs the movement of exchange rates is a core driver of financial and international economic research. Policymakers require an understanding of how exchange rates affect macroeconomic policy and on the basis of that understanding, flawed or otherwise, they may wish to initiate policy that attempts to influence the exchange rate's value. Investors, meanwhile, are concerned with currency movements in as much as they affect their decisions over portfolio allocation and risk. Forecasting exchange-rate movements is important. As is understanding interdependencies with other asset classes.

Over the years economics has volunteered a number of theories of exchange-rate determination. Equilibrium models (MacDonald, 2000), liquidity models (Grilli and Roubini, 1992), the portfolio balance approach (Dooley and Isard, 1979) and the flexible price monetary model (Frenkel, 1976)—overshadowed subsequently by the sticky-price model of overshooting—dominated research during the 1960s, 1970s and 1980s. Since the 1990s, following the work of Obstfeld and Rogoff (1995), new open economy macroeconomics has proposed that exchange-rate movements are best explained by dynamic general equilibrium models. Market microstructure approaches to exchange-rate determination have offered perhaps the best account of the high-frequency volatility of exchange rates. But for the purposes of policy, predictability and estimates of misalignment, market microstructure approaches fall short. Structural approaches continue to dominate both policy and research.[1]

Despite the progress made in the modelling of exchange rates, many questions regarding currency movement, misalignment and spill-over effects remain unanswered. This thesis tackles three questions. Firstly, how do policymakers respond to exchange-rate misalignments? Specifically, when policymakers intervene in the currency markets in response to misalignments, what influences the decision to intervene? This thesis offers a study of Japanese intervention in the currency markets in an effort to throw light on the determinants of intervention policy and to gauge the extent to which the intervention decision is driven by perceptions of exchange-rate misalignment. The second research question can be stated as, how does misalignment affect the decisions of financial-market participants? A partial answer is offered by investigating the role played by perceptions of misalignment, defined as deviations from covered and uncovered interest-rate parity, in shaping the decision to denominate debt in foreign currencies. The investigation incorporates a large sample of foreign bonds in a panel count model of currency choice. The third question asks, what do exchange-rate dynamics reveal about the choices of investors in the face of currency risk? To tackle this question this thesis undertakes an analysis of the extent to which currency dynamics and conditional correlations offer clues as to the suitability of other assets as hedging instruments. The empirical focus is on gold as a hedge against the US dollar.

A number of important results arise from this research. Those with the broadest implications can be summarised as follows. First, the perception of misalignment does indeed influence official intervention in the currency markets. Judging by Japan's history of official intervention, the larger the misalignment, the more likely the intervention. Second, perceptions of misalignment play an important role in shaping the borrowing decisions of corporate and public issuers of international debt: choice of issuance currency is sensitive to deviations from uncovered interest-rate parity. Third, the dynamics of the US dollar reveal that the suitability of other asset classes as hedging instruments varies over time. In recent years gold has become an increasingly suitable hedge against dollar volatility.

This chapter is organised as follows. Section 1.2 introduces the first study on misalignment and intervention policy. Section 1.3 previews the study on market response to perceived misalignment and Section 1.4 introduces the final study on currency dynamics and hedging. Some conclusions are offered in Section 1.5.

[1] For a recent survey of methodological advances in the estimation of equilibrium exchange rates see Bussiere et al. (2010).

1.2 Misalignment and intervention

This section introduces the research presented later in this thesis on currency misalignment and official intervention.

Despite falling out of fashion in the 1990s, official intervention in the currency markets has in recent years re-established itself as an important tool of exchange-rate policy for many countries. The Swiss National Bank revived its intervention policy in March 2009 in an attempt to prevent the Swiss franc from rising sharply in value. The aim, according to Swiss National Bank Chairman Philipp Hildebrand, was to "prevent an excessive appreciation" of the domestic currency. China continues to intervene heavily to stem the strength of the renminbi. Brazil, Poland, India, South Africa and South Korea all engaged in currency intervention in 2009 and 2010. America's monetary authorities offer clear advice regarding their stance on currency intervention:

> Since the breakdown of the Bretton Woods system in 1971, the United States has used currency intervention both to slow rapid exchange rate moves and to signal the US monetary authorities' view that the exchange rate did not reflect fundamental economic conditions. *Federal Reserve Bank of New York (May 2007)*

US intervention was considerable in the 1980s but became much less frequent in the 1990s. The American monetary authorities intervened in the currency markets on eight occasions in 1995, but only twice between August 1995 and December 2006.

Japan, meanwhile, has had an active intervention policy during recent decades. It has engaged in more than US$620bn-worth of interventions in the currency markets since 1991. In November 2009, comments from senior Japanese finance officials suggested the Bank of Japan was closer to currency intervention than at any time since it last intervened in March 2004.

However, while Japan ranks as perhaps the most prolific official intervener in currency markets, its reasons for intervening are understood barely, if at all.[2] Do Japan's monetary authorities intervene in the foreign-exchange markets in order to keep the yen close to a fixed, pre-determined value? To keep it within fixed bounds of tolerance? Within time-varying bounds of tolerance? Crucially, what role is played by perceptions of misalignment?

These questions form the motivating force behind the first study presented in this thesis, a study of the intervention policy of Japan between 1991 and 2006. The intention is that findings from this study offer valuable information,

[2] For surveys of intervention policy see Edison (1993), Dominguez and Frankel (1993) and Sarno and Taylor (2001).

both in approach and conclusions, that can be used to further research into intervention policy more widely.

The study runs as follows. First is a presentation of the stylised facts, describing the movement of the yen against the US dollar during the sample period and the timing and size of interventions. Two features are clear. (i) Intervention is infrequent. (ii) When intervention does occur, sales of yen are undertaken when the Japanese currency is strong relative to its unconditional mean and purchases of yen are undertaken when it is relatively weak. Next, the study offers a theoretical model of Japan's intervention policy. An intervention reaction function is derived. The reaction function proposes that intervention can be described adequately by assuming the central bank desires to minimise deviations from a time-varying currency target. The currency target is assumed to reflect, in this case, a capital-enhanced version of purchasing power parity (MacDonald and Marsh, 1997).

Following the theoretical model is a description of the empirical model in which the dependent variable, intervention, is represented as a qualitative dependent variable carrying a natural order, or rank, describing three categorical states: yen-selling intervention, no intervention, yen-buying intervention. The main contribution of the empirical model is to add flexibility over and above that present in other similar studies. Specifically, the empirical framework is a generalised ordinal logit model which accounts for asymmetry and, in particular, allows for the possibility that deviations from the currency target may have marginal effects that vary according to the intervention category. Results are then presented and conclusions drawn.

1.3 Misalignment and market response

This section introduces the second study in this thesis. The second study is an analysis of the role played by currency misalignment in affecting the decisions of financial-market participants, specifically issuers of international debt.

Issuance of foreign-currency-denominated debt securities has been an important feature in global financial markets for many years, with net issuance more than tripling in value during the past decade (measured at constant exchange rates), reaching USD 1.4 trillion in 2007. The choice of issuance currency is affected by a number of factors. One major factor is the issuer's desire to ensure its financial obligations are in currencies that match the currencies of its cash inflows. By doing so, the issuer creates a "natural hedge" against its currency risk. Another factor is strategy. The issuer's strategic considerations may include the desire to diversify its investor base and, for large-size bond issues, the opportunity to exploit fewer credit constraints in more liquid, foreign bond markets. A third factor affecting the choice of issuance currency (and a factor that is not well explored in the academic literature) is the scope for reductions in borrowing costs through issuing bonds in whichever currencies offer the lowest effective cost of capital. Lower effective borrowing costs can mean lower covered costs (incorporating the cost of covering against exchange-rate risk) or lower nominal costs, reflecting, simply, lower nominal interest rates. Anecdotally, participants in the international bond markets report that both covered and uncovered costs play important roles in the choice of issuance currency.

The second study in this thesis assesses the extent to which perceptions of currency misalignment, in the form of deviations from covered interest-rate parity and uncovered interest-rate parity, influence the decision to issue bonds denominated in foreign currencies. In other words, this study asks, does currency misalignment affect the choice of issuance currency?

Many existing studies of debt issuance offer plausible accounts of the motivating factors behind the issuance of international bonds. What they ignore, however, is the possibility that issuance in a foreign currency is driven largely by an opportunistic desire to lower costs. That is, they ignore the possibility that at the time of issuance, issuers choose to denominate their borrowing in one currency rather than another simply because the chosen currency offers lower effective borrowing costs.

The idea that cost savings can be secured by issuing bonds in low-interest-rate currencies does, of course, violate traditional interest-rate-parity conditions that seek to explain the short-term movement, and misalignment, of international exchange rates. The condition of uncovered interest-rate parity asserts that any discount in foreign interest rates will be offset exactly by the expected appreciation of the foreign currency. If this parity condition holds true, it leaves no scope for exploitable cost savings from opportunistic issuance. Empirically, however, uncovered interest-rate parity does not, in general, hold true.[3] Most empirical studies find that low-interest-rate currencies do not systematically appreciate over time as suggested by uncovered interest-rate parity. In fact, they tend to do the opposite: they depreciate. This suggests that in practice there are cost savings to be secured by leaving exchange-rate risk uncovered and issuing bonds in low-interest-rate currencies.

The second study in this thesis offers a closer examination of the responsiveness of international bond issuance to not just deviations from uncovered interest-rate parity but also from covered interest-rate parity. It draws on a large, unique dataset, employs a utility-consistent model, and adopts a novel empirical approach to tackle the question of currency choice in international bond issuance by focussing on the number, not the value, of bonds issued in

[3]See, for instance, Isard (1996).

international currencies.

The study takes the following format. First is presented a model of currency choice over time: a choice among major issuance currencies by issuers of international bonds. A description is given explaining how this model can be embedded within a utility-consistent framework. The complication here is that the dependent variable is chosen to be *number* of bonds issued rather than *value* of bonds issued. The reason for this is straightforward: there exists evidence to suggest that it is the number of issues, not the value, that responds to currency misalignment. This is because the issuer's decision over the *value* of any bond offering tends to be determined before the actual date of the offering, sometimes up to a year before. Irrespective of the value of the bond issue, a broker will advise the issuer of the most advantageous time to execute the bond offering. This advice will be based, for issuers of international bonds, on an evaluation of financial conditions including currency movements. At an aggregate level, therefore, the main, detectable response to deviations from covered and uncovered interest-rate parity, in any given period, will not, necessarily, be a change in total *value* of bonds issued in a certain currency, it will be a change in total *number* of bonds issued. The appropriate empirical model is, as such, a panel count model, a model that is shown to be consistent with utility theory.

Next the study provides a description of the dataset, compiled using thousands of individual records of bond issues dovetailed with the constructed measures of currency misalignment: deviations from covered and uncovered interest-rate parity. There is a description of how the dataset is split into three maturity brackets: short, medium and long. Also, there follows an overview of how the concepts of uncovered interest-rate parity and covered interest-rate parity can be made relevant for the types of time horizons that are applicable to bond issuance—namely, horizons of one year to ten years and beyond. Central is the role of the swaps market, allowing for a revised, non-arbitrage condition called swaps-covered interest-rate parity. Subsequent to this is a discussion of the empirical results, robustness checks and finally some concluding remarks.

1.4 Currency dynamics and hedging

This section introduces the third study in this thesis, which asks the question, what do currency dynamics reveal about the choices investors make when faced with currency risk?

The increasing role played by globalised financial markets in influencing the economic fortunes of the developed world offers a persuasive basis for investigating possible relationships between changes in the value of exchange rates and the returns on risky assets. Indeed, in the last ten years a strand of research has developed exploring the nature of the relationship between the dynamics of exchange rates and the returns on stocks and bonds. Important contributors to this work are Brandt et al. (2001), Pavlova and Rigobon (2003) and Hau and Rey (2006). Their findings suggest the relationships are strong and meaningful.

The third study in this thesis looks at the link between exchange-rate dynamics and commodity returns. In particular the focus is on the US dollar and changes in the price of gold. The main reason for this focus is that in financial markets the nature of the relationship between gold and the US dollar tends to be commented upon widely but understood little. Market wisdom has it that when the US dollar depreciates, the price of gold rises, and when the US dollar appreciates, the price of gold falls. When such price movements do coincide, market reports offer hazy rationalisations based on, among other things, substitution effects, pricing conventions and hedging motives. None of these offer convincing descriptions.

This thesis assesses the extent to which an inverse relationship between the price of gold and the value of the US dollar does, in fact, exist, and asks, does gold act as a hedge against the US dollar, as a safe haven, or neither? Specifically, the focus is on the association between movements in the price of gold and the US dollar using a model of dynamic conditional correlations covering 23 years of weekly data for 16 major US dollar-paired exchange rates.

The study runs as follows. First, definitions are established. What, exactly, is a hedge? What is a haven? After this is some background discussion regarding correlation models. Why does the concept of correlation feature heavily in the models of risk and return? What is required to ensure accurate estimation of correlations? Discussion centres on observability, on the need to incorporate dynamics and on the curse of dimensionality.

Next there is a description of the correlation model employed in the study: a model allowing for dynamic conditional correlations. The description outlines the model's origins and starts with the model of constant conditional correlations first proposed by Bollerslev (1990). Next is a discussion of estimation methodology. Estimation is a two-stage process (Engle, 2002). In the first stage, univariate GARCH models are estimated for each returns series. In the second stage, the first-stage residuals are taken and transformed by their standard deviations in order to estimate the parameters of the dynamic conditional correlation model. Following this is a discussion of the data, presentation of the empirical results and some concluding remarks.

1.5 Conclusions

The aim of this thesis is three-fold: to throw light on the link between currency misalignment and policy, between currency misalignment and market response, and between currency dynamics and hedging. This section summarises the main findings and highlights the contribution this thesis makes to the existing literature.

The first study, exploring the link between currency misalignment and intervention policy, estimates a reaction function for official Japanese intervention in the currency markets between April 1991 and March 2006. Estimation results show that intervention during the sample period conforms to a model in which the monetary authorities intervene in order to prevent the yen from straying too far from an equilibrium value defined by a capital-enhanced version of purchasing power parity. Predictability is good. A generalised ordered logit model provides the empirical framework and results show that studies of intervention that ignore violations of the *proportional odds assumption* are likely to suffer from specification error.

There are two primary contributions that this study makes to the existing literature. First, it tests the hypothesis that the aim of optimal intervention policy in Japan is to prevent the nominal exchange rate from straying too far from its medium-run equilibrium. *Medium-run equilibrium* is defined in terms of a capital-enhanced version of purchasing power parity.[4] By incorporating a measure of exchange-rate equilibrium explicitly within the intervention reaction function, this study improves over other studies that assume the monetary authorities desire nothing more than a backward-looking adjustment towards trend.

The other main contribution of this study is empirical: a *partial proportional odds model* of intervention is adopted that allows for asymmetry in the intervention objective function. Asymmetry is, indeed, shown to be present. By taking this flexible approach to estimation this study improves over other studies that do not allow for violations of the *proportional odds assumption.*

The second study contained in this thesis focuses on the market response of issuers of international debt to currency misalignment. Summarising the main results, this study finds that a significant response in terms of number of bonds issued in a given currency is, indeed, associated with deviations from uncovered interest-rate parity and, by extension, associated with perceptions of currency misalignment. If, in any given period, the basis-point measure of deviations from uncovered interest-rate parity for, say the euro, rises by 20 basis points, then the

[4]The capital-enhanced version of purchasing power parity is outlined by Juselius (1991, 1995), MacDonald and Marsh (1997, 1999) and Juselius and MacDonald (2000b,a).

expected number of international bonds issued in euros increases, on average, by almost 10%. Furthermore, in terms of number of bonds issued, financial corporations are even more responsive than the average issuer to deviations from uncovered interest-rate parity.

This study makes three main contributions to the existing literature. First, it employs a unique dataset that draws on the entire population of international bond issues during the sample period. Second, it presents an analysis of the issuance of foreign-currency bonds by *number* of issues rather than, as is customary in the literature, by *value* of issues (that is, this study draws on count-data techniques). Third, this study embeds its model of bond issuance within a framework of random utility maximisation.

The final study presented in this thesis investigates the relationship between currency dynamics and hedging. Specifically, the investigation assesses the role of gold as a hedge against the US dollar. Key findings are as follows. First, during the past 23 years gold has behaved as a hedge against the US dollar—that is, gold-price returns have, on average, been correlated negatively with US dollar returns. Second, there is no evidence to suggest that gold has acted as a consistent and effective safe haven. Third, in recent years gold has become an increasingly effective hedge against the US dollar, with conditional correlations more negative now than they have been at any point during the past two and a half decades.

The contribution of this study to the existing literature is two-fold. First, the study offers an empirical analysis of the relationship between gold-price returns and exchange-rate returns, modelling the time-varying correlations between a 17-variable system of returns using the correlation modelling techniques of Engle (2002). As far as the author is aware no other study offers such an analysis. Second, this study assesses the role of gold as both a hedge and a safe haven with respect to the US dollar. While other work has investigated the role of gold as a hedge and a haven for bonds and equities, no study has tackled the same subject with a specific focus on exchange rates.

In sum, the hope is that this thesis offers a useful contribution to the fields of international finance and applied econometrics. Findings are clear and well-defined and open up a number of potential avenues for future research.

Chapter 2

Misalignment and intervention

This chapter estimates a reaction function for official Japanese intervention in the currency markets between April 1991 and March 2006. The sample data is daily with intervention data provided by the Japanese Ministry of Finance. Estimation results show that intervention during the sample period conforms to a model in which the monetary authorities intervene in order to prevent the yen from straying too far from an equilibrium value defined by a capital-enhanced version of purchasing power parity. Predictability is good. A generalised ordered logit model provides the empirical framework and results show that studies of intervention that ignore violations of the proportional odds assumption are likely to suffer from specification error.

2.1 Introduction

Official intervention in the currency markets has in recent years been labelled variously as unsuccessful, ineffective, and even counterproductive,[1] and yet intervention remains an important tool of exchange-rate policy for many countries today. At the start of 2009, Russia, Brazil, Mexico, South Korea, India and Indonesia were all intervening actively in the currency markets. Meanwhile, Japanese authorities came under increasing pressure to intervene in the currency markets in what would represent Japan's first official intervention in five years.[2]

Of the world's biggest economies, Japan stands out for having had the most

[1] For recent studies of the effectiveness of intervention in the currency markets see, among others, Fatum and Hutchison (2003), Ito (2002), King and Fatum (2005), McLaren (2002), Neely (2005a), and Sarno and Taylor (2001).

[2] The Economist (2009) discusses recent pressure for Japanese intervention.

active intervention policy during the past 20 years, engaging in more than US$620bn-worth of interventions in the currency markets since 1991. US interventions over the same period amounted to less than a tenth of this value. Yet, despite Japan's activism in the currency markets, very little is known about what drives Japan's interventions. Do Japan's monetary authorities intervene in the foreign-exchange markets in order to keep the yen close to a fixed, predetermined value? To keep it within fixed bounds of tolerance? Within time-varying bounds of tolerance?

Common wisdom has it that most monetary authorities aim to maintain a stable exchange rate that is consistent with underlying economic fundamentals. But Japan's monetary authorities do not disclose publicly the precise aims of their intervention policy.[3] All we have is evidence of Japan's past interventions—in the form of the recorded dates of intervention, the amount of yen purchased or sold on the given dates, and the partner currencies involved in the intervention transactions—to offer us clues as to the ultimate aim of intervention policy. A small number of academic studies have used this information to construct plausible models, known as reaction functions, of Japan's intervention policy. None has been particularly successful.

Ito and Yabu (2007) estimate a reaction function for Japanese intervention with a model that assumes that the Bank of Japan, which has operational control of intervention policy in Japan, intervenes in order to keep the national currency, the yen, close to trend historical values.[4] Covering the period 1991 to 2002 and using in-sample prediction, the reaction function proposed by Ito and Yabu (2007) predicts at best 18% of actual interventions. At worst it predicts 18 instances of intervention when no intervention actually took place.

Frenkel et al. (2002) estimate a reaction function that assumes Japanese interventions in the foreign-exchange market respond to deviations of the yen-dollar exchange rate from a short-term, and a long-term, exchange-rate target. Their model anticipates correctly 52% of actual interventions. Ito (2002) estimates an intervention reaction function without appealing to any theoretical framework, while Almekinders and Eijffinger (1996) propose a friction model as the best description of Japan's intervention policy, whereby pursuit of an optimal intervention policy is compromised by friction costs that are associated with the political implementation of policy.

This study adopts the friction-model approach of Almekinders and Eijffinger (1996) but incorporates a number of additional features that add significantly to the empirical performance of the model. An intervention reaction function is

[3]For surveys of intervention policy see, for example, Edison (1993), Dominguez and Frankel (1993) and Sarno and Taylor (2001).

[4]The Bank of Japan acts as an agent for the implementation of intervention policy. Policy itself, and the intervention decision, is determined by Japan's Ministry of Finance.

estimated for Japan covering the period April 1991 to March 2006. The data is daily.

This study makes two main contributions to the existing literature. Firstly, it tests the hypothesis that the aim of optimal intervention policy in Japan is to prevent the nominal exchange rate from straying too far from its medium-run equilibrium. *Too far* is defined in relation to a tolerance zone for the exchange rate, while *medium-run equilibrium* is defined in terms of a capital-enhanced version of purchasing power parity.[5] By incorporating a measure of exchange-rate equilibrium explicitly within the intervention reaction function, this study improves over other studies that assume the monetary authorities desire nothing more than a backward-looking adjustment towards trend.

The second main contribution is empirical: this study adopts a *partial proportional odds model* of intervention that allows for asymmetry in the intervention objective function. That is, by employing a partial proportional odds model it is possible to test the idea that the monetary authorities in Japan do not react symmetrically to deviations in the value of the yen from its equilibrium value. By taking this flexible approach to estimation this study improves over other studies that do not allow for violations of the *proportional odds assumption*.

Key findings can be summarised as follows. (i) Between 1991 and 2006 Japan did, indeed, intervene in the currency markets in a manner that suggests its interventions were timed in order to prevent the yen from straying excessively from its medium-run equilibrium value against the US dollar. (ii) Medium-run equilibrium can be defined in terms of a capital enhanced version of purchasing power parity. (iii) The Japanese monetary authorities do not react symmetrically to deviations in the value of the yen from its target value. (iv) Studies of intervention that ignore violations of the proportional odds assumption are prone to specification error.

The rest of this chapter is organised as follows. Section 2.2 surveys the stylised facts: the movement of the yen against the US dollar during the sample period and the timing and size of intervention activity.[6] Section 2.3 derives a reaction function for Japanese intervention while Section 2.4 describes the estimation methodology. Section 2.5 describes the data. Section 2.6 presents the empirical results and Section 2.7 offers concluding remarks.

[5] The capital-enhanced version of purchasing power parity is outlined by Juselius (1991, 1995), MacDonald and Marsh (1997, 1999) and Juselius and MacDonald (2000b,a).

[6] Note that the sample period is determined solely by availability of data. Japan's Ministry of Finance discloses information on all daily intervention activities after 01 April 1991.

2.2 Stylised facts

This section presents an overview of both Japanese intervention in the currency markets and movements in the value of the yen against the US dollar during the sample period.

Between 1991 and 2006 the yen experienced a number of significant fluctuations against the US dollar. Figure 2.1 shows the major highs and lows. In April 1995 the yen hit a post-war high of 81 yen per US dollar as diplomatic frictions over US-Japanese trade policy sparked heavy selling of US dollars. Three years later, in August 1998, the Japanese currency slumped to 148 yen per US dollar, its weakest level during the 15-year sample period, amid knock-on effects from the financial crisis that struck Asia in the late 1990s. Throughout the sample period the yen averaged 115 yen per US dollar, and traded, mostly, between 100 yen and 140 yen per US dollar.

Figure 2.2 shows the extent to which the Bank of Japan intervened in the currency markets between 1991 and 2006 in order to either weaken, or strengthen, the yen. Positive amounts of intervention indicate purchases of yen. Negative amounts indicate sales of yen. The average value of a single intervention during the sample period is US$1.8bn.

Five main features characterise Japan's intervention behaviour during the sample period. First, intervention is infrequent: on most days during the sample period (91% of all days) there is no intervention.

Second, when intervention does occur, sales of yen are undertaken when the Japanese currency is strong relative to its average value and purchases of yen are undertaken when it is relatively weak. This chimes with the commonly encountered explanations for intervention being to either prevent too much appreciation or too much depreciation. Too much appreciation, the argument goes, would harm exporters, while too much depreciation would harm importers and confidence. More often than not, monetary authorities have justified intervention as a means of helping to maintain a stable exchange rate that is consistent with underlying economic fundamentals. For surveys see, for instance, Edison (1993), Dominguez and Frankel (1993) and Sarno and Taylor (2001).

Third, if intervention occurs on day $t-1$, the direction of intervention subsequently, on day t, is identical. That is, purchases of yen follow purchases of yen, and sales of yen follow sales of yen. There are no instances when a purchase is followed directly by a sale, or a sale by a purchase.[7]

Fourth, intervention policy is asymmetric: there are many more instances of yen-selling intervention, than yen-buying intervention.

[7] Neely (2000) provides a comprehensive discussion of the common practical features of currency intervention.

Figure 2.1: Japanese yen per US dollar

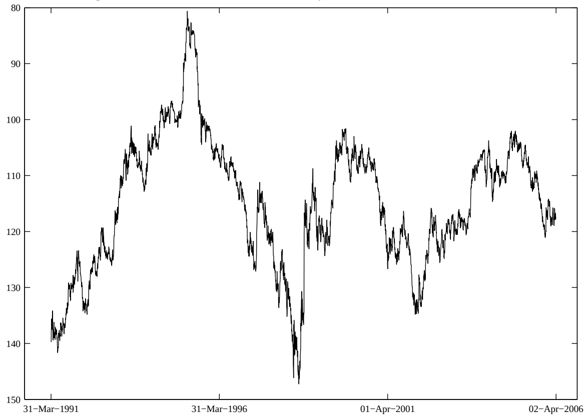

Notes: Japanese yen in terms of yen per US dollar. Reverse scale. Frequency is daily. Five-day week. New York Close. Source: Bloomberg.

Figure 2.2: Japanese intervention in the currency markets

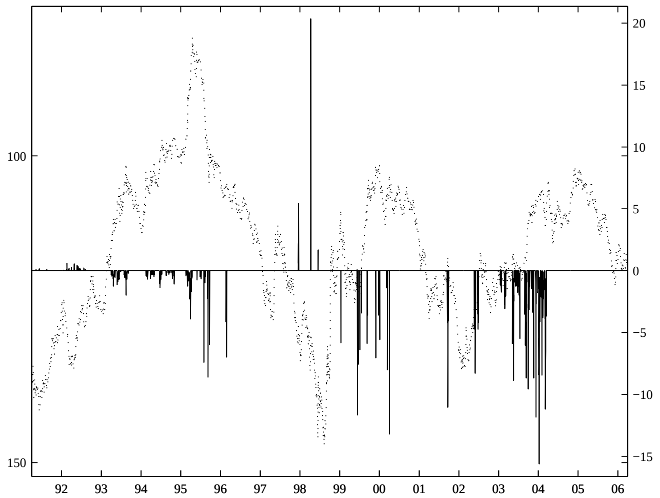

Notes: Dotted line shows Japanese yen in terms of yen per US dollar, reverse scale, measured on the left-hand axis. Solid line shows the amount, in US dollar billions, of Japanese intervention in the currency markets, measured on the right-hand axis. Frequency is daily. Positive amounts of intervention indicate purchases of yen. Negative intervention indicates sales of yen. Source: Japanese Ministry of Finance and Bloomberg.

Fifth, during the sample period the last recorded intervention occurred on 16 March 2004 when the Bank of Japan stepped in to sell Y68bn in exchange for US dollars. The sample records no subsequent instances of intervention. For a discussion of the policy debate involved in the the design of Japanese intervention policy in 2004, and the curtailment of interventions, see Taylor (2010).

Figure 2.3 illustrates the evolution of the value of the yen from 1991 to 2010. The figure shows the cessation of Japan's interventions in 2004 and the resumption in 2010 (with a single instance of yen-selling intervention valuing US$24bn occurring on 15 September 2010). The figure, on its own, gives no clear indication of whether the break in interventions after 2004 is consistent with previous breaks, or whether it represents a change in intervention regime. As such, this chapter looks only at the early period of interventions, allowing the sample to run to 2006, not 2010.

2.3 Model of intervention

This section presents a model of official intervention in the currency markets for Japan between 1991 and 2006. The model is presented in three stages. First, Section 2.3.1 outlines an intervention loss function for the central bank, whereby policy loss is driven by the central bank's desire to minimise deviations of the exchange rate from a time-varying target. Section 2.3.2 discusses the exchange-rate target. Section 2.3.3 introduces into the model a role for policy friction, which helps to explain why intervention occurs intermittently rather than continually.

2.3.1 Loss function

Most studies of the objectives of central-bank intervention, if they take a reaction-function approach to the subject, tend to construct these functions without appealing to any particular theory. Edison (1993) discusses many of these atheoretical approaches. A handful of investigations do, however, adopt intervention reaction functions that are derived from theory. Almekinders and Eijffinger (1996), for instance, combine a model of the exchange rate with a loss function for the central bank in order to derive an intervention reaction function. The loss function is fashioned around the idea that the central bank would prefer, if able, to minimise deviations of the exchange rate from a target level. The extent of policy loss is assumed to increase with both negative and positive deviations from the target level.

Another formal derivation of the intervention reaction function is provided by Frenkel et al. (2002). The authors assume, like Almekinders and Eijffinger (1996), that the central bank conforms to a policy loss function whereby it aims to minimise deviations of the exchange rate from a target level. They also assume that the central bank aims to minimise deviations from a target level of intervention, where the intervention target is set (for a flexible exchange-rate regime) at zero.

The assumption of an intervention target may, on the surface, seem fairly innocuous. However, the implication is that intervention policy is independent of exchange-rate developments. This would be the case if the central bank were to pursue an objective aimed at either adding to or depleting its stock of foreign-exchange reserves at a pre-determined rate. But in reality, such an objective is highly uncommon—or it is, at least, in developed countries. This suggests that there is little justification, here, for including an intervention target in the central bank's loss function. Indeed, this study takes the view, like Ito and Yabu (2004), that a more plausible loss function will include a target for the exchange rate and nothing else. More specifically, the loss function is assumed to take the following form

$$MinE_{t-1}(L_t^{CB}) = E_{t-1}(s_t - s_t^T)^2 \tag{2.1}$$

where s_t is the log of the yen-per-dollar spot exchange rate at date t (which in this case is the close of the New York trading day), where s_t^T represents the exchange-rate target at date t, and where the implication of the loss function in this form is that the central bank's expected policy loss increases more than proportionately with both positive and negative deviations from the exchange-rate target.[8]

Note, E_{t-1} implies that expectations are formed on the basis of information available to both the central bank and market agents on day $t-1$. This assumption is not without its faults. It has been criticised in particular by Sarno and Taylor (2001), who suggest it is not appropriate to assume that both the central bank and market agents base their expectations on the same information. If both the central bank and market agents use the same information to form expectations, then market agents have no incentive to monitor the central bank because monitoring will provide no additional information. Sarno and Taylor (2001) argue that in practice market agents do monitor central banks. Indeed, financial markets in developed countries subject their central banks to an immense amount of scrutiny.

On the surface, therefore, it seems that Sarno and Taylor (2001) have a

[8]Date t is centred on the New York closing rate because, as explained by Ito (2002), Japanese intervention on day t can be carried out during the Tokyo trading day, the European trading day, or the New York trading day.

point. It makes sense to think that that the central bank may have an informational advantage and that it will know more about its own future actions than will market agents. This will be the case if official interventions are not announced publicly but are, instead, undertaken secretly in order to increase effectiveness. Such behaviour would be in keeping with theories espoused by, for example, Balke and Haslag (1992), who suggest that in order for intervention to be effective the central bank must maintain an informational advantage.

The problem with this idea and, by association, the flaw in the argument put forward by Sarno and Taylor (2001), is that there is good evidence to suggest central banks do not operate with information that is any better than that available to market agents. Humpage (1997), for instance, finds that US intervention in the currency markets between 1990 and 1997 did not convey to market agents any information that they would not have possessed otherwise. The central bank did not, in short, possess an informational advantage.

It is the supposition of this study that the findings of Humpage (1997) are a fair description of the balance of information in the intervention process and that *neither the central bank nor market agents wield an informational advantage.* Expectations, as a result, are formed on the basis of information available to both the central bank and market agents at time $t-1$, and this behaviour is reflected in the formulation of the loss function represented by Eqn.(2.1).

Implicit in Eqn.(2.1) is the idea that the monetary authorities aim to use intervention to minimise the loss function. This does, of course, leave unanswered the question of just how, in the absence of intervention, does the exchange rate behave? It is assumed here that the central bank believes that the exchange rate behaves as a random walk and that intervention at date t, should it occur, has a contemporaneous effect on the exchange rate. The exchange rate can, therefore, be defined as

$$s_t = s_{t-1} + \lambda Int_t + u_t \tag{2.2}$$

where the implication is that the yen-per-dollar level of the exchange rate is determined by the exchange rate's own recent past s_{t-1}, by intervention Int_t, and by u_t, a white-noise error. Int_t takes a positive value to represent yen purchases and a negative value to represent yen sales.

If intervention is successful in causing not just a slowing of the exchange rate's movement, but an actual reversal, then λ should be negative. To see this, note that if yen-selling intervention by the monetary authorities (represented by a negative value for Int_t) causes, as intended, a depreciation in the value of the yen (with $s_t - s_{t-1} > 0$) then λ should, logically, take a negative sign.

In a survey of 22 monetary authorities, Neely (2000) found that 90% of authorities say they intervene sometimes or always in order to resist short-run trends in the exchange rate. Meanwhile, 67% of monetary authorities agreed with the premise that they intervene in order to return exchange rates to "fundamental values."

To treat the exchange rate as a random walk, or more generally as a martingale process, implies that the Japanese authorities accept the thesis of Meese and Rogoff (1983) who find that a random walk provides an adequate descriptive model of the behaviour of the exchange rate. Assuming random-walk behaviour implies that intervention may affect the exchange rate via the co-ordination channel, a channel of influence proposed by Taylor (1994, 2004, 2005). The co-ordination channel implies that central-bank intervention drives the exchange rate towards its fundamentals-based value by resolving a failure of co-ordination in the currency markets: if misalignments of the exchange rate are caused by non-fundamental factors (such as the influence of chartist traders) and it is only a failure of co-ordination among market participants that is preventing the exchange rate from returning to equilibrium, then official intervention may prove to be effective by acting as a co-ordinating signal that causes speculators to enter the market and resume true, fundamentals-based trading decisions that will return the exchange rate towards a level that is consistent with fundamental values.

Of course, the co-ordination channel is not the only channel through which intervention can influence the exchange rate. Two other channels are the signalling channel and the portfolio-balance channel. For further discussion see Sarno and Taylor (2001). It is also possible that intervention could influence the exchange rate via market-microstructure processes. See Lyons (2001). But these processes do not seem to account for prolonged effects on exchange rates—see Reitz and Taylor (2008)—and neither the signalling channel nor the portfolio-balance channel receive much empirical support in the current academic literature. The co-ordination channel does, therefore, seem to represent a plausible mechanism for the influence of intervention on the exchange rate where the exchange rate is assumed to behave like a random walk.

Autoregressive conditional heteroscedasticity

One drawback in adopting a random-walk model is that it is not the most flexible description of the exchange rate, especially when the empirical framework is one of daily data, as it is here. A more flexible model would, perhaps, be one that acknowledges the fact that a common form of heteroscedasticity in daily exchange-rate data is autoregressive conditional heteroscedasticity (meaning that large and small residuals tend to come in clusters). An obvious approach would be to adopt the model proposed by Bollerslev (1986), which captures generalised autoregressive conditional heteroscedasticity, otherwise known as a GARCH model.

It is well documented that GARCH models offer good descriptions of the returns from daily spot exchange rates. Baillie and Bollerslev (1989a), for instance, find that for six different currencies a GARCH model with daily dummy variables and conditionally t-distributed errors provides a good description of the kurtosis and time-dependent conditional heteroscedasticity in the exchange-rate data. Studies by Taylor (1986) and McCurdy and Morgan (1988), among many others, find that the most appropriate formulation of the GARCH model for daily exchange-rate data is a GARCH(1,1) model.

All this support for a GARCH approach when modelling exchange rates has led, not surprisingly perhaps, to GARCH techniques being employed by a number of intervention studies in recent years. Dominguez (1998) was the first to use GARCH conditional variances (alongside implied volatilities from currency options) to study the relationship between central-bank intervention and exchange-rate volatility, finding that intervention tends to increase volatility—or that it did, at least, for the US, Japan and Germany between 1977 and 1994. Using similar GARCH methods, Ito (2002), Frenkel et al. (2005), Hillebrand and Schnabl (2005), and Harada and Watanabe (2005), look specifically at Japan, investigating the impact of Japanese foreign-exchange intervention on the daily volatility of the yen's exchange rate against the US dollar.

Endogeneity

Unfortunately, using a GARCH framework to analyse central-bank intervention in the currency markets does have its drawbacks. One important drawback has been highlighted by Hillebrand et al. (2006), who note that GARCH approaches do not deal successfully with the problem (encountered in all intervention studies) of endogeneity.

Endogeneity is a problem in intervention studies because the close correlation that exists between intervention and exchange-rate movements does not necessarily imply that intervention is the cause of changes in the exchange rate. Correlation could imply the opposite—that exchange-rate movements cause the central bank to intervene. Hillebrand et al. (2006) suggest that GARCH approaches, especially in the GARCH mean equations, have failed to resolve this issue of endogeneity and point out that, partly as a result, the GARCH approach to intervention has lost some of its appeal among researchers in recent years and has been replaced, to a certain extent, by event studies.

Event studies

Event studies look at the behaviour of the exchange rate not over a continuous time series but over small windows of data clustered around periods of intervention. Event studies avoid the problem of endogeneity so long as two things hold true: first, so long as there is no error in the measurement of the timing of intervention; and second, so long as the frequency of the data is high enough (eg, intra-day data) to preclude the monetary authorities from reacting to market developments within the data interval.[9] If these two assumptions hold true, then there is no contemporaneous impact of the exchange rate on intervention and, as such, no endogeneity.

However, intra-day event studies are not without their limitations. One major limitation, as highlighted by Neely (2005b), is that only one country (namely Switzerland) has so far released official intra-day data on the precise timing of interventions, and as a result intra-day event studies are restricted to examining just one sample. Inferences cannot, therefore, be assumed to be particularly robust. A second limitation is the arbitrary choice of data-window size. The data window needs to be large enough to register the full effects of intervention and in intra-day event studies common wisdom has it that a two-hour period either side of the intervention should provide a window that is roughly large enough to capture all necessary effects (see, for instance, Payne and Vitale (2003)). But as Neely (2005a) notes, it is entirely possible that intervention has its full effect over days, if not weeks. Two fifths of central bankers surveyed by Neely (2000) said they believe that intervention takes at least a couple of days to have its full effect.

Structural approach to intervention

Perhaps the biggest limitation of event studies is, however, their inability to say anything useful about causality. The most that event studies can do is paint a statistical picture of the behaviour of the exchange rate around periods of intervention. Event studies reveal nothing about the reason for the observed behaviour. In order to throw light on the underlying causality, it is necessary to assume some structure for the system—and perhaps construct a *structural equation model* (Haavelmo, 1943).

Hillebrand et al. (2006) take a structural approach to intervention, using the concept of *realised volatility* as proposed by Andersen and Bollerslev (1998) to identify explicitly the effect of intervention on exchange rates. Kearns and Rigobon (2005) and Neely (2005b) also take structural approaches in their analyses of the exchange rate's response to intervention. All of these studies, since

[9] See Neely (2005b) for a full discussion.

they aim to create a system in which all structural parameters are identified, have the advantage of being able to deal directly with the problem of endogeneity.[10]

Single-equation approach

This study employs a single-equation approach, with structure, to the analysis of intervention. A number of other studies have employed a single-equation approach to intervention. One technique is to model intervention as a binary-choice dependent variable, with the dependent variable existing in one of two states: intervention or no intervention.[11]

Another valid single-equation technique is to approach intervention using a friction model, as first proposed by Rosett (1959). In Rosett's model the dependent variable takes the value of zero if the explanatory variables are close to their desired levels. If the dependent variable is intervention, the implication is that intervention occurs only when the explanatory variables stray outside an empirically-determined tolerance zone. The central bank, in other words, will maintain a policy of non-intervention so long as the factors that condition its decision to intervene do not breach pre-determined thresholds.

Almekinders and Eijffinger (1996) use a friction model to analyse the intervention policies of both the Bundesbank and the Federal Reserve during the late 1980s. An intervention reaction function is derived by combining a GARCH model of the exchange rate with a loss function for the central bank. The dependent variable in the reaction function is intervention amount. Ito and Yabu (2004) also adopt a friction-model approach to intervention, but the dependent variable in their reaction function is not intervention amount. Instead, it is an intervention indicator function, which can take one of three values (1, 0, or -1), representing either the sale of foreign currency (1), no intervention (0), or the purchase of foreign currency (-1). This seems to be a flexible approach and one that will be pursued in this study.

Endogeneity in the single-equation approach

Of course, the important question to ask is, Can a single-equation approach deal adequately with the problem of endogeneity? On the surface, using a single-equation approach might seem to free the economist from the problem of endogeneity because there will be less need to make assumptions about the structure of the economy. The single-equation approach of the event study, for instance,

[10] Since the structural parameters are identified, it is possible to estimate the parameters consistently.

[11] Baillie and Osterberg (1997), Dominguez (1998) and Frenkel et al. (2002) provide examples of this binary-choice approach.

makes precious few assumptions about the system's structure. However, Neely (2005a) suggests that event studies do, in fact, make hidden assumptions about structure that can lead to simultaneous-equations bias if intervention affects exchange-rate returns contemporaneously (with a contemporaneous interaction occurring if, for instance, daily data is employed rather than intra-day data). Furthermore, it is difficult to correct for this simultaneous-equation bias using an instrumental-variables approach because reliable instruments are hard to find.[12]

Clearly, event studies suffer from a number of weaknesses when it comes to dealing with endogeneity. While this study is not an event study, it still needs to contend with many of the same problems that cause event studies to be vulnerable to simultaneous-equations bias when deployed as tools for intervention analysis.

The obvious problem is that, as can be seen from equation Eqn.(2.2), intervention does, in this model, have a contemporaneous impact on the exchange rate. However, the reason why the results of this study are not tainted with simultaneous-equations bias is due to the fact that the intervention variable, Int_t, is not intervention amount, but is instead an indicator function.

Representing intervention with an indicator function means that using data of a daily frequency, as is done in this study, does not cause an endogeneity problem as would, normally, be expected. Under normal circumstances it would make sense to expect daily data on intervention to generate an endogeneity problem because during the course of a full day of trading hours, the central bank will be able to react to any given exchange-rate development and, as a result, there will be a contemporaneous interaction between intervention and exchange-rate returns.

However, this contemporaneous interaction does not result in an endogeneity problem so long as intervention is represented by an indicator function. The reason why boils down to the timing of the intervention decision: the decision to intervene on any given day occurs prior to the start of trading hours and so it is not possible for the intervention decision to react contemporaneously to events in the currency market during the day in question. If, however, intervention were to be represented by intervention amount rather than an indicator function then things would be very different. A contemporaneous interaction would, in fact, occur. The reason is that while the decision to intervene on any given day is taken before the start of market trading hours, the decision as to just how much to spend on intervention is taken during, not prior to, trading hours. Exchange-rate developments and intervention amount would, therefore, be free to interact

[12] As Neely (2005b) notes, it is difficult to find instruments that are correlated reliably with intervention but not with exchange-rate returns.

contemporaneously.

To support this argument it is necessary to be more explicit about the daily timeline of events that form the mechanics of the model in this study. Japan's intervention decision is assumed to take place during the three hours prior to the opening of the Tokyo currency market—that is, between 0600hrs and 0900hrs Tokyo time. The important thing about this three-hour window is that it represents the gap, in terms of time, between the close of New York trading day (at 0600hrs Tokyo time, or 0700hrs if daylight-saving time is being observed in the US) and the opening of the Tokyo trading day (at 0900hrs).

The implication of all this is that the decision to intervene on day t is made prior to the opening of market-trading hours and is based on all the information available at the end of day $t-1$. During the three hours prior to market opening, the decision is taken to intervene or not, but even if intervention is indeed sanctioned, there will not, necessarily, be a decision taken on the precise amount to be spent on intervention. The precise amount of intervention on day t will be decided during the course of trading on day t and will depend on the movement of the exchange rate.[13] Intervention amount will, as a result, be associated with an endogeneity problem. But there will be no such problem associated with an intervention indicator function.

Adopting an intervention indicator does unfortunately have its drawbacks, the biggest of which is a loss of efficiency. A loss of efficiency is experienced because by adopting an indicator function, which indicates only the direction of intervention, we are ignoring information on intervention amount which is both available and quantifiable. As a result the information set is only partial. This is not ideal. But the benefits of the indicator function—chiefly its usefulness as an aid to avoiding the endogeneity problem—are considered to be big enough to outweigh the drawbacks of forcing self-imposed limits on the information set. This study proceeds, therefore, with an intervention indicator function.[14]

Returning to the mathematical derivation of the model being used in this study, it is possible, using Eqn.(2.2) and Eqn.(2.1), to derive an intervention reaction function and an expression for optimal intervention, Int_t^*. The loss function Eqn.(2.2) is minimised subject to the constraint represented by the exchange rate Eqn.(2.1), leaving optimal intervention to be defined as

$$Int_t^* = -\frac{1}{\lambda}(s_{t-1} - s_t^T) \tag{2.3}$$

[13]In a survey of 22 monetary authorities, Neely (2000) reports that 21 authorities say that market reaction sometimes or always affects the size of any given intervention.

[14]One possible means of avoiding this loss of efficiency would be to construct an *intervention index*, whereby an intervention is defined to consist of those sales or purchases that occur over consecutive periods. This possibility is not pursued here but is left for future research.

2.3.2 Target exchange rate

On first inspection, the intervention reaction function represented by Eqn.(2.3) is no different to the expression for optimal intervention proposed by Ito and Yabu (2004). However, there is an important difference. The difference is that the target exchange rate, s^T, is constructed, in this study, in a manner that allows for a more sophisticated target.

Weighted target

In the intervention reaction function proposed by Ito and Yabu (2004), the monetary authorities use a target for the exchange rate that is calculated as a weighted average of past exchange rates. More specifically, the authors construct a composite measure of the target exchange rate that includes the average of the spot exchange rate during the previous day, s_{t-2}, the average of the spot exchange rate during the preceding four weeks (or in practice, 21 business days), s_M, and the average of the exchange rate during the preceding 12 months, s_Y. In this form, the target exchange rate can be represented as

$$s_t^T = \alpha_1 s_{t-2} + \alpha_2 s_M + \alpha_3 s_Y \tag{2.4}$$

where $\alpha_1 + \alpha_2 + \alpha_3 = 1$. The implication is that if the central bank is focused on long-run stability of the exchange rate then α_3 will take a value close to unity. When short-run stability is the priority, α_1 will take a value close to unity, and when the medium term is the main focus, α_2 will lie close to unity.

The weighted average approach to target exchange rates, represented by Eqn.(2.4), can be thought of as a generalisation of the construction used by Almekinders and Eijffinger (1996) where the central bank's target is a long-run target and α_3 is equal to unity. There are, of course, alternatives to the weighted-average approach. Baillie and Osterberg (1997), for instance, assume that the target exchange rate is a simple, static nominal value (specifically, the authors assume that between 1985 and 1990 the world's industrial countries agreed on target, nominal values for both the dollar-Deutschemark exchange rate and the dollar-yen exchange rate). Ito (2002) assumes that the long-run equilibrium exchange rate for Japan between 1991 and 2001 was 125 yen per US dollar.

Other studies mirror the approach of Almekinders and Eijffinger (1996), assuming that if the central bank does conduct intervention policy according to an exchange-rate target then that target will hold only over the long run, not over the short run or medium run. Artus (1977), Neumann (1984) and Knight and Mathieson (1983) all assume that the central bank pursues a single, long-run target. But this target is not, as per Ito and Yabu (2004), a twelve-month moving average. It is a target based on purchasing power parity.

Purchasing power parity

It is not illogical to think that that the monetary authorities may want to guide the exchange rate towards a value that brings into line international purchasing power. Indeed, Dominguez and Frankel (1993) find that purchasing power parity was an important part of intervention policy for America's Federal Reserve between 1982 and 1988 and Neely (2002) finds, similarly, that the US monetary authorities tend to intervene to support the US dollar when it is undervalued relative to PPP and sell it under opposite circumstances. Frenkel et al. (2002), meanwhile, use a real-exchange-rate target derived on the basis of purchasing-power-parity conditions for Japan between 1991 and 2001. What is more, even though neither Esaka (2000) nor Galati and Melick (1999) look specifically at the question of purchasing power parity as a target for Japanese exchange-rate policy, both studies conclude that Japan intervened in the currency markets during the 1990s in a manner consistent with there being some implicit target level for the yen-dollar exchange rate.

One advantage of a weighted-average target is that it is easy to define and test econometrically. Testing for the validity of PPP as a target is more challenging. The biggest challenge arises if the researcher is intending to devise a model based on daily data. Unfortunately, the components of any PPP measure of an exchange rate, namely domestic and foreign prices, tend to be reported on a monthly basis and, therefore, if monthly PPP data is used in a model of the daily exchange rate, it will lead to a target that proves, as noted by Ito and Yabu (2004), to be "sticky". One answer to this problem is to interpolate the monthly data into a daily format, as is done by Neely (2006) in his study of US intervention. But interpolation is far from ideal. Any type of interpolation, however sophisticated, means making strong assumptions about how the interpolated data behave when observed at higher frequencies, and this criticism does, perhaps, carry even more weight when the interpolated data is for such a notoriously inexact measure of exchange-rate equilibrium as PPP.

Purchasing power parity, as a presumed target for intervention policy, is not without its flaws. Although PPP is used regularly by private-sector economists and popular commentators as a rough-and-ready guide to a currency's equilibrium value, it would be rash to assume that central banks cannot afford the computational effort to come up with a better measure of exchange-rate equilibrium. It makes sense, then, to look for an alternative to PPP as a plausible target of intervention policy and, perhaps, to introduce more realistic formulations for both short-term and medium-term targets.

One aim of this study is to construct a model of intervention that has, embedded within it, an exchange-rate target that is more than just a short-hand expression of equilibrium devised solely in order to be computationally simple—or more, in other words, than just a weighted average of previous exchange-rate values, as per Ito and Yabu (2004). Instead, the aim is to allow for an exchange-rate target that mirrors as closely as possible the monetary authorities' perception of exchange-rate equilibrium.

What we mean by equilibrium exchange rate

Exchange-rate equilibrium is not a straightforward concept to model. There are short-, medium-, and long-run concepts of equilibrium. What is more, different measures of exchange-rate equilibrium are appropriate for different situations. A bewildering array of acronyms, representing different interpretations of exchange-rate equilibrium, have become established in the relevant literature, yet still there is a debate over the optimality of equilibrium, over its determination, over its evolution and even over its existence.[15]

For all practical purposes, however, the concept of exchange-rate equilibrium can, in fact, be employed successfully if the modeller takes into account the relevant time horizon. For instance, an equilibrium that pertains over the short run will not necessarily pertain also over the medium run or long run. For a full discussion of the relevance of the time horizon, see Driver and Westaway (2003), but for the purposes of this study it is necessary, here, to highlight just a handful of salient points about the long, medium and short run.

First, an exchange rate that is in equilibrium in the short run can be defined as being in an equilibrium that, in line with the thinking of Williamson (1983), satisfies the condition that all fundamental determinants are at their current values after netting out the influence of random effects such as bubbles.

A medium-run equilibrium is more difficult to define. In its simplest form, a medium-run equilibrium will exist when the economy is in balance both internally and externally. External balance implies that the current-account gap must be *sustainable*, in the sense that it must be consistent with convergence, eventually, to a stock-flow equilibrium. Unfortunately there is no hard-and-fast rule for defining what is meant by *sustainable*, which highlights the fact that a key feature of the internal-external-balance approach to exchange-rate equilibrium is that a large degree of judgment, or normative manipulation, is involved in defining external balance. The calculations involved in finding internal balance are, thankfully, slightly less prone to normative influence. Internal balance

[15] See Milgate (1998) for a discussion of the concept of equilibrium and its development over time.

occurs when demand is at its supply-potential and the economy is running at its natural speed limit. As such, an internal equilibrium can be defined as occurring when the economy is operating with no output gap and when unemployment stands at a steady-state level above which inflation will fall and below which inflation will rise (ie, at a non-accelerating-inflation rate of unemployment, or NAIRU).

Long-run equilibrium, meanwhile, can be defined as occurring when the economy reaches a state of stock-flow equilibrium. To get to this state may, of course, take years, or even decades. The important point is that stock-flow equilibrium in this context will occur when there is no reason for the level of asset stocks to change as a proportion of GDP. This is different from exchange-rate equilibrium in the medium-term when there is no stock-flow equilibrium. In the medium term, equilibrium can occur at any prevailing levels of national wealth. But in the long run, net wealth must be stock-flow consistent.

A modelling framework for exchange-rate equilibrium

Of course, these definitions of long-, medium- and short-run equilibrium are of little use unless they can be represented in a modelling framework. Clark and MacDonald (1999) outline a useful framework. Keeping to the spirit, if not the letter, of the approach taken by Clark and MacDonald (1999), the exchange rate can be represented as

$$s_t = \beta' Z_t + \tau' T_t + \varepsilon_t \tag{2.5}$$

where s_t is the exchange rate at time t; where Z_t represents a vector of fundamentals that are expected to have persistent effects on the exchange rate not just over the medium term (ie, over the business cycle) but also over the long term; where T_t is a vector of transitory, or short-run variables (including dynamic effects from the fundamentals, Z_t); where ε_t is a random error and where β' and τ' are vectors of coefficients.

Using this modelling framework, it is possible to describe, mathematically, what is meant by long-, medium- and short-run equilibrium for an exchange rate. A short-run equilibrium can, for instance, be defined as

$$s_t^{SEQ} = \beta' Z_t + \tau' T_t \tag{2.6}$$

where fundamentals are at their current values, transitory effects are present, but where there are no unanticipated shocks. Another valid way of representing short-run equilibrium, suggested by Driver and Westaway (2003), is to assume that equilibrium in the short run reflects fundamentals at their current values but precludes a role for transitory effects, such that

$$s_t^{SEQ*} = \beta' Z_t \tag{2.7}$$

and where, in the lexicon of Williamson (1983), the exchange rate is at a *current equilibrium.* In the same way, exchange-rate equilibrium in the medium run, s_t^{MEQ}, can be defined as

$$s_t^{MEQ} = \beta' \hat{Z}_t \tag{2.8}$$

where $\hat{Z}_t$ represents fundamentals at their trend, medium-run values (in the process of adjusting towards a long-run equilibrium). Ultimately, when fundamentals do reach their steady-state, long-run values $\bar{Z}_t$, exchange-rate equilibrium can be represented by

$$s_t^{LEQ} = \beta' \bar{Z}_t \tag{2.9}$$

What must be noted is that even though these models for long-, medium- and short-run equilibrium represent distinct concepts, at any given point in time they will all hold true.

Equilibrium as a target

While a nation's monetary authorities may, as part of normal operating procedure, measure the actual value of the exchange rate against estimates of its short-, medium- and long-run equilibrium values, what is less likely is that any intervention in the currency markets by the monetary authorities will aim to manipulate the exchange rate towards a value that will satisfy equilibrium at all time horizons. What seems more likely is that when intervening in order to drive the exchange rate towards a target value the monetary authorities will have in mind just one time horizon. The question is, which time horizon? Short, medium, or long?

On balance, it seems unlikely that any sensible central bank would be so bold as to think that with just a handful of daily interventions it could force the national currency into a position of long-run equilibrium. Most central bankers are amply aware that currencies can and do deviate from their long-rum equilibrium values due to the existence of persistent influences on the exchange rate over the business cycle that make it undesirable to pursue blindly a target consistent with Eqn.(2.9). For example, pursuing a long-run target of purchasing power parity without acknowledging the fact that real factors can affect the real exchange rate over the business cycle (as argued by, for instance, Mussa (1986)) would risk harming the economy. The productivity-bias effect on exchange rates outlined originally by Bela Belassa and Paul Samuelson is a well-known real determinant of any real exchange rate. In Japan it has, in fact, been an important influence on the exchange rate even in the long run.[16]

All this suggests that a central bank deploying currency-market intervention in order to pursue an exchange-rate target is likely to avoid choosing a target based on long-run equilibrium. On the surface, a target based on short-run equilibrium might seem more likely—and, indeed, more attainable. Pursuing a target that compensates for short-term bubbles (compensates in other words for random disturbances, with the target taking a form outlined by Eqn.(2.6)) would represent an exchange-rate policy aimed at eliminating misperceptions about fundamentals.[17] A policy such as this makes sense in principle. In practice, however, the flaw in such a strategy is that bubbles are hard to identify.[18] A policy aimed at offsetting the effect of bubbles cannot hope to be effective if the bubbles themselves cannot in fact be measured with sufficient accuracy.

An alternative to pursuing an exchange-rate target that compensates for bubbles might be, perhaps, to target a short-term equilibrium that aims to offset transitory influences on the exchange rate. In other words the central bank might choose to adopt a target such as Eqn.(2.7). The advantage of such a strategy would be that the central bank could focus on aligning the real exchange rate with its permanent, supply-side determinants while compensating for transitory determinants such as nominal shocks thought to have no bearing on the real exchange rate in the long run.[19] All this seems reasonable. The problem, however, is that transitory components can, for some currencies, explain a great deal of the movement of the real exchange rate. For Germany and Japan, for instance, Clarida and Gali (1994) find that around 70% and 60%, respectively, of the variances of these countries' real exchange rates are due to transitory components. To ignore these transitory, cyclical components when they play such a big part in the determination of the real exchange rates would be to ignore the fact that the driving fundamentals contain important transitory elements. Currency-market intervention aimed at compensating for these transitory elements could, therefore, be self-defeating. The implication is that no rational, forward-looking central bank would countenance such a policy.

If neither a target based on short-run equilibrium nor one based on long-run equilibrium is a plausible proposition for central-bank intervention policy,

[16] A number of economists have argued that during the second half of the twentieth century Japan experienced a prolonged Balassa-Samuelson effect. See for instance Marston (1987) and Koedijk et al. (1998).

[17] See Bernanke and Gertler (1999) for a discussion about policy responses to asset-price bubbles, and for rational speculative bubbles see, for instance, Buiter and Pesenti (1990).

[18] For a discussion about identifying exchange-rate bubbles see, for example, Norden (1986).

[19] For a flavour of the discussion about decomposing the real exchange rate into its permanent and transitory components, see, for instance, Clarida and Gali (1994) for the Beveridge-Nelson decomposition and for structural-vector-autoregression estimates and see Clark and MacDonald (2000) for cointegration-based estimates of permanent, equilibrium exchange rates.

the only alternative is to target a medium-run equilibrium. There are a number of reasons for thinking that a medium-run target, such as Eqn.(2.8), may be preferable. First, the target allows for the cyclical adjustment of the fundamentals towards their stock-flow-consistent values. This, surely, would be a desirable feature of any exchange-rate policy—it makes sense to adopt a policy that does not conflict with any endogenous, cyclical tendency for the economy to change. Indeed, survey results compiled by Neely (2000) suggest that most monetary authorities do not intervene in order to correct long- or medium-run misalignments, but instead aim to compensate for short-run volatility, with 90% of respondents saying they intervene sometimes or always to resist short-term trends in the exchange rate.

The second reason why central banks are more likely to target a medium-run equilibrium rather than a short- or long-run equilibrium is the appropriateness of the policy horizon. A medium-run target that allows for cyclical endogenous change is less likely to conflict with other policy objectives if those other objectives are equally sensitive to the business cycle, having forward-parameter-setting horizons of more than a year (eg, inflation-targeting monetary policy) but less than, say, seven years. In short, there seems to be ample support for supposing that those central banks that do intervene in the currency markets in order to pursue a target exchange rate do so in pursuit, frequently, of a medium-run target.

What type of medium-run equilibrium

If a medium-run equilibrium is to be used in the intervention reaction function, the next question is what form of medium-run equilibrium? There are a number of possibilities. Driver and Westaway (2003) highlight two main concepts of medium-run equilibrium: the Fundamental Equilibrium Exchange Rate (or FEER) and Desired Equilibrium Exchange Rate (or DEER). Both are models of internal-external balance, with equilibrium being defined as the level of the real exchange rate that is consistent with balance both internally and externally, permitting changes over time in net foreign assets.

Unfortunately neither FEERs nor DEERs are particularly simple to calculate. Estimates require either a fully specified macromodel or a partial-equilibrium model containing a subset of the relevant equations. The partial-equilibrium approach is more commonplace and, perhaps, simpler, but even with this approach the necessary calculations are lengthy. First, net-trade and net-income relationships are specified and the current-account trend is calculated on the assumption that real exchange rates are at their actual levels (while it is assumed that output levels at home and abroad are at their trend values). Cyclical factors will, therefore, account for the difference between the actual current account and the trend current account, and the FEER is the real exchange rate that brings into balance the trend current account with a sustainable level of savings and investment in each economy (sustainable, that is, according to some normative benchmark). Calculating DEERs is not much simpler. It involves the same essential elements of estimation as for FEERs, but the real exchange rate is conditioned upon some optimal trajectory for fiscal policy.

Complexity of estimation is just one reason why, in all likelihood, neither FEERs nor DEERs will be used as targets by any central bank aiming to intervene in the currency markets in an effort to guide the current exchange rate towards a medium-run equilibrium. Calculating both FEERs and DEERs, as has been explained, requires estimating a large number of variables, something that no central bank will be keen to entertain if hoping to respond quickly, with intervention, to daily movements in exchange rates. What is more, much of the data necessary for calculating FEERs and DEERs is available only with a long lag. Data for net foreign assets, for instance, is available sometimes with a delay of six months—and often longer.

Intuitively an exchange-rate target must, for operational purposes, possess a certain number of qualities: it must be relatively simple and quick to calculate, it must reflect an accepted measure of equilibrium, and estimates of its value must be timely (that is, any intervention in the currency markets must be conditional upon the target being accurate at the time of intervention rather than accurate at some arbitrary time in the past). In terms of simplicity and timeliness, it is clear that neither FEERs nor DEERs fit the bill as functional exchange-rate targets. There are, however, alternatives.

Purchasing power parity (PPP) is one obvious alternative. As a concept, it is simple, and as a calculable measure of equilibrium it can, to an extent, be timely if interpolation methods are used to decant monthly price data into a daily format. But as has been explained earlier, PPP is a measure of long-run equilibrium. On its own, it cannot be used as a reliable gauge of medium-run equilibrium.

However, if PPP is combined with other elements that are sensitive to the effects of the business cycle, then it can, in fact, be employed as a measure of medium-run equilibrium. One such PPP-hybrid measure of the medium-run equilibrium is the capital enhanced equilibrium exchange rate, or CHEER.

Capital-enhanced equilibrium exchange rate

The CHEER approach involves combining PPP with the theory of uncovered interest rate parity (UIP) and has been championed by, among others, Juselius

(1991, 1995), MacDonald and Marsh (1997, 1999) and Juselius and MacDonald (2000b,a). The basic premise of the approach is that while PPP may explain exchange-rate movements in the long run, in the medium run an exchange rate can diverge from PPP as a result of differences in interest rates across countries. The CHEER approach does, in other words, reflect a Casselian view of PPP (see MacDonald and Marsh (1997)). But it differs in one important respect. A strict Casselian approach would mean embracing the assumption that non-zero interest differentials have no more than a temporary impact on the real exchange rate. The CHEER approach, however, makes the assumption that there is persistence in both interest-rate differentials and the real exchange rate.

Assuming that there is persistence in interest-rate differentials and in the real exchange rate implies, here, that exchange rates can diverge from their PPP values as a result of, for instance, savings imbalances that manifest themselves as large current-account gaps. The important point here is that adjustment of any current-account gap in response to relative prices will be slow if the mean-reversion process of the real exchange rate towards PPP is also sluggish. A persistent current-account gap will, of course, be financed through the capital account, and so long as the capital account continues to function in this way, interest rates will diverge from uncovered interest parity and the movement of the exchange rate in the medium run will reflect a capital-enhanced equilibrium.[20]

The CHEER approach is fairly simple to represent statistically. First, purchasing power parity is defined, according to convention, as

$$s_t = p_t - p_t^* \tag{2.10}$$

where s_t is the log of the spot exchange rate (home currency per unit of foreign currency), p_t is the log of the domestic price level and where an asterisk denotes a foreign variable. Meanwhile, uncovered interest parity (UIP) can be expressed as

$$s_{t+k}^e - s_t = (i_t - i_t^*) \tag{2.11}$$

where i_t is the yield on bonds with maturity k and e represents expectations such that $s_{t+k}^e - s_t$ denotes the expected change in the exchange rate over the next k periods. The implication of Eqn.(3.1) is that the domestic interest rate is equal to the foreign interest rate plus the expected appreciation of the exchange rate at an annualized percentage rate.

When the forecast horizon lengthens, it is reasonable to expect, as highlighted by Juselius (1995), that the formation of expectations will be influenced increasingly by deviations from PPP. In this way, if the expected exchange rate

[20]Note that this capital-enhanced approach when applied, as it often is, to bilateral exchange rates, will reflect only partial equilibrium and will not reflect equilibrium in the whole economy.

is given as

$$s^e_{t+k} = p_t - p^*_t \tag{2.12}$$

then an expression for the spot exchange rate, s_t, can be derived, containing both PPP and UIP, by inserting Eqn.(2.12) into Eqn.(3.1), such that

$$(i_t - i^*_t) = \alpha(p_t - p^*_t) - s_t \tag{2.13}$$

which, in this formulation, represents a Casselian version of exchange-rate equilibrium. While this Casselian version of exchange-rate equilibrium does represent only a partial equilibrium, there is plenty of justification for its use here, not least the fact that, as highlighted by MacDonald and Marsh (1997), for recent sample periods covering the modern era of floating exchange rates (ie, since 1973), it does not seem appropriate to impose stock-flow consistency given that capital-account effects may have arisen as a result of productivity differentials or the impact of different monetary policies across countries. It seems right, therefore, to allow this Casselian formulation to be used, in this study, to represent medium-run equilibrium for the exchange rate and it will, as such, be employed as the exchange-rate target in the central-bank reaction function.

Optimal intervention

As has been explained earlier, by minimising the central bank's loss function in Eqn.(2.1) subject to the constraint Eqn.(2.2), it is possible to recover an expression for optimal intervention

$$Int^*_t = -\frac{1}{\lambda}(s_{t-1} - s^T_t) \tag{2.14}$$

Furthermore, assuming a Casselian version of exchange-rate equilibrium, the target exchange rate can be expressed, now, as

$$s^T_t = \alpha_1 p_t + \alpha_2 p^*_t + \omega_1 i_t + \omega_2 i^*_t \tag{2.15}$$

From Eqn.(2.15) it can be seen that if the central bank attaches importance to a strict version of purchasing power parity, then α_1 will take a value close to unity, and α_2 close to negative unity (that is, the monetary authorities will place emphasis on a purchasing-power-parity target that reflects price homogeneity). Meanwhile, if interest differentials play an important role in exchange-rate targeting, then negative ω_1 will be close in value to ω_2, whatever that value might be. However, if it is unconstrained interest rates that are important (ie, if it is unconstrained interest rates that are key to exchange-rate equilibrium in the medium run), then negative ω_1 need not be close in value to ω_2. Substituting Eqn.(2.15) into Eqn.(2.14) it is possible to derive an expression for optimal intervention, Int^*, where

$$Int_t^* = -\frac{1}{\lambda}(s_{t-1}) + \frac{1}{\lambda}(\alpha_1 p_t + \alpha_2 p_t^* + \omega_1 i_t + \omega_2 i_t^*) \tag{2.16}$$

or similarly,

$$Int_t^* = \beta_0 s_{t-1} + \beta_1 p_t + \beta_2 p_t^* + \beta_3 i_t + \beta_4 i_t^* \tag{2.17}$$

where $\beta_0 = -1/\lambda$, $\beta_1 = \alpha_1/\lambda$, $\beta_2 = -\alpha_2/\lambda$, $\beta_3 = -\omega_1/\lambda$ and where $\beta_4 = \omega_2/\lambda$. We estimate β but there is no way to identify α without additional assumptions. From Eqn(2.16) it is clear that optimal, daily intervention is a positive function of the difference between the value of the exchange rate on day $t-1$ and an expression for medium-run equilibrium represented by a target exchange rate containing prices and interest rates both at home and abroad. That is, optimal intervention will increase in line with the extent of the deviation of the actual exchange rate from its target level.

2.3.3 Friction

It must be noted that the formulation for optimal intervention that has been developed so far in this study, represented by equation Eqn.(2.16), implies that intervention is a continuous process, taking place every day. This does not, of course, chime with reality. The stylised facts show that intervention takes place on fewer than one trading day in every ten. There are long periods of no intervention followed by periods when intervention is large and sustained. There is, in other words, a high degree of serial correlation between interventions over time, with intervention on any given day making it more likely that intervention will be seen to occur again the next day.

This phenomenon of dynamic correlation between interventions has been suggested, by Ito and Yabu (2004), to be due to the fact that intervention policy is complicated by the existence of policy friction costs. With friction costs, the central bank, in executing any given decision to intervene, must first convince the nation's political authorities (in Japan's case, the Ministry of Finance) that intervention is necessary, and it is this process of negotiation that is assumed to represent, in itself, an intervention cost. If, on any given day, the central bank succeeds in securing consent to proceed with intervention, then on subsequent days intervention costs will be lower, causing intervention to be more likely. The result is a high degree of dynamic correlation between interventions.

A friction model such as this, based as it is on the notion that the pursuit of an optimal intervention strategy is hampered by the presence of intervention costs, does have obvious intuitive appeal. However, it also has flaws. Most importantly, there is no hard evidence to support the notion that significant intervention costs do exist.

Another explanation for dynamic correlation between interventions is the idea, suggested by Herrera and Ozbay (2005), that it is the objective of the central bank to minimise an intertemporal loss function that is non-time separable, whereby current-period interventions depend not just on current-period explanatory variables, but also on previous values and-or future values of the explanatory variables. It could be the case, for instance, that the central bank, in endeavouring to guide the nominal exchange rate towards a target level, chooses to respond not only to current values of the exchange rate but also to past values.

This idea has been explored by Bomfim and Rudebusch (1997) in application to inflation targeting and much of the logic of the approach can be used, without much difficulty, to explain dynamic correlation in a framework of exchange-rate targeting. However, this approach still fails to account for the long periods when there is no intervention. A friction model seems to represent one of the most convenient ways to account for extended periods of no intervention and the concept put forward by Ito and Yabu (2004) of policy frictions does seem to be tractable. It is, therefore, adopted in this study.

If intervention occurs after a period of no intervention, then on subsequent days the cost of persuading the political authorities to intervene again will be lower. As such, once intervention has occurred, then the chances of intervention occurring again, a day later, are higher and so we have a mechanism to explain why intervention tends to be correlated. A cost function for friction costs F at time t, where F^+ represents the friction cost associated with purchasing yen and F^- represents the friction cost associated with selling yen, can be defined in the style of Ito and Yabu (2004) as

$$F_t = \begin{cases} F^+_{F,t} - F_{D,t}[I(Int_{t-1} > 0)] & \text{if} \quad Int_t > 0 \\ F^-_{F,t} - F_{D,t}[I(Int_{t-1} < 0)] & \text{if} \quad Int_t < 0 \end{cases} \tag{2.18}$$

where F_t is the friction cost associated with intervention at time t. In addition, $F_{F,t}$ represents first-time friction costs when intervention occurs for the first time after a period of no intervention, and $F_{D,t}$ represents the size of the discount that reduces friction costs in any given period when political consent for intervention has been secured in the previous period.

Note that $F^j_{F,t} > 0$, $F_{D,t} > 0$, $F^j_{F,t} > F_{D,t}$ and $I(.)$ is the indicator function. If $Int_{t-1} > 0$ then $I(Int_{t-1} > 0)$ takes the value of unity. Otherwise it takes the value of zero. Furthermore, $F_{D,t} > 0$ implies that intervention on day $t-1$ is associated with a reduction in the friction costs of intervention on day t. Where

$F_{F,t}$ takes a positive superscript, this implies that the friction costs incurred are associated only with yen-buying interventions. Where $F_{F,t}$ takes a negative superscript, friction costs are those associated with yen-selling interventions. Also, recall that $Int_t > 0$ signifies yen-buying intervention and $Int_t < 0$ signifies yen-selling intervention.

One notable feature of Eqn.(2.18) is that it permits a role for asymmetry. That is, Eqn.(2.18) captures the fact that the political authorities may be more averse to, say, yen-buying intervention than yen-selling intervention, causing the friction costs associated with yen buying to be larger. In Eqn.(2.18) this would be represented by $F^{+}_{F,t} > F^{-}_{F,t}$.

2.3.4 Indicator function

To avoid the problem of endogeneity (as explained in Section 2.3.1), in this study the dependent variable, intervention, is modelled not as intervention *amount*, but as an intervention *indicator function*. The intervention indicator function takes one of three values: +1, 0, or -1, representing, respectively, either the purchase of yen, no intervention, or the sale of yen.

If the central bank faces friction costs that impede its pursuit of an optimal intervention strategy, then intervention takes place only if the benefits of intervention are higher than the costs. That is, intervention will occur only when optimal intervention, Int^*_t, exceeds certain threshold values, γ_{+1} (triggering purchases of yen) and γ_{-1} (triggering sales of yen), determined by the size of the friction costs. In this way, we can define optimal intervention as an unobserved latent variable, Int^*_t, such that

$$Int^*_t = X_t\beta + \epsilon_t \tag{2.19}$$

where we assume $\epsilon_t \sim i.i.d.N(0, \sigma^2)$, and where $X_t\beta$ is defined as

$$X_t\beta = \beta_0 s_{t-1} + \beta_1 p_t + \beta_2 p^*_t + \beta_3 i_t + \beta_4 i^*_t + \beta_5 Int_{t-1} \tag{2.20}$$

While Int^*_t remains unobservable, we can define its observable corollary Int_t, where

$$Int_t = \begin{cases} -1 & \text{if} \quad Int^*_t < \gamma_{-1} + \beta_5 I(Int_{t-1} < 0) \\ 0 & \text{if} \quad \gamma_{-1} + \beta_5 I(Int_{t-1} < 0) < Int^*_t < \gamma_{+1} - \beta_5 I(Int_{t-1} > 0) \\ +1 & \text{if} \quad \gamma_{+1} - \beta_5 I(Int_{t-1} > 0) < Int^*_t \end{cases} \tag{2.21}$$

where $\gamma_{-1} < 0$, $\gamma_{+1} > 0$ and $\beta_5 > 0$.[21] This expression can, however, be

[21]Recall that if $Int_{t-1} > 0$, then $I(Int_{t-1} > 0)$ takes the value of one. Otherwise it takes the value of zero. Similarly, if $Int_{t-1} < 0$, then $I(Int_{t-1} < 0)$ takes the value of one.

simplified further given that, in practice, the direction of intervention on date t is never different from the direction of intervention on date $t-1$. Given this, Eqn.(2.21) can be reformulated as

$$Int_t = \begin{cases} -1 & \text{if} \quad Int_t^* < \gamma_{-1} \\ 0 & \text{if} \quad \gamma_{-1} < Int_t^* < \gamma_{+1} \\ +1 & \text{if} \quad \gamma_{+1} < Int_t^* \end{cases} \tag{2.22}$$

In this form, the model is, empirically, a multiple-response model, and more specifically an ordered-response model, in which the underlying latent variable is optimal intervention, Int_t^*, and the observed variable is the intervention indicator function, Int_t. The form of this ordered-response model is determined by the distribution assumed for ϵ_i.

2.4 Estimation methodology

This section explains the empirical approach used to estimate Japan's intervention reaction function and, in particular, offers an outline of the ordered logit model, generalised ordered logit model and the partial proportional odds model.

2.4.1 Standard ordered logit model

Whenever the dependent variable in a regression is qualitative and can take more than two possible values, with these values having a natural order or rank, the standard approach is to employ either an ordered probit model or, as will be discussed here, an ordered logit model, as first proposed by McCullagh (1980). The ordered logit model allows for a qualitative dependent variable for which the categories have a natural order that reflects the magnitude of some continuous underlying variable. Here, the underlying variable is optimal intervention, Int^*. The qualitative dependent variable is the intervention indicator function, Int.

If the inherent ordering of the intervention indicator function were to be ignored, and if, instead of an ordered response model, a multinomial logit model were to be employed, the result would be mis-specification of the data-generating process. Inferences about the response variable would be erroneous. Ordinary least-squares estimation would be similarly inappropriate. Ordinary least squares assumes that differences between categories of the dependent variable are equal: the difference between a 1 and a 2 is the same as the difference between a 2 and a 3. But here, differences may not be equal: the intervention indicator function reflects only an ordinal ranking, not a cardinal ranking.

Otherwise it equals zero.

The dependent variable, Int^*, representing optimal intervention, is an unobserved latent variable and is defined according to

$$Int_t^* = X_t\beta + \epsilon_t \tag{2.23}$$

where ϵ_t is a disturbance term and X_t is a vector of explanatory variables such that

$$X_t\beta = \beta_0 s_{t-1} + \beta_1 p_t + \beta_2 p_t^* + \beta_3 i_t + \beta_4 i_t^* + \beta_5 Int_{t-1} \tag{2.24}$$

with s, p, p^*, i and i^* defined according the descriptions offered in Section 2.3.

An assumption is made that the dependent variable has three categories so that instead of observing Int^* we observe Int, where, as discussed previously,

$$Int_t = \begin{cases} -1 & \text{if} \quad Int_t^* < \gamma_{-1} \\ 0 & \text{if} \quad \gamma_{-1} < Int_t^* < \gamma_{+1} \\ +1 & \text{if} \quad \gamma_{+1} < Int_t^* \end{cases} \tag{2.25}$$

Here, the γs are unknown threshold parameters that must be estimated along with β in Eqn.(2.23). Estimation is undertaken by maximum likelihood, which in the case of the ordered logit model requires that the cumulative density function of ϵ is the logistic function. For more details see McCullagh (1980).

It is important to note that the estimation method is robust only if the disturbance term in Eqn.(2.23) satisfies certain regularity conditions that are consistent with asymptotic normality. For instance, one such regularity condition is that of increasing information: the amount of information in the data must increase indefinitely as the sample size increases. If, however, there is too much dependence in the data, then this condition will not hold. One potential source of dependence in the data is *day-of-the-week effects*. Day-of-the-week describe the tendency of daily asset returns to show repeated patterns from week to week. For instance, Damodaran (1989) finds that bad news tends to be released on Fridays and, due to the delayed release of information, Mondays tend to be associated with lower returns. If day-of-the-week effects are present in the daily yen-dollar exchange rate then dependence will be present and the maximum likelihood estimator may lack asymptotic normality. Researchers have attempted to measure the extent to which day-of-the-week effects are present in the yen-dollar exchange rate. Yamori and Mourdoukoutas (2003) and Yamori and Kurihara (2004) find that day-of-the-week effects for the yen-dollar exchange rate disappeared in 1990s. This study accepts this finding and continues on the assumption that day-of-the-week effects do not contaminate the maximum likelihood estimator.

The probability of observing $Int_t = -1$ is equal to

$$\begin{aligned} P\{Int_t^* = X_t\beta + \epsilon_t < \gamma_{-1}\} &= P\{\epsilon_t < \gamma_{-1} - X_t\beta\} \\ &= \int_{-\infty}^{\gamma_{-1}-X_t\beta} f(\epsilon)d\epsilon \end{aligned} \tag{2.26}$$

where $f(\epsilon)$ is the logistic function. Similarly, the probability of obtaining an observation with $Int_t = 0$ is equal to

$$\begin{aligned} P\{\gamma_{-1} < Int_t^* = X_t\beta + \epsilon_t < \gamma_{+1}\} &= P\{\gamma_{-1} - X_t\beta < \epsilon_t < \gamma_{+1} - X_t\beta\} \\ &= \int_{\gamma_{-1}-X_t\beta}^{\gamma_{+1}-X_t\beta} f(\epsilon)d\epsilon \end{aligned} \tag{2.27}$$

while the probability of obtaining an observation with $Int_t = 1$ is equal to

$$\begin{aligned} P\{\gamma_{+1} < Int_t^* = X_t\beta + \epsilon_t\} &= P\{\gamma_{+1} - X_t\beta < \epsilon_t\} \\ &= \int_{\gamma_{+1}-X_t\beta}^{\infty} f(\epsilon)d\epsilon \end{aligned} \tag{2.28}$$

The likelihood function is the product of all of these expressions. Maximising the likelihood function with respect to the γs and β gives the maximum likelihood estimates. Greene (2008) offers further discussion.

Key to the standard ordered logit model is the assumption that the slope coefficients are equal across all of the outcome equations, represented here by Eqn.(2.26), Eqn.(2.27) and Eqn.(2.28). This assumption is known as the *proportional odds assumption*. It can be illustrated more simply as follows.

First, state the current problem in more general terms, estimating the probability P that the event indicator Y takes the value $m = 1, \ldots, J$. In this manner Greene (2008) states the standard ordered logit model as,

$$P(Y > m) = \frac{exp(\tau_m + x\beta)}{1 + exp(\tau_m + x\beta)} \qquad \text{for} \quad m = 1, 2, \ldots, J-1 \tag{2.29}$$

where x is a vector of variables accounting for deviations from a currency target as defined in Section 2.3, and where τ and β are parameters to be estimated. Note that in Eqn.(2.29) the only component that differs across event outcomes is the cut-off parameter, τ. The link function's slope parameters, the βs, are assumed to be equal. That is, Eqn.(2.29) shows that the standard ordered logit model is equivalent to $J-1$ binary regressions with the critical assumption that the slope coefficients are identical across each regression. Here, for instance, where the event outcomes are intervention outcomes, identical slope coefficients imply that the effect of a change in x, such as a change in Japanese interest rates, will have the same effect on the probability of yen-selling intervention as it will on yen-buying intervention or no intervention. This may not, of course, hold true. Probabilities may differ. The intervention objective function may not be symmetric.

One way to allow for varying slope coefficients is to adopt the approach of Williams (2006) and specify a generalised ordered logit model.

2.4.2 Generalised ordered logit model

This section presents a brief overview of the *generalised ordered logit model*, which has been discussed elsewhere by Williams (2006), Clogg and Shihadeh (1994) and Fahrmeir and Tutz (1994).

The generalised ordered logit model allows not only for different intercepts for each event outcome. It allows also for different slope parameters. That is, the generalised ordered logit model can be specified as,

$$P(Y > m) = \frac{exp(\tau_m + x\beta_m)}{1 + exp(\tau_m + x\beta_m)} \qquad \text{for} \quad m = 1, 2, \ldots, J-1 \tag{2.30}$$

where β is allowed to differ for each of the event outcomes, $m = 1, \ldots, J-1$. The generalised ordered logit model nests the standard ordered logit model under the restriction $\beta_2 = \ldots = \beta_J$.

The probabilities that Y will take on values $m = 1, \ldots, J-1$ are given as,

$$\begin{aligned} P(Y = 1) &= 1 - \frac{exp(\tau_1 + x\beta_1)}{1 + exp(\tau_1 + x\beta_1)} \\ P(Y = m) &= \frac{exp(\tau_{m-1} + x\beta_{m-1})}{1 + exp(\tau_{m-1} + x\beta_{m-1})} - \frac{exp(\tau_m + x\beta_m)}{1 + exp(\tau_m + x\beta_m)} \quad m = 2, \ldots, J-1 \\ P(Y = J) &= \frac{exp(\tau_{J-1} + x\beta_{J-1})}{1 + exp(\tau_{J-1} + x\beta_{J-1})} \end{aligned} \tag{2.31}$$

As such, the generalised ordered logit model can be thought of as a series of simple, two-outcome logistic regressions where all multiple outcome categories $1, \ldots, J$ are allocated to one of two categorical states. For instance, if $J = 5$, then the first outcome category is compared with a grouped combination of categories $2, 3, 4$ and 5. For the second outcome category, both $J = 1$ and $J = 2$ are taken together and compared with a grouped combination of categories $3, 4$ and 5. For the third outcome category, categories $J = 1$, $J = 2$ and $J = 3$ are taken together and compared with a grouped combination of categories 4 and 5. And so on.

By allowing the βs to vary across values of m, the generalised ordered logit model allows for ordered outcome systems that violate the proportional odds assumption. In other words, the generalised ordered logit model is more flexible than the standard ordered logit model. However, the generalised ordered logit model does have a drawback. Because of its permissive structure it can lead to the estimation of more parameters than is necessary. That is, not all βs are necessarily different. Some may be equal. A special case of the generalised ordered logit model that can overcome these limitations is the partial proportional odds model.

2.4.3 Partial proportional odds model

The *partial proportional odds model* permits a relaxation of the proportional odds constraint for those variables that violate it. The partial proportional odds model is discussed further by Peterson and Harrell (1990).

In the partial proportional odds model, some of the slope coefficients, the βs, can be the same for all intervention outcomes, m, while others can vary. The model is less restrictive than the standard ordered logit model but offers more parsimony than the generalised ordered logit model. For instance, in the following model the βs for x_1 and x_2 are the same for all values of m but the βs for x_3 are free to vary.

$$P(Y > m) = \frac{exp(\tau_m + x_1\beta_1 + x_2\beta_2 + x_3\beta_{3m})}{1 + exp(\tau_m + x_1\beta_1 + x_2\beta_2 + x_3\beta_{3m})} \qquad \text{for} \quad m = 1, \ldots, J-1 \tag{2.32}$$

Estimation of the partial proportional odds model is discussed further in Section 2.6. Here it is sufficient to say that the partial proportional odds model offers an attractive and tractable method of estimating the model of intervention behaviour and is employed in this study in order to test the realised condition of symmetry in the intervention objective function.

2.5 Data

This section describes the data and offers some cursory data analysis.

The frequency of the data is daily. The sample period extends from 01 April 1991 to 31 March 2006, comprising 3,915 data points. The dependent variable, the intervention indicator function Int, is a constructed variable, and on any given day takes a value of either +1, representing Bank of Japan intervention in the currency markets to purchase yen (in exchange for US dollars), −1, representing intervention to sell yen, or 0, indicating no intervention. Table 2.1 shows the frequency of each type of intervention. Construction of the dependent variable is based on information provided by the Japanese Ministry of Finance.

Figure 2.3: Japanese intervention in the currency markets

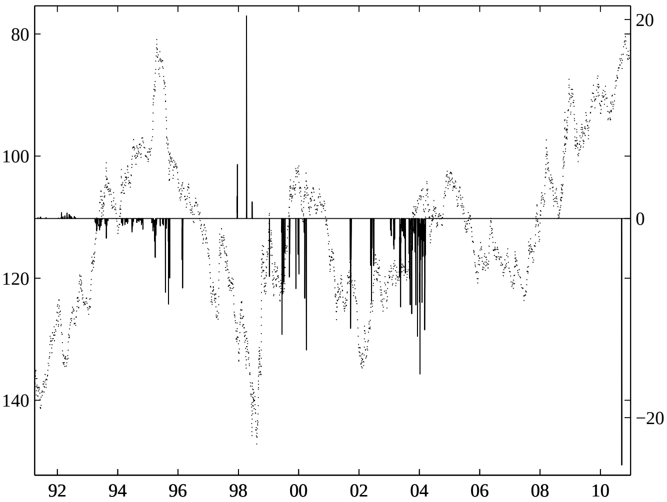

Notes: Dotted line shows Japanese yen in terms of yen per US dollar, reverse scale, measured on the left-hand axis. Solid line shows the amount, in US dollar billions, of Japanese intervention in the currency markets, measured on the right-hand axis. Frequency is daily. Positive amounts of intervention indicate purchases of yen. Negative intervention indicates sales of yen. Source: Japanese Ministry of Finance and Bloomberg.

Table 2.1: Intervention indicator function

Int	**Freq.**	**Per Cent**	**Cum.**
-1	311	7.9	7.9
0	3,571	91.2	99.2
+1	33	0.8	100.0
Total	3,915	100.0	

Notes: Intervention indicator function Int, where $Int = -1$ represents yen-selling intervention, $Int = 0$, no intervention and $Int = +1$, yen-buying intervention.

Table 2.2: Summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Int	3915	-0.07	0.29	-1.00	1.00
sl	3915	4.74	0.10	4.39	4.99
pus	3915	4.56	0.11	4.36	4.76
pjp	3915	4.59	0.01	4.55	4.62
ius	3915	0.06	0.01	0.03	0.08
ijp	3915	0.03	0.01	0.00	0.07
Intl	3915	-0.07	0.29	-1.00	1.00

Abbreviations: *Int*, intervention indicator function, *sl*, exchange rate, lagged one period (yen per US dollar), *pus*, US consumer prices, *pjp*, Japan consumer prices, *ius*, US interest rates, compounded continuously, *ijp*, Japan interest rates, compounded continuously, *Intl*, intervention indicator function, lagged one period.

The spot exchange rate, s, is defined in terms of Japanese yen per US dollar measured at the start of the New York trading day and transformed using the natural logarithm operator. The source for the daily exchange-rate data is Bloomberg.

Japanese prices p, and US prices p^*, are based on consumer price indices drawn from the International Monetary Fund's database of International Financial Statistics. The consumer price indices are recorded, at source, on a monthly basis. Cubic spline interpolation is used to transform the monthly series into daily series. Daily prices are in natural logarithms.

Japanese interest rates i, and US interest rates i^*, are daily one-month treasury bill rates, compounded continuously.[22] Augmented Dickey-Fuller statistics (not reported) suggest that each series is nonstationary in levels but stationary after first differencing. Table 2.2 offers a summary of the dataset.

2.6 Results

This section presents results from three empirical models of intervention in the currency markets: an ordered logit model, a generalised ordered logit model and a model that assumes partial proportional odds.

Estimation results are presented in Table 2.3. In the standard ordered logit model the parameter vector β is estimated together with two thresholds (cut points) separating the three categorical outcomes. Most parameters have the expected sign. For example, the association between intervention and the lagged value of the exchange rate (sl) is positive. That is, a weakening yen is associated with yen-buying intervention ($Int = +1$) and a strengthening yen with

[22] That is, where r is the interest rate in decimal form, the continuously compounded interest rate, i, is calculated as $i = \ln(r + 1)$.

Table 2.3: Ordinal outcome models of intervention

Variable	Ordered logit	Generalised ordered logit		Partial proportional odds	
Dep. var: *Int*	β	β_{-1}	β_0	β_{-1}	β_0
sl	8.15***	6.75***	16.95***	6.73***	16.08***
	(0.811)	(0.861)	(4.245)	(0.861)	(3.806)
pus	4.202**	4.02*	-54.73***	3.75*	-46.24***
	(1.964)	(2.061)	(19.823)	(2.008)	(12.724)
pjp	-3.51	5.20	128.64***	4.09	133.42***
	(10.146)	(11.741)	(41.099)	(11.564)	(39.135)
ius	109.663***	112.81***	128.05	114.04***	114.04***
	(19.571)	(20.919)	(96.924)	(20.435)	(20.435)
ijp	-35.24	-52.21**	-103.73	-55.63**	-55.63**
	(22.472)	(24.244)	(93.678)	(23.475)	(23.475)
Intl	3.535***	3.57***	2.08***	3.57***	2.06***
	(0.148)	(0.159)	(0.478)	(0.159)	(0.478)
const	-	-75.48	-437.05***	-69.05	-493.49***
	-	(58.825)	(165.831)	(57.791)	(137.794)
Observations	3,915	3,915		3,915	
Wald chi-square	986.6	1,031.6		1,031.3	
Log likelihood	-780.4	-757.9		-758.1	
Pseudo R-square	0.3873	0.4050		0.4048	

Notes: Table shows estimation results for three ordinal outcome models of intervention: ordered logit, generalised ordered logit and partial proportional odds model. Robust standard errors are in parentheses. (***) denotes significance at the 1% level, (**) at the 5% level, and (*) at the 10% level. Abbreviations: *Int*, intervention indicator function, *sl*, exchange rate, lagged one period (yen per US dollar), *pus*, US consumer prices, *pjp*, Japan consumer prices, *ius*, US interest rates, compounded continuously, *ijp*, Japan interest rates, compounded continuously, *Intl*, intervention indicator function, lagged one period. $Int = +1$ denotes yen-purchasing intervention, $Int = -1$ denotes yen-selling intervention and $Int = 0$ denotes a day on which no intervention occurred. Exchange rates and prices are in natural logarithms.

yen-selling intervention ($Int = -1$). The association between intervention and US interest rates (ius) is also positive. Higher US interest rates are associated with Japanese intervention to support the value of the yen. Meanwhile, the association between intervention and Japanese interest rates (ijp) is negative—higher Japanese interest rates are more likely to be associated with yen-selling intervention (aimed, perhaps, at stemming the strength of the yen). The statistical association between intervention today and intervention on day t and day $t-1$ is positive, in agreement with the stylised facts.

As appealing as these results from the ordered logit model may appear, there are two reasons to doubt them. First, it is misleading to rely on coefficient estimates alone in any evaluation of the economic significance of the covariates in an ordered logit model. Hosmer and Lemshow (2000) show that it is wiser to consult marginal effects. Second, as discussed earlier in Section 2.4, the ordered logit model may be too restrictive in the sense that it assumes equal parameter vectors in each part of the outcome distribution: that is, contrary to the assumptions of the ordered logit model, different outcome categories may carry different slope coefficients. For instance, it is possible that changes in interest rates may affect the respective likelihoods of intervention (yen-selling

Table 2.4: Brant test of proportional odds assumption

Variable	**all**	**sl**	**pus**	**pjp**	**ius**	**ijp**	**Intl**
	25.76***	5.68**	8.76***	8.35***	0.02	0.30	8.88**
	(0.000)	(0.017)	(0.003)	(0.004)	(0.879)	(0.587)	(0.003)

P-values in parenthases (p > chi-square). (***) denotes significance at the 1% level, (**) at the 5% level, and (*) at the 10% level. A significant test statistic provides evidence that the proportional odds assumption has been violated. Abbreviations: *all*, all covariates, *sl*, exchange rate, lagged one period (yen per US dollar), *pus*, US consumer prices, *pjp*, Japan consumer prices, *ius*, US interest rates, compounded continuously, *ijp*, Japan interest rates, compounded continuously, *Intl*, intervention indicator function, lagged one period.

and yen-buying) to different extents. The relationship may, in other words, be asymmetric, and as such, it may violate the *proportional odds* assumption.

Table 2.4 presents the results of a Wald test of the proportional odds assumption as first proposed by Brant (1990). Testing, first, the hypothesis that the slope coefficients for all covariates are simultaneously equal, a chi-square statistic of 25.76 suggests this hypothesis can be rejected at the 1% level. Further testing whether the proportional odds assumption holds for some covariates but not others, Table 2.4 shows that the largest violations are for prices and for the lagged values of the exchange rate and intervention. The implication is, then, that the ordered logit model is not appropriate for modelling the intervention behaviour of the Japanese monetary authorities. The proportional odds assumption does not hold. An appropriate model should relax the proportional odds assumption and allow all slope coefficients, β_j, to vary across intervention categories $j = -1, 0, +1$. In this respect, a more appropriate model is the the generalised ordered logit model.

In the generalised ordered logit model the estimated parameters are different for each outcome category: yen-buying intervention, yen-selling intervention and no intervention. Table 2.3 presents estimation results for the generalised ordered logit model.

Results for the generalised ordered logit model can be interpreted in a manner similar to results from a series of binary logistic regressions. That is, estimation results for β_{-1} compare the outcome category $Int = -1$ with, jointly, categories $Int = 0$ and $Int = +1$. Similarly, estimation results for β_0 compare outcome categories $Int = -1$ and $Int = 0$ with category $Int = +1$. Positive coefficients indicate that higher values for an explanatory variable make it more likely that intervention will match a higher outcome category (with the highest category being $Int = +1$, meaning yen-buying intervention, and the lowest category $Int = -1$, yen-selling intervention).

Returning to Table 2.3, the signs taken by each estimated coefficient for the generalised ordered logit model of intervention are the same as those for the ordered logit model. Exceptions are the coefficients for prices. The price coefficients for the generalised ordered logit model seem more intuitively plausible: a positive β_0 coefficient of 128.64 for Japanese prices implies that rising Japanese prices are more likely to be associated with yen-buying intervention than with no intervention or yen-selling intervention. Similarly, a negative β_0 coefficient of -54.73 for US prices suggests that rising US prices are more likely to be associated with yen-selling intervention or no intervention than with yen-buying intervention.

As discussed earlier in Section 2.4, the generalised ordered logit model has the advantage of freeing the ordinal system from the assumption of proportional odds. Its disadvantage is that it includes more parameters than are perhaps necessary. This is because it frees all variables from the proportional-odds constraint despite the fact that the constraint may not need to be relaxed for all variables. Some variables may conform to the constraint.

The *partial proportional odds model* overcomes these shortcomings by relaxing the proportional-odds constraint for those variables that violate it—and imposing the constraint for those that conform to it. One preliminary indication that the partial proportional odds model may be appropriate here is given by the fact that, for the two outcome categories β_{-1} and β_0 in the generalised ordered logit model, some parameters exhibit fairly small differences in estimated coefficients: for instance, interest rates.

Estimation of the partial proportional odds model is undertaken using a backwards, stepwise iterative procedure (Williams, 2006). First an unconstrained model is estimated. Then Wald tests are executed for each variable, testing whether coefficients differ across equations—that is, each Wald test is a test of the proportional odds assumption. If the test value is statistically insignificant for one or more variables, the variable with the least significant test statistic is constrained to have equal effects across equations. The model is then re-estimated with the constraints imposed, and the procedure is repeated until there are no more variables that conform to the proportional odds assumption. The final model with constraints is then subjected to a global Wald test, benchmarked against the unconstrained model. If the test statistic is insignificant, the final model, with constraints, is accepted as not in violation of the proportional odds assumption.

Table 2.3 displays the results for the partial proportional odds model. A global Wald test allows for two variables to be constrained: Japanese interest rates (ijp) and US interest rates (ius). Note that the parameter estimates for ijp and ius are identical for the equations in β_{-1} and β_0. In line with *a priori* assumptions, the parameter estimate for ijp is negative and for ius is positive. A negative coefficient for ijp suggests that higher Japanese interest rates tend to be more associated with yen-selling intervention than with no intervention or yen-buying intervention. A positive coefficient for *ius* suggests that higher US interest rates tend to be more associated with yen-buying intervention than with no intervention or yen-selling intervention.

For the variables that do not conform to the proportional odds assumption, an examination of the pattern of coefficients reveals information that would be hidden if a simple, ordered logit model were estimated in which all variables are forced to comply with the proportional odds constraint. For instance, a negative β_0 coefficient of -46.24 for US prices suggests that higher US prices—which could, conceivably, coincide with a strengthening yen—are more likely to be associated with yen-selling intervention or no intervention than with yen-buying intervention. This seems plausible if the Japanese monetary authorities aim to stem excessive yen strength. Similarly plausible is a positive β_0 coefficient of 16.08 for the lagged value of the exchange rate. It suggests a weakening yen is more likely to be associated with yen-buying intervention than with no intervention or yen-selling intervention. Finally, a positive β_0 coefficient for $Intl$ suggests that previous yen-buying intervention on day $t-1$ is more likely to be associated on day t with further yen-buying intervention than with no intervention or with yen-selling intervention. This chimes with *a priori* expectations. Political costs are lower when intervention is repeated.

Clearly the partial proportional odds model offers a valuable contribution to the estimation of intervention probabilities. But to measure the economic significance of the relationships under analysis, the estimated coefficients presented in Table 2.3 are not sufficient. It is necessary to consider marginal effects.

2.6.1 Marginal probability effects

To best interpret a probability model such as the partial proportional odds model for the estimation of ordered outcomes it is useful to appeal to marginal effects. Here, interpreting parameters in terms of marginal effects helps to answer the question, How does the probability of observing either yen-buying intervention, no intervention or yen-selling intervention change if one of the explanatory variables changes? In Table 2.5 marginal effects for each intervention outcome are evaluated at each variable's mean ($\bar{x}$). Effects for all variables except interest rates and the lagged value of intervention are reported as semi-elasticities.

Interpretation of Table 2.5 is straightforward. Consider, for instance, the marginal effect of a small, *ceteris paribus* change in Japanese interest rates, ijp, on the probability of observing yen-selling intervention, $Int = -1$. Table 2.5 reports a value of 1.57. The implication is that a one percentage point increase in

Table 2.5: Marginal effects for the partial proportional odds model

Variable	**Int = -1**	**Int = 0**	**Int = +1**	$\bar{x}$
sl	-0.19***	0.19***	0.00	4.741
	(0.023)	(0.023)	(0.001)	
pus	-0.11**	0.10*	-0.00	4.559
	(0.057)	(0.058)	(0.001)	
pjp	-0.16	0.11	0.00	4.592
	(0.328)	(0.327)	(0.004)	
ius	-3.23***	3.22***	0.00	0.056
	(0.560)	(0.560)	(0.004)	
ijp	1.57**	-1.57**	0.00	0.026
	(0.655)	(0.655)	(0.002)	
Intl	-0.10***	0.10***	0.00	-0.071
	(0.010)	(0.010)	(0.000)	

Notes: Marginal effects evaluated at the means as semi-elasticities for all variables except interest rates (ijp and ius) and lagged intervention ($Intl$). Standard errors in parenthases. (***) denotes significance at the 1% level, (**) at the 5% level, and (*) at the 10% level. Abbreviations: *sl*, exchange rate, lagged one period (yen per US dollar), *pus*, US consumer prices, *pjp*, Japan consumer prices, *ius*, US interest rates, compounded continuously, *ijp*, Japan interest rates, compounded continuously, *Intl*, intervention indicator function, lagged one period.

Japanese interest rates raises the probability of yen-selling intervention by 1.57 percentage points. Meanwhile, the estimated marginal probability effect of a change in American interest rates on the probability of yen-selling intervention suggests that intervention is more sensitive to US monetary conditions than Japanese conditions. The marginal effect is estimated as -3.23. By implication, a one percentage point increase in US interest rates reduces the probability of yen-selling intervention by more than three percentage points.

The most important inferences to be drawn from Table 2.5 come from comparing marginal effects across intervention outcomes. Comparison highlights a number of key features. First and foremost, the economic importance of the exchange-rate target, and its component parts, varies according to the type of intervention under consideration.

Take, for instance, US consumer prices, which form a key component of the exchange-rate target. Table 2.5 shows that US prices play an important role in the decision to intervene in the currency markets to sell yen (significant at the 5% level), are important in the decision to abstain from intervention (significant at the 10% level), but are of little or no importance, statistically, in the decision to intervene to support the value of the yen (statistically insignificant). Similarly, Japanese interest rates, US interest rates and the previous day's intervention and currency value are significant in the determination of yen-selling intervention and no intervention, but are not significant factors in shaping the decision to intervene to purchase yen.

In sum, the estimated model of intervention provides much help in explaining Japan's daily decision to either refrain from intervention in the currency markets or to intervene to sell yen: estimated marginal effects show that deviations from a capital-enhanced version of purchasing power parity play a significant role in the intervention decision. But the estimated model does little to explain the decision to intervene to buy yen. This may be due to a number of factors. For instance, yen-buying interventions account for less than 1% of all observed interventions during the sample period (see Table 2.1). The chance, therefore, of uncovering economic significance between the regressors and the regressand, is small. The problem is one of power.[23] The model may in fact offer an adequate description of yen-buying intervention but the number of yen-buying interventions in the sample may be too small to support adequate estimation.

Another potential explanation for the model's lack of success in explaining instances of yen-buying intervention is that the Japanese monetary authorities are happy to accept a depreciated currency and, in terms of intervention, happy to operate a policy of benign neglect. Such an explanation would fit with anecdotal reports that Japan's Ministry of Finance is more tolerant towards a weak yen than a strong yen (Ito and Yabu, 2007).

A look at the sign and magnitude of the marginal effects gives a more detailed picture. Consider, for instance, the marginal effects of a change in the previous day's value of the exchange rate (here, yen per US dollar in natural logarithms). A 1% drop in the value of the yen reduces the likelihood of *yen-selling intervention* by 0.2 percentage points. Further, a 1% drop in the value of the yen increases the likelihood of *no intervention* by a similar amount: 0.2 percentage points. The marginal effect on *yen-buying intervention* is negligible. There is no identifiable, statistical effect.

Table 2.5 shows that the likelihood of Japan intervening in the currency markets to sell yen is influenced more by US prices than Japanese prices. A 1% increase in US consumer prices, *pus*, is associated with a drop in the probability of yen-selling intervention of 0.11 percentage points. The marginal effect on yen-selling intervention of a change in Japanese prices is insignificant. Similarly, the marginal effect of a change in Japanese prices on the decision to refrain from intervention is insignificant. The marginal effect of a change in US prices on the decision to refrain from intervention is positive and significant, with a 1% increase in US consumer prices being associated with a rise of a tenth of a percentage point in the probability of no intervention.

Clearly, foreign and domestic prices play very different roles in the inter-

[23] As power increases, the chances of making a Type II error (failing to reject the null hypothesis when the null is in fact false) decrease.

vention decision. Foreign prices, namely US prices, play an important role, perhaps due to the important consequences for global inflation of price changes in the US. Meanwhile, estimated marginal effects suggest domestic prices play a limited role in Japan's decision to intervene.

Previous intervention has the expected effect. Previous yen-buying intervention reduces the likelihood of subsequent yen-selling intervention. The magnitude of the reduction is 0.1 percentage points. Previous yen-buying intervention increases the likelihood of no intervention by an equal magnitude. Meanwhile, the marginal effect of previous yen-buying intervention on subsequent yen-buying intervention is negligible.

Summarising briefly, the partial proportional odds model yields important information on the intervention behaviour of the Japanese monetary authorities. Target variables do exist and, furthermore, seem to reflect a form of capital-enhanced purchasing power parity. Within this targeting process, the Japanese monetary authorities respond to interest rates, prices and the recent history of intervention and yen strength. All of these target variables help to explain yen-selling intervention and periods of no intervention. But their explanatory power is less strong in terms of yen-buying intervention: episodes of yen-buying intervention are rare and statistical power is low. Even so, the estimates reported in Table 2.5 tell us that deviations from the exchange-rate target are associated with clear, economically significant marginal effects on the decision to intervene.

2.6.2 Prediction

This section presents a brief analysis of in-sample predictions generated by the partial proportional odds model of intervention. Predicted probabilities are estimated as

$$\hat{P}(y = m|x) = F(\hat{\tau}_m - x\hat{\beta}) - F(\hat{\tau}_{m-1} - x\hat{\beta}) \tag{2.33}$$

with cumulative probabilities calculated as

$$\hat{P}(y \leq m|x) = F(\hat{\tau}_m - x\hat{\beta}) \tag{2.34}$$

where F is the cumulative density function for ε (F is logistic with $Var(\varepsilon) = \pi^2/3$) and where the τs are the cutpoints between the J ordinal categories, $m = 1, \ldots, J$.

Figure 2.4 summarises the in-sample predictions of intervention generated by the partial proportional odds model. There are three points to note. First, the model is good at predicting days of no intervention: days of no intervention are predicted mostly with probability greater than 0.75. Second, the model is bad at predicting days of yen-buying intervention. Most predictions carry a probability of 0.25 or smaller. Third, the model is better at predicting yen-selling intervention. Many instances of yen-selling intervention are predicted with a probability of 0.75 or greater.

All of this confirms the findings presented in Section 2.6.1: the partial proportional odds model yields important, economically significant information on the intervention behaviour of the Japanese monetary authorities. Periods of no intervention are predicted well. As are periods of yen-selling intervention. But interventions to purchase yen, where the monetary authorities step in to support of the value of the Japanese currency against the US dollar, are not predicted with great accurately.

2.7 Conclusions

This study proposes a partial proportional odds model of Japanese intervention in the currency markets between 1991 and 2006. The sample data is daily with intervention data provided by the Japanese Ministry of Finance. Intervention is categorised each day into one of three outcome categories: yen-selling intervention, yen-buying intervention or no intervention.

Estimation results show that Japanese intervention during the sample period conforms to a model in which intervention is undertaken in order to prevent the yen from straying too far from an equilibrium value defined by a capital-enhanced version of purchasing power parity. The proposed model is poor at predicting yen-buying interventions—yen-buying interventions are rare—but good at predicting yen-selling interventions and periods of no intervention.

Marginal effects show that both interest rates and prices are significant economically in determining the yen's equilibrium value and, as such, significant in shaping the decision to intervene. A one percentage point increase in Japanese short-term interest rates raises the probability of yen-selling intervention by two percentage points. Meanwhile, a one percentage point increase in US interest rates reduces the probability of yen-selling intervention by more than three percentage points. Compared with the influence exerted by Japanese monetary conditions, US monetary conditions play a bigger role in influencing Japan's decision to intervene in the currency markets.

In line with findings elsewhere in the literature, results here show that there exists significant inertia in the intervention decision: intervention on day t is more likely to take place if intervention occurred previously on day $t-1$. Conceptually this phenomenon is captured by the idea that the intervention objective function is conditioned by friction costs. Political frictions impede the intervention decision.

Finally, results here show that a standard ordered logit model of intervention, as has been used elsewhere in the literature to characterise empirically Japan's interventions in the currency markets, yields misspecified results. A more flexible approach, in the form of a partial proportional odds model, offers a better characterisation of intervention behaviour, highlighting the fact that individual independent variables influence different intervention outcomes to different extents. This finding may be useful to those policymakers and market participants interested in anticipating future interventions.

In sum, evidence here suggests that Japan tends to intervene in the currency markets in a manner that is partially predictable and, as such, the model of intervention introduced in this study holds useful informational content for anticipating future interventions. The approach offered in this study allows for a more flexible assessment of intervention than has been practised elsewhere in the relevant literature. Future studies of intervention ought to, similarly, account for the differentiated response of intervention outcomes to variations in key covariates.

Figure 2.4: Predicted probabilities of intervention

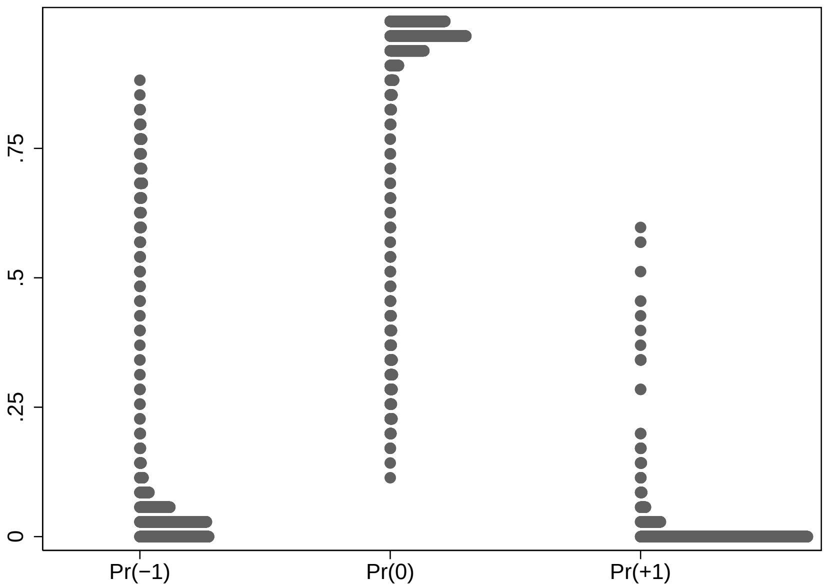

Notes: Figure shows predicted probabilities of intervention within the sample period for the partial proportional odds model. Probability is measured on the left-hand axis. Abbreviations: Probability of yen-buying intervention, Pr(+1), probability of no intervention, Pr(0), probability of yen-selling intervention, Pr(-1).

Chapter 3

Misalignment and market response

Using count-data techniques, this chapter studies the determinants of currency choice in the issuance of international bonds. In particular, this study investigates whether bond issuers choose their issuance currency in order to exploit the borrowing-cost savings associated with deviations from uncovered and covered interest-rate parity. Estimation results show that the choice of issuance currency is sensitive to deviations from uncovered interest-rate parity but insensitive, in general, to deviations from covered interest-rate parity. Furthermore, the influence of deviations from uncovered interest-rate parity is stronger for financial issuers than for nonfinancial issuers. In as much as the issuance of international bonds affects the relative international standing of world currencies, one implication of these findings is that monetary policy, through its influence on nominal interest rates, has a greater impact on the internationalisation of currencies than has been previously accounted for.

3.1 Introduction

This study investigates the aggregate behaviour of issuers of foreign-currency-denominated bonds—that is, bonds issued in a currency other than the currency of the country in which the borrower resides—and addresses the question of why issuers choose to issue bonds denominated in certain currencies and not others. Evidence is offered showing that issuers of foreign-currency-denominated bonds are sensitive to international differences in nominal interest rates and choose their currency of issuance at least partly in response to these differences. Put simply, macroeconomic factors matter, and contrary to much conventional wisdom, the choice of issuance currency is not immune to perceptions of misalignment or of uncovered yield.

Issuance of foreign-currency-denominated debt securities has been an important feature in global financial markets for many years, with net issuance more than tripling in value during the past decade (measured at constant exchange rates), reaching USD 1.4 trillion in 2007. The choice of issuance currency is affected by a number of factors. One major factor is the issuer's desire to ensure its financial obligations are in currencies that match the currencies of its cash inflows. By doing so, the issuer creates a "natural hedge" against its currency risk. Another factor is strategy. The issuer's strategic considerations may include the desire to diversify its investor base and, for large-size bond issues, the opportunity to exploit fewer credit constraints in more liquid, foreign bond markets. A third factor affecting the choice of issuance currency (and a factor that is not well explored in the academic literature) is the scope for reductions in borrowing costs through issuing bonds in whichever currencies offer the lowest effective cost of capital. Lower effective borrowing costs can mean lower covered costs (incorporating the cost of covering against exchange-rate risk) or lower nominal costs, reflecting, simply, lower nominal interest rates. Anecdotally, participants in the international bond markets report that both covered and uncovered costs play important roles in the choice of issuance currency.

It is this third factor, the scope for borrowing-cost savings, that forms the focus of this current study. Furthermore, the focus is firmly on the macroeconomic aspects that affect the cost of borrowing. This study presents an empirical assessment of the extent to which uncovered cost savings (defined as deviations from uncovered interest-rate parity) and covered cost savings (defined as deviations from covered interest-rate parity) influence the issuance of foreign-currency-denominated bonds.

This study makes three main contributions to the existing literature. First, it employs a unique dataset that draws on the entire population of international bond issues during the sample period. The second contribution is an analysis of the issuance of foreign-currency-denominated bonds by *number* of issues rather than, as is customary in the literature but less appropriate, by *value* of issues (that is, this paper draws on count-data techniques). Third, this study embeds its model of bond issuance within a framework of random utility maximisation.

The first contribution of this study is its dataset, which incorporates, as far as the author is aware, the largest sample of bond issues to ever have been used in a study of this kind, with the value, at issuance, of the final sample having an aggregate US dollar equivalent of $29 trillion. This study is the first to use this dataset. Perhaps the most important unique feature of the dataset, after its scale, is that it is constructed in a manner that allows for an assessment of bond-issuance behaviour by maturity. Bonds of a given maturity are matched with interest rates and swaps of the same maturity. Therefore this study avoids the inaccurate assumption (implicit in studies that pool all maturities together) that bond issuers make consistent errors of judgement in the term structure of their hedging strategies. The frequency of the data is quarterly and the sample includes foreign-currency-denominated bonds issued by all issuer types (eg, corporate, governmental, agency, financial, supranational) from a total of 116 countries over the period 1999 to 2008.[1] The sample covers bonds issued in the five main international currencies of issuance: the US dollar, the euro, the Japanese yen, the Swiss franc and the UK pound.

The second contribution of this study is the analysis of bond issuance by *number* of issues rather than by *value* of issues.[2] This study is the first to take this approach. The approach is adopted because there is evidence (both anecdotal from market participants and statistical from a cursory analysis of the data), that the *number* of issues is more responsive to changes in this paper's key variables: deviations from both uncovered interest rate parity and covered interest rate parity. This is because the issuer's decision over the *value* of any bond offering tends to be determined before the actual date of the offering, sometimes up to a year before, and is affected mostly by issuer-specific factors such as retained earnings, project finance, target-debt ratios and share-price valuation.[3] Irrespective of the value of the bond issue, a broker will advise the issuer of the most advantageous time to execute the bond offering. This advice will be based, for issuers of foreign-currency-denominated bonds, on an evaluation of international financial conditions.[4] At an aggregate level, therefore, the main, detectable response to deviations from covered and uncovered interest parity, in any given period, will not, necessarily, be a change in total *value* of bonds issued in a certain currency, it will be a change in total *number* of bonds issued.

In addition, there are two empirical advantages of conducting an analysis of bond issuance by number of bonds issued. First, it eliminates the problem of valuation-effects—that is, it eliminates the problem, inherent in an analysis of the nominal value of bond issuance, of interpreting a rise in the value of bond issuance in a given currency as a rise in issuance when, in fact, it may represent nothing more than a strengthening of the issuance currency. Second,

[1] Money market instruments and debt securities with a maturity of less than one year are not included in the sample.

[2] For comparison purposes, an analysis of value of issuance is also undertaken, as described in Section 3.6.

[3] See, for instance, Myers (2001).

[4] Descriptions, presented in this paper, of the mechanics of standard bond-issuance procedures are informed by the relevant literature and by market participants, including brokers, underwriters and representatives of a number of major bond issuers.

it permits the application of count-data techniques, which offer a number of advantages over other empirical approaches to choice behaviour, most notably a freedom from the assumption of independence of irrelevant alternatives.[5] This, alone, makes count-data techniques a particularly powerful tool for tackling the question of currency choice in the issuance of foreign-currency-denominated bonds.

The third main contribution of this study is the use of an econometric model of bond issuance that sits within a framework of random utility maximisation, making the model entirely consistent with utility theory. In particular, the analysis is interpretable as describing a population of heterogenous decision-makers (here, issuers of bonds choosing between a set of issuance currencies), each of whom chooses, at each point in time, the best available alternative. More formally, issuance behaviour is compatible with a random utility model of observed choices, where the probability of choosing, in any given period, issuance currency i is equal to the probability that an issuer chosen at random from the population has a utility function that makes i the utility-maximising alternative.

Summarising the main results, this study finds that while deviations from swap-covered interest rate parity do exist—implying that issuers of foreign-currency-denominated bonds do have the opportunity, in any given period, to achieve cost savings by issuing bonds in whichever currency offers the lowest covered cost of issuance—issuers are not responsive. That is, the availability of covered borrowing-cost savings does not trigger a statistically significant response in terms of number of bonds issued. A significant response is, however, associated with deviations from uncovered interest rate parity. If, in any given period, the basis-point measure of uncovered borrowing-cost savings for, say the euro, rises by 20 basis points, then the expected number of foreign-currency-denominated bonds issued in euros increases, on average, by almost 10%. The picture is very similar when issuance is examined in terms of number of bonds issued in each of the five main issuance currencies as a share of total number of bonds issued in all currencies. For two-year-maturity bonds, a 50-basis-point increase in uncovered borrowing-cost savings is associated with a rise in currency share of more than 2 percentage points. Furthermore, in terms of number of bonds issued, financial corporations are even more responsive than the average issuer to uncovered borrowing-cost savings.

[5]The assumption of independence of irrelevant alternatives implies that the relative probability of each option is independent and so does not change if other options are added or retracted. More simply, if, given a choice between the US dollar and the euro as a currency of issuance, a bond issuer prefers the US dollar, the assumption of independence of irrelevant alternatives implies that this preference for the US dollar will not change by introducing as an additional option, the yen. But in practice it may well change (see McFadden (1980) and Luce and Suppes (1965)).

The rest of this chapter is structured as follows. Section 3.2 surveys the relevant literature. Section 3.3 presents a model of the choice of issuance currency and then describes how this model can be embedded within a framework of random utility maximisation. Section 3.4 describes the empirical treatment of the model while Section 3.5 provides a description of the data on bond issuance and on this paper's constructed measures of covered and uncovered borrowing costs. Section 3.6 presents the empirical tests and results, while Section 3.7 offers concluding remarks.

3.2 Review of the literature

Most of the existing academic studies that tackle the question of currency choice in the issuance of foreign-currency-denominated bonds attempt to explain the choice of currency as being motivated, mainly, by a desire to provide a hedge against foreign cash inflows.[6] Allayanis and Ofek (2001), for instance, find that for a sample of US nonfinancial firms the issuance of foreign-currency-denominated debt is correlated positively with foreign sales and trade. Meanwhile, Kedia and Mozumdar (2003) and Aabo (2006) find that an issuer's probability of issuing debt in foreign currencies is influenced strongly by the presence of foreign operations.

Beyond this, other studies suggest that credit constraints in domestic bond markets provide an impetus for issuing bonds in foreign currencies. Kedia and Mozumdar (2003) and Siegfried et al. (2007) find that large corporations (assumed to be more likely to encounter credit constraints at home) tend to issue more in foreign currencies. Following a similar logic, Eichengreen and Hausmann (1999) suggest that for many emerging economies the domestic currency cannot be used to borrow abroad or to borrow long term, even domestically. Financial markets are, in other words, incomplete.[7]

All of these studies offer plausible explanations for the issuance of foreign-currency-denominated bonds. What they ignore, however, is the possibility that issuance in a foreign currency is driven by an opportunistic desire to lower costs. That is, they ignore the possibility that issuers choose to issue bonds in a foreign currency simply because the chosen currency offers, at the time of issuance, lower effective borrowing costs than any other currency.

A number of studies assess, indirectly, the question of cost-reduction in the issuance of foreign-currency-denominated bonds. Graham and Harvey (2001) find that 44% of the corporations in their survey cite lower borrowing costs as an important reason for issuing bonds denominated in foreign currencies. Johnson

[6]See Allayannis et al. (2003), Elliot et al. (2003), and Kedia and Mozumdar (2003).
[7]See also Hausmann and Panizza (2003).

(1988) finds that Canadian financial firms issue more debt in currencies that offer lower expected service costs, while Keloharju and Niskanen (2001) find that Finnish firms tend to issue bonds when the nominal interest rate for the loan currency, relative to other currencies, is lower than usual. East Asian non-financial firms are found by Allayannis et al. (2003) to react in a similar fashion to nominal interest-rate differentials. However, Henderson et al. (2006), investigating debt issues into the US, find only weak support for the proposition that companies issue debt overseas in order to profit from lower borrowing costs, while Cohen (2005) finds that bond issuance in a given currency tends to increase with higher, not lower, interest rates in that currency.

The idea that cost savings can be secured by issuing bonds in low-interest-rate currencies does, of course, violate traditional interest-rate-parity conditions. The condition of uncovered interest rate parity asserts that any discount in foreign interest rates will be offset exactly by the expected appreciation of the foreign currency. If this parity condition holds true, it leaves no scope for exploitable cost savings from opportunistic issuance.

Empirically, however, uncovered interest rate parity does not, in general, hold true (see, for instance, Isard (1996)). Most empirical studies find that low-interest-rate currencies do not systematically appreciate over time, as suggested by uncovered interest parity. In fact, they tend to do the opposite: they depreciate. This suggests that in practice there are cost savings to be secured by leaving exchange-rate risk uncovered and issuing bonds in low-interest-rate currencies.

Of course, issuers may opt against leaving exchange-rate risk uncovered. They may, if risk-averse, prefer to purchase cover. In this case, traditional parity conditions once again state that there is nothing to gain, in terms of cost-reduction, from issuing in one currency as opposed to another: covered interest rate parity states that foreign interest costs are always equal to domestic interest costs once the price of hedging against exchange-rate risk is taken into account. The implication is that if covered interest rate parity holds true in practice, there are no profitable arbitrage opportunities to be had from issuing in a rival currency in an attempt to lower covered borrowing costs.

Most empirical studies suggest that covered interest rate parity, unlike uncovered interest rate parity, does indeed hold true.[8] Transactions costs associated with the relevant arbitrage strategies tend to be small and so deviations from covered interest rate parity tend to be negligible (see Taylor (1987)). It is important to note, however, that most of this evidence in support of covered interest rate parity is based on empirical studies that look at time horizons of

[8] See, for instance, Taylor (1992) and Willet et al. (2002) for surveys of the literature on covered interest rate parity.

less than one year, with cover provided by the forward market. These horizons are too short to be relevant for the vast majority of international bond issuance, where bond maturities can range from one year to twenty years and beyond. The forward market becomes illiquid for time horizons much greater than a year and the potential cost of arbitrage strategies becomes, correspondingly, higher. Indeed, studies of long-term covered interest rate parity tend to reveal deviations from parity that are much larger and more persistent than those for short-term covered interest rate parity (Fletcher and Taylor, 1996).

For issuers of foreign-currency-denominated bonds, exchange-rate cover is provided not by the forward market, but by the swaps market. Specifically, issuers obtain cover by purchasing a currency swap (or an appropriate combination of currency swaps). By issuing a foreign-currency-denominated bond and combining it with a currency swap, an issuer can transform its fixed-rate *foreign* payments into fixed-rate *domestic* payments, remaining entirely free of exchange-rate risk. If swap-plus-bond yields are not constant across currencies (implying a violation of swap-covered interest rate parity), the issuer can reduce its total borrowing costs by issuing its bonds in whichever currency is associated with the lowest swap-inclusive yields.

The challenge is to verify this empirically. McBrady and Schill (2007) offer a recent empirical appraisal of deviations from both uncovered and covered interest rate parity for issuance of foreign-currency denominated bonds, concentrating on a small sample of issuers with no foreign subsidiaries or foreign-currency cash flows. They find that issuance of bonds responds to deviations from both uncovered and covered interest rate parity. This study offers a closer examination of the responsiveness of bond issuance to covered and uncovered cost savings, drawing on a large, unique dataset, employing a utility-consistent model, and adopting a novel empirical approach to tackle the question of currency choice in international bond issuance, focussing in particular on the number, not the value, of bonds issued in international currencies.

3.3 The Model

This study models currency choice in the issuance of foreign-currency-denominated bonds within a framework of random utility maximisation. In the model issuers of foreign-currency-denominated bonds choose, all else remaining equal, to issue bonds in currencies that offer the lowest cost of borrowing either including, or excluding, the cost of hedging against exchange-rate risk.

Furthermore, it is the central tenet of this study that when borrowing costs in a given currency are low, the main detectable response, in terms of issuance, is an increase in the *number* of bonds issued in that currency and not necessarily an increase in the *value* of bonds issued in that currency. This draws on the notion that any window of opportunity offering lower borrowing costs in a given currency will result in a greater number of entities issuing bonds in the low-cost currency irrespective of the total value of issuance.[9]

This section offers an outline of the model of currency choice focusing on a description of the main explanatory variables, "uncovered cost savings" and "covered cost savings", in Section 3.3.1 and Section 3.3.2. Thereafter follows a brief discussion of how this model of currency choice fits into a framework of utility maximisation when issuance is measured in terms of number of bonds issued rather than value of bonds issued (Section 3.3.3).

Consider first an issuer that chooses to issue bonds denominated in foreign currency for one reason only: to act as a natural hedge (an offsetting liability) against its foreign-currency cash inflows (inflows generated perhaps by foreign assets such as overseas subsidiaries). If such an issuer has h per cent of its cash-inflow-generating stock of assets denominated in foreign currency and, in each period t, the issuer is faced with the question of what proportion, b, of its borrowing to denominate in foreign currency, then in order to create a suitable natural hedge against its foreign-currency cash inflows, the issuer will choose to issue foreign-currency-denominated bonds such that $b_t = h_t + e_t$, where the random error $e_t \sim N(0, \sigma^2)$.[10] The main concern of this study is to test whether an issuer might choose to alter the currency composition of its foreign borrowing, and deviate from h, in order to reduce its borrowing costs.

By altering the currency composition of its foreign debt an issuer can bring about a reduction in its overall borrowing costs through two main channels. First, an issuer may decide to leave its foreign-exchange risk unhedged in an attempt to gain from favourable deviations from uncovered interest rate parity. In other words, an issuer can reduce its borrowing costs by issuing bonds in foreign currencies that *ex post* do not appreciate enough to offset the savings accrued through borrowing at lower interest rates.[11] This approach offers "uncovered cost savings". Second, an issuer can hedge its foreign-currency risk and look for arbitrage, risk-free, opportunities to lower borrowing costs when deciding the currency choice of issuance. In this case, the issuer can reduce its costs by issuing bonds in low-interest-rate currencies even after accounting for the additional cost of covering for (hedging against) exchange-rate risk. The next two sections discuss and explain these strategies.

[9]See, for instance, Fisher et al. (1989) and Graham and Harvey (2001).

[10]See also McBrady and Schill (2007) and Allayanis and Weston (2001).

[11]Alternatively, it is possible to reduce borrowing costs by issuing in foreign currencies for which interest rates are relatively higher, but that *ex post* depreciate so much as to offset the extra cost associated with higher interest rates.

3.3.1 Uncovered cost savings

In the absence of exchange-rate hedging, an issuer of foreign-currency-denominated bonds can realise savings on its borrowing costs if (i) it issues in a low-interest-rate currency that does not appreciate enough to offset the savings accrued from the favourable interest-rate differential, or (ii), it issues in a high-interest-rate currency that depreciates so much as to offset the extra cost incurred from the unfavourable interest-rate differential. Such savings are possible only if uncovered interest rate parity does not hold and, as has been discussed in Section 3.2, empirical evidence suggests that it does not. That is, most evidence suggests a failure of the standard expression of uncovered interest rate parity,

$$r_{t,t+k} = r^{\star}_{t,t+k} + (s^{e}_{t,t+k} - s_t) \tag{3.1}$$

where $r_{t,t+k}$ is the time t home interest rate (compounded continuously) that pertains over time interval $t + k$, where $r^{\star}_{t,t+k}$ is the time t foreign interest rate (again, compounded continuously) defined over the same interval, where s_t is the log of the spot exchange rate (defined in terms of home currency per foreign currency), and where $(s^{e}_{t,t+k} - s_t)$ is the expected rate of foreign-currency appreciation (compounded continuously) during the time interval $t + k$.

In Eqn.(3.1), the implication is that the domestic interest rate should, in frictionless markets with perfect foresight, equal the foreign interest rate plus the expected rate of foreign-currency appreciation. But if the empirical evidence is right and the foreign currency tends, in practice, to depreciate rather than appreciate when foreign interest rates are lower than domestic interest rates, then an issuer, by issuing in a low-interest currency while leaving its currency risk uncovered, can realise expected borrowing-cost savings equal to

$$\varepsilon^{u}_{t} \equiv (r_{t,t+k} - r^{\star}_{t,t+k}) - (s^{e}_{t,t+k} - s_t) \tag{3.2}$$

Of course, in Eqn.(3.2), ε^{u} is an "expected" cost saving and, as a result, risk aversion will reduce the sensitivity of an issuer to ε^{u}.

As McBrady and Schill (2007) suggest, the proportion, b, of debt that the issuer may decide to denominate in a foreign currency will be a positive function of ε^{u}. Combining this with the fraction of borrowing set aside as a natural hedge, h, against the issuer's foreign-currency cash inflows, gives

$$b_t = h_t + \beta^{u}\varepsilon^{u}_{t} + e_t \tag{3.3}$$

where any expected uncovered savings in borrowing costs will cause the issuer to increase b by an amount equal to $\beta^{u}\varepsilon^{u}$. Likewise, b will decrease by this amount when ε^u is negative.

3.3.2 Covered cost savings

In perfectly integrated and liquid financial markets, where it is possible to hedge foreign-exchange risk at a low cost, there is no opportunity for cost savings to be made by borrowing in one currency rather than another. In this case, the cost of borrowing is identical irrespective of the borrower's choice of currency (or equivalently, the issuer's choice of issuance currency). More explicitly, arbitrage will ensure the maintenance of covered interest parity, implying that interest rates across countries will be the same once the cost of hedging foreign-currency exposure is taken into account. Covered interest rate parity can, in the absence of a risk premium, be expressed as

$$r_{t,t+k} = r^{\star}_{t,t+k} + (f_{t,t+k} - s_t) \tag{3.4}$$

where $r_{t,t+k}$ and $r^{\star}_{t,t+k}$ are defined as before, where $f_{t,t+k}$ is the log of the forward exchange rate for k periods into the future and where s_t is the log of the spot exchange rate (defined in terms of home currency per foreign currency). The quantity $(f_{t+k} - s_t)$ is the forward premium, and represents the price paid in the forward market, over and above the spot exchange rate, to cover the foreign-currency exposure that is incurred by borrowing at foreign interest rate $r^{\star}_{t,t+k}$. If covered interest parity holds, the implication is that covered foreign borrowing is no cheaper, or more expensive, than uncovered home borrowing.

Theory suggests that Eqn.(3.4) will hold true in frictionless markets and empirical evidence suggests that covered interest rate parity is indeed the rule rather than the exception. However, most empirical studies of covered interest rate parity deal with time horizons of less than one year. These horizons are too short to be relevant for the vast majority of international bond issuance, where bond maturities can range from one year to twenty years and beyond. The forward market becomes illiquid for time horizons much greater than a year.

For issuers of foreign-currency-denominated bonds, forward cover is provided not by the forward market, but instead by the swaps market. Popper (1993) and Fletcher and Taylor (1996) explain how issuers of foreign-currency-denominated bonds cover exchange-rate risk using currency swaps.[12]

[12] Since these descriptions were first presented, in the 1990s, the swaps market has, to some extent, moved on, and covering for exchange-rate risk is no longer undertaken in precisely the same manner. Cover for an individual issue can now be acquired via a single, bespoke swap rather than a combination of standardised swaps in the manner suggested by Popper (1993). However, present-day methods of covering exchange-rate risk in the swaps market, and the pricing of this cover, are derived precisely from the underlying logic outlined by Popper (1993), and this logic is employed in this study with no known loss of accuracy.

Looking again at Eqn.(3.4), what matters for issuers seeking to cover the exchange-rate risk associated with foreign-currency-denominated bond issuance, is not the forward premium, but instead the difference in continuously-compounded currency swap yields, such that

$$r_{t,t+k} = r^{\star}_{t,t+k} + (c^{sw}_{t,t+k} - c^{sw^{\star}}_{t,t+k}) \tag{3.5}$$

where $c^{sw}_{t,t+k}$ is the domestic currency swap yield of the relevant maturity k, and $c^{sw^{\star}}_{t,t+k}$ is the foreign currency swap yield also of maturity k (with both yields compounded continuously).

A standard currency swap (known also as a cross-currency, interest-rate swap) transforms fixed-rate cash flows in one currency into floating-rate cash flows in US dollars. One important point to note is that a currency swap, unlike a forward contract, is not an agreement to exchange a *fixed* payment in two currencies. It is an agreement to exchange a *stream* of payments in two currencies.

An issuer of a foreign-currency-denominated bond pays the rate $r^{\star}_{t,t+k}$ to borrow in the debt securities market and then enters a swap transaction to transform its foreign-currency payment stream into a payment stream denominated in domestic currency. In the swap transaction, the issuer receives the foreign-currency swap rate, $c^{sw^{\star}}_{t,t+k}$, and pays the domestic-currency swap rate, $c^{sw}_{t,t+k}$. In this way, the issuer of the foreign-currency-denominated bond creates a "synthetic" domestic-currency bond, incurring a cost equal to $r^{\star}_{t,t+k}+(c^{sw}_{t,t+k}-c^{sw^{\star}}_{t,t+k})$.

Eqn.(3.5) indicates that the cost of this "synthetic" domestic-currency bond must be equal to the cost, $r_{t,t+k}$, of issuing directly in domestic currency. Covered borrowing-cost savings will exist if the spread between bond yields and currency-swap rates is not equal across currencies and is not arbitraged away. The magnitude of any covered borrowing-cost savings, ε^c, will equal

$$\varepsilon^c_t \equiv (r_{t,t+k} - c^{sw}_{t,t+k}) - (r^{\star}_{t,t+k} - c^{sw^{\star}}_{t,t+k}) \tag{3.6}$$

where the implication is that an issuer of bonds can achieve savings on its covered foreign-currency borrowing whenever the spread between foreign-currency-denominated bond yields and swap rates, $(r^{\star}_{t,t+k} - c^{sw^{\star}}_{t,t+k})$, is less than the spread between domestic bond yields and swaps rates, $(r_{t,t+k} - c^{sw}_{t,t+k})$. Put simply, an issuer can lower its borrowing costs by an amount ε^c if, rather than issue bonds in domestic currency, the issuer chooses instead to issue in foreign currency and swap its foreign-currency-denominated bond payments back into domestic currency. Since the complete currency-swap arrangement allows the domestic-currency principal and foreign-currency principal to be exchanged at maturity at the original exchange rate, the issuer accrues its cost saving, ε^c, with no exposure to exchange-rate risk.

Following McBrady and Schill (2007), incorporating ε^c into the issuance decision, it is now possible to hypothesize that an issuer chooses the foreign-currency share, b, of its total borrowing according to

$$b_t = h_t + \beta^c \varepsilon^c_t + e_t \tag{3.7}$$

whereby, in response to positive ε^c_t, the issuer is expected to increase b_t by an amount equal to $\beta^c \varepsilon^c_t$. The coefficient β^c measures the unit response of foreign-currency borrowing share to the percentage change in covered cost savings. For any period t, if ε^c_t takes a negative value, b_t will decrease rather than increase.

Multiple-currency model

Empirically, the next challenge is to construct measures of both covered cost savings and uncovered cost savings that accommodate a choice among multiple currencies. There are five currencies in the sample. So far in this study the two measures of borrowing-cost savings, ε^c and ε^u, allow for just two currencies: domestic and foreign.

To accommodate a multiple-currency framework in the calculation of covered cost savings, the foreign interest rate, $r^{\star}_{t,t+k}$, in Eqn.(3.6), is replaced by $r_{i(t,t+k)}$, representing the continuously compounded yield on the k-year-maturity benchmark government bond associated with issuance currency i (where i = euro, US dollar, yen, UK pound or Swiss franc), and where yields are calculated at the start of quarter t.[13]

Meanwhile, the domestic interest rate, $r_{t,t+k}$, in Eqn.(3.6), is redefined as $\overline{r}_{i(t,t+k)}$, the contemporaneous average of all benchmark government bond yields for currencies L ($l = 1, \ldots, L$), where L includes all currencies in the sample other than the currency of issuance, i (that is, L includes all currencies associated with the nationalities of the issuers, plus the issuance currencies other than the issuance currency selected, i).

The contemporaneous average yield is, in fact, a weighted average, where weights reflect the value (US-dollar equivalent) of bonds issued in each currency l at the end of the previous quarter, $t-1$. The logic behind weighting yields by value is straightforward. Value, in this case, is used as a proxy for liquidity.

[13]Government bond yields are used to proxy borrowing costs for a number of reasons. First, as highlighted by McBrady and Schill (2007), government bond yields, unlike corporate bond yields, are free of contamination from default-risk pricing, which may otherwise affect an issuer's choice of issuance currency. Second, yields on investment-grade corporate bonds (which may could be a better proxy for the borrowing costs faced by issuers of foreign-currency-denominated bonds) are unavailable for all currencies. Government bond yields are obtained from Bloomberg.

All else being equal, an issuer, in making a comparison between borrowing costs available in the issuance currency ($r_{i(t,t+k)}$) and in rival currencies ($\overline{r}_{i(t,t+k)}$), will, among the rival currencies, be more concerned about borrowing costs available in currencies associated with liquid markets for debt. The more liquid the market, the more attractive it will be as an alternative to the issuance-currency market. Weighting yields by value does, therefore, allow liquidity to be incorporated directly into the issuer's decision over currency choice.

In order to complete the adjustments necessary to reset Eqn.(3.6) into a multiple-currency framework, adjustments are made to the empirical treatment of fixed-for-floating currency swaps. The treatment adopted is identical to that outlined for interest rates, above. That is, $c^{sw^\star}_{t,t+k}$ is replaced with $c^{sw}_{i(t,t+k)}$, representing the currency-swap rate, continuously compounded, for currency i and maturity k, while $c^{sw}_{t,t+k}$ is replaced with $\overline{c}^{sw}_{t,t+k}$, the contemporaneous weighted average of all currency-swap rates for currencies L ($l = 1, \ldots, L$), where L includes all currencies in the sample other than the currency of issuance, i. The new, multiple-currency formulation, is,

$$\varepsilon^c_i t \equiv (\overline{r}_{t,t+k} - \overline{c}^{sw}_{t,t+k}) - (r_{i(t,t+k)} - c^{sw}_{i(t,t+k)}) \tag{3.8}$$

where, in a similar fashion to Eqn.(3.6), $(\overline{r}_{t,t+k} - \overline{c}^{sw}_{t,t+k})$ is the average weighted spread between bond yields and currency-swap rates for all currencies L and, likewise, $(r_{i(t,t+k)} - c^{sw}_{i(t,t+k)})$ is the spread for issuance currency i.

Unfortunately, while fixed-for-floating currency swaps are the appropriate measure of the cost of covering exchange-rate risk for issuance of foreign-currency-denominated bonds, consistent time-series data on currency swaps are unavailable. A proxy is required. One amenable proxy, for which data are available, is the *interest-rate swap*. An interest-rate swap is a mechanism that allows fixed-rate payments in one currency to be swapped into floating-rate payments in the same currency. It differs, in magnitude, from a fixed-for-floating currency swap by an amount equal, in basis points, to a *currency basis swap*, which, itself, represents a swap of floating-rate payments in one currency into floating-rate payments in US dollars. This relationship between the three swap transactions (currency swap, interest-rate swap and currency-basis swap) can be expressed as,

$$c^{sw}_{t,t+k} = c^{bsw}_{t,t+k} + i^{sw}_{t,t+k} \tag{3.9}$$

where $c^{sw}_{t,t+k}$, as before, is the domestic fixed-for-floating currency swap, where $c^{bsw}_{t,t+k}$ is the domestic currency basis swap and where $i^{sw}_{t,t+k}$ is the domestic interest-rate swap.

An interest-rate swap is a good proxy for a currency swap only if it can be established that currency basis swaps are small, in magnitude, compared with both $c^{sw}_{t,t+k}$ and $i^{sw}_{t,t+k}$. This is, in fact, the case. Although there is insufficient data upon which to conduct tests of measurement error, the data that are available for $c^{bsw}_{t,t+k}$, for the five main currencies of issuance in the sample, show that currency basis swaps vary by no more than 20 basis points throughout the sample period (that is, they are bounded above by positive 10 basis points, and below by negative 10 basis points).

Uncovered cost savings, ε^u_t, can also, like the concept of covered cost savings, be translated into a multiple-currency framework. Interest rates are dealt with as before. That is, the foreign interest rate, $r^\star_{t,t+k}$, in Eqn.(3.2), is replaced with $r_{i(t,t+k)}$, the continuously compounded yield on the k-year-maturity benchmark government bond associated with issuance currency i. The domestic interest rate, $r_{t,t+k}$, is redefined as $\overline{r}_{i(t,t+k)}$ the contemporaneous average of all benchmark government bond yields for all currencies other than the currency of issuance, i.

A full re-expression of uncovered cost savings, ε^u_{it}, as a multiple-currency variable requires that $(s^e_{t,t+k} - s_t)$, the expected appreciation of the foreign currency, be set in a new framework that gauges appreciation not as a bilateral concept, but as a multilateral concept, with appreciation of the issuance currency measured against all other currencies. In addition, a choice must be made regarding just how, empirically, to measure exchange-rate expectations.[14]

This study uses survey data to construct its measure of exchange-rate expectations. Surveyed exchange-rate expectations are obtained from Consensus Forecasts, a British-based surveyor of financial forecasters (including banks, economic consultancies and central banks). Bilateral forecasts for 14 major currencies, with two-year forecast horizons (the longest available horizons), are used to calculate implicit forecasts of the nominal effective exchange rates for each of the five currencies of issuance in the sample. This multiple-currency formulation of exchange-rate expectations permits a complete, re-expression of

[14]Typically, in empirical work, there are four different approaches available for modelling expected changes in the exchange rate. One approach is to assume perfect foresight and measure expected changes in the exchange rate by observing *ex post* changes. That is, assume $(s^e_{t,t+k} - s_t) = s_{t,t+k} - s_t$. The drawback with this approach is that when expectation horizons are lengthy, as is the case in this study, with horizons of up to ten years, then putting aside observations to be used as *ex post* measures of expected changes in the exchange rate causes the sample size to become prohibitively small. Two alternative approaches are to assume static expectations, letting $(s^e_{t,t+k} - s_t) = 0$, and extrapolative expectations, where $(s^e_{t,t+k} - s_t) = s_t - s_{t,t-k}$. The static-expectations approach is based on the idea that exchange rates follow a random walk, while extrapolative expectations assume a backward-looking behaviour. Although the theoretical basis for this seems unsound, in practice the difference in results between from an extrapolative-expectations model and a perfect-foresight model can be quite small (see, for instance, Cavaglia et al. (12) and MacDonald and Torrance (1990)). A fourth approach is to use surveys of exchange-rate expectations, letting $(s^e_{t,t+k} - s_t) = s^{survey}_{t,t+k} - s_t$, in an attempt to take a direct, as much as is possible, measurement of expectations.

uncovered cost savings, ε_t^u, where,

$$\varepsilon_{it}^u \equiv (\bar{r}_{t,t+k} - r_{i(t,t+k)}) - (sn_{i(t,t+8)}^e - sn_{it}) \tag{3.10}$$

with $(sn_{i(t,t+8)}^e - sn_{it})$ representing the expected appreciation, over $t+8$ quarters, of the nominal effective exchange rate for currency i.

One criticism of Eqn.(3.10) is that a forecast horizon of eight quarters matches only one of the three maturity brackets (where the brackets are two year, five year and ten year) that define the sample. However, the vast majority of financial forecasters do not calculate forecasts for time horizons greater than eight quarters, suggesting these two-year-ahead forecasts do, in fact, represent long-term forecasts suitable for both a five-year horizon and ten-year horizon. In addition, of those forecasters that do provide forecasts with horizons greater than two years, these forecasts deviate only marginally from two-year-ahead forecasts when compared with the extent of the deviations between two-year-ahead forecasts and forecasts of less than a year.

3.3.3 Random utility maximisation

This section describes how the model, outlined above, can be set in a framework of utility maximisation when the dependent variable, issuance of foreign-currency-denominated bonds, is measured in terms of number of bonds issued rather than value of bonds issued—that is, when the model is a count-data model, with the dependent variable having no upper bound but having a lower bound of zero.[15] As discussed earlier, there is an *a priori* basis for thinking that a count dependent variable should be more responsive to changes in covered and uncovered borrowing-cost savings and, indeed, estimation results presented later confirm that this is the case.

The question of currency choice in the issuance of foreign-currency-denominated bonds, when issuance is measured in terms of number of issues, has yet, in the limited literature that addresses this question, to be phrased within a framework of random utility maximisation.[16] The econometric approach that lends itself most readily to a utility-consistent treatment of choice is the polychotomous-dependent-variable approach, where estimation of a multiple-choice discrete variable is undertaken by a generalisation of the logit and probit models. McFadden (1974) provides one of the first lasting contributions to utility-consistent, econometric modelling of choice with polychotomous dependent variables by presenting a conditional logit model based on random utility maximisation. Carl-

[15]See Cameron and Trivedi (1998) for a full discussion of count-data models.

[16]Claessens (1992) studies the optimal currency composition of external debt using a utility-maximising approach where optimal means risk-minimising, and composition refers to currency composition by value.

ton (1979, 1983), among others, employs the techniques of McFadden (1974) to address the question of industrial location within a utility-consistent framework.

The conditional logit model does, however, have its limitations, the most notable of which is its assumption of independence of irrelevant alternatives, which, in the present context of currency choice among issuance currencies, states that issuers look at all currencies as similar after controlling for the observable characteristics tested in the model. If the assumption of independent errors is violated, it can lead to coefficient estimates that are biased.

Within the conditional logit model, no study has been able to fully accommodate the independence-of-irrelevant-alternatives problem. One solution, however, is to eschew the conditional logit model and employ, instead, a Poisson model—that is, a count-data model. Unlike a conditional logit model, which deals with choice among a number of alternatives, a count-data model allows the data to be treated in terms of number of non-negative integer events per period, per choice category (where choice category, in the present context, is issuance currency).

Only recently have attempts been made to set Poisson models within frameworks of random utility maximisation. Guimaraes et al. (2003) derive a Poisson model directly from random utility maximisation by finding an equivalence relation between the likelihood function of the conditional logit and the Poisson regression. This study exploits the same equivalence relation in order to cast its count-model-based analysis in a framework of random utility maximisation.

A count model of currency choice

This section gives an explicit description of how a count model of currency choice can be set in a framework of random utility maximisation. The starting point for this description is a statement of the equivalence relation derived by Guimaraes et al. (2003).

First, assume that issuers of bonds maximise profits by minimising their costs of issuance. Consider (without, for the moment, incorporating a time dimension) J issuers of bonds ($j = 1, \ldots, J$), each of which select independently an issuance currency i from a set of N potential currencies ($i = 1, \ldots, N$), then the profit the issuer will accrue, if it selects currency i, will be,

$$\pi_{ij} = \beta' \boldsymbol{x}_i + e_{ij} \tag{3.11}$$

where β is a vector of unknown parameters, $\boldsymbol{x}_i$ is a vector of currency-specific variables (including covered and uncovered borrowing-cost savings) and e_{ij} is an identically and independently distributed random term assumed to have an

Extreme Value Type I distribution.[17][18]

Using the approach employed by McFadden (1974), it is possible to show that issuer j's probability of choosing issuance currency i, is equal to,

$$p_i = \frac{\exp(\beta' \boldsymbol{x}_i)}{\sum_{i=1}^{N} \exp(\beta' \boldsymbol{x}_i)} \tag{3.12}$$

which is the common representation of the conditional logit model. If, now, the number of bonds issued in currency i is denoted by n_i, it is possible to estimate the parameters in Eqn.(3.12) by maximising the log-likelihood function,

$$\log L_{cl} = \sum_{i=1}^{N} n_i \log p_i \tag{3.13}$$

Guimaraes et al. (2003) show that this log-likelihood function is equivalent to a Poisson model with n_i as its dependent variable and $\boldsymbol{x}_i$ as its vector of explanatory variables. In other words, a Poisson model will yield the same results if n_i conforms to a Poisson distribution such that,

$$E(n_i) = \lambda_i = \exp(\beta' \boldsymbol{x}_i) \tag{3.14}$$

where λ_i is the Poisson mean parameter (in this case, the expected number of bond issues).

The count model as outlined thus far still fails to account adequately for the possibility of a violation of the assumption of independence of irrelevant alternatives. The most straightforward answer is to add to the profit function an additional effect, γ_i, specific to each alternative, which captures all the factors that may affect the choice of issuance currency but are unaccounted for by the explanatory variables, such that,

$$\pi_{ij} = \beta' \boldsymbol{x}_i + \gamma_i + e_{ij} \tag{3.15}$$

and further, if γ_i is a random variable, then the probability of an issuer choosing currency i is,

$$p_{i/\gamma} = \frac{\exp(\beta' \boldsymbol{x}_i + \gamma_i)}{\sum_{i=1}^{N} \exp(\beta' \boldsymbol{x}_i + \gamma_i)} \tag{3.16}$$

[17] The Extreme Value Type I distribution, also known as the Weibull distribution, has the property that the cumulative density of the difference between any two random variables with this distribution is given by the logistic function. This property makes it possible to link the random utility function with the logistic function. See Maddala (1983).

[18] Profit is total. The disturbance term indicates that a proportion of the issuer's profit is stochastic. Profit is not directly comparable with the accounting definition of profit. The profit function assumes issuer homogeneity, an assumption that can be questioned for its plausibility.

where choice of issuance currency is conditional on γ.

One option for estimation of Eqn.(3.16) is to exploit the relation between the conditional logit model and the Poisson model and estimate by means of a Poisson regression with random effects. More appropriate, however, for the purposes of modelling choice among issuance currencies (where there is no guarantee that the alternative specific effects are uncorrelated with the explanatory variables), is to assume that γ_i is a fixed effect, including a dummy variable for each currency of issuance, i, so that,

$$p_i = \frac{\exp(\beta' \boldsymbol{x}_i + \gamma_i)}{\sum_{i=1}^{N} \exp(\beta' \boldsymbol{x}_i + \gamma_i)} \tag{3.17}$$

Introducing, now, a longitudinal time dimension to Eqn.(3.17), sufficient time-series variation allows estimation of all parameters of interest. As such,

$$p_{it} = \frac{\exp(\beta' \boldsymbol{x}_{it} + \gamma_i)}{\sum_{i=1}^{N} \exp(\beta' \boldsymbol{x}_{it} + \gamma_i)} \tag{3.18}$$

where p_{it} is the probability of an issuer choosing, in time t, to issue debt denominated in currency i.

The formulation in Eqn.(3.18) is, as it stands, based on the Poisson regression model, which assumes that the mean number of events per period, $\lambda_i = \exp(\beta' \boldsymbol{x}_i)$, is equal to the variance λ_i. However, for most count data and for the sample employed in this study, the mean does not equal the variance. An alternative model that relaxes this assumption of equidispersion and allows instead for overdispersion (variance greater than the mean) is the negative binomial model, which represents a generalisation of the Poisson model. The Poisson model is generalised by introducing an individual, unobserved effect into the conditional mean (Greene, 2008). It is then assumed, as per Hall et al. (1984), that the conditional mean λ_{it} follows a gamma distribution with shape parameter ϕ and scale parameter δ, specified such that $\phi = e^{X_{it}\beta}$ with δ common both across issuance currencies and across time. Taking the gamma distribution for λ_{it} and integrating by parts gives,

$$\begin{aligned} p(n_{it}) &= \int_0^\infty \frac{1}{n_{it}} e^{-\lambda_{it}} \lambda_{it}^{n_{it}} f(\lambda_{it})\, \mathrm{d}\lambda_{it} && (3.19) \\ &= \frac{\Gamma(\phi_{it} + n_{it})}{\Gamma(\phi_{it})\Gamma(n_{it}+1)} \left(\frac{\delta}{(1+\delta)}\right)^{\phi_{it}} (1+\delta)^{-n_{it}} && (3.20) \end{aligned}$$

which is the negative binomial distribution with parameters (ϕ_{it}, δ), where $\Gamma(.)$ is the gamma function. In order to add issuance-currency-specific effects (that is, fixed effects) to the negative binomial model, the approach of Hall et al.

(1984) is adopted, allowing for the construction of the joint probability of bond issuance in a given currency conditional on the full-period total of bond issues, such that,

$$p(n_{i1}, \ldots, n_{iT}) = \frac{\Gamma\left(1 + \sum_{t=1}^{T_i} n_{it}\right) \Gamma\left(\sum_{t=1}^{T_i} \lambda_{it}\right)}{\Gamma\left(\sum_{t=1}^{T_i} n_{it} + \sum_{t=1}^{T_i} \lambda_{it}\right)} \prod_{t=1}^{T_i} \frac{\Gamma(n_{it} + \lambda_{it})}{\Gamma(1 + n_{it})\Gamma(\lambda_{it})} \quad (3.21)$$

which is the specification used for the empirical analysis presented in the next section.

3.4 Empirical Methodology

This section presents an overview of the empirical techniques used to estimate the model introduced above. The thesis that currency choice in bond issuance is affected by covered and uncovered cost savings is tested, first, in a model where the dependent variable is a count variable, defined as number of bonds issued in a given currency at time t (Section 3.4.1). In an extension of this approach, and as a robustness check, estimation is also undertaken with the dependent variable expressed as number of bonds issued in a given currency *as a share* of all bonds issued (Section 3.4.2).

3.4.1 Count model empirical methodology

This section presents the empirical counterpart to the discussion in Section 3.3.3 of a count-model approach to choice among issuance currencies. Allowing for unobserved heterogeneity across issuance currencies, fixed-effects panel regressions are estimated in a manner suitable for a dependent variable that behaves as a count variable, in this case the number of bonds issued in currency i. Recall that a count variable is bounded from below by zero and has no effective upper limit. Estimation is by means of a negative binomial model, which accounts for overdispersion in the data (that is, accounts for the fact that the variance of the dependent variable can, and often does, exceed its mean).[19]

Within this fixed-effects framework, the number of bonds issued in currency i is assumed to depend on a vector of explanatory variables such that,

$$B_{it}^c = \alpha_i + \beta K_{it} + \gamma R_{it} + e_{it} \quad (3.22)$$

where B_{it}^c is the dependent variable defined as number of bonds issued in currency i in period t, where α_i is a currency-specific fixed effect, where K_{it} is

[19]See Hall et al. (1986) for a discussion of the fixed effects model in a negative binomial setting.

a vector of variables representing the incentive to issue bonds denominated in currency i in order for these bonds to act as a natural hedge (as outlined in Section 3.3), and where R_{it} is a vector of variables representing both covered and uncovered cost savings.

In accounting for an empirical representation of aggregate tendency to issue bonds in currency i as a natural hedge, the vector K contains variables that reflect issuance-currency-country fundamentals plus variables that capture the scale of foreign-owned, cash-flow-generating assets in the issuance-currency region.[20] Specifically,

$$K_{it} = \beta_1 rgdp_{it} + \beta_2 ma_{it} + \beta_3 dinv_{it} + \beta_4 liq_{it} \tag{3.23}$$

where the frequency of all data is quarterly and where $rgdp$ is real GDP (in constant US dollar millions) in the issuance-currency country (or region) as a share of total GDP across all issuance-currency countries; where ma is the number of cross-border mergers and acquisitions into the issuance-currency country (or region), by acquirers that match, in nationality, the set of issuers in the given currency (again this is measured as a share of total mergers and acquisitions in all issuance-currency countries); where $dinv$ is direct investment in the issuance-currency country (or region) in US dollar millions as a share of total direct investment into all issuance-currency countries; and where liq is a proxy for financial depth, represented by total issuance of bonds and notes in the issuance currency (both domestic and foreign issues), divided by GDP (and, again, measured as a share of total liquidity in all issuance currencies). For further details see Table 3.1.

The vector of variables, R, in Eqn.(3.22), contains the two main variables of interest, namely, covered cost savings, ε^c_{it}, and uncovered cost savings, ε^u_{it}. These two variables are measured at the beginning of each quarter and expressed in terms of basis points. If issuers respond, as expected, to covered cost savings in currency i by issuing, in aggregate, more bonds denominated in currency i, then the parameter estimate for ε^c_{it} should be positive. Equally, if issuers respond, as expected, to uncovered cost savings in currency i, then the parameter estimate for ε^u_{it} should be positive also.

It should be note that one limitation of the empirical approach adopted is its vulnerability to omitted-variable bias. This bias stems from the use of issuer-country averages, such as $\overline{r}$. To see this, consider the unrestricted interest-rate differential for issuer l issuing in currency i, that is, $\sum_{j\neq l}\theta_j(r_j - r_i)$. This is

[20] Choice of these variables draws on the findings of other studies that account for the natural hedge, such as Cohen (2005) and Siegfried et al. (2007). Other variables, such as imports and investment in the issuance-currency region, were discarded when found to be statistically insignificant in all cases.

Table 3.1: Data Sources And Definitions

Variable	Definition	Source
B_i^c	Number of issues of foreign-currency bonds denominated in currency i, where foreign-currency bonds are defined as all bonds issued in a currency other than the currency of the country in which the borrower resides	Dealogic (Bondware)
B_i^s	Number of issues of foreign-currency bonds denominated in currency i as a share of all issues of foreign-currency bonds (share expressed as fraction of one)	Constructed variable
ε^u	Uncovered borrowing-cost savings, defined as deviations from uncovered interest rate parity, where $\varepsilon_{it}^u \equiv (\overline{r}_{t,t+k} - r_{i(t,t+k)}) - (sn_{i(t,t+8)}^e - sn_{it})$	Constructed variable
ε^c	Covered borrowing-cost savings, defined as deviations from swap-covered interest rate parity, where $\varepsilon_i^c t \equiv (\overline{r}_{t,t+k} - \overline{c}_{t,t+k}^{sw}) - (r_{i(t,t+k)} - c_{i(t,t+k)}^{sw})$	Constructed variable
$(\overline{r} - r_i)$	Interest rate differential, defined as home interest rate minus issuance-currency interest rate, where home interest rate $\overline{r}$ is expressed as a multiple-currency average	Constructed variable
$\overline{r}$	Contemporaneous average of all interest rates for currencies L $(l = 1, \ldots, L)$, where L includes all currencies in the sample other than the currency of issuance, i (that is, L includes all currencies associated with the nationalities of the issuers, plus the issuance currencies other than the issuance currency selected, i). The average is a weighted average, where weights reflect the value (US-dollar equivalent) of bonds issued in each currency l at the end of the previous quarter, $t-1$	Constructed variable
r	Yield on benchmark government bond, compounded continuously	Bloomberg
$(sn_i^e - sn_i)$	Expected appreciation of the nominal effective exchange rate for issuance currency i (index)	Constructed variable
sn_i^e	Expected value of the nominal effective exchange rate for issuance currency i, with weights calculated to match trade weights for sn (index)	Constructed variable
sn_i	Nominal effective exchange rate for issuance currency i	International Financial Statistics, IMF
s^e	Exchange-rate expectations, natural logarithm of, where expectations are proxied by two-year ahead consensus forecasts	Consensus Forecasts
s	Exchange rate, natural logarithm of, expressed in terms of home currency per foreign currency	Bloomberg
c	Benchmark currency-swap yield (proxied by interest-rate-swap yield)	Bloomberg
ma	Number of cross-border mergers and acquisitions into the issuance-currency region (by acquirers that match, in nationality, the set of issuers in the issuance currency) as a proportion of cross-border mergers and acquisitions into all issuance-currency countries (in percentage points)	Zephyr, Bureau Van Dijk
$rgdp$	Constant GDP in the issuance-currency region as a share of total constant GDP in all other issuance-currency countries (in percentage points)	International Financial Statistics, IMF
$dinv$	Direct investment into the issuance-currency region as a share of total direct investment into all sample issuance-currency regions (in percentage points)	International Financial Statistics, IMF
liq	Capitalisation of market for issuance-currency debt securities (both domestic bonds and foreign bonds) divided by issuance-currency GDP, as a share of total capitalisation of market for all debt securities in all issuance currencies divided by total GDP (in percentage points)	Dealogic (Bondware)

replaced with $(\overline{r}_l - r_i)$. The concern is that this could lead to omitted variable bias where

$$\sum_{j\neq l}\theta_j(r_j - r_i) = \sum_{j\neq l}\theta_j r_j - \Big(\sum_{j\neq l}\theta_j\Big)r_i \tag{3.24}$$

$$= \Big(\sum_{j\neq l}\theta_j\Big)(\overline{r}_l - r_i) + \sum_{j\neq l}\theta_j r_j - \Big(\sum_{j\neq l}\theta_j\Big)\overline{r}_i \tag{3.25}$$

$$= \Big(\sum_{j\neq l}\theta_j\Big)(\overline{r}_l - r_i) - \sum_{j\neq l}\theta_j(\overline{r}_l - r_j) \tag{3.26}$$

$$\tag{3.27}$$

which shows that using only $(\overline{r}_l - r_i)$ omits $\sum_{j\neq l}\theta_j(\overline{r}_l - r_j)$; at least some of the $(\overline{r}_l - r_j)$ could play significant role.

3.4.2 Currency share empirical methodology

An alternative to addressing the question of currency choice through a count-data approach is to adopt an approach wherein the dependent variable is transformed so as to represent the number of bonds issued in currency i as a *share* of total number of bonds issued in all currencies. Currency share is an alternative gauge of currency choice and, as such, an empirical analysis of currency share acts as a robustness check on the results from the model presented in Section 3.4.1.

For currency share, the count variable, B^c_{it} is replaced with a share variable, B^s_{it}, such that,

$$B^s_{it} = \alpha_i + \beta K_{it} + \gamma R_{it} + e_{it} \tag{3.28}$$

Transforming the dependent variable into a share variable is not without consequence. The dependent variable is, now, bounded between zero and one, and can, in theory, include both zero and one. The most appropriate estimator for an endogenous variable with such characteristics comes from the fractional logit approach developed by Papke and Wooldridge (1996). Proper application of this estimator in a panel requires, however, that the cross-sectional dimension of the panel is large (N greater than 100), but here, this is not the case ($N = 5$). For this reason, an alternative approach is adopted that assumes, as a starting point, that a standard Gaussian model is appropriate, and deals with departures from the Gauss-Markov conditions on an *ad hoc* basis.[21]

[21] As an empirical starting point, the Gaussian model does, in fact, seem valid for the dependent variable expressed as a share variable, since there are no zero observations in the two-year-maturity sample bracket, and just 3% of observations take the value zero in the five-year-maturity bracket and the ten-year-maturity bracket. The standard linear Gaussian model requires that the mean of the dependent variable is high enough so as not to be characterised by a preponderance of zero observations.

As a share variable, however, bond issuance exhibits a number of non-standard characteristics. One of these characteristics is contemporaneous correlation across error terms because, in any given period, currency shares sum almost to one. In addition, disturbances are likely to be heteroscedastic across issuance currencies. Furthermore, it is possible that currency-specific residuals are autocorrelated, with the autocorrelation parameter either constant for all issuance currencies or, perhaps, different for each currency.

More formally, if the disturbances in Eqn.(3.28) exhibit both heteroscedasticity and contemporaneous correlation, the disturbance covariance matrix will be represented by,

$$E[ee'] = \mathbf{\Omega} = \begin{pmatrix} \sigma_{11}\mathbf{I}_{11} & \sigma_{12}\mathbf{I}_{12} & \cdots & \sigma_{1n}\mathbf{I}_{1n} \\ \sigma_{21}\mathbf{I}_{21} & \sigma_{22}\mathbf{I}_{22} & \cdots & \sigma_{2n}\mathbf{I}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{n1}\mathbf{I}_{n1} & \sigma_{n2}\mathbf{I}_{n2} & \cdots & \sigma_{nn}\mathbf{I}_{nn} \end{pmatrix}$$

where σ_{ii} is the variance of the disturbances for issuance currency i, where σ_{ij} is the covariance of the disturbances between currency i and currency j when the periods are matched, where $i = 1, \ldots, n$, and where $\mathbf{I}$ is a T_i by T_i identity matrix, with T_i the number of periods.

Since our sample contains a limited number of heterogeneous units, the best approach is to use ordinary least squares to calculate unbiased parameter estimates in the absence of autocorrelation, and calculate Prais-Winsten estimates when autocorrelation is present.[22] For all regression specifications examined, Breusch-Pagan tests reject separate null hypotheses of cross-sectional independence of the residuals.

3.5 The Data

Data on international bond offerings are obtained from the Bondware database maintained by Dealogic, a financial-information provider. This database provides coverage of the world's debt markets with information, along numerous dimensions, on the entire population of bond offerings. The sample period extends from 1999 to the second quarter of 2008, with earlier data discarded in order, primarily, to permit an examination of the role of the euro as an issuance currency. Foreign-currency-denominated bonds are defined as all non-

[22]See Prais and Winsten (1954). An alternative estimation technique would be the application of feasible generalised least squares (FGLS). However, Beck and Katz (1995) have shown that FGLS variance-covariance estimates are unacceptably optimistic when dealing with panels where the number of heterogenous units is less than 20 and where there are 40 time periods per unit or less. The implication is that FGLS is inappropriate for the purposes of the present study.

convertible, fixed-coupon, investment-grade bonds denominated in a currency other than the currency of the nationality of the issuer.[23][24]

The sample is restricted to fixed-coupon bonds, which account for 70% of the total population of issues of foreign-currency-denominated bonds within the sample period. The final data-set includes 172,352 bond offerings with an aggregate US-dollar-equivalent principal value of $29 trillion (gross issuance).

Table 3.2 displays aggregate statistics for the world's major issuance currencies ranked by outstanding amount of foreign-currency-denominated bonds issued in each currency throughout the sample period. It can be seen that a small number of issuance currencies, namely the US dollar, the euro, the yen, the UK pound and the Swiss franc, dominate aggregate offerings, with the top five accounting for 93% of total value of announced bond issuances and 87% of the total number of issuances. In the empirical work that follows, the sample is restricted to these five top issuance currencies.

Table 3.3 shows how the distribution of bond maturities, which range from 1 year to 100 years, is not uniform across the different currencies. For this reason, the sample of bonds is partitioned into three maturity brackets (two year, five year and ten year) in order to match bonds with interest rates and swap rates of corresponding maturity along the yield curve. All bonds in the sample are allocated to one of these three maturity brackets.[25] Table 3.4 presents a comparison of aggregate annual offerings of foreign-currency-denominated bonds both by value and by number. One interesting observation is that during 2007 and 2008 the share of foreign-currency-denominated bonds issued in euros dropped sharply in terms of number of bonds issued, but not in terms of value. Over the same period, interest rates were falling elsewhere in the world (most notably in the US) but not in the euro area.

[23] Issuer nationality is defined, in a manner consistent with the Bank for International Settlements, as the nationality of the upper-most level of corporate responsibility, which, as a definition, accommodates the possibility that the issuer may be part of a multinational company, eg, a subsidiary, or a branch plant.

[24] In order to ensure that the issuers in the sample are, in fact, able to exercise a reasonable choice among the five currencies in the sample, included are only those issuers that are observed to issue bonds in at least three of the five issuance currencies during the sample period. This sorting procedure is conducted by nationality rather than by individual issuer, so that if one issuer of a given nationality is observed to issue in three or more different currencies, then all issuers of the same nationality are included in the sample.

[25] Securities with maturities of one year or less are excluded because for securities with such short maturities the forward market can provide cover for exchange-rate risk. Bonds with maturities greater than 15 years are omitted in order to reduce the scope for matching errors generated by inexact matching of maturities between bonds, swap yields and interest rates. The two-year-maturity bracket includes all bonds with maturities greater than one year but less than or equal to three years. The five-year-maturity bracket includes all bonds with maturities greater than three years but less than or equal to seven years. The ten-year-maturity bracket includes all bonds with maturities greater than seven years but less than or equal to 15 years.

Table 3.2: Aggregate Issuance By Currency, 1999-2008*

	Principal Amount, US$bn	(%)	Number of offerings	(%)
US dollar	13,755	47.1	96,533	56.9
Euro	8,646	29.6	36,852	21.7
Yen	3,810	13.0	9,979	5.9
Pound sterling	700	2.4	2,075	1.2
Swiss franc	350	1.2	2,449	1.4
Australian dollar	211	0.7	1,800	1.1
Other	1,759	6.0	19,969	11.8
Total	29,231	100	169,657	100

Notes: Principal amount (value in US$bn equivalent) and number of foreign-currency-denominated bonds issued during 1999-2008 ranked according to principal amount. Percentages refer to issuance (by principal amount and by number of bonds issued) in given currency as a per cent of all foreign-currency-denominated bonds issued. Foreign-currency-denominated bonds are defined as those bonds issued in a currency other than the currency of the country in which the borrower resides. Includes only fixed-interest-rate debt securities (ie, straight bonds). Excludes debt securities with maturities of less than one year and more than 15 years. (*) Data for 2008 is for the first half of 2008. Source is Bondware.

Table 3.3: Aggregate Issuance By Maturity, 1999-2008*

	Number of offerings (per cent) in issuance currency		
	2yr maturity	5yr maturity	10yr maturity
US dollar	22.6	40.8	36.6
Euro	55.1	24.6	20.3
Yen	10.8	40.7	48.5
Pound sterling	21.9	43.3	34.8
Swiss franc	12.6	42.4	45.0

Notes: Foreign-currency-denominated bonds of specified maturity issued during 1999-2008 as a share of total foreign-currency-denominated bonds issued in selected currencies. Maturity here refers to maturity “brackets”, as described in the text. Foreign-currency-denominated bonds are defined as those bonds issued in a currency other than the currency of the country in which the borrower resides. Sample includes only fixed-interest-rate securities. (*) Data for 2008 is for the first half of 2008. Source is Bondware.

Table 3.4: Aggregate Issuance By Year Of Offering, 1999-2008*

Panel A. Total number of foreign-currency-denominated bonds offered										
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
US dollar	9,046	7,428	14,479	12,985	15,589	14,299	9,710	7,709	4,502	786
Euro	2,828	2,370	2,113	2,153	3,751	6,405	7,937	6,765	2,258	272
Yen	1,027	1,252	1,287	1,201	1,077	948	1,178	1,069	791	149
Pound sterling	160	201	136	139	171	282	263	305	239	179
Swiss franc	257	303	214	255	228	208	312	251	277	144
Other	673	1,485	2,330	2,883	3,174	3,348	3,637	3,794	2,452	688
Total	13,991	13,039	20,559	19,616	23,990	25,490	23,037	19,893	10,519	2,218
Panel B. Currency share (per cent) of total number of foreign-currency-denominated bonds offered										
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
US dollar	64.7	57.0	70.4	66.2	65.0	56.1	42.1	38.8	42.8	35.4
Euro	20.2	18.2	10.3	11.0	15.6	25.1	34.5	34.0	21.5	12.3
Yen	7.3	9.6	6.3	6.1	4.5	3.7	5.1	5.4	7.5	6.7
Pound sterling	1.1	1.5	0.7	0.7	0.7	1.1	1.1	1.5	2.3	8.1
Swiss franc	1.8	2.3	1.0	1.3	1.0	0.8	1.4	1.3	2.6	6.5
Other	4.8	11.4	11.3	14.7	13.2	13.1	15.8	19.1	23.3	31.0
Total	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Panel C. Total principal amount, gross issuance (US$bn)										
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
US dollar	1,056.7	919.7	1,842.8	1,794.5	2,060.4	1,955.6	1,601.5	1,738.3	574.6	210.8
Euro	756.2	557.2	662.7	587.3	781.7	909.1	1,956.5	1,778.5	461.8	195.0
Yen	77.4	128.8	433.3	457.7	473.8	632.5	752.0	735.8	102.4	16.7
Pound sterling	38.0	34.3	31.2	39.1	79.2	94.3	113.0	135.6	82.3	52.9
Swiss franc	34.0	34.8	26.7	34.0	34.0	32.6	49.0	41.6	39.5	23.4
Other	87.0	92.7	125.9	164.8	197.4	321.9	392.9	433.7	176.3	52.3
Total	2,049.4	1,767.5	3,122.6	3,077.4	3,626.5	3,946.0	4,864.9	4,863.5	1,436.8	551.0
Panel D. Currency share (per cent) of total principal amount, gross issuance										
	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
US dollar	51.6	52.0	59.0	58.3	56.8	49.6	32.9	35.7	40.0	38.3
Euro	36.9	31.5	21.2	19.1	21.6	23.0	40.2	36.6	32.1	35.4
Yen	3.8	7.3	13.9	14.9	13.1	16.0	15.5	15.1	7.1	3.0
Pound sterling	1.9	1.9	1.0	1.3	2.2	2.4	2.3	2.8	5.7	9.6
Swiss franc	1.7	2.0	0.9	1.1	0.9	0.8	1.0	0.9	2.7	4.2
Other	4.2	5.2	4.0	5.4	5.4	8.2	8.1	8.9	12.3	9.5
Total	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0

Notes: Foreign-currency-denominated bonds issued in selected issuance currencies during 1999-2008. Number of foreign-currency-denominated bonds issued (Panel A) in selected currencies and number issued as a share (in per cent) of all foreign-currency-denominated bonds issued (Panel B). Gross issuance total principal amount (Panel C) in US-dollar equivalents at current exchange rates. Foreign-currency-denominated bonds are defined as those bonds issued in a currency other than the currency of the country in which the borrower resides. Includes only fixed-interest-rate debt securities. Excludes debt securities with maturities of one year or less, and with maturities greater than 15 years. (*) Data for 2008 is for the first half of 2008. Source is Bondware.

3.6 Results

This section presents results from empirical tests of the hypothesis that issuers of foreign-currency-denominated bonds choose, all else being equal, to issue in currencies that offer the lowest available uncovered and covered borrowing costs. Results from count-model panel regressions are presented first, followed by results from an empirical model of currency share.

3.6.1 Count-model results

Results from empirical testing of the count model of currency choice among the five currencies of issuance (the US dollar, the euro, the yen, the UK pound and the Swiss franc) are presented in Table 3.5. Coefficient estimates and standard errors (corrected for the overdispersion in the data) are displayed for the three separate maturity brackets under analysis—two years, five years and ten years. Table 3.5 reports a number of different specifications of the basic model, in particular testing the impact of covered and uncovered borrowing costs separately (columns 1 and 2) and jointly (column 3). In addition, columns 4 and 5 isolate the separate contributions to uncovered borrowing costs of interest-rate differentials and expected exchange-rate appreciation.

Likelihood-ratio tests indicate that all specifications outperform a pooled estimator (where the negative binomial estimator takes a constant dispersion). Parameter estimates suggest issuers of foreign-currency-denominated bonds do not respond to covered cost savings, with ε_{it}^{c} proving to be statistically insignificant across all three maturity brackets. While the availability of covered cost savings appears to play a negligible role in the issuance decision, this is not the case with uncovered cost savings. Issuers appear to be responsive to uncovered cost savings when issuing bonds of all maturities. In all three maturity brackets the estimated coefficient on ε_{it}^{u} carries the expected sign, namely positive, and its magnitude is similar (around 0.3), implying that a 20 basis-point increase in uncovered borrowing-cost savings (the average absolute change in ε_{it}^{u} for bonds of all maturities during the sample period is 25 basis points) is associated with a 7% increase in the expected number of bonds issued in the issuance currency.[26]

Table 3.5 also presents a decomposition of ε_{it}^{u} into its two component parts, the interest-rate differential ($\overline{r}_{t,t+k} - r_{i(t,t+k)}$) and the expected appreciation of the issuance currency ($sn_{i(t,t+8)}^{e} - sn_{it}$).[27] Examining these two component

[26]Note that percentage change in the expected number of bonds issued for a unit change in each explanatory variable, holding other variables constant, is calculated as 100*[exp(*estimated coefficient*)-1].

[27]Regressions were also estimated with alternative approximations of "expected appreciation" (based, for example, on backward-looking extrapolative expectations), but the results were not materially different.

Table 3.5: Fixed effects negative binomial estimation

Panel A. Issuance of foreign-currency bonds, two-year maturity					
	(1)	(2)	(3)	(4)	(5)
ε^c	0.063		0.042		
	(0.12)		(0.12)		
ε^u		0.332**	0.404**		
		(0.11)	(0.11)		
$(\overline{r} - r_i)$				0.301**	0.335**
				(0.11)	(0.11)
$(sn_i^e - sn_i)$					-3.036**
					(1.07)
$rgdp$	0.024	0.032**	0.039**	0.031**	0.032**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$liq(t-1)$	0.008	0.006	0.005	0.006	0.004
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$dinv$	-0.001	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ma	0.008	0.013	0.015*	0.013	0.009
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Log likelihood	-978.703	-974.267	-972.424	-975.059	-971.041
Likelihood ratio	152.795	159.068	160.095	157.859	164.849
Panel B. Issuance of foreign-currency-denominated bonds, five-year maturity					
ε^c	-0.073		0.022		
	(0.01)		(0.11)		
ε^u		0.212**	0.219**		
		(0.08)	(0.08)		
$(\overline{r} - r_i)$				0.219**	0.218**
				(0.08)	(0.08)
$(sn_i^e - sn_i)$					0.519
					(0.75)
$rgdp$	0.030**	0.020*	0.019	0.020*	0.021*
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$liq(t-1)$	0.003	0.001	0.001	0.001	0.002
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$dinv$	0.000	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ma	0.018**	0.013**	0.013**	0.013**	0.013**
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log likelihood	-1005.457	-1002.062	-1002.041	-1001.823	-1001.585
Likelihood ratio	146.374	146.281	125.814	148.517	141.225
Panel C. Issuance of foreign-currency-denominated bonds, ten-year maturity					
ε^c	-0.129		-0.012		
	(0.09)		(0.09)		
ε^u		0.299**	0.296**		
		(0.08)	(0.08)		
$(\overline{r} - r_i)$				0.311**	0.308**
				(0.08)	(0.08)
$(sn_i^e - sn_i)$					1.054
					(0.75)
$rgdp$	0.014	0.004	0.004	0.003	0.003
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$liq(t-1)$	0.004	0.001	0.001	0.001	0.003
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$dinv$	0.000	0.000	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ma	0.009*	0.005	0.005	0.004	0.003
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Log likelihood	-1001.585	-995.025	-995.017	-994.387	-993.413
Likelihood ratio	129.332	122.63	119.45	124.336	125.762

Notes: Fixed effects, negative binomial, count-data model accommodating overdispersion. Dependent variable is number of foreign-currency-denominated bonds issued in currency i at time t. Regressions include fixed effects and year dummies. All explanatory variables are measured at the beginning of the quarter and are expressed as proportions (as explained in the text), measured in percentage points. The sample period is from 1999 to the second quarter of 2008. Standard errors are in parenthases. (**) and (*) denote significance at the 1% and 5% levels respectively. All p-values for the likelihood ratio tests are smaller than 0.001.

variables, it becomes clear that what drives the overall significance of ε_{it}^{u} is not the expected appreciation, or depreciation, of the issuance currency, but the nominal interest-rate differential.[28] The interest-rate-differential parameter is significant for bonds of all maturities and is similar, in magnitude, to parameter estimates for ε_{it}^{u}. The implication is, according to these results, that nominal interest rates do matter, whereas exchange-rate expectations are not generally relevant for the choice of currency in international bond issuance. Figure 3.1 illustrates these findings graphically for euro-denominated bonds carrying a five-year maturity: issuance can be seen to correlate strongly with interest-rate differentials, less strongly with uncovered cost savings, and correlate hardly at all with covered cost savings.

Table 3.5 shows that relative financial depth of the bond market associated with each issuance currency (*liq*), relative share of direct investment into each issuance-currency region (*dinv*) and relative share of cross-border mergers and acquisitions into each issuance-currency region (*ma*) are found in general, for bonds of all maturities, to be statistically insignificant as drivers of currency choice among issuance currencies.[29] According to our evidence, these variables fail to capture the potential incentive among issuers to issue foreign-currency-denominated bonds in order for these bonds to act as a natural hedge against foreign cash inflows. The role of the natural hedge, if present, is captured by fixed effects or, potentially, *rgdp*.

Indeed, Table 3.5 shows that economic activity within the issuance-currency region (*rgdp*) acts as a significant driver of issuance for all bonds in the sample except for those with the longest maturities. For all bonds other than those that fall into the ten-year-maturity bracket, the estimated coefficients imply that a one percentage point increase in the share of economic activity in the issuance-currency region (the average absolute change in share throughout the sample period is indeed one percentage point) is associated with an increase of roughly 4% in the number of bonds offered in the issuance currency.[30]

[28]Expected appreciation is significant as an explanatory variable for only short-maturity bonds (Panel A), where the estimated coefficient is of the expected sign, namely negative (suggesting that issuers prefer to issue bonds in currencies that they expect, broadly, to depreciate over time), and where the magnitude of the estimated coefficient implies that a one-basis-point increase in expected appreciation (the average absolute change in $(sn^{e}_{i(t,t+8)} - sn_{it})$ during the sample period is 1.5 basis points) is associated with a 9% drop in the expected number of bonds issued in the issuance currency.

[29]Recall that variables *liq*, *dinv*, *ma* and *rgdp* are expressed as shares relative to total amounts in all issuance-currency regions. These variables are expressed as relative shares in order to facilitate comparability with results presented in subsequent sections, Section 3.6.1 and Section 3.6.2 Other formulations of these variables (for instance, relative rates of change), yield similar results.

[30]Note that the variables *rgdp*, *liq*, *dinv* and *ma* are expressed in terms of percentage points.

Figure 3.1: Issuance, interest-rate differentials and cost savings

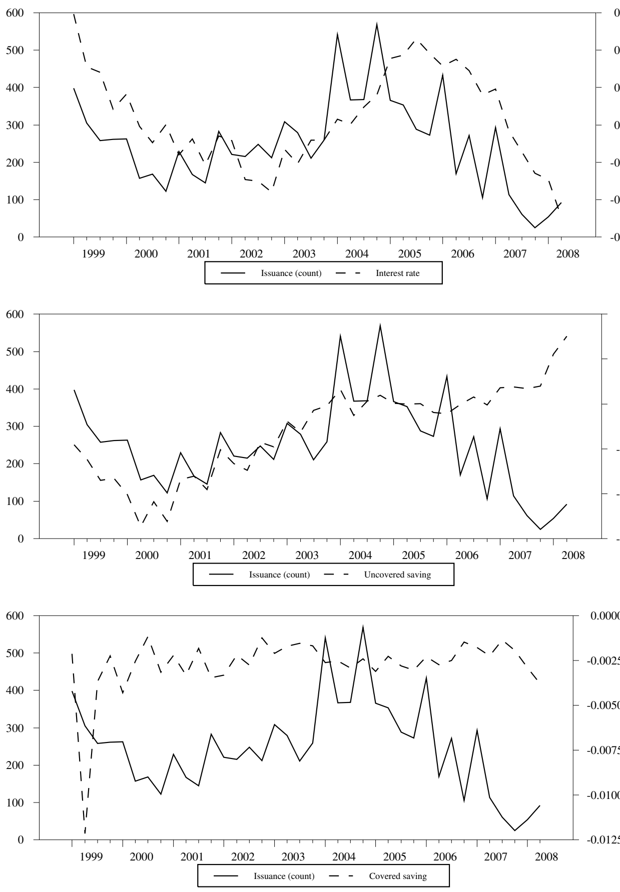

Notes: Top chart shows the number of foreign-currency-denominated bonds carrying a maturity of five years (for maturity details see text) issued in euros versus the interest-rate differential (weighted average of other interest rates minus euro interest rates) for the euro, where all interest rates are of a five-year maturity. Centre chart shows the number of foreign-currency-denominated bonds carrying a maturity of five years issued in euros versus uncovered borrowing-cost savings, as defined in the text. Bottom chart shows the number of foreign-currency-denominated bonds carrying a maturity of five years issued in euros versus covered borrowing-cost savings, as defined in the text. Left scale corresponds to number of bonds issued. Right scale is decimal scale for interest-rate differential, uncovered cost savings and covered cost savings.

3.6.2 Currency-share-model results

Table 3.6 reports coefficient estimates, standard errors (in brackets) and goodness-of-fit measures for panel estimation of the currency share of issuance of foreign-currency-denominated bonds for the five sample currencies. More precisely, the dependent variable is number of bonds issued in currency i as a *share* of number of bonds issued in all issuance currencies. The results are broadly consistent with those presented in Section 3.6.1. Uncovered cost savings, ε^u, play an important role in the choice of issuance currency for bonds of all maturities: the total share of the number of bonds issued in currency i tends to increase in tandem with an increase in the magnitude of uncovered borrowing-cost savings associated with currency i. For two-year-maturity bonds, a 50-basis-point increase in uncovered borrowing-cost savings is associated with in an increase in currency share of issuance of more than 2 percentage points.[31] For five-year and ten-year-maturity bonds, a 50-basis-point increase in uncovered borrowing-cost savings is associated with an increase in currency share of around 0.8 percentage points.

Estimates of the two component parts of ε^u, namely the interest-rate differential, $(\overline{r} - r_i)$, and the expected appreciation of the issuance currency, $(sn_i^e - sn_i)$, again, indicate that the nominal interest-rate differential is the biggest factor influencing the statistical significance of uncovered borrowing-cost savings.[32] The expected change in the value of the issuance currency appears to play no role in the choice of issuance currency for bonds across all maturities. Similarly, covered borrowing-cost savings do not appear to exert an economically important influence on currency choice during the sample period.

Relative economic activity, as in the count model, is found to be a significant driver of currency choice for issuance of bonds of all maturities. For bonds with a maturity of roughly two years, a one percentage point rise in real output in the issuance-currency region relative to all other issuance-currency countries is associated with a rise in currency share of issuance of around three percentage points. For bonds with longer maturities, the influence of relative economic activity is less strong, but still significant. The main difference with results from the count model in Section 3.6.1 is the statistical significance of the coefficient associated with relative financial depth, *liq*. Financial depth exerts a small but significant influence on currency choice among issuance currencies. A one percentage point increase in relative financial depth (total capitalisation of both domestic plus foreign announced issues denominated in issuance currency i)

[31]The average absolute quarterly change in currency share of issuance during the sample period (for bonds that fall into the two-year-maturity bracket) is 3 percentage points. Recall also that the average absolute change in ε_{it}^u during the sample period is 25 basis points.

[32]In tests of parameter equality, unreported, we are unable to reject the null hypothesis of equality of coefficients for ε^u and $(\overline{r} - r_i)$.

Table 3.6: Fixed effects Prais-Winsten estimation

Panel A. Currency share of foreign-currency-denominated bonds, two-year maturity					
	(1)	(2)	(3)	(4)	(5)
ε^c	-0.316		3.903		
	(2.19)		(2.53)		
ε^u		4.370**	5.094**		
		(0.6)	(0.75)		
$(\overline{r} - r_i)$				4.354**	4.266**
				(0.61)	(0.58)
$(sn_i^e - sn_i)$					9.375
					(27.93)
$rgdp$	3.204**	3.438**	3.294**	3.448**	3.466**
	(0.33)	(0.30)	(0.32)	(0.3)	(0.31)
$liq(t-1)$	0.428**	0.412**	0.387**	0.416**	0.424**
	(0.11)	(0.11)	(0.10)	(0.11)	(0.11)
$dinv$	0.001	-0.009	-0.011	-0.008	-0.008
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ma	0.654**	0.618**	0.568**	0.620**	0.625**
	(0.19)	(0.18)	(0.17)	(0.18)	(0.18)
Adj. R^2	0.872	0.892	0.894	0.892	0.892
RMSE	9.611	9.157	9.112	9.153	9.175
Panel B. Currency share of foreign-currency-denominated bonds, five-year maturity					
ε^c	-1.201		-0.044		
	(1.98)		(2.30)		
ε^u		1.405**	1.397*		
		(0.43)	(0.59)		
$(\overline{r} - r_i)$				1.331**	1.652**
				(0.44)	(0.41)
$(sn_i^e - sn_i)$					-34.068
					(22.08)
$rgdp$	2.216**	2.239**	2.240**	2.238**	2.172**
	(0.26)	(0.24)	(0.25)	(0.24)	(0.24)
$liq(t-1)$	0.257**	0.246**	0.246**	0.247**	0.217**
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
$dinv$	0.005	0.001	0.001	0.001	0.000
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ma	0.097	0.072	0.073	0.074	0.055
	(0.11)	(0.11)	(0.10)	(0.11)	(0.11)
Adj. R^2	0.925	0.942	0.943	0.942	0.944
RMSE	6.530	6.473	6.491	6.479	6.421
Panel C. Currency share of foreign-currency-denominated bonds, ten-year maturity					
ε^c	0.523		2.082		
	(1.69)		(1.88)		
ε^u		1.496**	1.882**		
		(0.37)	(0.44)		
$(\overline{r} - r_i)$				1.508**	1.391**
				(0.37)	(0.36)
$(sn_i^e - sn_i)$					12.425
					(17.78)
$rgdp$	1.788**	1.899**	1.822**	1.903**	1.927**
	(0.21)	(0.19)	(0.2)	(0.19)	(0.2)
$liq(t-1)$	0.236**	0.234**	0.221**	0.235**	0.246**
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
$dinv$	0.012	0.009	0.008	0.009	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ma	0.049	0.044	0.017	0.045	0.051
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
Adj. R^2	0.955	0.968	0.968	0.968	0.968
RMSE	5.378	5.287	5.267	5.284	5.287

Notes: Fixed-effects panel estimation with panel-corrected standard errors, corrected for heteroscedasticity and contemporaneous correlation across panels (ie, across issuance currencies). Accommodation for first-order autocorrelation (common to all panels) where present. Dependent variable is number of foreign-currency-denominated bonds issued in currency i at time t as a share of all foreign-currency-denominated bonds issued (expressed in percentage points). Regressions include fixed effects and panel-specific time trends. All explanatory variables are measured at the beginning of the quarter. The sample period is from 1999 to the second quarter of 2008. Standard errors are in parenthases. (**) and (*) denote significance at the 1% and 5% levels respectively.

corresponds to a rise in currency share of issuance of around 0.3 percentage points. Neither relative direct investment nor relative share of cross border mergers and acquisitions influence the choice of issuance currency for bonds of all maturities.

Finally, it is important to note that all of these findings are consistent with this study's underlying assumption, outlined in Section 3.1, that it is more appropriate to measure the *number* of bonds issued rather than the *value* of bonds issued when assessing the responsiveness of issuance to changes in covered and uncovered borrowing-cost savings. When Prais-Winsten panel regressions are estimated with the dependent variable (in Eqn.(3.28)) defined in terms of currency share of issuance *value*, the explanatory power of the key variables in this study drops significantly. Adjusted R-square statistics for each regression are, on average, 20 percentage points lower than those reported in Panels A, B and C in Table 3.6. In addition, both ε^u and ε^c are found to be insignificant in general as determinants of choice among issuance currencies. Relative economic activity is the only variable found to be consistently significant across all specifications.

Financial versus non-financial issuers

One common message from the count-model estimates in Section 3.6.1 and the currency-share estimates, in Section 3.6.2, is that uncovered borrowing costs do play an important role in currency choice for bond issuance. If there is no omitted-variable bias then a significant number of issuers must be responding positively to signals indicating cheaper uncovered borrowing costs in given issuance currencies. The purpose of this section is to examine whether this result is sensitive to the distinction between the type of issuer, in particular differentiating between financial and non-financial issuers.

Before doing so, it is necessary to discuss the implications of allowing for heterogeneity among issuers. The basic model of this paper is premised on issuer homogeneity—the individual issuer is assumed to be representative of the population of issuers as a whole. Relaxing this assumption runs contrary to the underlying model of utility-maximising choice. The best way to introduce issuer heterogeneity would be to construct a new panel count model of currency choice that allows explicitly for issuer heterogeneity. This is beyond the scope of this paper. As far as the author is aware there is no study, to date, that has succeeded in incorporating agent heterogeneity into a panel count model. It has, of course, been achieved in mixed (also known as heterogeneous) multinomial logit models of choice. But agent heterogeneity has not been accounted for in count models of choice. This subject is left for future research. Here, heterogeneity is accounted for in a manner that is based on nothing more than expediency. The approach is discussed below.

In terms of number of bonds issued, financial corporations (ie, investment banks, commercial banks, credit institutions and international banks) dominate global issuance of foreign-currency-denominated bonds. Table 3.7 shows the extent to which financial issuers dominate issuance in all major currencies, in particular, in the issuance of shorter-maturity bonds. For instance, financial issuers account for 97% of all bonds issued in euros with an average maturity of two years.

If uncovered cost savings are an important influence on the issuance decision, then it is conceivable that, of all potential issuers, financial corporations will be most responsive to these cost savings because, firstly, they have a greater speculative motive, and secondly, they have the market knowledge necessary to exploit such savings. Meanwhile, the empirical literature shows that non-financial issuers are concerned mainly with the need to find a natural hedge when issuing foreign-currency-denominated bonds (see Section 3.2).

In order to assess the difference in issuance behaviour, if any, between financial issuers and non-financial issuers, the full sample is split according to Standard Industrial Classification codes (SIC codes) so as to separate all those issuers operating in the financial sector (coinciding, mostly, with the 6000-7000 SIC classification codes) from the rest. The same Prais-Winsten regressions, as above, are run on the two sample subgroups for each of the three maturity brackets. Table 3.8 reports the results for financial issuers, and Table 3.9 for non-financial issuers.

The results suggest that financial issuers do, indeed, respond more strongly than non-financial issuers to uncovered borrowing-cost savings. For longer maturities (beyond two years), coefficient estimates for ε^u are larger and more significant for financial issuers. Coefficient estimates suggest that financial issuers are most responsive when the bonds they are issuing carry maturities of roughly five years in length. For five-year-maturity bonds, a 50-basis-point increase in uncovered borrowing-cost savings is associated with an increase in currency share of around 1.8 percentage points. For short-maturity bonds, uncovered borrowing-cost savings are a statistically significant driver of issuance for financial issuers but not for nonfinancial issuers.

Coefficients associated with control variables accounting for the natural hedge (*rgdp*, *liq*, *dinv* and *ma*) are in general consistent with estimates returned for the full sample in Section 3.6.2. Relative share of economic activity is in general important for both financial and nonfinancial issuers. However, relative share of direct investment and relative share of mergers and acquisitions exert an unexpected influence on issuance (ie, negative rather than positive) in a few specifications. Overall, the role of uncovered cost savings remains a consistent feature in the issuance of foreign-currency-denominated bonds of all maturities.

3.7 Conclusions

This study examines the determinants of currency choice in the issuance of international bonds, focussing on the presence of opportunistic behaviour by bond issuers in response to deviations from covered and uncovered interest-rate parity. Count-data techniques are used to study the number of bonds issued across five major currencies during the period 1999 to 2008. In a robustness check, this study also examines the number of bonds issued in each issuance currency as a share of total number of bonds issued in all currencies. Results are robust across all specifications.

The main finding is that the scope for uncovered borrowing-cost savings, defined as deviations from uncovered interest-rate parity, exerts a significant influence on the choice of issuance currency. These uncovered borrowing-costs savings are assessed in terms of their two main component parts: nominal interest-rate differentials and expected exchange-rate depreciation of the issuance currency. Interest-rate differentials are shown to have a statistically significant impact on currency choice across different empirical specifications, consistent with the findings of other studies. The implication is that issuers prefer to borrow in currencies that offer low nominal interest rates. Meanwhile, issuance does not respond in a consistent manner to expected depreciation of the issuance currency, suggesting that issuers do not, at the aggregate level, attempt to lower borrowing costs by issuing bonds in currencies that are expected to fall in value.

Assessing issuance behaviour by maturity of the bonds being issued reveals that the influence of nominal interest-rate differentials is similar for bonds of all maturities—that is, the influence is no stronger for long-maturity bonds than it is for short-maturity bonds. However, the influence is stronger for financial issuers (eg, investment banks, commercial banks and credit institutions), suggesting that, perhaps, financial issuers are driven by a stronger speculative motive than non-financial issuers when choosing their currency of issuance and have greater access to the type of market information that is necessary to exploit such cost-saving opportunities.

This study finds no robust evidence that covered cost savings systematically affect the number of bonds issued in a given issuance currency. Arbitrage opportunities do seem to be present in the swaps markets, but are not taken up by bond issuers. It is possible that the frequency of our dataset—quarterly data—may introduce a measurement error that impairs a proper assessment of the impact of this variable.

Overall, our findings offer a useful contribution to the understanding of currency choice in the issuance of foreign-currency-denominated bonds by highlighting the importance of uncovered borrowing-cost savings and nominal interest rates. Furthermore, in as much as the issuance of foreign-currency-denominated bonds affects the relative international standing of world currencies, these findings suggest that monetary policy, through its influence on nominal interest rates, has a greater impact on the internationalisation of currencies than has been previously accounted for.

Table 3.7: Financial issuers of international bonds and notes, 1999-2008*

	Number of offerings by financial entities as share (per cent) of all offerings		
	2yr maturity	5yr maturity	10yr maturity
US Dollar	85.2	73.7	72.3
Euro	97.1	71.0	53.6
Yen	60.3	51.3	85.7
UK Pound	81.5	79.8	64.0
Swiss Francs	88.2	70.8	76.6

Notes: Foreign-currency-denominated bonds of specified maturity issued by financial entities during 1999-2008 as a share of total foreign-currency-denominated bonds issued in selected currencies. Maturity here refers to maturity "brackets", as described in the text. Securities with maturities of less than one year are excluded. Foreign-currency-denominated bonds are defined as those bonds issued in a currency other than the currency of the country in which the borrower resides. Sample includes only fixed-interest-rate securities. (*) Data for 2008 is for the first half of 2008. Source is Bondware.

Table 3.8: Fixed effects Prais-Winsten estimation: Financial issuers

Panel A. Currency share of foreign-currency-denominated bonds, two-year maturity					
	(1)	(2)	(3)	(4)	(5)
ε^c	0.403		1.934		
	(3.16)		(3.89)		
ε^u		1.489*	1.848		
		(0.68)	(1.27)		
$(\overline{r} - r_i)$				1.345*	2.013*
				(0.61)	(0.82)
$(sn_i^e - sn_i)$					-70.949
					(46.25)
$rgdp$	1.819**	1.923**	1.852**	1.919**	1.782**
	(0.50)	(0.46)	(0.49)	(0.46)	(0.48)
$liq(t-1)$	0.176	0.173	0.161	0.175	0.112
	(0.17)	(0.17)	(0.16)	(0.17)	(0.17)
$dinv$	-0.051	-0.054	-0.056	-0.054	-0.058
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
ma	0.872**	0.866**	0.841**	0.867**	0.829**
	(0.27)	(0.28)	(0.27)	(0.28)	(0.28)
Adj. R^2	0.776	0.782	0.781	0.782	0.785
RMSE	13.964	13.919	13.947	13.925	13.811
Panel B. Currency share of foreign-currency-denominated bonds, five-year maturity					
ε^c	-0.302		3.104		
	(2.67)		(3.13)		
ε^u		3.536**	4.112**		
		(0.72)	(0.92)		
$(\overline{r} - r_i)$				3.522**	3.457**
				(0.73)	(0.70)
$(sn_i^e - sn_i)$					6.822
					(34.56)
$rgdp$	1.213**	1.401**	1.286**	1.409**	1.422**
	(0.38)	(0.35)	(0.37)	(0.35)	(0.36)
$liq(t-1)$	0.268*	0.255*	0.235	0.258*	0.264*
	(0.13)	(0.13)	(0.12)	(0.13)	(0.13)
$dinv$	-0.040	-0.048*	-0.050**	-0.048*	-0.047*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ma	0.353	0.323	0.284	0.325	0.329
	(0.19)	(0.19)	(0.18)	(0.19)	(0.19)
Adj. R^2	0.827	0.837	0.837	0.837	0.836
RMSE	10.871	10.612	10.603	10.609	10.638
Panel C. Currency share of foreign-currency-denominated bonds, ten-year maturity					
ε^c	-2.617		-2.601		
	(1.39)		(1.40)		
ε^u		1.940*	1.957*		
		(0.94)	(0.89)		
$(\overline{r} - r_i)$				1.868*	2.145*
				(0.94)	(0.95)
$(sn_i^e - sn_i)$					20.035
					(20.33)
$rgdp$	-0.158	-0.254	-0.158	-0.251	-0.212
	(0.16)	(0.16)	(0.16)	(0.16)	(0.16)
$liq(t-1)$	0.086	0.069	0.085	0.069	0.087
	(0.10)	(0.10)	(0.10)	(0.09)	(0.10)
$dinv$	0.003	0.001	0.003	0.001	0.002
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ma	0.085*	0.052	0.085*	0.052	0.063*
	(0.04)	(0.03)	(0.04)	(0.03)	(0.03)
Adj. R^2	0.856	0.861	0.861	0.861	0.860
RMSE	7.697	7.685	7.669	7.684	7.683

Notes: Fixed-effects panel estimation with panel-corrected standard errors, corrected for heteroscedasticity and contemporaneous correlation across panels (ie, across issuance currencies). Accommodation for first-order autocorrelation (common to all panels) where present. Dependent variable is number of foreign-currency-denominated bonds issued in currency i at time t as a share of all foreign-currency-denominated bonds issued. Regressions include fixed effects and panel-specific time trends. All explanatory variables are measured at the beginning of the quarter. The sample period is from 1999 to the second quarter of 2008. Standard errors are in parenthases. (**) and (*) denote significance at the 1% and 5% levels respectively.

Table 3.9: Fixed effects Prais-Winsten estimation: Nonfinancial issuers

Panel A. Currency share of foreign-currency-denominated bonds, two-year maturity					
	(1)	(2)	(3)	(4)	(5)
ε^c	-2.272		-3.076		
	(2.92)		(3.36)		
ε^u		1.456	-0.971		
		(0.80)	(0.96)		
$(\overline{r} - r_i)$				1.432	1.495
				(0.81)	(0.80)
$(sn_i^e - sn_i)$					-6.654
					(38.31)
$rgdp$	1.663**	1.162**	1.646**	1.164**	1.151**
	(0.36)	(0.36)	(0.36)	(0.36)	(0.37)
$liq(t-1)$	-0.13	-0.206	-0.122	-0.204	-0.21
	(0.12)	(0.13)	(0.12)	(0.13)	(0.13)
$dinv$	-0.067**	-0.072**	-0.064**	-0.072**	-0.072**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ma	-0.214	-0.365*	-0.198	-0.364*	-0.367*
	(0.14)	(0.15)	(0.14)	(0.15)	(0.15)
Adj. R^2	0.843	0.823	0.842	0.823	0.823
RMSE	10.738	11.377	10.752	11.378	11.409
Panel B. Currency share of foreign-currency-denominated bonds, five-year maturity					
ε^c	-3.704		-3.284		
	(1.97)		(2.29)		
ε^u		1.116*	0.507		
		(0.47)	(0.64)		
$(\overline{r} - r_i)$				1.173*	0.847
				(0.48)	(0.44)
$(sn_i^e - sn_i)$					34.519
					(21.83)
$rgdp$	0.565*	0.453*	0.574*	0.459*	0.525*
	(0.24)	(0.22)	(0.24)	(0.22)	(0.23)
$liq(t-1)$	0.212**	0.187*	0.208**	0.188*	0.218**
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
$dinv$	-0.004	-0.007	-0.005	-0.007	-0.005
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ma	-0.091	-0.142	-0.1	-0.141	-0.123
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
Adj. R^2	0.923	0.931	0.931	0.931	0.932
RMSE	6.979	6.975	6.929	6.971	6.918
Panel C. Currency share of foreign-currency-denominated bonds, ten-year maturity					
ε^c	-3.557		-3.029		
	(2.16)		(2.49)		
ε^u		1.200*	0.638		
		(0.51)	(0.70)		
$(\overline{r} - r_i)$				1.243*	0.986*
				(0.52)	(0.49)
$(sn_i^e - sn_i)$					27.242
					(22.65)
$rgdp$	1.425**	1.325**	1.437**	1.330**	1.383**
	(0.26)	(0.24)	(0.26)	(0.24)	(0.24)
$liq(t-1)$	0.054	0.03	0.049	0.031	0.055
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
$dinv$	0.026	0.022	0.024	0.022	0.024
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
ma	-0.231*	-0.280**	-0.241*	-0.280**	-0.265*
	(0.11)	(0.10)	(0.11)	(0.10)	(0.10)
Adj. R^2	0.931	0.936	0.937	0.936	0.937
RMSE	7.386	7.378	7.346	7.374	7.352

Notes: Fixed-effects panel estimation with panel-corrected standard errors, corrected for heteroscedasticity and contemporaneous correlation across panels (ie, across issuance currencies). Accommodation for first-order autocorrelation (common to all panels) where present. Dependent variable is number of foreign-currency-denominated bonds issued in currency i at time t as a share of all foreign-currency-denominated bonds issued. Regressions include fixed effects and panel-specific time trends. All explanatory variables are measured at the beginning of the quarter. The sample period is from 1999 to the second quarter of 2008. Standard errors are in parenthases. (**) and (*) denote significance at the 1% and 5% levels respectively.

Chapter 4

Currency dynamics and hedging

This chapter offers a study of the association between movements in the price of gold and the US dollar using a model of dynamic conditional correlations covering 23 years of weekly data for 16 major dollar-paired exchange rates. The aim is to address a practical investment question: Does gold act as a hedge against the US dollar, as a safe haven, or neither? Key findings are as follows. (i) During the past 23 years gold has behaved as a hedge against the US dollar. (ii) There is no evidence to suggest that gold has acted as a consistent and effective safe haven. (iii) In recent years gold has become increasingly correlated (negatively) with the US dollar, more so than at any point during the past two and a half decades, suggesting that gold has, during the recent financial turmoil, acted as a particularly effective hedge against currency risk associated with the movements in the US dollar.

4.1 Introduction

For many years gold as a tradable financial asset has had a reputation as a safe haven from market turbulence. Market reports often refer to gold as a *safe-haven asset*. But very few academic studies have addressed the role of gold as a safe-haven asset and even fewer have examined gold's safe-haven status with respect, specifically, to currency movements. Further, to date, the work that has addressed this issue suffers from shortcomings that offer scope for improvement. This paper examines gold's ability to act as a financial safe haven and improves on other work by addressing correlation rather than dependence, allowing for system feedback and by focusing on the link between changes in the price of gold and the US dollar. Specifically, this paper asks, Does gold act as a safe haven against the US dollar, as a hedge, or neither? Movements in the price of gold and the US dollar are analysed using a model of dynamic conditional correlations covering 23 years of weekly data for 16 major US dollar-paired exchange rates.

Studies relevant to this paper are few. Perhaps most relevant is the work of Capie et al. (2005) which assesses the role of gold as a *hedge* against the US dollar by estimating elasticities for a model of the responsiveness of gold to changes in the exchange rate. Capie et al. (2005) find that gold has in the past acted as an effective hedge. But their approach takes the form of a single-equation model in which the independent variable, the exchange rate, is assumed to be unaffected by the time path of the dependent variable, the price of gold. That is, the authors assume no feedback.[1] Improving over the work of Capie et al. (2005), this paper focuses on correlation, employing a dynamic model of conditional correlations in which all variables are treated symmetrically.[2]

Baur and Lucey (2006) address the specific question of gold's role as a *safe-haven asset*. They find evidence in support of gold providing a haven from losses incurred in the bond and stock markets. However, their approach includes generated regressors, neglects interactions with the currency market and, like Capie et al. (2005), permits no explicit role for feedback in its model of returns. The work of Baur and McDermott (2010), similarly, neglects feedback in its principal regression model even after allowing for it in a number of constructed parameters.[3]

A handful of studies investigate the financial concept of a *safe-haven asset* without reference to gold. Ranaldo and Soderlind (2010) and Kaul and Sapp (2006) examine safe-haven currencies while Upper (2000) examines German government bonds as safe-haven investments. Other studies look at the wider financial properties of gold without focusing on its role as a safe haven.[4] None of these examine gold's ability to act as a safe haven with respect to the US dollar.

In addressing the question of gold's use as a safe haven from currency risk, it is instructive to ask, is currency risk large enough, in general, to elicit the

[1] Chen and Rogoff (2003), Clements and Fry (2008) and Swift (2004) among others highlight the importance of allowing for feedback and co-determination in the analysis of currency and commodity markets.

[2] In modelling exchange rates and the price of gold simultaneously this study adopts, implicitly, the idea that exchange rates can be considered as asset prices (Engel and West, 2005) and, therefore, directly comparable with prices for commodities such as gold. Indeed, there is a large literature studying the link between exchange rates and commodities—see for instance Chen and Rogoff (2003)—and this chapter draws directly from that literature.

[3] Baur and McDermott (2010) and Baur and Lucey (2006) also construct GARCH models that include dummy variables, causing standard inference on their estimated coefficients to be potentially invalid (Doornik and Ooms, 2003).

[4] See for instance Cheung and Lai (1993), Faugere and Erlach (2004), Sherman (1982), Sjaastad (2008) and Worthington and Pahlavani (2007).

pursuit of safe-haven assets? Existing research suggests that it is. Santis and Gerard (1998), for instance, show that currency risk is economically significant and represents a large fraction of the total risk faced when investing overseas. Andersen et al. (2007) show that exchange-rate volatility outstrips bond-market volatility in their sample of futures prices for US, British and German markets while Hau and Rey (2006) find that the ratio of exchange-rate volatility to equity return volatility is close, but less than one, in line with their equilibrium model of exchange rates, stock prices and capital flows.

There are of course various ways to hedge against currency risk. Hedging mechanisms can be financial or operational.[5] Given the options available, why might gold be used as a hedge or safe haven? The reasons are many. Gold, as a financial asset, is liquid, available, priced in US dollars and can be traded on a futures market. Further, while gold as a hedge cannot be designed for purpose in the same way as foreign-exchange derivatives, even bespoke hedging techniques are less than perfect in their effectiveness (Huffman and Makar, 2004). Gold, as a *natural* hedge or haven, may be useful if effective. It is the effectiveness of gold as a hedge and safe haven that this paper aims to examine.

Any discussion of investment safe havens and hedges requires clear definitions. What, exactly, is a haven? What is a hedge? This study adopts the definitional approach of Baur and Lucey (2006) and Kaul and Sapp (2006): If an investor holds a given asset, γ, then a haven is defined as any other asset that does not co-move with γ in times of stress. That is, a haven is uncorrelated or correlated negatively with γ if γ experiences sharp changes in value. A hedge, meanwhile, is an asset that is uncorrelated or correlated negatively with γ not just in times of stress, but *on average*. The definitional difference between a hedge and a haven is subtle but important: an asset that functions as a haven is uncorrelated or correlated negatively with γ in times of stress only, and not necessarily on average.

The contribution of this study to the existing literature is two-fold. First, this study assesses the role of gold as both a hedge and a safe haven with respect to the US dollar. While other work has investigated the role of gold as a hedge and a haven for bonds and equities, no study has tackled the same subject with a specific focus on exchange rates.[6] Second, using the correlation modelling techniques of Engle (2002), this study offers an empirical analysis of a 17-variable system of returns, considering a larger number of currencies than Capie et al. (2005).

This study's key findings are as follows. (i) During the past 23 years gold

[5] See Allayannis and Ofek (2001), Allayannis et al. (2001), Elliot et al. (2003) and Habib and Joy (2010).

[6] For a discussion of gold's relationship with bonds and equities see Baur and Lucey (2006).

has behaved as a hedge against the US dollar—that is, gold-price returns have, on average, been correlated negatively with US dollar returns. (ii) There is no evidence to suggest that gold has acted as a consistent and effective safe haven. (iii) In recent years gold has become an increasingly effective hedge against the US dollar, with conditional correlations more negative now than they have been at any point during the past two and a half decades.

The rest of this study is structured as follows. Section 4.2 presents some background discussion of correlation models. Section 4.3 gives an empirical outline of models of constant conditional correlations and dynamic conditional correlations. Section 4.4 discusses the dataset, Section 4.5 presents results of the empirical analysis, and finally, Section 4.6 offers some conclusions.

4.2 Why Correlation Models?

Correlation models have many important financial applications. Accurate estimates of the correlations of asset returns and, in turn, their volatilities (second moments), are required for asset pricing, hedging, capital allocation and risk management.[7]

Correlation features heavily in the models of risk and return first introduced by Markowitz (1952) and now used widely in financial markets. Both the Capital Asset Pricing Model and Arbitrage Pricing Theory use correlation as a measure of dependence between different assets in order to estimate an optimal portfolio. However, the estimation and forecasting of correlation is not always straightforward. It is complicated by a number of factors.

First, correlation between asset returns is not directly observable: daily correlation today is not observable because there is only one observation in a trading day. The conventional approach is to estimate the correlation matrix using realised data on daily asset returns (Andersen and Bollerslev, 1998). More recent techniques use higher frequency data to estimate realised correlations (Andersen et al., 2001). Second, for many years correlations in finance were assumed to be constant and were modelled as such. But during the past 20 years studies have shown that correlations are time-varying.[8] Further, they often increase during periods of high volatility and market stress.[9] Third, as the number of assets increases, so the estimation of the correlation matrix becomes increasingly difficult—the curse of dimensionality becomes a major obstacle.

[7] For a recent discussion of asset pricing and capital allocation, see Cochrane (2005). See Tsay (2005) for an introduction to correlation models in risk management. For the use of correlation models in the estimation of hedge ratios, see, for instance, Bos and Gould (2007) and Lien et al. (2002).

[8] For empirical evidence see, for instance, Furstenberg and Jeon (1989) and Koch and Koch (1991).

[9] See, among others, Ang and Chen (2002) and Ramchand and Susmel (1998).

Given these difficulties, the modelling of correlations has branched into a number of competing fields of research, all drawing on the seminal research into financial volatility undertaken by Engle (1982) and Bollerslev (1986). Alternative approaches include stochastic volatility models, implied volatility models (with information extracted from options prices), and models that accommodate generalised autoregressive conditional heteroscedasticity (GARCH effects).[10] In recent years multivariate GARCH models have been developed that allow for time-varying correlations (Engle, 2002; Lien et al., 2002). In these approaches, conditional variances are modelled as univariate GARCH processes—that is, past innovations and variances of one variable are precluded from affecting the conditional variances of other variables. This is a limitation. But in their favour, these models cope well with the curse of dimensionality and can be augmented to accommodate a variety of empirical dynamic phenomena.[11]

In this study a model of dynamic conditional correlations is used to examine gold-price returns and exchange-rate returns. The next section, Section 4.3, offers an econometric description of the model. It describes the dynamic model and outlines its origins in the model of constant conditional correlations first proposed by Bollerslev (1990).

4.3 Empirical Methodology

Correlation models attract as much attention from financial-market practitioners as they do from academics. Models vary from the simple (eg, rolling historical correlations) to the complex (based on varieties of stochastic volatility or on models of multivariate generalised autoregressive conditional heteroscedasticity, known as GARCH models). Correlation models have been popularised by, among others, Bollerslev (1990), Kroner and Claessens (1991), Engle et al. (1990) and Ding and Engle (2001).

At the root of every correlation model is a structure of conditional correlations. The conditional correlations between two random variables, r_1 and r_2 (here, for example, exchange-rate returns), both of which have a mean of zero, are defined as

$$\rho_{r_1 r_2,t} = \frac{E_{t-1}(r_{1t} r_{2t})}{\sqrt{E_{t-1}(r_{1t}^2) E_{t-1}(r_{2t}^2)}} \tag{4.1}$$

All conditional correlations, $\rho_{r_1 r_2,t}$, lie in the interval $[-1, +1]$ and are based on information from the previous period $(t-1)$. One important point to note is the nature of the link between conditional correlations and conditional covariances.

[10] See Campbell et al. (1997).

[11] For asymmetric dynamics see Cappiello et al. (2006), for smooth transition dynamics see Silvennoinen and Tersvirta (2009).

To see this link, first express each returns series, r_i, as the product of the conditional standard deviation ($\sqrt{h_{it}}$) and the standardised error term (ϵ_{it}), such that

$$r_{it} = \sqrt{h_{it}}\epsilon_{it} \qquad h_{it} = E_{t-1}(r_{it}^2) \qquad i = 1, 2 \tag{4.2}$$

where ϵ is a standardised error term that has mean zero and variance one for each series. Substituting Eqn.(4.2) into Eqn.(4.1) gives

$$\rho_{r_1 r_2, t} = \frac{E_{t-1}(\epsilon_{1t}\epsilon_{2t})}{\sqrt{E_{t-1}(\epsilon_{1t}^2)E_{t-1}(\epsilon_{2t}^2)}} = E_{t-1}(\epsilon_{1t}\epsilon_{2t}) \tag{4.3}$$

From Eqn.(4.3) it can be seen that the conditional correlation, $\rho_{r_1 r_2, t}$, is equal to the conditional covariance between the standardised error terms, $E_{t-1}(\epsilon_{1t}\epsilon_{2t})$.

As will be discussed later, there are many alternative approaches to estimating correlation models for multivariate systems, but one of the most popular approaches, and the one pursued here, is to assume that the variables under analysis exhibit GARCH effects.

GARCH models, despite being introduced more than two decades ago by Bollerslev (1986) as an extension to the ARCH model of Engle (1982), continue to provide an important research tool for modelling the dynamics of asset prices and, in particular, the phenomenon of volatility clustering.[12] GARCH models draw on the idea that the volatility clustering of asset prices can be captured by allowing the variance of ϵ_t, the error term, to depend upon its history. In particular, Engle (1982) proposes that the variance of the error term at time t depends upon the squared error terms from the previous period. Since many financial variables, not least exchange rates, are interrelated and affected by the same market information, it makes sense to extend univariate GARCH models to their multivariate equivalents in order to capture common dynamics. However, extending a univariate GARCH system to its multivariate counterpart within a correlation model is not easy. The number of parameters to be estimated is large, and the construction of the conditional variance-covariance matrix is complicated (Engle and Kroner, 1995).

In an effort to side-step these difficulties in the pursuit of a well-specified correlation model, Bollerslev (1990) introduces a bivariate GARCH system that assumes all conditional correlations are constant. An overview of this model of constant correlations is presented in the next section, Section 4.3.1. Section 4.3.2, meanwhile, presents a model in which the conditional correlations are not constant, but dynamic.

[12] Bauwens et al. (2006) provide a good recent survey of the application of GARCH models.

4.3.1 Constant conditional correlations

This section offers an overview of a multivariate GARCH model of correlations proposed by Bollerslev (1990) wherein conditional correlations between the variables of interest are constant over time. By presenting this model, and highlighting its drawbacks, the intention is to provide a good starting point for the discussion of a multivariate GARCH model of conditional correlations that are not constant over time, but are time-varying.

Bollerslev (1990) measures the closeness of association between movements in exchange rates by constructing a multivariate time-series model that allows for time-varying conditional variances and covariances but permits only constant, not time-varying, conditional correlations.[13] Constant conditional correlations do, of course, imply that all conditional correlations between variables are fixed over time. This is a restrictive assumption. Nonetheless, a brief discussion, here, of the constant conditional correlation model will act as a useful introduction to the dynamic conditional correlation model presented in Section 4.3.2.[14]

The constant conditional correlation model, in its construction as a tool to analyse exchange rates, builds on the observation that the short-run dynamics of exchange rates are contaminated with heteroscedasticity.[15] Therefore, any model that seeks to measure the closeness of association between the movements of a number of exchange rates should, argues Bollerslev (1990), take the form of a multivariate time-series model allowing for heteroscedasticity.

To set up the constant conditional correlation model, let r_t denote a $N \times 1$ times series vector (where r_t represents, for instance, a series of exchange-rate returns), with time-varying conditional variance-covariance matrix H_t such that

$$r_t = E(r_t|\psi_{t-1}) + \epsilon_t \tag{4.4}$$

where ψ_{t-1} is the measurable space generated by all the available information up to and including time $t-1$. Meanwhile, the time-varying conditional variance-covariance matrix can be defined as

$$H_t = Var(\epsilon_t|\psi_{t-1}) \tag{4.5}$$

where H_t is positive definite for all t (that is, the characteristic roots of H are positive for all t). Together, Eqn.(4.4) and Eqn.(4.5) form a model of general heteroscedasticity, allowing for both conditional and unconditional heteroscedas-

[13]That is, the variance and covariances of the current error term are *time-varying* functions of the previous period's error terms, whereas the current-period correlations are *constant* functions of the previous period's error terms.

[14]Bollerslev (1990) states that the constant conditional correlation model can be interpreted as an extension of a Seemingly Unrelated Regression model that allows for heteroscedasticity.

[15]See for instance Domowitz and Hakkio (1985) and Bollerslev and Ghysels (1996).

ticity (a model of generalised autoregressive conditional heteroscedasticity, or GARCH model).

Next, let h_{ijt} denote the ij^{th} element in H_t, corresponding, in the case of Bollerslev (1990), to currency i and currency j. Also let r_{it} denote the i^{th} element in r_t and let ϵ_{it} denote the i^{th} element in ϵ_t, the error vector. Then, recalling that a correlation coefficient is given by the covariance of the two random variables standardised by the two standard deviations ($\rho_{xy} = cov(X,Y)/\sqrt{var(X)var(Y)}$), it can be seen that a scale-invariant measure of the degree of co-movement between r_{it} and r_{jt} evaluated at time $t-1$ is given by the conditional correlation

$$\rho_{ijt} = \frac{h_{ijt}}{\sqrt{h_{iit}h_{jjt}}} \tag{4.6}$$

where the correlation coefficient ρ_{ijt} is between -1 and $+1$ for all t and is not affected by the scaling of the variables.

This measure ρ_{ijt} of co-movement will vary through time because H_t varies through time. However, Bollerslev (1990) notes that ρ_{ijt} will be constant over time ($\rho_{ijt} = \rho_{ij}$) if the time-varying conditional covariances h_{ijt} are proportional to the square root of the product of the corresponding two conditional variances such that

$$h_{ijt} = \rho_{ij}\sqrt{h_{iit}h_{jjt}} \qquad j = 1,\ldots,N, \quad i = j+1,\ldots,N. \tag{4.7}$$

Whether or not Eqn.(4.7) holds true is an empirical matter (as it would be for any other parameterisation of the conditional heteroscedasticity). Bollerslev (1990) suggests that Eqn.(4.7) provides an adequate characterisation of the co-movement of many financial series—while other studies, for instance those by Longin and Solnik (1995) and Boyer et al. (1997), suggest that this model does not hold well for all financial data.

Perhaps the most attractive feature of the constant conditional correlation model is its tractability when it comes to estimation and inference. To illustrate this tractability it is useful to re-express the conditional variances, h_{iit}, as

$$h_{iit} \equiv \omega_i \sigma_{it}^2 \qquad i = 1,\ldots,N \tag{4.8}$$

where ω_i is a positive, time-invariant scalar and where σ_{it}^2 is greater than zero for all t. Note that the decomposition in Eqn.(4.8) is only unique up to scale. Given Eqn.(4.7) and Eqn.(4.8), the full conditional variance-covariance matrix, H_t, may be partitioned as

$$H_t = D_t \Gamma D_t \tag{4.9}$$

where D_t denotes the $N \times N$ stochastic diagonal matrix with elements $\sigma_{1t}, \ldots, \sigma_{Nt}$ and where Γ is an $N \times N$ time-invariant matrix with typical element $\rho_{ij}\sqrt{\omega_i\omega_j}$. That is, Γ is a correlation matrix containing the conditional correlations. As such,

$$\Gamma = \begin{pmatrix} \rho_{11}\omega_1 & \rho_{12}\sqrt{\omega_1\omega_2} & \cdots & \rho_{1N}\sqrt{\omega_1\omega_N} \\ \rho_{21}\sqrt{\omega_2\omega_1} & \rho_{22}\omega_2 & \cdots & \rho_{2N}\sqrt{\omega_2\omega_N} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{N1}\sqrt{\omega_N\omega_1} & \rho_{N2}\sqrt{\omega_N\omega_2} & \cdots & \rho_{NN}\omega_N \end{pmatrix} \tag{4.10}$$

$$D_t = \begin{pmatrix} \sigma_{1t} & 0 & \cdots & 0 \\ 0 & \sigma_{2t} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_{Nt} \end{pmatrix} \tag{4.11}$$

with it following that H_t will be positive definite for all t if and only if each of the N conditional variances $(h_{iit} \equiv \omega_i\sigma_{it}^2)$ are well defined and Γ is positive definite (that is, if Γ is positive definite then so too will be H_t because D_t is nothing more than a time-varying scaling matrix). These conditions, argues Bollerslev (1990), are easy to impose and verify compared with the conditions implied by other parameterisations of the time-varying conditional variance-covariance matrix H_t.[16]

Estimation: constant conditional correlations

Maximum-likelihood estimation of the constant conditional correlations model requires the assumption of conditional normality (that is, requires the assumption that, conditional on information up to and including period $t-1$, the error term ϵ_t is distributed normally with mean zero and variance H_t). Recall, next, that the log-likelihood function for the simple linear regression model can be stated as

$$\log L(\theta) = -\frac{N}{2}\log(2\pi\sigma^2) - \frac{1}{2}\sum_{i=1}^{N}\frac{\epsilon_i^2}{\sigma^2} \tag{4.12}$$

where θ is a K-dimensional vector of unknown parameters. Assuming conditional normality, the log-likelihood function for the general heteroscedastic model represented by Eqn.(4.4) and Eqn.(4.5) can be expressed as

$$\log L(\theta) = -\frac{TN}{2}\log(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\log|H_t| - \frac{1}{2}\sum_{t=1}^{T}\epsilon_t' H_t^{-1}\epsilon_t \tag{4.13}$$

[16]See Baba et al. (1989) for an overview of similar conditions in the context of a multivariate linear GARCH(p,q) model.

where θ denotes all the unknown parameters in ϵ_t and H_t.

Bollerslev (1990) states that under standard regularity conditions, the maximum likelihood estimate for θ is asymptotically normal and it is possible to employ traditional inference procedures. However, evaluation of the likelihood function in Eqn.(4.13) requires the inversion of one $N \times N$ matrix for each time period t (that is, for evaluation it is necessary to find H_t^{-1}). As a result, the maximisation of $L(\theta)$ by iterative methods can be costly, in terms of computational effort, even when T and N are not large. Computational effort is reduced dramatically, however, by the assumption in Eqn.(4.7) that the time-varying conditional covariances are proportional to the square root of the corresponding two conditional variances.

Substituting Eqn.(4.9) into Eqn.(4.13), gives the amended log-likelihood function

$$\log L(\theta) = -\frac{TN}{2}\log(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\log|D_t\Gamma D_t| - \frac{1}{2}\sum_{t=1}^{T}\epsilon_t'(D_t\Gamma D_t)^{-1}\epsilon_t \qquad (4.14)$$

$$= -\frac{TN}{2}\log(2\pi) - \frac{T}{2}\log|\Gamma| - \sum_{t=1}^{T}\log|D_t| - \frac{1}{2}\sum_{t=1}^{T}\tilde{\epsilon}_t'\Gamma^{-1}\tilde{\epsilon}_t \qquad (4.15)$$

where $\tilde{\epsilon}_t = D_t^{-1}\epsilon_t$ denotes the $N \times 1$ vector of standardised residuals. Bollerslev (1990) notes that, apart from the term $-\sum_{t=1}^{T}\log|D_t|$, a Jacobian term that arises as a result of the transformation from ϵ_t to $\tilde{\epsilon}_t$, the likelihood function in Eqn.(4.15) is easier to evaluate, requiring only one inversion of an $N \times N$ matrix.[17] Eqn.(4.13) requires T inversions.

The log-likelihood function in Eqn.(4.15) is not linear in the parameters. As a result, algebraic maximisation is infeasible and, instead, it is necessary to adopt a numerical search technique.[18]. This increases the complexity of the approach but even so, it is still much easier to evaluate Eqn.(4.15) than Eqn.(4.13).[19]

In the same manner as for seemingly unrelated regressions (SUR), consistent estimates of the variances and covariances (that is, of Γ), conditional on $\tilde{\epsilon}_t$ ($t = 1,\ldots,T$) is given by Γ's sample analogue, $\hat{\Gamma} = T^{-1}\sum_t \tilde{\epsilon}_t\tilde{\epsilon}_t'$, which is nonsingular by construction and hence positive definite. Drawing on the fact that maximum likelihood estimators are invariant under strict monotone transformations (here, the transformation is from ϵ_t to $\tilde{\epsilon}_t$), then the maximum likelihood estimate of

[17] Here, $-\sum_{t=1}^{T}\log|D_t|$ is a *Jacobian term* in the sense that it is a multivariate adjustment factor arising from the transformation from ϵ_t to $\tilde{\epsilon}_t$ (Wilks, 1962).

[18] A numerical search technique would be necessary, also, for maximisation of the general heteroscedastic model in Eqn.(4.13)

[19] Note also that $\log|D_t|$, the natural logarithm of the absolute value of the determinant of D_t, is simply the sum of $\log\sigma_{1t},\ldots,\log\sigma_{Nt}$.

each of the conditional correlations will be given by

$$\hat{\rho}_{ijt} = \frac{\sum_t \tilde{\epsilon}_{it}\tilde{\epsilon}_{jt}}{\sqrt{\sum_t \tilde{\epsilon}_{it}^2 \sum_t \tilde{\epsilon}_{jt}^2}} \tag{4.16}$$

Meanwhile, the parameters in Γ (note that Γ contains $1/2N(N+1)$ parameters) can be concentrated out of the likelihood function.[20] From Eqn.(4.15), concentrating out gives,

$$\begin{aligned}
\log L(\theta) =& -\frac{TN}{2}\log 2\pi - \sum_{t=1}^{T}\log|D_t| - \frac{T}{2}\log\left|\frac{1}{T}\sum_t \tilde{\epsilon}_t\tilde{\epsilon}_t'\right| \\
&- \frac{1}{2}\sum_{t=1}^{T}\tilde{\epsilon}_t'\left(\frac{1}{T}\sum_t \tilde{\epsilon}_t\tilde{\epsilon}_t'\right)^{-1}\tilde{\epsilon}_t \\
=& -\frac{TN}{2}\log 2\pi - \sum_{t=1}^{T}\log|D_t| - \frac{T}{2}\log T - \frac{T}{2}\log\left|\sum_t \tilde{\epsilon}_t\tilde{\epsilon}_t'\right| \\
&- \frac{TN}{2}\sum_{t=1}^{T}\tilde{\epsilon}_t'\left(\sum_t \tilde{\epsilon}_t\tilde{\epsilon}_t'\right)^{-1}\tilde{\epsilon}_t \\
=& -\frac{TN}{2}(1+\log 2\pi - \log T) - \sum_{t=1}^{T}\log|D_t| - \frac{T}{2}\log\left|\sum_{t=1}\tilde{\epsilon}_t\tilde{\epsilon}_t'\right|
\end{aligned} \tag{4.17}$$

Despite the simplified nature of Eqn.(4.17), it must be noted that the information matrix (that is, the variance of the score vector, which is calculated by multiplying the Hessian of the log-likelihood function by -1) between the parameters in D_t and Γ is not block diagonal. As a result, in order to obtain an estimate of the asymptotic variance-covariance matrix (the inverse of the information matrix) by standard maximum-likelihood techniques, we require the derivatives of the full likelihood function in Eqn.(4.15). The maximisation of the log-likelihood function requires iterative methods. Bollerslev (1990) employs the algorithm proposed by Berndt et al. (1974) together with numerical first-order derivatives to approximate $\delta \log L(\theta)/\delta\theta$.

In summary, the constant conditional correlation model is a bivariate generalised autoregressive conditional heteroscedasticity (GARCH) model whereby univariate GARCH processes are estimated for each variable. The correlation matrix is then estimated using the standard, closed form, maximum likelihood correlation estimator (by transforming the residuals using their estimated conditional standard deviations).[21] It is the assumption of constant correlations

[20] By *concentrated out*, we mean that Γ can be expressed as a function of other parameters in the likelihood function.

[21] Estimation of a generalised autoregressive conditional heteroscedasticity (GARCH) model

that makes possible the estimation of large models. This assumption also ensures that the estimator is positive definite, with the only requirements being that each univariate conditional variance is not zero and that the correlation matrix is of full rank.

Bollerslev (1990), studying the movements of the German mark, the French franc, the Italian lira, the Swiss franc and the British pound between 1973 and 1985, claims that the assumption of constant conditional correlations is valid. However, others disagree. Tsui and Yu (1999), for instance, find that there is no constancy of conditional correlations for share prices in two major Chinese stock markets.

Empirical rejection of the model of constant conditional correlations is allied with a number of technical objections, the most prominent of which is the fact that the constant correlation estimator, as proposed by Bollerslev (1990), does not yield consistent standard errors. As an alternative to the constant conditional correlation model, Engle (2002) proposes a model of *dynamic* conditional correlations.

The dynamic model of Engle (2002) is set within a framework that preserves the advantages of the Bollerslev (1990) model (preserves the empirical benefit of estimating fewer parameters than those necessary in rival models such as the BEKK formulation of Engle and Kroner (1995)) and yet is able to incorporate non-constant (ie, dynamic) correlations. Section 4.3.2 outlines in detail the model of dynamic conditional correlations and discusses its estimation.

4.3.2 Dynamic conditional correlations

This section introduces a multivariate GARCH model of dynamic conditional correlations first proposed by Engle (2002), which itself builds upon the constant-conditional-correlation model presented in Section 4.3.1. The advantage of the dynamic conditional correlation model (also known as the DCC-GARCH model) is that it offers a tractable way of modelling, simultaneously, both time-varying conditional volatilities and time-varying conditional correlations.

The DCC-GARCH model can be best understood by recalling, firstly, that the conditional correlation between two random variables r_1 and r_2 (where r_1 and r_2 represent, here, asset-price returns), each with mean zero, can be defined as

$$\rho_{r_1 r_2,t} = \frac{E_{t-1}(r_{1t} r_{2t})}{\sqrt{E_{t-1}(r_{1t}^2) E_{t-1}(r_{2t}^2)}} \tag{4.18}$$

is frequently carried out using numerical optimisation and quasi maximum-likelihood techniques because closed-form estimates of the parameters are often not available. With numerical optimisation, the resulting estimator depends on the implementation, with different optimisation techniques leading to potentially different estimators. A closed-form estimator is, therefore, preferable.

The relation between the conditional correlations, $\rho_{r_i r_j t}$, and the conditional variances, h_{it}, can be clarified by expressing each returns series, r_{it}, as the product of the conditional standard deviation, $\sqrt{h_{it}}$, and the standardised error term, ϵ_{it}, such that

$$r_{it} = \sqrt{h_{it}}\epsilon_{it} \qquad h_{it} = E_{t-1}(r_{it}^2) \qquad i = 1, 2 \tag{4.19}$$

where ϵ is a standardised error term that has mean zero and variance one for each series. Substituting Eqn.(4.19) into Eqn.(4.18) gives

$$\rho_{r_1 r_2, t} = \frac{E_{t-1}(\epsilon_{1t}\epsilon_{2t})}{\sqrt{E_{t-1}(\epsilon_{1t}^2)E_{t-1}(\epsilon_{2t}^2)}} = E_{t-1}(\epsilon_{1t}\epsilon_{2t}) \tag{4.20}$$

From Eqn.(4.20) it can be seen that the conditional correlation, $\rho_{r_1 r_2, t}$, is equal to the conditional covariance between the standardised error terms, $E_{t-1}(\epsilon_{1t}\epsilon_{2t})$.

Defining the conditional variance-covariance matrix of returns as

$$H_t \equiv E_{t-1}(r_t r_t') \tag{4.21}$$

allows us, next, to highlight the key features of the DCC-GARCH model.

Dynamic conditional correlation estimators have a number of identifying characteristics. First, and most obviously, they are designed to capture conditional correlations that are time-varying in nature, and specifically, designed to capture GARCH-like, time-varying correlations. Second, like constant conditional correlation estimators—and unlike other multivariate GARCH approaches such as the parameter-restricted model proposed by Bollerslev et al. (1988)—they can handle systems involving a large number of parameters. This is because the number of parameters to be estimated in the correlation process is independent of the number of series to be correlated. Third, dynamic conditional correlation models retain the flexibility of univariate GARCH models in that univariate GARCH models are estimated for each variable and then, using standardised residuals from this first phase of univariate estimates, a time-varying correlation matrix is estimated. As a result, this *two-stage* estimation process preserves the simple logic of interpretation associated with univariate GARCH models.

In the multivariate GARCH model of *constant* conditional correlations described in Section 4.3.1, the conditional variance-covariance matrix of returns, H_t, can be partitioned as

$$H_t = D_t \Gamma D_t \tag{4.22}$$

where $D_t = \text{diag}\{\sqrt{h_t}\}$ and where Γ is a correlation matrix containing the conditional correlations. That Γ contains the conditional correlations can be seen directly from noting that $\epsilon_t = D_t^{-1} r_t$, which allows us to re-express Eqn.(4.22) as

$$\begin{aligned}\Gamma &= D_t^{-1} H_t D_t^{-1} \\ &= E_{t-1}(\varepsilon_t \varepsilon_t')\end{aligned} \tag{4.23}$$

Eqn.(4.23) tells us that the conditional covariance between the standardised residuals, $E_{t-1}(\varepsilon_t \varepsilon_t')$, is equal to the matrix of conditional correlations, Γ.

In the model of dynamic conditional correlations, as in the model of constant conditional correlations, the elements of D_t are modelled as univariate GARCH processes. That is, all time-varying conditional volatilities are assumed to be represented adequately well by GARCH processes such that

$$h_{it} = \omega_i + \sum_{p=1}^{P_i} \alpha_{ip} r_{i(t-p)}^2 + \sum_{q=1}^{Q_i} \beta_{iq} h_{i(t-q)} \tag{4.24}$$

for $i = 1, 2, \ldots, k$ with the usual GARCH restrictions for non-negativity of variances and stationarity.[22] Lag lengths for P and Q need not be the same. In this way, with all time-varying conditional volatilities modelled as GARCH processes, D_t becomes a time-varying diagonal matrix of standard deviations from univariate GARCH models.

While in Eqn.(4.22) the conditional correlations are assumed to be *constant*, Engle (2002) proposes the DCC-GARCH model in which correlations are *dynamic*. That is,

$$H_t = D_t \Gamma_t D_t \tag{4.25}$$

where Γ_t is, as before, the correlation matrix, but where this correlation matrix is now allowed to vary over time. The conditional variances of Γ_t must be equal to one. Other than this, requirements for the parameterisation of Γ_t are the same as for H_t.

Typical elements of Γ_t will be of the form

$$\rho_{ijt} = \frac{q_{ijt}}{\sqrt{q_{iit} q_{jjt}}} \tag{4.26}$$

with the aim being to define q_{ijt} in such a way as to provide a dynamic correlation structure (provide a parameterisation of Γ_t) that is both useful and tractable.

Engle (2002) notes that perhaps the simplest parameterisation of Γ_t is the

[22] For stationarity, $\sum_{p=1}^{P_i} \alpha_{ip} + \sum_{q=1}^{Q_i} \beta_{iq} < 1$.

exponential smoother, whereby ρ_{ijt} is a geometrically weighted average of the standardised residuals. Specifying the correlation matrix in the form of an exponential smoother gives,

$$\rho_{ijt} = \frac{\sum_{s=1}^{t-1} \lambda^s \epsilon_{i(t-s)} \epsilon_{j(t-s)}}{\sqrt{(\sum_{s=1}^{t-1} \lambda^s \epsilon_{i(t-s)}^2)(\sum_{s=1}^{t-1} \lambda^s \epsilon_{j(t-s)}^2)}} = (\Gamma_t)_{ij} \tag{4.27}$$

or similarly,

$$q_{ijt} = (1-\lambda)(\epsilon_{i(t-1)}\epsilon_{j(t-1)}) + \lambda(q_{ij(t-1)}) \tag{4.28}$$

where $\rho_{ijt} = \frac{q_{ijt}}{\sqrt{q_{iit} q_{jjt}}}$. While the exponential smoother offers simplicity, there are of course other available specifications. Engle (2002) suggests one natural option is to relax the parameter restriction of λ in Eqn.(4.28) and allow q_{ijt} to follow a GARCH(1,1) model, such that

$$q_{ijt} = \bar{\rho}_{ij} + \alpha(\epsilon_{i(t-1)}\epsilon_{j(t-1)} - \bar{\rho}_{ij}) + \beta(q_{ij(t-1)} - \bar{\rho}_{ij}) \tag{4.29}$$

where $\bar{\rho}_{ij}$ is the unconditional correlation between ϵ_{it} and ϵ_{jt}. Rearranging and substituting lags of Eqn.(4.29) recursively into itself, gives

$$\begin{aligned} q_{ijt} &= \bar{\rho}_{ij} + \alpha(\epsilon_{i(t-1)}\epsilon_{j(t-1)} - \bar{\rho}_{ij}) + \beta(q_{ij(t-1)} - \bar{\rho}_{ij}) \\ &= \bar{\rho}_{ij}(1-\alpha-\beta) + \alpha(\epsilon_{i(t-1)}\epsilon_{j(t-1)}) + \beta q_{ij(t-1)} \\ &= \bar{\rho}_{ij}(1-\alpha-\beta)(1+\beta+\beta^2+\ldots) + \alpha \sum_{s=1}^{\infty} \beta^{s-1} \epsilon_{i(t-1)}\epsilon_{j(t-1)} + \beta^{\infty} q_{ij(t-\infty)} \\ &= \bar{\rho}_{ij} \frac{1-(\alpha+\beta)}{1-\beta} + \alpha \sum_{s=1}^{\infty} \beta^{s-1} \epsilon_{i(t-1)}\epsilon_{j(t-1)} \end{aligned} \tag{4.30}$$

Eqn.(4.30) captures the main features of the model of dynamic conditional correlations (the DCC-GARCH model) proposed by Engle (2002). The mean of q_{ijt} will be $\bar{\rho}_{ij}$. That is, $\bar{q}_{it} \cong \bar{\rho}_{ij}$. The mean variance will equal one. Meanwhile, the correlation estimator, $\rho_{ijt} = \frac{q_{ijt}}{\sqrt{q_{iit} q_{jjt}}}$, will be positive definite because the variance-covariance matrix $Q_t \equiv \{q_{ijt}\}$ is a weighted average of a positive-definite matrix, $\varepsilon_{t-1}\varepsilon_{t-1}'$, and a positive semi-definite matrix, Q_{t-1}. The unconditional expectation of q_{ijt} is $\bar{\rho}_{ij}$, while both q_{iit} and q_{jjt} each have an expected value of one (mean variance is one). The model in Eqn.(4.30) will be mean-reverting so long as $\alpha + \beta < 1$.[23]

A more flexible representation of the DCC-GARCH model in Eqn.(4.30),

[23]When $\alpha+\beta = 1$, Eqn.(4.30) reduces to a correlation process characterised by exponential smoothing, such that, $q_{ijt} = (1-\lambda)(\epsilon_{i(t-1)}\epsilon_{j(t-1)}) + \lambda(q_{ij(t-1)})$.

allowing for a GARCH(m, n) process in the dynamics of q_{ijt}, can be given by,

$$
\begin{aligned}
Q_t &= \bar{Q}(1 - \sum_{m=1}^{M} \alpha_m - \sum_{n=1}^{N} \beta_n) + \sum_{m=1}^{M} \alpha_m (\epsilon_{t-m}\epsilon'_{t-m}) + \sum_{n=1}^{N} \beta_n Q_{t-n} \\
\Gamma_t &= \text{diag}\{Q_t\}^{-1} Q_t \text{diag}\{Q_t\}^{-1}
\end{aligned}
\tag{4.31}
$$

where $Q_t \equiv \{q_{ijt}\}$ is the conditional variance-covariance matrix of residuals, where $\bar{Q}$ is the time-invariant (unconditional) variance-covariance matrix found by estimating Eqn.(4.24) in what is the *first stage* of the estimation process. Meanwhile, diag$\{Q_t\}$ is a diagonal matrix composed of the square root of the diagonal elements of Q_t such that

$$
\text{diag}\{Q_t\} = \begin{pmatrix} \sqrt{q_{11}} & 0 & \cdots & 0 \\ 0 & \sqrt{q_{22}} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sqrt{q_{kk}} \end{pmatrix} \tag{4.32}
$$

implying that, as before, a typical element of Γ_t will take the form $\rho_{ijt} = q_{ijt}/\sqrt{q_{iit}q_{jjt}}$. Eqn.(4.31) tells us that Q_t can be thought of as an autoregressive, moving-average process capturing deviations in the correlations around their unconditional values $(\bar{Q})$.

For the purposes of this study, the focus of interest is Γ_t, and in particular, $\rho_{1jt} = q_{1jt}/\sqrt{q_{11t}q_{jjt}}$, which represents the conditional correlation between the price of gold and each exchange-rate pair, j, in the dataset.

Estimation: dynamic conditional correlations

This section describes a two-stage estimation process proposed by Engle (2002) to estimate the multivariate GARCH model of dynamic conditional correlations outlined above.

In the first stage, univariate GARCH models are estimated for each returns series. In the second stage, the first-stage residuals are taken and transformed by their standard deviations in order to estimate the parameters of the dynamic conditional correlation model.

The multivariate GARCH model of dynamic conditional correlations can be specified as,

$$
\begin{aligned}
r_t|\psi_{t-1} &\sim N(0, H_t) \\
H_t &\equiv D_t \Gamma_t D_t \\
D_t^2 &= \operatorname{diag}\{\omega_i\} + \operatorname{diag}\{\kappa_i\} \otimes r_{t-1}r_{t-1}' + \operatorname{diag}\{\lambda_i\} \otimes D_{t-1}^2 \\
\epsilon_t &= D_t^{-1} r_t \\
Q_t &= \bar{Q} \otimes (\iota\iota' - A - B) + A \otimes \varepsilon_{t-1}\varepsilon_{t-1}' + B \otimes D_{t-1} \\
\Gamma_t &= \operatorname{diag}\{Q_t\}^{-1} Q_t \operatorname{diag}\{Q_t\}^{-1}
\end{aligned}
\tag{4.33}
$$

The assumption of multivariate normality in $r_t|\psi_{t-1} \sim N(0, H_t)$ permits maximum likelihood estimation. Without the assumption of normality (that is, when the returns have non-Gaussian innovations), the estimator described by Eqn.(4.33) can be interpreted as a Quasi Maximum Likelihood (QML) estimator, resulting in estimated parameters that are both consistent and asymptotically normal.

The returns, r_t, can be either mean zero or the residuals from a filtered series. The standard errors of the model will not depend on the filtering method because, as Engle and Sheppard (2001) note, the cross partial derivative of the log-likelihood with respect to the mean and the variance parameters has expectation zero when using the normal likelihood.

The third relationship in Eqn.(4.33), describing the behaviour of D_t^2, indicates that each variable follows a univariate GARCH process, where ω_i, κ_i and λ_i are the familiar non-negative coefficients of a traditional GARCH specification. If each variable were to follow something other than a univariate GARCH process, this would not alter the formulation of the rest of the statistical model in Eqn.(4.33).

Estimation of this multivariate GARCH model of dynamic conditional correlations can, as suggested above, be done in two stages. In the *first stage*, univariate GARCH models are estimated for each returns series, r_t. The *second stage* involves a transformation of the first-stage residuals by their standard deviations (where estimates of the standard deviations are drawn from the first-stage results), such that $\varepsilon_t = D_t^{-1} r_t$, with the transformed residuals being used to estimate the parameters of the dynamic conditional correlation model in Q_t.

The log-likelihood function for the estimator in Eqn.(4.33), where $r_t|\psi_{t-1} \sim$

$N(0, H_t)$, can be expressed as

$$\begin{aligned}
\log L(\theta) &= -\frac{1}{2}\sum_{t=1}^{T}\left(n\log 2\pi + \log|H_t| + r_t' H_t^{-1} r_t\right) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\left(n\log 2\pi + \log|D_t\Gamma_t D_t| + r_t' D_t^{-1}\Gamma_t^{-1} D_t^{-1} r_t\right) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\left(n\log 2\pi + 2\log|D_t| + \log|\Gamma_t| + \varepsilon_t'\Gamma_t^{-1}\varepsilon_t\right) \\
&= -\frac{1}{2}\sum_{t=1}^{T}\left(n\log 2\pi + 2\log|D_t| + r_t' D_t^{-1} D_t^{-1} r_t - \varepsilon_t'\varepsilon_t + \log|\Gamma_t| + \varepsilon_t'\Gamma_t^{-1}\varepsilon_t\right)
\end{aligned} \tag{4.34}$$

which can be maximised over the parameters of the model, where θ is the vector of unknown parameters.

On its own, the formulation in Eqn.(4.34) does not offer much in terms of computational efficiency when it comes to estimating large variance-covariance matrices (that is, when it comes to dealing with many variables). It is the two-step estimation process that offers the computational gains.

To help outline the two-step process, let the parameters of the model, θ, be separated into two groups, such that $(\phi_1, \phi_2, \ldots, \phi_n, \xi) = (\phi, \xi)$, where ϕ represents the parameters in D (the stochastic diagonal matrix of conditional standard deviations), and where ξ denotes the additional parameters contained in Γ (the correlation matrix). The elements of ϕ_i correspond to the parameters of the univariate GARCH model for the ith asset, such that, $\phi_i = (\omega, \kappa_{1i}, \ldots, \kappa_{p_i i}, \lambda_{1i}, \ldots, \lambda_{q_i i})$.

Engle (2002) shows that the log-likelihood can be expressed as the sum of a volatility component and a correlation component, where

$$\log L(\phi, \xi) = \log L_v(\phi) + \log L_c(\phi, \xi) \tag{4.35}$$

The *volatility* component of the log-likelihood, $\log L_v(\phi)$, contains the parameters of D from Eqn.(4.34) and can be written as

$$\log L_v(\phi) = -\frac{1}{2}\sum_{t=1}^{T}\left(n\log 2\pi + \log|D_t|^2 + r_t' D_t^{-2} r_t\right) \tag{4.36}$$

The *correlation* component of the log-likelihood, $\log L_c(\phi, \xi)$, containing the additional parameters in Γ, can be expressed as

$$\log L_c(\phi, \xi) = -\frac{1}{2} \sum_{t=1}^{T} \Big(\log |\Gamma_t| + \varepsilon_t' \Gamma_t^{-1} \varepsilon_t - \varepsilon_t' \varepsilon_t \Big) \tag{4.37}$$

Note that the volatility component of the log-likelihood is just the sum of the individual log-likelihoods for the GARCH models for each of the i assets. That is,

$$\begin{aligned} \log L_v(\phi) &= -\frac{1}{2} \sum_{t=1}^{T} \Big(n \log 2\pi + \log |D_t|^2 + r_t' D_t^{-2} r_t \Big) \\ &= -\frac{1}{2} \sum_{t=1}^{T} \Big(n \log 2\pi + \sum_{i=1}^{n} \Big(\log h_{it} + \frac{r_{it}^2}{h_{it}} \Big) \Big) \\ &= -\frac{1}{2} \sum_{i=1}^{n} \Big(T \log 2\pi + \sum_{t=1}^{T} \Big(\log h_{it} + \frac{r_{it}^2}{h_{it}} \Big) \Big) \end{aligned} \tag{4.38}$$

which is jointly maximised by maximising each term separately. Estimation of Eqn.(4.38) represents the first step of the two-step estimation process.

The second component of the log-likelihood, the correlation component, $\log L_c(\phi, \xi)$, is used to estimate the correlation parameters. The squared error terms are not dependent on the correlation parameters and are not, therefore, included in the first-order conditions. They can be ignored.

The two-step approach to maximising the log-likelihood in Eqn.(4.35) can be summarized as follows. First, find estimates of the parameters in D_t, the stochastic diagonal matrix of conditional standard deviations, which maximise the log-likelihood $\log L_v(\phi)$. That is, in the *first step*, find

$$\hat{\phi} = \arg\max\{\log L_v(\phi)\} \tag{4.39}$$

and then take this value as given in the *second step*, which involves maximising the log-likelihood of the correlation component, such that

$$\max_{\xi}\{\log L_c(\phi, \xi)\} \tag{4.40}$$

Engle (2002) explains that the maximum of the second step will be a function of the first-step estimates of the parameters, and as such, if the first step is consistent the second step will be consistent.

4.4 Data

This section presents an overview of the data and some preliminary analysis offering motivation for the empirical work undertaken in Section 4.5.

The dataset consists of the price of gold (US dollars per Troy ounce) and 16 US dollar exchange-rate pairings (expressed in terms of home currency per US dollar). The frequency of the data is weekly. The sample period extends from 10 January 1986 to 29 August 2008, comprising $t = 1,182$ observations per variable. Exchange rates are from Datastream. Gold prices are from Bloomberg.

In constructing the dataset the intention is to include as many exchange-rate pairings as possible. Exchange rates are excluded from the dataset only if data is unavailable at the selected frequency or if the exchange rate is fixed against the US dollar during the sample period. The 16 currencies included in the sample, all expressed in terms of home currency per US dollar, are the euro, yen, Indian rupee, Taiwan dollar, Australian dollar, Canadian dollar, Danish krone, Israeli Shekel, Maltese lira, New Zealand dollar, Norwegian krone, Singapore dollar, South African rand, Swedish krona, Swiss franc, and the UK pound.

Demeaned continuously compounded percentage returns of the exchange rates are calculated by taking the weekly difference of the natural logarithm of each exchange rate, subtracting the sample mean, then multiplying by 100. Demeaned returns for gold is calculated in a similar fashion.

Figure 4.1, Figure 4.2 and Figure 4.3 show the returns for the price of gold and for the 16 US dollar exchange rates included in the sample. Table 4.1 suggests that gold-price returns, like other commodity-price returns, are more variable than exchange-rate returns, mirroring the standard findings of other studies.[24] Variance for the price of gold (2.8) is more than double the average variance of the 16 nominal exchange rates in the sample (which have mean variance of 1.3).

All series seem to exhibit two common features of financial time series: excess kurtosis and volatility clustering. Indeed, Table 4.1 shows that in Jarque-Bera tests of normality, the null hypothesis of normality can be rejected in all cases for both exchange-rate returns and gold-price returns.

Financial time series, and in particular exchange rates, often exhibit little correlation in the mean processes (the returns) but significant correlation in the variance processes (the square of the returns). See for instance Baillie and Bollerslev (1989b) and Diebold and Nerlove (1989). Under such conditions, GARCH modelling is particularly appropriate. The next step here, then, is to quantify the correlation present in the returns and the square of the returns. This is done by employing the Ljung and Box (1978) portmanteau test for serial

[24] See for instance Clements and Fry (2008).

Figure 4.1: Gold-price returns and exchange-rate returns

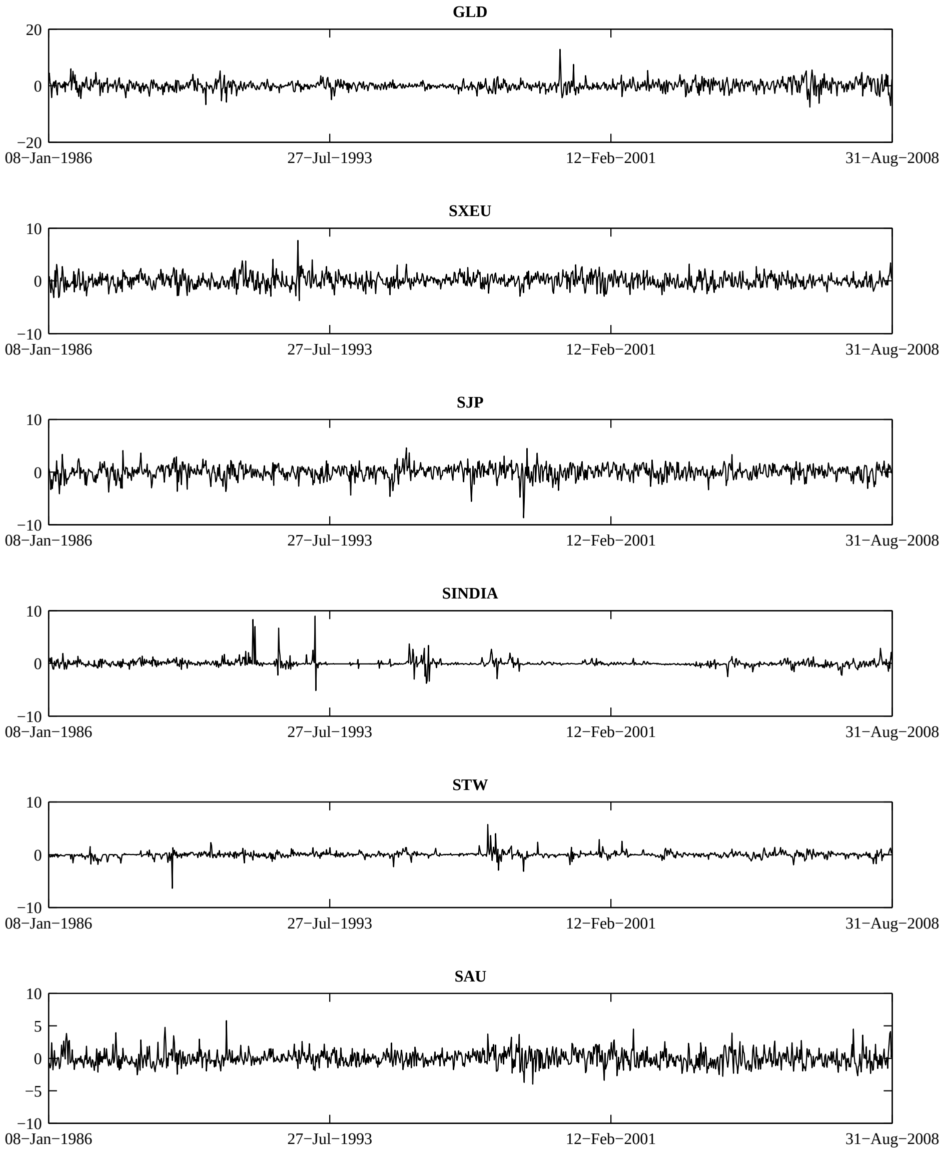

Notes: Percentage demeaned nominal returns. Exchange rates expressed as home currency per US dollar. Frequency is weekly. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU).

Figure 4.2: Exchange-rate returns

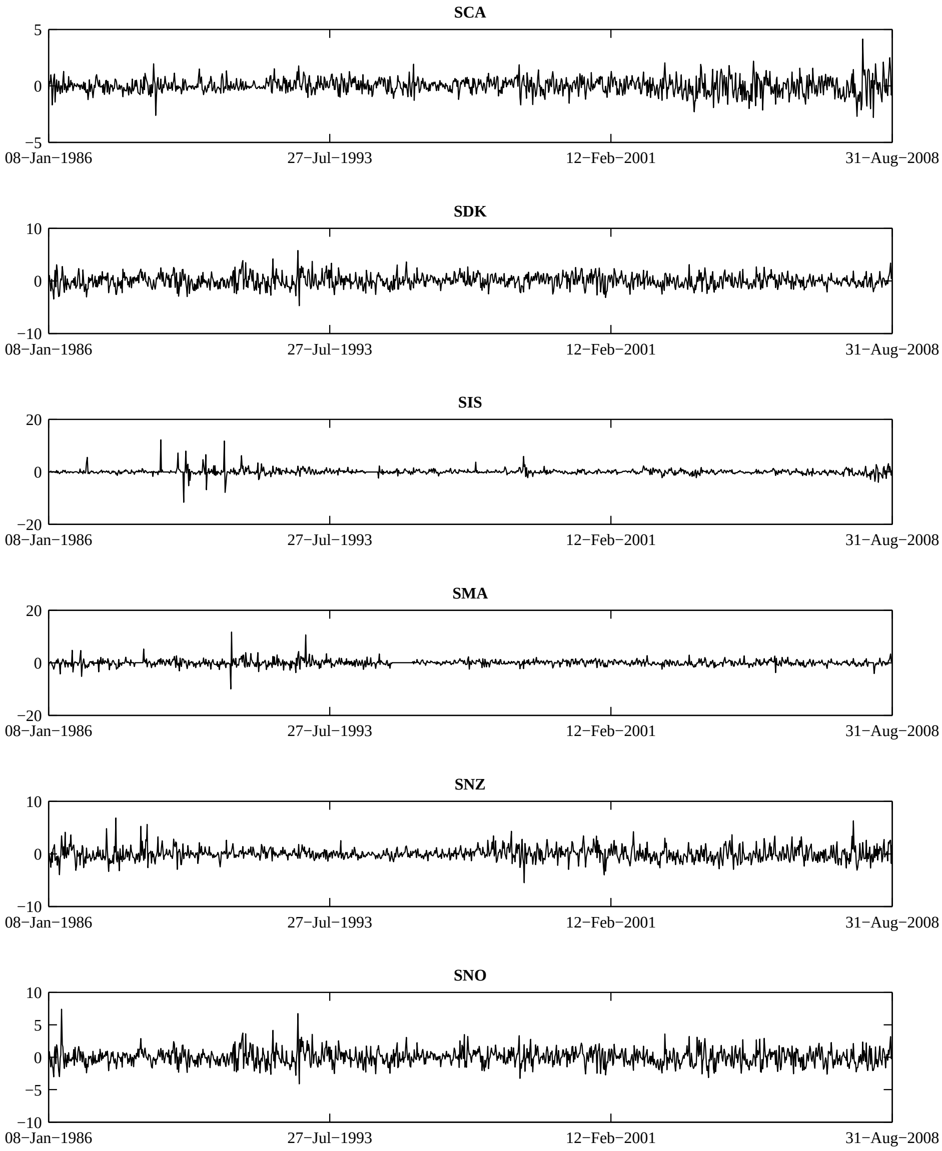

Notes: Percentage demeaned nominal commodity-price returns. Exchange rates expressed as home currency per US dollar. Frequency is weekly. Abbreviations: Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO).

Figure 4.3: Exchange-rate returns

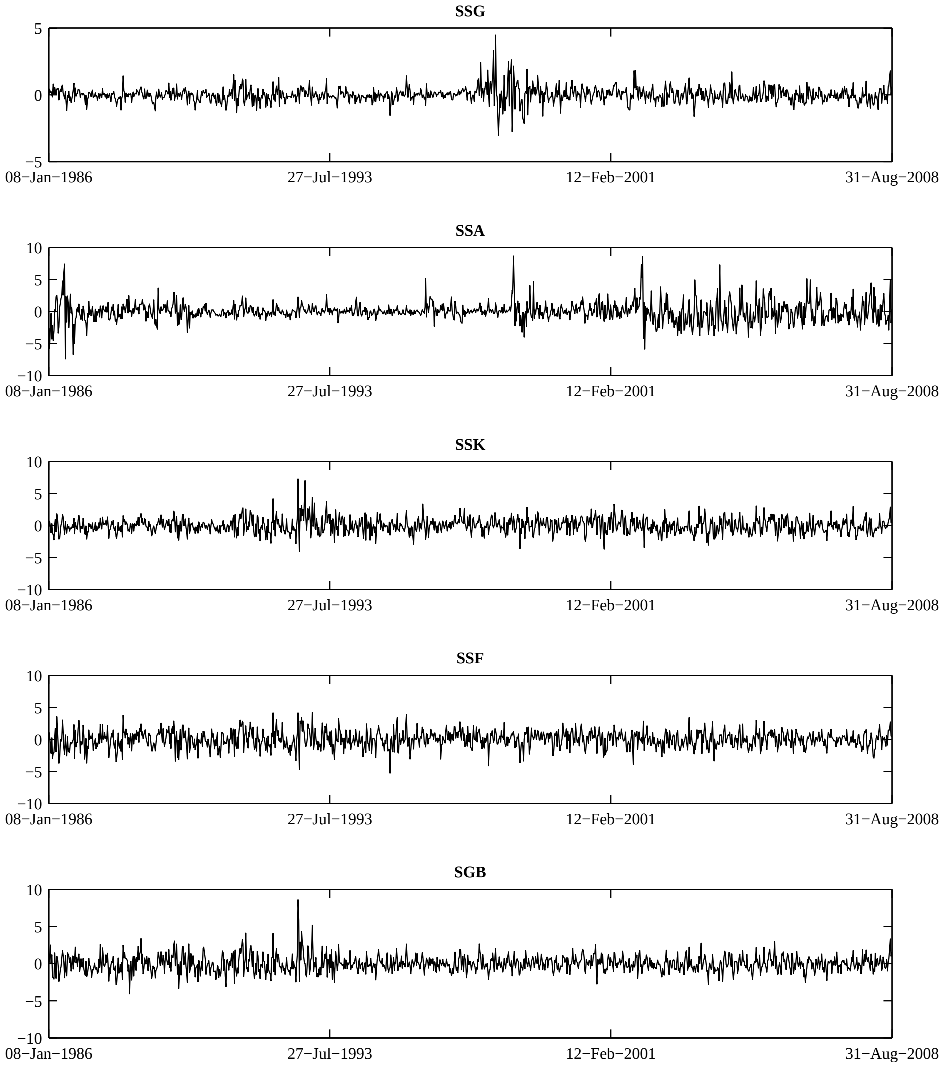

Notes: Percentage demeaned nominal currency returns. Exchange rates expressed as home currency per US dollar. Frequency is weekly. Abbreviations: Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

Table 4.1: Descriptive statistics

	GLD	SXEU	SJP	SINDIA	STW	SAU
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. dev.	1.660	1.147	1.229	0.755	0.580	1.135
Variance	2.757	1.317	1.511	0.571	0.336	1.288
Jarque-Bera	1,184	190	628	101,813	32,552	178
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SCA	SDK	SIS	SMA	SNZ	SNO
Mean	0.000	0.000	0.000	0.000	0.000	0.000
Std. dev.	0.672	1.167	1.154	1.164	1.205	1.182
Variance	0.452	1.363	1.332	1.355	1.451	1.397
Jarque-Bera	209	37	65,142	19,984	502	280
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SSG	SSA	SSK	SSF	SGB
Mean	0.000	0.000	0.000	0.000	0.000
Std. dev.	0.551	1.503	1.169	1.299	1.105
Variance	0.304	2.259	1.368	1.686	1.221
Jarque-Bera	2,899	1,284	375	17	1,056
Probability	0.000	0.000	0.000	0.000	0.000

Notes: All returns are demeaned. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

correlation. Under the null hypothesis of no serial correlation, the test statistic is asymptotically chi-square distributed.

Testing for serial correlation in the *square* of the returns, Table 4.2 shows that the Ljung and Box (1978) portmanteau test for up to twentieth-order correlation breaches the relevant critical value (31.401) for the 95% fractile in the asymptotic chi-square distribution for nearly all the currencies and the commodities in the dataset. That is, the null of no serial correlation is, for nearly all of the series, rejected.

Clearly, the returns here are not independent through time. Large returns tend to be followed by large returns, and small returns tend to be followed by small returns. Furthermore, positive returns are not necessarily followed by positive returns, nor are negative returns necessarily followed by negative returns: sign carries no predictability. These features are typical of the empirical characteristics, first formalised by Mussa (1979), of many financial series, and the model perhaps best able to capture this pattern of time dependence is the ARCH(q) model developed by Engle (1982), or more parsimoniously, the GARCH(p, q) model developed by Bollerslev (1986). Indeed, in the empirical analysis that follows, all conditional variances are assumed to behave in a manner consistent with GARCH(p, q) processes.

Table 4.2: Ljung-Box-Pierce Q-Test for Serial Correlation

	GLD	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
Ljung-Box (Mean)	120.043	81.245	122.472	86.258	134.654	94.589
Probability	0.000	0.000	0.000	0.000	0.000	0.000
Ljung-Box (Variance)	204.314	106.438	154.668	133.749	77.560	61.482
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SCA	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
Ljung-Box (Mean)	86.529	81.228	31.255	25.161	81.812	59.321
Probability	0.000	0.000	0.052	0.195	0.000	0.000
Ljung-Box (Variance)	503.304	149.507	57.006	129.643	108.678	80.187
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SSG	**SSA**	**SSK**	**SSF**	**SGB**
Ljung-Box (Mean)	110.883	132.219	88.959	75.842	97.630
Probability	0.000	0.000	0.000	0.000	0.000
Ljung-Box (Variance)	739.293	450.492	344.206	65.062	235.119
Probability	0.000	0.000	0.000	0.000	0.000

Notes: All returns are demeaned. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

4.5 Results

This section presents results from an empirical analysis of the association between exchange rates and the price of gold using a model of dynamic conditional correlations.

First, for the purposes of completeness, Section 4.5.1 presents results from a model in which the conditional correlations are constant. Section 4.5.2 presents results from a model in which the conditional correlations are dynamic.

4.5.1 Results: Constant conditional correlations

This subsection presents results from a constant-correlations model of the complete set of returns introduced in Section 4.4. Guided by the preliminary analysis presented in Section 4.4, the following analysis adopts a model of conditional correlations that is assumed to be characterised by conditional variances that follow a GARCH(p, q) structure.

The GARCH(p, q) model specifies the conditional variance as a linear function of the past q squared residuals and the past p conditional variances. As outlined in Section 4.3, in its general form the GARCH(p, q) model can be represented as

$$\begin{aligned}\mathrm{Var}_t(\epsilon_{it}) &= h_{iit}\\ &= \omega_i + \sum_{k=1}^{q} \alpha_{ik}\epsilon_{it-k}^2 + \sum_{l=1}^{p} \beta_{il} h_{iit-l}\end{aligned} \tag{4.41}$$

given the time-varying conditional variance-covariance matrix $H_t = Var(\epsilon_t|\psi_{t-1})$, where h_{ijt} denotes the ij^{th} element in H_t, where ψ_{t-1} is the measurable space generated by all the available information up to and including time $t-1$, and where ϵ_t is the error vector.

Eqn. (4.41) shows that the GARCH model can be thought of as a univariate, autoregressive, moving-average model of the conditional second moments, ϵ_{it-k}^2 and h_{iit-l}. This model has been shown in past studies to be particularly suitable for capturing the short-run movements of international exchange rates. See for instance Baillie and Bollerslev (1989b), Diebold and Nerlove (1989), Domowitz and Hakkio (1985) and Engle and Bollerslev (1986). Furthermore, of all GARCH(p, q) models available, a GARCH(1,1) model has been shown to offer a particularly useful and parsimonious description of short-run, currency dynamics. Implicit in the GARCH(1,1) model is an assertion that the best predictor of the variance in the next period is a weighted average of the long-run average variance (ω_i), the variance predicted for this period (h_{iit}), and the new information in this period that is captured by the most recent squared residual (ϵ_{it}^2).

As a step in testing the suitability, here, of the GARCH(1,1) structure, Table 4.3 displays estimates for *univariate* GARCH(1,1) models for the observed data. The three coefficients listed in Table 4.3 come from the variance equation, Eqn. (4.41). They are the intercept, ω_i, the coefficient on the first lag of the squared return, α_{i1}, and the coefficient on the first lag of the conditional variance, β_{i1}. The coefficients α_{i1} and β_{i1} sum to less than one, which is required in order to have a mean-reverting variance process. Where the sum is close to one, the reversion process is slow.

With few exceptions nearly all the parameters in the time-varying conditional variances are individually significant at the 5% level of significance and in likelihood-ratio tests for the absence of conditional heteroscedasticity (that is, in tests of $\alpha_{i1} = \beta_{i1} = 0$ for all i), the test statistic, which takes a chi-square distribution under the null hypothesis, is at its minimum, 44.9, with a critical value of 6.0, leading us to reject, for all i, the premise that the data can be modelled adequately with a homoscedastic, seemingly-unrelated-regression

Table 4.3: Estimates for time-varying conditional variances

i	**GLD**	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
ω_i	0.031	0.035	0.093	0.004	0.062	0.026
	(0.015)	(0.017)	(0.046)	(0.001)	(0.015)	(0.015)
α_{i1}	0.895	0.914	0.866	0.576	0.426	0.929
	(0.020)	(0.026)	(0.051)	(0.022)	(0.058)	(0.024)
β_{i1}	0.101	0.058	0.065	0.424	0.574	0.051
	(0.021)	(0.017)	(0.023)	(0.048)	(0.146)	(0.016)

i	**SCA**	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
ω_i	0.000	0.037	0.209	0.033	0.005	0.079
	(0.001)	(0.018)	(0.049)	(0.013)	(0.004)	(0.033)
α_{i1}	0.858	0.911	0.457	0.913	0.950	0.853
	(0.018)	(0.026)	(0.054)	(0.020)	(0.011)	(0.041)
β_{i1}	0.142	0.061	0.543	0.068	0.048	0.089
	(0.020)	(0.017)	(0.148)	(0.019)	(0.011)	(0.025)

i	**SSG**	**SSA**	**SSK**	**SSF**	**SGB**
ω_i	0.008	0.018	0.032	0.026	0.021
	(0.003)	(0.007)	(0.017)	(0.018)	(0.011)
α_{i1}	0.868	0.861	0.920	0.950	0.938
	(0.027)	(0.021)	(0.025)	(0.019)	(0.020)
β_{i1}	0.106	0.139	0.055	0.033	0.041
	(0.023)	(0.024)	(0.016)	(0.012)	(0.013)

Notes: Asymptotic standard errors are in parentheses. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

model (SUR model). See Table 4.4.

By employing the GARCH(1,1) model, the aim is to capture all the dynamic features of the mean and the variance. The estimated residuals should be serially uncorrelated and should contain no remaining conditional volatility. To test for this, the first step is to create a set of standardised residuals $\hat{s}_{it}$ (where $\hat{s}_{it} = \hat{\epsilon}_{it}/\hat{h}_{it}^{1/2}$). The standardised residuals will have a mean of zero and a variance of one.

If there is any serial correlation in the standardised residuals, $\hat{s}_{it}$, then the implication is that the model of the mean is not properly specified. Here the model of the mean is specified as, $r_{it} = E(r_{it}|\psi_{t-1}) + \epsilon_{it}$. To test the suitability of this model, it is necessary to calculate Ljung-Box test statistics for $\hat{s}_{it}$. Rejection of the null hypothesis would imply that the various test statistics are significantly different from zero and the model of the mean has been poorly specified. Table 4.5 shows that in tests for serial correlation in the standardised residuals, there is little evidence of serial correlation. Mostly the mean is adequately specified.

If the GARCH(1,1) model specified in Eqn. (4.41) captures adequately the variance characteristics of the data under analysis then the residuals of the GARCH(1,1) model should be free of any remaining GARCH effects. This can be tested by calculating Ljung-Box test statistics for the *squared* standardised

Table 4.4: Likelihood ratio test for absence of conditional heteroscedasticity

	GLD	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
Likelihood	142.251	51.925	39.258	380.801	129.378	44.930
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SCA	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
Likelihood	126.243	56.867	144.446	68.594	130.123	45.628
Probability	0.000	0.000	0.000	0.000	0.000	0.000

	SSG	**SSA**	**SSK**	**SSF**	**SGB**	
Likelihood	154.661	316.120	64.442	28.245	70.431	
Probability	0.000	0.000	0.000	0.000	0.000	

Notes: Probability values are in parentheses. The likelihood-ratio test statistic is asymptotically chi-square distributed with degrees of freedom equal to the number of restrictions imposed (here, two, where $\alpha_{i1} = \beta_{i2} = 0$). Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

residuals, $\hat{s}^2_{it}$. If $\hat{s}^2_{it}$ is a good estimate of $v^2_{it} = \epsilon^2_{it}/h_{it}$ then $\hat{s}^2_{it}$ should have the characteristics of a white-noise process. If there are no remaining GARCH effects then it will not be possible to reject the null hypothesis that the sample values of the test statistics are equal to zero. Indeed, the Ljung-Box statistics reported in Table 4.5 indicate that for the overwhelming majority of the returns series the GARCH(1,1) model adequately captures all relevant volatility dynamics.

Although the GARCH(1,1) model appears to offer an adequate description of the volatility dynamics, it is useful to consider longer lag lengths. Table 4.6, shows estimates for GARCH(4,4) models while Table , Table and Table present the estimation results from GARCH(12,12) models. The results for the GARCH(4,4) models show that for nearly all of the univariate series the second lags and subsequent lags are not significant—the effect of lagged shocks dies out fairly rapidly. Results from the GARCH(12,12) models are similar. Long lags do not capture any particularly valuable dynamics.

Given the univariate analysis above, it is assumed that the GARCH(1,1) model offers an adequate representation of the conditional variances under analysis. That is, in the correlation model that follows, the conditional variances are assumed to follow a GARCH(1,1) structure while the conditional correlations between the returns are assumed to take constant, non-zero values as outlined previously in Section 4.3.1.

Incorporating the GARCH(1,1) structure for conditional variances into a model of constant conditional correlations using the notation of Eqn. (4.4),

Table 4.5: Model adequacy: tests for serial correlation in the standardised residuals and squared standardised residuals

	GLD	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
Ljung-Box (Std resids)	40.130	10.337	16.802	51.036	38.156	18.237
Probability	0.005	0.962	0.666	0.000	0.008	0.572
Ljung-Box (Std resids sqrd)	18.867	15.108	33.631	19.145	15.974	9.494
Probability	0.531	0.770	0.029	0.512	0.718	0.976

	SCA	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
Ljung-Box (Std resids)	16.454	11.365	13.764	19.723	23.400	11.300
Probability	0.688	0.936	0.842	0.475	0.270	0.938
Ljung-Box (Std resids sqrd)	14.397	21.068	0.943	20.378	22.882	14.564
Probability	0.810	0.393	1.000	0.434	0.295	0.801

	SSG	**SSA**	**SSK**	**SSF**	**SGB**
Ljung-Box (Std resids)	22.779	29.206	12.902	9.599	17.361
Probability	0.300	0.084	0.882	0.975	0.629
Ljung-Box (Std resids sqrd)	14.328	7.910	20.501	12.729	15.665
Probability	0.813	0.992	0.427	0.889	0.737

Notes: Table reports Ljung-Box lack-of-fit hypothesis tests for model misspecification. Under the null hypothesis that the model fit is adequate, the test statistic is asymptotically chi-square distributed. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

Eqn. (4.5), Eqn. (4.6) and Eqn. (4.7), gives

$$\begin{aligned}
r_{it} &= E(r_{it}|\psi_{t-1}) + \epsilon_{it} \\
\mathrm{Var}_t(\epsilon_{it}) &= h_{iit} \\
h_{iit} &= \omega_i + \alpha_{i1}\epsilon^2_{it-1} + \beta_{i1}h_{iit-1} \\
& \qquad\qquad i,j = \text{oil,euro,}\ldots \quad i \neq j \\
h_{ijt} &= \rho_{ij}\sqrt{h_{iit}h_{jjt}}
\end{aligned} \tag{4.42}$$

where, as before, r_{it} defines the returns, h_{iit} defines the conditional variances and h_{ijt} defines the conditional covariances. Estimation of this model is undertaken by maximum likelihood in line with the discussion in Section 4.3.1.

Maximum likelihood estimates for the model in Eqn. (4.42) are presented in Table 4.10 and Table 4.11 under the assumption of conditional normality (that is, under the assumption that, conditional on information up to and including period $t-1$, the error term ϵ_t is distributed normally with mean zero and variance H_t).

Conditional normality is, of course, a strong assumption. Bollerslev (1990) notes that the assumption of conditional normality, in the context of modelling asset prices even after accounting for ARCH effects, does not make for a good approximation when using daily data because of the tendency of daily data

Table 4.6: GARCH(4,4) estimates for time-varying conditional variances

i	**GLD**	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
ω_i	0.136	0.113	0.412	0.048	0.104	0.119
	(0.051)	(0.054)	(0.102)	(0.009)	(0.005)	(0.069)
α_{i1}	0.236	0.051	0.138	0.336	0.681	0.085
	(0.054)	(0.037)	(0.035)	(0.039)	(0.079)	(0.033)
α_{i2}	0.006	0.026	0.066	0.175	0.014	0.000
	(0.021)	(0.013)	(0.031)	(0.052)	(0.042)	(0.023)
α_{i3}	0.070	0.076	0.000	0.096	0.275	0.000
	(0.029)	(0.035)	(0.022)	(0.027)	(0.023)	(0.025)
α_{i4}	0.035	0.000	0.085	0.163	0.030	0.081
	(0.067)	(0.026)	(0.033)	(0.036)	(0.020)	(0.075)
β_{i1}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.043)	(0.060)	(0.021)	(0.085)	(0.064)	(0.033)
β_{i2}	0.000	0.450	0.000	0.229	0.000	0.099
	(0.029)	(0.031)	(0.010)	(0.076)	(0.032)	(0.032)
β_{i3}	0.289	0.310	0.000	0.000	0.000	0.418
	(0.172)	(0.032)	(0.009)	(0.056)	(0.031)	(0.031)
β_{i4}	0.352	0.000	0.433	0.000	0.000	0.228
	(0.163)	(0.076)	(0.095)	(0.030)	(0.025)	(0.025)

i	**SCA**	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
ω_i	0.021	0.104	0.066	0.029	0.010	0.221
	(0.008)	(0.078)	(0.011)	(0.014)	(0.008)	(0.063)
α_{i1}	0.132	0.099	0.120	0.091	0.048	0.093
	(0.030)	(0.035)	(0.026)	(0.021)	(0.018)	(0.021)
α_{i2}	0.077	0.034	0.000	0.000	0.060	0.036
	(0.037)	(0.025)	(0.005)	(0.038)	(0.019)	(0.027)
α_{i3}	0.115	0.042	0.000	0.000	0.040	0.150
	(0.031)	(0.042)	(0.008)	(0.049)	(0.018)	(0.034)
α_{i4}	0.000	0.000	0.150	0.000	0.000	0.000
	(0.029)	(0.025)	(0.041)	(0.014)	(0.015)	(0.043)
β_{i1}	0.000	0.000	0.000	0.000	0.000	0.164
	(0.113)	(0.076)	(0.034)	(0.022)	(0.053)	(0.322)
β_{i2}	0.000	0.135	0.000	0.185	0.000	0.000
	(0.023)	(0.023)	(0.019)	(0.242)	(0.021)	(0.043)
β_{i3}	0.484	0.266	0.364	0.239	0.000	0.001
	(0.256)	(0.256)	(0.048)	(0.127)	(0.028)	(0.060)
β_{i4}	0.157	0.347	0.366	0.463	0.852	0.403
	(0.235)	(0.235)	(0.048)	(0.203)	(0.042)	(0.280)

i	**SSG**	**SSA**	**SSK**	**SSF**	**SGB**
ω_i	0.021	0.073	0.146	0.045	0.072
	(0.018)	(0.077)	(0.061)	(0.039)	(0.017)
α_{i1}	0.154	0.335	0.069	0.069	0.092
	(0.077)	(0.050)	(0.036)	(0.028)	(0.023)
α_{i2}	0.027	0.000	0.036	0.000	0.000
	(0.014)	(0.388)	(0.022)	(0.023)	(0.026)
α_{i3}	0.088	0.000	0.113	0.000	0.077
	(0.015)	(0.150)	(0.033)	(0.026)	(0.027)
α_{i4}	0.000	0.000	0.000	0.000	0.000
	(0.175)	(0.499)	(0.034)	(0.039)	(0.105)
β_{i1}	0.000	0.314	0.000	0.000	0.000
	(0.048)	(0.405)	(0.080)	(0.044)	(0.116)
β_{i2}	0.000	0.100	0.000	0.000	0.000
	(0.098)	(0.601)	(0.100)	(0.068)	(0.023)
β_{i3}	0.507	0.000	0.288	0.288	0.263
	(0.511)	(0.385)	(0.157)	(0.102)	(0.295)
β_{i4}	0.152	0.251	0.387	0.613	0.508
	(0.445)	(0.256)	(0.151)	(0.136)	(0.234)

Notes: Asymptotic standard errors in parenthases. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

Table 4.7: GARCH(12,12) estimates for time-varying conditional variances

i	**GLD**	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
ω_i	0.208	0.247	0.574	0.054	0.095	0.274
	(0.081)	(0.090)	(0.223)	(0.007)	(0.000)	(0.094)
α_{i1}	0.208	0.064	0.118	0.365	0.609	0.090
	(0.054)	(0.020)	(0.042)	(0.046)	(0.034)	(0.032)
α_{i2}	0.013	0.034	0.047	0.205	0.013	0.000
	(0.032)	(0.024)	(0.032)	(0.056)	(0.013)	(0.022)
α_{i3}	0.079	0.033	0.000	0.216	0.268	0.000
	(0.031)	(0.028)	(0.046)	(0.044)	(0.028)	(0.019)
α_{i4}	0.089	0.000	0.084	0.039	0.033	0.147
	(0.039)	(0.023)	(0.032)	(0.047)	(0.040)	(0.045)
α_{i5}	0.050	0.038	0.078	0.030	0.000	0.000
	(0.036)	(0.021)	(0.026)	(0.012)	(0.047)	(0.019)
α_{i6}	0.025	0.011	0.012	0.034	0.000	0.000
	(0.025)	(0.019)	(0.039)	(0.008)	(0.037)	(0.024)
α_{i7}	0.050	0.000	0.000	0.000	0.000	0.018
	(0.034)	(0.018)	(0.030)	(0.014)	(0.030)	(0.038)
α_{i8}	0.012	0.089	0.000	0.000	0.000	0.038
	(0.010)	(0.036)	(0.015)	(0.007)	(0.050)	(0.027)
α_{i9}	0.000	0.000	0.000	0.000	0.000	0.021
	(0.043)	(0.038)	(0.043)	(0.008)	(0.023)	(0.027)
α_{i10}	0.000	0.043	0.023	0.112	0.000	0.000
	(0.015)	(0.035)	(0.028)	(0.005)	(0.011)	(0.021)
α_{i11}	0.052	0.117	0.000	0.000	0.030	0.044
	(0.044)	(0.043)	(0.035)	(0.002)	(0.022)	(0.028)
α_{i12}	0.067	0.000	0.027	0.000	0.048	0.000
	(0.034)	(0.015)	(0.045)	(0.013)	(0.029)	(0.023)
β_{i1}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.046)	(0.044)	(0.033)	(0.030)	(0.024)	(0.013)
β_{i2}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.054)	(0.034)	(0.010)	(0.102)	(0.006)	(0.020)
β_{i3}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.038)	(0.020)	(0.009)	(0.051)	(0.006)	(0.020)
β_{i4}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.074)	(0.042)	(0.000)	(0.029)	(0.052)	(0.037)
β_{i5}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.019)	(0.019)	(0.006)	(0.015)	(0.065)	(0.029)
β_{i6}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.050)	(0.024)	(0.033)	(0.028)	(0.037)	(0.011)
β_{i7}	0.069	0.000	0.216	0.000	0.000	0.000
	(0.068)	(0.020)	(0.006)	(0.008)	(0.008)	(0.009)
β_{i8}	0.287	0.000	0.000	0.000	0.000	0.000
	(0.089)	(0.052)	(0.009)	(0.013)	(0.013)	(0.046)
β_{i9}	0.000	0.000	0.000	0.000	0.000	0.438
	(0.012)	(0.024)	(0.017)	(0.011)	(0.015)	(0.092)
β_{i10}	0.000	0.000	0.002	0.000	0.000	0.000
	(0.033)	(0.025)	(0.009)	(0.003)	(0.031)	(0.020)
β_{i11}	0.000	0.054	0.000	0.000	0.000	0.000
	(0.021)	(0.076)	(0.010)	(0.010)	(0.026)	(0.011)
β_{i12}	0.000	0.327	0.000	0.000	0.000	0.000
	(0.053)	(0.068)	(0.011)	(0.008)	(0.015)	(0.023)

Notes: Asymptotic standard errors in parenthases. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

Table 4.8: GARCH(12,12) estimates for time-varying conditional variances

i	**SCA**	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
ω_i	0.042	0.215	0.073	0.077	0.028	0.337
	(0.017)	(0.093)	(0.013)	(0.024)	(0.016)	(0.144)
α_{i1}	0.139	0.100	0.155	0.113	0.041	0.095
	(0.027)	(0.036)	(0.029)	(0.022)	(0.029)	(0.046)
α_{i2}	0.078	0.039	0.000	0.000	0.021	0.044
	(0.038)	(0.027)	(0.004)	(0.031)	(0.023)	(0.029)
α_{i3}	0.090	0.024	0.000	0.000	0.025	0.123
	(0.026)	(0.028)	(0.003)	(0.006)	(0.024)	(0.042)
α_{i4}	0.033	0.000	0.000	0.000	0.000	0.023
	(0.037)	(0.036)	(0.015)	(0.007)	(0.007)	(0.044)
α_{i5}	0.020	0.034	0.038	0.021	0.024	0.028
	(0.015)	(0.033)	(0.026)	(0.013)	(0.023)	(0.031)
α_{i6}	0.121	0.016	0.000	0.000	0.000	0.046
	(0.041)	(0.036)	(0.003)	(0.013)	(0.037)	(0.036)
α_{i7}	0.014	0.000	0.000	0.083	0.000	0.034
	(0.025)	(0.029)	(0.008)	(0.009)	(0.033)	(0.032)
α_{i8}	0.000	0.088	0.000	0.000	0.038	0.058
	(0.047)	(0.033)	(0.008)	(0.012)	(0.025)	(0.041)
α_{i9}	0.079	0.000	0.000	0.006	0.024	0.000
	(0.042)	(0.047)	(0.020)	(0.012)	(0.019)	(0.007)
α_{i10}	0.000	0.048	0.000	0.000	0.094	0.000
	(0.035)	(0.030)	(0.005)	(0.013)	(0.038)	(0.038)
α_{i11}	0.033	0.084	0.293	0.052	0.000	0.118
	(0.035)	(0.019)	(0.037)	(0.047)	(0.016)	(0.047)
α_{i12}	0.000	0.018	0.000	0.000	0.000	0.000
	(0.002)	(0.007)	(0.019)	(0.006)	(0.031)	(0.071)
β_{i1}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.037)	(0.052)	(0.023)	(0.037)	(0.016)	(0.073)
β_{i2}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.114)	(0.074)	(0.006)	(0.025)	(0.017)	(0.017)
β_{i3}	0.000	0.000	0.051	0.000	0.000	0.000
	(0.041)	(0.026)	(0.013)	(0.021)	(0.018)	(0.088)
β_{i4}	0.000	0.000	0.067	0.000	0.000	0.000
	(0.044)	(0.092)	(0.022)	(0.061)	(0.014)	(0.019)
β_{i5}	0.000	0.000	0.000	0.000	0.609	0.000
	(0.022)	(0.030)	(0.022)	(0.027)	(0.107)	(0.057)
β_{i6}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.066)	(0.050)	(0.003)	(0.018)	(0.040)	(0.051)
β_{i7}	0.323	0.000	0.000	0.008	0.000	0.000
	(0.064)	(0.056)	(0.004)	(0.023)	(0.007)	(0.096)
β_{i8}	0.000	0.000	0.000	0.000	0.000	0.000
	(0.031)	(0.029)	(0.006)	(0.031)	(0.018)	(0.077)
β_{i9}	0.000	0.000	0.000	0.664	0.000	0.000
	(0.028)	(0.040)	(0.010)	(0.023)	(0.040)	(0.041)
β_{i10}	0.000	0.000	0.264	0.000	0.000	0.000
	(0.015)	(0.059)	(0.019)	(0.037)	(0.021)	(0.045)
β_{i11}	0.000	0.000	0.130	0.000	0.000	0.000
	(0.071)	(0.012)	(0.022)	(0.025)	(0.059)	(0.054)
β_{i12}	0.000	0.391	0.000	0.000	0.112	0.200
	(0.030)	(0.024)	(0.015)	(0.035)	(0.000)	(0.106)

Notes: Asymptotic standard errors in parenthases. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

Table 4.9: GARCH(12,12) estimates for time-varying conditional variances

i	**SSG**	**SSA**	**SSK**	**SSF**	**SGB**
ω_i	0.040	0.139	0.332	0.108	0.152
	(0.011)	(0.060)	(0.121)	(0.069)	(0.069)
α_{i1}	0.164	0.368	0.037	0.075	0.087
	(0.040)	(0.065)	(0.019)	(0.032)	(0.032)
α_{i2}	0.037	0.117	0.031	0.013	0.000
	(0.029)	(0.013)	(0.024)	(0.017)	(0.017)
α_{i3}	0.114	0.083	0.108	0.000	0.115
	(0.038)	(0.036)	(0.029)	(0.033)	(0.033)
α_{i4}	0.044	0.000	0.000	0.000	0.000
	(0.026)	(0.021)	(0.024)	(0.022)	(0.022)
α_{i5}	0.000	0.073	0.064	0.062	0.038
	(0.025)	(0.020)	(0.025)	(0.018)	(0.036)
α_{i6}	0.034	0.034	0.011	0.000	0.000
	(0.044)	(0.033)	(0.035)	(0.022)	(0.025)
α_{i7}	0.030	0.000	0.015	0.000	0.000
	(0.028)	(0.021)	(0.028)	(0.010)	(0.034)
α_{i8}	0.077	0.000	0.031	0.059	0.003
	(0.027)	(0.029)	(0.022)	(0.005)	(0.010)
α_{i9}	0.000	0.000	0.027	0.000	0.000
	(0.058)	(0.015)	(0.032)	(0.026)	(0.043)
α_{i10}	0.029	0.000	0.085	0.000	0.000
	(0.039)	(0.019)	(0.014)	(0.013)	(0.015)
α_{i11}	0.000	0.000	0.021	0.027	0.000
	(0.032)	(0.042)	(0.005)	(0.037)	(0.044)
α_{i12}	0.000	0.026	0.000	0.000	0.073
	(0.020)	(0.026)	(0.029)	(0.034)	(0.034)
β_{i1}	0.000	0.000	0.000	0.000	0.000
	(0.031)	(0.022)	(0.003)	(0.011)	(0.046)
β_{i2}	0.000	0.000	0.002	0.000	0.000
	(0.009)	(0.028)	(0.029)	(0.008)	(0.054)
β_{i3}	0.000	0.000	0.001	0.000	0.000
	(0.045)	(0.066)	(0.052)	(0.037)	(0.038)
β_{i4}	0.000	0.000	0.000	0.000	0.000
	(0.051)	(0.027)	(0.043)	(0.044)	(0.074)
β_{i5}	0.000	0.000	0.000	0.000	0.000
	(0.023)	(0.091)	(0.035)	(0.016)	(0.019)
β_{i6}	0.000	0.000	0.000	0.000	0.000
	(0.034)	(0.067)	(0.037)	(0.019)	(0.050)
β_{i7}	0.034	0.173	0.000	0.000	0.177
	(0.118)	(0.056)	(0.089)	(0.131)	(0.068)
β_{i8}	0.000	0.000	0.000	0.000	0.252
	(0.063)	(0.038)	(0.071)	(0.090)	(0.089)
β_{i9}	0.000	0.000	0.000	0.283	0.000
	(0.054)	(0.032)	(0.091)	(0.011)	(0.012)
β_{i10}	0.122	0.036	0.317	0.000	0.000
	(0.020)	(0.017)	(0.018)	(0.026)	(0.033)
β_{i11}	0.000	0.090	0.000	0.000	0.124
	(0.013)	(0.026)	(0.013)	(0.170)	(0.021)
β_{i12}	0.174	0.000	0.000	0.410	0.001
	(0.060)	(0.000)	(0.107)	(0.025)	(0.053)

Notes: Asymptotic standard errors in parenthases. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

Table 4.10: Estimates: model of constant conditional correlations

i	**GLD**	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**	**SCA**	**SDK**	**SIS**
ω_i	0.063	0.038	0.138	0.020	0.059	0.044	0.008	0.041	0.023
	(0.039)	(0.023)	(0.055)	(0.015)	(0.026)	(0.044)	(0.003)	(0.025)	(0.013)
α_{i1}	0.127	0.058	0.101	0.351	0.545	0.060	0.121	0.062	0.039
	(0.050)	(0.021)	(0.037)	(0.049)	(0.231)	(0.033)	(0.027)	(0.021)	(0.024)
β_{i1}	0.863	0.913	0.807	0.649	0.455	0.907	0.867	0.907	0.945
	(0.044)	(0.033)	(0.057)	(0.090)	(0.141)	(0.062)	(0.025)	(0.035)	(0.024)
ρ_{GLDi}	1.000	—	—	—	—	—	—	—	—
	(—)								
ρ_{SXEUi}	-0.309	1.000	—	—	—	—	—	—	—
	(0.010)	(—)							
ρ_{SXJPi}	-0.214	0.500	1.000	—	—	—	—	—	—
	(0.008)	(0.017)	(—)						
$\rho_{SINDIAi}$	-0.087	0.173	0.152	1.000	—	—	—	—	—
	(0.020)	(2.935)	(0.030)	(—)					
ρ_{STWi}	-0.155	0.262	0.268	0.085	1.000	—	—	—	—
	(0.054)	(0.021)	(0.018)	(0.032)	(—)				
ρ_{SAUi}	-0.341	0.248	0.148	0.097	0.261	1.000	—	—	—
	(0.020)	(0.037)	(0.025)	(0.018)	(0.024)	(—)			
ρ_{SCAi}	-0.159	0.155	0.065	0.074	0.174	0.380	1.000	—	—
	(0.008)	(0.124)	(0.079)	(0.011)	(0.008)	(0.009)	(—)		
ρ_{SDKi}	-0.320	0.976	0.501	0.173	0.260	0.246	0.160	1.000	—
	(0.012)	(0.027)	(0.039)	(0.006)	(0.006)	(0.003)	(0.004)	(—)	
ρ_{SISi}	-0.057	0.197	0.092	0.047	0.061	0.087	0.109	0.193	1.000
	(0.016)	(0.011)	(0.014)	(0.006)	(0.006)	(0.003)	(0.002)	(0.005)	(—)
ρ_{SMAi}	-0.287	0.603	0.297	0.124	0.203	0.228	0.148	0.581	0.179
	(0.007)	(0.012)	(0.008)	(0.006)	(0.005)	(0.003)	(0.006)	(0.003)	(0.005)
ρ_{SNZi}	-0.256	0.319	0.231	0.116	0.252	0.676	0.287	0.323	0.097
	(0.027)	(0.182)	(0.009)	(0.007)	(0.006)	(0.003)	(0.002)	(0.004)	(0.002)
ρ_{SNOi}	-0.300	0.859	0.435	0.144	0.241	0.274	0.202	0.863	0.183
	(0.017)	(0.039)	(0.035)	(0.016)	(0.013)	(0.008)	(0.005)	(0.004)	(0.003)
ρ_{SSGi}	-0.261	0.572	0.518	0.183	0.339	0.306	0.131	0.561	0.140
	(0.020)	(0.010)	(0.011)	(0.011)	(0.009)	(0.007)	(0.004)	(0.003)	(0.005)
ρ_{SSAi}	-0.257	0.431	0.236	0.153	0.205	0.272	0.165	0.430	0.133
	(0.019)	(0.014)	(0.012)	(0.015)	(0.012)	(0.005)	(0.003)	(0.007)	(0.004)
ρ_{SSKi}	-0.268	0.855	0.433	0.153	0.253	0.303	0.215	0.847	0.198
	(0.021)	(0.021)	(0.348)	(0.026)	(0.021)	(0.010)	(0.005)	(0.005)	(0.004)
ρ_{SSFi}	-0.317	0.911	0.542	0.145	0.227	0.194	0.102	0.919	0.156
	(0.047)	(0.156)	(2.444)	(0.022)	(0.032)	(0.009)	(0.005)	(0.006)	(0.005)
ρ_{SGBi}	-0.262	0.741	0.381	0.205	0.213	0.279	0.166	0.730	0.156
	(0.026)	(0.013)	(0.018)	(0.020)	(0.021)	(0.009)	(0.004)	(0.004)	(0.005)

Notes: Table reports maximum likelihood estimates for a multivariate GARCH model of constant conditional correlations where $h_{it} = \omega_i + \alpha_{i1}\epsilon^2_{i(t-1)} + \beta_{i1}h_{i(t-1)}$ and where $h_{ijt} = \rho_{ij}\sqrt{h_{iit}h_{jjt}}$. The upper panel reports parameters in the time varying conditional variances. The lower panel reports conditional correlations. Asymptotic standard errors are in parenthases. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS). Exchange rates priced in terms of home currency per US dollars.

Table 4.11: Estimates: model of constant conditional correlations

i	**SMA**	**SNZ**	**SNO**	**SSG**	**SSA**	**SSK**	**SSF**	**SGB**
ω_i	0.014	0.003	0.078	0.009	0.063	0.042	0.020	0.015
	(0.013)	(0.003)	(0.035)	(0.005)	(0.033)	(0.019)	(0.035)	(0.017)
α_{i1}	0.041	0.039	0.111	0.104	0.211	0.067	0.023	0.039
	(0.025)	(0.012)	(0.033)	(0.034)	(0.068)	(0.021)	(0.024)	(0.018)
β_{i1}	0.951	0.960	0.837	0.864	0.776	0.902	0.965	0.949
	(0.029)	(0.012)	(0.044)	(0.043)	(0.059)	(0.028)	(0.044)	(0.032)
ρ_{GLDi}	—	—	—	—	—	—	—	—
ρ_{SXEUi}	—	—	—	—	—	—	—	—
ρ_{SXJPi}	—	—	—	—	—	—	—	—
$\rho_{SINDIAi}$	—	—	—	—	—	—	—	—
ρ_{STWi}	—	—	—	—	—	—	—	—
ρ_{SAUi}	—	—	—	—	—	—	—	—
ρ_{SCAi}	—	—	—	—	—	—	—	—
ρ_{SDKi}	—	—	—	—	—	—	—	—
ρ_{SISi}	—	—	—	—	—	—	—	—
ρ_{SMAi}	1.000	—	—	—	—	—	—	—
	(—)							
ρ_{SNZi}	0.241	1.000	—	—	—	—	—	—
	(0.007)	(—)						
ρ_{SNOi}	0.518	0.331	1.000	—	—	—	—	—
	(0.005)	(0.018)	(—)					
ρ_{SSGi}	0.362	0.360	0.527	1.000	—	—	—	—
	(0.003)	(0.007)	(0.013)	(—)				
ρ_{SSAi}	0.288	0.251	0.436	0.337	1.000	—	—	—
	(0.007)	(0.009)	(0.011)	(0.019)	(—)			
ρ_{SSKi}	0.521	0.348	0.816	0.523	0.425	1.000	—	—
	(0.006)	(0.023)	(0.012)	(0.008)	(0.024)	(—)		
ρ_{SSFi}	0.526	0.286	0.797	0.534	0.368	0.777	1.000	—
	(0.008)	(0.021)	(0.007)	(0.012)	(0.014)	(0.028)	(—)	
ρ_{SGBi}	0.479	0.332	0.692	0.495	0.390	0.676	0.684	1.000
	(0.004)	(0.017)	(0.008)	(0.013)	(0.021)	(0.013)	(0.020)	(—)

Notes: Table reports maximum likelihood estimates for a multivariate GARCH model of constant conditional correlations where $h_{it} = \omega_i + \alpha_{i1}\epsilon^2_{i(t-1)} + \beta_{i1}h_{i(t-1)}$ and where $h_{ijt} = \rho_{ij}\sqrt{h_{iit}h_{jjt}}$. The upper panel reports parameters in the time varying conditional variances. The lower panel reports constant conditional correlations. Asymptotic standard errors are in parenthases. Abbreviations: Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB). Exchange rates priced in terms of home currency per US dollars.

to be characterised by distributions that have fat tails. That is, measures of kurtosis tend to be high. However, when using *weekly* data, the assumption of conditional normality tends to be more reasonable.[25] Coefficients of kurtosis tend to be much smaller.[26] Weekly data is, indeed, employed in this study and, as such, it is considered acceptable here to assume conditional normality.

Turning to the estimates in Table 4.10 and Table 4.11, all the parameters in the time-varying conditional variances (that is, ω_i, α_{i1} and β_{i1}) are significant at the 5% significance level. Similarly, estimates for the conditional correlations (ρ_{ij}) are all highly significant. The price of gold is negatively correlated with all US dollar exchange rate pairs. Meanwhile, all exchange rate pairs are positively correlated with each other.

Estimates of the unconditional variances, $\hat{\omega}_i(1 - \hat{\alpha}_{i1} - \hat{\beta}_{i1})^{-1}$, for gold-price returns and for all US dollar exchange-rate returns in the sample show that unconditional variances are lower for nearly all the exchange-rate pairs than for gold. The unconditional variance for gold price returns (6.28) is roughly three times as large as for the returns on most US dollar exchange-rate pairs in the sample.

Table 4.10 also shows that while gold price returns are correlated negatively with the returns on all US dollar exchange-rate pairs during the sample period, the strongest negative correlation is with returns on US dollars priced in terms of Australian dollars (-0.341). Gold price returns are also correlated strongly with returns on US dollars priced in terms of Danish kroner (-0.320), Swiss francs (-0.317) and euros (-0.309).

A negative conditional correlation between the price of gold and the US dollar would, if constant and stable, support an argument in favour of the relationship between gold and the dollar being driven by a *numeraire effect*, whereby a drop in the value of the dollar is associated with a rise in the price of gold simply because gold is priced in dollars. If conditional correlations are, indeed, constant and stable, then a numeraire effect would seem plausible. The presence of a numeraire effect would imply an absence of pricing-to-market behaviour among sellers of gold and would, instead, imply a role for pass-through pricing. For aggregate prices, existing evidence suggests that pass-through effects are small: prices of tradeable goods tend to respond incompletely to variations in exchange rates.[27] However, gold is unusual in its characteristics as a tradeable commodity. It is homogenous and highly liquid. The presence of a durable nu-

[25] See Baillie and Bollerslev (1989b).

[26] Indeed, Enders (2004) notes that it is acceptable to ignore the issue of fat-tailed distributions when the sample is large. Quasi-maximum likelihood estimates use the normal distribution even though the actual distribution of the residuals is fat-tailed. The reason is that under weak assumptions the parameter estimates for the model of the mean and the conditional variance are consistent and normally distributed.

[27] See for instance Berman et al. (2009) and Campa and Goldberg (2005).

meraire effect linking gold-price returns and US dollar returns is not, therefore, inconceivable.

Summarising the results from this section, the analysis above presents cautious evidence in favour of the multivariate GARCH(1,1) model in Eqn. (4.42) with constant conditional correlations as offering a reasonable representation of the short-run dynamics of gold-price returns and exchange-rate returns for those exchange rates in the sample. However, the analysis above suffers from clear limitations. To assume that conditional correlations are constant throughout the sample period may be at best misleading and at worst erroneous. Interrelationships may have changed. Correlations may have strengthened. Or weakened. Indeed, conditional correlations between gold-price returns and US dollar returns, despite appearing to be characterised, in the results above, by a negative correlation, may have experienced periods of positive correlation during the sample period. Only an analysis of the *dynamics* of the conditional correlations can throw light on these issues. The next section, Section 4.5.2, offers an analysis of dynamic conditional correlations.

4.5.2 Results: Dynamic conditional correlations

This subsection presents results from a dynamic model of conditional correlations of the asset returns described in Section 4.4. The model offers an accurate, tractable way of modelling, simultaneously, both time-varying conditional volatilities and time-varying conditional correlations.

Recall that, in Section 4.4, preliminary data analysis finds evidence of conditional heteroscedasticity which lends support to the use of ARCH-type models and, in particular, to the use of GARCH(p, q) models to capture the volatility behaviour of the data series. Analysis in Section 4.4 suggests that, here, the most appropriate form of GARCH(p, q) model is the GARCH(1,1) model.

The multivariate GARCH model of dynamic conditional correlations is to be estimated using maximum likelihood. However, since the series in our dataset are consistently non-normal (see Section 4.4), the remedy here is to use the quasi maximum likelihood method (Bollerslev et al., 1988) in order to generate consistent standard errors that are robust to non-normality.

A comparison of the loglikelihood values among alternative lag specifications suggests that our data are best captured by a DCC(1,1) with each of the conditional variances captured by a univariate GARCH(1,1) model.

Table 4.12 displays estimation results for the DCC(1,1)-GARCH(1,1) model for the 17 asset prices under analysis: the price of gold, the US dollar against the euro, the yen, the Indian rupee, the Australian dollar, the Canadian dollar, the Danish krone, the Israeli shekel, the Maltese lira, the New Zealand dollar,

Table 4.12: DCC-GARCH model estimation results, 1986-08

i	**GLD**	**SXEU**	**SJP**	**SINDIA**	**STW**	**SAU**
GARCH parameters						
ω_i	0.063	0.038	0.138	0.020	0.059	0.044
	(0.039)	(0.023)	(0.055)	(0.015)	(0.026)	(0.044)
α_{i1}	0.127	0.058	0.101	0.351	0.545	0.060
	(0.050)	(0.021)	(0.037)	(0.049)	(0.231)	(0.033)
β_{i1}	0.863	0.913	0.807	0.649	0.455	0.907
	(0.039)	(0.033)	(0.057)	(0.090)	(0.141)	(0.062)

i	**SCA**	**SDK**	**SIS**	**SMA**	**SNZ**	**SNO**
GARCH parameters						
ω_i	0.008	0.041	0.023	0.014	0.003	0.078
	(0.003)	(0.025)	(0.013)	(0.013)	(0.003)	(0.035)
α_{i1}	0.121	0.062	0.039	0.041	0.039	0.111
	(0.027)	(0.021)	(0.024)	(0.025)	(0.012)	(0.033)
β_{i1}	0.867	0.907	0.945	0.951	0.960	0.837
	(0.025)	(0.035)	(0.024)	(0.029)	(0.012)	(0.044)

i	**SSG**	**SSA**	**SSK**	**SSF**	**SGB**
GARCH parameters					
ω_i	0.009	0.063	0.042	0.020	0.015
	(0.005)	(0.033)	(0.019)	(0.019)	(0.035)
α_{i1}	0.104	0.211	0.067	0.023	0.039
	(0.034)	(0.068)	(0.021)	(0.021)	(0.024)
β_{i1}	0.864	0.776	0.902	0.965	0.949
	(0.043)	(0.059)	(0.028)	(0.044)	(0.032)

DCC parameters	
a_1	0.010
	(0.000)
b_1	0.988
	(0.013)
Diagnostics	
$\chi^2 - test : R_t = R$	317.085
	(0.000)
Log-likelihood	-13,108

Notes: Parameter estimates are based on the DCC-GARCH model: $h_{it} = \omega_i + \alpha_{i1}\epsilon^2_{i(t-1)} + \beta_{i1}h_{i(t-1)}$ and $Q_t = (1 - a_1 - b_1)\bar{Q} + a_1(\varepsilon_{t-1}\varepsilon'_{t-1}) + b_1 Q_{t-1}$. Probability values ($p$-values) are in parenthases. All estimation is undertaken with MATLAB using the author's proprietary code. Abbreviations: Gold (GLD), Euro (SXEU), Yen (SJP), Indian rupee (SINDIA), Taiwan dollar (STW), Australian dollar (SAU), Canadian dollar (SCA), Danish krone (SDK), Israeli Shekel (SIS), Maltese lira (SMA), New Zealand dollar (SNZ), Norwegian krone (SNO), Singapore dollar (SSG), South African rand (SSA), Swedish krona (SSK), Swiss franc (SSF), UK pound (SGB).

the Norwegian krone, the Singapore dollar, the South African rand, the Swedish krona, the Swiss franc and the pound sterling. Probability values reflect t-stats calculated with robust standard errors.

First, note from Table 4.12 that the GARCH parameters estimated under the assumption of dynamic conditional correlation are identical to those estimated for constant conditional correlation: the first-stage of the estimation process, estimating the univariate GARCH models, is identical for both models. All univariate GARCH processes show a high degree of persistence. That is, the sums of α_i and β_i are all close to one. Meanwhile, the estimated DCC parameters, a_1 and b_1, imply a highly persistent correlation, with a half-life innovation of six years.[28] However, results of a test of parameter constancy indicate strong

[28]Half-life is defined as the time it takes for a shock to correlation to reduce by half. Half-life

evidence against the assumption of constant conditional correlations: the test, developed by Engle and Sheppard (2001), uses a χ^2-statistic to test the null of $R_t = R$. The resulting test statistic, 317.1, is highly significant, rejecting the null hypothesis of constant conditional correlations.

Figure 4.4 and Figure 4.5 show the evolution over time of ρ_{1i}, the dynamic conditional correlation between the price of gold and the 16 major exchange rate pairs in the sample. The sign of all correlation coefficients over the sample period is consistently negative. That is, there is a negative relationship between gold-price returns and US dollar returns. Further, most if not all of the correlation coefficients in the sample grow in magnitude from the early 1990s onwards, becoming increasingly negative, reaching their most negative at the end of 2008. The negative relationship between the price of gold and the dollar's value in terms of euros is particularly well-defined, with ρ_{12} reaching -0.6 in August 2008.

Figure 4.4 and Figure 4.5 also plot the bootstrapped 95% confidence intervals for the estimated constant conditional correlations. The dynamic conditional correlations vary widely, and for large periods of time stray beyond the limits of the estimated confidence bounds. By implication, it is not descriptively useful to assume that the conditional correlations between gold-price returns and US dollar returns are constant. The conditional correlations display a great deal of time-variation. Significantly, for nearly all currencies, the dynamic conditional correlations turn increasingly negative during the final six years of the sample period and remain outside the confidence intervals for the entirety of this period.

If the only link between the price of gold and the US dollar is a numeraire effect—with a weak dollar implying, by pricing convention, more dollars per troy ounce of gold—then we would expect the conditional correlation to be stable. We would not expect a sharp change in magnitude. But Figure 4.4 and Figure 4.5 suggest that the negative relationship between gold returns and US dollar returns has grown stronger during the past decade and a half, becoming particularly acute in the last five years.

Figure 4.4 and Figure 4.5 also offer an insight into whether gold has, during the sample period, behaved as a *hedge* against the dollar or whether it has behaved as a *haven*.

Recall the definitions of hedge and haven. An asset that functions as a haven for another asset will not co-move with the other asset *in times of stress*. That is, an asset acts as a haven if it is uncorrelated or correlated negatively with another asset that is experiencing sustained losses. Meanwhile, an asset that acts as a hedge is one that is uncorrelated or correlated negatively with another asset *on average*. Note the difference: An asset that functions as a haven is

is computed as $\ln(0.5) + \ln(a_1 + b_1)$.

Figure 4.4: Dynamic conditional correlation: gold and exchange-rate returns

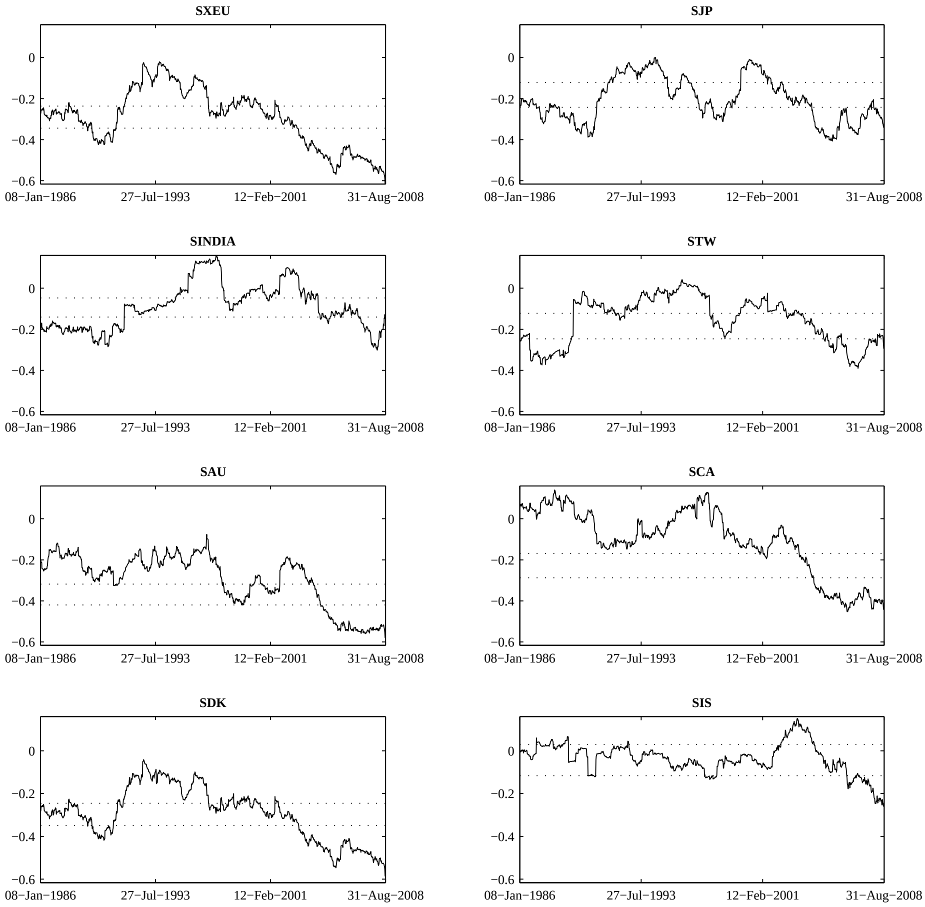

Notes: Figure shows the dynamic conditional correlation between innovations in the price of gold and eight major exchange rate pairs: US dollar against the euro (SXEU), the yen (SJP), the Indian rupee (SINDIA), the Taiwan dollar (STW), the Australian dollar (SAU), the Canadian dollar (SCA), the Danish krone (SDK) and the Israeli shekel (SIS). The dashed lines are the confidence bands (bootstrapped) for the estimated constant conditional correlations. Returns are percentage demeaned nominal currency returns. Exchange rates expressed as home currency per US dollar. Frequency is weekly. The dynamic conditional correlations have a wide range and are often outside the confidence bands resulting from the estimated constant conditional correlations.

Figure 4.5: Dynamic conditional correlation: gold and exchange-rate returns

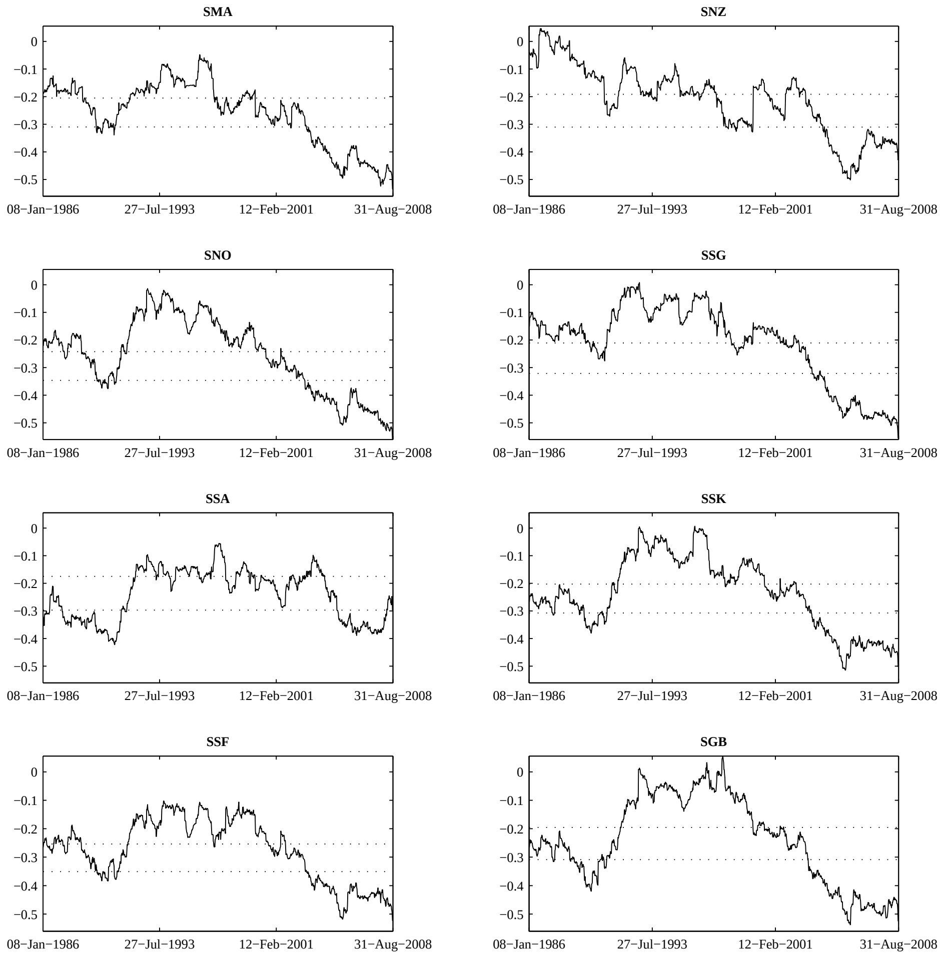

Notes: Figure shows the dynamic conditional correlation between innovations in the price of gold and eight major exchange rate pairs: US dollar against the Maltese lira (SMA), New Zealand dollar (SNZ), the Norwegian krone (SNO), the Singapore dollar (SSG), the South African rand (SSA), the Swedish krona (SSK), the Swiss franc (SSF) and the pound sterling (SGB). The dashed lines are the confidence bands (bootstrapped) for the estimated constant conditional correlations. Returns are percentage demeaned nominal currency returns. Exchange rates expressed as home currency per US dollar. Frequency is weekly. The dynamic conditional correlations have a wide range and are often outside the confidence bands resulting from the estimated constant conditional correlations.

uncorrelated or negatively correlated with another asset in times of stress only, and not necessarily on average.

Using these definitions, Figure 4.4 and Figure 4.5 show that gold has acted as a hedge against the US dollar throughout the sample period. That is, gold-price returns have been correlated negatively with US dollar returns, for all 16 exchange-rate pairs, not only in times of stress but also on average throughout the 22-year sample period.

Supporting evidence is offered in Table 4.13. The table gives maximum likelihood estimates of the constant conditional correlations between gold-price returns and returns for the 16 US dollar exchange-rate pairings. All are negative and significant and are consistent with the hypothesis that gold provides an effective *hedge* against the US dollar. The most consistently negative relationships are between gold-price returns and US dollar returns in terms of euros, Swiss francs, Australian dollars and Danish kroner.

The third, fourth and fifth columns in Table 4.13 show mean dynamic correlations during periods of market stress, defined according to the 10%, 5% and 1% quantiles of most negative exchange-rate returns. The smaller the size of the quantile the more extreme the market stress. The quantile correlations define the extent to which gold acts as a *safe haven* from US dollar volatility. That is, for any given US dollar exchange-rate pair, if gold acts as an effective safe haven, then quantile correlations will be more negative than the corresponding constant correlations. Or they will be uncorrelated. Table 4.13 shows that, in fact, neither is true. For most of the US dollar exchange-rate pairings the quantile correlations are less negative, not more negative, than the constant conditional correlations. The yen and the UK pound are exceptions at the 1% quantile. However, the difference is not statistically significant: quantile correlations for the yen and UK pound are more negative by less than two asymptotic standard errors. All of this suggests that gold's role as a safe haven from US dollar movements is negligible.[29] Gold's only effective role, in terms of offering investment protection from movements in the US dollar, is as a hedge.

Indeed, Figure 4.4 and Figure 4.5 show that since 2001 gold's efficacy as a hedge has become more pronounced. The negative conditional correlation between gold-price returns and US dollar returns has grown increasingly strong. The reason why is unclear.[30] This period does coincide with a steady downward

[29] Baur and McDermott (2010) similarly find no consistent role for gold as a safe haven from share-price movements for data of the same frequency, ie, weekly. They do find evidence in favour of gold's role as safe haven for daily data, but the evidence is partial (supported only at the 1% quantile) and economically trivial (marginal effects are very small). Baur and Lucey (2006) find that gold is not a safe haven for bonds. For stocks, the only economically meaningful role gold plays as a safe haven is, they find, for the UK. For other markets the marginal effects are small.

[30] Potential explanations based on the increasing role of the derivatives market during the

Table 4.13: Constant conditional correlations and quantiles

	Const. corr.	**10% quantile**	**5% quantile**	**1% quantile**
US dollar exchange rate				
Euro	-0.309	-0.284	-0.269	-0.252
	(0.010)	(0.125)	(0.110)	(0.113)
Yen	-0.214	-0.210	-0.218	-0.219
	(0.008)	(0.092)	(0.087)	(0.095)
Indian rupee	-0.087	-0.132	-0.123	-0.290
	(0.020)	(0.097)	(0.102)	(0.133)
Taiwan dollar	-0.155	-0.217	-0.230	-0.243
	(0.054)	(0.107)	(0.106)	(0.093)
Australian dollar	-0.341	-0.337	-0.351	-0.364
	(0.020)	(0.127)	(0.125)	(0.111)
Canadian dollar	-0.159	-0.218	-0.246	-0.270
	(0.008)	(0.165)	(0.160)	(0.188)
Danish krone	-0.320	-0.280	-0.266	-0.241
	(0.012)	(0.110)	(0.100)	(0.107)
Israeli shekel	-0.057	-0.056	-0.056	-0.096
	(0.016)	(0.103)	(0.117)	(0.105)
Maltese lira	-0.287	-0.260	-0.242	-0.249
	(0.007)	(0.102)	(0.091)	(0.123)
New Zealand dollar	-0.256	-0.223	-0.202	-0.173
	(0.027)	(0.141)	(0.142)	(0.149)
Norwegian krone	-0.300	-0.281	-0.277	-0.255
	(0.017)	(0.138)	(0.137)	(0.139)
Singapore dollar	-0.261	-0.232	-0.222	-0.204
	(0.020)	(0.139)	(0.121)	(0.062)
South African rand	-0.257	-0.248	-0.240	-0.241
	(0.019)	(0.087)	(0.083)	(0.067)
Swedish krona	-0.268	-0.233	-0.221	-0.182
	(0.021)	(0.135)	(0.125)	(0.130)
Swiss franc	-0.317	-0.258	-0.256	-0.219
	(0.047)	(0.095)	(0.095)	(0.084)
UK pound	-0.262	-0.249	-0.248	-0.280
	(0.026)	(0.144)	(0.121)	(0.112)

Notes: Table shows constant conditional correlations (Const. corr.) for gold returns versus the returns of 16 US dollar exchange rates pairings estimated over the full sample period (10 January 1986 to 29 August 2008); table also shows mean dynamic conditional correlations for selected quantiles (10%, 5%, 1%) of the most negative exchange-rate returns. Asymptotic standard errors are in parenthases for constant conditional correlations. Quantile standard deviations are in parenthases for mean dynamic conditional correlations.

spiral in the value of the US dollar. But other periods of dollar depreciation have not gone hand in hand with strengthening negative correlations with gold. Figure 4.4 and Figure 4.5 show that between 1985 and 1988, when the US dollar lost 12% of its trade-weighted value, the correlation between gold returns and US dollar returns did not turn increasingly negative.

4.6 Conclusions

This study investigates the nature of the relationship between the price of gold and the US dollar, how it has changed during the past 25 years, and how these changes cast light upon the role gold plays as an investment hedge and a haven. Empirical results based on a multivariate GARCH model of dynamic conditional correlations show that the conditional correlation between changes in the price of gold and changes in the US dollar's exchange rate is broadly negative. That is, increases in the price of gold tend to be associated with decreases in the value of the US dollar. This correlation has not, however, remained constant over time. During the past 7 years the correlation has turned increasingly negative. In 2008 it was more negative than at any point during the past three decades. The implication is that gold's role as an investment *hedge* against the US dollar is much stronger and more durable than suggested by Capie et al. (2005).

Analysis of gold's role as a *safe haven* provides very different conclusions. Quantile correlations show gold does not act as an effective safe haven from market stress. These results chime with those of Baur and McDermott (2010), who find no evidence that gold acts as a consistent safe haven with respect to weekly movements in international share prices. Baur and Lucey (2006) find no evidence that gold acts as a safe haven for bonds.

Given these findings, identifying the factors that have contributed to gold's strengthening role as a hedge against the US dollar offers plenty of scope for further research.

1990s (Kearney and Lombra, 2008) or feedback trading (Campbell and Kyle, 1993; Cutler et al., 1990; Delong et al., 1990; Kirman, 1993; Shleifer, 2000) are beyond the scope of this paper.

Chapter 5

Conclusion

This chapter concludes the thesis, offering some discussion, implications for policy and outlining scope for future research.

5.1 Discussion

This section gathers together the results from the three analytical studies that form the core of this thesis and offers some broad conclusions.

In the last two decades there have been a number of important developments in the field of exchange-rate economics, with major steps forward taken in both the empirical and theoretical treatment of exchange-rate determination. Econometrics has been a significant driving force behind this progress. As has greater access to high quality data. Consequently, the academic literature on exchange-rate determination has proliferated. However, while the academic world has done much to improve our understanding of exchange rates and exchange-rate movements, many questions remain unanswered. This thesis examines three unresolved issues.

Chapter 2 addresses the question of misalignment and policy response. How do policymakers respond to exchange-rate misalignment? In particular, how is intervention policy designed? What are the aims? The focus here is on official Japanese intervention in the currency markets. Does Japan intervene in a manner that suggests it pursues a currency target? If so, what type of target? Fixed? Time-varying? Does the target reflect an assessment of exchange-rate equilibrium? The study presented in Chapter 2 addresses these questions by estimating an intervention reaction function and testing the hypothesis that Japan intervenes in order to guide the exchange rate towards equilibrium. Equilibrium is defined as being consistent with a capital-enhanced version of purchasing power parity. Estimation results show that Japan's interventions do indeed suggest the pursuit of a currency target that is compatible with an augmented form purchasing power parity. Further, results suggest that the monetary authorities are not equally tolerant towards currency weakness and currency strength. Monetary authorities are more responsive to a strong currency—strong, that is, relative to its equilibrium value.

In addition to these findings, which highlight the continuing importance of partial-equilibrium models in shaping modern exchange-rate policy, Chapter 2 also offers an advance in the empirical treatment of intervention policy. The advance is to discard the standard assumption of proportional odds that is embedded in the ordered logit model of intervention. To assume proportional odds is to assume that the determinants of intervention carry identical weight in the policy process whatever the type of intervention under consideration: sale of domestic currency, purchase, or abstinence. Weights, however, may not be identical. In practice they may vary. To allow for this a flexible approach is proposed in the form of a generalised ordered logit model, which permits misalignment to affect intervention differently depending on the type of intervention. This empirical innovation turns out to be critical: results show that the monetary authorities do indeed react differently to misalignment depending on the type of intervention under consideration.

The next issue to be addressed in this thesis, in Chapter 3, is market response to exchange-rate misalignment. Exchange-rate misalignment is defined in terms of deviation from covered interest-rate parity and uncovered interest-rate parity, two of the oldest conceptual pillars of the neoclassical approach to global capital flows and exchange-rate economics. The aim, in Chapter 3, is to measure market response to deviations from these parity conditions and the analysis focuses on one specific form of market response: the issuance of international bonds. How does the issuance of international bonds respond to exchange-rate misalignment? In particular, does misalignment affect the choice of issuance currency? The analytical study presented in Chapter 3 addresses these questions by constructing a utility-consistent model of currency choice, focussing on the number, rather than the value, of bonds issued, and by drawing on a large, unique dataset of international debt securities. Results show that while deviations from swap-covered interest-rate parity do exist, issuers do not seem to respond to them. However, deviations from uncovered interest-rate parity do affect the choice of issuance currency. Issuers, on aggregate, issue more bonds in currencies that are associated with low and falling interest rates. Further, financial institutions are particularly responsive to deviations from uncovered interest-rate parity.

Chapter 4 addresses the final research question tackled in this thesis: what do exchange-rate dynamics reveal about hedging behaviour in the face of currency risk? Specifically, does gold act as an effective hedge against changes in the value of the US dollar? As an effective safe haven? As neither? Definitions here are important. If an investor holds a given asset γ, then a safe haven is defined to be any asset that is either uncorrelated or correlated negatively with γ in times of market stress. A hedge is defined to be any asset that is either uncorrelated or correlated negatively with γ on average. Studying weekly data for 16 major dollar-paired exchange rates, Chapter 4 shows that for the past 23 years gold has behaved as a consistent and effective hedge against the US dollar. However, the evidence suggests that gold does not provide an effective safe haven: outside periods of market stress, gold remains correlated with the US dollar.

The studies contained in Chapter 2, Chapter 3 and Chapter 4 offer findings that are well-defined and actionable. But more importantly they have clear implications for economic policy. The next section discusses these implications.

5.2 Policy implications

This section discusses the policy implications stemming from the three analytical studies that make up this thesis.

Studying official Japanese intervention in the currency markets yields a number of policy-relevant results. The first result is that currency targets matter: Japan's intervention policy since 1991 has been consistent with the active pursuit of an exchange-rate target. Interventions are timed and directed in a manner that suggests Japan's monetary authorities aim to keep the national currency, the yen, close to a time-varying target. Other central banks that have active intervention policies may follow similar strategies. This can be verified empirically. What is important is that intervention policy in Japan, either explicitly or implicitly, is compatible with the pursuit of an exchange-rate target and this result can and should be explored elsewhere, for other central banks. Indeed, it is more than possible, given the globalised nature of international financial markets and evidence elsewhere of coordination among central banks, that other monetary authorities do pursue similar strategies based on currency targets.[1] Of course, the findings here should not be taken as proof of *effectiveness*. The effectiveness of an intervention policy based on the pursuit of a time-varying currency target is beyond the scope of this thesis. Here, the scope is only to model behaviour. Intervention may or may not be effective depending on factors such frequency, surprise, coordination and credibility.

Japanese intervention in the currency markets can, this thesis shows, be

[1] For discussion of coordinated intervention and the coordination channel of intervention effectiveness, see Sarno and Taylor (2001), Taylor (2005) and Reitz and Taylor (2008).

modelled adequately as a process whereby the central bank seeks to minimise deviations, both positive and negative, from a currency target. The precise nature of this target is important. It is not a fixed target. Nor is it a volatility target. Chapter 2 shows that Japan intervenes in the currency markets in order to drive the exchange rate towards a target that reflects an assessment of partial equilibrium. The central bank's assessment is consistent with a partial equilibrium model allowing for persistence in both the real exchange rate and interest-rate differentials. The critical policy implication here is that, as the experience of Japan shows, it is possible, operationally, for a central bank to make *daily* calculations of exchange-rate equilibrium that are transparent, tractable and that convey information that is sufficient and rich enough to provide an adequate platform for a daily intervention policy.

Chapter 3 investigates choice of issuance currency in the issuance of international bonds and the findings offer a number of policy implications. The first is that nominal interest rates have a significant effect on the choice of issuance currency. Low interest-rate currencies attract more issuance despite the fact that standard interest-rate parity conditions suggest there ought to be nothing to gain, in terms of borrowing-cost savings, from issuing debt in low-interest-rate currencies. The empirical evidence in Chapter 3 shows that low interest rates exert a positive influence on issuance. What this entails for policy is straightforward. Changes in interest rates can, at the margin, have a significant effect on the use of the national currencies in international transactions. Further, monetary policy and the term structure of domestic interest rates can affect the internationalisation of world currencies. The opposite is also true: a policy of non-internationalisation of the domestic currency, such as that pursued by Singapore, can be either undermined or strengthened by domestic interest-rate policy.

The question of currency internationalisation is important for many countries and for many existing and fledgling monetary unions. A country or monetary union that allows international debt to be denominated in its own currency generates a series of potential welfare gains. First, it experiences a welfare gain because total demand for its securities will increase due to foreign demand, and the return for holding these securities will fall. There will also be a related welfare gain for the rest of the world in the form of an increase in choice of securities to invest in. Second, there will be a general welfare gain as a result of the expansion of the pool of investors. A bigger pool of investors will increase trade in the securities of the issuance-currency country, boosting liquidity and reducing the impact of demand shocks on prices. Third, for the issuance-currency country, an increase in the use of its currency in international transactions will expand the size of its foreign-exchange market, cutting transaction costs involved in trade in both goods and assets.

Chapter 4 asks, what do exchange-rate dynamics reveal about hedging behaviour in the face of currency risk? The findings throw up a number of policy implications. Perhaps the most important is that gold acts as hedge against changes in the value of the US dollar and, in as much as current thinking in financial economics suggests that the availability of hedging instruments increases welfare and reduces market volatility, then the use of gold as a hedge should not be restricted. Hedging practices should be allowed to continue. This recommendation does, however, come with one caveat. In recent years gold has become increasingly correlated (negatively) with the US dollar, more so than at any point in the past two and a half decades. This could, conceivably, represent a source of systemic risk. Systemic risk is present whenever a wide range of assets become highly correlated. Systemic risk is also associated with infrequent events. But the high conditional correlation between exchange-rate returns and gold returns revealed in Chapter 4 is not infrequent. It is persistent. The risk, therefore, is not systemic, but it may, to some extent, be *systematic*, where systematic risk refers to correlation between assets (and a common factor) with no requirement that the correlated changes be infrequent.[2] To this extent, the tight relationship between the price of gold and the US dollar should be monitored closely and future research should be directed towards investigating its risk characteristics.

5.3 Future research

This section offers some thoughts on the extent to which future research can build on this thesis and draw on its findings.

The findings presented here on intervention policy offer a number of avenues for further research. One avenue is to investigate the extent to which implicit currency targets play a role in intervention policy for other developed nations. To what extent is the experience of Japan mimicked elsewhere? Do assessments of partial equilibrium drive intervention policy in other central banks? If so, this would suggest that basic economic aggregates, or *fundamentals* as they are often called, can and do play a bigger role in exchange-rate policy formation than is commonly assumed. Another intervention issue that offers scope for future research is that of effectiveness. For nations that pursue a currency target in their intervention policy, to what extent can policy be shown to be effective? That is, to what extent does intervention succeed in driving the exchange rate towards its target, as intended? The existing literature on intervention effectiveness is

[2] For further discussion regarding systemic risk and systematic risk, see, for instance, the survey by Bandt and Hartmann (2000) and also Das and Uppal (2004).

already large.[3] But no attempt has been made to explore the effectiveness of a policy based on implicit currency targeting.

One of the key findings in this thesis is that the incentives to intervene in the currency markets can change depending on the type of intervention under consideration: intervention to buy foreign currency, to sell foreign currency or to abstain from intervention in the current period. Future research should explore alternative theoretical frameworks for such behaviour. That is, future research should aim to explore the nature of the nonlinearities in the intervention reaction function. How, exactly, do these asymmetric preferences manifest themselves? A suitable starting point for any future research would be to draw from the growing literature on nonlinearities in monetary policy reaction functions. See for instance Orphanides and Wieland (2000), Ruge-Murcia (2004), Dolado et al. (2002) and Cukierman and Muscatelli (2002).

Of the findings presented in Chapter 3 on currency choice in the issuance of international bonds, one area that offers particular scope for further research is the finding that, while there is no evidence that bond issuers respond to deviations from long-term covered interest rate parity, quantitatively these deviations do exist. Future research should look further into these unexploited arbitrage opportunities. For how long do they remain unexploited? Why is round-trip arbitrage so costly as to cause these arbitrage opportunities to remain unexploited? Perhaps a more sophisticated description of issuance behaviour is required to explain why these cost-saving opportunities are not arbitraged away rapidly. The introduction of dynamics may help. Empirically dynamics can be introduced with a linear feedback model, allowing for a generalisation of the Poisson model to an autoregressive process. Windmeijer (2006) provides a useful overview of the literature on dynamics in panel count data models.

Another avenue for further research is to build on the findings in Chapter 3 by asking, what causes borrowers to expose themselves to currency risk by denominating their debt, unhedged, in foreign currency? In emerging markets it is argued that incomplete financial markets limit the ability of firms to hedge their foreign-currency exposure. Meanwhile, currency pegs give an implicit guarantee against short-term movements in the exchange rate. But analysis of unhedged foreign borrowing under flexible exchange rates remains scarce. McKinnon and Pill (1999) and Burnside et al. (1999) analyse the desire of banks to take on unhedged foreign debt. They assume that this desire is a consequence of moral hazard arising from deposit insurance and other bailout guarantees. However, this does little to explain the desire by the *non-bank* sector to issue unhedged foreign debt. The aim should be to explain how currency-risk premiums arise en-

[3] For a survey of the literature see Neely (2005a). Also for recent treatments see Neely (2005b), Fatum and Hutchison (2003), and Reitz and Taylor (2008).

dogenously from the interest-rate differentials between countries and how these premiums act as an incentive for borrowers to denominate their debt, unhedged, in foreign currency.

Findings from the work in this thesis on hedging and exchange-rate dynamics offer substantial scope for further research. Secondary to the key results in Chapter 4 is the finding that in recent years gold returns have become increasingly correlated (negatively) with returns on the US dollar, more so than at any point in the past two and a half decades. This trend demands further analysis. One useful approach would be to investigate the role of herd behaviour. Rational herd behaviour theories suggest that investors act in a herd-like manner either because they receive similar or correlated information or because they infer, rationally, information from the actions of other investors.[4] Statistical measures of herd behaviour gauge the average tendency of a group of money managers to buy or sell the same assets at the same time. Measures such as these could, potentially, be employed as time series in models of dynamic conditional correlations designed to model the dependencies between herding, returns and correlations for commodity prices, like gold, and US dollar exchange-rate pairs. The models of conditional correlations presented in Chapter 4 offer a suitable, empirical starting point.

While this thesis offers plenty of scope for analytical extensions and further research, the intention is that in itself, this thesis represents a useful contribution to the fields of applied econometrics and international finance.

[4]See Froot et al. (1990), Hirshleifer et al. (1994) and Bikhchandani et al. (1992).

Bibliography

Aabo, T. (2006). The importance of corporate foreign debt in managing exchange rate exposures in non-financial companies. *European Financial Management 12*(4), 633–649.

Allayanis, G. and E. Ofek (2001). Exchange rate exposure, hedging and the use of foreign currency derivatives. *Journal of International Money and Finance 20*, 273–296.

Allayanis, G. and J. Weston (2001). The use of foreign-currency derivatives and firm market value. *Review of Financial Studies 14*, 243–276.

Allayannis, G., G. Brown, and L. Klapper (2003). Capital structure and financial risk: evidence from foreign debt use in east asia. *Journal of Finance 58*, 2667–2709.

Allayannis, G., J. Ihrig, and J. Weston (2001). Exchange-rate hedging: Financial versus operational strategies. *American Economic Review 91*(2), 391–395.

Allayannis, G. and E. Ofek (2001). Exchange rate exposure, hedging, and the use of foreign currency derivatives. *Journal of International Money and Finance 20*(2), 273–296.

Almekinders, G. and S. Eijffinger (1996). A friction model of daily bundesbank and federal reserve intervention. *Journal of Banking and Finance 20*(8), 1365–1380.

Andersen, T. and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review 39*(4), 885–905.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association 96*, 42–55.

Andersen, T., T. Bollerslev, F. Diebold, and C. Vega (2007). Real time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics 73*(2), 251–277.

Ang, A. and J. Chen (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics 3*(63), 443–494.

Artus, J. (1977). Exchange-rate stability and managed floating: the experience of the federal republic of germany. *IMF Staff Papers 23*, 312–33.

Baba, Y., R. Engle, D. Kraft, and K. Kroner (1989). Multivariate simultaneous generalised arch. *UCSD Department of Economics Manuscript*.

Baillie, R. and T. Bollerslev (1989a). The message in daily exchange rates: A conditional-variance tale. *Journal of Business and Economic Statistics 7*(3), 297–305.

Baillie, R. and T. Bollerslev (1989b). The message in daily exchange rates: a conditional variance tale. *Journal of Business and Economic Statistics 7*, 297–305.

Baillie, R. and W. Osterberg (1997). Why do central banks intervene. *Journal of International Money and Finance 16*(6), 909–919.

Balke, N. and J. Haslag (1992). A theory of fed watching in a macroeconomic policy game. *International Economic Review 33*(3), 619–28.

Bandt, O. and P. Hartmann (2000). Systemic risk: a survey. *European Central Bank Working Paper Series* (35).

Baur, D. and B. Lucey (2006). Is gold a hedge or a safe haven? an analysis of stocks, bonds and gold. *Institute for International Integration Studies Discussion Paper* (198).

Baur, D. and T. McDermott (2010). Is gold a safe haven? international evidence. *Journal of Banking and Finance*, 1886–1898.

Bauwens, L., S. Laurent, and J. Rombouts (2006). Multivariate garch models: a survey. *Journal of Applied Econometrics 21*, 79–109.

Beck, N. and J. Katz (1995). What to do (and not to do) with time-series cross-section data. *American Political Science Review 89*, 634–647.

Berman, N., P. Martin, and T. Mayer (2009). How do different exporters respond to exchange rate changes? theory, empirics and aggregate implications. *CEPR Discussion Paper 7493*.

Bernanke, B. and M. Gertler (1999). Monetary policy and asset price volatility. *NBER Working Paper 7559*.

Berndt, E., B. Hall, R. Hall, and J. Hausman (1974). Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement 3*, 653–65.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Bollerslev, T. (1986). Generalised autoregressive conditional heteroscedasticity. *Journal of Econometrics 31*, 307–27.

Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalised arch model. *Review of Economics and Statistics 72*, 498–505.

Bollerslev, T., R. Engle, and J. Wooldridge (1988). A capital asset pricing model with time varying covariances. *Journal of Political Economy 96*, 116–131.

Bollerslev, T. and E. Ghysels (1996). Periodic autogregressive conditional heteroscedasticity. *Journal of Business and Economic Statistics 14*, 139–151.

Bomfim, A. and G. Rudebusch (1997). Opportunistic and deliberate disinflation under imperfect credibility. *Journal of Money, Credit and Banking 32*(4), 707–721.

Bos, C. and P. Gould (2007). Dynamic correlations and optimal hedge ratios. *Tinbergen Institute Discussion Papers* (07-025/4).

Boyer, B., S. Michael, and M. Loretan (1997). Pitfalls in tests for changes in correlations. *Federal Reserve System International Finance Discussion Papers 597*.

Brandt, M., J. Cochrane, and P. Santa-Clara (2001). International risk sharing is better than you think (or exchange rates are much too smooth). *National Bureau of Economic Research Working Papers* (8404).

Brant, R. (1990). Assessing proportionality in the proportional odds model for the ordered logistic regression. *Biometrics 46*(4), 1171–1178.

Buiter, W. and P. Pesenti (1990). Rational speculative bubbles in an exchange rate target zone. *NBER Working Papers 3467*.

Burnside, C., M. Eichenbaum, and S. Rebelo (1999). Hedging and financial fragility in fixed exchange rate regimes. *National Bureau of Economic Research Working Papers* (7143).

Bussiere, M., M. Cazorzi, A. Chudik, and A. Dieppe (2010). Methodological advances in the assessment of equilibrium exchange rates. *European Central Bank Working Paper Series* (1151).

Cameron, C. and P. Trivedi (1998). *Regression Analysis of Count Data*. Number 30 in Econometric Society Monograph. Cambridge University Press.

Campa, J. and L. Goldberg (2005). Exchange rate pass through into import prices. *Review of Economics and Statistics*.

Campbell, J. and A. Kyle (1993). Smart money, noise trading and stock price behaviour. *Review of Economic Studies 60*, 1–34.

Campbell, J., A. Lo, and A. MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton, New Jersey: Princeton University Press.

Capie, F., T. Mills, and G. Wood (2005). Gold as a hedge against the dollar. *International Financial Markets, Institutions and Money 15*, 343–352.

Cappiello, L., R. Engle, and K. Sheppard (2006). Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics 4*(4), 537–572.

Carlton, D. (1979). *Why New Firms Locate Where They Do: An Econometric Model*. The Urban Institute.

Carlton, D. (1983). The location and employment choices of new firms: an econometric model with discrete and continuous endogenous variables. *Review Of Economics And Statistics 65*(3), 440–449.

Cavaglia, S., W. Verschoor, and C. Wolff (12). Further evidence on exchange rate expectations. *Journal of International Money and Finance*, 78–98.

Chen, Y. and K. Rogoff (2003). Commodity currencies and empirical exchange-rate puzzles. *DNB Staff Reports (discontinued)* (76).

Cheung, Y. and K. Lai (1993). Do gold market returns have long memory? *The Financial Review 28*(2), 181–202.

Claessens, S. (1992). The optimal currency composition of external debt: theory and applications to mexico and brazil. *The World Bank Economic Review 6*, 503–528.

Clarida, R. and J. Gali (1994). Sources of real exchange rate fluctuations: How important are nominal shocks. *CEPR Discussion Papers 951.*

Clark, P. and R. MacDonald (1999). Exchange rates and economic fundamentals: A methodological comparison of beers and feers. *IMF Working Paper* (98-67).

Clark, P. and R. MacDonald (2000). Filtering the beer: A permanent and transitory decomposition. *Global Finance Journal 15*(1), 29–56.

Clements, K. and R. Fry (2008). Commodity currencies and currency commodities. *Resources Policy 33*, 55–73.

Clogg, C. and E. Shihadeh (1994). *Statistical models for ordinal variables.* Thousand Oaks, CA: Sage.

Cochrane, J. (2005). Financial markets and the real economy. *National Bureau of Economic Research Working Paper Series* (11193).

Cohen, B. (2005). Currency choice in international bond issuance. *Bank For International Settlements Quarterly Review*.

Cukierman, A. and A. Muscatelli (2002). Do central banks have precautionary demands for expansions and for price stability? theory and evidence. *CESifo Working Paper Series* (764).

Cutler, D., J. Poterba, and L. Summers (1990). Speculative dynamics and the role of feedback traders. *American Economic Review 80*, 63–68.

Damodaran, A. (1989). The weekend effect in information releases: a study of earnings and dividends. *Review of Financial Studies 2*(4), 607–623.

Das, S. and R. Uppal (2004). Systemic risk and international portfolio choice. *Journal of Finance 59*(6), 2809–2834.

Delong, B., A. Shleifer, L. Summers, and R. Waldmann (1990). Noise trader risk in financial markets. *Quarterly Journal of Economics 108*, 137–156.

Diebold, F. and M. Nerlove (1989). The dynamics of exchange rate volatility: A multivariate latent factor arch model. *Journal of Applied Econometrics 4*(1), 1–21.

Ding, Z. and R. Engle (2001). Large scale conditional covariance matrix modelling, estimation and testing. *Academia Economic Papers* (29), 157–184.

Dolado, J., R. Mara-Dolores, and F. Ruge-Murcia (2002). Non-linear monetary policy rules: Some new evidence for the us. *CEPR Discussion Papers* (3405).

Dominguez, K. (1998). Central bank intervention and exchange rate volatility. *Journal of International Money and Finance 17*, 161–190.

Dominguez, K. and J. Frankel (1993). Does foreign exchange-rate intervention work? *Institute for International Economics*.

Domowitz, I. and C. Hakkio (1985). Conditional variance and the risk premium in the foreign exchange market. *Journal of International Economics 19*, 47–66.

Dooley, M. and P. Isard (1979). The portfolio-balance model of exchange rates. *Federal Reserve System International Finance Discussion Papers* (141).

Doornik, J. and M. Ooms (2003). Multimodality in the garch regression model. *University of Oxford Nuffield College Economics Papers* (2003-W20).

Driver, R. and P. Westaway (2003). Concepts of equilibrium exchange rates. *Bank of England Working Paper 248*.

Economist, T. (2009). Japan's currency: up and away. *The Economist 5 February 2009*.

Edison, H. (1993). The effectiveness of central-bank intervention: A survey of the litterature after 1982. *Princeton Studies in International Economics 18* (18). International Economics Section.

Eichengreen, B. and R. Hausmann (1999). *Exchange rates and financial fragility.* Kansas City: Federal Reserve Bank of Kansas City.

Elliot, W., S. Huffman, and S. Makar (2003). Foreign-denominated debt and foreign currency derivatives: Complements or substitutes in hedging foreign currency risk? *Journal of Multinational Financial Management 13*, 123–139.

Enders, W. (2004). *Applied Econometric Time Series* (Second ed.). Hoboken, NJ: John Wiley and Sons.

Engel, C. and K. West (2005). Exchange rates and fundamentals. *Journal of Political Economy 113*, 485–517.

Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*(4), 987–1007.

Engle, R. (2002). Dynamic conditional correlation: a sumple class of multivariate generalised autoregressive conditional heteroscedasticity models. *Journal of Business and Economic Statistics 20*(3), 339–350.

Engle, R. and T. Bollerslev (1986). Modelling the persistence of conditional variances. *Econometric Reviews 5*, 1–50.

Engle, R. and F. Kroner (1995). Multivariate simultaneous generalised arch. *Econometric Theory 11*, 122–150.

Engle, R., V. Ng, and M. Rothschild (1990). Asset pricing with a factor arch covariance structure: empirical estimates for treasury bills. *Journal of Econometrics* (45), 213–238.

Engle, R. and K. Sheppard (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate garch. *NBER Working Papers* (8554).

Esaka, T. (2000). The louvre accord and central bank intervention: Was there a target zone? *Japan and the World Economy 12*(2), 107–126.

Fahrmeir, L. and G. Tutz (1994). *Multivariate Statistical Modelling Based on Generalized Linear Models.* New York, NY: Springer.

Fatum, R. and M. Hutchison (2003). Is sterilised foreign exchange intervention effective after all? an event study approach. *The Economic Journal 113*, 390–411.

Faugere, C. and J. Erlach (2004). The price of gold: A global required yield theory. *EconWPA Series in Finance* (0403003).

Fisher, E., R. Heinkel, and J. Zechner (1989). Dynamic capital structure choice: theory and tests. *Journal of Finance 44*, 19–40.

Fletcher, D. and L. Taylor (1996). Swap covered interest parity in long-date capital markets. *Review of Economics and Statistics 78*, 530–538.

Frenkel, J. (1976). A monetary approach to the exchange rate: doctrinal aspects and empirical evidence. *Scandinavian Journal of Economics 78*, 200–224.

Frenkel, M., C. Pierdzioch, and G. Stadtmann (2002). Bank of japan and federal reserve interventions in the yen/us dollar market: Estimating the central banks' reaction functions. *Kiel Institute For World Economics*.

Frenkel, M., C. Pierdzioch, and G. Stadtmann (2005). Bank of japan and federal reserve interventions in the yen/us dollar market: Estimating the central banks' reaction functions. *Kiel Institute For World Economics*.

Froot, K., D. Scharfstein, and J. Stein (1990). Herd on the street: Informational inefficiencies in a market with short-term speculation. *National Bureau of Economic Research Working Papers* (3250).

Furstenberg, G. and B. Jeon (1989). International stock price movements: Links and messages. *Brookings Papers on Economic Activity 20*(1989-1), 125–180.

Galati, G. and W. Melick (1999). Perceived central bank intervention and market expectations: an empirical study of the yen-dollar exchange rate, 1993-96. *BIS Working Paper 77*.

Graham, J. and C. Harvey (2001). The theory and practice of corporate finance: Evidence from the field. *Journal of Financial Economics 61*, 187–243.

Greene, W. (2008). *Econometric Analysis* (6 ed.). Englewood Cliffs: Prentice Hall.

Grilli, V. and N. Roubini (1992). Liquidity and exchange rates. *Journal of International Economics 32*(3), 339–352.

Guimaraes, P., O. Figueiredo, and D. Woodward (2003). A tractable approach to the firm location decision problem. *Review of Economics and Statistics 85*(1), 201–204.

Haavelmo, T. (1943). The statistical implications of a system of simultaneous equations. *Econometrica 11*(1), 1–12.

Habib, M. and M. Joy (2010). Foreign currency bonds: currency choice and the role of covered and uncovered interest rate parity. *Applied Financial Economics 20*(8), 601–626.

Hall, B., Z. Griliches, and J. Hausman (1984). Econometric models for count data with an application to the patents-research-and-development relationship. *Econometrica 52*(4), 909–983.

Hall, B., Z. Griliches, and J. Hausman (1986). Patents and r and d: Is there a lag? *International Economic Review 27*(2), 265–283.

Harada, K. and T. Watanabe (2005). Effects of the bank of japan's intervention on yen-dollar exchange rate volatility. *Journal of Japanese and International Economies*, 1–13.

Hau, H. and H. Rey (2006). Exchange rates, equity prices, and capital flows. *Review of Financial Studies 19*(1), 273–317.

Hausmann, R. and U. Panizza (2003). On the determinants of original sin: an empirical investigation. *Journal of International Money and Finance 22*, 957–990.

Henderson, B., N. Jegadeesh, and M. Weisbach (2006). World markets for raising new capital. *Journal of Financial Economics 82*, 63–102.

Herrera, A. and P. Ozbay (2005). A dynamic model of central bank intervention. *Central Bank of the Republic of Turkey Working Paper 0501*.

Hillebrand, E. and G. Schnabl (2005). A structural break in the effects of japanese foreign exchange intervention on yen-dollar exchange volatility. *Working Paper*.

Hillebrand, E., G. Schnabl, and Y. Ulu (2006). Japanese foreign exchange intervention and the yen/dollar exchange rate: A simultaneous equations approach using realized volatility. *Mimeo*.

Hirshleifer, D., A. Subrahmanyam, and S. Titman (1994). Security analysis and trading patterns when some investors receive information before others. *Journal of Finance 49*(5), 1665–98.

Hosmer, D. and S. Lemshow (2000). *Applied Logistic Regression* (2nd ed.). New York, NY: Wiley.

Huffman, S. P. and S. D. Makar (2004). The effectiveness of currency-hedging techniques over multiple return horizons for foreign-denominated debt issuers. *Journal of Multinational Financial Management 14*(2), 105–115.

Humpage, O. (1997). Recent us intervention: Is less more? *Federal Reserve Bank of Cleveland Economic Review 33*(3), 2–10.

Isard, P. (1996). Uncovered interest parity. *IMF Working Papers 96*(06). International Monetary Fund.

Ito, T. (2002). Is foreign exchange intervention effective? the japanese experience in the 1990s no. 8914. *NBER Working Paper* (8914).

Ito, T. and T. Yabu (2004). What prompts japan to intervene in the forex market? a new approach to a reaction function. *NBER Working Papers* (10456).

Ito, T. and T. Yabu (2007). What prompts japan to intervene in the forex market? a new approach to a reaction function. *Journal of International Money and Finance 26*, 193–212.

Johnson, D. (1988). The currency denomination of long-term debt in the canadian corporate sector: an empirical analysis. *Journal of International Money and Finance 7*, 77–90.

Juselius, K. (1991). *Econometric Decision Models: New Methods Of Modelling And Applications*, Chapter Long-run relations in a well-defined statistical model for the data-generating process: Cointegration analysis of the PPP and UIP relations between Denmark and Germany. New York: Springer Verlag.

Juselius, K. (1995). Do purchasing power parity and uncovered interest rate parity hold in the long run? an example of likelihood inference in a multivariate time-series model. *Journal of Econometrics 69*, 211–240.

Juselius, K. and R. MacDonald (2000a). The international parities between the usa, germany and japan: A tripolar analysis. *European University Institute Unpublished Report*.

Juselius, K. and R. MacDonald (2000b). International parity relationships between germany and the united states. *European University Institute Unpublished Report*.

Kaul, A. and S. Sapp (2006). Y2k fears and safe haven trading of the us dollar. *Journal of International Money and Finance 25*(5), 760–779.

Kearney, A. and L. Lombra (2008). The non-neutral short run effects of derivatives on gold prices. *Applied Financial Economics 18*, 985–994.

Kearns, J. and R. Rigobon (2005). Identifying the efficacy of central bank interventions: Evidence from australia and japan. *Journal of International Economics 66*(1), 31–48.

Kedia, S. and A. Mozumdar (2003). Foreign currency-denominated debt: An empirical examination. *Journal of Business 76*(4), 521–546.

Keloharju, M. and M. Niskanen (2001). Why do firms raise foreign currency denominated debt? evidence from finland. *European Financial Management 7*(4), 481–496.

King, M. and R. Fatum (2005). The effectiveness of official foreign exchange intervention in a small open economy: the case of the canadian dollar. *Bank of Canada Working Paper*.

Kirman, A. (1993). Ants, rationality and recruitment. *Quarterly Journal of Economics 108*, 137–156.

Knight, M. and D. Mathieson (1983). *Economic Interdependence and Flexible Exchange Rates*, Chapter Economic change and policy response in Canada under fixed and flexible exchange rates. Cambridge, Mass.: MIT Press.

Koch, P. and T. Koch (1991). Evolution in dynamic linkages across daily national stock indexes. *Journal of International Money and Finance 10*(2), 231–251.

Koedijk, K., P. Schotman, and M. V. Dijk (1998). The re-emergence of ppp in the 1990s. *Journal Of International Money And Finance 17*(1), 51–61.

Kroner, K. and S. Claessens (1991). Optimal dynamic hedging portfolios and the currency composition of external debt. *Journal of International Money and Finance 10*, 131–148.

Lien, D., Y. Tse, and A. Tsui (2002). Evaluating the hedging performance of the constant-correlation garch model. *Applied Financial Economics 12*(11), 791–98.

Ljung, G. and G. Box (1978). On a measure of lack of fit in time series models. *Biometrika 65*, 297–303.

Longin, F. and B. Solnik (1995). Is the correlation in international equity returns constant: 1960 to 1990? *Journal of International Money and Finance 14*, 3–26.

Luce, R. and P. Suppes (1965). *Handbook of Mathematical Psychology*, Volume 3, Chapter Preference, utility and subjective probability, pp. 249–410. New York: John Wiley and Sons.

Lyons, R. (2001). *The Microstructure Approach To Exchange Rates.* Cambridge, Mass.

MacDonald, R. (2000). Concepts to calculate equilibrium exchange rates: an overview. *Deutsche Bundesbank Discussion Paper 3*(00), 1–70.

MacDonald, R. and I. Marsh (1997). On fundamentals and exchange rates: a casselian perspective. *Review Of Economics and Statistics 79*(4), 655–664.

MacDonald, R. and I. Marsh (1999). *Exchange-Rate Modelling.* Amsterdam: Kluwer.

MacDonald, R. and T. Torrance (1990). Expectations formation and risk in four foreign exchange markets. *Oxford Economic Papers 42*, 544–561.

Maddala, G. (1983). *Limited-Dependent and Qualitative Variables In Econometrics.* Cambridge: Cambridge University Press.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance 7*(1), 77–91.

Marston, R. (1987). *Real-Financial Linkages Among Open Economies*, Chapter Real Exchange Rates And Productivity Growth In The United States And Japan, pp. 71–96. MIT Press, Cambridge, MA.

McBrady, M. and M. Schill (2007). Foreign currency-denominated borrowing in the absence of operating incentives. *Journal of Financial Economics 86*, 145–177.

McCullagh, P. (1980). Regression models for ordinal data. *Journal of the Royal Statistical Society* (42), 109–142.

McCurdy, T. and I. Morgan (1988). Testing the martingale hypothesis in deutsche-mark futures with models specifying the form of the heteroskedasticity. *Journal of Applied Econometrics 3*, 187–202.

McFadden, D. (1974). *Conditional logit analysis of qualitative choice behaviour*, Volume 4. New York: Academic Press.

McFadden, D. (1980). Econometric models for probabilistic choice among products. *Journal of Business 53*(3), 13–29.

McKinnon, R. and H. Pill (1999). Exchange rate regimes for emerging markets: Moral hazard and international overborrowing. *Stanford University Department of Economics Working Paper Series* (99018).

McLaren, Z. (2002). The effectiveness of reserve bank of australia foreign exchange intervention. *University of Melbourne Economics Department Research Paper* (849).

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

Milgate, M. (1998). *The New Palgrave Dictionary Of Economics*, Chapter Equilibrium: Development Of The Concept. MacMillan Press.

Mussa, M. (1979). Empirical regularities in the behaviour of exchange rates and the theories of the foreign exchange market. *Carnegie-Rochester Conference Series on Public Policy 11*(9-57).

Mussa, M. (1986). Nominal exchange rate regimes and the behaviour of real exchange rates: Evidence and implications. *Carnegie-Rochester Conference Series On Public Policy* (26).

Myers, S. (2001). Capital structure. *Journal of Economic Perspectives 15*(2), 81–102.

Neely, C. (2000). The practice of central bank intervention: looking under the hood. *Federal Reserve Bank of St Louis, October,* (2000-028A).

Neely, C. (2002). The temporal pattern of trading rule returns and exchange rate intervention: intervention does not generate technical trading profits. *Journal of International Economics 58*(1), 211–232.

Neely, C. (2005a). An analysis of recent studies of the effect of foreign exchange intervention. *Federal Reserve Bank of St Louis* (2005-030B).

Neely, C. (2005b). Identifying the effects of us intervention on the levels of exchange rates. *Federal Reserve Bank of St. Louis* (2005-031).

Neely, C. (2006). Identifying the effects of us intervention on the levels of exchange rates. *Federal Reserve Bank of St Louis Working Paper Series* (2005-031).

Neumann, M. (1984). *Devisenmarktinterventionen der Zentral-banken*, Volume 139 of *Schriften des Veriens fur Social politik*, Chapter Auf der Suce nach der Interventionsfunktion der Deutschen Bundesbank. Berlin: Duncker und Humblot.

Norden, S. V. (1986). Regime switching as a test for exchange rate bubbles. *Journal of Applied Econometrics 11*(3), 219–51.

Obstfeld, R. and K. Rogoff (1995). Exchange rates dynamics redux. *Journal of Political Economy 103*, 624–660.

Orphanides, A. and V. Wieland (2000). Inflation zone targeting. *European Economic Review 44*(7), 1351–1387.

Papke, L. and J. Wooldridge (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics 11*(4), 619–632.

Pavlova, A. and R. Rigobon (2003). Asset prices and exchange rates. *National Bureau of Economic Research Working Papers* (9834).

Payne, R. and P. Vitale (2003). A transaction level study of the effects of central bank intervention on exchange rates. *Journal of International Economics 61*(2), 331–52.

Peterson, B. and F. Harrell (1990). Partial proportional odds models for ordered response variables. *Journal of the Royal Statistical Society 39*(2), 205–217.

Popper, H. (1993). Long-term covered interest parity: evidence from currency swaps. *Journal of International Money And Finance 12*, 439–438.

Prais, S. and C. Winsten (1954). Trend estimators and serial correlation. *Cowles Commission Discussion Papers* (383).

Ramchand, L. and R. Susmel (1998). Volatility and cross correlation across major stock markets. *Journal of Empirical Finance 5*, 397–416.

Ranaldo, A. and P. Soderlind (2010). Safe haven currencies. *Review of Finance 14*(2), 1–23.

Reitz, S. and M. Taylor (2008). The coordination channel of foreign exchange intervention: A nonlinear microstructural analysis. *European Economic Review 52*(1), 55–76.

Rosett, R. (1959). A statistical model of friction in economics. *Econometrica 27*(2), 263–267.

Ruge-Murcia, F. (2004, February). The inflation bias when the central bank targets the natural rate of unemployment. *European Economic Review 48*(1), 91–107.

Santis, G. and B. Gerard (1998). How big is the premium for currency risk?1. *Journal of Financial Economics 49*(3), 375–412.

Sarno, L. and M. Taylor (2001). Official intervention in the foreign exchange market: Is it effective and, if so, how does it work. *Journal of Economic Literature 39*(3), 839–868.

Sherman, E. (1982). Gold: a conservative, prudent, diversifier. *Journal of Portfolio Management 8*(3), 21–27.

Shleifer, A. (2000). *Inefficient Markets: An Introduction To Behavioural Finance*. Oxford: Oxford University Press.

Siegfried, N., E. Simeonova, and C. Vespro (2007). Choice of currency in bond issuance and the international role of currencies. *European Central Bank Working Paper Series* (814).

Silvennoinen, A. and T. Tersvirta (2009). Modeling multivariate autoregressive conditional heteroskedasticity with the double smooth transition conditional correlation garch model. *Journal of Financial Econometrics 7*(4), 373–411.

Sjaastad, L. (2008). The price of gold and the exchange rates: Once again. *Resources Policy 33*(2), 118–124.

Swift, R. (2004). Exchange rate changes and endogenous terms of trade effects in a small open economy. *Journal of Macroeconomics 26*(4), 737–745.

Taylor, J. (2010). Not a repeat of the great intervention. *Economics One Online Weblog*.

Taylor, M. (1987). Covered interest parity: A high frequency, high quality data study. *Economica 54*(216), 429–28.

Taylor, M. (1992). *The New Palgrave Dictionary of Money and Finance*, Chapter Covered interest parity. London: MacMillan.

Taylor, M. (1994). *Understanding Interdependence: The Macroeconomics of the Open Economy*, Chapter Exchange-rate behaviour under alternative exchange-rate regimes. Princeton: Princeton University Press.

Taylor, M. (2004). Is official exchange rate intervention effective. *Economica 71*, 1–11.

Taylor, M. (2005). Official foreign exchange intervention as a coordinating signal in the dollar-yen market. *Pacific Economic Review 10*(1), 73–82.

Taylor, S. (1986). *Modelling Financial Time Series.* New York: Wiley.

Tsay, R. (2005). *Analysis of Financial Time Series* (2 ed.). New York, NY: Wiley.

Tsui, A. and Q. Yu (1999). Constant conditional correlation in a bivariate garch model: evidence from the stock market in china. *Mathematics and Computers in Simulation 48*, 503–509.

Upper, C. (2000). How safe was the safe haven? financial market liquidity during the 1998 turbulence. *Deutsche Bundesbank Research Centre Discussion Paper Series* (2000,01).

Wilks, S. (1962). *Mathematical Statistics.* New York: John Wiley and Sons.

Willet, T., M. Keil, and Y. Ahn (2002). Capital mobility for developing countries may not be so high. *Journal of Development Economics 68*, 421–434.

Williams, R. (2006). Generalised ordered logit / partial proportional odds models for ordinal dependent variables. *Stata Journal 6*(1), 58–82.

Williamson, J. (1983). The exchange rate system. *Policy Analysis In International Economics 5.* Institute Of International Economics.

Windmeijer, F. (2006). Gmm for panel count data models. *University of Bristol Discussion Paper Series 06*(591).

Worthington, A. and M. Pahlavani (2007). Gold investment as an inflationary hedge: cointegration evidence with allowance for endogenous structural breaks. *Applied Financial Economics Letters 3*(4), 259–262.

Yamori, N. and Y. Kurihara (2004). The day of the week effect in foreign exchange markets: multi currency evidence. *Research in International Business and Finance 18*, 51–57.

Yamori, N. and P. Mourdoukoutas (2003). Does day of the week effect in foreign currency markets disappear? evidence from the yen-dollar market. *International Finance Review 4*, 443–459.