The London School of Economics and Political Science

Essays on Foreign Exchange Risk

Lukas Kremens

Thesis submitted to the Department of Finance of the London School of Economics and Political Science for the degree of *Doctor of Philosophy*

April 2019

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

Statement of conjoint work
I confirm that chapter 2 is jointly co-authored with Ian Martin. I contributed 50% of the work for chapter 2.

I declare that my thesis consists of 33,361 words.

Acknowledgement

I am forever indebted to my advisors Daniel Ferreira and Ian Martin for their inspiration and advice. My countless interactions with them have fuelled my enjoyment of the work that has gone into this dissertation and their influence on my thinking pervades every chapter. I could not possibly have attempted this journey without their guidance.

I thank Martin Oehmke and Christian Julliard for their invaluable advice, as well as Ulf Axelson, Mike Burkart, Thummim Cho, Vicente Cuñat, Juanita Gonzalez-Uribe, Dirk Jenter, Dong Lou, Daniel Paravisini, Christopher Polk, Andrea Tamoni, Philippe Mueller, Andrea Vedolin, David Webb, and many others for the countless insightful interactions that have made my time at the LSE so immensely instructive. I gratefully acknowledge financial support from the London School of Economics in the form of the Adeline and Karl Goeltz Scholarship.

I am grateful to Lorenzo Bretscher, Bernardo Ricca, Jesús Gorrin, Brandon Han, Dimitris Papadimitriou, Francesco Nicolai, Su Wang, Gosia Ryduchowska, Marco Pelosi, Fabrizio Core, Simona Risteska, Alberto Pellicioli, James Guo, Petar Sabtchevsky, Una Savic, Olga Obizhaeva, Friedrich Geiecke, Matteo Benetton, and Gianluca Rinaldi for their insight and friendship over the years.

My deepest gratitude is to my family for the life-long support that has enabled me to pursue this path, and to Léa for her endless tolerance, loving encouragement, and abundant patience with me throughout and beyond this endeavour.

Abstract

This dissertation consists of three chapters examining three different dimensions of foreign exchange risk. In chapter one, I deal with currency redenomination risk in the Eurozone, that is, the risk that euros held in a particular Eurozone country are converted into a new national currency once the country leaves the currency union. This conversion exposes holders of euro-denominated assets in that country to foreign exchange risk. I extract empirical measures of redenomination risk from asset prices and show that redenomination risk in France and Italy spikes around plebiscites in 2017 and 2018. French redenomination risk is associated with redenomination risk in other Eurozone countries, while Italian redenomination risk does not co-move with similar risks in other countries. These results are consistent with the interpretation that a French—unlike an Italian—exit from the Eurozone is associated with a Eurozone break-up.

Chapter two is conjoint work with Ian Martin. We present a new identity that relates expected exchange rate appreciation to the currency's (risk-neutral) covariance with equity markets, and use it to motivate a currency forecasting variable based on the prices of quanto index contracts. We show via panel regressions that the quanto forecast variable is an economically and statistically significant predictor of currency appreciation and of excess returns on currency trades. Out of sample, the quanto variable outperforms predictions based on uncovered interest parity, on purchasing power parity, and on a random walk as a forecaster of differential (dollar-neutral) currency appreciation.

In the third chapter, as in chapter two, I examine the exposures of different currencies to equity market risk. In contrast to the second chapter, chapter three analyses the link between risk exposures and speculative trading patterns, rather than measuring conditional risk exposures to forecast returns. I find, in a post-crisis sample, that currencies are more positively correlated with equity market returns, when hedge funds are long the currency future and vice versa for short positions.

Contents

List of Tables

List of Figures

1. Currency Redenomination Risk

Lukas Kremens[1]

How stable is the Eurozone? The debate about its composition is older than the currency union itself. However, while this debate has focused on potential new members for most of its history, more recent controversy has evolved around potential exits by current members. Current members could, in principle, re-introduce their own currencies to restore monetary sovereignty. This option exposes holders of outstanding sovereign bonds to currency redenomination risk: the risk of receiving the fixed payments of interest and principal in a different currency from the original numéraire. In this chapter, I present a quantitative measure of this redenomination risk in French, Italian, and German government bonds, which is forward-looking, based solely on asset prices, and observable in real time.

Greece came close to exiting the Eurozone in the summer of 2015, when the Greek government imposed severe capital controls and bank closures for almost 20 days, during negotiations with its public international creditors regarding extended loan facilities. Three years earlier, Greece had restructured a large portion of its outstanding bonds, but remained a member of the currency union. The two episodes highlight the distinction between redenomination and 'conventional' credit risk. While bondholders face losses in either scenario, other stakeholders throughout the Eurozone (e.g., depositors or banks) experience vastly different shocks. In addressing the empirical question of spillovers from a Eurozone exit, it is therefore crucial to distinguish exit risk from other forms of default.

[1] I am grateful to Martin Oehmke, Christian Julliard, Ian Martin, Daniel Ferreira, Gianluca Rinaldi, Lorena Keller, Patrick Augustin, Victor Lyonnet, Matteo Benetton, Andrea Vedolin, Francesco Nicolai, Pete Zimmerman, Bernard Dumas, and seminar participants at LSE, Dauphine, EFA-DT, HEC, Collegio Carlo Alberto (LTI), Chicago, LBS, HBS, BC, UT Austin, UW, UIUC, Bocconi, and UNC for helpful comments. I thank the Systemic Risk Centre at the LSE for their support, and for providing access to data sourced from Markit under license.

A key determinant of spillovers from redenomination risk is whether or not an initial single exit from the Eurozone is contagious in the sense that it is followed by a Eurozone break-up. The fear of such contagion spreading from a 'Grexit' (or 'Graccident') was undoubtedly an important factor in the decision taken by Greece's public creditors to restructure and extend their loans in 2015.

An empirical approach to assessing the risk of such contagion requires a clean, observable measure of redenomination risk for different Eurozone countries. I construct a measure of redenomination risk for three large Eurozone members (France, Italy, and Germany), using a change in the standardized terms of sovereign credit default swaps (CDS). CDS contracts issued before September 2014 effectively allow for a redenomination into a new national currency for any G7 country without triggering payouts from the CDS, irrespective of any exchange rate losses incurred by bondholders in the process. Contracts signed under the new terms, implemented in September 2014, are triggered by a redenomination out of euros into a new and depreciating national currency.

To illustrate the different economics of the two CDS contract types, suppose that France and Spain both decide to leave the Eurozone and change the currency of payments on outstanding bonds into new francs and new pesetas, respectively, each with a conversion rate of 1-to-1. Suppose further that once the new currencies are traded, the first freely determined market exchange rate is 0.8€ per franc and 0.75€ per peseta, and, that French and Spanish bonds trade at par in the new currencies following redenomination. This last assumption ensures that the only loss to bondholders stems from the initial depreciation of the new currency. In this simple example, redenomination triggers payouts for French CDS contracts issued under the 2014 definitions, and these pay out 20% of the notional value, that is, the loss to bondholders from redenomination ($= 1 - 0.8$). However, French CDS based on the previous definitions, are not triggered and therefore make no payout at all, since the new currency of denomination (the new franc) is that of a G7 country (France) and the redenomination itself, therefore, does not constitute a credit event. In contrast, both contract types are triggered in a CDS written on Spain, and CDS holders receive 25% ($= 1 - 0.75$) of the notional value.

The pricing difference between French contracts under the two definitions reflects how market participants asses the likelihood of a redenomination and the contingent losses incurred by bondholders. I account for other contractual changes and potential liquidity differences by subtracting the same difference measure for a matched synthetic control country constructed from non-G7 Eurozone countries.

The resulting time-series measure is analogous to other commonly used difference-in-difference approaches, with the important feature that it uses contemporaneous differences rather than before-after relations. I will refer to this measure as the *redenomination spread*. This daily time series reflects the cost of insurance against losses from currency redenomination and is directly observable from sovereign CDS spreads.

Armed with the measure, I establish various empirical facts. French redenomination risk is economically small for most of the sample period, but spikes to 25 basis points per year in the run-up to the first round of the French presidential elections in April 2017. At its peak, redenomination risk accounts for 40% of the total French CDS spread. With the first-round victory of pro-EU candidate Emmanuel Macron, redenomination risk drops sharply. In the case of Italy, redenomination risk spikes sharply to around 80 basis points following the formation of a Eurosceptic government in May 2018, accounting for almost one third of the total Italian CDS spread. In both cases, redenomination risk is driven by political shocks: as such, the measure reveals shocks to a country's political willingness to remain a Eurozone member. German redenomination risk is close to zero throughout the sample period, consistent with the interpretation that a redenomination into a new German currency is not expected to cause losses for bondholders and/or such a redenomination is highly unlikely.

If such a measure were available for all members of the currency union, the covariance matrix of these time series would directly reveal contagion of redenomination risk. In the absence of a broader cross-section of redenomination risk measures, I will document signs of contagion in other asset prices. Prior to an exit, the prospect of initial depreciation for some of the new national currencies induces capital flight out of weaker and into stronger countries.[2] The distinctive feature of a currency union is that exchange rates cannot adjust to such flows. Instead, the adjustment works through the yields of the assets targeted by such flows, such as sovereign bonds. Investors demand higher nominal yields on assets that are likely to be redenominated into a depreciating new currency, and yields in other Eurozone countries fall if their bonds remain denominated in euros or repay in a new, stronger currency.[3]

A Eurozone exit by, say, France can either be *isolated* or *contagious*. The signa-

[2]The dissolution of Czechoslovakia in February 1993 provides a historical example of such capital movements from the subsequently weaker currency area (Slovakia) into the stronger one.

[3]Brunnermeier et al. (2016b, p. 226) make a similar point, arguing that "*as [...] redenomination risk does not exist for 'German euros', a Greek euro will necessarily be worth less than a German euro. As long as Greek euros can be converted one-to-one into German euros, Greeks may thus decide to withdraw their deposits [...] and buy German Bunds...*".

ture of the latter type is that pricing spillovers to other Eurozone bonds are heterogeneous: if each country's bonds are ultimately repaid in a new national currency, bond prices reflect the expected gains or losses that the currency redenomination imposes on bondholders. In contrast, an isolated exit implies no such consequences for the remaining Eurozone bonds. All other bonds repay in euros, and function as potential substitutes to French debt for euro-investors tilting their portfolios away from redenominatable French bonds. I model the two cases formally in Section 1.4.1.

I find that sovereign yields drop significantly with increases in Italian redenomination risk for all Eurozone countries other than Italy. In contrast, Eurozone yields co-move heterogeneously with French redenomination risk. German and Austrian yields fall as French redenomination risk rises. However, yields in Italy, Portugal, and—naturally—France rise with French redenomination risk. Corporate credit spreads paint a similar picture: spreads tend to drop with increases in Italian redenomination risk for financial and non-financial firms outside Italy. Similarly to German Bund yields, US Treasury yields fall with rising French redenomination risk, but the effect is weaker in magnitude than for Bunds. In relation to Italian redenomination risk, however, Treasury yields do not fall like Eurozone yields. The euro-dollar exchange rate tends to depreciate with French, but not with Italian redenomination risk. This set of findings is consistent with the interpretation that a French exit from the monetary union is contagious and expected to be associated with further redenominations and a broader break-up of the Eurozone, while an Italian exit is expected to remain isolated. The heterogeneity in responses to French (i.e., contagious) redenomination risk corresponds to heterogeneity in the countries' fiscal positions, labor productivities, and current account balances, consistent with the interpretation that these responses reflect expected post-Eurozone appreciation and depreciation of national shadow currencies.

Literature.—On the surface, my empirical measure of redenomination risk is related to the measure of De Santis (forthcoming), who uses *quanto* CDS, that is, the difference between dollar-denominated CDS and contracts denominated in euros. As Mano (2013) shows, this difference measures the (risk-neutral) expected depreciation of the euro against the dollar in the event that CDS payouts for a given country are triggered. Similarly, Augustin et al. (2018) disentangle expected depreciation from the default event risk in a structural model, using the term structure of quanto CDS. While this provides an important measure of euro currency risk and its connection to sovereign default risk, it does not distinguish between credit risk and redenomination risk. Instead, my measure isolates the currency redenomination event as a particular form of default and relates directly to the depreciation of the new national currency versus the euro, as opposed to the euro versus the US dollar. The wide-spread view that a sovereign default by a Eurozone member is likely to lead to a euro depreciation has also led to the gradual disappearance of euro-denominated CDS contracts for Eurozone sovereigns. My redenomination risk measure uses only the more liquid dollar-denominated contracts.

My empirical analysis of sovereign CDS spreads adds to a wide literature, including Pan and Singleton (2008), Augustin (2014), and Fontana and Scheicher (2016). Augustin et al. (2014) provide a broad survey of sovereign CDS markets. Longstaff et al. (2011) show strong co-movement in sovereign CDS. Beyond the study of CDS, this chapter links to the extensive literature on sovereign risk and contagion (e.g. Reinhart and Rogoff, 2011). Arellano et al. (2018) look at the financial linkages responsible for such spillovers of sovereign risk in the Eurozone. Aguiar et al. (2015) show how debt crises in one member country impact other members through centralized monetary policy. The distinction between credit risk and redenomination risk in a currency union is analogous to the question of local currency sovereign risk, studied by Du and Schreger (2016). In addressing the impact of political risk on asset prices, my approach also relates to the work of Pastor and Veronesi (2013) and Kelly et al. (2016). Neuberg et al. (2018) exploit other differences relating to government intervention and bail-in events between CR14 and CR restructuring clauses in CDS contracts written on financial institutions. In analogy to my approach, Berndt et al. (2007) distinguish between restructuring events and default events and estimate restructuring risk premia in US corporate debt by comparing CDS contracts with, and without, restructuring clauses.

Redenomination risk has also been identified by the ECB as a risk to the transmission of monetary policy and an explicit target of policy measures.[4] Krishnamurthy et al. (2018) assess the effect of three specific ECB policy measures launched in 2011-2012 on bond yields and redenomination risk. They quantify redenomination risk in sovereign bonds by decomposing a panel of sovereign and corporate yields. The key identifying assumptions are that (i) default affects bonds under foreign law and bonds under domestic law in the same way; and (ii) corporate and sovereign bonds are affected in the same way by redenomination. Bayer et al. (2018) construct a term structure of redenomination risk. Importantly, my measure does not rely on combinations of bonds CDS and is therefore robust to variation in the

[4]See, for instance, Benoît Cœuré's speech on the objectives of the OMT program (03/09/2013): ecb.europa.eu/press/key/date/2013/html/sp130902.en.html. See also Leombroni et al. (2017).

so-called CDS-bond basis (see, e.g., Bai and Collin-Dufresne (forthcoming)). Redenomination risk affects, through deposit redenomination, the portfolio choice of banks holding euro-denominated sovereign debt. The well-documented home bias of banks in euro sovereign bonds has sparked a large literature on sovereign-bank feedback loops.[5] The simple model I present in Section 1.4.1 features home bias as a natural equilibrium outcome of redenomination risk.

1.1 Redenomination and credit default swaps

A credit default swap is a bilateral financial contract wherein one party (the *protection seller*) provides insurance to the other party (the *protection buyer*) against losses to the holders of bonds issued by a particular entity (the *reference entity* or *issuer*). In the event (referred to as a *credit event*) that the reference entity fails to honor its contractual obligations as the issuer of its outstanding bonds, the protection buyer receives from the protection seller a payment of a prespecified face value (*notional*) minus the recovery on this face value. This recovery is typically set at the market value of defaulted bonds, which is determined in an auction of such defaulted bonds arranged by the International Swaps and Derivatives Association (ISDA). In exchange, the protection buyer pays the protection seller a periodic (typically quarterly) insurance premium: the so-called CDS *spread*.

Denote the spread today for a swap with maturity T by $S_{0,T}$. Swaps are quoted such that the market value of the swap is zero and no money is exchanged at initiation, i.e., the expected discounted value of payments to the protection seller equals that of payments to the protection buyer. For expositional purposes, consider the simplified case of a hypothetical single-period CDS:

$$\begin{aligned} S_{0,T} =& e^{-rT}\,\mathbb{E}_0^{\mathbb{Q}}\left(\mathbb{1}_T(1-R_T)\right) \\ =& e^{-rT} q_T\,\mathbb{E}_0^{\mathbb{Q}}(1-R_T \mid \mathbb{1}_T = 1), \end{aligned} \tag{1.1}$$

where the indicator denotes the occurrence of a credit event between 0 and T, q_T the probability of said credit event, and R_T denotes the contingent recovery rate. While I will not, without further assumptions, be able to disentangle the probability q_T from the conditional loss $(1-R_T)$, the CDS spread is economically meaningful in itself, as it reflects the economic cost of insurance against losses (net of recovery) to creditors of a certain entity from a range of credit events. To facilitate the trading of CDS, the

[5] See, for instance, Acharya et al. (2014), Farhi and Tirole (forthcoming), and Brunnermeier et al. (2016a).

contract terms typically follow a standardized set of definitions, governed by ISDA, including the precise circumstances, which constitute a credit event and trigger the insurance payout. Currency redenomination may be one of these circumstances, but the insurance premium will also reflect other risks, such as the bankruptcy filing of the issuer, the failure to make a contractual interest or principal payment, or the restructuring of the bonds to the detriment of bondholders.

For the purposes of this chapter, the reference entities of the CDS will predominantly be sovereign countries, and—for lack of established bankruptcy procedures for such borrowers—'defaults' typically occur in the form of a restructuring. However, a restructuring itself may take many different forms and is not limited to currency redenomination: for instance, Greece restructured a large part of its outstanding debt in 2012 by exchanging existing bonds for a package of new securities with longer maturity, lower face value, and lower coupon rate, while, at the same time, keeping the euro as the currency of denomination. The CDS spread reflects the risk of all of these credit events, rather than isolate the risk of currency redenomination. For the remainder of this chapter, and in a slight abuse of terminology, I will use 'default' to refer to any credit event that does *not* involve a change in the currency of denomination. In contrast, I will use 'redenomination' to refer to a restructuring involving *only* this change of currency.

1.1.1 Credit event definitions – 2003 versus 2014

ISDA periodically updates the standardized definitions. The most recent update was implemented in September 2014. Many of the revisions from the earlier definitions (released in 2003) address problems in CDS on corporate issuers, some responding directly to events unfolding over the Eurozone sovereign debt crisis (particularly relating to financial institutions and government interventions such as bail-outs or bail-ins). However, a few changes relate specifically to sovereign reference entities. One of the new terms refers to the set of events that constitute a restructuring, defined in Section 4.7 of the ISDA definitions. Subsection (a)(v) specifies a number of "permitted currencies" into which an obligation may be redenominated without triggering the CDS payout. Under the 2003 definitions,

"Permitted Currency" means (1) the legal tender of any Group of 7 country (or any country that becomes a member of the Group of 7 if such Group of 7 expands its membership) or (2) the legal tender of any country which, as of the date of such change, is a member of the Organization for Economic Cooperation and Development

and has a local currency long-term debt rating of AAA or higher [...].[6]

The Group of 7 (G7) consists of Canada, France, Germany, Italy, Japan, the UK, and the US. The three current Eurozone members France, Germany, and Italy would therefore—without triggering CDS payouts—be able to leave the Eurozone, issue a national currency, and redenominate any existing debt into this new currency, regardless of any market value losses that such a redenomination may imply for bondholders. During the Eurozone sovereign debt crisis, the potential consequences for CDS contracts of a member country exiting the currency union, as well as the distinction between G7 countries and other Eurozone members, became a widely debated topic among market participants.[7] In response to the unwanted special status of French, German, and Italian debt, ISDA amended Section 4.7(a) (v) in its 2014 definitions to define the relevant redenomination event as

... any change in the currency of any payment of interest, principal or premium to any currency other than the lawful currency of Canada, Japan, Switzerland, the United Kingdom and the United States of America and the euro and any successor currency to any of the aforementioned currencies (which in the case of the euro, shall mean the currency which succeeds to and replaces the euro in whole).[8]

Therefore, redenomination into a new French, German, or Italian currency triggers CDS contracts under the 2014 definitions (if such a redenomination leads to market value losses for bondholders), but not for contracts under the 2003 definitions. Accordingly, the two contracts are quoted separately in financial markets, specifying the applicable restructuring clause as either 'CR14' for 2014 definitions or 'CR' for 2003.

For an illustrative example of the pricing consequences, we revisit the case of the potential exits of France and Spain from the Eurozone, as well as the simplified pricing equation (1.1). Consider at time t the pricing of single-period CDS contracts with maturity $t+1$. Suppose that the net risk-free interest rate, r, is equal to zero, and that the risk-neutral probability of either exit at time $t+1$ is $q^R_{i,t+1} = 0.1$ for $i = \{FRA, ESP\}$. As previously, the expected depreciations of the new national currencies against the euro, are $\mathbb{E}^{\mathbb{Q}}_t R^R_{FRA,t+1} = 0.8$ and $\mathbb{E}^{\mathbb{Q}}_t R^R_{ESP,t+1} = 0.75$. The loss from redenomination, $1 - R^R$, may stem from a number of sources: in the absence of further amendments to the debt contract, depreciation of the new cur-

---

[6]ISDA (2003, p. 32-33) Credit Derivatives Definitions

[7]See, e.g., ftalphaville.ft.com/2010/02/12/148481/euro-breakup-not-necessarily-a-credit-event/.

[8]ISDA (2014) Credit Derivatives Definitions, p. 42

rency is likely to be responsible for a large part of the losses suffered by bondholders. Upon introduction, the leaving country chooses an initial 'conversion rate' of its new national currency against the euro, for the purposes of redenominating various contracts within the economy, such as sovereign debt. In the above example of France and Spain, both conversion rates are 1-to-1. However, this rate does not represent a market exchange rate. With the split from the euro and the re-nationalization of monetary policy, the new currency obtains its own risk characteristics and risk premium as well as its own future interest rate path. If the new currency differs from the euro in either of these two dimensions, the new market exchange rate has to deviate from the initially chosen conversion rate. This is an incarnation of the *Mundell-Fleming trilemma*: keeping the exchange rate at the level of the conversion rate amounts to fixing the exchange rate against the euro, while the monetary policy path is allowed to deviate from that of the currency union, both of which are not jointly attainable in the absence of capital controls.

In addition, the prices of the redenominated bonds may also reflect changes in credit risk, if the country's fiscal position changes following the Eurozone exit. At the same time, suppose that the risk-neutral probability of either country restructuring its debt without a change of currency, i.e., 'defaulting' at $t+1$ is $q^D_{i,t+1} = 0.1$, with an expected recovery of $\mathbb{E}^{\mathbb{Q}}_t R^D_{i,t+1} = 0.5$. Also suppose that the events of redenomination and default are independent. This assumption may not seem innocuous, but is in line with the contractual differences between CR and CR14 clauses: a default occurring *simultaneously* with redenomination would constitute a credit event under either contract. As such, my approach of looking at the difference between CR and CR14 spreads neglects such an event of simultaneous redenomination and default. To the extent that redenomination is likely to be accompanied by simultaneous default, my measure of redenomination risk in isolation from default provides a lower bound on the true magnitude of redenomination risk.

Returning to the illustrative example, denote by S^{CR14}_i and S^{CR}_i country i's single-period CDS spread under CR14 and CR restructuring clauses, respectively. For France, only CR14 contracts recognize redenomination into new francs as a credit event, so

$$\begin{aligned} S^{CR14}_{FRA} &= q^D_{FRA,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^D_{FRA,t+1}) + q^R_{FRA,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^R_{FRA,t+1}) = 0.07 \\ S^{CR}_{FRA} &= q^D_{FRA,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^D_{FRA,t+1}) = 0.05. \end{aligned}$$

The spread-difference between CR14 and CR contracts is sometimes referred to as the 'ISDA basis'. If there are no other pricing differences between CR and CR14 contracts, the ISDA basis directly measures the insurance premium due to redenomination risk.

However, the clause on permitted currencies is not the only difference between the two contract types relevant to sovereign issuers. A clause referred to as 'Asset Package Delivery' (APD) affects the calculation of the recovery value and may, therefore, lead to differential pricing of the two contracts. The clause is described in more detail in Subsection 1.5.2. Unlike the clause on permitted currencies, APD does not distinguish issuers based on G7 membership. Similarly, liquidity may differ between the newer CR14 CDS and the superseded CR contracts. Therefore, a diff-in-diff approach is well-suited to isolate the pricing impact of redenomination risk. Suppose that all potential pricing differences between the two contract types, which are unrelated to redenomination risk (notably APD or liquidity), are captured by an extra spread $\lambda_t = 0.02$. The pricing equation (1.1) for French contracts under 2003 definitions becomes $S^{CR}_{FRA} = e^{-r} q_{i,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R_{i,t+1}) - \lambda_t = 0.03$. For Spain, both restructuring clauses are triggered by redenomination into new pesetas, and therefore

$$\begin{aligned} S^{CR14}_{ESP} &= q^{D}_{ESP,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^{D}_{ESP,t+1}) + q^{R}_{ESP,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^{R}_{ESP,t+1}) = 0.075 \\ S^{CR}_{ESP} &= S^{CR14}_{ESP} - \lambda_t = 0.055. \end{aligned}$$

While simply taking the difference between S^{CR14}_{FRA} and S^{CR}_{FRA} jointly reveals redenomination risk and liquidity- or APD-driven components of the spread, the diff-in-diff measure isolates the component of the spread that is due to redenomination risk:

$$\left(S^{CR14}_{FRA} - S^{CR}_{FRA}\right) - \left(S^{CR14}_{ESP} - S^{CR}_{ESP}\right) = q^{R}_{FRA,t+1}\,\mathbb{E}^{\mathbb{Q}}_t(1 - R^{R}_{FRA,t+1}) = 0.02.$$

Of course, λ_t may itself be a function of other variables and therefore differ across countries. For the diff-in-diff measure, I construct a synthetic control country to match the time-variation in several characteristics of French and Italian CDS and bond markets, such as yield levels and bid-ask spreads.

1.2 The redenomination spread

I collect daily CDS spreads for dollar-denominated contracts with a maturity of five years for the Eurozone member countries Austria, Belgium, France, Germany, Ireland, Italy, the Netherlands, Portugal, and Spain. The CDS time series range from September 2014 when the CR14 contracts were launched, to June 2018. I focus on the five-year maturity, because these CDS contracts tend to be the most liquid. Since CDS are traded over-the-counter, transaction prices are difficult to observe.[9] However, Markit collects quotes from a range of market makers and financial intermediaries and reports consensus measures obtained from these quotes. These consensus measures are then widely used by derivatives market participants as an external valuation of their accounting positions as well as to fulfil regulatory requirements. I assess the liquidity and reliability of the quotes provided to Markit in Subsection 1.5.1.

Figure A.1 reports the time series of outstanding notionals by country. Net notionals are shown in Panel A and gross notionals in Panel B. Unfortunately, volumes are only available on an aggregated basis rather than by contract type (CR14 and CR).[10] The Italian CDS market is by far the largest in the Eurozone with over $12bn outstanding net notional as of February 2018, followed by the French, German, and Spanish CDS markets with over $6bn, $5bn, and $4bn aggregate net notional, respectively. Among the sampled Eurozone economies, CDS markets are smallest for Austria ($1.6bn), the Netherlands ($1.6bn), and Belgium ($2.5bn). Outstanding notionals have been trending downwards across all countries since the height of the European sovereign debt crisis in 2012. Overall CDS market volumes have been declining since 2008 (Oehmke and Zawadowski, 2017) reflecting a reduction in inter-dealer volumes; relative to corporate single-name CDS, the share of sovereign reference entities has risen steadily and quadrupled to 16% in June 2015, from 4% in December 2008 (BIS, 2015). Total outstanding volumes rose slightly in late 2014 to early 2015 for French and Italian CDS, consistent with the introduction of the new CR14 contracts.

Table A.1 reports summary statistics on the different CDS spreads. Spreads on CR and CR14 contracts are strongly positively correlated for all countries in the sample. However, since the contracts differ in their treatment of currency redenomination for France and Italy, the correlation is much weaker ($\rho_{FRA} = 0.86$ and $\rho_{ITA} = 0.75$, respectively) than for non-G7 countries, where correlation coefficients are in excess of 0.97. Similarly, the difference between CR and CR14 spreads (i.e., the ISDA basis) is more volatile relative to its mean for France and Italy than for other countries. Based on summary statistics, Germany resembles the control group countries: the difference between CR14 and CR spreads is close to zero, as is its volatility, and the correlation of the two CDS spreads is close to perfect at

[9]See Oehmke and Zawadowski (2017) for an overview of trading in (corporate) CDS markets.

[10]Outstanding notional data are obtained from swapsinfo.org.

$\rho_{GER} = 0.97$.

The CDS spread does not measure the probability of a redenomination event, but rather the cost of insurance against *losses* from the event. For any restructuring event to trigger the CDS payout, the restructuring must be to the detriment of bondholders. For a currency redenomination, this means that the exchange rate must depreciate from its conversion rate at redenomination. In the case of a newly issued national currency, there is no established market exchange rate. Broadly speaking, the new exchange rate will depreciate from the conversion rate fixed for redenomination if (i) market participants expect monetary policy at the national level to be more inflationary than the previously centralized policy in the currency union, and/or (ii) the risk characteristics of the newly issued currency are such that investors demand a higher risk premium to hold the new currency than the euro.[11] If market participants expect, say, a new German mark to appreciate upon introduction, a potential redenomination would not cause losses for bondholders and therefore not trigger CDS payouts. Consequently, the ISDA basis for Germany would be (close to) zero in this case.

Figure A.2 plots CDS spreads for Austria, Belgium Spain, Ireland, the Netherlands, and Portugal. Despite not being affected by the change to permitted redenomination currencies, the ISDA basis is positive for all control group members and widens slightly over the last year of the sample. Figure A.3 plots the different spreads for France, Italy, and Germany. Buying protection via a CR14 contract (solid) is consistently more expensive than via a CR contract (dashed). The difference is indeed close to zero, but positive, for Germany throughout the sample period. The sign of the basis in the control group suggests that the liquidity- or APD-driven component of the off-the-run CDS spread is positive, i.e., $\lambda_t > 0$. Consequently, the ISDA basis itself is not a clean measure of redenomination risk since it compares older and newer CDS contracts which are subject to different levels of liquidity and different calculations of recovery values.

To isolate redenomination risk, I construct—in the spirit of Abadie and Gardeazabal (2003)—synthetic controls from the different control group countries, which match the treated countries as closely as possible on relevant dimensions. Since treatment (i.e., being a G7 country) affects the economics of the (old) CR contract, the goal of the synthetic control is to construct a counterfactual CR CDS spread for France, Italy, and Germany without their respective G7-membership. For each

[11] Hassan et al. (2016) discuss to what extent these risk characteristics are chosen by policy makers.

trading day of the sample period, the synthetic control for each treated country is a convex combination of control group countries, matched on variables, which are successful contemporaneous predictors of the counterfactual CR spread. The four variables I choose for this matching are: (i) the CR14 CDS spread (which does not distinguish between issuers based on G7-membership); (ii) the bid-ask spread for the CR14 CDS spread; (iii) the five-year sovereign bond yield; and (iv) the bid-ask spread of the five-year sovereign bond yield. Daily time series for the three latter variables are obtained from Bloomberg.

For days, where no observations are available for a particular control group country on one or more of the matching variables, that country is excluded from the control group for that day, and the synthetic control is formed as a convex combination of the remaining control group countries. Similarly, if on any given day, observations are missing on any particular matching variable for more than one control group country, that variable is omitted from the matching process for that day.

Similarly to Abadie et al. (2010), I pick the weights of the matching variables by optimizing the fit of the synthetically constructed CR spread for a control group country to the observed CR spread of that country. Figure A.4 plots the observed (solid) and synthetic (dashed) CR spreads for Belgium, Spain, and Ireland over the sample period, showing that the synthetic control procedure generates a close fit in these 'placebo' countries. Across the three different placebo countries, the optimal weights are similar, and I will use the median set of optimal weights (Spain) to generate the synthetic controls for the three G7 countries. The resulting optimal matching procedure places the largest weight on the two CDS variables: the CR14 spread plays the dominant role (with a weight of 0.8626), followed by its bid-ask spread (0.1332). The two bond market variables do not contribute sizeably to the matching, with the optimal weight on the five-year sovereign yield and bond market bid-ask spread close to zero at 0.0013 and 0.0029, respectively. These matching weights are constant over the sample period.

The time-varying weights of each control group country in the synthetic control are then chosen each day to minimize the weighted sum of squared deviations of the matching variables for the synthetic control from the observed matching variables for, respectively, France, Italy, and Germany for that day. Using these time-varying country-weights, I then compute the time series of credit spreads $CR14_{s(i),t}$ and $CR_{s(i),t}$ for the synthetic control country as convex combinations of the control group observations for the respective CDS spread. The final diff-in-diff measure is then computed as $RS_{i,t} = CR14_{i,t} - CR_{i,t} - \left(CR14_{s(i),t} - CR_{s(i),t}\right)$ for $i = \{FRA, ITA, GER\}$. The diff-in-diff measure is designed to eliminate the confounding factors contained in the raw difference between CR14 and CR spreads, such as differential liquidity between older and newer contracts. I discuss in Section 1.5 two empirical concerns and outline why my diff-in-diff methodology is appropriate in this particular setting.

Table A.2 reports summary statistics by country for the diff-in-diff measure. The redenomination spread measures are distinct from conventional credit risk: the correlation coefficients between the redenomination spreads and CR CDS spreads are -0.03 for France, -0.01 for Italy, and 0.08 for Germany. Figure A.5 plots the diff-in-diff measure for France and Germany. As discussed above, the new currency must be expected to depreciate for redenomination to render sovereign bonds *risky* (with respect to redenomination) and for this risk to show up in CDS spreads. The lower plot shows the German redenomination spread against RS_{FRA}: German redenomination risk is close to zero throughout the sample period, consistent with the interpretation that either (i) the probability of redenomination is close to zero, or (ii) conditional on redenomination, the new currency is not expected to depreciate against the euro.

The French redenomination spread hovers around zero for most of the sample, but spikes dramatically to 25 basis points in the run-up to the presidential elections in spring 2017: the two red asterisks indicate the Fridays before each of the two election rounds (Sunday, April 23rd, and Sunday, May 7th, 2017). In the two-round system, a president is elected by absolute majority in the first round. If—as is commonly the case—no candidate receives an absolute majority, the two candidates with the highest vote move to the second round, in which one candidate will attain more than 50% of the votes. In 2017, pre-election polls saw four candidates as potential contenders in the decisive second round, including far-left candidate Jean-Luc Mélenchon and far-right candidate Marine Le Pen, both vocal critics of the European Union and widely considered potential supporters of a French exit from the Eurozone. Figure A.6 shows the combined vote share of Mélenchon and Le Pen from February through April against RS_{FRA}.[12] On Sunday, April 23rd, the results of the first round eliminated the possibility of a run-off between these two candidates, since pro-European candidate Emmanuel Macron placed first. The following Monday, the redenomination spread drops sharply to 7 from 21 basis points. According to

[12]Polling results are obtained from various sources. A convenient summary is available at `en.wikipedia.org/wiki/Opinion_polling_for_the_French_presidential_election,_2017`.

polls, first-round runner-up Le Pen was expected to lose the run-off to Macron and the remaining uncertainty ahead of Macron's eventual second-round win only raises redenomination risk by 2 basis points to 5 basis points over the second-round election weekend in May.

At its peak, the redenomination spread accounts for approximately 40% of the French CR14 CDS spread. Suppose, for illustrative purposes, that the (risk-neutral) expected recovery of bondholders in a redenomination scenario is 90% of face value (implying a 10% depreciation against the euro) and that the risk-free rate is 1%. Under these two assumptions, the simplified pricing equation (1.1) for q_T translates the redenomination spread into a back-of-the-envelope estimate of the redenomination probability. On April 21st, 2017, just before the first presidential election round, the redenomination spread of 0.21% translates into a risk-neutral probability of 2.21% that France will change the currency-denomination of its outstanding bonds within the following five years.

Figure A.7 plots the redenomination spread for Italy. The possibility of an Italian exit from the Eurozone (termed 'Italexit', 'Italeave', or—domestically—'Euroscita') has received a lot of attention during the formation of the coalition government supported by the populist *Five Star Movement* and the right-wing *League*. Ahead of the March 2018 elections, both parties had been strictly opposed to any form of cooperation, resulting in a hung parliament post-election. The election period itself is associated with mildly elevated levels of the redenomination spread, consistent with all relevant parties confirming Italy's Eurozone membership during their campaigns. However, during coalition negotiations in May, the question was raised, and a draft coalition agreement was leaked to the media, citing as objectives the *"introduction of specific technical procedures for single states to leave the Eurozone and regain monetary sovereignty"*, along with a request for €250bn debt relief from the ECB, and a radical reform of the *Stability and Growth Pact*.[13] While both parties immediately claimed the document was "outdated", the redenomination spread rises from 13 to 18bps on the day the draft leaked, and rises further over the following week of negotiations. The spread then jumps to 85 basis points at the end of May, amid further uncertainty surrounding the government formation, including the possibility of repeat elections within a few months. It stays above 60 basis points following the appointment and inauguration of the cabinet under Prime Minister Giuseppe Conte.[14] For both France and Italy, the difference-in-difference measure

[13]The draft document was published by HuffingtonPost.it on May15th, and is available here.

[14]Figure A.16 in Appendix A.4 plots the French and Italian *RS* against the respective *Economic*

evidently identifies redenomination-relevant political events. Figure A.8 plots the ratio of redenomination spread and the total CR14 CDS spread for France and Italy. While low on average, redenomination risk is at times economically large in magnitude, contributing up to 40% of the total CDS spread for France, and up to 32% for Italy. Having established an observable quantitative measure of redenomination risk, I now examine its association with yields and asset prices both in the country at risk of redenomination and elsewhere.

1.3 Redenomination risk and asset prices

This section documents a set of empirical results about the co-movement of different asset prices with the redenomination risk measure identified in the previous Section.

Eurozone sovereign debt.—To examine the relationship between redenomination risk and the cross-section of Eurozone yields, I collect yields for Austria (AUT), Belgium (BEL), Spain (ESP), France (FRA), Germany (GER), Ireland (IRE), Italy (ITA), the Netherlands (NED), Portugal (POR), and Denmark (DEN). I restrict attention to Eurozone countries, for which I have daily yield and CDS data (including bid-ask spreads), adding Denmark as a country with a fixed exchange rate against the euro throughout the sample period. I subtract the maturity-matched euro overnight swap rate (OIS) and then regress these yield spreads on the French and Italian redenomination spreads. Country j's yield with maturity T, $y_{j,T,t}$ is observed daily and the sample period ranges from September 2014 to June 2018:

$$y_{j,T,t} - OIS_{€,T,t} = \alpha_{j,T} + \beta_{FRA,j,T} RS_{FRA,t} + \beta_{ITA,j,T} RS_{ITA,t} + \varepsilon_{j,T,t}, \tag{1.2}$$

for $j = \{$AUT, BEL, ESP, FRA, GER, IRE, ITA, NED, POR, DEN$\}$ and $T = 5$ years. Table A.3 reports the results and Figure A.9 plots the β-coefficients with their 95% confidence intervals for the Eurozone countries.

The left panels show large cross-sectional variation in yield responses to French redenomination risk: the estimates are negative for German and Austrian government bond yields. Dutch, Irish, and Belgian responses are close to zero. Spanish yields rise, but not significantly at 5%. Portuguese and French yields rise sharply, with the French coefficient statistically indistinguishable from 1. The near-zero coefficient for Italian yields is misleading, since Regression (1.2) directly controls for Italian redenomination risk. Dropping RS_{ITA} from the regressors produces a strongly significant β_{FRA}-estimate of 1.698. Similarly, the coefficient is positive and

Policy Uncertainty index created by Baker et al. (2016), available at policyuncertainty.com.

strongly significant at 1.099 regressing the observable RS_{ITA} directly on French redenomination risk.

In comparison, the coefficients on the Italian redenomination spread—shown in the right panels—are significantly negative for all countries other than Italy. Further, the responses of other countries' sovereign yields to Italian redenomination risk, unlike for French risk, are also similar in magnitude. The Italian coefficient on Italian redenomination risk is indistinguishable from 1, just like in the case of France.

Another interesting place to look for responses to Eurozone redenomination risk is Denmark. Denmark has the right to opt-out of the eventual adoption of the euro under the Maastricht Treaty, and in 2000, the introduction of the euro was rejected in a public referendum (with 53.2% of votes in favor of retaining the krone). Nonetheless, the Danish krone (DKK) has been pegged to the euro under the European Exchange Rate Mechanism (ERM II) since 1999, which requires it to trade within 2.25% of 7.46038 kroner per euro. ERM II membership is one of criteria for a country to join the Eurozone, and, hence, the peg allows Denmark to keep the option of euro membership despite the opt-out. The tight peg de facto makes Denmark a Eurozone member as far as the risk and return characteristics of its sovereign bonds are concerned, with the crucial distinction that an 'exit' (i.e., abandoning the peg) is substantially simpler to implement for Denmark than for de jure Eurozone members. As a quasi-Eurozone member, Danish yields behave similarly to Austrian yields, with significantly negative coefficients on French and Italian redenomination risk.

Exchange rates.—Regarding the patterns of responses to French and Italian risk, a similar discrepancy exists in the response of the euro in currency markets. Denote by $e_{\$/€}$ the natural logarithm of the euro-dollar exchange rate defined as the \$-price of 1€. Consequently, an increase in this variable reflects an appreciation of the euro against the dollar. Similarly, $e_{€}$ and $e_{\$}$ denote, respectively, the logarithms of the euro index constructed by Bloomberg and the ICE US dollar index, each measuring the respective currency's value against a trade- and liquidity-weighted basket of global currencies. I obtain daily exchange rates from September 2014 to June 2018 and run the following time-series regressions:

$$e_{€,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{€} OIS_{€,t} + \varepsilon_t, \tag{1.3}$$

$$e_{\$/€,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{€} OIS_{€,t} + \gamma_{\$} OIS_{\$,t} + \varepsilon_t, \tag{1.4}$$

$$e_{\$,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{\$} OIS_{\$,t} + \varepsilon_t. \tag{1.5}$$

I report the results in Table A.4. The euro depreciates significantly against the dollar and a broader currency basket in response to higher French redenomination risk. The magnitudes of the coefficients indicate that a 1 basis point increase in the French redenomination spread is associated with a 0.3% lower euro exchange rate against the currency basket (0.5% against the dollar directly). In contrast, the euro exchange rate appreciates slightly but significantly by 0.1% on average against the currency basket for one basis point higher Italian redenomination risk. The US dollar appreciates significantly against the currency basket in response to French redenomination risk. The dollar index is not significantly correlated with Italian redenomination risk.

US Treasuries.—Next, I compare the sensitivity of German yields to redenomination risk to that of yields outside of the universe of €-denominated assets (or pegged, as in the case of Denmark), specifically US Treasury yields. I obtain daily Bund yields for maturities of 1, 2, 3, 5, and 10 years from Bloomberg, matching the sample period of the redenomination spread from September 2014 until June 2018, and run Regression (1.2) for Bund yields for maturities $T = \{1, 2, 3, 5, 10\}$. The coefficient estimates are reported in Panel A of Table A.5. The estimates for β are significantly negative for all but the 10-year maturity for the French redenomination spread, and for all maturities for the Italian redenomination spread. I then run the same time-series regression with US Treasury yields as the dependent variable, replacing euro swap rates with US dollar swap rates:

$$y_{US,T,t} = \alpha_T + \beta_{FRA,T} RS_{FRA,t} + \beta_{ITA,T} RS_{ITA,t} + \gamma_T OIS_{\$,T,t} + \varepsilon_{i,T,t}. \tag{1.6}$$

Panel B of Table A.5 reports the results for US Treasuries. Figure A.10 visualizes the comparison between Bunds and Treasuries from Table (A.5) by plotting the regression coefficients and their 95% confidence intervals across the term structure. The response of US Treasuries to French redenomination risk is similar to that of German Bunds. The coefficients for German yields are more negative than those for US yields, but the 95% confidence intervals overlap slightly across most of the term structure, so the distinction in magnitudes resides in the margins of statistical significance. Focusing next on the estimates for $\beta_{ITA,T}$ in the right panel of Figure A.10, dollar-denominated US Treasuries behave differently from euro-denominated Bunds. Much like most other euro-denominated sovereign yields, Bund yields tend to fall significantly in times of high Italian redenomination risk. This is not true for US Treasuries of any maturity.

Corporate credit spreads.—I now extend the above examination of redenomination risk to corporate credit spreads across Europe. To this end, I collect five-year CDS spreads (denominated in euros, and with CR14 restructuring clauses) for the 125 European companies included in the iTraxx Europe Index—a tradable CDS-index of the most liquid European corporates with an investment-grade rating. These credit spreads refer to senior unsecured bonds issued by these 125 corporates. In addition, I collect five-year subordinated credit spreads for 30 financial corporates (banks and insurance companies included in the 125 sampled companies, for which these subordinated CDS are traded separately). I repeat Regression (1.2), substituting as the dependent variable (i) the portfolio of 125 senior corporate CDS (iTraxx Europe), (ii) the portfolio of 30 senior financial CDS (iTraxx Financials Senior), (iii) the portfolio of 30 subordinated financial CDS, and (iv) 10 portfolios of corporate CDS spreads sorted by country and split into financial and non-financial companies. These country portfolios span five Eurozone countries (GER, NED, ITA, ESP, and FRA), with at least one financial company within the original set of 125. These countries cover 71 of the original 125 individual companies. All portfolios are equally weighted. The results are reported in Table A.6 and illustrated in Figure A.11.

Broadly speaking, the results are weaker than those for sovereign yields in Table A.3: corporate credit spreads across Europe are positively associated with French redenomination risk, but this association lacks statistical significance for Spain and non-financial companies in Italy. The point estimate for German non-financial corporates is negative and insignificant. Among non-financial companies, the response to French risk is strongest for French corporates.

Just like sovereign yields, corporate credit spreads are negatively associated with Italian redenomination risk, and the results are significant for the non-financial companies in Germany, the Netherlands, Spain, and France. However, in contrast to Italian sovereign yields, non-financial Italian corporate spreads react only marginally, and insignificantly positively to Italian redenomination risk.

The coefficients are generally more positive for the CDS spreads of financial companies, and even more so for their subordinated CDS spreads. Credit spreads of Italian financial companies are significantly positively associated with French redenomination risk, unlike their non-financial counterparts. With respect to Italian redenomination risk, only Dutch and French banks have significant negative coefficients.

1.3.1 Redenomination risk vs. credit risk

Ultimately, I consider the co-movement of Eurozone sovereign yields and the two redenomination risk measures to document signs of redenomination risks in sovereign debt beyond France and Italy. In regressing sovereign yields on redenomination risk, the idea is to think of the yield as the sum of different components:

$$y_{j,T,t} \approx \text{risk-free rate}_{€,T,t} + \text{credit risk}_{j,T,t} + \text{redenomination risk}_{j,T,t}, \tag{1.7}$$

where *credit risk* is meant to capture all default risk unrelated to redenomination. The risk-free rate is meant to include all euro-wide return components (e.g., a term premium). I account for the latter using maturity-matched swap rates on the right-hand side of the yield regressions. For the three G7-Eurozone members, the redenomination spread presented in Section 1.2 measures [redenomination risk]$^+$, that is, the positive part of redenomination risk. The asymmetry stems from the fact that CDS contracts cover only losses from credit events. A redenomination to the benefit of bondholders would therefore not trigger CDS payouts.

However, regarding the credit risk component, finding a suitable measure is more difficult: for all non-G7 Eurozone members, CDS spreads measure the sum of credit risk and [redenomination risk]$^+$, irrespective of which contract type I consider. An isolated observable measure of credit risk only exists for the three G7-countries. Since the CR contracts do not cover redenomination into a G7-currency, these spreads are a suitable measure for conventional credit risk, *excluding* redenomination. As a sense-check for this decomposition, I regress French and Italian five-year yields on the five-year swap rate, each country's respective redenomination spread, and its CR spread:

$$y_{j,T,t} = \alpha_{j,T} + \beta_{j,T} RS_{j,t} + \psi_{j,T} CR_{j,t} + \gamma_{j,T} OIS_{€,T,t} + \varepsilon_{j,T,t}, \tag{1.8}$$

for $j = \{FRA, ITA\}$. The results are reported in Table A.8. For both countries, the ψ-coefficients are indistinguishable from one and R^2 are high at 0.93 and 0.85, respectively. Since $RS_{ITA} > 0$ for the vast majority of the sample period, the limitation that the CDS-based measure only captures positive redenomination risk has little bearing. Accordingly, the coefficient β_{ITA} is indistinguishable from one. For France, this coefficient is statistically below one over the full sample, but becomes indistinguishable from one once I drop the earlier part of the sample (pre-February 2017), when RS_{FRA} hovers around zero but appears to be noisy.

The only other country, for which credit and redenomination risk are directly observable in isolation is Germany. For Germany, however, the limitation that $RS_{GER} = [\text{redenomination risk}_{GER}]^+$ becomes more problematic, as RS_{GER} is essentially zero throughout the sample. This raises the concern that the CDS-based measure fails to capture a negative redenomination risk component in Bund yields arising from expected currency gains conditional on redenomination. Table A.8 reports the results for a variant of Regression (1.2), now including a direct control for German credit risk (the CR spread). The coefficients on French and Italian redenomination risk remain significantly negative, and increase in magnitude relative to those reported in Table A.3. Adding credit risk raises the R^2 of the yield regression by 15 percentage points to 0.59. For all other Eurozone members, credit spreads or other observable variables do not control for credit risk in isolation from redenomination risk.

Taking a different approach, I compare the response coefficients of yields to redenomination risk to those of credit risk, each measured in isolation for France and Italy. To this end, I add to the redenomination risk measures in Regression (1.2) the French and Italian CR credit spreads and run an analogous set of time-series regressions.

$$\begin{aligned} y_{j,T,t} - OIS_{€,T,t} = \alpha_{j,T} + \beta_{FRA,j,T} RS_{FRA,t} + \beta_{ITA,j,T} RS_{ITA,t} + \\ + \psi_{FRA,j,T} CR_{FRA,t} + \psi_{ITA,j,T} CR_{ITA,t} + \varepsilon_{j,T,t}, \end{aligned} \tag{1.9}$$

Table A.7 and the middle and lower panels of Figure A.9 report the results. The pattern in the β-coefficients remains. Unlike for redenomination risk, sovereign yields are positively and significantly associated with French credit risk. The exceptions to this are Danish and Portuguese yields, which both have insignificant coefficients. In stark contrast to the results from Regression (1.2), German Bund yields show the largest positive response among five-year yields. Similarly, the coefficients on Italian credit risk differ from those on Italian redenomination risk. While the responses to redenomination risk are significantly negative for all countries other than Italy, the credit-risk correlations are close to zero, with the exceptions of Italy and Portugal. The latter yields rise significantly with the reaction even larger in magnitude than that of Italian yields. These results confirm the notion that redenomination risk—as measured by the approach introduced in this chapter—is genuinely distinct from non-redenomination credit risk.

Since CR spreads and redenomination spreads are close to uncorrelated within-country for France and Italy, the β-estimates for Regressions (1.2) and (1.9) do not differ substantially.

As a further robustness check, I limit the sample to the years 2017 and 2018, where the most substantial variation in redenomination risk occurs. The subsample results in Table A.9 confirm the findings from the headline regressions: loadings on French redenomination risk vary widely across countries. The β_{FRA} estimates are significantly negative for Germany and Austria, and significantly positive for France and Portugal. Again, the Italian estimate for β_{FRA} is misleadingly low, since the regression controls for RS_{ITA} directly. Once the Italian regressors are dropped from the regressors, the coefficient jumps to 0.71. In comparison, the estimates for β_{ITA} are close together, ranging from -0.51 (ESP) to 0.46 (BEL); the Portuguese coefficient is more negative at -1.65. Crucially, the β_{ITA}-estimates do not share the cross-sectional dispersion seen for β_{FRA}.

1.3.2 Negative redenomination risk

I return to the previous observation that—for all countries—CR14 credit spreads measure the sum of credit risk and the positive part of redenomination risk, i.e.,

$$CR14_{j,T,t} = \text{credit risk}_{j,T,t} + [\text{redenomination risk}_{j,T,t}]^{+}$$

To examine the asymmetry in redenomination risk more closely, I repeat Regression (1.2) with each country's five-year CR14 spread as the dependent variable, instead of its yield:

$$CR14_{j,t} = \alpha_j + \beta_{FRA,j} RS_{FRA,t} + \beta_{ITA,j} RS_{ITA,t} + \varepsilon_{j,t}, \tag{1.10}$$

for $j = \{\text{AUT, BEL, ESP, GER, IRE, NED, POR}\}$. I drop the observations for France and Italy from the dependent variables, as both variables are mechanically included in the construction of the redenomination spreads on the right-hand side. If the negative yield responses in other Eurozone government bonds reflect negative redenomination risk, these responses will be absent from the CDS spreads and the resulting coefficients are bounded below by zero. I report the results in Table A.10 and Figure A.12 illustrates the comparison of the point estimates for credit spreads with those for yields. As shown in the left panel, the point estimates for $\beta_{FRA,j}$ are indeed non-negative. While bond yields for Germany, Austria, and the Netherlands have negative point estimates for their respective correlations with French redenomination risk, the coefficients for their respective CDS spreads are all significantly positive. In stark contrast, just like the coefficients for yields, all coefficients on the Italian redenomination spread are significantly negative, albeit generally smaller in magnitude than the yield coefficients.[15] This result suggests that the negative β_{ITA} coefficients in Regression (A.3) do not reflect negative redenomination risk across other Eurozone bonds, but instead stem from changes in credit risk premia. The comparison between the two panels points once more to the systematic distinction between French and Italian redenomination risk and their associations with asset prices. In the next section, I provide an economic rationale for the joint set of the above results.

1.4 Contagion, safety, and substitution

When measuring redenomination risk and its co-movement with asset prices, the imminent question is whether a hypothetical Eurozone exit of any given country is likely to be associated with a break-up of the currency union, or whether such an exit would remain isolated. To make this question more empirically tangible, this section lays out a simple model to sketch possible spillover effects across a range of asset prices in response to the risk of each type of exit—contagious or isolated. The redenomination risk measure presented in Subsection 1.2 of this chapter quantifies exit risk for France and Italy and I interpret the empirical results presented in Section 1.3 as symptoms of spillovers from this exit risk to other asset prices.

The prospect of a Eurozone break-up and re-introduction of national currencies sparks capital flight out of countries with expected weaker national currencies and into those with stronger ones. As outlined at the start of Section 1.2, the redenomination risk measure presented in this chapter only captures the downside of redenomination, that is, if the national shadow currency of, say, France is expected to depreciate against the euro. As a consequence of the repricing of redenominatable bonds at risk of such depreciation, the nominal yields on such bonds rise. As a first consistency check, a positive redenomination spread for France should therefore be associated with higher yields for French sovereign bonds. The results reported in Table A.3 and Figure A.9 show that this is true for both France and Italy (with yields rising close to 1-for-1 with redenomination risk).

Looking beyond France or Italy and towards spillover effects, the question of contagion versus isolation becomes crucial: if an exit of, say, Italy becomes more likely, but this event is not expected to lead to a redenomination of bonds issued by, say, Spain, this may lead investors in Eurozone government bonds to shift their investments out of Italian bonds and into Spanish ones, regardless of how the Spanish

[15] The findings in Table A.10 are robust to estimating regression (1.10) as a Tobit model.

national shadow currency would fare, because that currency remains hypothetical in the absence of break-up risk. To illustrate this channel more formally, consider the simple model presented in the following subsection.

1.4.1 A simple model

There are two dates, *today* and *tomorrow*, and the model describes the bond market in a currency union of three countries, A, B, and H. On the supply side of the bond market, the asset universe consists of four zero-coupon bonds: a risk-free bond with a net supply of σ_S in nominal face value, and three redenominatable government bonds issued by countries A, B, and H in nominal net supplies σ_A, σ_B, and σ_H, respectively. Today, all bonds are denominated in the common numéraire (let's call it 'euro') and their prices are determined by market clearing. Prices are expressed in terms of gross yield denoted by y_J for bond J, such that its price per unit of face value is $P_J = 1/y_J$. Countries A and B are individually at risk of exiting the currency union and redenominating their bonds into a national currency. Country H, however, only redenominates its bonds if both A and B jointly exit, that is, if the currency union ceases to exist. Consequently, there are four possible states of the world tomorrow, denoted by $s \in \{1, 2, 3, 4\}$:

(1) Stability: no exit, no bond is redenominated,

(2) Isolated exit A: only A is redenominated,

(3) Isolated exit B: only B is redenominated,

(4) Break-up: all bonds, A, B, and H, are redenominated.

In case of redenomination, the face value is repaid in the new currency worth a euro-equivalent of $(1 - \delta_J)$ per unit, such that the gross return on bond J in case of redenomination is $y_J(1 - \delta_J)$. $\delta_A > 0$ and $\delta_B > 0$, meaning that currencies A and B depreciate against the euro-numéraire once they are introduced. In contrast, the new currency of country H (for *haven*) appreciates, $\delta_H < 0$, resulting in exchange rate gains from redenomination for bondholders. The risk-free bond denoted by subscript S repays one unit of the numéraire per unit of face value in all states of the world. This bond can be thought of as a privately issued euro-denominated security with sufficient collateral to be default-free and remote from redenomination. With four linearly independent assets and four states of the world, markets are complete.

The demand side of the asset market consists of two risk-averse banks, a and b operating in countries A and B, respectively. Adding a third bank operating in country H does not change any of the model results in a meaningful way. For clarity of notation, I use lower case superscripts to refer to banks, and upper case subscripts to refer to countries/bonds. Today, the right-hand side of each bank's balance sheet consists of deposits, d raised from households in the respective country, and bank equity, e, such that the total endowment of each bank amounts to one unit of the common numéraire as shown below. Households are passive, and their decisions are not modelled.

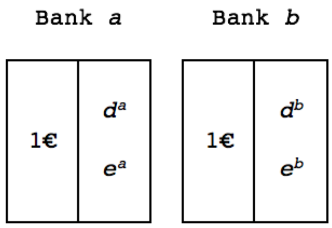

Crucially, redenomination also extends to bank deposits: the euro-equivalent of deposits taken by bank a falls to $d^a(1-\delta_A)$ after redenomination by country A, and equivalently for bank b and country B.

Today, banks choose a portfolio of the four assets in order to maximize expected log utility over their respective equity tomorrow. Let w_J^i be the euro-investment of bank i in bond J, and by e_s^i the value of bank i's equity in state s. State-probabilities are denoted by p_s. I assume that deposits, d, and redenomination losses, δ, are sufficiently small, such that bank equity is strictly positive in all states and utility is well-defined:

$$max_{\{w_A^i, w_B^i, w_H^i, w_S^i\}} \sum_s p_s \log\left(e_s^i\right) \quad \text{s. t.} \quad w_A^i + w_B^i + w_H^i + w_S^i = 1$$

Rather than to *generate* contagion in redenominations, the purpose of the model is to formally examine the relationships of the different asset prices *given* contagion or the lack thereof. Starting with the latter case, *isolation*, suppose that exits by A and B are independent, and the redenomination probabilities are ρ_A and ρ_B, respectively. The probabilities of the four possible states in the isolation case are:

(1) Stability: $p_1 = (1-\rho_A)(1-\rho_B)$,

(2) Isolated exit A: $p_2 = \rho_A(1-\rho_B)$,

(3) Isolated exit B: $p_3 = (1-\rho_A)\rho_B$, and

(4) Break-up: $p_4 = \rho_A \cdot \rho_B$.

Next, I consider the other extreme case: the *contagion* case with perfect correlation in redenominations. To this end, suppose that B exits and redenominates if and only if A does, such that the state probabilities in the contagion case are:

(1) Stability: $p_1 = (1 - \rho_A)$,

(2) Isolated exit A: $p_2 = 0$,

(3) Isolated exit B: $p_3 = 0$, and

(4) Break-up: $p_4 = \rho_A$.

In this configuration of the contagion case, country A drives the disintegration of the currency union and the model therefore studies spillover effects from A to B, but not vice versa. The equilibrium in this model exhibits the following relationships between redenomination risk (ρ_A) and bond yields.

(i) Spillover effects.—In the isolation case (with independent redenominations), an increase in A's redenomination probability, ρ_A, *lowers* the yield on country B's bonds. This result is illustrated in terms of comparative statics of equilibrium yields with respect to ρ_A in Panel A of Figure A.13. It is an indirect spillover effect through portfolio substitution: rising risk in country A lowers yields in country B, because absent a change in yields, both banks shift portfolio weight from country A's bonds to those of country B (and those of country H, and the risk-free bond). Yields on bond B therefore need to fall to restore market clearing. In the contagion case, however, the sign and magnitude of spillover effects on another country's bond yield from an increase in ρ_A are dictated by, respectively, the sign and magnitude of the other country's δ: since $\delta_B > 0$, country B's bond yield increases with redenomination risk in country A, while the yield on the bonds of country H falls ($\delta_H < 0$). Panel B of Figure A.13 illustrates the yield spillovers in the contagion case. Aside from the spillover effects, the model delivers two additional results, which are notable in the empirical context of the Eurozone.

(ii) Home bias.—Sovereign bonds are predominantly held by domestic banks. Battistini et al. (2014) note that the redenomination of liabilities gives domestic banks a "comparative advantage" in holding domestic sovereign debt.[16] This is precisely the mechanism behind this model result, which is a direct consequence of deposit redenomination: the losses from a redenomination of domestic government bonds on the bank's asset side are partially offset by the redenomination of its deposits. Accordingly, domestic bonds are less risky to domestic banks than to

[16] Alongside redenomination risk, they note two primary motives for home bias in Eurozone banks: (i) "moral suasion" by authorities in order to raise demand for domestic sovereign debt; and (ii) "carry trade" investments into particularly high-yield euro-denominated sovereign debt, funded with low-yield euro borrowing (see also Acharya and Steffen (2015)).

foreign banks, resulting in home bias in bank bond holdings. The proof is left to Appendix A.3. I document home bias in Table A.11 using data as of year-end 2015, provided by the EBA: banks domiciled in most European countries hold a larger fraction of their liquid sovereign debt holdings in domestic government debt, in which most of their deposit-taking activity occurs. For non-Eurozone countries, such as Poland (98.8% of net sovereign bond exposure of Polish banks is to the Polish government) or Norway (63.2%), this likely represents a straight-forward currency matching between assets and liabilities. For Eurozone-domiciled banks, for instance in Italy (65.8%), Ireland (67.6%), Spain (50.5%), or Germany (44.4%), redenomination risk makes this currency matching more subtle and currency bias implies home bias even in a currency union.[17]

(iii) Sub-zero lower bound.—The nominal yield on bond H is below that of the risk-free asset. This effect is straight-forward if redenomination leads to exchange rate gains. Bonds from a country whose currency is expected to appreciate in the break-up scenario carry a yield below the risk-free rate. Even if the risk-free rate is bounded below (say, by zero), 'haven' bond yields are not. This intuitive notion is important for the assessment of monetary policy and its transmission in the presence of redenomination risk and negative bond yields. Again, the proof is left to Appendix A.3.

1.4.2 Interpreting the empirical results

I now compare the model results to the empirical results in Section 1.3. The spillovers through portfolio substitution described above apply to all other assets in the model (all sharing the common numéraire). The right-hand side panels of Figure A.9 show that the statistical relationship of the Italian redenomination spread with Eurozone government yields outside Italy is indeed significantly negative and homogeneous in the cross-section. As shown in Table A.6, the same is true for corporate credit spreads in Germany, the Netherlands, Spain, and France. Dollar-denominated US Treasuries do not exhibit the same behavior as Eurozone yields and remain flat across most of the term structure with respect to increases in Italian redenomination risk (Figure A.10). Furthermore, the euro-dollar exchange rate is uncorrelated with Italian redenomination risk. Against a broad currency basket, the euro appreciates slightly with Italian redenomination risk. The results for Trea-

[17]Among Eurozone-domiciled banks, home bias is relatively low for the two Austrian banks included in the EBA stress tests. Both have relatively large exposures to central and eastern European sovereigns, consistent with their prominent consumer banking presence and deposit base in that region.

suries and exchange rates speak against the notion that Italian redenomination risk is associated with a broader flight-to-safety phenomenon.

In contrast, spillovers from contagious redenomination risk separate the remaining Eurozone members—observably, through the reaction of their sovereign bonds—into those with expected strong currencies and those with expected weak national currencies. Bonds that are redenominated into a stronger national currency (or stronger miniature-currency unions) are more desirable, and these bonds appreciate. If, say, the new German currency is expected to appreciate against other national euro-successor currencies, then 'German euros', which are converted in the event of a Eurozone break-up, provide an effective hedge against the break-up event and exhibit 'safe haven' properties. In a similar way, safe haven candidate assets denominated in other currencies, such as US Treasuries, might benefit similarly, if the future of the euro is at risk. Instead, countries, for which national currencies are expected to be weaker exhibit rising yields. The implied cross-sectional heterogeneity in bond yield responses is evident in the left panels of Figure A.9, which plots the yield-coefficients with respect to the French redenomination spread. Similarly to German Bunds, US Treasury yields—a plausible safe-haven asset in the case of a Eurozone break-up—drop with rising French redenomination risk, as shown in Figure A.10. The interpretation that the negative response of German and Austrian yields to the French RS-measure reflects negative redenomination risk in these countries is corroborated by the absence of this response in German and Austrian CDS spreads: since CDS contracts only cover losses from credit events, their prices reflect redenomination risk asymmetrically, unlike bond yields which reflect both expected losses and gains from redenomination.[18] Table A.4 further points to the negative association of bilateral euro exchange rates with the French redenomination spread: the euro depreciates significantly, consistent with the interpretation that a French redenomination would put the existence of the euro at risk.

The results for corporate credit spreads in Table A.6 are weaker in magnitude and significance than those for sovereign yields. Credit spreads of Italian companies (financial as well as non-financial) are not significantly correlated with Italian redenomination risk with coefficients close to zero, while their sovereign counterparts show significantly positive coefficients close to one. This result suggests that an isolated Eurozone exit would not necessarily imply currency redenomination for the

[18]An important caveat in the comparison between yields and CDS spreads is the large literature on the CDS-bond basis (Bai and Collin-Dufresne, forthcoming, e.g.) and the potential disconnect between sovereign yields and CDS spreads due to financial regulation and the price impact of financial institutions in CDS markets (Antón et al., 2017; Klingler and Lando, 2018).

debt of large domestic corporate borrowers. Therefore, the ability of the diff-in-diff measure to identify sovereign redenomination risk without assumptions on corporate redenomination marks an important contribution of this chapter relative to Krishnamurthy et al. (2018) and Bayer et al. (2018). Redenomination of corporate debt is more likely in a break-up scenario, where the common currency ceases to exist. In line with this interpretation, French corporate credit spreads rise significantly with French redenomination risk. Cross-sectional patterns in the coefficients on credit spreads outside France or Italy are difficult to interpret, due to the differing size and industry composition of the country portfolios. With this caveat, I note that the β_{FRA}-coefficients are smallest for German corporates (financial as well as non-financial), mirroring some of the cross-sectional pattern observed for sovereign debt. Overall, a within-country comparison of the β_{FRA} and β_{ITA} coefficients suggests once more that the risk of a French exit from the Eurozone is priced more severely than that of an Italian exit in European corporate CDS markets.

As an additional test, the risk of a contagious redenomination in one country should—by virtue of being contagious—also be correlated with redenomination risk in other countries, and, therefore, with the observable redenomination spread. Throughout the sample period over which I can observe redenomination spreads, the Italian measure is high whenever the French redenomination spread is high, but not vice versa, consistent with contagious French risk and isolated Italian risk (after accounting for the French component in Italian risk). The German spread is essentially zero throughout, and this exception is consistent with the inability of the measure to capture negative redenomination risk, that is an expected appreciation of a country's national currency following the euro break-up. The hypothesis that this applies to a new German mark is further consistent with the behavior shown by German sovereign yields and CDS spreads with respect to French redenomination risk.

All of the empirical results presented in this chapter are, therefore, consistent with the interpretation that the redenomination risk in French CDS around the presidential elections in 2017 was deemed contagious by market participants, while the risk measured from Italian CDS immediately after the French election, ahead of the Italian elections in March 2018, and particularly following the formation of the coalition government in May 2018, was not expected to spill over into other Eurozone countries. Interpreting the exposures of sovereign yields to French redenomination risk as indicators of the strength of each national shadow currency, I next relate these coefficients to fundamental country-level variables.

1.4.3 National shadow currencies

Which factors explain the cross-section of sovereign yield reactions to contagious redenomination risk? In the simple model set-up, this reaction is pinned down by paramenter δ_J, the new national currency's depreciation relative to the euro (or relative to the new national currencies of other currency union members). To be more precise, the change in bond prices reflects the change in expected losses or gains from currency redenomination. This expectation is the product of the probability of redenomination in country B, conditional on redenomination in country A (= 1 in the contagion case of the model), and the expected depreciation of country B's national currency after redenomination (= δ_B). It is natural to ask which factors determine the heterogeneity in δ, that is heterogeneity in national exchange rates immediately following the break-up of the currency union.[19] In line with the evidence and interpretation presented above, suppose that the time series for French redenomination risk reflects the risk of a Eurozone break-up. This simplifying assumption echoes the contagion case of the model, such that for all countries the probability of redenomination conditional on redenomination in France is equal to 1, and the cross-sectional heterogeneity in yield responses to French redenomination risk is driven by heterogeneity in δ_J across countries, that is, heterogeneity in the performance of the different national currencies immediately following the dissolution of the currency union.

Consider the β_{FRA}-estimates in Table A.7 (from Regression (1.9), which controls directly for credit risk): German and Austrian yields have the most negative coefficients, followed by Danish and Dutch yields. As in the baseline Regression (1.2), Portuguese and French yields have significantly positive β_{FRA}-estimates. Dropping the Italian regressors, the β_{FRA}-estimate for Italy jumps to 0.71. If these coefficients reflect an expected appreciation of, say, a new German currency or of a de-pegged Danish krone relative to a new Italian or Portuguese currency, then what is it about Germany or Denmark (Italy or Portugal) that promises a strong (weak) national currency after the break-down of the peg enforced by the currency union?

To relate the regression coefficients to country-fundamentals, I run univariate cross-sectional regressions of the β_{FRA} coefficients from Regression (1.9) on the (i) debt-to-GDP ratio, (ii) budget surplus/deficit, (iii) labor productivity, and (iv) cur-

[19] As outlined in Subsection 1.1.1, losses from redenomination may stem from an increase in credit risk premia for the respective country outside of the Eurozone, alongside the depreciation of the new numéraire. Without imposing strong further assumptions on these only indirectly observable quantities, it is impossible to disentangle the different sources of redenomination losses.

rent account balance. Data on all four variables are obtained from the OECD. Debt-to-GDP ratios and labor productivity data (GDP per hour worked) are as of 2016, due to incomplete data for 2017. All other data are for 2017. Budget surplus, and current account balance are expressed as a percentage of GDP.

For each country, Figure A.14 plots the point estimates for $\beta_{FRA,j}$ from Regression (1.9) against each of the five fundamental variables, along with the univariate R^2 of the respective cross-sectional regression. Sovereign debt and the 2017 budget surplus (both scaled by GDP) each linearly account for large shares of the cross-country variation in β_{FRA}: 0.66, and 0.72, respectively. (These two fundamental variables are also strongly correlated across countries.) Labor productivity and the current account balance deliver univariate R^2 of 0.25 and 0.32, respectively. Since a negative β_{FRA} suggests a strong national currency, it is not surprising that the only variable that is positively related to the coefficient is the debt-to-GDP ratio. A strong link between a country's fiscal position and the value of its currency is in line with the long literature on the fiscal theory of the price level (e.g., Sargent and Wallace (1984) and Sims (1994), or—in more recent applications—Jiang (2018) and Bolton and Huang (2017)), and the high univariate R^2 are consistent with the interpretation of β_{FRA} as a weakness-gauge for the national shadow currencies of Eurozone members. In contrast, the limited variation in the coefficients on Italian redenomination risk, $\beta_{ITA,j}$, does not relate to variation in macro fundamentals. Sovereign debt yields a modest R^2 of 0.10; the R^2 for the three other fundamental variables are essentially zero.

Taking the interpretations from Subsection 1.4.2 at face value, I now proceed to a back-of-the-envelope calculation of the fiscal costs to different national treasuries that are attributable to the periods of heightened redenomination risk in France and Italy.

1.4.4 Fiscal contagion or 'exorbitant' privilege?

The quantitative measure of redenomination risk can be used to estimate the overall fiscal cost of French and Italian redenomination risk on these two countries as well as the cost of spillovers on other Eurozone members. As discussed above, the sign of yield spillovers varies by country, such that German taxpayers benefit from the risk of redenomination in France, while sovereign yields for most other Eurozone member countries rise. As the de facto provider of safe assets for the Eurozone, the German treasury collects insurance premia in the form of interest savings on newly issued debt. The role of Germany as an insurance provider against redenomination risk can be viewed in analogy to the role of the United States as a provider of safe assets and the US dollar as the reserve currency within the global financial system, which has been described as an "exorbitant privilege" by French then-Minister of Finance, Valéry Giscard d'Estaing in the 1960s, and is interpreted as that of an insurance provider by Gourinchas et al. (2010).

To quantify the impact of positive and negative spillovers on yields, I compute counterfactual yield curves for each sample country on each day from 2017 until the end of the sample period in June 2018 as if redenomination spreads of France and Italy were zero throughout. I restrict attention to the sample period with most of the time-variation in redenomination risk. Specifically, I compute the counterfactual yield $\tilde{y}_{j,T,t} = \hat{\alpha}_{j,T,t} + \hat{\gamma}_{j,T,t} OIS_{€,T,t} + \hat{\varepsilon}_{j,T,t}$ using estimates from rolling-window regression analogous to Regression (1.2) for maturities $T = \{1, 2, 3, 5, 10, 30\}$ years. Estimating rolling coefficients allows for time-variation in the sensitivity of bond yields to redenomination risk. I choose a window-length of 250 daily observations up to and including observation t. I then compute the estimated spillover costs (or cost savings, if negative) as $c_{j,T,t} = y_{j,T,t} - \tilde{y}_{j,T,t} = \hat{\beta}_{FRA,j,T,t} RS_{FRA,t} + \hat{\beta}_{ITA,j,T,t} RS_{ITA,t}$, which reflects an estimate of the yield component that is due to French and Italian redenomination risk. For each bond issuance, I then multiply $c_{j,T,t}$ by the issuance volume, $v_{j,T,t}$ and capitalize the differential interest costs over the maturity of the bond with an annuity factor to obtain $C_{j,T,t} = c_{j,T,t} \cdot v_{j,T,t} \cdot a(T, y_{j,T,t})$. Since the coefficients are estimated for nominal yields, I exclude inflation-linked bond issuances.

Crucially, this exercise assumes that the (plausibly endogenous) choice of issuance volume and maturity is fixed. If national treasuries adjust issuance volume and/or maturity to changes in yields, my back-of-the-envelope cost estimate will be biased downwards, since the unobservable counterfactual issuance choice would have resulted in higher interest costs than the observable optimized issuance. Given the high frequency of changes in redenomination risk, which is characterized by sudden jumps and few periods of sustained elevated levels over the sample period, an adjustment of issuance by national treasuries would have to occur rather quickly. The idea that issuance volumes are chosen in response to changes in redenomination risk is also at odds with the overall low amount of issuance by the German treasury, which benefits most from redenomination risk in this sample. Figure A.15 plots these costs for each bond auction of France (Panel A), Italy (B), Spain (C), and Germany (D). To visualize the time-variation in the intensity of redenomination risk, I include the two redenomination spreads in each plot. The aggregate measure can be decomposed into a French and an Italian component, $C_j = C_j^{FRA} + C_j^{ITA}$.

The result of this back-of-the-envelope calculation is that, from 2017 until June 2018, French taxpayers incur a substantial fiscal cost from redenomination risk: the risk surrounding the presidential election in 2017 resulted in substantially increased interest costs of around €400m. However, these costs to taxpayers are more than offset by the benefits newly issued French bonds have subsequently reaped as a substitute to Italian bonds during periods of Italian redenomination risk. The net benefit estimate to French taxpayers from redenomination risk in 2017-2018 amounts to €858m.

With the exception of a few very short-term issuances, Italian debt issues carried higher interest rates over the period, both during the tumultuous run-up to the French election in March/April 2017, but particularly following the Italian elections in March 2018 that led to the formation of the coalition between the far-left Five Star Movement and the far-right League in late May 2018. The estimated interest cost from redenomination risk to Italian taxpayers totals €3.5bn.

Spain is a net beneficiary over the period, despite negative spillovers and higher yields ahead of the French election (with costs totaling around €40m). Following the Italian elections, Spanish yields were negatively correlated with rising Italian redenomination risk, leading to "cheap" debt issuances and an estimated net fiscal benefit of €499m over the entire period.

As a provider of 'safe' assets, the German treasury benefitted sizeably from the risks surrounding the French election, with interest savings of around €280m in early 2017. The total estimated net benefit to German taxpayers over the period from January 2017 through June 2018 amounts to €565m. To put this number into perspective, I note that nominal bond issuance by the German treasury over those 18 months totaled €205bn. It is important to note that the direct fiscal costs, which may appear minuscule in relation to the trillions of euros of outstanding sovereign debt, are computed on the basis of newly issued debt only. At the same time, the risks measured over the sample period suggest event probabilities of a few (single-digit) percentage points under the risk-neutral measure (i.e., an upper bound on the true, physical probability). The fact that such small probabilities have consequences of economically meaningful magnitude highlights the need for investors, policy makers, and electorates alike to understand the full ramifications a Eurozone exit, not to mention a break-up of the currency union.

Next, I discuss potential empirical concerns with the difference-in-difference approach used to quantify redenomination risk.

1.5 Concerns in measuring redenomination risk

The difference-in-difference measure of redenomination risk is obtained from the relative behavior of CDS contracts based on differing definitions of credit events. The most pressing concern when comparing new contracts to old ones is one of liquidity differences: similar to on-the-run/off-the-run premia in the US Treasury market, the potential lack of liquidity for off-the-run CDS contracts may result in different CDS spreads, and the consistently positive difference between CDS spreads for the largely unaffected issuers in the control group suggests that such a liquidity component exists. I also describe the other important contract change that applies to sovereign CDS and why my approach deals with it appropriately.

1.5.1 Liquidity

The diff-in-diff will account for liquidity-driven differences between old and new contracts, as long as such differences are common across treatment and control groups. Liquidity differences are likely to be more severe in smaller markets. Since both France and especially Italy are among the largest European CDS markets, the control group is more likely to overstate the correction due to market-size driven liquidity.

However, the additional distinction between the two types in treated issuers may create clientele effects that generate price differences between CR and CR14 contracts as some investors shift holdings from CR to CR14 contracts and the market for CR adjusts to the new clientele. Such adjustments in market clientele are not purely driven by an on-the-run versus off-the-run phenomenon, and would, therefore, be systematically different between treatment and control groups. However, such adjustments are also likely to be temporary, if the launch of CR14 contracts was widely anticipated. Transitory price effects driven by the adjustment of market clientele to the newly bifurcated market may be responsible for the elevated Italian redenomination spread in October to November 2014 following the introduction of CR14 contracts. The spread then goes back to hover around zero.

Anticipation plays a problematic role in the interpretation of many difference-in-difference measures. In this case, however, anticipation does not threaten the validity of the diff-in-diff, since both treated and untreated variables are observed simultaneously (i.e., the diff-in-diff is not across time). On the contrary, for my diff-in-diff measure to reveal redenomination risk, it is necessary that market participants are immediately and fully aware of the differences between the two contract types and price them accordingly. To show that this was likely the case, I briefly summarize the timeline of the revision process.

ISDA began the revision of its CDS definitions in May 2012, following the restructuring credit event in Greece. In November 2013, ISDA published a draft of the revised definitions to review comments from market participants ahead of the final release of the new definitions on February 21st, 2014. Trading in the new contracts began 7 months later on September 22nd.[20] The release of the new definitions in February also announced the implementation process for the new set of definitions. For the vast majority of reference entities, the changes were retroactively applied to existing contracts on October 6th, but due to the expected pricing impact of the sovereign-specific changes (see also Subsection 1.5.2), most sovereign issuers were excluded from this adjustment such that CR contracts remained widely outstanding alongside the newly issued CR14 contracts. Among sovereign issuers, existing contracts were only migrated to the new 2014 definitions for emerging market sovereign issuers because ISDA was concerned that the resulting lack of liquidity in legacy CR contracts would be insufficient to support efficient trading in a bifurcated market (Simmons & Simmons, 2016). At the same time, there were no such liquidity concerns for developed-market sovereign issuers, including all Eurozone countries studied in this chapter. Due to the broad consultation of market participants in the revision process and the long lead time between the release of the final new definitions and the beginning of trading, it is reasonable to assume that, at the time the CR14 contracts were launched, market participants were immediately and fully aware of the differences, and prices reflect these differences throughout my sample period. ISDA's decision to exclude most sovereign reference entities from a retroactive activation of the new definitions suggests that market liquidity in the remaining CR contracts was viewed as sufficient for price discovery in both markets.

This view is consistent with the market depth of available quotes for both contract types: Table A.12 reports the market 'depth' as the number of quote submissions from dealers used in Markit's computation of the consensus quote. Differences in market depth between the two contracts are small for all countries: for the average country, 4.90 CR14 quotes are reported on the average day, versus 5.05 for the older CR contracts—the older contract type receives slightly *more* quotes on average. Excluding Portugal, for which this difference is the largest in favor of the older CR contracts, the remaining average difference is zero. Similarly, the volatility, maxima, and minima of market depth are comparable across both contract types for

[20] see ISDA release dated February 21st, 2014 and June 30th, 2014, respectively, here and here.

all countries. The absolute market depth of around five intermediary submissions is consistent with the large concentration of these OTC markets among a few dealers (Giglio, 2014; Siriwardane, forthcoming).

1.5.2 Asset package delivery

A second change in the CR14 restructuring clause relative to the CR clause that relates particularly to sovereign issuers is the introduction of 'asset package delivery' (APD). This reform in the calculation of the recovery value is a direct response to the Greek debt restructuring of 2012. When Greece restructured its debt in 2012, existing bonds with 1€ in face value were exchanged into a *package* of new securities: (i) 15 cents of face value in short-term notes to be repaid by the European Financial Stability Facility (EFSF), (ii) 46.5 cents of face value in new Greek bonds with 30 years to maturity and a coupon rate of 2%, and (iii) detachable GDP-warrants which pay a capped amount if Greek GDP growth exceeds certain projections. Greek CDS payouts were triggered, but since old bonds were exchanged, the recovery had to be determined in an auction of the new 30 year bonds, which traded at approximately 30% of par value. As Duffie and Thukral (2012) outline, the *true* recovery is derived from the value (relative to the face value of the original bond) of the total asset package that is received in exchange for the original bonds rather than the value of just the single security, which is determined to be the 'deliverable obligation' and auctioned by ISDA. The APD clause addresses this flaw in the original CDS terms and specifies that recovery be based on the market value of the full asset package. Since the APD clause may impact the recovery offset against the CDS payout, the change in this clause potentially introduces another difference between CR and CR14 CDS spreads. As seen in equation (1.1), the recovery value interacts with the default probability in determining the fair insurance premium. If the APD term is responsible for differences between CR and CR14 spreads, this difference should therefore scale with the level of the spread. Table A.1 shows that, in the control group, this is true in the cross-section: countries with higher average CR14 spreads show a larger difference between CR14 and CR spreads. However, this correlation does not show up within-country. The correlation is negative and/or close to zero for all sampled countries, indicating that the ISDA basis is unlikely to stem from the presence of the APD clause in CR14 contracts. Nonetheless, while necessary, the difference-in-difference method is well-suited to eliminate APD-driven pricing effects, since the introduction of APD in the CR14 relative to the CR restructuring clause applies to all sovereign issuers regardless of G-7 membership.

1.6 Conclusion

This chapter presents a directly observable quantitative measure of redenomination risk in French, Italian, and German government bonds. The measure uses CDS spreads on contracts, which, respectively, do and do not cover bondholders' losses from a redenomination into a newly issued French, Italian, or German national currency. I construct a difference-in-difference measure to account for potential liquidity differences and other contractual discrepancies between the two CDS types.

French redenomination risk is economically large before the 2017 presidential elections, when it accounts for 40% of the total French CDS spread. Italian redenomination risk is elevated around and immediately following the French presidential election, and spikes to 80 basis points (close to one third of the total CDS spread) during coalition negotiations in late May 2018. German redenomination risk is close to zero throughout the sample period, consistent with the interpretation that a redenomination into a new German currency (i) does not cause losses for bondholders, and/or (ii) is very unlikely.

French redenomination risk is associated with a statistically and economically significant drop in yields on German and Austrian government bonds, while many other sovereign Eurozone yields rise—particularly those on Portuguese debt. The German Bund response to French risk is similar to that of US Treasuries. In contrast, all Eurozone sovereign yields other than Italian yields are negatively correlated with Italian redenomination risk: higher redenomination risk in Italy is associated with lower sovereign yields elsewhere. I do not find a similar association with Italian redenomination risk for dollar-denominated US Treasuries.

Sovereign yields for most European countries, European corporate credit spreads, US Treasury yields, and the euro exchange rate react differently to French and Italian redenomination risk changes. French redenomination risk appears to have heterogeneous spillover effects on Eurozone assets, while Italian redenomination risk is associated with homogeneously lower yields on most other euro-denominated assets. This discrepancy is consistent with the interpretation that a French exit from the Eurozone is expected to lead to further redenominations in other European countries. In contrast, an Italian exit is expected to remain isolated, and benefits other euro-denominated sovereign and corporate debt, which serve as substitutes to Italian bonds.

I relate the co-movement of Eurozone yields with the presumably contagious French risk to fundamental variables. I find that the heterogeneity lines up with cross-sectional variation in the countries' fiscal positions, trade balances, and labor productivities.

I do not address the question why a French exit is associated with a Eurozone break-up, while an Italian exit is not. I leave it to further research to uncover the political, macroeconomic, and/or financial channels, which may or may not generate a 'contagious' cross-country correlation in withdrawals from the Eurozone.

2. The Quanto Theory of Exchange Rates

LUKAS KREMENS AND IAN MARTIN[1]

It is notoriously hard to forecast movements in exchange rates. A large part of the literature is organized around the principle of uncovered interest parity (UIP), which predicts that expected exchange rate movements offset interest rate differentials and therefore equalise expected returns across currencies. Unfortunately many authors, starting from Hansen and Hodrick (1980) and Fama (1984), have shown that this prediction fails: returns have historically been larger on high interest rate currencies than on low interest rate currencies.[2]

Given its empirical failings, it is worth reflecting on why UIP represents such an enduring benchmark in the FX literature. The UIP forecast has three appealing properties. First, it is determined by asset prices alone rather than by, say, infrequently updated and imperfectly measured macroeconomic data. Second, it has no free parameters: with no coefficients to be estimated in-sample or "calibrated," it is perfectly suited to out-of-sample forecasting. Third, it has a straightforward interpretation as the expected exchange rate movement perceived by a risk-neutral investor. Put differently, UIP holds if and only if the *risk-neutral* expected appreci-

[1]We thank the Systemic Risk Centre and the Paul Woolley Centre at the LSE for their support, and for providing access to data sourced from Markit under license. We are grateful to Christian Wagner, Tarek Hassan, John Campbell, Mike Chernov, Gino Cenedese, Anthony Neuberger, Dagfinn Rime, Urban Jermann, Bryn Thompson-Clarke, Adrien Verdelhan, Bernard Dumas, Pierpaolo Benigno, Alan Taylor, Daniel Ferreira, Ulf Axelson, Scott Robertson, and to participants in seminars at the LSE, Imperial College, Cass Business School, LUISS, BI Business School, Boston University, and Queen Mary University of London, for their comments; and to Lerby Ergun for research assistance. Ian Martin is also grateful for support from the ERC under Starting Grant 639744.

[2]Some studies (e.g. Sarno et al., 2012) find that currencies with high interest rates appreciate on average, exacerbating the failure of UIP; this has become known as the forward premium puzzle. Others, such as Hassan and Mano (forthcoming), find that exchange rates move in the direction predicted by UIP, though not by enough to offset interest rate differentials.

ation of a currency is equal to its *real-world* expected appreciation, the latter being the quantity relevant for forecasting exchange rate movements.

There is, however, no reason to expect that the real-world and risk-neutral expectations should be similar. On the contrary, the modern literature in financial economics has documented that large and time-varying risk premia are pervasive across asset classes, so that risk-neutral and real-world distributions are very different from one another: in other words, the perspective of a risk-neutral investor is not useful from the point of view of forecasting. Thus, while UIP has been a useful organizing principle for the empirical literature on exchange rates, its predictive failure is no surprise.[3]

In this chapter we propose a new predictor variable that also possesses the three appealing properties mentioned above, but which does not require that one takes the perspective of a risk-neutral investor. This alternative benchmark can be interpreted as the expected exchange rate movement that must be perceived by a risk-averse investor with log utility whose wealth is invested in the stock market. (To streamline the discussion, this description is an oversimplification and strengthening of the condition we actually need to hold for our approach to work, which is based on a general identity presented in Result 1.) This approach has been shown by Martin (2017) and Martin and Wagner (forthcoming) to be successful in forecasting returns on the stock market and on individual stocks, respectively.

It turns out that such an investor's expectations about currency returns can be inferred directly from the prices of so-called *quanto contracts*. For our purposes, the important feature of such contracts is that their prices are sensitive to the correlation between a given currency and some other asset price. Consider, for example, a quanto contract whose payoff equals the level of the S&P 500 index at time T, denominated in euros (that is, the exchange rate is fixed—in this example, at 1 euro per dollar—at initiation of the trade). The value of this contract is sensitive to the correlation between the S&P 500 index and the dollar/euro exchange rate. If the euro appreciates against the dollar at times when the index is high, and depreciates when the index is low, then this quanto contract is more valuable than a conventional, dollar-denominated, claim on the index.[4]

[3] Various authors have fleshed out this point in the context of equilibrium models: see for example Verdelhan (2010), Hassan (2013), and Martin (2013a). On the empirical side, authors including Menkhoff et al. (2012), Barroso and Santa-Clara (2015) and Della Corte et al. (2016a) have argued that it is necessary to look beyond interest rate differentials to explain the variation in currency returns.

[4] A different type of quanto contract—specifically, quanto CDS contracts—is used by Mano (2013) and Augustin et al. (2018) to study the relationship between currency depreciation and

We show that the relationship between currency-i quanto forward prices and conventional forward prices on the S&P 500 index reveals the risk-neutral covariance between currency i and the index. Quantos therefore signal which currencies are risky—in that they tend to depreciate in bad times, i.e., when the S&P 500 declines—and which are hedges; it is possible, of course, that a currency is risky at one point in time and a hedge at another. Intuitively, one expects that a currency that is (currently) risky should, as compensation, have higher expected appreciation than predicted by UIP, and that hedge currencies should have lower expected appreciation. Our framework formalizes this intuition. It also allows us to distinguish between variation in risk premia across currencies and variation over time.

It is worth emphasizing various assumptions that we do *not* make. We do not require that markets are complete (though our approach remains valid if they are). We do not assume the existence of a representative agent, nor do we assume that all economic actors are rational: the forecast in which we are interested reflects the beliefs of a rational investor, but this investor may coexist with investors with other, potentially irrational, beliefs. We do not assume lognormality, nor do we make any other distributional assumptions: our approach allows for skewness and jumps in exchange rates. This is an important strength of our framework, given that currencies often experience crashes or jumps (as emphasized by Brunnermeier et al. (2008), Jurek (2014), Della Corte et al. (2016c), Chernov et al. (2018) and Farhi and Gabaix (2016), among others), and are prone to structural breaks more generally. The approach could even be used, in principle, to compute expected returns for currencies that are currently pegged but that have some probability of jumping off the peg. To the extent that skewness and jumps are empirically relevant, this fact will be embedded in the asset prices we use as forecasting variables.

Our approach is therefore well adapted to the view of the world put forward by Burnside et al. (2011), who argue that the attractive properties of carry trade strategies in currency markets may reflect the possibility of peso events in which the stochastic discount factor takes extremely large values. Investor concerns about such events, if present, should be reflected in the forward-looking asset prices that we exploit, and thus our quanto predictor variable should forecast high appreciation for currencies vulnerable to peso events even if no such events turn out to happen in sample.

We derive these and other theoretical results in Section 2.1, and test them in Section 2.2 by running panel currency-forecasting regressions. The estimated coefficient

sovereign default.

on the quanto predictor variable is economically large and statistically significant: in our headline regression (2.20), we find t-statistics of 3.2 and 2.3 respectively with and without currency fixed effects. (Here, as throughout the chapter, we compute standard errors—and more generally the entire covariance matrix of coefficient estimates—using a nonparametric block bootstrap to account for heteroskedasticity, cross-sectional correlation across currencies, and autocorrelation in errors induced by overlapping observations.) The quanto predictor outperforms forecasting variables such as the interest rate differential, average forward discount, and the real exchange rate as a univariate forecaster of currency excess returns. On the other hand, we find that some of these variables—notably the real exchange rate and average forward discount—interact well with our quanto predictor variable, in the sense that they substantially raise R^2 above what the quanto variable achieves on its own. We interpret this fact, through the lens of the identity (2.6) of Result 1, as showing that these variables help to measure deviations from the log investor benchmark. We also show that the quanto predictor variable—that is, forward-looking risk-neutral covariance—predicts future realized covariance and substantially outperforms lagged realized covariance as a forecaster of exchange rates.

An important challenge is that our dataset spans a relatively short time period. If we assess the significance of joint hypothesis tests by using p-values based on the *asymptotic* distributions of test statistics (with bootstrapped covariance matrices, as always), we find, in our pooled regressions, that the estimated coefficients on the quanto predictor variable and interest rate differential are consistent with the predictions of the log investor benchmark, but we can reject the hypothesis that, in addition, the intercept is zero. This rejection can be attributed to US dollar appreciation, during our sample, that was not anticipated by our model. But using asymptotic distributions of test statistics to assess p-values risks giving a false impression of precision, in view of our short sample period. In Section 2.2.6, we bootstrap the small-sample distributions of the relevant test statistics to account for this issue. When we use the associated, more conservative, small-sample p-values, we do not reject even the most optimistic hypothesis in any of the specifications, though the individual significance of the quanto predictor becomes more marginal, with p-values ranging from 5.1% to 9.7%.

In Section 2.3 we show that the quanto variable performs well out of sample. We focus on forecasting differential returns on currencies in order to isolate the cross-sectional forecasting power of the quanto variable in a dollar-neutral way, in the spirit of Lustig et al. (2011), and independent of what Hassan and Mano (forthcoming) refer to as the dollar trade anomaly. (As noted in the preceding paragraph, the dollar strengthened against almost all other currencies over our relatively short sample, so quantos are not successful in forecasting the average performance of the dollar itself. Our findings are therefore complementary to Gourinchas and Rey (2007), who use a measure of external imbalances to forecast the appreciation of the dollar against a trade- or FDI-weighted basket of currencies.)

In a recent survey of the literature, Rossi (2013) emphasizes that the exchange-rate forecasting literature has struggled to overturn the frustrating fact, originally documented by Meese and Rogoff (1983), that it is hard even to outperform a random walk forecast out of sample. Our out-of-sample forecasts exploit the fact that our theory makes an a priori prediction for the coefficient on the quanto predictor variable. When the coefficient is fixed at the level implied by the theory, we end up with a forecast of currency appreciation that has no free parameters, and which is therefore—like the UIP and random walk forecasts—perfectly suited for out-of-sample forecasting. Following Meese and Rogoff (1983) and Goyal and Welch (2008), we compute mean squared errors for the differential currency forecasts made by the quanto theory and by three competitor models: UIP, which predicts currency appreciation through the interest rate differential; PPP, which uses past inflation differentials (as a proxy for expected inflation differentials) to forecast currency appreciation; and the random walk forecast. The quanto theory outperforms all three competitors. We also show that it outperforms on an alternative performance benchmark, the correct classification frontier, that has been proposed by Jordà and Taylor (2012).

2.1 Theory

We start with the fundamental equation of asset pricing,

$$\mathbb{E}_t \left(M_{t+1} \widetilde{R}_{t+1} \right) = 1, \tag{2.1}$$

since this will allow us to introduce some notation. Today is time t; we are interested in assets with payoffs at time $t+1$. We write $\mathbb{E}_t$ for the (real-world) expectation operator, conditional on all information available at time t, and M_{t+1} for a stochastic discount factor (SDF) that prices assets denominated in dollars. (We do not assume complete markets, so there may well be other SDFs that also price assets denominated in dollars. But all such SDFs must agree with M_{t+1} on the prices of the payoffs in which we are interested, since they are all tradable.) In equation

(2.1), $\widetilde{R}_{t+1}$ is the gross return on some arbitrary dollar-denominated asset or trading strategy. If we write $R^{\$}_{f,t}$ for the gross one-period dollar interest rate, then the equation implies that $\mathbb{E}_t\, M_{t+1} = 1/R^{\$}_{f,t}$, as can be seen by setting $\widetilde{R}_{t+1} = R^{\$}_{f,t}$; thus (2.1) can be rearranged as

$$\mathbb{E}_t\, \widetilde{R}_{t+1} - R^{\$}_{f,t} = -R^{\$}_{f,t} \operatorname{cov}_t\left(M_{t+1}, \widetilde{R}_{t+1}\right). \tag{2.2}$$

Consider a simple currency trade: take a dollar, convert it to foreign currency i, invest at the (gross) currency-i riskless rate, $R^{i}_{f,t}$, for one period, and then convert back to dollars. We write $e_{i,t}$ for the price in dollars at time t of a unit of currency i, so that the gross return on the currency trade is $R^{i}_{f,t} e_{i,t+1}/e_{i,t}$; setting $\widetilde{R}_{t+1} = R^{i}_{f,t} e_{i,t+1}/e_{i,t}$ in (2.2) and rearranging,[5] we find that

$$\mathbb{E}_t\, \frac{e_{i,t+1}}{e_{i,t}} = \underbrace{\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{UIP forecast}} \underbrace{- R^{\$}_{f,t} \operatorname{cov}_t\left(M_{t+1}, \frac{e_{i,t+1}}{e_{i,t}}\right)}_{\text{residual}}. \tag{2.3}$$

This (well known) identity can also be expressed using the risk-neutral expectation $\mathbb{E}^*_t$, in terms of which the time t price of any payoff, X_{t+1}, received at time $t+1$ is

$$\text{time } t \text{ price of a claim to } X_{t+1} = \frac{1}{R^{\$}_{f,t}}\, \mathbb{E}^*_t\, X_{t+1} = \mathbb{E}_t\left(M_{t+1} X_{t+1}\right). \tag{2.4}$$

The first equality is the defining property of the risk-neutral probability distribution. The second equality (which can be thought of as a dictionary for translating between risk-neutral and SDF notation) can be used to rewrite (2.3) as

$$\mathbb{E}^*_t\left(\frac{e_{i,t+1}}{e_{i,t}}\right) = \frac{R^{\$}_{f,t}}{R^{i}_{f,t}}. \tag{2.5}$$

From an empirical point of view, the challenging aspect of the identity (2.3) is the presence of the unobservable SDF M_{t+1}. If M_{t+1} were constant conditional on time t information then the covariance term would drop out and we would recover the UIP prediction that $\mathbb{E}_t\, e_{i,t+1}/e_{i,t} = R^{\$}_{f,t}/R^{i}_{f,t}$, according to which high-interest-rate currencies are expected to depreciate. Thus, if the UIP forecast is used to predict exchange rate appreciation, the implicit assumption being made is that the covariance term can indeed be neglected.

[5] Unlike most authors in this literature, we prefer to work with true returns, $\widetilde{R}_{t+1}$, rather than with log returns, $\log \widetilde{R}_{t+1}$, as the latter are only "an approximate measure of the rate of return to speculation," in the words of Hansen and Hodrick (1980).

Unfortunately, as is well known, the UIP forecast performs poorly in practice: the assumption that the covariance term is negligible in (2.3) (or, equivalently, that the risk-neutral expectation in (2.5) is close to the corresponding real-world expectation) is not valid. This is hardly surprising, given the existence of a vast literature in financial economics that emphasizes the importance of risk premia, and hence shows that the SDF M_{t+1} is highly volatile (Hansen and Jagannathan, 1991). The risk adjustment term in (2.3) therefore cannot be neglected: expected currency appreciation depends not only on the interest rate differential, but also on the covariance between currency movements and the SDF. Moreover, it is plausible that this covariance varies both over time and across currencies. We therefore take a different approach that exploits the following observation:

Result 1. *Let R_{t+1} be an arbitrary gross return. We have the identity*

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} = \underbrace{\frac{R^{\$}_{f,t}}{R^{i}_{f,t}}}_{\text{UIP forecast}} + \underbrace{\frac{1}{R^{\$}_{f,t}} \operatorname{cov}^*_t \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right)}_{\text{quanto-implied risk premium}} - \underbrace{\operatorname{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right)}_{\text{residual}}. \tag{2.6}$$

The asterisk on the first covariance term in (2.6) *indicates that it is computed using the risk-neutral probability distribution.*

Proof. Setting $\widetilde{R}_{t+1} = R^{i}_{f,t} e_{i,t+1}/e_{i,t}$ in (2.1) and rearranging, we have

$$\mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) = \frac{1}{R^{i}_{f,t}}. \tag{2.7}$$

We can use (2.4) and (2.7) to expand the risk-neutral covariance term that appears in the identity (2.6) and express it in terms of the SDF:

$$\begin{aligned} \frac{1}{R^{\$}_{f,t}} \operatorname{cov}^*_t \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right) &\overset{(2.4)}{=} \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - R^{\$}_{f,t} \, \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) \\ &\overset{(2.7)}{=} \mathbb{E}_t \left(M_{t+1} \frac{e_{i,t+1}}{e_{i,t}} R_{t+1} \right) - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}}. \end{aligned} \tag{2.8}$$

Note also that

$$\operatorname{cov}_t \left(M_{t+1} R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}} \right) = \mathbb{E}_t \left(M_{t+1} R_{t+1} \frac{e_{i,t+1}}{e_{i,t}} \right) - \mathbb{E}_t \left(\frac{e_{i,t+1}}{e_{i,t}} \right). \tag{2.9}$$

Subtracting (2.9) from (2.8) and rearranging, we have the result. □

As (2.3) and (2.6) are identities, each must hold for all currencies i in any economy that does not exhibit riskless arbitrage opportunities. Nor do they make any assumptions about the exchange rate regime. If currency i is perfectly pegged then the covariance terms in (2.6) are zero, and we recover the familiar fact that countries with pegged currencies must either lose control of their monetary policy (that is, set $R^i_{f,t} = R^\$_{f,t}$) or restrict capital flows to prevent arbitrageurs from trading on the interest rate differential. More generally, the covariance terms should be small if a currency has a low probability of jumping off its peg.

The identity (2.6) generalizes (2.3), however, by allowing R_{t+1} to be an arbitrary return. To make the identity useful for empirical work, we want to choose a return R_{t+1} with two aims in mind. First, the residual term should be small. Second, the middle term should be easy to compute.

These two goals are in tension. If we set $R_{t+1} = R^\$_{f,t}$, for example, then (2.6) reduces to (2.3), which achieves the second of the goals but not the first. Conversely, one might imagine setting R_{t+1} equal to the return on an elaborate portfolio exposed to multiple risk factors and constructed in such a way as to minimise the volatility of $M_{t+1}R_{t+1}$: this would achieve the first but not necessarily the second, as will become clear in the next section.

To achieve both goals simultaneously, we want to pick a return that offsets a substantial fraction of the variation[6] in M_{t+1}; but we must do so in such a way that the risk-neutral covariance term can be measured empirically. For much of this chapter, we will take R_{t+1} to be the return on the S&P 500 index. (We find similar—and internally consistent—results if R_{t+1} is set equal to the return on other stock indexes, such as the Nikkei, Euro Stoxx 50, or SMI: see Sections 2.1.2 and 2.2.1.) It is highly plausible that this return is negatively correlated with M_{t+1}, consistent with the first goal; in fact we provide conditions below under which the residual is exactly zero. We will now show that the second goal is also achieved with this choice of R_{t+1} because we can calculate the quanto-implied risk premium directly from asset prices without any further assumptions—specifically, from *quanto forward prices* (hence the name).

2.1.1 Quantos

An investor who is bullish about the S&P 500 index might choose to go long a *forward contract* at time t, for settlement at time $t+1$. If so, he commits to pay F_t at time $t+1$ in exchange for the level of the index, P_{t+1}. The dollar payoff on

[6]More precisely, all we need is to pick a return that offsets *the component of the variation in M_{t+1} that is correlated with currency movements.* But as this component will in general vary according to the currency in question, it is sensible simply to choose R_{t+1} to offset variation in M_{t+1} itself.

the investor's long forward contract is therefore $P_{t+1} - F_t$ at time $t+1$. Market convention is to choose F_t to make the market value of the contract equal to zero, so that no money needs to change hands initially. This requirement implies that

$$F_t = \mathbb{E}_t^* P_{t+1}. \tag{2.10}$$

A *quanto forward contract* is closely related. The key difference is that the quanto forward commits the investor to pay $Q_{i,t}$ *units of currency* i at time $t+1$, in exchange for P_{t+1} *units of currency* i. (At each time t, there are N different quanto prices indexed by $i = 1, \ldots, N$, one for each of the N currencies in our data set. Other than in Section 2.1.2, the underlying asset is always the S&P 500 index, whatever the currency.) The payoff on a long position in a quanto forward contract is therefore $P_{t+1} - Q_{i,t}$ units of currency i at time $t+1$; this is equivalent to a time $t+1$ dollar payoff of $e_{i,t+1}(P_{t+1} - Q_{i,t})$. As with a conventional forward contract, the market convention is to choose the quanto forward price, $Q_{i,t}$, in such a way that the contract has zero value at initiation. It must therefore satisfy

$$Q_{i,t} = \frac{\mathbb{E}_t^* e_{i,t+1} P_{t+1}}{\mathbb{E}_t^* e_{i,t+1}}. \tag{2.11}$$

(We converted to dollars because $\mathbb{E}_t^*$ is the risk-neutral expectations operator that prices *dollar* payoffs.) Combining equations (2.5) and (2.11), the quanto forward price can be written

$$Q_{i,t} = \frac{R^i_{f,t}}{R^\$_{f,t}} \mathbb{E}_t^* \frac{e_{i,t+1} P_{t+1}}{e_{i,t}},$$

which implies, using (2.5) and (2.10), that the gap between the quanto and conventional forward prices captures the conditional risk-neutral covariance between the exchange rate and stock index,

$$Q_{i,t} - F_t = \frac{R^i_{f,t}}{R^\$_{f,t}} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, P_{t+1} \right). \tag{2.12}$$

We will make the simplifying assumption that dividends earned on the index between time t and time $t+1$ are known at time t and paid at time $t+1$. It then follows from (2.12) that

$$\frac{Q_{i,t} - F_t}{R^i_{f,t} P_t} = \frac{1}{R^\$_{f,t}} \operatorname{cov}_t^* \left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1} \right), \tag{2.13}$$

so the quanto forward and conventional forward prices are equal if and only if currency i is uncorrelated with the stock index under the risk-neutral measure. This allows us to measure the risk-neutral covariance term that appears in (2.6) directly from the gap between quanto and conventional index forward prices (which, as noted, we will refer to as the quanto-implied risk premium).

We still have to deal with the final covariance term in the identity (2.6). The next result exhibits a case in which this covariance term is exactly zero.

Result 2 (The log investor)**.** *If we take the perspective of an investor with log utility whose wealth is fully invested in the stock index then* $M_{t+1} = 1/R_{t+1}$*, so that* $\text{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$ *is identically zero. The expected appreciation of currency* i *is then given by*

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{\frac{R^{\$}_{f,t}}{R^{i}_{f,t}} - 1}_{IRD_{i,t}} + \underbrace{\frac{Q_{i,t} - F_t}{R^{i}_{f,t} P_t}}_{QRP_{i,t}}, \tag{2.14}$$

and the expected excess return[7] *on currency* i *equals the quanto-implied risk premium:*

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \frac{Q_{i,t} - F_t}{R^{i}_{f,t} P_t}.$$

Equation (2.14) splits expected currency appreciation into two terms. The first is the UIP prediction which, as we have seen in equation (2.5), equals *risk-neutral* expected currency appreciation. We will often refer to this term as the *interest rate differential* (IRD); and as above we will generally convert to net rather than gross terms by subtracting 1. (We choose to refer to a high-interest-rate currency as having a *negative* interest rate differential because such a currency is forecast to depreciate by UIP.) The second is a risk adjustment term: by taking the perspective of the log investor, we have converted the general form of the residual that appears in (2.3) into a quantity that can be directly observed using the gap between a quanto forward and a conventional forward.[8] Since it captures the risk premium perceived by the log investor, we refer to this term as the *quanto-implied risk premium* (QRP). Lastly, we refer to the sum of the two terms as *expected currency appreciation* (ECA = IRD + QRP).

[7]Formally, $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^{i}_{f,t}$ is an excess return because it is a tradable payoff whose price is zero, by (2.5).

[8]More generally, we can allow for the case in which the log investor chooses a portfolio $R_{p,t+1} = wR_{t+1} + (1-w)R^{\$}_{f,t}$. (The case in the text corresponds to $w = 1$.) The identity (2.6) then reduces to

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} = \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} + \frac{w}{R^{\$}_{f,t}} \text{cov}^*_t\left(\frac{e_{i,t+1}}{e_{i,t}}, R_{t+1}\right).$$

We thank Scott Robertson for pointing this out to us. See footnote 13 for more discussion.

Results 1 and 2 link expected currency returns to *risk-neutral* covariances, so deviate from the standard CAPM intuition (that risk premia are related to *true* covariances) in that they put more weight on comovement in bad states of the world. This distinction matters, given the observation of Lettau et al. (2014) that the carry trade is more correlated with the market when the market experiences negative returns. Even more important, risk-neutral covariance is directly measurable, as we have shown.[9] In contrast, forward-looking true covariances are *not* directly observed so must be proxied somehow, typically by historical realized covariance. In Section 2.2.3, we show that risk-neutral covariance drives out historical realized covariance as a predictor variable.

Lastly, we emphasize that while Result 2 represents a useful benchmark and is the jumping-off point for our empirical work, in our analysis below we will also allow for the presence of the final covariance term in the identity (2.6). Throughout the chapter, we do so in a simple way by reporting regression results with (and without) currency fixed effects, to account for any currency-dependent but time-independent component of the covariance term. In Section 2.2.5, we consider further proxies that depend both on currency and time.

2.1.2 Alternative benchmarks

Our choice to think from the perspective of an investor who holds the US stock market is a pragmatic one. From a purist point of view, it might seem more natural to adopt the perspective of an investor whose wealth is invested in a globally diversified portfolio;[10] unfortunately global-wealth quantos are not traded, whereas S&P 500 quantos are. Our approach implicitly relies on an assumption that the US stock market is a tolerable proxy for global wealth. We think this assumption makes sense; it is broadly consistent with the 'global financial cycle' view of Miranda-Agrippino and Rey (2019).

Nonetheless, one might wonder whether the results are similar if one uses other countries' stock markets as proxies for global wealth.[11] For, just as the forward price

[9] While it is well known from the work of Ross (1976) and Breeden and Litzenberger (1978) that risk-neutral expectations of functions of a single asset price can typically be inferred from the price of options on that asset, Martin (2018) shows that it is in general considerably harder to infer risk-neutral expectations of functions of multiple asset prices. It is something of a coincidence that precisely the assets whose prices reveal these risk-neutral covariances are traded.

[10] This perspective is suggested by the analysis of Solnik (1974) and Adler and Dumas (1983), for example.

[11] In practice, many investors do choose to hold home-biased portfolios (French and Poterba (1991), Tesar and Werner (1995), and Warnock (2002); and see Lewis (1999) and Coeurdacier and

of the US stock index quantoed into currency i reveals the expected appreciation of currency i versus the dollar, as perceived by a log investor whose portfolio is fully invested in the US stock market, so the forward price of the currency-i stock index quantoed into dollars reveals the expected appreciation of the dollar versus currency i, as perceived by a log investor whose portfolio is fully invested in the currency-i market.

Recall Result 2 for the expected appreciation of currency i versus the dollar,

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - 1 = \underbrace{\mathrm{IRD}_{i,t} + \mathrm{QRP}_{i,t}}_{\mathrm{ECA}_{i,t}}. \tag{2.15}$$

(To reiterate, a positive value indicates that currency i is expected to strengthen against the dollar.) The corresponding expression for the expected appreciation of the dollar versus currency i, from the perspective of a log investor whose wealth is fully invested in the currency-i stock market, is

$$\mathbb{E}_t^i \frac{1/e_{i,t+1}}{1/e_{i,t}} - 1 = \underbrace{\mathrm{IRD}_{1/i,t} + \mathrm{QRP}_{1/i,t}}_{\mathrm{ECA}_{1/i,t}}, \tag{2.16}$$

where we write $\mathrm{IRD}_{1/i,t} = R^i_{f,t}/R^\$_{f,t} - 1$, and where $\mathrm{QRP}_{1/i,t}$ is obtained from conventional forwards and *dollar*-denominated quanto forwards on the currency-i stock market. When the left-hand side of the above equation is positive, the dollar is expected to appreciate against currency i.

In Section 2.2.1 below, we show that the two perspectives captured by (2.15) and (2.16) are broadly consistent with one another (for those currencies for which we observe the appropriate quanto forward prices). If, say, the forward price of the S&P 500 quantoed into euros implies that the euro is expected to appreciate against the dollar by 2% (using equation (2.15)), then the forward price of the Euro Stoxx 50 index quantoed into dollars typically implies that the dollar is expected to *de*preciate against the euro by about 2% (using equation (2.16)). To be more precise, we need to take into account Siegel's "paradox" (Siegel, 1972) that, by Jensen's inequality,

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} \geq \left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}}\right)^{-1}. \tag{2.17}$$

(The corresponding inequality with $\mathbb{E}_t$ replaced by any other expectation operator also holds.) If the US and currency-i investors have the same expectations about

Rey (2013) for surveys).

currency appreciation then (2.15)–(2.17) imply that

$$\log\left(1+\mathrm{ECA}_{i,t}\right) \geq -\log\left(1+\mathrm{ECA}_{1/i,t}\right). \tag{2.18}$$

In practice $\log(1+\mathrm{ECA}) \approx \mathrm{ECA}$, so the above inequality is essentially equivalent to $\mathrm{ECA}_{i,t} \geq -\mathrm{ECA}_{1/i,t}$: thus (continuing the example) if the euro is expected to appreciate by 2% against the dollar, then the dollar should be expected to depreciate against the euro by at most 2%.

The difference between the two sides of (2.18) reflects a convexity correction whose size is determined by the amount of conditional variation in $e_{i,t+1}$:

$$\begin{aligned}
\log\left(1+\mathrm{ECA}_{i,t}\right) - \left(-\log\left(1+\mathrm{ECA}_{1/i,t}\right)\right) &= \log \mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \log\left[\left(\mathbb{E}_t \frac{1/e_{i,t+1}}{1/e_{i,t}}\right)^{-1}\right] \\
&= \boldsymbol{c}_t(1) + \boldsymbol{c}_t(-1) \\
&= 2\sum_{n \text{ even}} \frac{\kappa_{n,t}}{n!},
\end{aligned}$$

where $\boldsymbol{c}_t(\cdot)$ and $\kappa_{n,t}$ denote, respectively, the conditional cumulant-generating function and the nth conditional cumulant of log exchange rate appreciation at time t. In particular, $\kappa_{2,t} = \sigma_t^2$ is the conditional variance and $\kappa_{4,t}/\sigma_t^4$ the excess kurtosis of $\log e_{i,t+1}$. (For more on cumulants, see Backus et al. (2001) and Martin (2013b).)

To get a sense of the size of the convexity correction, note that if the exchange rate is lognormal then all higher cumulants are zero: $\kappa_{n,t} = 0$ for $n > 2$. Thus if exchange rate volatility, σ_t, is on the order of 10%, the two perspectives should disagree by about 1% (so in the example above, expected euro appreciation of 2% would be consistent with expected dollar depreciation of 1%). In Section 2.2.1, we show that the convexity gap observed in our data is consistent with this calculation.

2.2 Empirics

We obtained forward prices and quanto forward prices on the S&P 500, together with domestic and foreign interest rates, from Markit; the maturity in each case is 24 months. The data is monthly and runs from December 2009 to October 2015 for the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish krone (DKK), Euro (EUR), British pound (GBP), Japanese yen (JPY), Korean won (KRW), Norwegian krone (NOK), Polish zloty (PLN), and Swedish krona (SEK). As these quantos are used to forecast exchange rates over a 24-month horizon, our forecasting sample runs from December 2009 to October 2017. Markit reports consensus prices based on quotes received from a wide range of financial intermediaries. These prices are used by major OTC derivatives market makers as a means of independently verifying their book valuations and to fulfil regulatory requirements; they do not necessarily reflect transaction prices. Accounting for missing entries in our panel, we have 656 currency-month observations. (Where we do not observe a price, we treat the observation as missing. Larger periods of consecutive missing observations occur only for DKK, KRW, and PLN and are shown as gaps in Figure B.12.)

Since the financial crisis of 2007-2009, a growing literature (including Du et al. (2018)) has discussed the failure of covered interest parity (CIP)—the no-arbitrage relation between forward exchange rates, spot exchange rates and interest rate differentials—and established that since the financial crisis, CIP frequently does not hold if interest rates are obtained from money markets. For each maturity, we observe currency-specific discount factors directly from our Markit data set. The implied interest rates are consistent with the observed forward prices and the absence of arbitrage. Our measure of the interest rate differentials therefore does not violate the no-arbitrage condition we require for identity (2.6) to hold.

The two building blocks of our empirical analysis are the currencies' quanto-implied risk premia (QRP, which measure the risk-neutral covariances between each currency and the S&P 500 index, as shown in equation (2.13)), and their interest rate differentials vis-à-vis the US dollar (IRD, which would equal expected exchange rate appreciation if UIP held). Our measure of expected currency appreciation (the quanto forecast, or ECA) is equal to the sum of IRD and QRP, as in equation (2.14).

Figure B.1 plots each currency's QRP over time; for clarity, the figure drops two currencies for which we have highly incomplete time series (PLN and DKK). The QRP is negative for JPY and positive for all other currencies (with the partial exception of EUR, for which we observe a sign change in QRP near the end of our time period).

Figure B.12 shows the evolution over time of ECA (solid) and of the UIP forecast (dashed) for each of the currencies in our panel. The gap between the two lines for a given currency is that currency's QRP. Table B.1 reports summary statistics of ECA. The penultimate line of the table averages the summary statistics across currencies; the last line reports summary statistics for the pooled data. Table B.2 reports the same statistics for IRD and QRP.

The volatility of QRP is similar to that of interest rate differentials, both currency-by-currency and in the panel. There is considerably more variability in IRD and

QRP when we pool the data than there is in the time series of a typical currency: this reflects substantial dispersion in IRD and QRP across currencies that is captured in the pooled measure but not in the average time series.

Table B.3 reports volatilities and correlations for the time series of individual currencies' ECA, IRD, and QRP. The table also shows three aggregated measures of volatilities and correlations. The row labelled "Time series" reports time-series volatilities and correlations for a typical currency, calculated by averaging time-series volatilities and correlations across currencies. Conversely, the row labelled "Cross section" reports cross-currency volatilities and correlations of time-averaged ECA, IRD, and QRP. Lastly, the row labelled "Pooled" averages on both dimensions: it reports volatilities and correlations for the pooled data.

All three variables (ECA, IRD, and QRP) are more volatile in the cross section than in the time series. This is particularly true of interest rate differentials, which exhibit far more dispersion across currencies than over time.

The correlation between IRD and QRP is negative when we pool our data ($\rho = -0.696$). Given the sign convention on IRD, this indicates that currencies with high interest rates (relative to the dollar) tend to have high risk premia; thus the predictions of the quanto theory are consistent with the carry trade literature and the findings of Lustig et al. (2011). The average time-series (i.e., within-currency) correlation between IRD and QRP is more modestly negative ($\rho = -0.331$): a typical currency's risk premium tends to be higher, or less negative, at times when its interest rate is high relative to the dollar, but this tendency is fairly weak. The disparity between these two facts is accounted for by the strongly negative cross-sectional correlation between IRD and QRP ($\rho = -0.798$). If we interpret the data through the lens of Result 2, these findings suggest that the returns to the carry trade are more the result of persistent *cross-sectional* differences between currencies than of a *time-series* relationship between interest rates and risk premia. This prediction is consistent with the empirical results documented by Hassan and Mano (forthcoming).

We see a corresponding pattern in the time-series, cross-sectional, and pooled correlations of ECA and QRP. The time-series (within-currency) correlation of the two is substantially positive ($\rho = 0.393$), while the cross-sectional correlation is negative ($\rho = -0.305$). In the time series, therefore, a rise in a given currency's QRP is associated with a rise in its expected appreciation; whereas in the cross-section, currencies with relatively high QRP on average have relatively *low* expected currency appreciation on average (reflecting relatively high interest rates on average). Putting the two together, the pooled correlation is close to zero ($\rho = -0.026$). That is, Result 2 predicts that there should be no clear relationship between currency risk premia and expected currency appreciation; again, this is consistent with the findings of Hassan and Mano (forthcoming).

These properties are illustrated graphically in Figure B.2. We plot confidence ellipses centred on the means of QRP and IRD in panel (a), and of QRP and ECA in panel (b), for each currency. The sizes of the ellipses reflect the volatilities of IRD and QRP (or ECA): under joint normality, each ellipse would contain 50% of its currency's observations in population. (Our interest is in the relative sizes of the ellipses: the choice of 50% is arbitrary.) The orientation of each ellipse illustrates the within-currency time series correlation, while the positions of the different ellipses reveal correlations across currencies. The figures refine the discussion above. QRP and IRD are negatively correlated within currency (with the exceptions of CAD, CHF, and KRW) and in the cross-section. QRP and ECA are positively correlated in the time series for every currency, but exhibit negative correlation across currencies; overall, the pooled correlation between the two is close to zero.

Our empirical analysis focuses on contracts with a maturity of 24 months because these have the best data availability. But in one case—the S&P 500 index quantoed into euros—we observe a range of maturities, so can explore the term structure of QRP. Figure B.13 plots the time series of annualized euro-dollar QRP for horizons of 6, 12, 24, and 60 months. On average, the term structure of QRP is flat over the sample period, but QRP is slightly more volatile at shorter horizons, so that the term structure is downward-sloping when QRP spikes and upward-sloping when QRP is low.

2.2.1 A consistency check

Our data also includes quanto forward prices of certain other stock indexes, notably the Nikkei, Euro Stoxx 50, and SMI. We can use this data to explore the predictions of Section 2.1.2, which provides a consistency check on our empirical strategy.

Figure B.3 implements (2.15) and (2.16) for the EUR/USD, JPY/USD, EUR/JPY, and EUR/CHF currency pairs. In each of the top-left, bottom-left and bottom-right panels, the solid line depicts the expected appreciation of the euro against the US dollar, yen, and Swiss franc, respectively, while the dashed line shows the expected *de*preciation of the three currencies against the euro (that is, we flip the sign on the "inverted" series for readability). In the top-right panel, the solid and dashed lines show the expected appreciation of the yen against the US dollar and expected depreciation of the US dollar against the yen, respectively. In every case, the two measures are strongly correlated over time and the solid line is above the dashed line, as they should be according to (2.18). The gaps between the measures are therefore consistent with the Jensen's inequality correction one would expect to see if our currency forecasts measured expected currency appreciation perfectly. Moreover, given that annual exchange rate volatilities are on the order of 10%, the sizes of the gaps between the measures are quantitatively consistent with the Jensen's inequality correction derived at the end of Section 2.1.2.

The EUR/CHF pair in the bottom-right panel represents a particularly interesting case study. The Swiss national bank instituted a floor on the EUR/CHF exchange rate at CHF1.20/€ in September 2011 and consequently also reduced the conditional volatility of the exchange rate. Following this, the two lines converge and the gap remains narrow, at around 0.2%, until January 2015 when the sudden removal of the floor prompted a spike in the volatility of the currency pair, visible in the figure as the point at which the two lines diverge.

2.2.2 Return forecasting

We run two sets of panel regressions in which we attempt to forecast, respectively, currency excess returns and currency appreciation. The literature on exchange rate forecasting has found it substantially more difficult to forecast pure currency appreciation than currency excess returns, so the second set of regressions should be considered more empirically challenging. In each case, we test the prediction of Result 2 via pooled panel regressions. We also report the results of panel regressions with currency fixed effects; by doing so, we allow for the more general possibility that there is a currency-dependent—but time-independent—component in the second covariance term that appears in the identity (2.6).

To provide a sense of the data before turning to our regression results, Figures B.4 and B.5 represent our baseline univariate regressions graphically in the same manner as in Figure B.2. Figure B.4 plots realized currency excess returns (RXR) against QRP and against IRD.[12] Excess returns are strongly positively correlated with QRP both within currency and in the cross-section, suggesting strong predictability with a positive sign. The correlation of RXR with IRD is negative in the cross-section but close to zero, on average, within currency.

[12]As noted in Section 2.1, we work with true returns as opposed to log returns. Engel (2016) points out that it may not be appropriate to view log returns as approximating true returns, as the gap between the two is a similar order of magnitude as the risk premium itself.

Figure B.5 shows the corresponding results for realized currency appreciation (RCA). Panel (a) suggests that the within-currency correlation with the quanto predictor ECA is predominantly positive (with the exceptions of AUD and CHF), as is the cross-sectional correlation. In contrast, panel (b) suggests that the correlation between realized currency appreciation and interest rate differentials is close to zero both within and across currencies, consistent with the view that interest rate differentials do not help to forecast currency appreciation.

We first run a horse race between the quanto-implied risk premium and interest rate differential as predictors of currency excess returns:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{2.19}$$

Here (and from now on) the length of the period from t to $t+1$ over which we measure our return realizations is 24 months, corresponding to the forecasting horizon dictated by the maturity of the quanto contracts we observe in our data.

We also run two univariate regressions. The first of these,

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1}, \tag{2.20}$$

is suggested by Result 2. The second uses interest rate differentials to forecast currency excess returns, as a benchmark:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{2.21}$$

We also run all three regressions with currency fixed effects α_i in place of the shared intercept α.

Table B.4 reports the results. We report coefficient estimates and R^2 for each regression, with and without currency fixed effects; standard errors are shown in parentheses. These standard errors are computed via a nonparametric bootstrap to account for heteroskedasticity, cross-sectional and serial correlation in our data. (The serial correlation arises due to overlapping observations: we make forecasts of 24-month excess returns at monthly intervals.) For comparison, these nonparametric standard errors exceed those obtained from a parametric residual bootstrap by up to a factor of 2, and Hansen–Hodrick standard errors by a factor of around 1.3. We provide a detailed description of our bootstrap procedure and address potential small-sample concerns in Section 2.2.6.

The estimated coefficient on the quanto-implied risk premium is positive and economically large in every specification in which it occurs. Moreover, the R^2 values are substantially higher in the two regressions (2.19) and (2.20) that feature the quanto-implied risk premium than in the regression (2.21) in which it does not occur. The estimate for β in our headline regression (2.20) is 2.604 (standard error 1.127) in the pooled regression and 4.995 (standard error 1.565) in the regression with fixed effects. The fact that these estimates are above 1 raises the possibility that beyond its direct importance in (2.6), the quanto-implied risk premium may also proxy for the second covariance term.[13] We explore this issue in Section 2.2.5. Another noteworthy qualitative feature of our results is the consistently negative intercept, which reflects an unexpectedly strong dollar over our sample period; we discuss the statistical interpretation of this fact in Section 2.2.6.

Following Fama (1984), we can also test how the theory fares at predicting currency appreciation ($e_{i,t+1}/e_{i,t} - 1$). To do so, we run the regression

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{2.22}$$

We do so not because we are interested in the coefficient estimates, which are mechanically related to those of regression (2.19), but because we are interested in the R^2.

To explore the relative importance of the quanto-implied risk premium and interest rate differentials for forecasting currency appreciation, we run univariate regressions of currency appreciation onto the quanto-implied risk premium,

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1}, \tag{2.23}$$

and onto interest rate differentials,

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}. \tag{2.24}$$

As previously, we also run the three regressions (2.22)–(2.24) with fixed effects.

The regression results are shown in Table B.5, which is structured similarly to

[13]Another possibility is that it is more reasonable to think of a log investor as wishing to hold a levered position in the market (so $w > 1$ in the notation of footnote 8). If so, we should find a coefficient on QRP that is larger than one. We are cautious about suggesting this as an explanation, however, because a log investor would never risk bankruptcy. To match the point estimate for specification (2.20), we would need $w = 2.604$ or $w = 4.995$ (respectively without and with fixed effects). In the latter case, the investor would go bankrupt if the market dropped by 20% over the two year horizon.

Table B.4. There is little evidence that the interest rate differential helps to forecast currency appreciation on its own; this is consistent with the previous set of results and with the large literature that documents the failure of UIP. In the pooled panel, the estimated γ in regression (2.24) is close to 0, and the R^2 is essentially zero. With fixed effects, the estimate of γ is marginally negative, providing weak evidence that currencies tend to appreciate against the dollar when their interest rate relative to the dollar is higher than its time-series mean.

More strikingly, the quanto-implied risk premium makes a very large difference in terms of R^2, which increases by two orders of magnitude when moving from specification (2.24) to (2.22) in both the pooled regressions (0.16% to 16.01%) and the fixed-effects regressions (0.20% to 20.56%). It is also interesting that when QRP is included in the regressions (with or without fixed effects) the coefficient estimate on IRD, γ, increases toward the value of 1 predicted by Result 2.

For completeness, Table B.14 reports the results of running regressions (2.20), (2.21), (2.22), and (2.24) separately for each currency at the 24-month horizon, and at 6- and 12-month horizons for the euro. Consistent with the previous literature (for example Fama (1984) and Hassan and Mano (forthcoming)), the coefficient estimates are extremely noisy. A further appealing feature of Result 2 is that it provides a justification for constraining all the coefficient on the quanto-implied risk premium to be equal across currencies, as we have done above.

2.2.3 Risk-neutral covariance vs. true covariance

We have emphasized the importance of risk-neutral covariances of currencies with stock returns, as captured by quanto-implied risk premia, and below we will show that risk-neutral covariance performs well empirically. But it is natural to wonder whether this empirical success merely reflects the fact that currency returns line up with *true* covariances, as studied by Lustig and Verdelhan (2007), Campbell et al. (2010), Burnside (2011) and Cenedese et al. (2016), among others. More formally, from the perspective of the log investor we can conclude, from (2.3), that

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = R^{\$}_{f,t} \operatorname{cov}_t \left(\frac{e_{i,t+1}}{e_{i,t}}, -\frac{1}{R_{t+1}} \right). \tag{2.25}$$

Note that it is the true, not the risk-neutral, covariance that appears in this equation.

The fundamental challenge for a test of this prediction is that forward-looking true covariance is not directly observed. This is the major advantage of our approach: risk-neutral covariance *is* directly observed via the quanto-implied risk premium.

That said, we attempt to test (2.25) by using lagged realized covariance, RPCL, as a proxy for true forward-looking covariance.

The results are shown in Table B.6 of the Appendix. RPCL is positively related to subsequently realized currency excess returns, as suggested by (2.25), but it is not statistically significant in our sample, and is driven out as a predictor by risk-neutral covariance (QRP), consistent with Result 2.

In principle, this might simply indicate that lagged realized covariance is an imperfect proxy for true forward-looking covariance: perhaps the success of QRP simply reflects its superiority as a forecaster of realized covariance? Table B.6 shows that risk-neutral covariance *is*, individually, a statistically significant forecaster of future realized covariance. But it is driven out when lagged realized covariance and the interest-rate differential are included in the multivariate regression (B.4). Moreover, the optimal covariance forecast generated by this multivariate regression is driven out by QRP in the excess-return-forecasting regression (B.5).

The relationship between risk-neutral covariance and true covariance is interesting in its own right. Figure B.6 illustrates the empirical relationship between the covariance forecast obtained from regression (B.4) (our proxy for forward-looking true covariance) and forward-looking risk-neutral covariance (obtained from quanto contracts). The two are positively correlated in the cross-section and in the time-series, but risk-neutral covariance is generally larger (smaller) than future realized covariance for currencies with positive (negative) risk-neutral covariances. This is consistent with the observation of Lettau et al. (2014) that carry trade returns are more correlated with the market at times of negative market returns. As we will now see, it is problematic for lognormal models.

2.2.4 Lognormal models

Lognormal models impose a tight connection between the covariance risk premium and the market and currency risk premium. Define the equity premium $\mathrm{ERP}_t = \log \mathbb{E}_t \frac{R_{t+1}}{R^{\$}_{f,t}}$ and currency risk premium $\mathrm{CRP}_{i,t} = \log \mathbb{E}_t \frac{\widetilde{R}_{i,t+1}}{R^{\$}_{f,t}}$ where $\widetilde{R}_{i,t+1} = R^i_{f,t} e_{i,t+1}/e_{i,t}$ is the return on the currency trade defined earlier.

Result 3 (The covariance risk premium in lognormal models)**.** *Suppose that the market return, exchange rate, and SDF are conditionally jointly lognormal. Then we have*

$$\log \frac{\mathrm{cov}_t(R_{t+1}, e_{i,t+1}/e_{i,t})}{\mathrm{cov}^*_t(R_{t+1}, e_{i,t+1}/e_{i,t})} = \mathrm{ERP}_t + \mathrm{CRP}_{i,t} \tag{2.26}$$

or equivalently

$$\operatorname{cov}_t(r_{t+1}, \Delta e_{i,t+1}) = \operatorname{cov}_t^*(r_{t+1}, \Delta e_{i,t+1}), \tag{2.27}$$

where $r_{t+1} = \log R_{t+1}$ *and* $\Delta e_{i,t+1} = \log(e_{i,t+1}/e_{i,t})$.

Proof. See Appendix B.2. □

Empirically, it is plausible that the right-hand side of (2.26) is positive for most currencies (the yen being a possible exception). But we find that the left-hand side is typically negative in our data. No lognormal model can match these patterns.

It is nonetheless an interesting exercise to see how the quanto risk premium (and the residual covariance term, which would be zero from the perspective of the log investor) behaves inside an equilibrium model. As QRP has a simple characterization in terms of risk-neutral covariance, this is an easy exercise to carry out in any equilibrium model; we suggest that it makes an interesting diagnostic for future generations of international finance models. In that spirit, we have calculated the currency risk premium, QRP, IRD and the residual covariance term within the model of Colacito and Croce (2011).

The results are shown in Appendix B.4. We deviate from the symmetric baseline calibration of Colacito and Croce in order to generate a non-trivial currency risk premium. The comparative statics of their long-run risk model are such that our calibrations which yield a positive asymmetric currency risk premium generate positive risk-neutral covariance (QRP) and a positive residual. In this model, the residual covariance term therefore adds to the prediction of the quanto forecast, as opposed to offsetting it. This positive relationship between risk-neutral covariance and the residual is consistent with our finding that the slope coefficients on QRP in the predictive regressions in Section 2.2.2 are generally larger than 1.

2.2.5 Beyond the log investor

The identity (2.6) expresses expected currency appreciation as the sum of IRD, QRP, and a covariance term, $-\operatorname{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$. Thus far, we have either assumed that this term is constant across currencies and over time (so is captured by the constant in our pooled regressions) or that it has a currency-dependent but time-independent component (so is captured by fixed effects).

To get a sense of what these assumptions may leave out, we conduct a principal components analysis on unexpected currency excess returns: that is, on the difference between realized currency excess returns and the corresponding ex ante expected returns. We calculate these unexpected excess returns in two ways. *Regression residuals* are defined as the estimated residuals $\varepsilon_{i,t+1}$ in the specification of regression (2.20) that includes currency fixed effects. *Theory residuals* are defined similarly, except that we impose $\alpha = 0$, $\beta = 1$ in (2.20).

These residuals reflect both the ex ante residual from the identity (2.6) and the ex post realizations of unexpected currency returns. The identity implies that the predictable component of the realized residuals—if there is one—reveals the covariance term, $-\operatorname{cov}_t(M_{t+1}R_{t+1}, e_{i,t+1}/e_{i,t})$.

We decompose the theory and regression residuals into their respective principal components (dropping DKK, KRW, and PLN from the panel to minimize the impact of missing observations). Table B.11 shows the principal component loadings. The first principal component, which explains just under two thirds of the variation in residuals, can be interpreted as a level, or 'dollar,' factor since it loads positively on all currencies (with the exception of GBP, in the case of the regression residuals).

Motivated by this fact, we now include an additional predictor variable, $\overline{\mathrm{IRD}}_t$, which is calculated as the cross-sectional average of the interest rate differentials in our balanced panel of eight currencies (i.e., excluding DKK, KRW, and PLN); Lustig et al. (2014) interpret this average interest rate differential (which they refer to as the 'average forward discount') as a dollar factor and show that it helps to forecast currency returns. We also include the logarithm of the real exchange rate, which Dahlquist and Penasse (2017) have shown to be a successful forecaster of currency returns.

Table B.7 reports the results of regressions of currency excess returns onto currency fixed effects and subsets of four forecasting variables: the quanto-implied risk premium (QRP), the interest rate differential (IRD), the real exchange rate (RER), and the average interest rate differential ($\overline{\mathrm{IRD}}$). The table reports the univariate, bivariate, 3-variate, and 4-variate specifications with the highest R^2. (Table B.12 reports the R^2 for all $2^4 - 1 = 15$ subsets of the four explanatory variables, though not—for lack of space—the estimated coefficients.) The quanto-implied risk premium features in all R^2-maximizing regressions. The estimates of β are larger than 1 in every specification, suggesting that, over and above its relevance as a direct measure of risk-neutral covariance, the quanto-implied risk premium helps to capture the physical covariance term in (2.6). As we increase from one to two to three explanatory variables, R^2 increases from 22.03% (using QRP alone) to 35.40% (adding the real exchange rate) to 43.56% (adding the dollar factor $\overline{\mathrm{IRD}}$). The interest rate differential itself, IRD, contributes almost no further explanatory power when it is then added as a fourth variable.

As the real exchange rate performs well, we report further results relating to it in Table B.13 of the Appendix.

2.2.6 Joint hypothesis tests and finite-sample issues

We now consider the joint hypothesis tests that are suggested by Result 2. In our three main specifications (2.19), (2.20), and (2.22), equation (2.14) predicts an intercept $\alpha = 0$, and a slope coefficient on QRP $\beta = 1$. For the excess return forecast in regression (2.19), it predicts that the interest rate differential should have no predictive power, i.e. $\gamma = 0$; whereas it predicts that $\gamma = 1$ in the currency-appreciation regression (2.22).

Here, as elsewhere, we use a nonparametric bootstrap procedure to compute the covariance matrix of coefficient estimates. A detailed exposition of the bootstrap methodology is provided in Politis and White (2004) and Patton et al. (2009). In the bootstrap procedure, we resample the data by drawing with replacement blocks of 24 time-series observations from the panel while ensuring that this time-series resampling is synchronized in the cross-section. The length of the time-series blocks is chosen to equal the forecasting horizon of 24 months. The resulting panel is then resampled with replacement in the cross-sectional dimension by drawing blocks of uniformly distributed width (between 2 and 11, the latter being the width of the full cross-section). Since currencies which are adjacent in the panel are more likely to be included together in any given one of these cross-sectional blocks, we permute the cross-section of our panel randomly before each resampling. We then compute the point estimates of the coefficients from the two-dimensionally resampled panel and repeat this procedure 100,000 times. The standard errors are then computed as the standard deviations of the respective coefficients across the 100,000 bootstrap repetitions.

Table B.8 reports p-values for tests of various hypotheses about our baseline regressions. In addition to conventional p-values calculated using the asymptotic (chi-squared) distribution of the Wald test statistic, the table also reports more conservative small-sample p-values obtained from a bootstrapped test statistic distribution. We compute these small-sample p-values by constructing a small-sample distribution of the Wald test-statistic for each regression: We simulate 5,000 sets of monthly data for the LHS variable under the null hypothesis of no predictability, such that the simulated data matches the monthly autocorrelation and covariance matrix of the realized, observed LHS data. We then aggregate the simulated monthly data into 24-month horizon data, like the LHS data used in our regressions (e.g. excess returns over 24 months). As we aim to measure the small-sample performance of our bootstrap routine, the simulated data sets each have the same number of data points as the observed LHS data. For each specification, we then regress the 5,000 simulated LHS data on the respective observed RHS variable(s). Where we run the regression with currency fixed effects, we use the demeaned RHS variable(s). We obtain the point estimates of the coefficients and their covariance matrix from the bootstrap routine outlined above and use the test statistics from these 5,000 regressions to construct the empirical small-sample distribution of the respective Wald statistic under the respective null hypothesis. This procedure also accounts for the potential small-sample Stambaugh bias in the p-values.

Figure B.14 illustrates by plotting the histograms of the bootstrapped distribution of test statistics for various hypotheses on regression (2.22). Panels a and b show the finite-sample bootstrapped distributions of the test statistic for the hypothesis that Result 2 holds, respectively in the pooled and fixed-effects regressions. The value of the test statistic in the data is indicated with an asterisk in each panel. The finite-sample and asymptotic (shown with a solid line) distributions are strikingly different: the asymptotic distribution suggests that we can reject the hypothesis that Result 2 holds, but this conclusion is overturned by the finite-sample distribution. (In the pooled case, the discrepancy is largely due to the intercept, as becomes clear on comparing the asymptotic p-values for tests of hypotheses H_0^1 and H_0^2 in Table B.8: the asymptotic distribution penalizes the fact that the US dollar was strong over our sample period, whereas the finite-sample distribution does not.)

In contrast, the asymptotic and finite-sample distributions tell more or less the same story in panels c and d, which show the corresponding results for tests (without and with fixed effects) of the hypothesis H_0^3 that $\beta = 0$, i.e., that QRP is not useful in forecasting currency appreciation. While the small-sample distributions of the test statistics exhibit fatter tails than the asymptotic χ^2 distribution, the discrepancy between the two is small by comparison with panels a and b, and even using the finite-sample distribution we can reject the hypothesis with some confidence (with p-values of 0.082 and 0.051 in the pooled and fixed-effects cases, respectively).

We reach similar conclusions for regressions (2.19) and (2.20): we do not reject the predictions of Result 2 in the joint Wald tests for any of the three baseline regressions using the small-sample distribution of the test statistic; and QRP remains individually significant as a predictor at the 10% level in all three specifications, with and without currency fixed effects, even if we take the most conservative approach to computing p-values that relies on the empirical small-sample test statistic distribution.

2.3 Out-of-sample prediction

We now test the quanto theory out of sample. Since the dollar strengthened strongly over the relatively short time period spanned by our data (as reflected in the negative intercept in our pooled panel regression (2.22)), we focus on forecasting differential currency appreciation: that is, we seek to predict, for example, the *relative* performance of dollar-yen versus dollar-euro.

In the previous section, we estimated the loadings on the quanto-implied risk premium, QRP, and interest rate differential, IRD, via panel regressions. These deliver the best in-sample coefficient estimates in a least-squares sense. But for an out-of-sample test we must pick the loadings a priori. Here we can exploit the distinctive feature of Result 2 that it makes specific quantitative predictions for the loadings: each should equal 1, as in the formula (2.14). We therefore compute out-of-sample forecasts by fixing the coefficients that appear in (2.22) at their theoretical values: $\alpha = 0$, $\beta = 1$, $\gamma = 1$.

We compare these predictions to those of three competitor models: UIP (which predicts that currency appreciation should offset the interest rate differential, on average), a random walk without drift (which makes the constant forecast of zero currency appreciation, and which is described in the survey of Rossi (2013) as "the toughest benchmark to beat"), and relative purchasing power parity (which predicts that currency appreciation should offset the inflation differential, on average). These models are natural competitors because, like our approach, they make a priori predictions without requiring estimation of parameters, and so avoid in-sample/out-of-sample issues.

To compare the forecast accuracy of the model to those of the benchmarks, we define a dollar-neutral R^2-measure similar to that of Goyal and Welch (2008):

$$R^2_{OS} = 1 - \frac{\sum_i \sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_i \sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2},$$

where $\varepsilon^Q_{i,t+1}$ and $\varepsilon^B_{i,t+1}$ denote forecast errors (for currency i against the dollar) of the quanto theory and the benchmark, respectively, so our measure compares the accuracy of differential forecasts of currencies i and j against the dollar. We hope to find that the quanto theory has lower mean squared error than each of the competitor models, that is, we hope to find positive R^2_{OS} versus each of the benchmarks.

The results of this exercise are reported in Table B.9. The quanto theory outperforms each of the three competitors: when the competitor model is UIP, we find that $R^2_{OS} = 10.91\%$; and when it is relative PPP, we find $R^2_{OS} = 26.05\%$. In our sample, the toughest benchmark is the random walk forecast, consistent with the findings of Rossi (2013). Nonetheless, the quanto theory easily outperforms it, with $R^2_{OS} = 9.57\%$.

To get a sense for whether our positive results are driven by a small subset of the currencies, Table B.9 also reports the results of splitting the R^2 measure currency-by-currency: for each currency i, we define

$$R^2_{OS,i} = 1 - \frac{\sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2}.$$

This quantity is positive for all i and all competitor benchmarks B, indicating that the quanto theory outperforms all three benchmarks for all 11 currencies. We run Diebold–Mariano tests (Diebold and Mariano, 1995) of the null hypothesis that the quanto theory and competitor models perform equally well for all currencies, using a small-sample adjustment proposed by Harvey et al. (1997), and find that the outperformance is strongly significant.

Jordà and Taylor (2012) have argued that assessments of forecast performance based solely on mean squared errors may not fully reflect the economic benefits of a forecasting model. In Appendix B.3, we use the approach they suggest, which essentially asks whether a predictor variable is more or less successful at predicting whether a currency will appreciate or depreciate than competitor predictors. (This is an oversimplification; full details are in Appendix B.3.) Our approach also outperforms on their metric, both in forecasting currency excess returns and in forecasting currency appreciation.

2.4 Conclusion

UIP forecasts that high interest rate currencies should depreciate on average: it reflects the expected currency appreciation that a genuinely risk-neutral investor would perceive in equilibrium. Unsurprisingly—given that the financial economics literature has repeatedly documented the importance of risk premia—the UIP forecast performs extremely poorly in practice.

We have proposed an alternative forecast, the quanto-implied risk premium, that can be interpreted as the expected excess return on a currency perceived by an investor with log utility whose wealth is fully invested in the stock market. Like the UIP forecast, the quanto forecast has no free parameters and can be computed directly from asset prices. Unlike the UIP forecast, the quanto forecast performs well empirically both in and out of sample. Its main deficiency is its failure to predict the strength of the dollar itself on average against other currencies over our sample period: time will tell if this is a small-sample issue or something more fundamental.

We find that currencies tend to have high quanto-implied risk premia if they have high interest rates on average, relative to other currencies (a cross-sectional statement), or if they currently have unusually high interest rates (a time-series statement); and that there is more cross-sectional than time-series variation in quanto-implied risk premia. These facts explain both the existence of the carry trade and the empirical importance of persistent cross-currency asymmetries, as documented by Hassan and Mano (forthcoming).

The interpretation of the quanto-implied risk premium as revealing the log investor's expectation of currency excess returns is a special case of the identity (2.6), which decomposes expected currency appreciation into the interest rate differential (the UIP term), risk-neutral covariance (the quanto-implied risk premium), and a real-world covariance term that, we argue, is likely to be small—and in particular, smaller than the corresponding covariance term in the well-known identity (2.3). In the log investor case, this real-world covariance term is exactly zero, a fact we use to provide intuition and to motivate our out-of-sample analysis. But we also allow for deviations from the log investor benchmark—that is, for a nontrivial real-world covariance term—by running regressions including currency fixed effects, realized covariance, interest rate differentials, the average forward discount of Lustig et al. (2014), and the real exchange rate, as in Dahlquist and Penasse (2017), in addition to the quanto-implied risk premium itself. The quanto-implied risk premium is the best performing univariate predictor, and features in every R^2-maximizing multivariate specification.

Although we have argued that quanto-implied risk premia should (in theory) and do (in practice) predict currency excess returns, we have said nothing about why a particular currency should have a high or low quanto-implied risk premium at a given time. Analogously, the CAPM predicts that assets' betas should forecast their returns but has nothing to say about why a given asset has a high or low beta. Connecting quanto-implied risk premia to macroeconomic fundamentals is an interesting topic for future research.

3. Bets and Betas: Market Risk in Foreign Exchange

Lukas Kremens[1]

Market participants often identify market environments with large price movements as "risk-off" or "risk-on". In a risk-off market, global equity markets, high yield bonds, and emerging market currencies lose value, while so-called safe-haven assets like US Treasuries, gold, or typically the Japanese yen gain, along with implied volatilities across asset classes. A recent literature has linked the co-movement across asset classes to the role of financial intermediaries as marginal investors in different markets.[2] In this chapter, I look at currency futures markets to examine a slightly different transmission channel of cross-asset co-movement in risk-off market environments.

Hedge funds and other typically highly leveraged market participants place directional bets in asset markets. In currency futures markets, the counterparty to these positions is typically a large intermediary, which hedges the currency exposure from these futures positions in other markets. I find that, since the financial crisis, hedge funds reduce the scale of their directional currency positions when the S&P 500 falls and the VIX rises, that is, when markets are hit by a risk-off shock. On the other side of the market, intermediaries scale down their hedges as they unwind the futures positions vis-à-vis the hedge funds.

At times when a given bet is particularly popular amongst hedge funds, the unwinding of this bet in a risk-off scenario will unleash substantial capital flows in the opposite direction of the original bet. In the short run, the liquidity required to accommodate the large scale of capital flows out of these assets may transcend

[1] I thank Daniel Ferreira, Ian Martin, Andrea Vedolin, Philippe Mueller, Christopher Polk, Dong Lou, Thummim Cho, Jonathan Berk and seminar participants at LSE and HEC for helpful comments.

[2] See, for instance, He et al. (2017) and Miranda-Agrippino and Rey (2019).

even the depth of the most liquid foreign exchange spot markets. An asset that is held long as part of the bet experiences selling pressure and depreciates as a result, and vice versa for assets previously held short. As a result, assets held long become positively correlated with equity markets and other risky assets.

An anecdotal example of such an asset is the US dollar vis-à-vis the euro in 2014-2015, when many macro hedge funds were involved in a particular form of the carry trade, commonly referred to as the "divergence trade". This bet on diverging monetary policies between the Federal Reserve and the ECB involved a long position in the dollar against a short position in the euro. Over the course of 2015, the euro-value of the US dollar—historically a typical safe haven asset which would rise during times of market stress—became positively correlated with equity markets. This is to say that the euro started to behave like a "safe-haven" currency relative to the US dollar, which was commonly attributed to the flows of speculative capital out of the dollar and into the euro during particularly bad times for equity markets:

"Is the euro the new safe haven?"

(CNBC, August 2015)[3]

"The euro is looking like the yen – where money tends to come home when the world is a scary place"

(Société Générale, September 2015)[4]

"The euro isn't a haven, but is acting like one because of its role in the carry trade. The distinction is important because it means the link will diminish as these positions, or shorts, are unwound."

(Pioneer Investments, September 2015)[5]

Currency futures are a particularly interesting setting to study the dynamics of speculative positions, because these markets are dominated by hedge funds on one side of the market and broker-dealers on the other, intermediating the hedge funds' demand and hedging their position in other markets, including the spot market.[6] The net positions of the two groups are strongly negatively correlated and account for the vast majority of open interest in most USD FX futures. I document that in the period following the financial crisis—i.e., since 2010—times of negative S&P

[3] cnbc.com/2015/08/24/is-the-euro-the-new-safe-haven.html

[4] http://www.independent.ie/business/world/euro-is-gaining-safehaven-status-among-traders-at-worst-time-for-draghi-31559999.html

[5] ibid.

[6] See Figure C.1. I provide a more detailed description of the data in Section 3.1

returns and increases in the VIX are associated with unwinding positions held by hedge funds and intermediaries in the FX futures market. To illustrate the logic of price impact in the spot market from changes in futures (or forward) positions, Figure 3.1 depicts the possible chain of flows across the two different markets. In

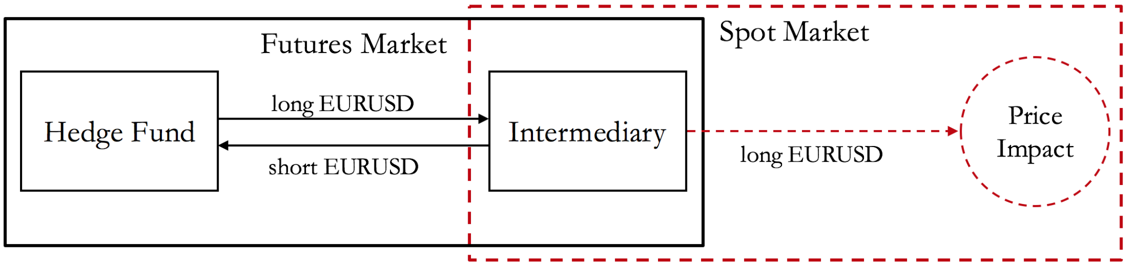

Figure 3.1: Trading flows across FX markets

the example, an investor enters into an unhedged position in the futures market (long €, short \$). Since futures are in zero net-supply, the investor requires a counterparty, and if no other investor wants to take the opposite unhedged position, an intermediary will step in. The intermediary can then hedge its futures position (long \$, short €) with the opposite position in the spot market, thereby passing the futures market flow from the hedge fund position directly on to the spot market, where this flow may have an impact on the spot exchange rate.

Across a sample of 9 currencies in the post-crisis period, currencies exhibit higher equity betas when they are subject to long positioning from hedge funds. Long positioning by hedge funds predicts stronger covariation with the S&P (more positive correlations) and the VIX (more negative) at weekly horizons. The interpretation that this relationship is driven by the price impact from unwinding hedge fund positions is consistent with the finding that the positions of hedge funds contract in size over weeks when S&P returns are negative and the VIX rises. A similar effect does not exist for the positions of institutional investors. The conjectured link between these two sets of results is further consistent with their synchronicity: both occur exclusively in the post-crisis period since 2010, but not in the years 2006 through 2009. The predictive power of hedge fund positioning for subsequently realized betas is independent from interest rates, which are often used to proxy for unobservable risk premia and could therefore potentially drive both hedge fund positions and realized betas. I test the economic significance of the effect by designing a trading strategy that takes positions in the opposite direction of unwinding hedge fund positions following a negative S&P return, and find that this trading strategy is highly profitable with an annualized Sharpe ratio of up to 2. A strategy that provides liquidity to accommodate large flows out of popular hedge fund strategies in times of low equity market returns is compensated with high returns as the temporary price dislocations from those flows tend to correct within the next week of trading. I outline this trading strategy in more detail in Subsection 3.2.3.

Related literature.—The price impact of intermediated cross-currency flows is consistent with the model proposed by Gabaix and Maggiori (2015): capital-constrained intermediaries are counterparty to a net cross-currency flow from investors and therefore bear exchange rate risk. In equilibrium, they are compensated for providing scarce risk-bearing capacity through an expected appreciation in the currency in which they are long, and in which the currency investors are short. To achieve the expected appreciation, the currency must first depreciate when the intermediary enters into the position, that is, net cross-currency flows that require intermediation, as they are "imbalanced", have price impact. Their model speaks to two strands of literature which address "global imbalances" in exchange rate determination: the early literature on portfolio balance models (see Branson and Henderson (1985)), and the more recent work on the valuation channel to external adjustment (see e.g. Gourinchas and Rey (2007)). Della Corte et al. (2016b) confirm empirically the theoretical prediction of Gabaix and Maggiori (2015) that, in equilibrium, currencies of net debtor countries (i.e., currencies that are held by foreigners or foreign intermediaries) perform poorly in bad times when global risk aversion is high, and therefore carry a positive risk premium. This chapter addresses this literature by holding a microscope to the relationship between "imbalances" and exposures to global risk by studying a setting where the imbalances—i.e., cross-currency flows originating from the directional bets of speculators—are highly time-varying and can be measured at higher frequencies (weekly) than macro-variables.

The vast literature on endogenous risks from the unwinding of positions by capital-constrained traders goes back to the limits-of-arbitrage models of Shleifer and Vishny (1997) and Gromb and Vayanos (2002). The two papers, which are closest in spirit to mine are Brunnermeier et al. (2008) and Cho (2018). The latter explores a limits-of-arbitrage argument of endogenous risk in the context of cross-sectional equity strategies and finds that the strategies which are likely to be subject to heavy hedge fund trading are exposed to shocks to the leverage of the broker-dealers financing the hedge funds. I use a setting in currency markets with observable futures positions to explore the link of hedge fund positioning and the economically broader notion of equity market risk in contrast to broker-dealer funding risk, which is not observable at higher than quarterly frequencies.

Brunnermeier et al. (2008) link the profitability of the carry trade to crash risk in high-interest currencies and find—using CFTC futures data similar to mine—that futures positions predict realized skewness and that positioning in high-interest currencies is reduced following increases in the VIX. I separate the impact of unwinding positions from the interest-rate-based carry trade risk premium—as well as any other information contained in interest rates—and explore how the unwinding of positions exposes currencies to spillovers from negative shocks to other markets. In contrast to potentially idiosyncratic skewness, I explore the implications of unwinding on the joint distribution of currency returns and equity markets.

As an example for the importance of equity market risk in other asset classes, particularly in relation to hedge fund strategies, Duarte et al. (2007) show that the returns to swap spread arbitrage predominantly constitute compensation for equity market risk. The relevance of "market" risk for currency risk premia is the subject of a long literature covering consumption risk (Lustig and Verdelhan, 2007; Verdelhan, 2010; Burnside, 2011) as well as equity market risk (Campbell et al., 2010; Lettau et al., 2014; Kremens and Martin, 2019) and equity market volatility (Lustig et al., 2011).

I find that the transmission of contracting positions following "risk-off" shocks into currency betas emerges particularly in more recent data since 2010. A rapidly growing literature on the deviations from covered interest parity, which have emerged in currency markets since the financial crisis, discusses the effects of post-crisis financial regulation on the risk-bearing capacity of financial intermediaries.[7]

3.1 Data

I obtain data on futures positions in foreign exchange markets from the U.S. Commodity Futures Trading Commission (CFTC) which reports weekly commitments of traders in financial futures traded on the Chicago Mercantile Exchange. The data span observations from June 2006 to June 2017 for USD futures and exchange rates versus 9 currencies: Australian dollar (AUD), Brazilian real (BRL), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), pound Sterling (GBP), Japanese yen (JPY), Mexican peso (MXN), and New Zealand dollar (NZD). Each exchange rate is expressed in terms of USD per unit of foreign currency, such that a positive net return reflects an appreciation of the respective currency against the US dollar. Traders are categorized by the CFTC into four categories:

[7]See Du et al. (2018) as an example, and Levich (2017) for a more representative overview of this active literature documenting and discussing CIP deviations.

1. "Dealer/Intermediary" (I will refer to these as *intermediaries*),
2. "Asset Manager/Institutional" (*institutional investors*),
3. "Leveraged Funds" (*hedge funds*),
4. "Other reportables" (*others*).

Group 1 (intermediaries) includes large banks (the so-called "sell-side"), which "typically [...] are dealers and intermediaries that earn commissions on selling financial products, capturing bid/offer spreads and otherwise accommodating clients."[8] The remaining three groups form the "buy-side". The second group (institutional investors) comprises "institutional investors, including pension funds, endowments, insurance companies, mutual funds and those portfolio/investment managers whose clients are predominantly institutional." Group 3 includes hedge funds, and their strategies "involve taking outright positions". These traders "may be engaged in [...] conducting proprietary futures trading and trading on behalf of speculative clients." The final group (others), naturally, contains any remaining types of traders, which mostly use markets "to hedge business risk, whether that risk is related to foreign exchange, equities or interest rates, including [...] corporate treasuries, central banks, smaller banks, mortgage [and] credit unions".

I denote by $nlf_{i,t}$ the net exposure—long positions minus short positions—of *Leveraged Funds* (who I will refer to as hedge funds) to currency i versus the US dollar at time t. The net exposures of institutional investors and intermediaries are analogously denoted by $nam_{i,t}$ and $ndi_{i,t}$, respectively. I will also use a scaled version of this variable, denoted by $\widetilde{nlf}_{i,t} = nlf_{i,t}/oi_{i,t}$ (and again analogously for nam and ndi), where $oi_{i,t}$ denotes the open interest in currency i reported by the CFTC. Table C.1 reports all cross-correlations between the four groups. The positions of hedge funds and intermediaries are strongly negatively correlated ($\rho = -0.88$ and $\rho = -0.87$, respectively, for absolute and relative positions), suggesting that the intermediaries act as counterparties for the directional bets of hedge funds. In contrast, institutional asset managers appear to account for a substantially smaller part of intermediaries' positions with a correlation of -0.43.

Table C.2 reports the average net position of all four groups by currency, along with the respective standard deviations and autocorrelations. Group 4 ("Others") is by far the one with the smallest and least volatile positions for most currencies.

[8]The full CFTC explanatory notes are available at cftc.gov/idc/groups/public/@commitmentsoftraders/documents/file/tfmexplanatorynotes.pdf.

Hedge funds and intermediaries account for the majority of trading in most currencies, but institutional investors hold average positions of notable size in some currencies, such as GBP, JPY, and MXN. While the positioning of all four groups is reasonably strongly autocorrelated from week to week, they all exhibit a considerable amount of within-currency variation. The bottom three rows of each panel in Table C.2 show that this time-series variation far outweighs the cross-sectional variation in the composition of the total panel variation. Figure C.1 shows the scale and time-variation of each group (omitting group 4 to make the graph more readable). The graphs show the considerable time-variation in the direction of speculative trade: hedge funds change from being net long to net short several times over the sample period for all currencies.

3.2 Currency betas

Does the positioning of hedge funds in a given currency help describe time-series variations in currency betas? I run time-series regressions of daily currency returns on S&P returns and an interaction of S&P returns and the $\widetilde{nlf}$ variable, that is, the positioning of hedge funds scaled by open interest. Denote by $r_t^{S\&P} = S\&P_t/S\&P_{t-1} - 1$ the return on the index from day $t-1$ to day t. As $\widetilde{nlf}$ is only observed weekly, I use the last available observation prior to day t to interact with the daily return. Introducing some further notation, I will denote by $r_{i,t} = e_{i,t}/e_{i,t-1} - 1$, the net currency return i versus the US dollar from time $t-1$ to t. I then run the regressions

$$r_{i,t} = \alpha_i + \beta_i r_t^{S\&P} + \beta_i^* r_t^{S\&P} \cdot \widetilde{nlf}_{i,t} + \varepsilon_{i,t}, \tag{3.1}$$

$$r_{i,t} = \alpha_i + \beta_i r_t^{S\&P} + \beta^* r_t^{S\&P} \cdot \widetilde{nlf}_{i,t} + \varepsilon_{i,t}. \tag{3.2}$$

While the conventional beta is estimated as a time-invariant characteristic of each currency, the coefficient on the interaction term, β_i^*, reflects time variation in the beta that is related to the positioning of hedge funds in the currency. Since $\widetilde{nlf}$ captures the sign of the positioning, the coefficient on the interaction term, β_i^* picks up the component of currency betas that is related to hedge fund positioning. If such a component is *caused* by hedge fund positioning, the effect on equity market exposure goes in the same direction for all currencies, and regression (3.2), estimates a *joint* coefficient β^* for the entire cross section of currencies.

Results are reported in Table C.3: Over the full sample, β_i^* is positive for 6 out of 9 currencies and significantly so for 4 (CAD, CHF, JPY, NZD). While the pooled coefficient is positive, the null of $\beta_0^* = 0$ is not rejected at conventional levels with a p-value of 0.12. However, the effect predominantly occurs in the years following the financial crisis (2010-2017). In this post-crisis sample, $\hat{\beta}_i^*$ is positive for all currencies except GBP, and statistically significant for AUD, CAD, CHF, JPY, and NZD. (Incidentally, these are the currencies most commonly associated with speculative trading strategies.) The joint estimate $\hat{\beta}^*$ is positive at 0.186 and statistically significant at the 5% level. In each sample period, the standard market betas take signs consistent with conventional wisdom and previous literature on currency risk: all currencies are "risky" relative to the US dollar in terms of their positive covariance with equity markets, with the exception of the Japanese yen—commonly seen as a "safe haven"—with a significantly negative beta, and the Swiss franc (zero market beta), which is on-par with the US dollar in terms of its equity market risk exposure.

These regressions split the conventionally estimated univariate beta into a baseline exposure and an exposure that can be (statistically) explained by *hedge fund positioning*. For instance, $\beta_€$ may be positive, but nonetheless, the euro may correlate negatively with equity markets in times when $\widetilde{nlf}_€$ is particularly negative—i.e., leveraged market participants are substantially short the euro, as in 2015—as long as $\beta_€^*$ is sufficiently large. It is worth noting that the association of currency-equity co-movement with hedge fund positioning emerges particularly in the post-crisis sample from 2010 onwards.

3.2.1 Futures positions and equity market shocks

The direct link from hedge fund positions to betas requires that positions are unwound in response to negative shocks. I divide realizations of market risk into "risk-off" (bad) and "risk-on" (good) shocks: Other than the S&P 500 as a headline equity market gauge, the VIX is a natural proxy for "risk-on" and "risk-off" movements in financial markets, so let $\Delta VIX_t = VIX_t - VIX_{t-1}$. The two market risk measures, $r_t^{S\&P}$ and ΔVIX_t, are strongly negatively correlated with $\rho(r^{S\&P}, \Delta VIX) = -0.80$ over the full sample period.

To capture the asymmetric effects of positive and negative shocks, I define two truncated weekly S&P-return variables, $r_t^{S\&P^+} = \max(0, r_t^{S\&P})$ and $r_t^{S\&P^-} = \min(0, r_t^{S\&P})$. Similarly, let $\Delta VIX_t^+ = \max(0, \Delta VIX_t)$ and $\Delta VIX_t^- = \min(0, \Delta VIX_t)$ denote the weekly changes in the VIX, truncated at 0. Note that $r^{S\&P^-}$ and ΔVIX^+ represent the proxies for "risk-*off*" shocks. To measure the unwinding of futures positions—both long and short—, I consider the weekly change in the absolute net exposure of the different trader groups, i.e., $\Delta|ndi|$, $\Delta|nlf|$, and for comparison $\Delta|nam|$. I then run the following regressions:

$$\Delta|ndi|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.3}$$

$$\Delta|ndi|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.4}$$

$$\Delta|nlf|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.5}$$

$$\Delta|nlf|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.6}$$

$$\Delta|nam|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.7}$$

$$\Delta|nam|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.8}$$

I report the results in Table C.4 separately for the 2006-2009 period and for the post-crisis period (2010-2017). In the 2006-2009 period, hedge fund and intermediary positions are not significantly exposed to the "risk-off" shocks, $r^{S\&P^-}$ and ΔVIX^+. The coefficient for negative S&P returns is negative (but lacks significance), indicating that positions *expand* over weeks with negative S&P returns.

For the post-crisis period, however, the positions of intermediaries and hedge funds *contract* with risk-off shocks, that is, with negative S&P returns and spikes in the VIX, and do so with statistical significance at conventional levels. The magnitudes of the effects are several times larger than during the earlier sample period. The discrepancies between the two periods line up with the above results for currency betas: unwinding is associated with risk-off shocks in the post-crisis period, when—as shown in Table C.3—hedge fund positions help explain the time-variation in the equity market risk exposures of currencies.[9] This sensitivity to risk-off shocks is not present for institutional investors. The results in Table C.4 support the interpretation that the relationship in Table C.3 between currency betas and hedge fund positions in the post-crisis period is driven by the unwinding of hedge fund positions in response to risk-off shocks.

3.2.2 Realized betas are predictable

If the relationship between positioning and betas is the result of a direct mechanism, a simple test is to use the hedge fund positions as a predictor of subsequently realized currency exposures to the S&P and to the VIX.

I define the following variables measuring the co-movement of currency i with

[9] Futures positions are not available for BRL in the 2006-2009 sample. The results for the 2010-2017 period do not change materially once BRL is dropped from the sample.

equity markets over the week following the observed positioning at date t: $\rho^{SPX}_{i,t\to t+1}$, the correlation of daily exchange rate movements with daily S&P 500 returns; $\beta^{MKT}_{i,t\to t+1} = \frac{\rho^{SPX}_{i,t\to t+1}\sigma_{i,t\to t+1}}{\sigma^{SPX}_{t\to t+1}}$, the beta of daily exchange rate movements with respect to the S&P 500; and $\rho^{VIX}_{i,t\to t+1}$, the correlation of daily exchange rate movements with daily VIX changes between t and $t+1$. Table C.5 reports the averages, standard deviations, and autocorrelations of these variables by currency. At the weekly horizon, the average autocorrelations are low at 0.24, 0.11, and 0.19, respectively.[10] To test the predictive power of futures positions for these beta measures of a currency's equity market risk exposure, I run the following forecasting regressions:

$$y_{i,t\to t+1} = \alpha_i + \eta\,\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, fd^w_{i,t} + \varepsilon_{i,t+1} \tag{3.9}$$

$$y_{i,t\to t+1} = \alpha + \eta\,\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, fd^w_{i,t} + \lambda\, y_{i,t-1} + \varepsilon_{i,t+1} \tag{3.10}$$

The results for regressions (3.9) and (3.10) for $y_{i,t\to t+1} = \{\rho^{SPX}_{i,t\to t+1}, \beta^{MKT}_{i,t\to t+1}, \rho^{VIX}_{i,t\to t+1}\}$ are reported in Table C.6 for the two subsamples 2006-2009 and 2010-2017 (right panels). The relative positioning of hedge funds is a strongly significant and positive predictor of equity market risk as captured by all three measures: currencies that are heavily bought by hedge funds in the futures market, have higher correlations and betas with the S&P 500 over the subsequent week, and more negative correlations with the VIX. Neither of these relationships is present in the earlier sample period, when—recalling Table C.4—hedge fund and intermediary positions are largely invariant to equity market shocks. The result holds within-currency (i.e., with currency fixed effects, Panel A) and across currencies (Panel B).

In the pre-crisis period, the forward discount predicts risk exposures in the cross section, in the sense that high-interest currencies are more positively (negatively) correlated with the S&P (VIX), but this relationship vanishes after 2009. Cross-sectionally, hedge fund positions are also significant predictors of β^{MKT} in the early sample, but the economic magnitude of the coefficient (0.069) is much smaller than in the later sample, and does not hold within-currency. Table C.7 reports a robustness check to the above results, running regressions (3.9) and (3.10) in weekly changes. The different η-coefficients remain statistically significant with the exception of predictions for $\rho^{VIX}_{i,t\to t+1}$, where the p-values rise to 0.127. The point estimates rise slightly in comparison to the estimates in levels, but the order of economic mag-

[10] Computing the beta variables over a time horizon as short as one week with daily data inevitably renders these measures noisy. I choose the weekly horizon in order to avoid overlapping observations and make better use of the weekly futures data, rather than, for instance, forecasting monthly correlations.

nitude is unchanged. On average, a unit change in $\widetilde{nlf}$—which amounts to around three standard deviations—in a given currency is associated with an increase in the currency's risk exposure—as measured by its correlation/beta with the S&P—of around 0.3. To put this number into perspective, the unconditional average risk exposures (shown in Table C.5) are between roughly -0.3 (JPY) and 0.5 (MXN). These economic magnitudes are similar for the VIX-correlations (with the opposite sign).

3.2.3 A profitable contrarian trading strategy

Following the above regressions and *statistical* results, I now construct a trading strategy, which performs a test of the *economic* significance of the link between equity market returns and the unwinding of futures positions by hedge funds. The strategy is implementable in real time, rebalanced weekly, and designed to exploit the temporary price dislocations that result from such unwinding of hedge fund positions. To this end, it takes a position following weeks of bad equity returns, in the opposite direction of changes in hedge fund positioning during that past week.

Specifically, the strategy takes positions at time t if $r_t^{S\&P} < x$, i.e., the S&P return over the week between $t-1$ and t is below a certain threshold x. It then takes a position in currency i against the dollar if two conditions are jointly satisfied: (i) hedge funds have reduced their positions in currency i—i.e., $|nlf_{i,t}| < |nlf_{i,t-1}|$—, and (ii) currency i has moved in the direction of the change in the net hedge fund position—i.e., $\text{sign}(r_{i,t}) = \text{sign}(\Delta nlf_{i,t})$—to separate plausibly flow-induced currency movements from, say, fundamental exchange rate movements.

I formulate two versions of this strategy. In the *contract-weighted* version of this strategy, the positions taken in different currencies in any given week are scaled to be proportional in size to the change in hedge fund positions: Let $\Omega_{i,t}^{CW}$ denote the number of futures contracts in currency i against the dollar, included in the strategy at time t:

$$\Omega_{i,t}^{CW} = \frac{\omega_{i,t}^{CW}}{\sum_j |\omega_{j,t}^{CW}| e_{j,t} s_j}, \text{ and } \omega_{i,t}^{CW} = -\Delta nlf_{i,t} \; \mathbb{1}_{\{r_t^{S\&P} < x\}} \; \mathbb{1}_{\{\Delta |nlf|_{i,t} < 0\}} \; \mathbb{1}_{\{\text{sign}(\Delta nlf_{i,t}) = \text{sign}(r_{i,t})\}}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function which takes value 1 if $\cdot$ is true, and 0 otherwise. Given the exchange rate $e_{i,t}$ and contract size s_i in units of foreign currency, $e_{i,t}s_i$ expresses the dollar notional of each contract and the term in the denominator therefore scales the positions to ensure that the gross notional of the total position is constant through time at \$1.

The *equal-weighted* version of the strategy fixes the dollar notional of each individual position, such that all non-zero positions taken at any point in time have the same absolute dollar exposure. Denote by $\Omega_{i,t}^{EW}$ the dollar notional amount in futures contracts of currency i against the dollar, then

$$\Omega_{i,t}^{EW} = \frac{\omega_{i,t}^{EW}}{\sum_j |\omega_{j,t}^{EW}|}, \text{ and } \omega_{i,t}^{EW} = -\text{sign}(\Delta nlf_{i,t})\ \mathbb{1}_{\{r_t^{S\&P}<x\}}\ \mathbb{1}_{\{\Delta|nlf|_{i,t}<0\}}\ \mathbb{1}_{\{\text{sign}(\Delta nlf_{i,t})=\text{sign}(r_{i,t})\}}.$$

Table C.8 reports the key return characteristics of this strategy for the threshold levels $x = 0$, $x = -3\%$, and no threshold at all. Since the strategy relies on a previous deterioration in the market environment (that is, a negative S&P return), it is only active in a subset of the weeks in the sample from January 2010 until June 2017. For the zero-threshold, this subset includes 138 out of 389 weeks in the sample. The strategy enters positions based on 1-week forward exchange rates (obtained from Bloomberg), and accounts for transaction costs by implementing long (short) positions at the ask (bid) price. Measuring performance over the active weeks, both versions of the strategy are economically profitable, with unlevered mean returns of 12bps (contract-weighted) and 15bps (equal-weighted) per week and annualized Sharpe ratios of 0.74 and 0.97, respectively. Over those 138 weeks, the strategy achieves a cumulative unlevered return of 17.22% (contract-weighted) and 22.29% (equal-weighted).

Under a stricter conditioning rule, where the strategy only becomes active if the previous week's S&P return was below −3%, the strategy only trades in 18 weeks, average returns rise to, respectively, 38bps and 35bps per week, and the annualized Sharpe ratios rise to 2.06 and 1.97, respectively.[11] For comparison, the unconditional strategy—which takes positions regardless of S&P returns—yields weekly excess returns close to 0, with Sharpe ratios of only 0.23 and 0.04, respectively. This comparison suggests, that hedge fund flows are only associated with temporary price dislocations (which revert within the next week), when these flows occur during times, when risky assets sell off.

3.2.4 Joint bets and currency co-movement

Next, I take a look at the impact of hedge fund positioning on exchange rate variation at even higher frequencies than the daily data used in the previous subsections. In the spirit of Antón and Polk (2014) and Lou and Polk (2013), I examine the correlations of currencies that are similarly exposed to trading by hedge funds, in

[11] Accounting for inactive periods in the annualization, Sharpe ratios range from 0.42 to 0.58.

intraday data. I ask whether currencies, which are traded in the same direction (long or short) by hedge funds, share common variation at intraday frequencies.

Let $\rho_{i,j,t}$ be the correlation in the exchange rate movements of currencies i and j over the week following date t. To capture a richer picture of variation over the course of one week, I obtain intraday data on exchange rates for AUD, CAD, CHF, EUR, GBP, JPY, MXN, and NZD spanning the period from July 2010 to January 2017. Correlations are computed over 15-minute intervals on the 7 days following the observation of hedge fund positions. To capture whether two currencies are traded in the same direction by hedge funds, I consider the differential in the scaled positioning variable $\widetilde{nlf}$. For currencies which are either both bought or both sold by hedge funds, and where hedge fund trading accounts for a similar share of overall trading activity, the $\widetilde{nlf}$ variable will take similar values and the differential will be small. If currencies are "connected" through hedge fund trading even at intraday frequencies, the observed correlation will be decreasing in the $\widetilde{nlf}$-differential.

In order to absorb other (time-invariant) sources of correlation, I include currency-pair fixed effects. For comparison, I can compute the same positioning-differential for the positions of institutional investors. To account for other time-varying currency characteristics that may lead to high-frequency correlation, I add the differential forward discount against the dollar, fd^w, as a competing predictor and estimate the following regressions:

$$\rho_{i,j,t} = \alpha_{i,j} + \beta \mid \widetilde{nlf}_{i,t} - \widetilde{nlf}_{j,t} \mid + \gamma \mid \widetilde{nam}_{i,t} - \widetilde{nam}_{j,t} \mid + \phi \mid fd^w_{i,t} - fd^w_{j,t} \mid + \varepsilon_{i,j,t} \tag{3.11}$$

$$\rho_{i,j,t} = \alpha_{i,j} + \beta \mid \widetilde{nlf}_{i,t} - \widetilde{nlf}_{j,t} \mid + \varepsilon_{i,j,t} \tag{3.12}$$

$$\rho_{i,j,t} = \alpha_{i,j} + \gamma \mid \widetilde{nam}_{i,t} - \widetilde{nam}_{j,t} \mid + \varepsilon_{i,j,t} \tag{3.13}$$

$$\rho_{i,j,t} = \alpha_{i,j} + \phi \mid fd^w_{i,t} - fd^w_{j,t} \mid + \varepsilon_{i,j,t}. \tag{3.14}$$

Table C.9 reports the results. Out of $\widetilde{nlf}$, $\widetilde{nam}$, and fd^w, the only variable that predicts intraday correlation—with a p-value of 0.1 at the margins of conventional significance—is $\widetilde{nlf}$: two currencies are more strongly correlated, when they are both included on the same side of hedge fund positioning, be it long or short. The same is not true for the positions of institutional investors. Brunnermeier et al. (2008) run a regression similar to (3.14) and show that currencies with similar interest rates have higher pairwise correlations in daily data over the next quarter. I find that this does not hold for shorter horizons and higher frequencies, while the direct statistical relationship between hedge fund positioning and correlation is discernible even at intraday frequencies.

3.3 Alternative explanations and empirical concerns

Is this strong relationship between hedge fund positioning and market risk exposures driven by the price impact of unwinding positions in "risk-off" states, when equity markets are down and the VIX is up? For instance, one might expect hedge fund positions to reflect unobservable signals about conditional risk premia. If these risk premia are related to market risk, positions would predict betas, not because one is driving the other but because both are driven by underlying fundamental risk profiles. Given the size of the futures market relative to total FX trading, another important question is whether futures positions are a relevant measure of the overall positioning of different groups of market participants. This section tries to address these questions in turn, starting with the latter.

3.3.1 Flows and returns

For hedge fund positions to be the driver of time-varying currency betas, the unwinding of these positions must have price impact. An important concern when looking at futures data is that the futures market covers only a small fraction of currency trading—most trades are done over-the-counter in the forward market and will therefore not be reported to the CFTC. Nonetheless, the futures data are the best publicly available indication of the overall positioning of market participants and I see no reason to expect the lack of complete coverage to bias the empirical results in a particular direction.

As empirical support for the relevance of the futures data, I can test whether exchange rates move with the hedge fund positions in the futures market. Price impact implies a systematic contemporaneous association between changes in the positions of traders (i.e., portfolio flows of hedge funds) and exchange rate movements. I regress net currency movements, $r_{i,t} = e_{i,t}/e_{i,t-1} - 1$, on the changes in nlf and $\widetilde{nlf}$. This first set of regressions merely serves as a simple test of whether or not the futures data are consistent with this prediction and a potential lack of contemporaneous association between flows and returns would cast doubt on the interpretation of the link between currency betas and futures positions. The contemporaneous regressions take the form

$$r_{i,t} = \alpha_i + \eta_i \Delta nlf_{i,t} + \gamma_i \Delta nam_{i,t} + \varepsilon_{i,t} \tag{3.15}$$

$$r_{i,t} = \alpha_i + \eta_i \Delta \widetilde{nlf}_{i,t} + \gamma_i \Delta \widetilde{nam}_{i,t} + \varepsilon_{i,t}, \tag{3.16}$$

where $\Delta nlf_{i,t} = nlf_{i,t} - nlf_{i,t-1}$ is the week-on-week change in the net long position of "Leveraged Funds" in currency i versus the US dollar (similarly for the scaled variable $\widetilde{nlf}$). The results are reported in Table C.10: for both regressions the estimates for η_i are positive and significant at 1% for all currencies. The statistical significance of this result is less pervasive for institutional investors, where most estimates, $\hat{\gamma}_i$, are positive, but only significant for 5 out of 9 currencies, at conventional confidence levels. The observable futures data are therefore consistent with price impact from hedge fund portfolio flows.[12]

3.3.2 Anticipation of risk premia

The key empirical concern in this setting is reverse causality. Do positions predict betas because unwinding has price impact, or do the anticipated betas lead hedge funds to position themselves in the way they do? The question boils down to whether or not the weekly variations in *fundamental* currency correlations with equity markets and the VIX are predictable.

Hedge fund positions are not driven by interest rate differentials.—Prompted by the profitability of the carry trade, the variable that is most traditionally associated with currency risk premia is the interest rate differential (Verdelhan, 2010; Lustig et al., 2011). Consider the one-week forward discount of currency i versus the dollar, denoted by $fd^w_{i,t}$: Figure C.2 plots the means of $\widetilde{nlf}$ against fd^w by currency. The confidence ellipses represent the joint distribution of the two variables under normality, and are scaled to contain 20% of the observations in population to make the figure readable. The relative sizes of the ellipses visualize the relative volatilities of the variables, and their orientation captures the time-series correlation of $\widetilde{nlf}$ and fd^w for the given currency. The cross-sectional correlation is negative, suggesting that hedge funds tend to be long in high interest currencies *on average*, but the within-currency correlations span a wide range from -0.61 (JPY) to 0.48 (NZD) and

[12]The results from these contemporaneous regressions are equally consistent with the reverse interpretation that hedge fund portfolio flows "chase" returns rather than "driving" them. The result does not present evidence of causality but is meant to provide a sense-check of the futures data as these are used as a proxy for the unobservable OTC positions, which make up the majority of total FX trading.

the slope coefficient for the within-currency regression is not significantly different from zero. Another way of looking for carry trade activity is to consider the time-series correlation in the $\widetilde{nlf}$ variable for currency pairs that are conventionally on opposite sides of the carry trade: The yen and the Swiss franc are among the most

	NZD	AUD	CAD
JPY	0.029	0.184	0.170
CHF	0.208	0.089	-0.136

commonly used funding currencies of the carry trade, while NZD, AUD, and CAD have been commonly used as investment currencies, so carry trade activity would suggest a negative correlation between the net hedge fund exposures in the funding and investment currencies.[13] However, 5 out of 6 of these pairwise correlations are positive (and sizeable in some cases), with CAD-CHF the only negative correlation.

Beta predictability is not driven by scheduled FOMC announcements.—One particular example of a setting in which risk premia (and potentially also currency betas) may reasonably be predictable even at a weekly horizon is the exposure to scheduled FOMC announcements. Savor and Wilson (2014) and Mueller et al. (2017) find that these announcements are accompanied by substantial excess returns of foreign currencies against the US dollar, and interpret these returns as compensation for monetary policy uncertainty. To test whether the results reported so far are driven by a strong relationship between hedge fund positioning and subsequently realized currency betas around a scheduled FOMC announcement, I consider those weeks separately from non-announcement weeks. Table C.11 shows the results for regression (3.10) splitting the sample into FOMC announcement weeks and non-announcement weeks. The results are larger in magnitude and more significant over non-announcement weeks, indicating that the headline results presented in this chapter are not driven by reverse causality stemming from anticipated risk exposures ahead of FOMC announcements.

[13] As CAD interest rates have been low in historical comparison after the financial crisis, the currency may have lost much of its previous status as a common investment currency in the carry trade.

3.4 Conclusion

This chapter documents a close link between the positioning of leveraged market participants and currency betas since the financial crisis. A currency is more exposed to movements in the S&P or the VIX when hedge funds hold long positions in that currency. Hedge fund positions in the futures market, which are observable at weekly frequencies, predict realized exposures over the subsequent week, both in the cross section and within-currency. At the same time, the size of open positions held by hedge funds and intermediaries, decreases following a negative shock to the S&P or an increase in the VIX. A similar exposure of futures positions to these adverse shocks is not present in the years 2006-2009, and at the same time, the strong link between hedge fund positions and currency betas does not exist in this earlier period.

My findings are consistent with the interpretation that (i) the balance sheets of hedge funds have become more responsive to market risk, and hedge fund positions in the futures market are unwound in response to adverse equity market shocks, (ii) hedged intermediaries as the counterparties to hedge funds' futures positions pass on the resulting flows to the spot market, as they unwind hedges to their futures positions, and (iii) the price impact from these unwinding flows in the spot market endogenously exposes the respective currency to equity market risk.

I design a trading strategy to exploit the temporary price dislocations in currency markets caused by unwinding hedge fund positions in response to an adverse shock to the market environment. The strategy is highly economically attractive with an annualized Sharpe ratio of up to 2. I find no evidence for alternative explanations for the link between hedge fund positions and currency betas: the association of futures positions and currency betas is not driven by a common correlation with expected returns compensating investors for monetary policy uncertainty around scheduled FOMC announcements.

Bibliography

A. Abadie and J. Gardeazabal. The economic costs of conflict: A case study of the Basque Country. *American Economic Review*, 93(1):113–132, 2003.

A. Abadie, A. Diamond, and J. Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control. *Journal of the American Statistical Association*, 105(490):493–505, 2010.

V. Acharya and S. Steffen. The "greatest" carry trade ever? Understanding eurozone bank risks. *Journal of Financial Economics*, 115:215–236, 2015.

V. V. Acharya, I. Drechsler, and P. Schnabl. A pyrrhic victory? Bank bailouts and sovereign credit risk. *Journal of Finance*, 69:2689–2739, 2014.

M. Adler and B. Dumas. International portfolio choice and corporation finance: A synthesis. *Journal of Finance*, 38(3):925–984, 1983.

M. Aguiar, M. Amador, E. Farhi, and G. Gopinath. Coordination and crisis in monetary unions. *Quarterly Journal of Economics*, 130(4):1727–1779, 2015.

M. Antón and C. Polk. Connected stocks. *Journal of Finance*, 69:1009–1127, 2014.

M. Antón, S. Mayodormo, and M. Rodriguez-Moreno. Dealing with dealers: Sovereign CDS comovements in Europe. *Working Paper*, 2017.

C. Arellano, Y. Bai, and S. Lizarazo. Sovereign risk contagion. *Working Paper*, 2018.

P. Augustin. Sovereign credit default swap premia. *Journal of Investment Management*, 12(2):65–102, 2014.

P. Augustin, M. G. Subrahmanyam, D. Y. Tang, and S. Q. Wang. Credit default swaps: A survey. *Foundations and Trends in Finance*, 9(1-2):1–196, 2014.

P. Augustin, M. Chernov, and Dongho Song. Sovereign credit risk and exchange rates: Evidence from CDS quanto spreads. *NBER Working Paper No. 24506*, 2018.

D. K. Backus, S. Foresi, and C. I. Telmer. Affine term structure models and the forward premium anomaly. *Journal of Finance*, 56(1):279–304, 2001.

J. Bai and P. Collin-Dufresne. The CDS-bond basis. *Financial Management*, forthcoming.

S. Baker, N. Bloom, and S. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

P. Barroso and P. Santa-Clara. Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 50(5):1037–1056, 2015.

N. Battistini, M. Pagano, and S. Simonelli. Systemic risk, sovereign yields and bank exposures in the euro crisis. *Economic Policy*, 29:203–51, 2014.

C. Bayer, C. H. Kim, and A. Kriwoluzky. The term structure of redenomination risk. *CEPR Discussion Paper*, (No. 12965), 2018.

A. Berndt, R. Jarrow, and C. O. Kang. Restructuring risk in credit default swaps: An empirical analysis. *Stochastic Processes and their Applications*, 117:1724–1749, 2007.

Bank for International Settlements BIS. OTC derivatives statistics at end-June 2015. *Statistical Release, Monetary and Economic Department*, 2015.

P. Bolton and H. Huang. The capital structure of nations. *Review of Finance*, 22 (1):45–82, 2017.

W. Branson and D. Henderson. The specification and influence of asset markets. *Handbook of International Economics*, 2:749–805, 1985.

D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978.

M. K. Brunnermeier, S. Nagel, and L. H. Pedersen. Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23:313–347, 2008.

M. K. Brunnermeier, L. Garicano, P. Lane, M. Pagano, R. Reis, T. Santos, D. Thesmar, S. van Nieuwerburgh, and D. Vayanos. The sovereign-bank diabolic loop and esbies. *American Economic Review Papers and Proceedings*, 106:508–512, 2016a.

M. K. Brunnermeier, H. James, and J.-P. Landau. The euro and the battle of ideas. *Princeton University Press*, 2016b.

C. Burnside. The cross section of foreign currency risk premia and consumption growth risk: Comment. *American Economic Review*, 101(7):3456–3476, 2011.

C. Burnside, M. Eichenbaum, I. Kleshchelski, and S. Rebelo. Do peso problems explain the returns to the carry trade? *Review of Financial Studies*, 24(3):853–891, 2011.

J. Y. Campbell, K. Serfaty-De Medeiros, and L. M. Viceira. Global currency hedging. *Journal of Finance*, 65(1):87–121, 2010.

G. Cenedese, R. Payne, L. Sarno, and G. Valente. What do stock markets tell us about exchange rates? *Review of Finance*, 20(3):1045–1080, 2016.

M. Chernov, J. Graveline, and I. Zviadadze. Crash risk in currency returns. *Journal of Financial and Quantitative Analysis*, 53(1):137–170, 2018.

T. Cho. Turning alphas into betas: Arbitrage and endogenous risk. *Working Paper*, 2018.

N. Coeurdacier and H. Rey. Home bias in open economy financial macroeconomics. *Journal of Economic Literature*, 51(1):63–115, 2013.

R. Colacito and M. M. Croce. Risks for the long run and the real exchange rate. *Journal of Political Economy*, 119(1):153–182, 2011.

M. Dahlquist and J. Penasse. The missing risk premium in exchange rates. *Working Paper*, 2017.

R. A. De Santis. Redenomination risk. *Journal of Money, Credit and Banking*, forthcoming.

P. Della Corte, T. Ramadorai, and L. Sarno. Volatility risk premia and exchange rate predictability. *Journal of Financial Economics*, 20(1), 2016a.

P. Della Corte, S. Riddiough, and L. Sarno. Currency premia and global imbalances. *Review of Financial Studies*, 2016b.

P. Della Corte, L. Sarno, M. Schmeling, and C. Wagner. Sovereign risk and currency returns. *Working Paper*, 2016c.

F. X. Diebold and R. S. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3):253–263, 1995.

W. Du and J. Schreger. Local currency sovereign risk. *Journal of Finance*, 71(3): 1027–1070, 2016.

W. Du, A. Tepper, and A. Verdelhan. Deviations from covered interest parity. *Journal of Finance*, 73(3):915–957, 2018.

J. Duarte, F. A. Longstaff, and F. Yu. Risk and return in fixed-income arbitrage: Nickels in front of a steamroller? *Review of Financial Studies*, 20(3):769–811, 2007.

D. Duffie and M. Thukral. Redesigning credit derivatives to better cover sovereign default risk. *Working Paper No. 3107*, 2012.

C. Engel. Exchange rates, interest rates, and the risk premium. *American Economic Review*, 106(2):436–474, 2016.

E. F. Fama. Forward and spot exchange rates. *Journal of Monetary Economics*, 14: 319–338, 1984.

E. Farhi and X. Gabaix. Rare disasters and exchange rates. *The Quarterly Journal of Economics*, 131(1):1–52, 2016.

E. Farhi and J. Tirole. Deadly embrace: Sovereign and financial balance sheets doom loops. *Review of Economic Studies*, forthcoming.

A. Fontana and M. Scheicher. An analysis of euro area sovereign CDS and their relation with government bonds. *Journal of Banking and Finance*, 62:126–140, 2016.

K. R. French and J. M. Poterba. Investor diversification and international equity markets. *American Economic Review*, 81(2):222–226, 1991.

X. Gabaix and M. Maggiori. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420, 2015.

S. Giglio. Credit default swap spreads and systemic financial risk. *Working Paper*, 2014.

P.-O. Gourinchas and H. Rey. International financial adjustment. *Journal of Political Economy*, 115(4):665–703, 2007.

P.-O. Gourinchas, H. Rey, and N. Govillot. Exorbitant privilege and exorbitant duty. *IMES Discussion Paper Series*, (10-E-20), 2010.

A. Goyal and I. Welch. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.

D. Gromb and D. Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66:361–407, 2002.

L. P. Hansen and R. J. Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5): 829–853, 1980.

L. P. Hansen and R. Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991.

D. Harvey, S. Leybourne, and P. Newbold. Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13:281–291, 1997.

T. A. Hassan. Country size, currency unions, and international asset returns. *The Journal of Finance*, 68(6):2269–2308, 2013.

T. A. Hassan and R. C. Mano. Forward and spot exchange rates in a multi-currency world. *Quarterly Journal of Economics*, forthcoming.

T. A. Hassan, T. M. Mertens, and T. Zhang. Currency manipulation. *NBER Working Paper 22790*, 2016.

Z. He, B. T. Kelly, and A. Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126:1–35, 2017.

ISDA. 2003 credit derivatives definitions. *International Swaps and Derivatives Association, Inc.*, 2003.

ISDA. 2014 credit derivatives definitions. *International Swaps and Derivatives Association, Inc.*, 2014.

Z. Jiang. Fiscal cyclicality and currency risk premia. *Working Paper*, 2018.

O. Jordà and A. M. Taylor. The carry trade and fundamentals: Nothing to fear but FEER itself. *Journal of International Economics*, 88:74–90, 2012.

J. W. Jurek. Crash neutral currency carry trades. *Journal of Financial Economics*, 113(3):325–347, 2014.

B. Kelly, L. Pastor, and P. Veronesi. The price of political uncertainty: Theory and evidence from the option market. *Journal of Finance*, 71:2417–2480, 2016.

S. Klingler and D. Lando. Safe haven CDS premiums. *Review of Financial Studies*, 31(5):1855–1895, 2018.

L. Kremens and I. W. R. Martin. The quanto theory of exchange rates. *American Economic Review*, 109(3):810–843, 2019.

A. Krishnamurthy, S. Nagel, and A. Vissing-Jorgensen. ECB policies involving government bond purchases: Impact and channels. *Review of Finance*, 2018.

M. Leombroni, A. Vedolin, G. Venter, and P. Whelan. Central bank communication and the yield curve. *Working Paper*, 2017.

M. Lettau, M. Maggiori, and M. Weber. Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, 114(2):197–225, 2014.

R. M. Levich. CIP then and now: A brief survey of measuring and exploiting deviations from covered interest parity. *Working Paper*, 2017.

K. K. Lewis. Trying to explain home bias in equities and consumption. *Journal of Economic Literature*, 37(2):571–608, 1999.

F. A. Longstaff, J. Pan, L. H. Pedersen, and K. J. Singleton. How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3:75–103, 2011.

D. Lou and C. Polk. Comomentum: Inferring arbitrage activity from return correlations. *Working Paper*, 2013.

H. Lustig and A. Verdelhan. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1):89–117, 2007.

H. N. Lustig, N. L. Roussanov, and A. Verdelhan. Common risk factors in currency markets. *Review of Financial Studies*, 24(11), 2011.

H. N. Lustig, N. L. Roussanov, and A. Verdelhan. Countercyclical currency risk premia. *Journal of Financial Economics*, 111(3):527–553, 2014.

R. C. Mano. Exchange rates upon sovereign default. *Working Paper*, 2013.

I. W. R. Martin. The forward premium puzzle in a two-country world. *Working Paper*, 2013a.

I. W. R. Martin. Consumption-based asset pricing with higher cumulants. *Review of Economic Studies*, 80(2):745–773, 2013b.

I. W. R. Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433, 2017.

I. W. R. Martin. Options and the gamma knife. *Journal of Portfolio Management*, 44(6):47–55, 2018.

I. W. R. Martin and C. Wagner. What is the expected return on a stock? *Journal of Finance*, forthcoming.

R. Meese and K. Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24, 1983.

L. Menkhoff, L. Sarno, M. Schmeling, and A. Schrimpf. Carry trades and global foreign exchange volatility. *Journal of Finance*, 67(2):681–718, 2012.

S. Miranda-Agrippino and H. Rey. US monetary policy and the global financial cycle. *Working Paper*, 2019.

P. Mueller, A. Tahbaz-Salehi, and A. Vedolin. Exchange rates and monetary policy uncertainty. *Journal of Finance*, 72(3):1213–1252, 2017.

R. Neuberg, P. Glasserman, B. Kay, and S. Rajan. The market-implied probability of government support for distressed european banks. *Working Paper*, 2018.

M. Oehmke and A. Zawadowski. The anatomy of the CDS market. *Review of Financial Studies*, 30(1):80–119, 2017.

J. Pan and K. Singleton. Default and recovery implicit in the term structure of sovereign CDS spreads. *Journal of Finance*, 63:2345–2384, 2008.

L. Pastor and P. Veronesi. Political uncertainty and risk premia. *Journal of Financial Economics*, 110(3):520–545, 2013.

A. Patton, D. N. Politis, and H. White. Correction to automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 28:372–375, 2009.

D. N. Politis and H. White. Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23:53–70, 2004.

C. M. Reinhart and K. S. Rogoff. From financial crash to debt crisis. *American Economic Review*, 101(5):1676–1706, 2011.

S. A. Ross. Options and efficiency. *Quarterly Journal of Economics*, 90(1):75–89, 1976.

B. Rossi. Exchange rate predictability. *Journal of Economic Literature*, 51(4):1063–1119, 2013.

T. J. Sargent and N. Wallace. Some unpleasant monetarist arithmetic. *Monetarism in the United Kingdom*, pages 15–41, 1984.

L. Sarno, P. Schneider, and C. Wagner. Properties of foreign exchange risk premiums. *Journal of Financial Economics*, 105(2):279–310, 2012.

P. Savor and M. Wilson. Asset pricing: A tale of two days. *Journal of Financial Economics*, 113:171–201, 2014.

A. Shleifer and R. W. Vishny. The limits of arbitrage. *The Journal of Finance*, 52 (1):35–55, 1997.

J. J. Siegel. Risk, interest rates, and the forward exchange. *Quarterly Journal of Economics*, 86(2):303–309, 1972.

Simmons & Simmons. The aftermath of the 2014 ISDA credit derivatives definitions. *available at http://www.elexica.com/en/resources/podcast/ /media/5A19213DC80D4B3DA6EC068617C0F438.ashx*, 2016.

C. A. Sims. A simple model for study of the determination of the price level and the interaction of monetary and fiscal policy. *Economic Theory*, 4:381–399, 1994.

E. N. Siriwardane. Limited investment capital and credit spreads. *Journal of Finance*, forthcoming.

B. H. Solnik. An equilibrium model of the international capital market. *Journal of Economic Theory*, 8:500–524, 1974.

L. L. Tesar and I. M. Werner. Home bias and high turnover. *Journal of International Money and Finance*, 14(4):467–492, 1995.

A. Verdelhan. A habit-based explanation of the exchange rate risk premium. *Journal of Finance*, 65(1):123–145, 2010.

F. E. Warnock. Home bias and high turnover reconsidered. *Journal of International Money and Finance*, 21(6):795–805, 2002.

A. Appendix to Currency Redenomination Risk

A.1 Tables

Table A.1: Summary statistics – CDS spreads

This table reports the summary statistics of CDS spreads collected from Markit for a cross-section of Eurozone countries. The maturity in each case is five years, which is typically the most liquid maturity in CDS markets. CDS spreads for contracts under, respectively, 2003 and 2014 ISDA definitions are denoted by CR and $CR14$. The daily data run from September 22nd, 2014 to June 19th, 2018. CDS spreads are annualized and reported in basis points. Rows with $\mu(\cdot)$, $\sigma(\cdot)$, and $\rho(\cdot,\cdot)$ report, respectively, the mean, standard deviation, and correlation.

Country	FRA	ITA	GER	AUT	BEL	ESP	IRE	NED	POR
$\mu(CR14)$	32.69	132.26	15.95	23.79	34.52	81.03	48.60	20.17	195.36
$\sigma(CR14)$	11.38	26.95	3.79	5.80	12.54	19.08	14.66	5.28	70.39
$\mu(CR)$	27.57	106.91	14.21	21.61	30.89	68.25	42.83	18.55	172.90
$\sigma(CR)$	10.53	23.52	3.94	6.28	12.73	22.07	16.05	5.51	71.37
$\rho(CR14, CR)$	0.88	0.74	0.97	0.99	0.99	0.97	0.99	0.98	0.99
$\mu(CR14 - CR)$	5.10	26.30	1.81	2.27	3.65	13.04	5.86	1.67	22.71
$\sigma(CR14 - CR)$	5.41	18.43	0.92	1.01	1.28	6.34	2.60	1.03	9.14
$\rho(CR14 - CR, CR14)$	0.39	0.52	-0.05	-0.41	-0.09	-0.36	-0.47	-0.13	-0.04

Table A.2: Summary statistics – Redenomination spreads (RS)

This table reports the summary statistics of the French, Italian, and German redenomination spreads (RS) constructed as a difference-in-difference measure: $RS_{i,t} = CR14_{i,t} - CR_{i,t} - (CR14_{s(i),t} - CR_{s(i),t})$, where $s(i)$ denotes variables relating to a synthetic control country constructed to match country i. The daily data run from September 22nd, 2014 to June 19th, 2018. Redenomination spreads are annualized and reported in basis points. Rows denoted by $\mu(\cdot)$, $\sigma(\cdot)$, $\rho(\cdot,\cdot)$, and Max$(\cdot)$ report, respectively, mean, standard deviation, correlation and maximum.

Country	FRA	ITA	GER
$\rho(RS, CR)$	-0.03	-0.01	0.08
$\mu(RS)$	1.46	8.23	0.10
$\sigma(RS)$	3.99	10.81	0.99
$\mu(RS/CR14)$	0.04	0.06	0.00
Max$(RS/CR14)$	0.40	0.32	0.42

Table A.3: Regression of Eurozone sovereign yields on RS

This table reports the results for time-series regressions of Eurozone (plus Denmark) government bond yields on French and Italian redenomination spreads, controlling for €-denominated overnight swap rates (OIS).

$$y_{j,T,t} - OIS_{€,T,t} = \alpha_{j,T} + \beta_{FRA,j,T} RS_{FRA,t} + \beta_{ITA,j,T} RS_{ITA,t} + \varepsilon_{j,T,t}, \qquad (1.2)$$

for maturity $T = 5$ years. Newey–West standard errors are reported in parentheses. The daily data run from September 2014 to June 2018. Yields, swap rates, and redenomination spreads are measured in %-points.

Country	GER	AUT	DEN†	NED	IRE	BEL	ESP	ITA	FRA	POR
RS_{FRA}	-0.965	-0.939	-0.613	-0.210	0.514	0.042	0.805	0.365	1.230	5.410
	(0.434)	(0.221)	(0.249)	(0.335)	(0.364)	(0.216)	(0.520)	(0.522)	(0.210)	(0.757)
RS_{ITA}	-1.316	-0.441	-0.590	-0.770	-1.109	-0.433	-1.357	1.213	-0.678	-1.898
	(0.356)	(0.165)	(0.181)	(0.268)	(0.294)	(0.169)	(0.389)	(0.331)	(0.152)	(0.414)
Intercept	-0.390	-0.058	-0.033	-0.094	0.205	-0.021	0.612	0.615	0.045	1.479
	(0.028)	(0.013)	(0.024)	(0.020)	(0.025)	(0.015)	(0.031)	(0.026)	(0.012)	(0.077)
R^2	0.386	0.370	0.173	0.267	0.351	0.149	0.263	0.319	0.359	0.116
Obs.	967	965	967	965	818	967	966	967	967	967

†: included as a quasi-Eurozone member.

Table A.4: Regression of € and $ FX rates on French and Italian *RS*

This table reports the results for time-series regressions of exchange rate variables on French and Italian redenomination spreads, controlling for overnight swap rates.

$$e_{€,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{€} OIS_{€,t} + \varepsilon_t, \tag{A.1}$$
$$e_{\$/€,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{€} OIS_{€,t} + \gamma_{\$} OIS_{\$,t} + \varepsilon_t, \tag{A.2}$$
$$e_{\$,t} = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma_{\$} OIS_{\$,t} + \varepsilon_t. \tag{A.3}$$

where $e_{€,t}$, $e_{\$/€,t}$, and $e_{\$,t}$ denote, respectively, the natural logarithms of the Bloomberg euro spot index, the euro-dollar exchange rate, and the ICE US-dollar spot index. The euro-dollar exchange rate is defined such that an increase reflects an appreciation of the euro against the dollar. For the two indices, an increase in e reflects an appreciation of the respective currency against a trade- and liquidity-weighted basket of other currencies. Newey–West standard errors (max. 10 lags) are reported in parentheses. The daily data run from September 2014 to June 2018. Redenomination spreads are measured in basis points.

Currency	EUR index	EURUSD	USD index
RS_{FRA}	−0.003	−0.005	0.004
	(0.000)	(0.000)	(0.000)
RS_{ITA}	0.001	0.000	0.000
	(0.000)	(0.000)	(0.000)
$OIS_{€}$	0.069	0.053	
	(0.015)	(0.044)	
$OIS_{\$}$		0.036	−0.036
		(0.019)	(0.006)
Intercept	6.770	0.081	4.601
	(0.005)	(0.027)	(0.008)
R^2	0.307	0.347	0.255
Obs.	969	969	972

Table A.5: Regression of German and US government bond yields on RS

This table reports the results for time-series regressions of German (Panel A) and US (Panel B) government bond yields on French and Italian redenomination spreads, controlling for both €- and \$-denominated overnight swap rates.

$$y_{GER,T,t} = \alpha + \beta_{FRA,T} RS_{FRA,t} + \beta_{ITA,T} RS_{ITA,t} + \gamma_T OIS_{€,T,t} + \varepsilon_{i,T,t}, \tag{1.2}$$
$$y_{US,T,t} = \alpha + \beta_{FRA,T} RS_{FRA,t} + \beta_{ITA,T} RS_{ITA,t} + \gamma_T OIS_{\$,T,t} + \varepsilon_{i,T,t}, \tag{1.6}$$

for maturities $T = \{1, 2, 3, 5, 10\}$ years. Newey–West standard errors (max. 10 lags) are reported in parentheses. The daily data run from September 2014 to June 2018. Yields, swap rates, and redenomination spreads are measured in percentage points.

Maturity	1y	2y	3y	5y	10y
			Panel A: Bund yield, FRA		
RS_{FRA}	−1.266	−1.321	−1.087	−1.234	−0.186
	(0.182)	(0.251)	(0.269)	(0.401)	(0.154)
RS_{ITA}	−0.369	−0.612	−0.725	−1.101	−0.507
	(0.142)	(0.190)	(0.195)	(0.329)	(0.099)
OIS EUR	1.535	1.355	1.181	0.467	1.024
	(0.070)	(0.083)	(0.069)	(0.076)	(0.028)
Intercept	−0.013	−0.065	−0.124	−0.410	−0.115
	(0.012)	(0.018)	(0.020)	(0.027)	(0.011)
R^2	0.853	0.789	0.753	0.445	0.897
Obs	970	970	969	967	970
			Panel B: US Treasury yield, FRA		
RS_{FRA}	−0.723	−0.616	−0.490	-0.146	0.183
	(0.126)	(0.122)	(0.116)	(0.104)	(0.146)
RS_{ITA}	−0.073	−0.002	0.117	0.283	0.487
	(0.096)	(0.090)	(0.083)	(0.089)	(0.120)
OIS USD	1.070	1.007	0.959	0.892	0.768
	(0.011)	(0.017)	(0.021)	(0.025)	(0.026)
Intercept	−0.038	0.056	0.158	0.374	0.735
	(0.010)	(0.022)	(0.027)	(0.034)	(0.041)
R^2	0.989	0.975	0.966	0.952	0.913
Obs	970	970	970	970	970

Table A.6: Regression of corporate CDS spreads on French and Italian *RS*

This table reports the results for time-series regressions of corporate CDS spreads on French and Italian redenomination spreads, controlling for overnight swap rates.

$$S_t = \alpha + \beta_{FRA} RS_{FRA,t} + \beta_{ITA} RS_{ITA,t} + \gamma OIS_{€,t} + \varepsilon_t. \quad (A.4)$$

where S_t denotes the spread of different equally weighted portfolios of corporate CDS contracts: The first portfolio contains CDS written on the senior debt of the 125 investment grade corporates included in the iTraxx Europe CDS Index. The second and third portfolios contain, respectively, senior and subordinated CDS contracts for 30 investment grade financial companies (i.e., banks and insurance companies). The next ten portfolios split the original set of 125 CDS by industry into financial and non-financial companies, and by country into the five Eurozone countries with at least one financial company (Germany, the Netherlands, Italy, Spain, and France). Newey–West standard errors (max. 10 lags) are reported in parentheses. The row entitled $\mu(S_t)$ reports the time-series averages of the credit spreads in the respective portfolio. The daily data run from September 2014 to June 2018. All variables are measured in percentage points.

	All	Financials		GER		NED		ITA		ESP		FRA	
	Senior	Senior	Subord.	Fin.	Non-F.	Fin.	Non-F.	Fin.	Non-F.	Fin.	Non-F.	Fin.	Non-F.
RS_{FRA}	0.357	0.705	1.640	0.300	−0.012	1.229	0.297	1.510	0.179	0.794	0.017	1.475	0.460
	(0.183)	(0.266)	(0.465)	(0.209)	(0.254)	(0.217)	(0.119)	(0.358)	(0.192)	(0.588)	(0.484)	(0.357)	(0.166)
RS_{ITA}	−0.157	−0.157	−0.138	−0.071	−0.431	−0.514	−0.192	0.337	0.112	−0.750	−0.720	−0.705	−0.208
	(0.144)	(0.222)	(0.365)	(0.166)	(0.198)	(0.182)	(0.091)	(0.298)	(0.141)	(0.455)	(0.355)	(0.300)	(0.123)
$OIS_{€}$	−0.310	−0.638	−1.621	−0.724	−0.336	−0.505	−0.103	−1.421	−0.557	−1.177	−0.873	−0.428	−0.317
	(0.058)	(0.068)	(0.140)	(0.058)	(0.074)	(0.069)	(0.040)	(0.119)	(0.072)	(0.121)	(0.146)	(0.079)	(0.051)
Interc.	0.669	0.749	1.625	0.648	0.712	0.637	0.536	1.128	0.740	0.995	0.989	0.655	0.609
	0.0157	(0.020)	(0.038)	(0.017)	(0.020)	(0.019)	(0.010)	(0.034)	(0.016)	(0.040)	(0.039)	(0.023)	(0.014)
R^2	0.255	0.469	0.557	0.569	0.273	0.473	0.146	0.575	0.427	0.479	0.336	0.379	0.312
$\mu(S_t)$	0.664	0.753	1.656	0.655	0.679	0.618	0.525	1.193	0.758	0.957	0.939	0.622	0.602
Comp.	125	30	30	5	16	3	8	4	3	2	4	4	22
Obs..	967	967	967	967	967	967	967	967	967	967	967	967	967

Table A.7: Regression of Eurozone sovereign yields on RS and CR CDS spreads

This table reports the results for time-series regressions of Eurozone (plus Denmark) net government bond yields on French and Italian redenomination risk, controlling for credit risk through CR CDS spreads.

$$y_{j,T,t} - OIS_{€,T,t} = \alpha_{j,T} + \beta_{FRA,j,T} RS_{FRA,t} + \beta_{ITA,j,T} RS_{ITA,t} + \\ + \psi_{FRA,j,T} CR_{FRA,t} + \psi_{ITA,j,T} CR_{ITA,t} + \varepsilon_{j,T,t}, \tag{1.9}$$

for maturity $T = 5$ years. Newey–West standard errors are reported in parentheses. The daily data run from September 2014 to June 2018. Yields, swap rates, and CDS spreads are measured in %-points.

Country	GER	AUT	DEN†	NED	IRE	BEL	ESP	ITA	FRA	POR
RS_{FRA}	-1.650	-1.103	-0.916	-0.773	-0.243	-0.087	-0.024	-0.710	0.874	1.399
	(0.325)	(0.166)	(0.223)	(0.201)	(0.192)	(0.157)	(0.332)	(0.406)	(0.086)	(0.853)
RS_{ITA}	-0.596	-0.103	-0.413	-0.187	-0.478	0.009	-0.541	1.735	-0.284	-1.397
	(0.178)	(0.104)	(0.142)	(0.116)	(0.108)	(0.101)	(0.170)	(0.155)	(0.055)	(0.513)
CR_{FRA}	1.579	0.805	0.333	1.273	1.363	1.090	1.768	0.909	0.870	-0.350
	(0.176)	(0.088)	(0.248)	(0.107)	(0.120)	(0.094)	(0.144)	(0.152)	(0.085)	(0.452)
CR_{ITA}	-0.001	-0.108	0.092	0.007	0.069	-0.201	0.036	0.397	-0.013	2.428
	(0.114)	(0.046)	(0.092)	(0.005)	(0.070)	(0.047)	(0.099)	(0.065)	(0.035)	(0.194)
Intercept	-0.875	-0.190	-0.234	-0.492	-0.277	-0.143	0.030	-0.088	-0.209	-1.004
	(0.088)	(0.042)	(0.069)	(0.046)	(0.047)	(0.047)	(0.089)	(0.051)	(0.020)	(0.151)
R^2	0.718	0.637	0.243	0.792	0.871	0.563	0.690	0.771	0.840	0.723
Obs.	967	965	967	965	818	967	966	967	967	967

†: included as a quasi-Eurozone member.

Table A.8: Decomposition of Eurozone-G7 sovereign yields

This table reports the results for time-series regressions of five-year Eurozone sovereign yields on redenomination risk, credit risk (CR CDS spread), and five-year swap rates.

$$y_{j,t} = \alpha_j + \beta_j RS_{j,t} + \psi_j CR_{j,t} + \gamma_j OIS_{€,t} + \varepsilon_{j,t}, \quad (1.8)$$

for $j = \{FRA, ITA, GER\}$. The right panel uses the redenomination spreads of France and Italy as regressors, instead of the German redenomination spread. Newey–West standard errors are reported in parentheses. The daily data run from September 2014 to June 2018. Yields, swap rates, and CDS spreads are measured in %-points. The bottom panel reports t-statistics for the null hypothesis that the two β-coefficients are equal to one.

Country	FRA		ITA	GER
Subsample	full	02/17-06/18	full	full
RS_{FRA}	0.552	0.856		−1.827
	(0.091)	(0.128)		(0.412)
RS_{ITA}			1.128	−0.797
			(0.155)	(0.285)
CR_j	0.976	0.952	0.852	3.041
	(0.046)	(0.122)	(0.068)	(0.682)
$OIS_{€}$	1.021	1.110	1.447	0.803
	(0.029)	(0.062)	(0.067)	(0.080)
Intercept	−0.270	−0.302	−0.279	−0.855
	(0.013)	(0.023)	(0.071)	(0.097)
R^2	0.930	0.867	0.855	0.592
Obs.	967	351	969	966
t-stat: $\beta^{RS} = 1$	-4.92	-1.13	0.82	
t-stat: $\beta^{CR} = 1$	-0.53	-0.39	-2.19	2.99

Table A.9: Subsample regression of Eurozone sovereign yields

This table reports the results for time-series regressions of Eurozone (plus Denmark) sovereign bond yields on French and Italian redenomination risk, for the subsample January 2017 to June 2018 (controlling for French and Italian credit risk, regression (1.9)). Newey–West standard errors are reported in parentheses. The daily data run from January 2017 to June 2018. Yields, swap rates, and CDS spreads are measured in %-points.

Country	GER	AUT	DEN†	NED	IRE	BEL	ESP	ITA	FRA	POR
RS_{FRA}	-0.742	-0.558	-0.089	-0.271	0.370	0.243	1.165	-0.444	0.870	2.822
	(0.175)	(0.076)	(0.188)	(0.080)	(0.077)	(0.187)	(0.177)	(0.287)	(0.075)	(0.722)
RS_{ITA}	-0.339	0.357	0.387	0.044	-0.218	0.459	-0.507	1.141	-0.265	-1.650
	(0.090)	(0.071)	(0.131)	(0.047)	(0.054)	(0.106)	(0.106)	(0.131)	(0.052)	(0.535)
CR_{FRA}	-0.696	0.934	1.920	0.249	0.708	1.583	-1.674	-2.533	0.802	-2.139
	(0.280)	(0.295)	(0.529)	(0.193)	(0.183)	(0.375)	(0.284)	(0.486)	(0.166)	(1.466)
CR_{ITA}	0.520	-0.337	-0.639	0.220	0.103	-0.459	0.902	1.468	0.046	2.729
	(0.099)	(0.104)	(0.159)	(0.057)	(0.062)	(0.137)	(0.100)	(0.157)	(0.060)	(0.459)
Intercept	-1.090	-0.121	-0.039	-0.602	-0.286	-0.098	-0.270	-0.391	-0.262	-1.017
	(0.047)	(0.046)	(0.055)	(0.022)	(0.018)	(0.067)	(0.051)	(0.064)	(0.023)	(0.193)
R^2	0.566	0.409	0.262	0.780	0.887	0.330	0.798	0.932	0.913	0.813
Obs.	373	373	373	373	364	373	373	373	373	373

†: included as a quasi-Eurozone member.

Table A.10: Regression of Eurozone sovereign CDS spreads on RS

This table reports the results for time-series regressions of Eurozone sovereign CDS spreads on French and Italian redenomination risk.

$$CR14_{j,t} = \alpha_j + \beta_{FRA,j} RS_{FRA,t} + \beta_{ITA,j} RS_{ITA,t} + \varepsilon_{j,t}, \tag{1.10}$$

Newey–West standard errors are reported in parentheses. The daily data run from September 2014 to June 2018. CDS spreads are measured in basis points.

	5-year CR14 CDS spreads						
Country	GER	AUT	NED	IRE	BEL	ESP	POR
RS_{FRA}	0.317	0.285	0.491	1.043	0.151	0.296	5.619
	(0.058)	(0.090)	(0.068)	(0.249)	(0.246)	(0.368)	(0.760)
RS_{ITA}	−0.103	−0.278	−0.145	−0.615	−0.601	−0.835	−2.002
	(0.022)	(0.062)	(0.026)	(0.158)	(0.188)	(0.297)	(0.423)
Intercept	16.35	25.68	33.39	52.18	39.31	87.53	203.74
	(0.403)	(0.600)	(1.214)	(1.627)	(1.508)	(2.267)	(7.975)
R^2	0.118	0.223	0.136	0.181	0.249	0.202	0.116
Obs.	970	970	970	970	970	970	970

Table A.11: Bank-sovereign home bias

This table reports the relative exposures of banks to different sovereign issuers within liquid asset holdings. I consider net direct exposures in assets held as *available-for-sale (AFS)*, *held-for-trading (HFT)*, and *held-to-maturity (HTM)*. The data refer to balance sheet exposures as of December 31st, 2015, and are obtained from the European Banking Authority (EBA) and its reports on the stress tests conducted in 2016.

	Bank Country												
Sovereign	AUT	BEL	DEN	ESP	FRA	GER	IRE	ITA	NED	NOR	POL	SWE	UK
AUT	**20.1**	0.9	4.3	0.0	2.1	2.1	0.1	3.5	4.0	0.0	0.0	1.2	0.5
BEL	0.4	**40.9**	5.5	0.1	8.0	2.4	2.1	0.5	8.8	0.0	0.0	3.7	0.9
DEN	0.0	0.0	**12.2**	0.0	0.0	0.0	0.0	0.0	0.2	0.0	0.0	5.6	0.5
ESP	0.3	5.2	6.3	**50.5**	3.8	3.2	7.5	7.8	2.5	0.0	0.0	0.0	0.5
FRA	1.3	9.0	12.8	0.7	**38.8**	4.5	6.1	3.7	11.3	0.0	0.0	8.3	4.0
GER	1.8	0.3	10.2	0.0	5.3	**44.4**	1.0	3.9	14.6	5.5	0.0	19.1	6.7
IRE	0.0	1.8	3.1	0.0	0.5	0.8	**67.6**	0.1	0.0	0.0	0.0	0.1	0.2
ITA	1.3	10.2	4.8	9.0	8.3	4.6	7.4	**65.8**	3.0	0.0	0.0	0.3	1.3
NED	1.1	0.5	6.3	0.2	2.6	4.2	2.0	0.4	**29.4**	0.0	0.0	3.9	1.6
NOR	0.0	0.0	0.4	0.1	0.0	0.0	0.0	0.1	0.0	**63.2**	0.0	4.9	0.1
POL	5.5	1.8	0.0	2.0	0.9	3.8	0.6	2.6	3.9	0.0	**98.8**	0.0	0.2
SWE	0.2	0.0	7.9	0.0	0.1	0.3	0.0	0.2	0.6	7.0	0.0	**18.0**	0.4
UK	0.0	0.0	13.1	2.4	1.7	1.4	4.0	0.1	0.2	0.0	0.0	0.0	**26.8**
POR†	0.0	0.5	1.0	3.7	0.7	0.5	0.7	0.1	0.0	0.0	0.0	0.0	0.1
US	2.7	1.0	0.0	5.4	10.9	12.5	0.0	1.0	5.8	1.6	0.0	20.8	24.8
CH	0.0	0.0	0.0	0.0	0.2	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.5
CAN	0.0	0.5	0.0	0.0	1.3	0.7	0.0	0.3	0.7	20.7	0.0	0.9	2.5
Others	65.2	27.5	12.1	25.9	14.6	13.8	0.8	10.0	15.0	2.1	1.2	13.1	28.4

†: No Portuguese banks were included in the 2016 EBA stress test, due to a size threshold.

Table A.12: Summary statistics – Market depth

This table reports the market depth by contract type as the number of quote submissions from financial intermediaries to Markit. Rows denoted by $\mu(\cdot)$, $\sigma(\cdot)$, $\rho(\cdot,\cdot)$, Max$(\cdot)$, and Min$(\cdot)$ report, respectively, the mean, standard deviation, correlation, maximum, and minimum of the respective variables.

Country	FRA	ITA	GER	AUT	BEL	ESP	IRE	NED	POR
μ(Depth CR14)	4.74	6.32	4.29	4.48	4.50	5.93	4.44	3.77	5.63
μ(Depth CR)	4.60	5.62	4.16	5.04	4.30	5.84	4.91	4.11	6.90
Min(Depth CR14)	2.00	2.00	2.00	2.00	2.00	3.00	2.00	2.00	2.00
Min(Depth CR)	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00
Max(Depth CR14)	10.00	12.00	9.00	9.00	8.00	11.00	10.00	11.00	11.00
Max(Depth CR)	10.00	12.00	9.00	9.00	8.00	11.00	10.00	11.00	11.00
σ(Depth CR14)	1.48	1.62	1.23	1.24	1.15	1.90	1.28	1.33	2.03
σ(Depth CR)	1.46	2.21	1.33	1.24	1.23	1.97	1.14	1.29	1.54
ρ(Depth CR14, Depth CR)	0.86	0.79	0.87	0.62	0.76	0.94	0.77	0.86	0.63

A.2 Figures

Figure A.1: Aggregate outstanding notionals by country

Panel A: Net *Panel B: Gross*

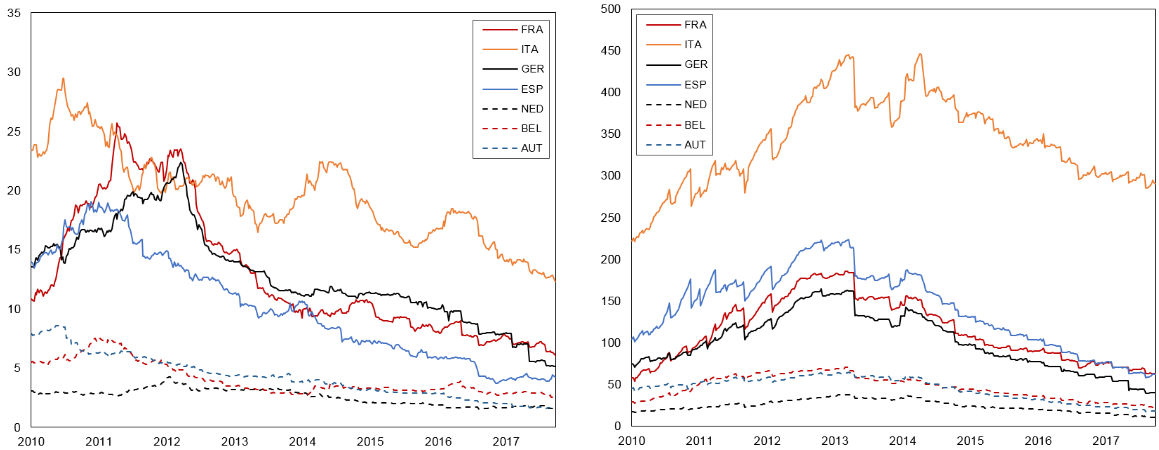

Outstanding notionals in $bn from swapsinfo.org.

Figure A.2: CDS spreads (in bps) under 2003- and 2014 ISDA definitions (control group)

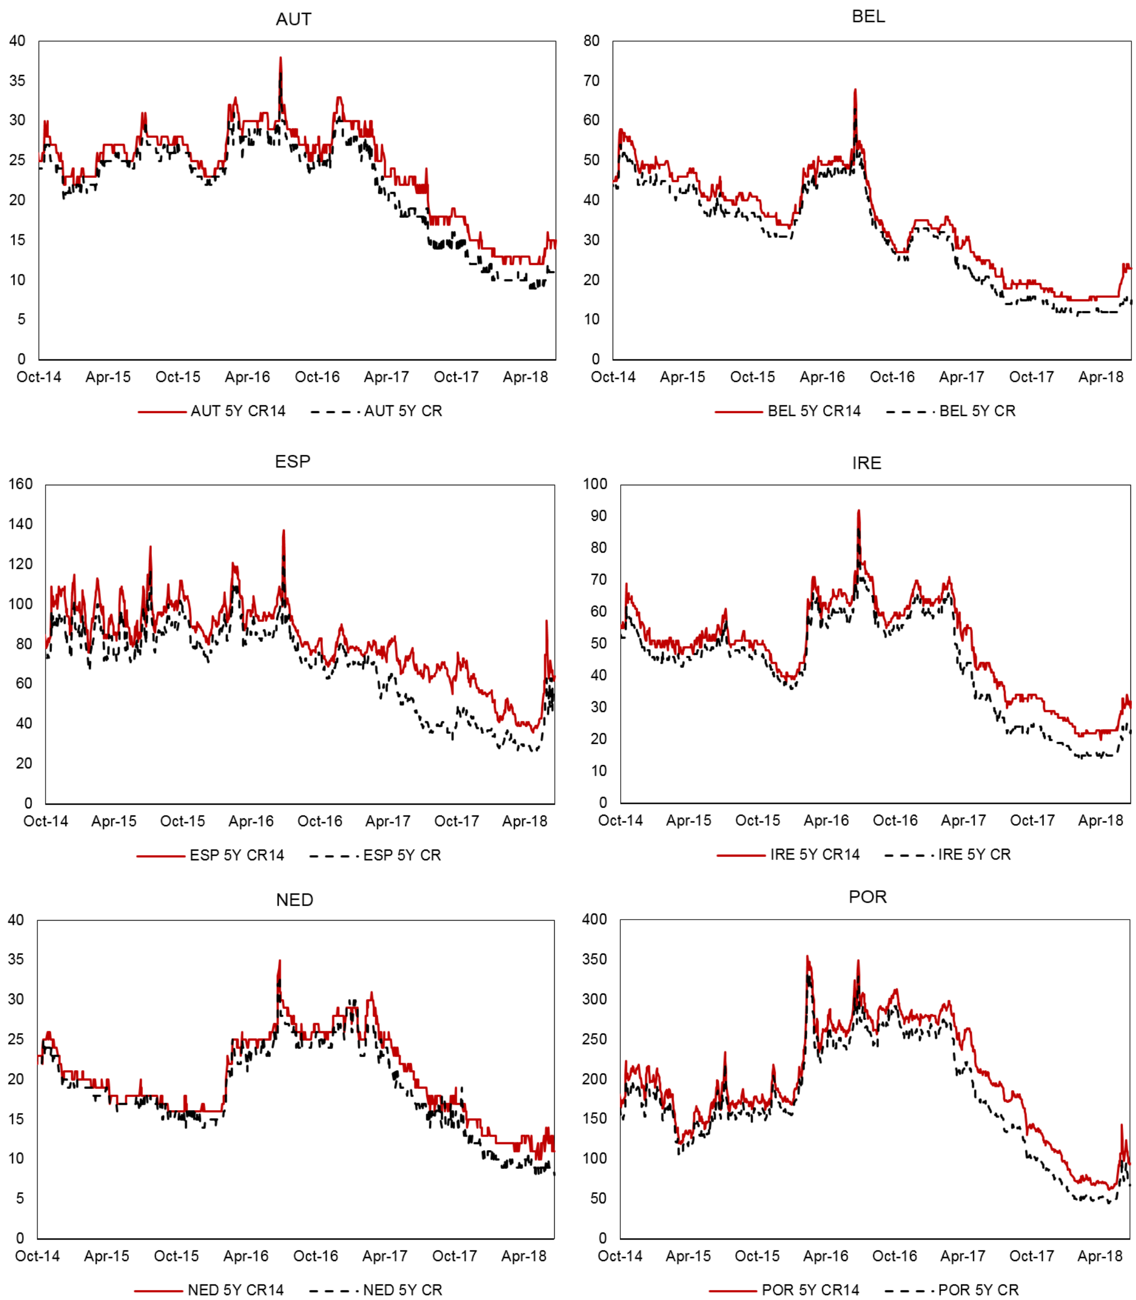

Figure A.3: CDS spreads (in bps) under 2003- and 2014 ISDA definitions (treatment group)

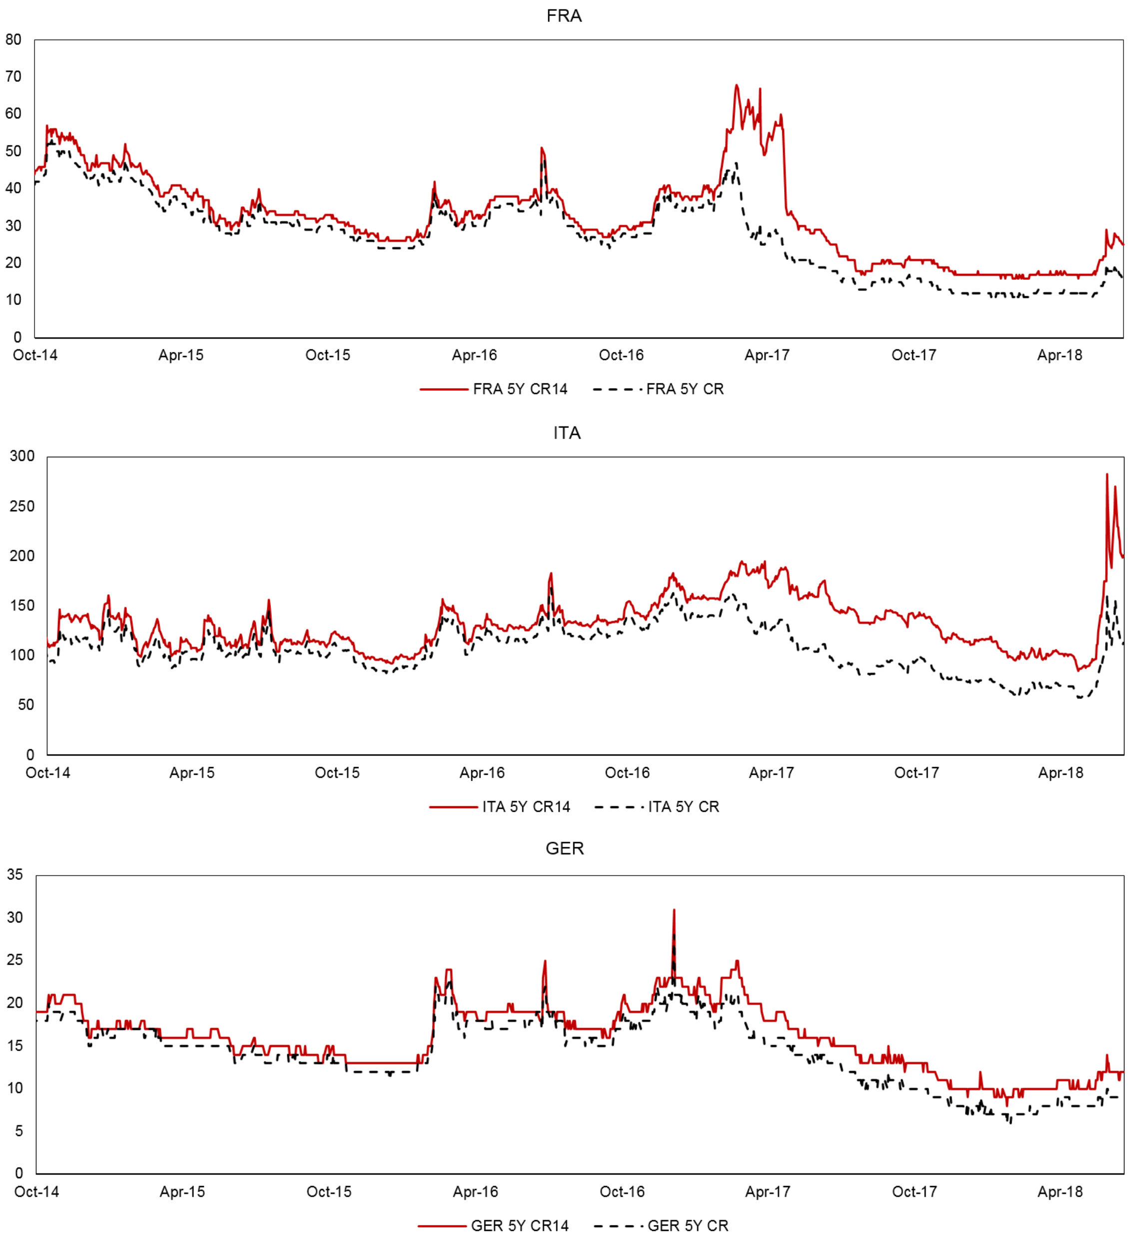

Figure A.4: Synthetic control

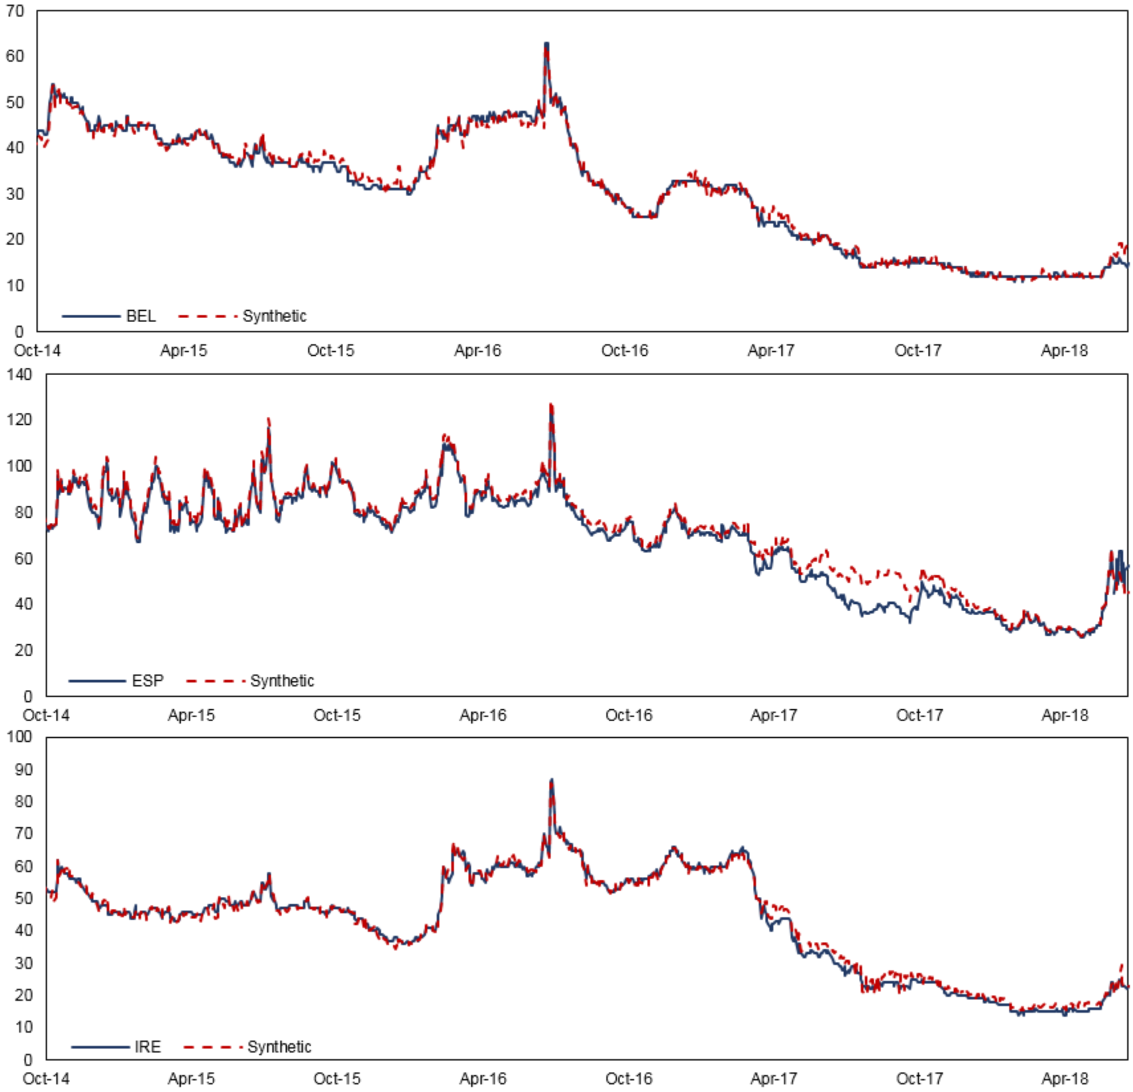

Observed (solid, blue) versus synthetically constructed (dashed, red) CR spread for Belgium (top), Spain (middle), and Ireland (bottom).

Figure A.5: Redenomination spreads for France and Germany (bottom)

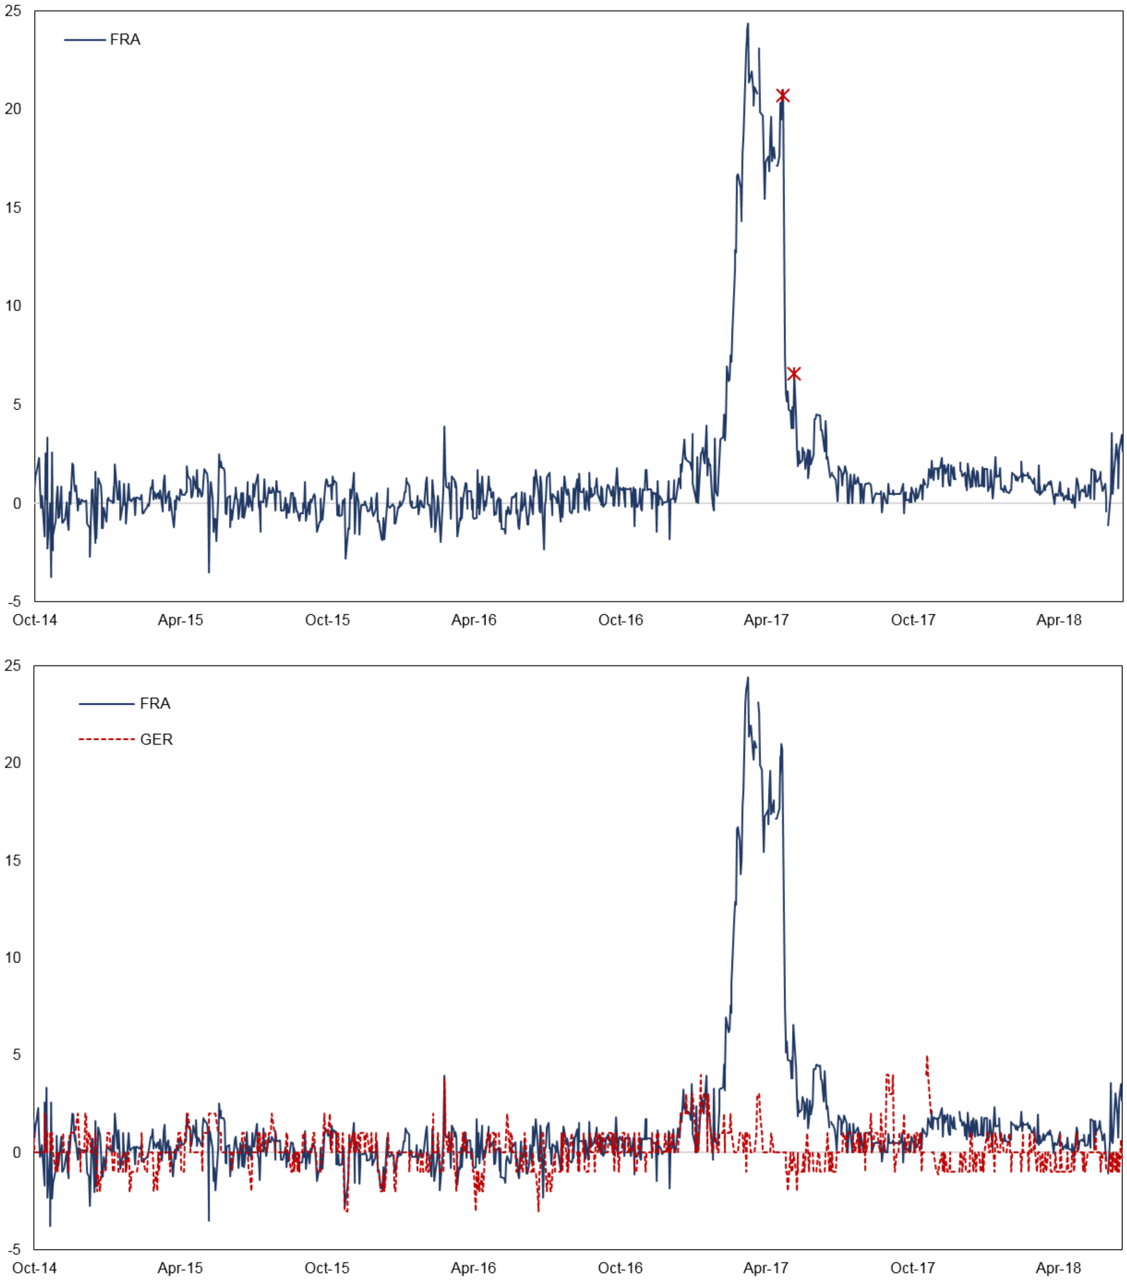

Asterisks denote major plebiscites: 1st and 2nd round of the French presidential elections on April 23rd and May 7th, 2017. In each case, the asterisk marks the observation for the Friday preceding the Sunday plebiscite.

Figure A.6: RS_{FRA} (LHS) and polls (RHS)

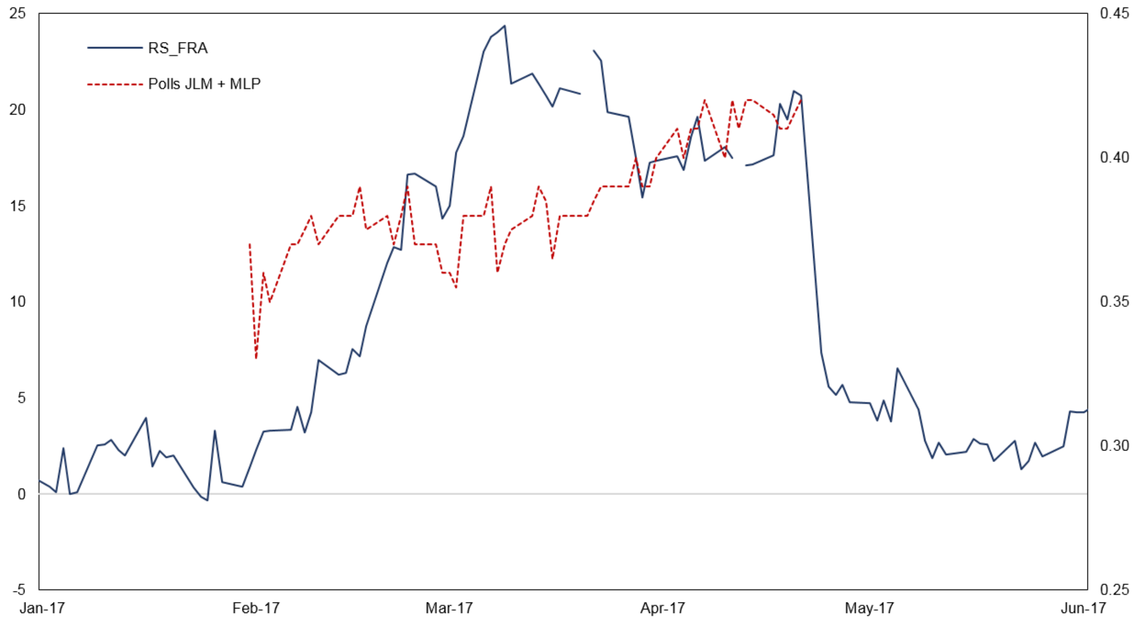

Combined polling share for J.-L. Mélenchon & M. Le Pen.

Figure A.7: Redenomination spread for Italy

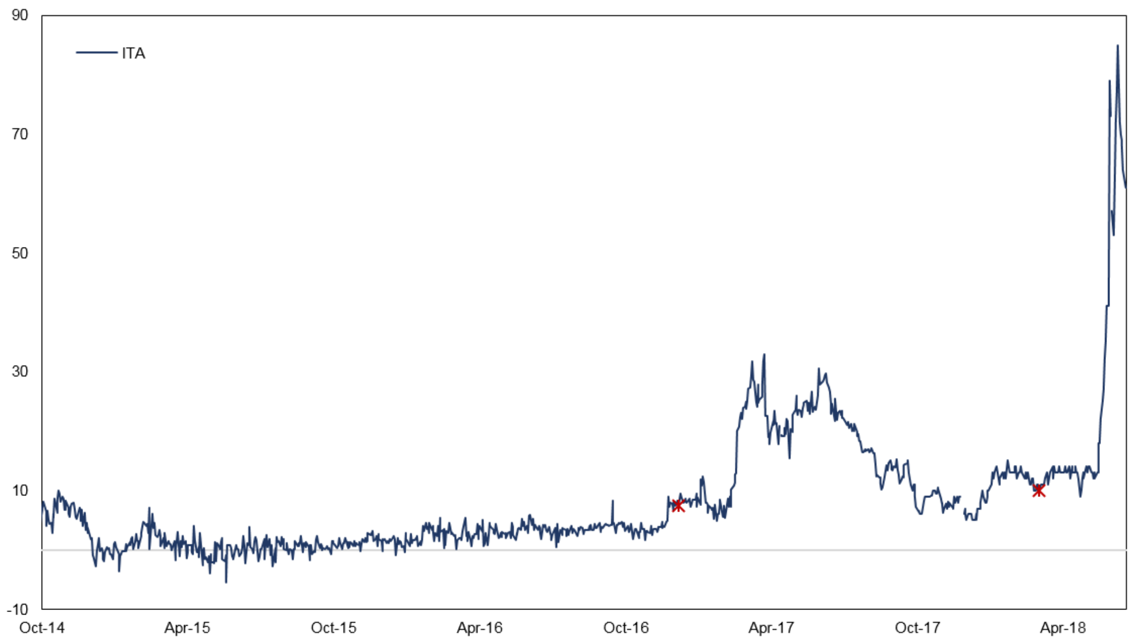

Asterisks denote major plebiscites: the constitutional referendum in Italy, held on December 4th, 2016, as well as the general elections on March 4th, 2018.

Figure A.8: Redenomination spreads for France and Italy (fraction of the CR14 spread)

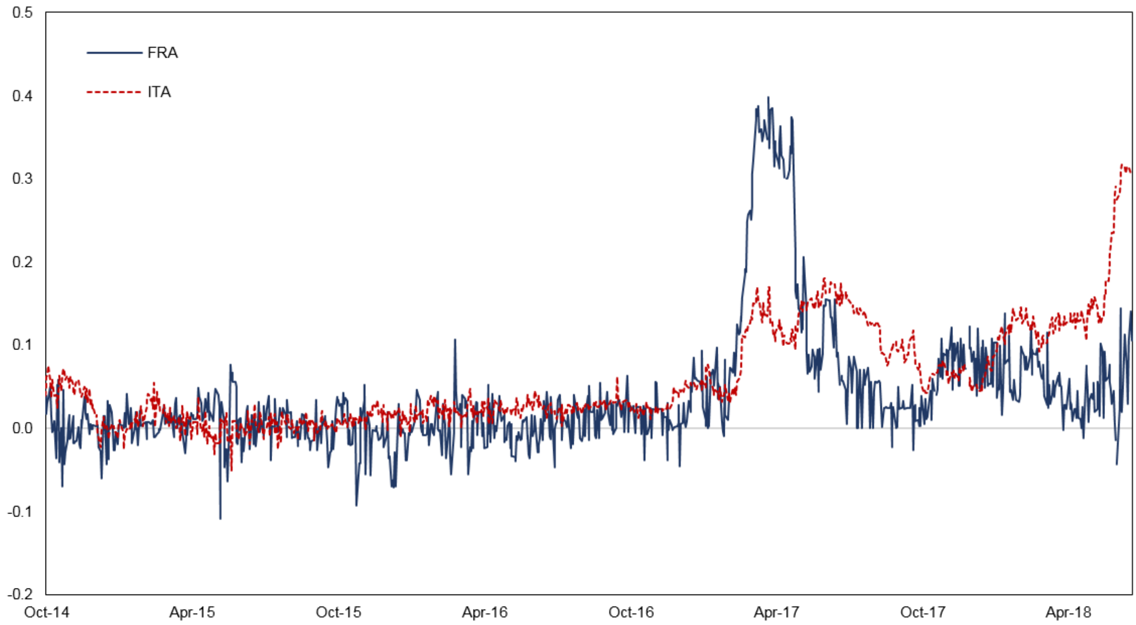

Figure A.9: Slope coefficients from Regression (1.2)

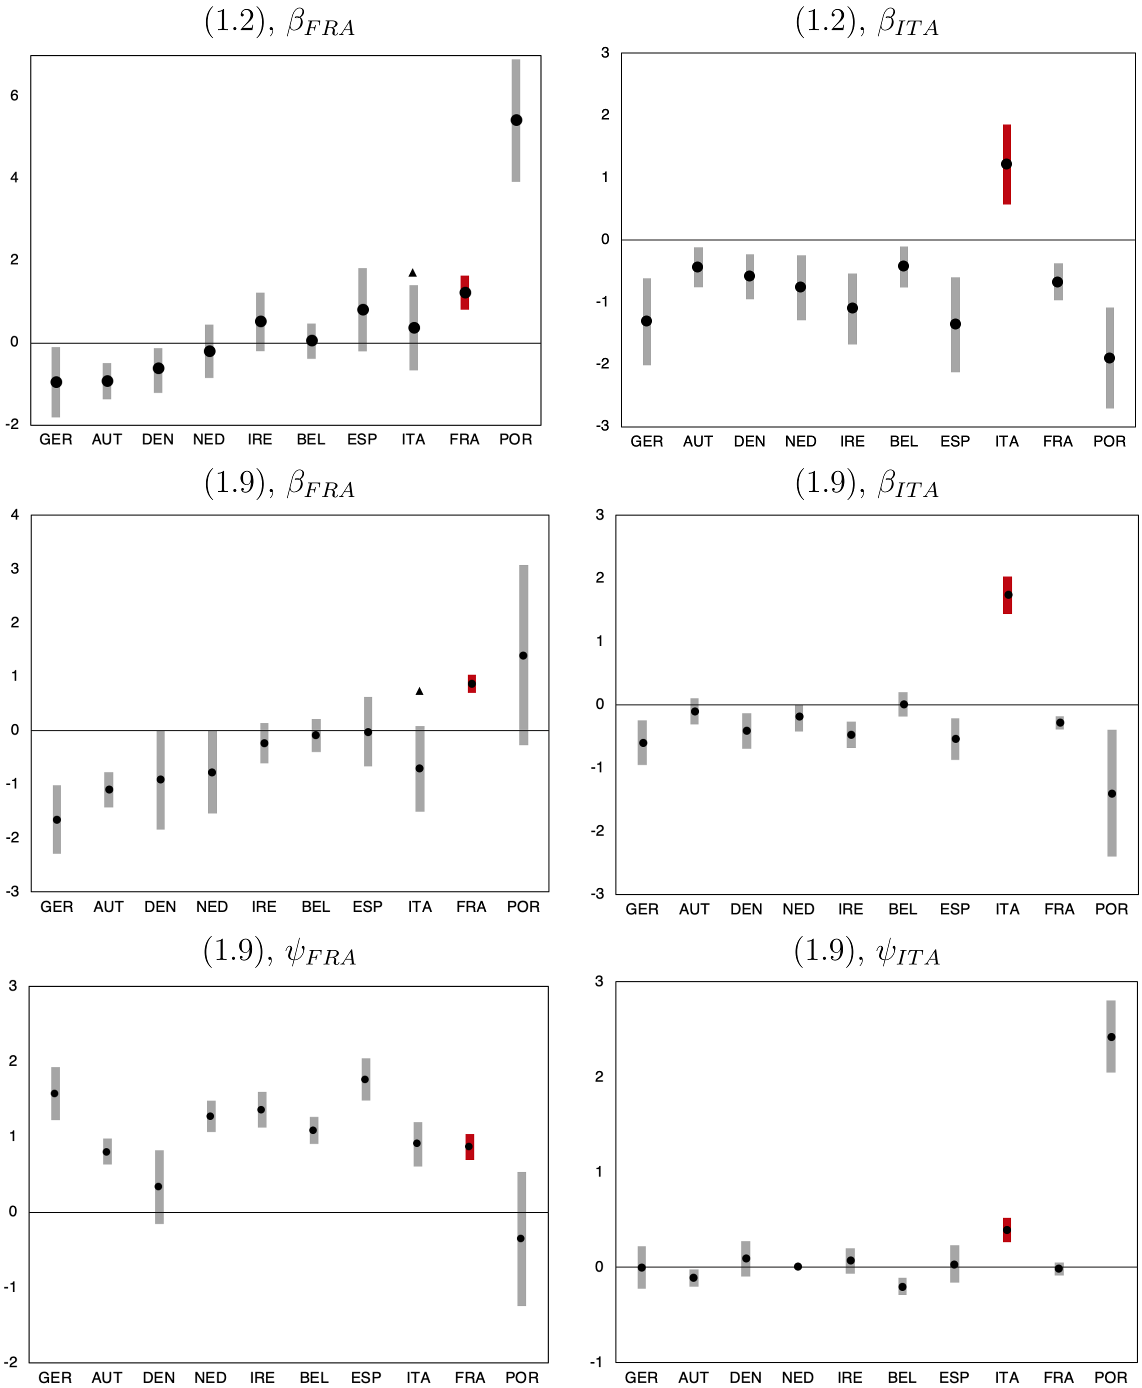

Coefficients for the regression of five-year sovereign yields on French (left) and Italian (right) redenomination spreads and 95% confidence bars (Newey–West). Middle panels: coefficients from Regression (1.9), which controls for French and Italian credit risk. Lower panels: coefficients on credit risk. The triangular marker in the left-hand side panels indicates the β_{FRA} estimate for Italian yields, once RS_{ITA} is dropped from the regressors.

Figure A.10: Slope coefficients of German and US government bond yields on *RS*

Panel A: β_{FRA} *Panel B:* β_{ITA}

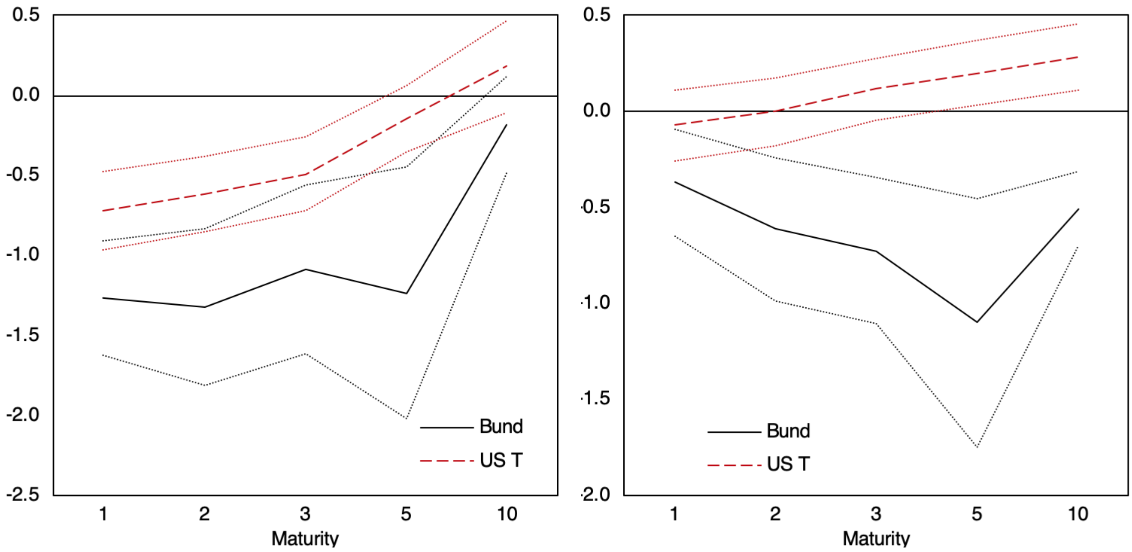

Slope coefficients from Table A.5.

Figure A.11: Slope coefficients from Regression (A.4)

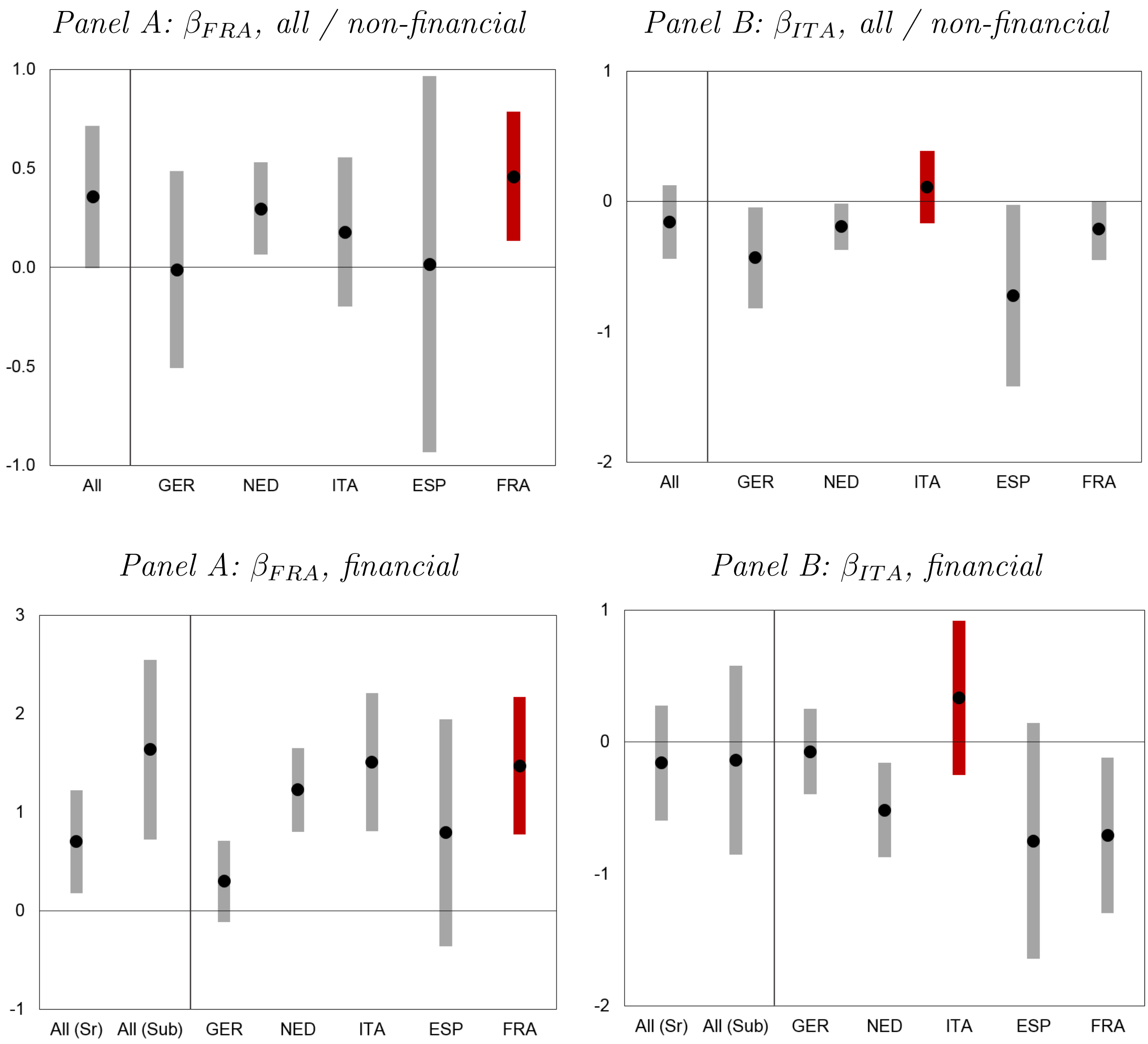

Coefficients from regression of Corporate five-year CDS spreads on French (left) and Italian (right) redenomination spreads with 95% confidence bars (Newey–West).

Figure A.12: Slope coefficients from Regressions (1.2) and (1.10)

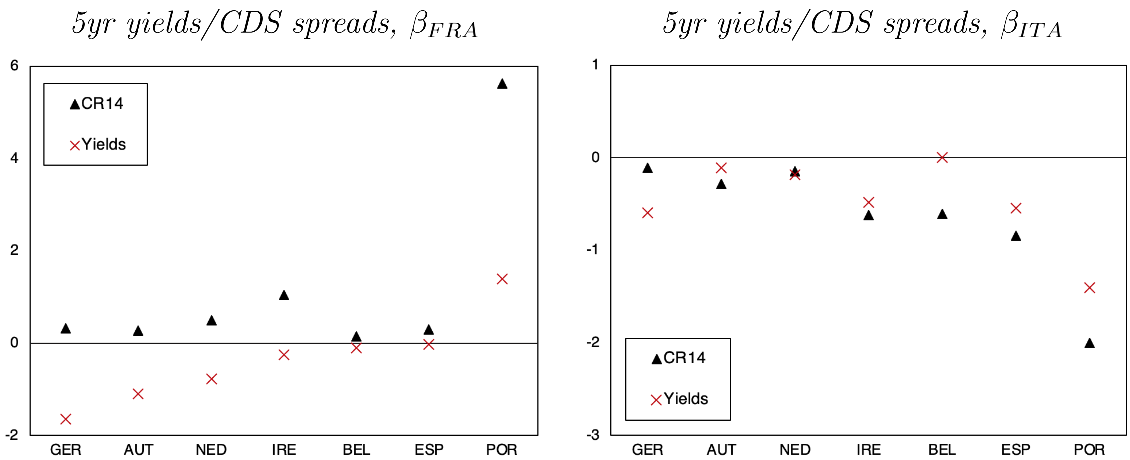

Coefficients from regression of, respectively, five-year Eurozone sovereign bond yields and 5-year CR14 CDS spreads on French (left) and Italian (right) redenomination risk. I omit confidence bars in the interest of readability.

Figure A.13: Comparative statics

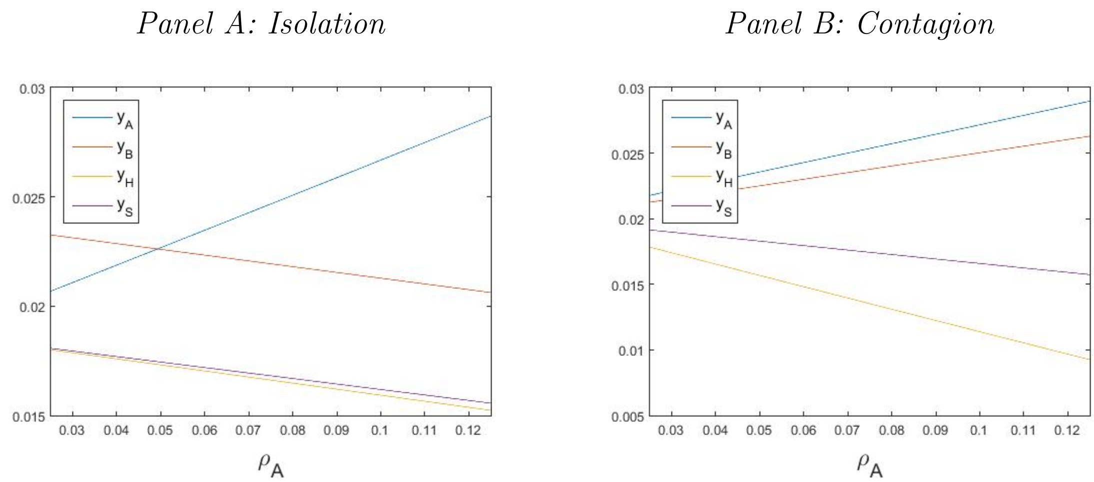

Parameters: $d^a = d^b = 0.02$, $\rho_B = 0.05$, $\delta_A = 0.1$, $\delta_B = 0.08$, $\sigma_A = \sigma_B = \sigma_H = \sigma_S = 0.51$.

Comparative statics of risky bond investments by banks a and b, and (net) bond yields with respect to redenomination probability in country A.

Figure A.14: Redenomination risk and macro fundamentals

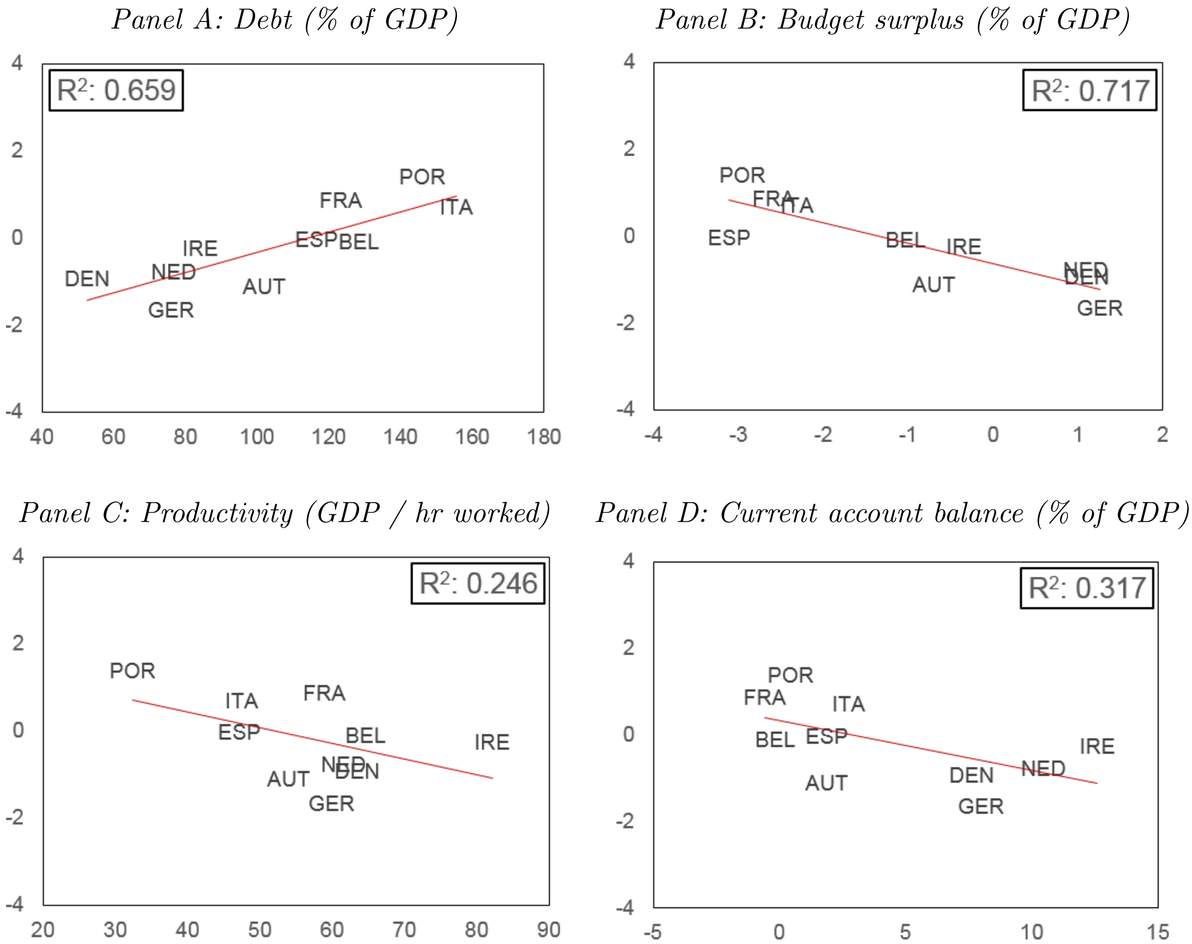

Regression coefficients $\beta_{FRA,j}$ from regression (1.9) (horizontal axis) versus fundamental variables (vertical axis, from OECD) by country, and univariate R^2.

Figure A.15: Cost of redenomination risk

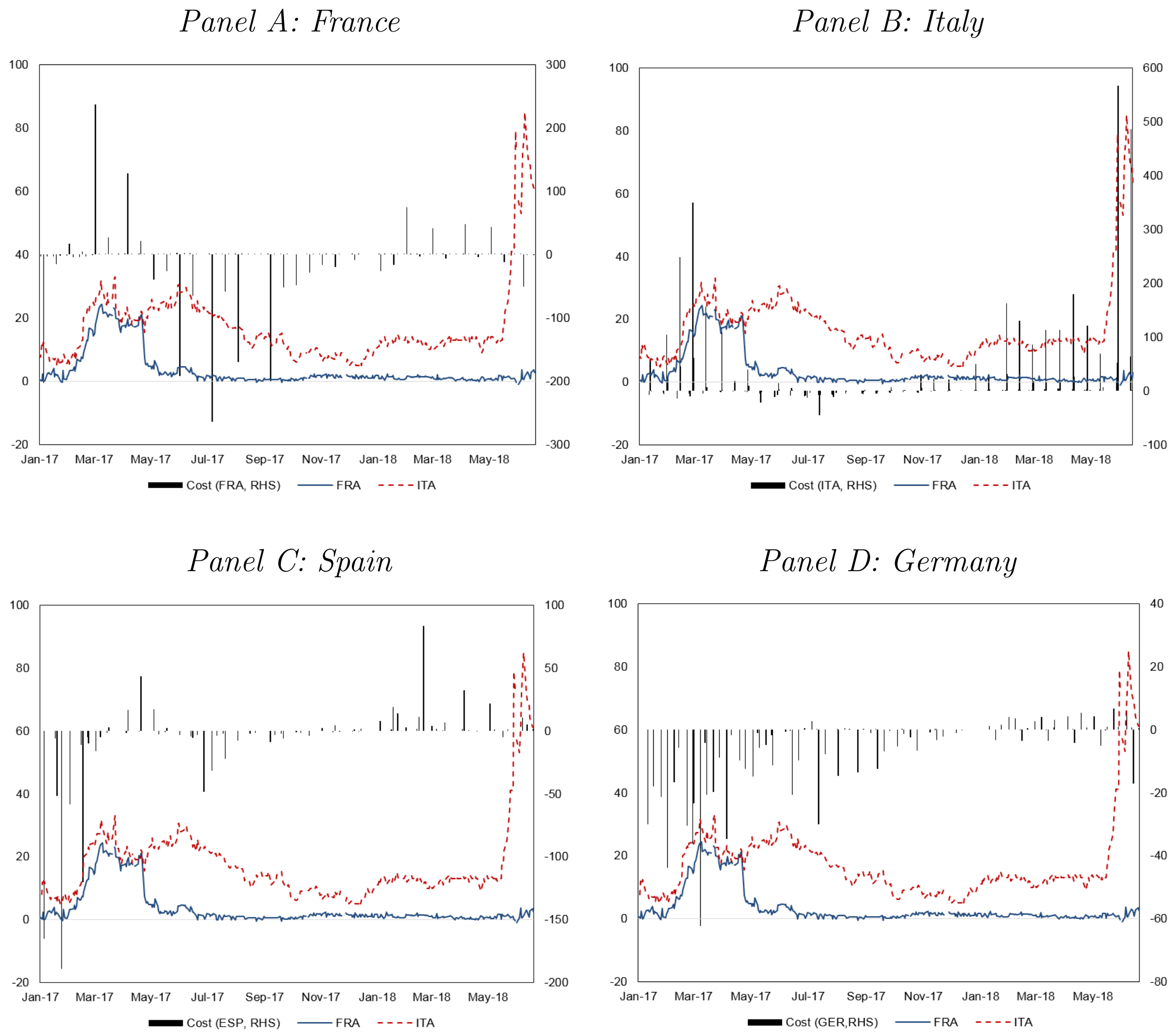

Additional funding cost attributable to French and Italian redenomination risk on debt issued between January 2017 and June 2018. Differential interest rates on the issuance by country j on date t are computed as $c_{j,t} = \hat{\beta}_{j,t}^{FRA} RS_{FRA,t} + \hat{\beta}_{j,t}^{ITA} RS_{ITA,t}$, using 250-trading-day rolling windows up to date t. Differential interest rates are then multiplied by the observed issuance volume, $v_{j,t}$, capitalized with annuity factor $a(T, y_{j,t,T})$ as $C_{j,t} = c_{j,t} \cdot v_{j,t} \cdot a(T, y_{j,t,T})$, and plotted on the RHS axes, in €m.

A.3 Proofs

In the interest of parsimony, I omit bulky closed-form expressions for equilibrium yields and comparative statics w.r.t. ρ_A. Instead, Figure A.13 provides a visual exposition of the different spillover effects in the isolation and contagion cases. The equilibrium objects, i.e., four bond yields and six portfolio weights, are determined by the four market clearing conditions and six first-order conditions w.r.t. bond investments:

$$\begin{aligned}
\sigma_J =&(w_J^a + w_J^b) \cdot y_J && \text{for } J = \{A, B, S, H\} \\
0 =& \left(\frac{p_1}{e_1^i} + \frac{p_3}{e_3^i}\right)(y_A - y_S) + \left(\frac{p_2}{e_2^i} + \frac{p_4}{e_4^i}\right)(y_A(1-\delta^A) - y_S) && \text{for } i = \{a, b\} \\
0 =& \left(\frac{p_1}{e_1^i} + \frac{p_2}{e_2^i}\right)(y_B - y_S) + \left(\frac{p_3}{e_3^i} + \frac{p_4}{e_4^i}\right)(y_B(1-\delta^B) - y_S) && \text{for } i = \{a, b\} \\
0 =& \left(\frac{p_1}{e_1^i} + \frac{p_2}{e_2^i} + \frac{p_3}{e_3^i}\right)(y_H - y_S) + \frac{p_4}{e_4^i}(y_H(1-\delta^H) - y_S) && \text{for } i = \{a, b\}
\end{aligned}$$

Due to market completeness, the equilibrium can be determined alternatively based on the prices for the four Arrow-Debreu securities. I go on to prove the two remaining effects, namely home bias in redenominatable sovereign bond holdings and the sub-zero lower bound of redenominatable haven bond yields.

Home bias: $w_A^a > w_A^b$ and $w_B^a < w_B^b$.

Proof. With four linearly independent assets, and four states of the world, markets are complete. Due to market completeness, marginal utilities (and hence equity values) are equalized state-by-state across agents in equilibrium: $u'(e_s^a) = u'(e_s^b) \Leftrightarrow e_s^a = e_s^b \; \forall \, s$.

$$\begin{aligned}
e_1^a = e_1^b \Leftrightarrow & \, w_A^a y_A + w_B^a y_B + w_S^a y_S + w_H^a y_H - d^a \\
=& w_A^b y_A + w_B^b y_B + w_S^b y_S + w_H^b y_H - d^b && \text{(A.5)} \\
e_2^a = e_2^b \Leftrightarrow & \, w_A^a y_A(1-\delta_A) + w_B^a y_B + w_S^a y_S + w_H^a y_H - d^a(1-\delta_A) \\
=& w_A^b y_A(1-\delta_A) + w_B^b y_B + w_S^b y_S + w_H^b y_H - d^b && \text{(A.6)} \\
e_3^a = e_3^b \Leftrightarrow & \, w_A^a y_A + w_B^a y_B(1-\delta_B) + w_S^a y_S + w_H^a y_H - d^a \\
=& w_A^b y_A + w_B^b y_B(1-\delta_B) + w_S^b y_S + w_H^b y_H - d^b(1-\delta_B) && \text{(A.7)} \\
e_4^a = e_4^b \Leftrightarrow & \, w_A^a y_A(1-\delta_A) + w_B^a y_B(1-\delta_B) + w_S^a y_S + w_H^a y_H(1-\delta_H) - d^a(1-\delta_A) \\
=& w_A^b y_A(1-\delta_A) + w_B^b y_B(1-\delta_B) + w_S^b y_S + w_H^b y_H(1-\delta_H) - d^b(1-\delta_B) && \text{(A.8)}
\end{aligned}$$

Combining (A.5)–(A.7) yields $w_A^a - w_A^b = d^a/y_A > 0$ and $w_B^b - w_B^a = d^b/y_B > 0$. Home bias, i.e., the difference between risky bond holdings by the domestic and the foreign bank, is positive and proportional to the domestic bank's redenominatable deposits. □

Note that this proof does not cover the extreme case of perfectly correlated redenominations, where states 2 and 3 have probability 0. In this case, bond payoffs A, B, and H are no longer linearly independent and the bond holdings are indeterminate. Assigning $\varepsilon > 0$ probability to states 2 and 3 restores the proof.

Sub-zero lower-bound: $y_H - y_S < 0$.

Proof. By each bank's Euler equation, the price of an asset with payoff X is given by $\mathbb{E}(y_S^{-1} u'(e^i) X)$. The Arrow-Debreu security that pays off in state (4) consists of $-\delta_H^{-1}$ units of bond H, and δ_H^{-1} units of bond S. Bond prices are y_H^{-1} and y_S^{-1} , respectively, and therefore

$$\begin{aligned} \delta_H^{-1}(y_S^{-1} - y_H^{-1}) &= \frac{1}{y_S} \cdot \frac{p_4}{e_4^i} \\ \Rightarrow \delta_H^{-1}(1 - y_S/y_H) &= \frac{p_4}{e_4^i} \end{aligned} \tag{A.9}$$

Equity is strictly positive by assumption and $p_4 \in (0,1)$. The RHS of (A.9) is therefore strictly positive, which, together with $\delta_H < 0$, implies that $y_H < y_S$. □

A.4 Supplementary Tables and Figures

Figure A.16: Redenomination risk and *Economic Policy Uncertainty*

Panel A: France

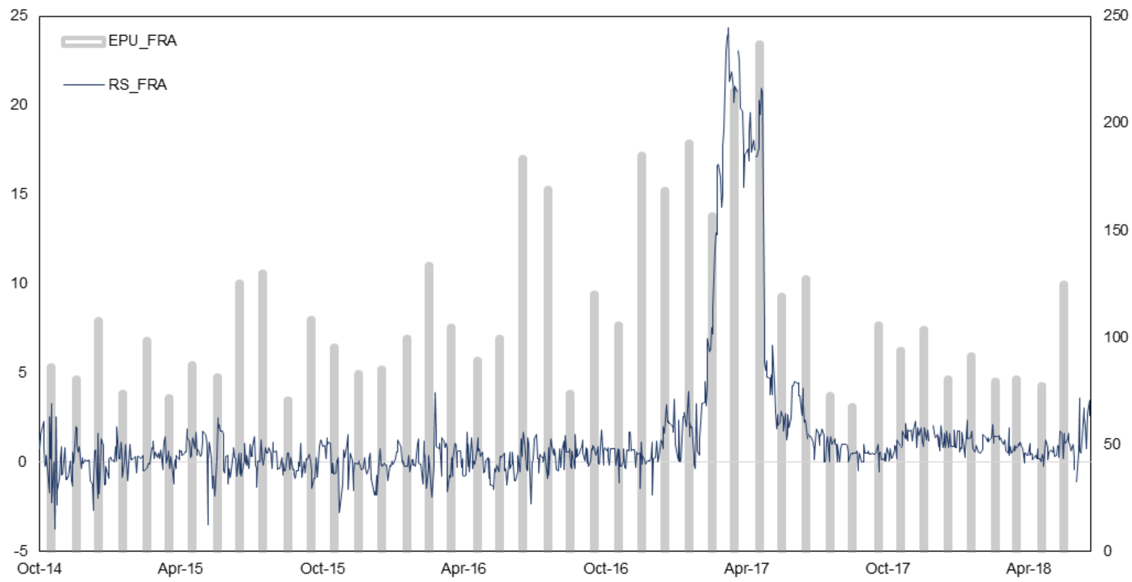

Panel B: Italy

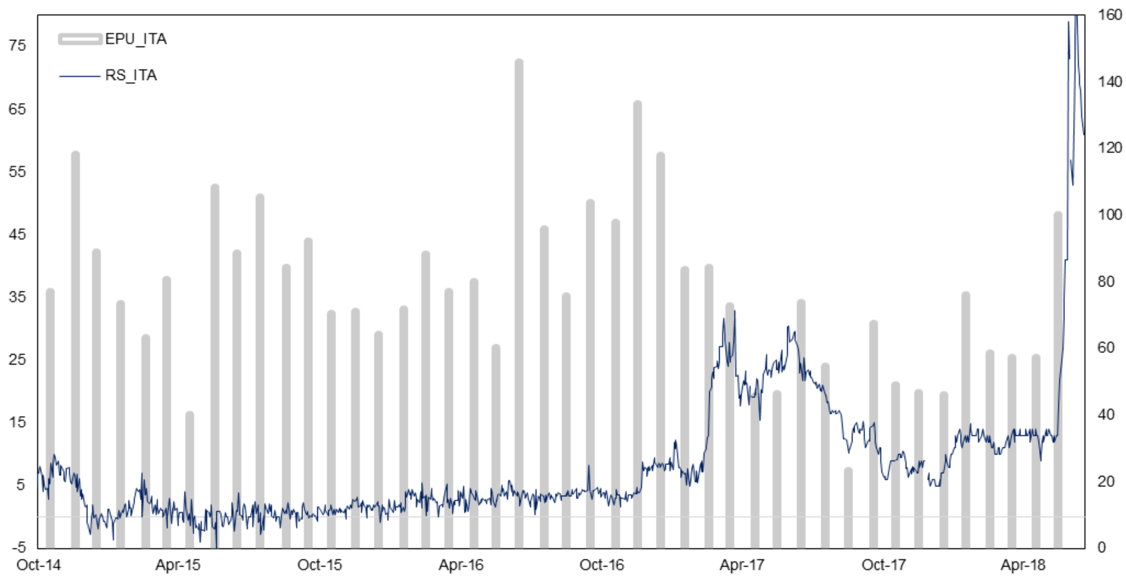

Economic policy uncertainty (RHS) obtained from (Baker et al., 2016).

B. Appendix to The Quanto Theory

B.1 Tables and Figures

Figure B.1: The time series of QRP

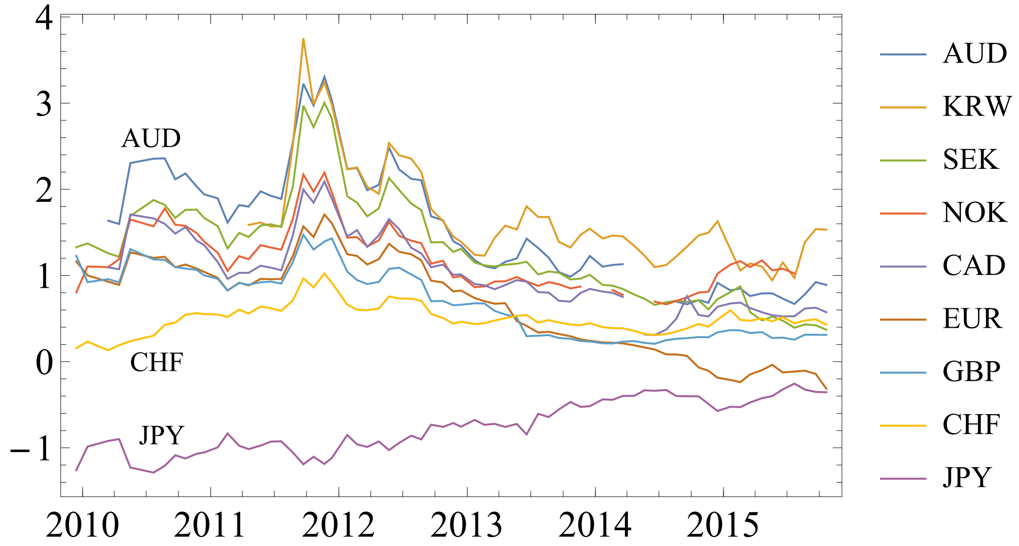

The figure drops two currencies (PLN and DKK) for which we have highly incomplete time series.

Table B.1: Summary statistics of ECA

This table reports annualized summary statistics (in %) of quanto-based expected currency appreciation (ECA).

	Mean	**Std Dev.**	**Skew**	**Kurtosis**	**Min**	**Max**	**Autocorr.**
				Expected currency appreciation, ECA			
AUD	−1.231	0.723	−0.114	−0.577	−2.550	0.450	0.864
CAD	0.327	0.526	0.909	0.494	−0.526	1.835	0.845
CHF	1.064	0.472	1.147	0.210	0.422	2.176	0.934
DKK	0.331	0.487	−0.097	−0.606	−0.587	1.172	0.762
EUR	0.587	0.398	−0.725	0.799	−0.493	1.300	0.877
GBP	0.326	0.350	−0.103	−0.517	−0.444	1.077	0.894
JPY	−0.337	0.412	0.484	−0.989	−0.978	0.555	0.953
KRW	0.706	0.724	1.455	2.922	−0.182	3.387	0.770
NOK	−0.398	0.622	0.624	0.040	−1.474	0.991	0.877
PLN	−1.340	0.892	0.759	−0.479	−2.554	0.436	0.881
SEK	0.574	0.656	−0.143	−0.340	−0.907	1.885	0.885
Average	0.056	0.569	0.382	0.087	−0.934	1.388	0.867
Pooled	0.056	0.908	−0.500	0.630	−2.554	3.387	

Table B.2: Summary statistics of IRD and QRP

This table reports annualized summary statistics (in %) of UIP forecasts (IRD, top panel), and quanto-implied risk premia (QRP, bottom).

	Mean	**Std Dev.**	**Skew**	**Kurtosis**	**Min**	**Max**	**Autocorr.**
				Interest rate differential, IRD			
AUD	−2.815	1.007	−0.104	−1.081	−4.533	−1.168	0.979
CAD	−0.712	0.353	1.121	0.204	−1.133	0.195	0.890
CHF	0.560	0.441	1.501	1.137	0.013	1.690	0.953
DKK	−0.821	0.470	0.298	−0.794	−1.596	0.005	0.915
EUR	−0.056	0.622	−0.282	−0.509	−1.377	0.983	0.977
GBP	−0.352	0.223	−0.098	−0.745	−0.865	0.082	0.925
JPY	0.410	0.206	0.476	−1.229	0.133	0.809	0.909
KRW	−0.973	0.443	0.587	−1.017	−1.614	−0.116	0.877
NOK	−1.596	0.690	0.587	−0.286	−2.798	−0.107	0.955
PLN	−3.422	1.030	2.010	2.733	−4.215	−0.806	0.967
SEK	−0.715	0.905	0.430	−0.421	−2.354	1.105	0.981
Average	−0.954	0.581	0.593	−0.183	−1.849	0.243	0.939
Pooled	−0.954	1.265	−0.952	0.657	−4.533	1.690	
				Quanto-implied risk premium, QRP			
AUD	1.584	0.692	0.546	−0.454	0.666	3.306	0.941
CAD	1.039	0.441	0.509	−0.572	0.309	2.090	0.926
CHF	0.504	0.171	0.663	1.405	0.131	1.023	0.900
DKK	1.153	0.275	0.400	0.336	0.643	1.768	0.788
EUR	0.643	0.556	−0.104	−1.274	−0.315	1.708	0.978
GBP	0.678	0.389	0.270	−1.318	0.207	1.472	0.959
JPY	−0.746	0.295	−0.033	−1.287	−1.287	−0.255	0.945
KRW	1.679	0.589	1.605	2.582	0.944	3.752	0.859
NOK	1.198	0.359	0.876	0.462	0.665	2.194	0.890
PLN	2.083	0.650	0.814	0.026	1.194	3.509	0.868
SEK	1.289	0.616	0.801	0.620	0.371	3.004	0.938
Average	1.009	0.457	0.577	0.048	0.321	2.143	0.908
Pooled	1.009	0.857	−0.107	0.658	−1.287	3.752	

Table B.3: Volatilities and correlations of ECA, IRD, and QRP

This table presents the standard deviations (in %) of, and correlations between, the interest rate differential (IRD), the quanto-implied risk premium (QRP), and expected currency appreciation (ECA), calculated from (2.14) for each currency i:

$$\begin{aligned}
\mathrm{IRD}_{i,t} &= \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} - 1 \\
\mathrm{QRP}_{i,t} &= \frac{Q_{i,t} - F_t}{R^{i}_{f,t} P_t} \\
\mathrm{ECA}_{i,t} &= \mathrm{QRP}_{i,t} + \mathrm{IRD}_{i,t}.
\end{aligned}$$

The row labelled "Time series" reports means of the currencies' time-series standard deviations and correlations. The row labelled "Cross section" reports cross-sectional standard deviations and correlations of time-averaged ECA, IRD, and QRP. The row labelled "Pooled" reports standard deviations and correlations of the pooled data. All quantities are expressed in annualized terms.

	$\sigma(ECA)$	$\sigma(IRD)$	$\sigma(QRP)$	$\rho(ECA, IRD)$	$\rho(ECA, QRP)$	$\rho(IRD, QRP)$
AUD	0.723	1.007	0.692	0.727	−0.013	−0.696
CAD	0.526	0.353	0.441	0.558	0.748	−0.134
CHF	0.472	0.441	0.171	0.932	0.355	−0.007
DKK	0.487	0.470	0.275	0.835	0.342	−0.231
EUR	0.398	0.622	0.556	0.476	0.183	−0.777
GBP	0.350	0.223	0.389	0.137	0.822	−0.451
JPY	0.412	0.206	0.295	0.738	0.882	0.333
KRW	0.724	0.443	0.589	0.582	0.792	−0.036
NOK	0.622	0.690	0.359	0.855	0.090	−0.439
PLN	0.892	1.030	0.650	0.780	0.135	−0.514
SEK	0.656	0.905	0.616	0.733	−0.013	−0.690
Time-series	0.569	0.581	0.457	0.669	0.393	−0.331
Cross-section	0.786	1.242	0.751	0.817	−0.305	−0.798
Pooled	0.908	1.265	0.857	0.736	−0.026	−0.696

Figure B.2: QRP, IRD, and ECA

(a) The relationship between QRP and IRD

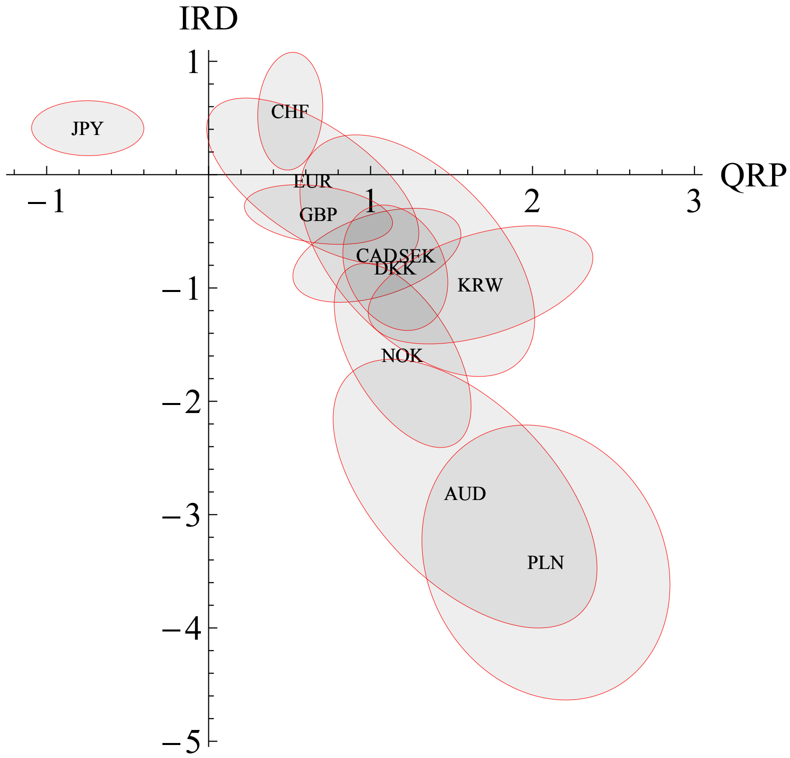

(b) The relationship between QRP and ECA

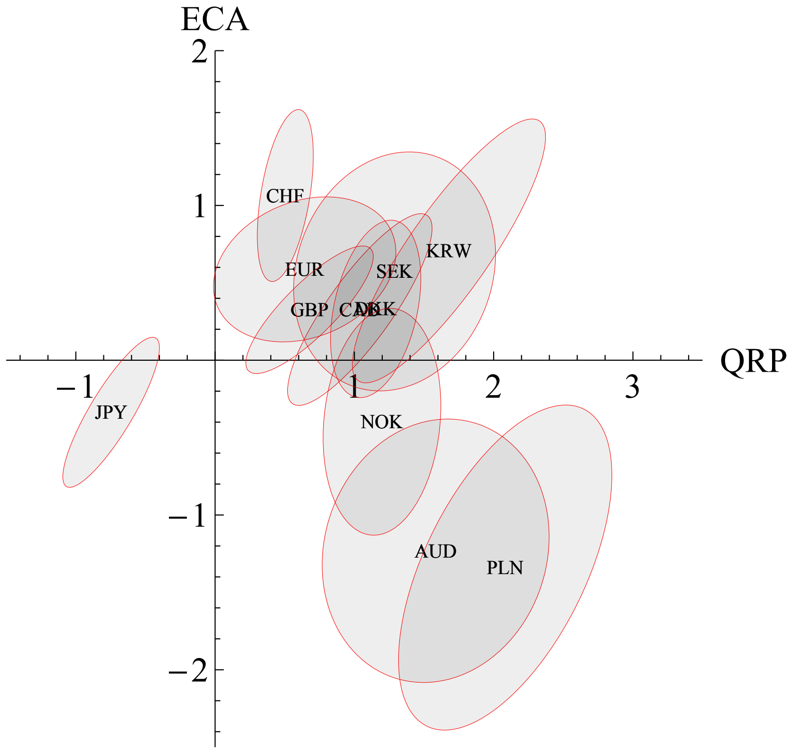

For each currency, the figures plot mean QRP and IRD (or ECA) surrounded by a confidence ellipse whose orientation reflects the time-series correlation between QRP and IRD (or ECA), and whose size reflects their volatilities. The location and orientation of the ellipses in panel (a) indicate that high interest rates are associated with high quanto-implied risk premia in the cross section and in the time series.

Figure B.3: Expected currency appreciation over a 24-month horizon (annualized)

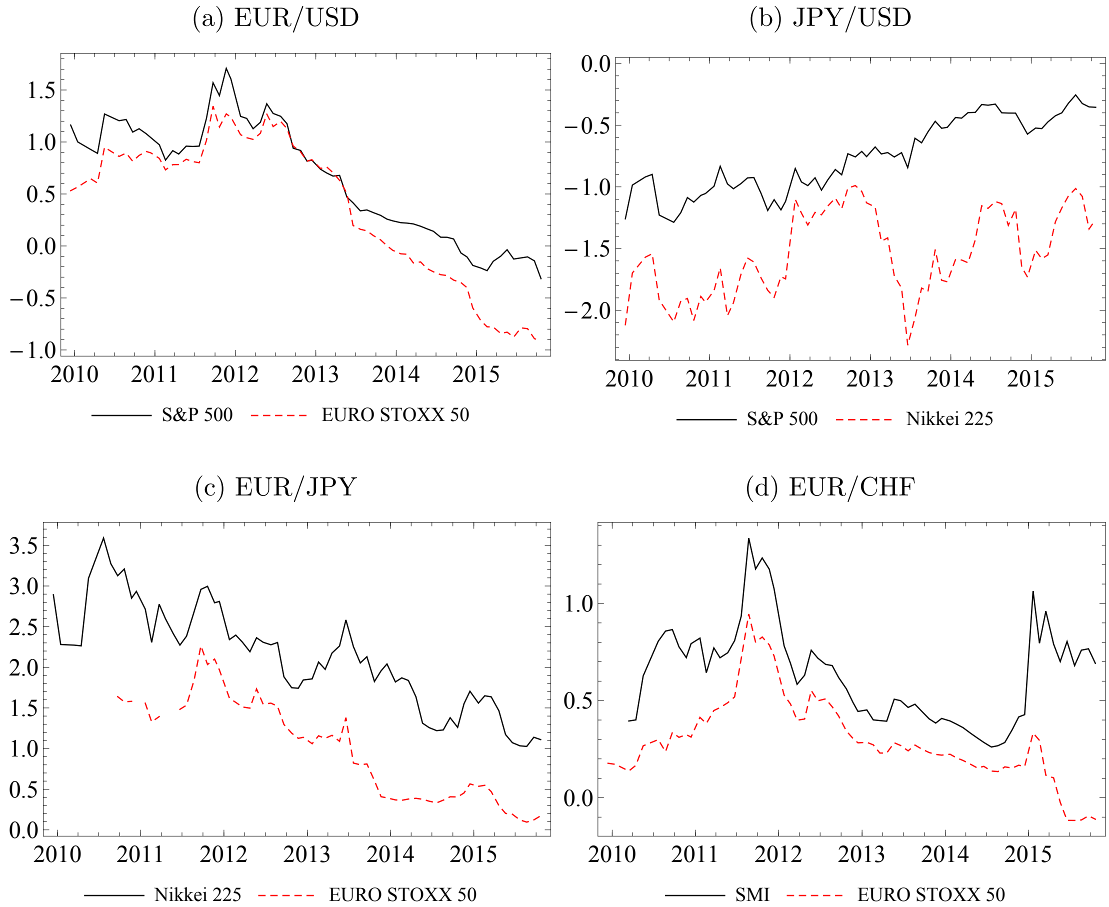

Expected currency appreciation as measured by ECA from equation (2.14), for the EUR/USD, JPY/USD, EUR/JPY, and EUR/CHF currency pairs. Each panel plots ECA for the respective currency pair from the two national perspectives, using quanto contracts on the respective domestic index denominated in the respective foreign currency. The solid line plots ECA as perceived by a log investor fully invested in the S&P 500 (top two panels), Nikkei 225 (bottom left panel), and SMI (bottom right panel), respectively. The dashed line plots the negative of ECA for the same currency pair (inverting the exchange rate) from the perspective of a log investor fully invested in the respective foreign equity index.

Figure B.4: Realized and expected currency excess return

(a) Realized currency excess return against QRP, computed from (2.14)

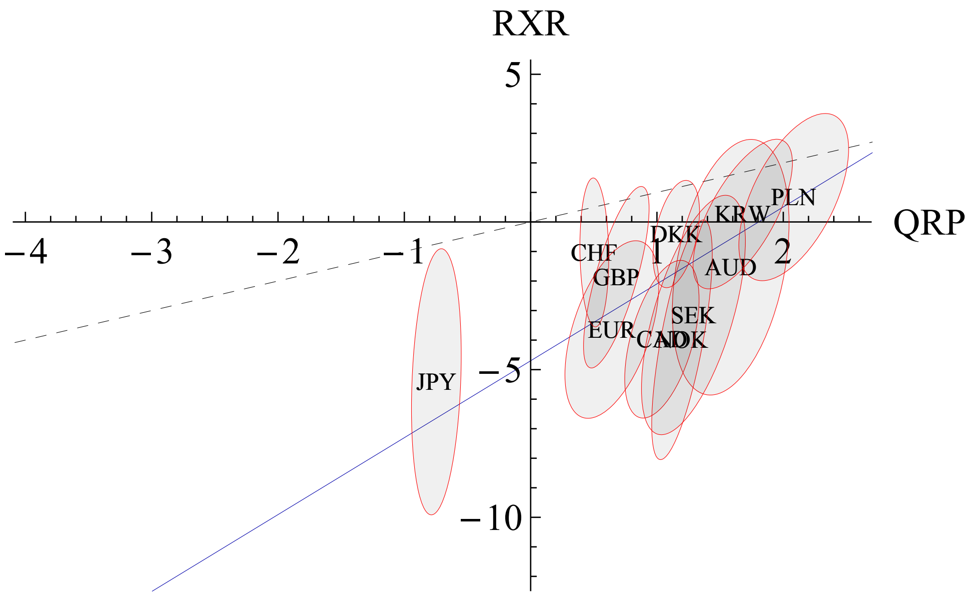

(b) Realized currency excess return against IRD

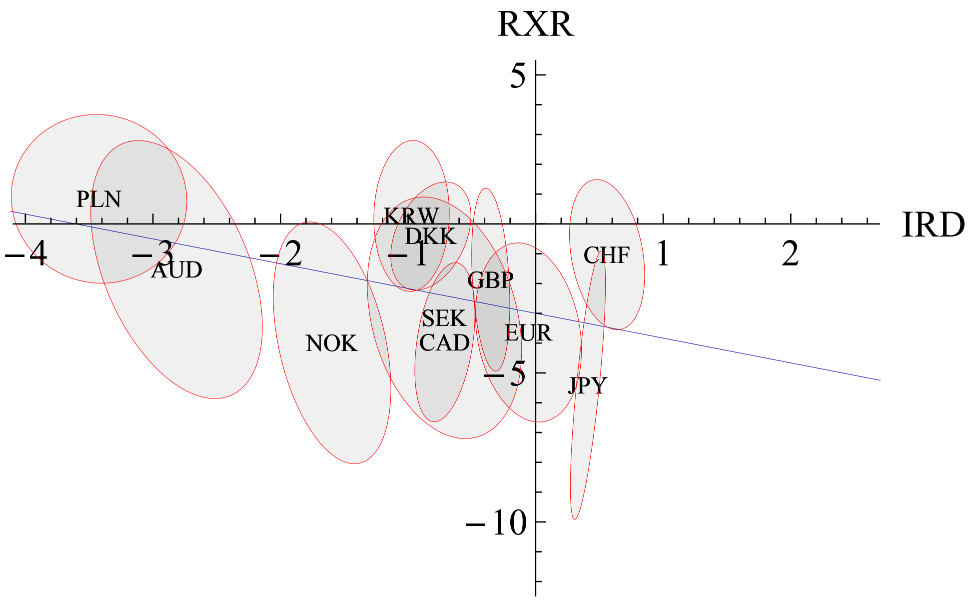

Expected currency excess return according to (a) the quanto theory and (b) UIP. The centre of each confidence ellipse represents a currency's mean expected and realized currency excess return. In population, each ellipse would contain 20% of its currency's data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse's size reflects their volatilities. Panel (a) shows a dotted 45° line for comparison.

Figure B.5: Realized and expected currency appreciation

(a) Realized currency appreciation against ECA, computed from (2.14)

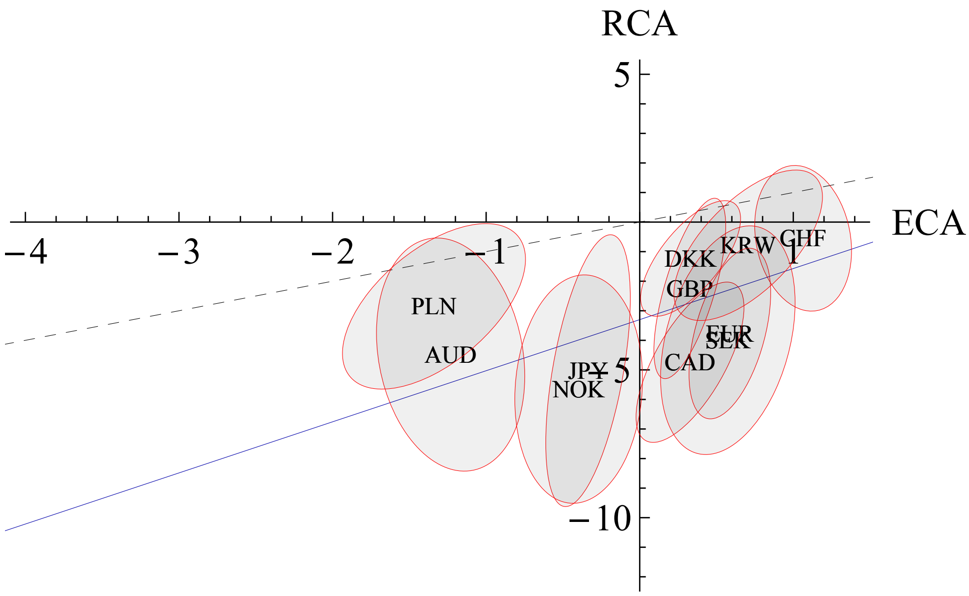

(b) Realized currency appreciation against IRD

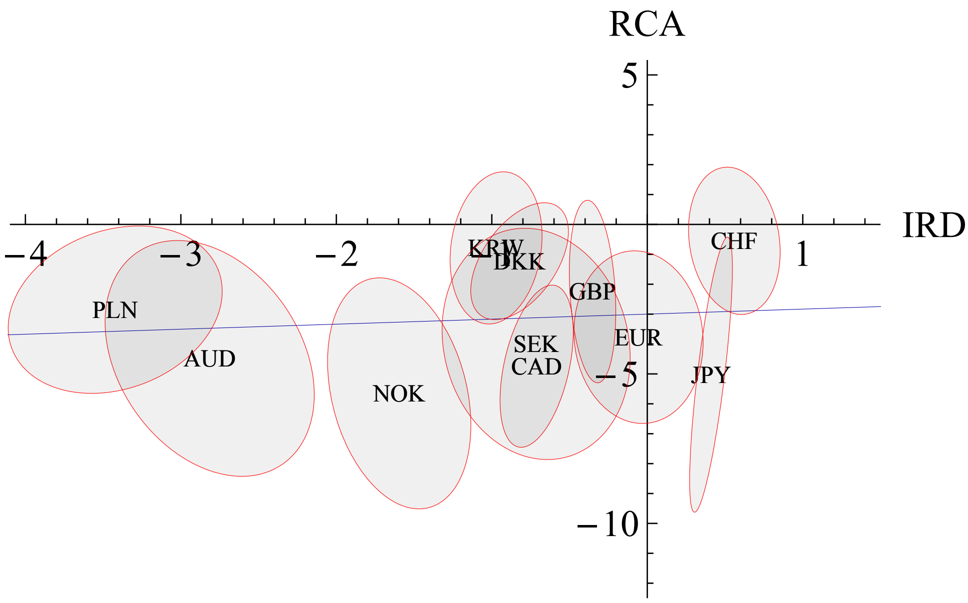

Expected currency appreciation according to (a) the quanto theory and (b) UIP. The centre of each confidence ellipse represents a currency's mean expected and realized currency appreciation. In population, each ellipse would contain 20% of its currency's data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and forecast appreciation for the given currency, while the ellipse's size reflects their volatilities. Panel (a) shows a dotted 45° line for comparison.

Table B.4: Currency excess return forecasting regressions

This table presents results from three currency excess return forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{2.19}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \tag{2.20}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{2.21}$$

Return realizations correspond to the forecasting horizon of 24 months. The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, as well as R^2 (in %).

Panel A: Pooled panel regressions			
Regression	(2.19)	(2.20)	(2.21)
α (p.a.)	−0.048	−0.047	−0.030
	(0.020)	(0.019)	(0.014)
β	3.394	2.604	
	(1.734)	(1.127)	
γ	0.769		−0.832
	(1.040)		(0.651)
R^2	19.13	17.43	3.88
Panel B: Panel regressions with currency fixed effects			
β	5.456	4.995	
	(2.046)	(1.565)	
γ	0.717		−1.363
	(1.411)		(1.001)
R^2	22.60	22.03	2.77

Table B.5: Currency forecasting regressions

This table presents results from three currency forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{2.22}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \tag{2.23}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{2.24}$$

Return realizations correspond to the forecasting horizon of 24 months. The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, as well as R^2 (in %).

Panel A: Pooled panel regressions			
Regression	(2.22)	(2.23)	(2.24)
α (p.a.)	−0.048	−0.045	−0.030
	(0.020)	(0.019)	(0.014)
β	3.394	1.576	
	(1.726)	(1.172)	
γ	1.769		0.168
	(1.045)		(0.651)
R^2	16.01	6.63	0.16
Panel B: Panel regressions with currency fixed effects			
β	5.456	4.352	
	(2.047)	(1.682)	
γ	1.717		−0.363
	(1.414)		(1.007)
R^2	20.56	17.16	0.20

Table B.6: Realized covariance regressions

This table presents results of regressions using the lagged realized covariance of exchange rate movements with the negative reciprocal of the S&P 500 return (RPCL) as a proxy for the currency beta:

$$\mathrm{RPCL}_{i,t} = R^{\$}_{f,t} \left(\sum_{t-h}^{t} \left[\frac{e_{i,s}}{e_{i,s-1}} \left(-\frac{1}{R_s} \right) \right] - \frac{1}{h} \sum_{t-h}^{t} \left(-\frac{1}{R_s} \right) \sum_{t-h}^{t} \frac{e_{i,s}}{e_{i,s-1}} \right),$$

where the summation is over daily returns on trading days s preceding t over a time-frame corresponding to our forecasting horizon, h, so that $\mathrm{RPCL}_{i,t}$ is observable at time t. We also define a realized covariance measure $\mathrm{RPC}_{i,t}$ that is analogous to the above definition except that the summation is over trading days *following* t over the appropriate time-frame (so that it is not observable until time $t+h$). We test whether risk-neutral covariance forecasts realized covariance, in a univariate regression as well as in the presence of lagged realized covariance and IRD as competing predictors. Lastly, we denote by $\widehat{\mathrm{RPC}}_{i,t}$ the optimal forecast of $\mathrm{RPC}_{i,t}$ from regression (B.4) and test whether it forecasts excess returns.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \gamma\,\mathrm{RPCL}_{i,t} + \varepsilon_{i,t+1} \tag{B.1}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{RPCL}_{i,t} + \varepsilon_{i,t+1} \tag{B.2}$$

$$\mathrm{RPC}_{i,t} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \tag{B.3}$$

$$\mathrm{RPC}_{i,t} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{RPCL}_{i,t} + \delta\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \tag{B.4}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\widehat{\mathrm{RPC}}_{i,t} + \varepsilon_{i,t+1} \tag{B.5}$$

Return realizations correspond to the forecasting horizon of 24 months. We report coefficient estimates for each regression, with standard errors (computed using a nonparametric block bootstrap) in brackets. See Section 2.2.6 for more details.

Panel A: Pooled panel regression					
Regression	(B.1)	(B.2)	(B.3)	(B.4)	(B.5)
α (p.a.)	-0.034	-0.047	-0.000	0.000	-0.047
	(0.017)	(0.018)	(0.001)	(0.001)	(0.018)
β		2.798	0.447	-0.026	3.096
		(1.366)	(0.158)	(0.126)	(1.639)
γ	1.307	-0.213		0.370	-1.103
	(1.111)	(1.193)		(0.123)	(3.206)
δ				-0.131	
				(0.061)	
R^2	7.37	17.52	36.56	66.44	17.94
Panel B: Panel regression with currency fixed effects					
β		4.643	0.330	-0.107	4.988
		(2.006)	(0.168)	(0.017)	(2.073)
γ	1.967	0.387		0.313	0.023
	(1.474)	(1.384)		(0.125)	(3.300)
δ				-0.237	
				(0.138)	
R^2	9.14	22.27	9.43	45.69	22.03

Figure B.6: Risk-neutral and predicted covariances of exchange rates and S&P returns

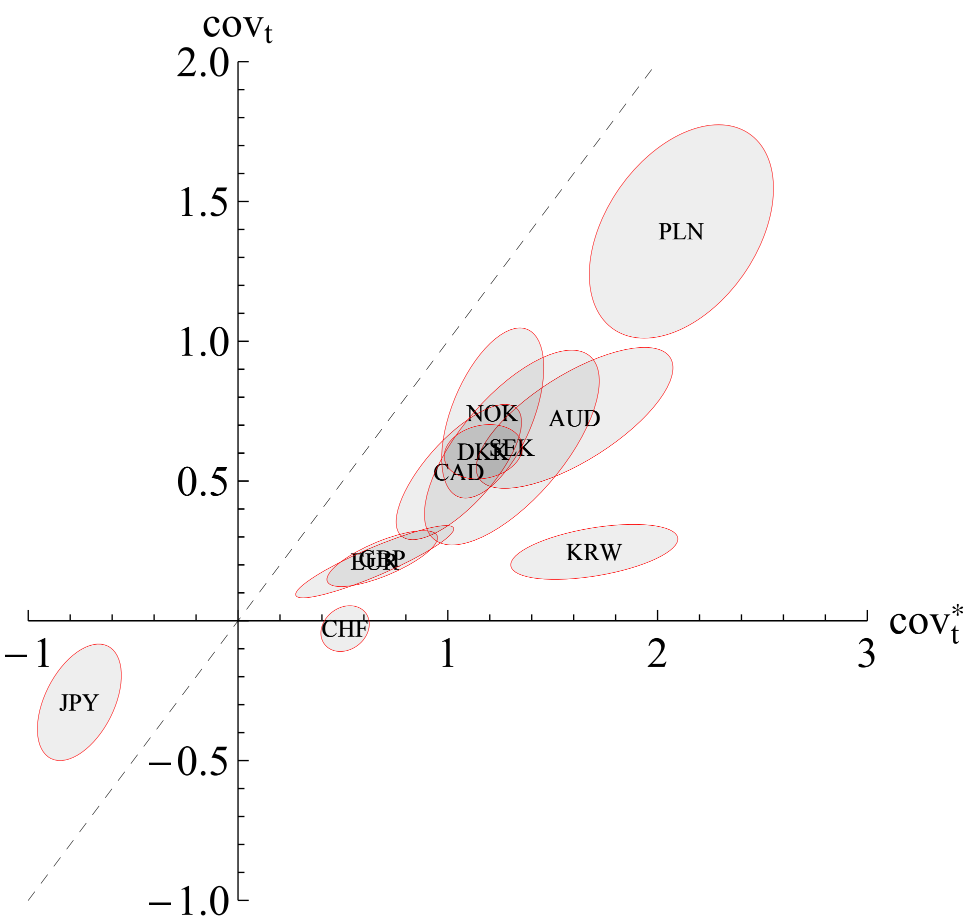

The centre of each confidence ellipse represents a currency's average risk-neutral and realized covariance. In population, each ellipse would contain 20% of its currency's data points under normality. The orientation of each ellipse reflects the time-series correlation between realized and risk-neutral covariance for the given currency, while the ellipse's size reflects their volatilities. We plot a dotted 45° line for comparison.

Table B.7: Beyond the log investor

This table reports the R^2-maximizing univariate, bivariate, 3-variate, and 4-variate specifications in regressions of 24-month realized currency excess returns onto combinations of QRP, IRD, the average forward discount $\overline{\text{IRD}}$, and the real exchange rate, q. The table reports standard errors (computed using a nonparametric block bootstrap) in brackets. See Section 2.2.5 for more detail. The last line reports R^2 in %.

Panel regressions with currency fixed effects				
Regressor	univariate	bivariate	3-variate	4-variate
QRP, β	4.995	5.654	3.799	3.541
	(1.565)	(1.402)	(1.657)	(1.836)
IRD, γ				-1.059
				(1.573)
$\overline{\text{IRD}}$, δ			-5.060	-4.266
			(1.605)	(1.538)
RER, ζ		-0.413	-0.780	-0.804
		(0.136)	(0.159)	(0.188)
R^2	22.03	35.40	43.56	44.09

Table B.8: Joint tests of statistical significance

This table presents results from three currency forecasting regressions:

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \qquad (2.19)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1} \qquad (2.20)$$

$$\frac{e_{i,t+1}}{e_{i,t}} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1} \qquad (2.22)$$

Realizations for excess returns and currency appreciation correspond to the forecasting horizon of 24 months. The Table reports p-values of Wald tests of various hypotheses on the regression coefficients. H_0^1 is the hypothesis suggested by Result 2: $\alpha = \gamma = 0$ and $\beta = 1$ in regression (2.19), $\alpha = 0$ and $\beta = 1$ in regression (2.20), and $\alpha = 0$ and $\beta = \gamma = 1$ in regression (2.22). Hypothesis H_0^2 drops the constraint that $\alpha = 0$, and therefore tests our model's ability to predict differences in currency returns but not its ability to predict the absolute level of (dollar) returns. Hypothesis H_0^3 is that QRP is not useful for forecasting. For each Wald test, we report both the asymptotic p-values obtained from the χ^2 distribution and p-values from a bootstrapped small-sample distribution (in the format asymptotic p-value / small-sample p-value).

Panel A: Pooled panel regression			
Regression	(2.19)	(2.20)	(2.22)
H_0^1: $\alpha = \gamma = 0, \beta = 1$	0.029 / 0.357		
H_0^1: $\alpha = 0, \beta = 1$		0.039 / 0.342	
H_0^1: $\alpha = 0, \beta = \gamma = 1$			0.030 / 0.340
H_0^2: $\beta = 1, \gamma = 0$	0.342 / 0.546		
H_0^2: $\beta = 1$		0.155 / 0.299	
H_0^2: $\beta = 1, \gamma = 1$			0.339 / 0.493
H_0^3: $\beta = 0$	0.050 / 0.088	0.021 / 0.097	0.049 / 0.082
Panel B: Panel regression with currency fixed effects			
H_0^2: $\beta = 1, \gamma = 0$	0.029 / 0.256		
H_0^2: $\beta = 1$		0.011 / 0.163	
H_0^2: $\beta = 1, \gamma = 1$			0.029 / 0.238
H_0^3: $\beta = 0$	0.008 / 0.051	0.001 / 0.089	0.008 / 0.051

B.2 Lognormal models

Suppose that the SDF, X_{t+1} and Y_{t+1} are conditionally jointly lognormal, and write lower-case variables for logs of the corresponding upper-case variables. Assume further that X_{t+1} and Y_{t+1} are tradable. Then we have the following three facts:

$$\begin{aligned}
\operatorname{cov}_t(X_{t+1}, Y_{t+1}) &= \mathbb{E}_t\, X_{t+1}\, \mathbb{E}_t\, Y_{t+1} \left(e^{\operatorname{cov}_t(x_{t+1}, y_{t+1})} - 1\right) \\
\operatorname{cov}^*_t(X_{t+1}, Y_{t+1}) &= \operatorname{cov}_t(X_{t+1}, Y_{t+1}) e^{\operatorname{cov}_t(m_{t+1}, x_{t+1}+y_{t+1})} \\
\operatorname{cov}^*_t(X_{t+1}, Y_{t+1}) &= \mathbb{E}^*_t\, X_{t+1}\, \mathbb{E}^*_t\, Y_{t+1} \left(e^{\operatorname{cov}^*_t(x_{t+1}, y_{t+1})} - 1\right).
\end{aligned}$$

These follow by direct calculation because $\log \mathbb{E}_t\, Z_{t+1} = \mathbb{E}_t \log Z_{t+1} + \frac{1}{2} \operatorname{var}_t \log Z_{t+1}$ for any conditionally lognormal random variable Z_{t+1} (and using the definition (2.4) of the risk-neutral measure to derive the second and third facts).

The first fact implies that equation (2.2) can be rewritten (in the lognormal case) as

$$\log \mathbb{E}_t \frac{\widetilde{R}_{t+1}}{R^{\$}_{f,t}} = -\operatorname{cov}_t(m_{t+1}, \widetilde{r}_{t+1}),$$

and in particular that $\mathrm{ERP}_t = -\operatorname{cov}_t(m_{t+1}, r_{t+1})$ and $\mathrm{CRP}_{i,t} = -\operatorname{cov}_t(m_{t+1}, \Delta e_{i,t+1})$, where ERP_t and $\mathrm{CRP}_{i,t}$ are defined in the main text and we write $r_{t+1} = \log R_{t+1}$ and $\Delta e_{i,t+1} = \log(e_{i,t+1}/e_{i,t})$. Combined with the second fact, this gives (in the lognormal case) equation (2.26) in the main text:

$$\log \frac{\operatorname{cov}_t(R_{t+1}, e_{i,t+1}/e_{i,t})}{\operatorname{cov}^*_t(R_{t+1}, e_{i,t+1}/e_{i,t})} = \mathrm{ERP}_t + \mathrm{CRP}_{i,t}\,.$$

To see that this is equivalent to (2.27), exponentiate both sides and use the definitions of ERP_t and $\mathrm{CRP}_{i,t}$, together with the first and third facts above, to conclude that

$$\frac{\mathbb{E}_t\, R_{t+1}\, \mathbb{E}_t\, e_{i,t+1}/e_{i,t} \left\{e^{\operatorname{cov}_t(r_{t+1}, \Delta e_{i,t+1})} - 1\right\}}{\mathbb{E}^*_t\, R_{t+1}\, \mathbb{E}^*_t\, e_{i,t+1}/e_{i,t} \left\{e^{\operatorname{cov}^*_t(r_{t+1}, \Delta e_{i,t+1})} - 1\right\}} = \mathbb{E}_t \frac{R_{t+1}}{R^{\$}_{f,t}}\, \mathbb{E}_t \frac{R^i_{f,t} e_{i,t+1}}{R^{\$}_{f,t} e_{i,t}}\,.$$

By the definition (2.4) of the risk-neutral measure, we have $\mathbb{E}^*_t\, R_{t+1} = R^{\$}_{f,t}$; and similarly we have $\mathbb{E}^*_t\, e_{i,t+1}/e_{i,t} = R^{\$}_{f,t}/R^i_{f,t}$ by equation (2.5). Equation (2.27) follows.

Table B.9: Out-of-sample forecast performance

We define a dollar-neutral out-of-sample R^2 similar to Goyal and Welch (2008):

$$R^2_{OS} = 1 - \frac{\sum_i \sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_i \sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2},$$

where $\varepsilon^Q_{i,t+1}$ and $\varepsilon^B_{i,t+1}$ denote forecast errors (for currency i against the dollar) of the quanto theory and the benchmark, respectively. We use the quanto theory and three competitor benchmarks to forecast currency appreciation as follows:

$$\begin{aligned}
\text{Theory: } \mathbb{E}^Q_t \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \text{QRP}_{i,t} + \text{IRD}_{i,t} \\
\text{UIP: } \mathbb{E}^U_t \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \text{IRD}_{i,t} \\
\text{Constant: } \mathbb{E}^C_t \frac{e_{i,t+1}}{e_{i,t}} - 1 &= 0 \\
\text{PPP: } \mathbb{E}^P_t \frac{e_{i,t+1}}{e_{i,t}} - 1 &= \left(\frac{\pi^\$_t}{\pi^i_t}\right)^2 - 1
\end{aligned}$$

We also report results for the following decomposition of R^2_{OS}, which focusses on dollar-neutral forecast performance for currency i:

$$R^2_{OS,i} = 1 - \frac{\sum_j \sum_t (\varepsilon^Q_{i,t+1} - \varepsilon^Q_{j,t+1})^2}{\sum_j \sum_t (\varepsilon^B_{i,t+1} - \varepsilon^B_{j,t+1})^2}.$$

The second panel reports R^2_{OS} measures by currency. (All R^2_{OS} measures are reported in %.) The last line of the table reports p-values for a small-sample Diebold–Mariano test of the null hypothesis that the quanto theory and competitor model perform equally well for all currencies.

Benchmark	**IRD**	**Constant**	**PPP**
R^2_{OS}	10.91	9.57	26.05
$R^2_{OS,AUD}$	9.71	0.93	11.42
$R^2_{OS,CAD}$	6.24	6.55	21.31
$R^2_{OS,CHF}$	1.40	16.37	11.43
$R^2_{OS,DKK}$	10.22	7.71	23.36
$R^2_{OS,EUR}$	7.65	5.36	24.56
$R^2_{OS,GBP}$	2.98	9.74	32.35
$R^2_{OS,JPY}$	19.21	9.59	33.74
$R^2_{OS,KRW}$	21.98	17.09	34.71
$R^2_{OS,NOK}$	3.43	12.86	18.97
$R^2_{OS,PLN}$	13.25	8.32	19.62
$R^2_{OS,SEK}$	7.68	5.88	28.22
DM p-value	0.039	0.000	0.000

B.3 Binary forecast accuracy

In this section, we follow the approach of Jordà and Taylor (2012) by computing a *correct classification frontier* (CCF) to assess the forecast performance of the quanto theory.

Denote by $f^Q_{i,j,t} = \mathrm{QRP}_{i,t} - \mathrm{QRP}_{j,t}$ and $f^B_{i,j,t}$ the forecasts obtained, respectively, from the quanto variable and a competitor benchmark for currency pair (i,j) at time t. Similarly, $r_{i,t} = e_{i,t+1}/e_{i,t} - R^\$_t/R^i_t$ denotes the realized excess return of currency i against the dollar, and $r_{i,j,t} = r_{i,t} - r_{j,t}$ represents the dollar-neutral return in currency pair (i,j). We calculate the *true positive* (TP) and *true negative* (TN) rates for each forecasting model as a function of a threshold, c. For the quanto forecast, for instance,

$$\mathrm{TP}(c) = \frac{\sum_{i,j:\, f^m_{i,j,t}>c \text{ and } r_{i,j,t}>0} 1}{\sum_{i,j:\, r_{i,j,t}>0} 1} \qquad \text{and} \qquad \mathrm{TN}(c) = \frac{\sum_{i,j:\, f^Q_{i,j,t}<c \text{ and } r_{i,j,t}<0} 1}{\sum_{i,j:\, r_{i,j,t}<0} 1}.$$

These represent, respectively, the fractions of ex post positive long and short returns that were correctly identified ex ante as profitable by the forecasting model. For the same 55 dollar-neutral currency pairs used above, we find that $\mathrm{TP}(0) = 0.50$, $\mathrm{TN}(0) = 0.64$, with a weighted average correct classification of 0.57 for the quanto forecast.

As binary accuracy does not reflect the magnitudes of returns from the signal, we follow Jordà and Taylor (2012) and compute the corresponding *return-weighted* true positive (TP^*) and true negative (TN^*) rates as

$$\mathrm{TP}^*(c) = \frac{\sum_{i,j:\, f^Q_{i,j,t}>c \text{ and } r_{i,j,t}>0} r_{i,j,t}}{\sum_{i,j:\, r_{i,j,t}>0} r_{i,j,t}} \qquad \text{and} \qquad \mathrm{TN}^*(c) = \frac{\sum_{i,j:\, f^Q_{i,j,t}<c \text{ and } r_{i,j,t}<0} r_{i,j,t}}{\sum_{i,j:\, r_{i,j,t}<0} r_{i,j,t}}.$$

We find $\mathrm{TP}^*(0) = 0.58$, $\mathrm{TN}^*(0) = 0.67$, with a weighted average of 0.63. Both rates increase relative to the equally-weighted classifications, which implies that the direction of excess return realizations is more likely to have been predicted by the quanto variable when these realizations are large.

The CCF (and analogously CCF*) is defined as the set of pairs $\{\mathrm{TP}(c), \mathrm{TN}(c)\}$ for all possible values of c between $-\infty$ and ∞. Varying the threshold level, c, trades off true positives against true negatives by shifting the direction of the forecast. For instance, for $c = \infty$, the true negative rate is maximized at $\mathrm{TN} = 1$, at the cost of $\mathrm{TP} = 0$. Since $\mathrm{TN}(c)$ and $\mathrm{TP}(c)$ must lie between 0 and 1, we can plot the resulting CCF in the unit square, and compute the *area under the CCF* (AUC).

Intuitively, the AUC can be interpreted as the probability that the forecast for a randomly chosen positive return realization will be higher than that for a randomly chosen negative return realization. Under the UIP forecast the excess return on any currency is 0, so the CCF is the diagonal with slope -1 in the unit square and, accordingly, AUC $= 0.5$.

Figure B.7: Correct classification frontier (CCF) and AUC statistics

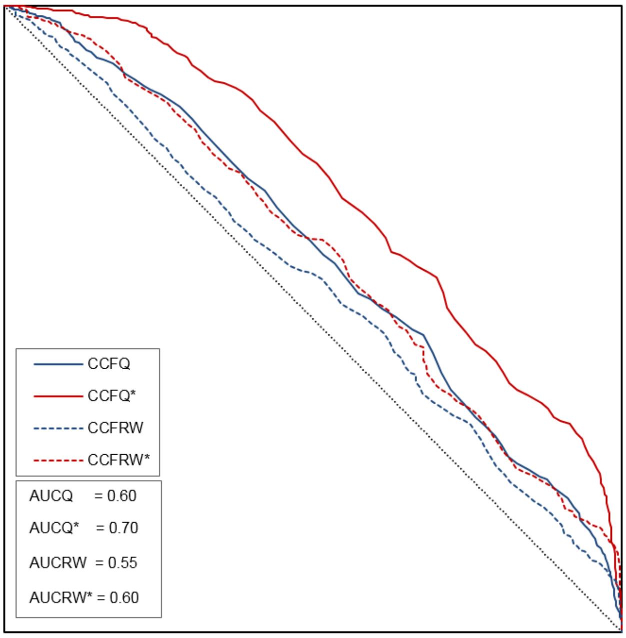

CCF and AUC statistics for the quanto excess return forecast, and a competitor excess return forecast under which exchange rates follow a random walk.

We benchmark the quanto forecast against the driftless random walk model considered above (which forecasts the currency excess return as being equal to the interest rate differential). Figure B.7 shows the resulting CCFs. The quanto forecast outperforms the random walk model for equally-weighted and return-weighted classifications. For the quanto forecast, AUCQ $= 0.60$ and AUCQ* $= 0.70$, while the random walk model achieves AUCRW $= 0.55$ and AUCRW* $= 0.60$. Both forecasts correctly identify large returns more often than small returns, as the CCF* (red) lies above the CCF (blue) in both cases.

We also reverse the conditioning in the true positive and true negative rates, to calculate how likely a forecast is to signal the correct direction of trade, and denote these by PT(c) and NT(c), respectively. In the case of the quanto theory,

$$\mathrm{PT}(c) = \frac{\sum_{i,j:\, f^Q_{i,j,t}>0 \text{ and } r_{i,j,t}>c} 1}{\sum_{i,j:\, f^Q_{i,j,t}>0} 1} \qquad \text{and} \qquad \mathrm{NT}(c) = \frac{\sum_{i,j:\, f^Q_{i,j,t}<0 \text{ and } r_{i,j,t}<c} 1}{\sum_{i,j:\, f^Q_{i,j,t}<0} 1}.$$

Figure B.8: Reverse-conditioned CCF and AUC statistics

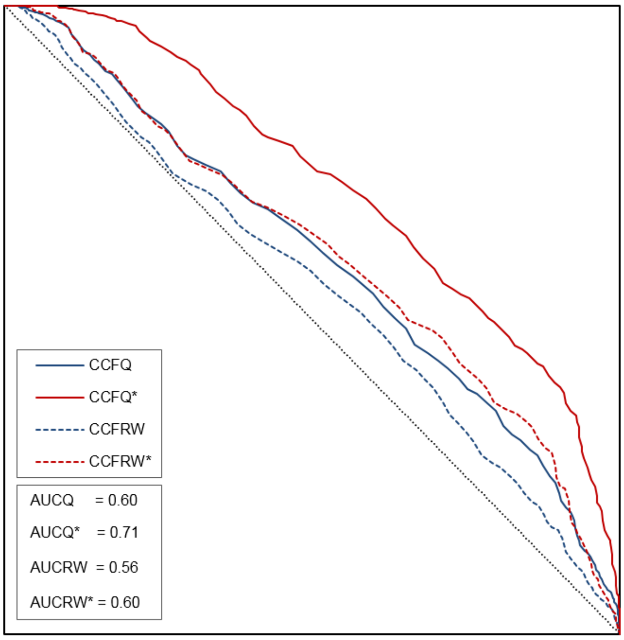

CCF and AUC statistics for the quanto excess return forecast, and a competitor excess return forecast under which exchange rates follow a random walk.

Figure B.9: CCF and AUC statistics for currency appreciation forecasts

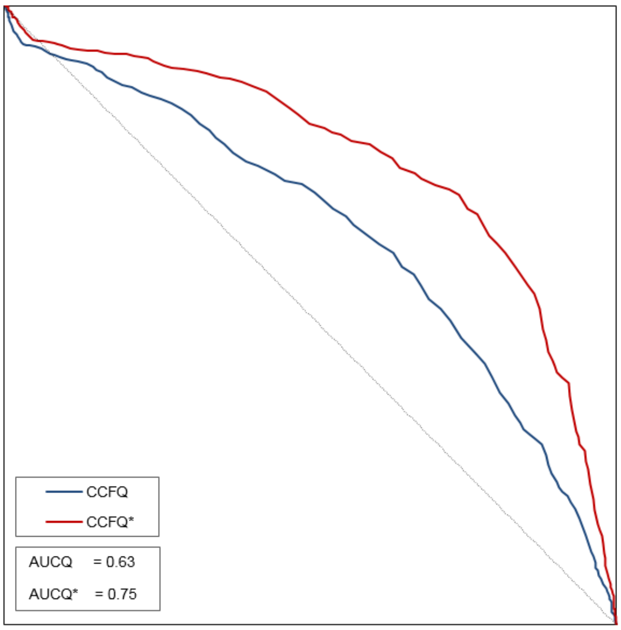

Figure B.10: Reverse-conditioned CCF and AUC statistics for currency appreciation forecasts

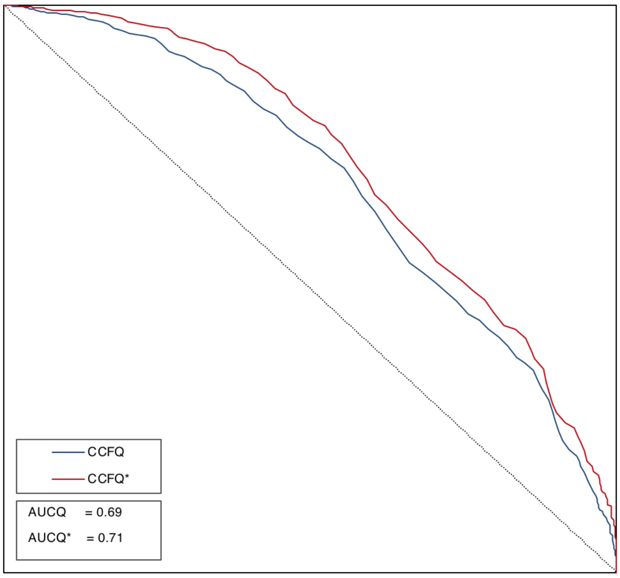

We find PT(0) = 0.60, NT(0) = 0.54, PT*(0) = 0.65, and NT*(0) = 0.63. Plotting the resulting CCFs, Figure B.8 shows that the quanto variable outperforms the random walk forecast with AUC measures of AUCQ = 0.60 and AUCQ* = 0.71, as against the random walk model with AUCRW = 0.55 and AUCRW* = 0.60.

Figures B.9 and B.10 repeat this exercise, but now the goal is to forecast currency appreciation, as opposed to currency excess returns. In this case, the random walk forecast is represented by the diagonal with slope -1 in the unit square, and AUC = 0.5. As the figures show, the quanto forecast outperforms the random walk model, with AUCQ = 0.63 and AUCQ* = 0.75. The outperformance persists under reverse conditioning, with AUCQ = 0.69 and AUCQ* = 0.71.

B.4 Quantos in Colacito and Croce (2011)

This section studies the relationship between the currency risk premium, QRP, and the residual covariance term in the two-country long-run risk model of Colacito and Croce (2011). Log consumption growth, log dividend growth, the long-run risk variable, the log SDF, the log market return, and the log risk-free rate follow these processes:

$$\begin{aligned}
\Delta c_t &= \mu_c + x_{t-1} + \varepsilon_{c,t}, \\
\Delta d_t &= \mu_d + \lambda x_{t-1} + \varepsilon_{d,t}, \\
x_t &= \rho x_{t-1} + \varepsilon_{x,t}, \\
m_{t+1} &= \log\delta - \psi^{-1} x_t + \kappa_c \frac{1-\gamma\psi}{\psi(1-\rho\kappa_c)} \varepsilon_{x,t+1} - \gamma \varepsilon_{c,t+1}, \\
r_{d,t+1} &= \overline{r}_d + \psi^{-1} x_t + \kappa_d \frac{\lambda - 1/\psi}{1-\rho\kappa_d} \varepsilon_{x,t+1} + \varepsilon_{d,t+1}, \\
r_f &= \overline{r}_f + \psi^{-1} x_t.
\end{aligned}$$

The representative agent has Epstein–Zin preferences with risk aversion γ and elasticity of intertemporal substitution ψ. Shocks are i.i.d. Normal over time, with mean zero and (diagonal) covariance matrix Σ, with diagonal $[\sigma^2, \varphi_d^2\sigma^2, \varphi_x^2\sigma^2]$. Thus returns and the SDF are jointly lognormal and subject to the issues described in Subsection 2.2.4. Between-country correlations of shocks are ρ_c^{hf}, ρ_d^{hf}, and ρ_x^{hf}, respectively. The exchange rate satisfies $e_{t+1}/e_t = M_{t+1}^f / M_{t+1}$, where M^f denotes the foreign SDF (which is uniquely determined, as markets are complete).

The baseline calibration is symmetric, so both currencies are equally "risky." To generate a currency risk premium, we vary—one-by-one—the parameter values for (i) the volatility of the foreign long-run risk shock, governed by φ_x^f, (ii) its persistence, ρ^f, (iii) the cross-country correlation of long-run risk shocks, ρ_x^{hf}, and (iv) the cross-country correlation of consumption shocks, ρ_c^{hf}. We plot the resulting comparative statics in Figure B.11 below. We use the baseline calibration of Colacito and Croce (2011) for all other model parameters. With the exception of ρ_x^{hf}, which is equal to 1 in the baseline calibration, we vary the parameters of interest in a symmetric window around their baseline values.

Through the lens of this model, we now consider the identity (2.6), which decomposes the currency risk premium into risk-neutral covariance (QRP) and the residual covariance term:

$$\mathbb{E}_t \frac{e_{i,t+1}}{e_{i,t}} - \frac{R_{f,t}^{\$}}{R_{f,t}^{i}} = \mathrm{QRP}_{i,t} - \underbrace{\mathrm{cov}_t\left(M_{t+1}R_{t+1}, \frac{e_{i,t+1}}{e_{i,t}}\right)}_{\text{residual covariance term}}.$$

As shown in panel (a), a lower long-run risk volatility generates a positive risk premium on the foreign currency, positive QRP, a positive residual, and a negative interest rate differential. (The calibration is monthly, but we annualize by multiply-

Figure B.11: Comparative statics

(a) Long-run risk volatility, $\varphi_x^f \sigma^f$ (b) Long-run risk persistence, ρ^f

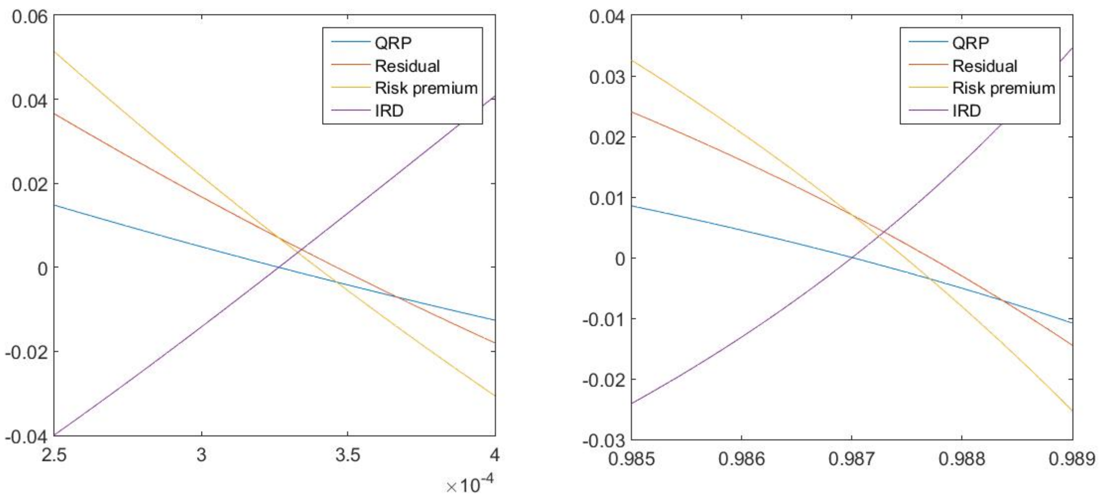

(c) Long-run risk cross-country correlation, ρ_x^{hf} (d) Consumption risk cross-country correl., ρ_c^{hf}

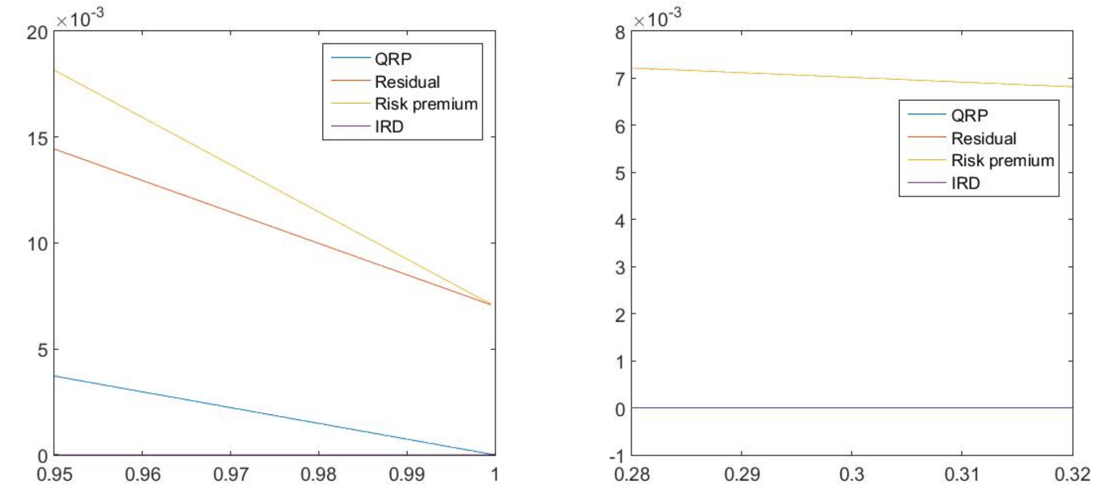

Each panel plots the comparative statics of the risk premium, risk-neutral covariance (QRP), the residual covariance, and the interest rate differential (IRD) with respect to a single model parameter (varied on the horizontal axis). In panel (d), QRP and IRD are both zero so the risk premium coincides with the residual.

ing all quantities by 12, so the y-axis is in annual terms in all four panels.) As the residual scales with QRP, we would expect to find that the coefficient on QRP in a forecasting regression is larger than 1. Qualitatively, the same holds for a lower persistence of the foreign long-run risk process in panel (b). The risk premia in panels (c) and (d) are symmetric, in the sense that they increase the expected appreciation of *both* currencies in another manifestation of Siegel's paradox (see Section 2.1.2). In the case of a less-than-perfect cross-country correlation of long-run risk shocks, the resulting risk premium is captured proportionately by QRP and the residual, and would lead to a β coefficient larger than 1 in our forecasting regressions.

B.5 Evidence from other quanto contracts

Due to the limited availability of time-series data on quanto forwards, we look at USD-denominated futures on the Nikkei 225 index, which have started trading on the CME prior to the beginning of our OTC sample. We collect prices for USD-denominated Nikkei 225 futures traded on CME, and JPY-denominated Nikkei 225 futures traded on JPX (Osaka) for a sample period from 2004 through 2017. (JPY-denominated futures are also traded on CME, but at much lower volumes than the JPX-traded contracts.) Contracts expire each quarter, in March, June, September, and December, and we use contracts with the latest available expiration, which have a maturity ranging from 9-12 months. To calculate the QRP and IRD measures, we use dollar- and yen-denominated LIBOR rates matched to the maturity of the respective pair of futures. Table B.10 below reports the results for our baseline regressions.

Table B.10: Forecasting regressions with exchange traded quanto-futures

This table reports the results of running regressions (2.20), (2.21), (2.22), and (2.24) for the USD/JPY currency pair at the 12-month horizon, based on dollar-denominated quanto futures on the Nikkei 225 (traded on CME). Since this setting essentially takes the perspective of a log investor who holds the Nikkei, the exchange rate is defined as ¥1 = \$$e$. We report the OLS estimates along with Hansen–Hodrick standard errors. R^2 are reported in %.

Regression	(2.20)	(2.21)	(2.22)	(2.24)
α (p.a.)	0.018	0.026	0.022	0.026
	(0.027)	(0.036)	(0.036)	(0.036)
β	0.339	0.366	0.274	1.366
	(0.720)	(1.917)	(0.587)	(1.917)
γ			1.293	1.366
			(1.912)	(1.917)
R^2	0.26	0.26	3.60	3.44

We also calculate the out-of-sample R^2 based on mean-squared forecast errors as in Section 2.3. The quanto-based forecast outperforms the random walk and the UIP forecast by 1.96% and 3.25%, respectively, over the given period.

There are two important caveats. First, the available futures only provide information about a single currency-pair, dollar-yen. One of the strengths of the quanto data used in this chapter lies in the cross-sectional dimension, which allows us to compute dollar-neutral forecasts in isolation from any base-currency effects. Table B.14 suggests that the yen is not representative of the remaining panel. (USD-denominated futures on the FTSE 100 are also traded on the CME, which would provide information about dollar-sterling, but these contracts have only been traded since late 2015.) Second, the theory calls for quanto forward prices rather than quanto futures prices. If interest-rate movements are correlated with the underlying assets (as is plausibly true both of exchange rates and of the Nikkei 225) the two will differ. It is not clear how the pricing discrepancies between futures and forwards would affect the predictive power of our theory when applied to futures contracts.

B.6 Supplementary Tables and Figures

Table B.11: Principal components analysis of residuals

This table reports the loadings on the principal components of realized residuals obtained from the quanto theory (top panel) and the fixed-effects specification of regression (2.20) (bottom panel). In order to limit the impact of missing observations, the residuals are only obtained for the balanced panel of currencies (excluding DKK, KRW, and PLN).

	PC1	**PC2**	**PC3**	**PC4**	**PC5**	**PC6**	**PC7**	**PC8**
				Theory residuals				
AUD	0.520	0.160	0.108	−0.443	−0.273	0.235	0.578	−0.183
CAD	0.311	−0.015	−0.107	−0.257	−0.090	0.458	−0.490	0.606
CHF	0.194	−0.124	0.644	0.344	−0.534	−0.270	−0.067	0.228
EUR	0.243	−0.265	−0.308	0.688	−0.119	0.490	0.127	−0.179
GBP	0.083	−0.471	0.579	−0.104	0.552	0.296	−0.046	−0.176
JPY	0.353	0.741	0.200	0.325	0.397	0.009	−0.145	−0.055
NOK	0.472	−0.194	−0.190	−0.147	−0.099	−0.334	−0.527	−0.532
SEK	0.427	−0.283	−0.238	0.093	0.382	−0.472	0.324	0.446
Explained	61.26%	26.49%	7.26%	2.80%	0.93%	0.53%	0.39%	0.34%
				Regression residuals				
AUD	0.532	0.138	0.019	−0.261	0.665	−0.025	−0.368	−0.227
CAD	0.276	−0.057	−0.175	−0.271	0.248	0.057	0.657	0.566
CHF	0.177	−0.243	0.662	0.273	0.070	−0.594	0.052	0.193
EUR	0.178	−0.291	−0.430	0.732	0.248	−0.004	0.205	−0.244
GBP	−0.086	−0.440	0.489	0.024	0.195	0.714	0.073	−0.082
JPY	0.558	0.539	0.243	0.289	−0.372	0.303	0.154	−0.050
NOK	0.369	−0.451	−0.060	−0.399	−0.409	−0.148	0.229	−0.506
SEK	0.351	−0.384	−0.209	0.068	−0.295	0.144	−0.555	0.516
Explained	65.70%	16.33%	10.65%	3.10%	2.12%	1.20%	0.54%	0.34%

Table B.12: R^2 of different variable combinations

This table reports the R^2 (in %) from currency excess return forecasting regressions (with currency fixed effects) using all possible univariate, bivariate, 3-variate and 4-variate combinations of the quanto-implied risk premium (QRP), the interest rate differential (IRD), the average interest rate differential ($\overline{\text{IRD}}$), and the real exchange rate (RER).

	univariate	bivariate	3-variate	4-variate
QRP	22.03			
RER	7.97			
IRD	2.77			
$\overline{\text{IRD}}$	2.06			
QRP, RER		35.40		
$\overline{\text{IRD}}$, RER		34.47		
IRD, RER		28.22		
QRP, $\overline{\text{IRD}}$		22.77		
QRP, IRD		22.60		
IRD, $\overline{\text{IRD}}$		2.79		
QRP, $\overline{\text{IRD}}$, RER			43.56	
QRP, IRD, RER			39.89	
IRD, $\overline{\text{IRD}}$, RER			36.77	
QRP, IRD, $\overline{\text{IRD}}$			22.80	
QRP, IRD, $\overline{\text{IRD}}$, RER				44.09

Table B.13: Quantos and the real exchange rate

This table presents results from currency excess return forecasting regressions that extend the baseline results in Table B.4 by adding the log real exchange rate to the regressors on the right-hand side. Following Dahlquist and Penasse (2017), we compute the log real exchange rate as $\mathrm{RER}_{i,t} = \log\left(e_{i,t}\frac{P_{i,t}}{P_{\$,t}}\right)$, where $P_{i,t}$ and $P_{\$,t}$ are consumer price indices for country i and the US, respectively, obtained from the OECD.

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{B.6}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \beta\,\mathrm{QRP}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{B.7}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \gamma\,\mathrm{IRD}_{i,t} + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{B.8}$$

$$\frac{e_{i,t+1}}{e_{i,t}} - \frac{R^{\$}_{f,t}}{R^{i}_{f,t}} = \alpha_i + \zeta\,\mathrm{RER}_{i,t} + \varepsilon_{i,t+1} \tag{B.9}$$

The two panels report coefficient estimates for each pooled and fixed effects regression, respectively, with standard errors (computed using a nonparametric block bootstrap) in parentheses, see Section 2.2.6 for more detail.

Panel regressions with currency fixed effects

Regression	(B.6)	(B.7)	(B.8)	(B.9)
QRP, β	4.292	5.654		
	(1.843)	(1.402)		
IRD, γ	-2.624		-4.791	
	(1.547)		(1.242)	
RER, ζ	-0.616	-0.413	-0.729	-0.314
	(0.205)	(0.136)	(0.201)	(0.162)
R^2	39.89	35.40	28.22	7.97

Table B.14: Separate return forecasting regressions using QRP and IRD predictors

This table reports the results of running regressions (2.20), (2.21), (2.22), and (2.24) separately for each currency at the 24-month horizon, and at 6- and 12-month horizons for the euro. We report the OLS estimates along with Hansen–Hodrick standard errors. R^2 are reported in %.

Currency	AUD	CAD	CHF	DKK	EUR	EUR	EUR	GBP	JPY	KRW	NOK	PLN	SEK
Horizon	24m	24m	24m	24m	6m	12m	24m	24m	24m	24m	24m	24m	24m
Panel A: Regression (2.20): $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^{i}_{f,t} = \alpha + \beta\,\mathrm{QRP}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.062	−0.085	−0.003	−0.052	−0.040	−0.071	−0.060	−0.086	−0.012	−0.068	−0.180	−0.065	−0.106
	(0.071)	(0.042)	(0.038)	(0.022)	(0.056)	(0.052)	(0.030)	(0.031)	(0.090)	(0.034)	(0.061)	(0.026)	(0.048)
β	3.258	4.754	−1.657	4.125	3.702	6.361	4.148	9.217	4.750	4.227	11.860	3.580	5.930
	(3.991)	(3.546)	(6.903)	(1.723)	(6.263)	(5.527)	(3.367)	(3.791)	(10.959)	(1.757)	(4.698)	(0.956)	(3.316)
R^2	12.15	25.39	0.60	17.42	3.17	17.98	25.93	57.48	4.06	46.59	49.96	33.01	38.00
Panel B: Regression (2.21): $e_{i,t+1}/e_{i,t} - R^{\$}_{f,t}/R^{i}_{f,t} = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.091	−0.006	0.001	0.014	−0.015	−0.019	−0.034	−0.043	−0.152	0.007	−0.091	0.005	−0.042
	(0.084)	(0.030)	(0.027)	(0.023)	(0.083)	(0.040)	(0.025)	(0.034)	(0.046)	(0.034)	(0.065)	(0.045)	(0.035)
γ	−2.859	4.135	−2.246	2.147	2.626	1.869	−1.439	−5.564	25.539	0.312	−3.310	−0.118	−1.765
	(2.743)	(3.543)	(3.067)	(2.036)	(7.375)	(6.349)	(3.255)	(6.779)	(8.318)	(3.011)	(3.698)	(1.211)	(2.730)
R^2	19.82	12.30	7.33	13.77	1.23	1.31	3.90	6.93	57.26	0.14	14.39	0.09	7.28
Panel C: Regression (2.22): $e_{i,t+1}/e_{i,t} - 1 = \alpha + \beta\,\mathrm{QRP}_{i,t} + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.093	−0.055	0.010	−0.041	−0.055	−0.092	−0.078	−0.082	−0.165	−0.063	−0.185	−0.041	−0.117
	(0.087)	(0.044)	(0.035)	(0.021)	(0.053)	(0.043)	(0.027)	(0.033)	(0.079)	(0.046)	(0.070)	(0.032)	(0.043)
β	0.698	5.291	−1.698	5.252	10.008	12.916	7.321	9.760	−1.348	4.241	11.230	4.736	7.895
	(3.130)	(2.984)	(6.621)	(1.260)	(7.198)	(4.771)	(2.895)	(3.519)	(7.485)	(1.719)	(3.491)	(0.848)	(2.552)
γ	−1.525	6.019	−1.250	3.857	11.447	11.992	4.651	3.094	27.182	1.514	0.253	2.419	2.938
	(2.429)	(2.637)	(3.050)	(1.671)	(8.450)	(4.880)	(2.175)	(3.124)	(8.344)	(2.149)	(2.402)	(1.003)	(1.683)
R^2	9.79	46.74	3.04	48.62	14.42	45.19	33.51	57.29	59.41	48.22	46.61	45.28	39.00
Panel D: Regression (2.24): $e_{i,t+1}/e_{i,t} - 1 = \alpha + \gamma\,\mathrm{IRD}_{i,t} + \varepsilon_{i,t+1}$													
α (p.a.)	−0.091	−0.006	0.001	0.014	−0.007	−0.019	−0.034	−0.043	−0.152	0.007	−0.091	0.005	−0.042
	(0.084)	(0.030)	(0.027)	(0.023)	(0.041)	(0.040)	(0.025)	(0.034)	(0.046)	(0.034)	(0.065)	(0.045)	(0.035)
γ	−1.859	5.135	−1.246	3.147	3.626	2.869	−0.439	−4.564	26.539	1.312	−2.310	0.882	−0.765
	(2.743)	(3.543)	(3.067)	(2.036)	(7.375)	(6.349)	(3.255)	(6.779)	(8.318)	(3.011)	(3.698)	(1.211)	(2.730)
R^2	9.47	17.78	2.38	25.54	2.32	3.03	0.38	4.77	59.13	2.48	7.57	4.79	1.45

Figure B.12: Time series of annualized expected currency appreciation

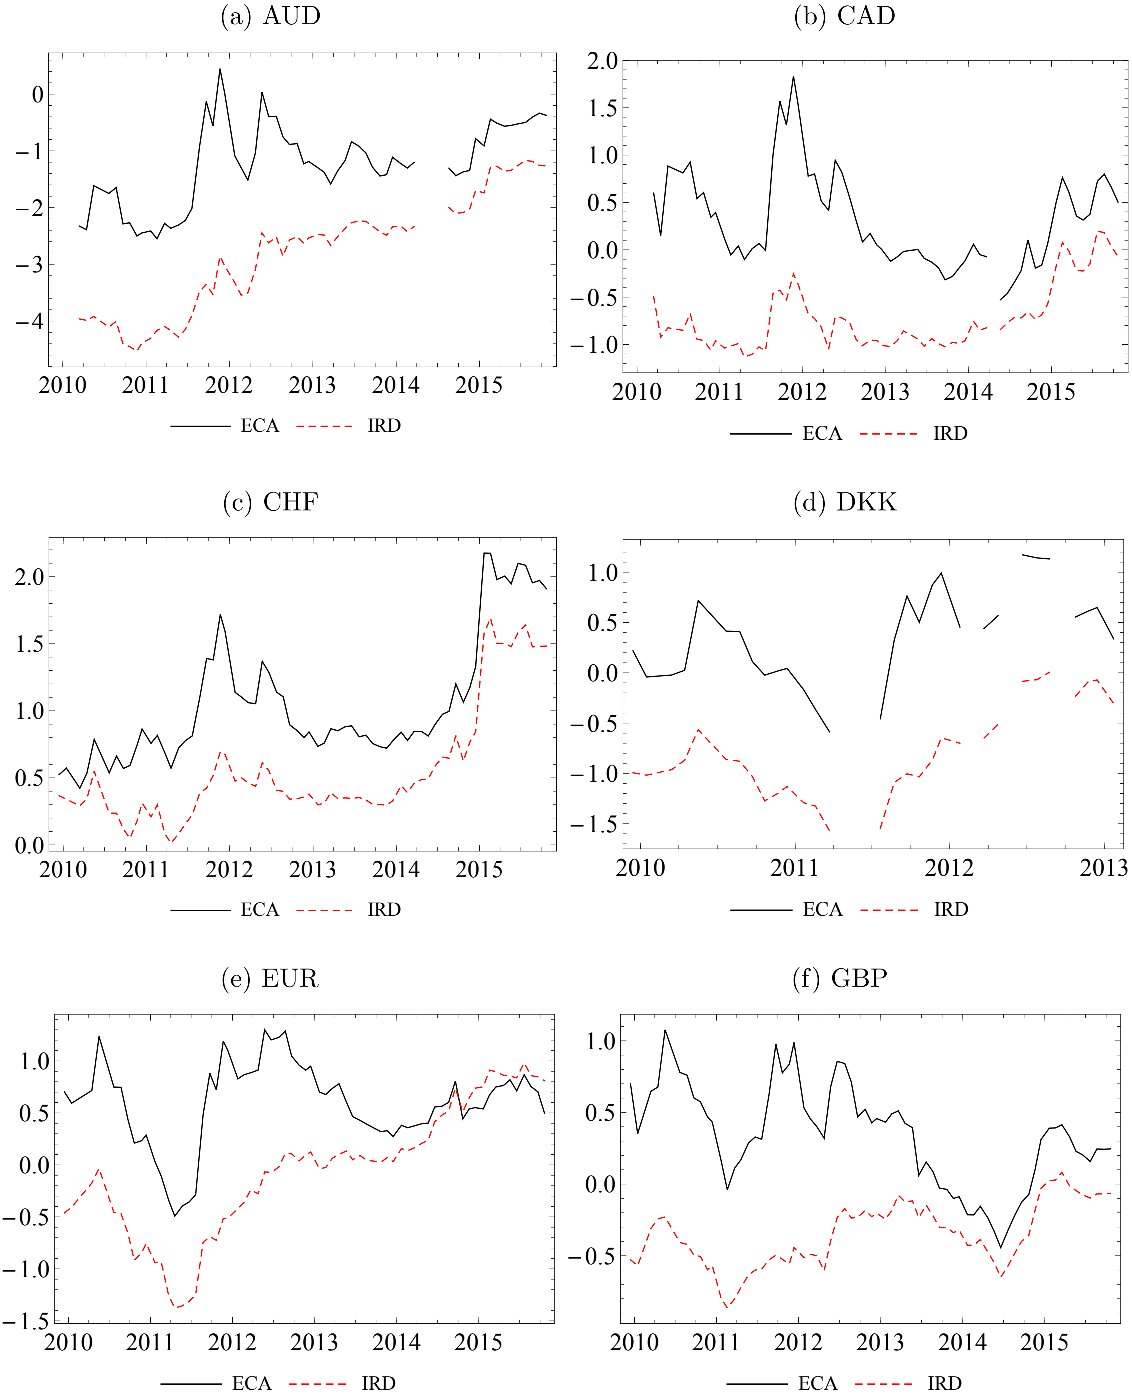

Expected currency appreciation implied by the quanto theory (ECA) and by UIP (IRD).

Figure B.12: Time series of annualized expected currency appreciation

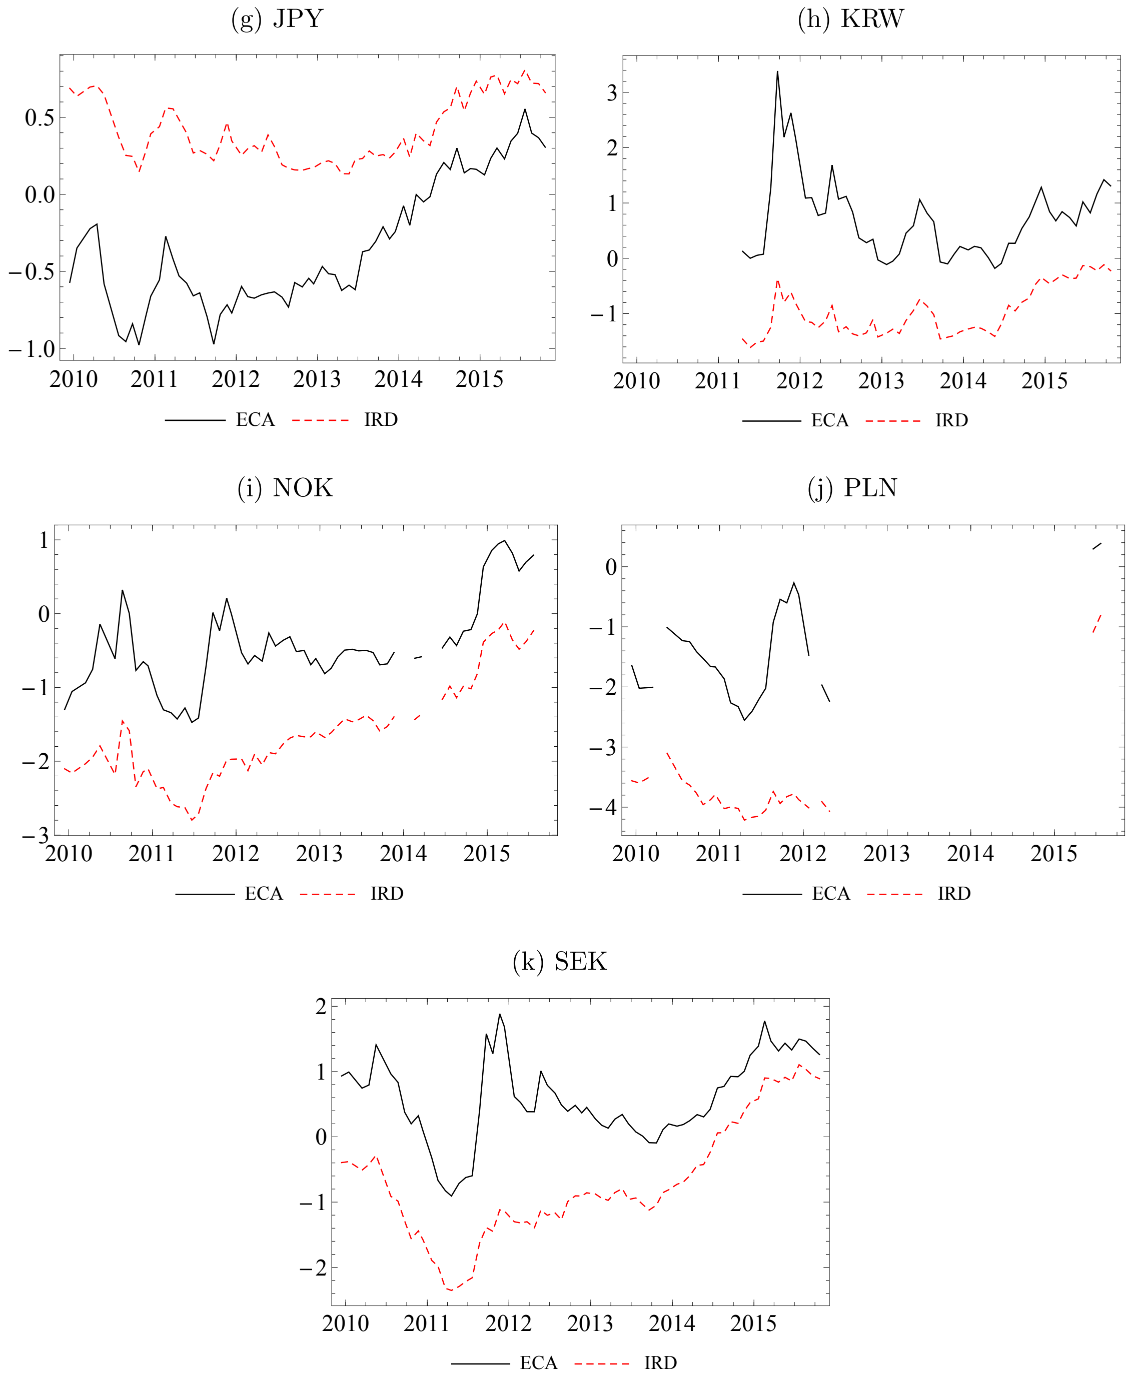

Expected currency appreciation implied by the quanto theory (ECA) and by UIP (IRD).

Figure B.13: Term structure of the euro-dollar risk premium

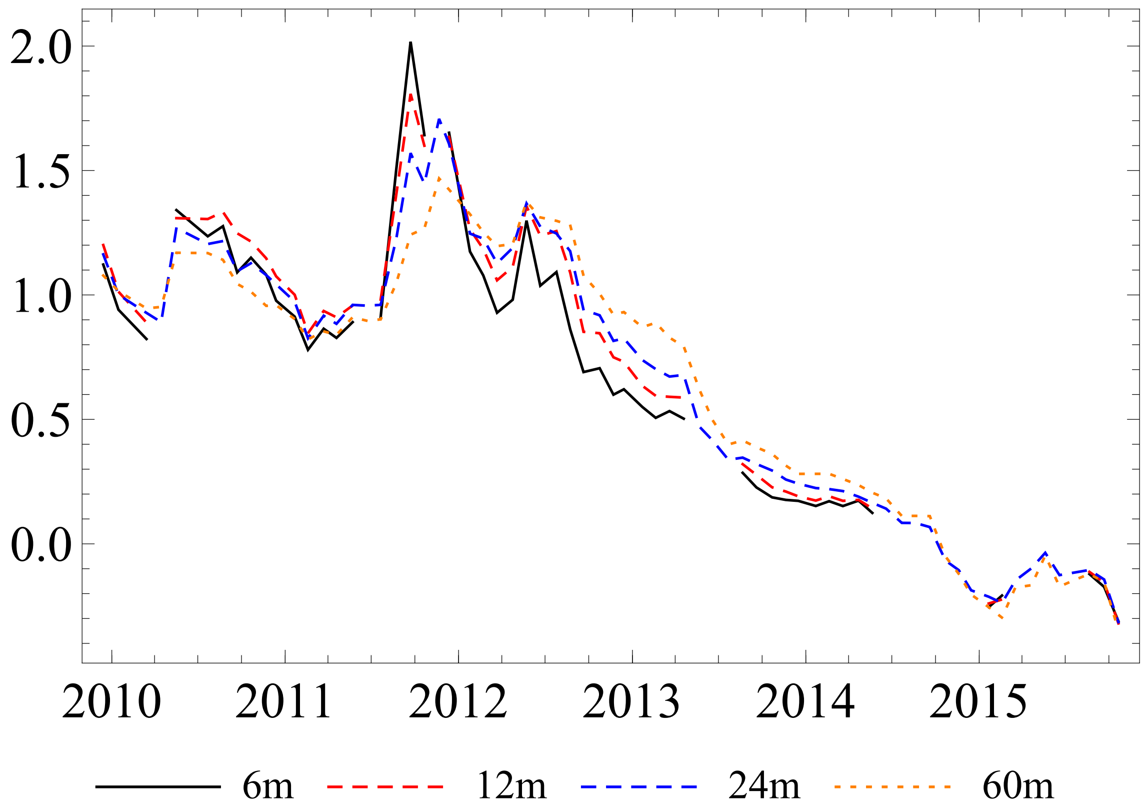

Term structure as measured by QRP, in the time series for horizons of 6, 12, 24, and 60 months.

Figure B.14: Histogram of the small-sample distributions

(a) Pooled, H_0^1

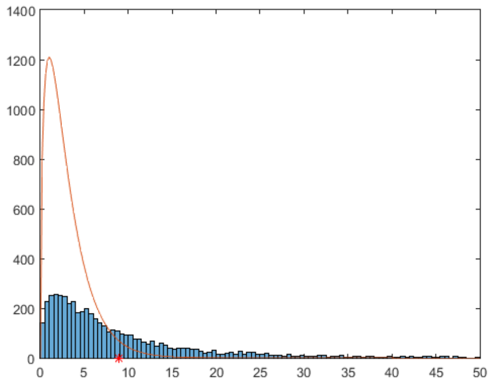

(b) Fixed effects, H_0^2

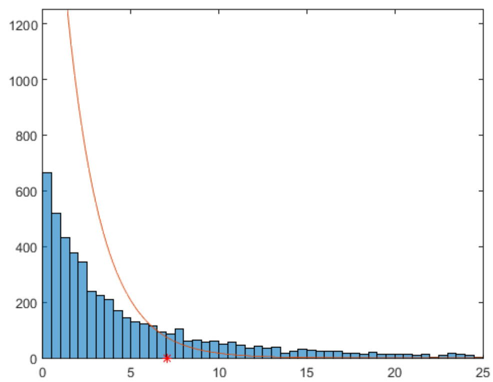

(c) Pooled, H_0^3

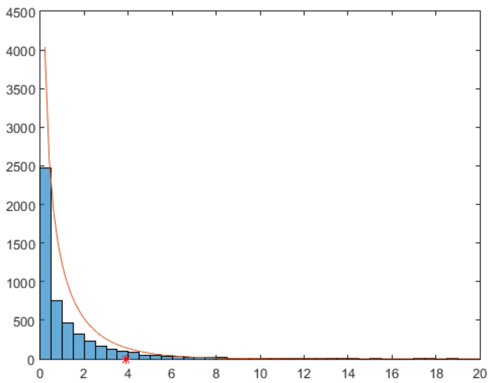

(d) Fixed effects, H_0^3

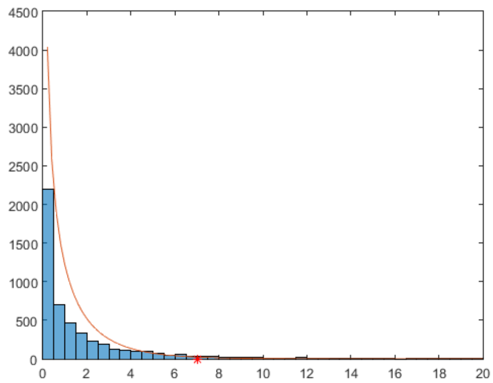

Small-sample distributions of the test statistics for various hypotheses on regression (2.22). The asymptotic distribution is shown as a solid line. Asterisks indicate the test statistics for the original sample.

C. Appendix to Bets and Betas

C.1 Tables

Table C.1: Correlations of futures positions

This table reports the correlations of net positions of different trader groups, measured in contracts (shown in the left panel), and percentage of open interest (right panel).

Correlations of net positions (in contracts and shares of open interest)									
	ndi	nlf	nam	no	$\widetilde{ndi}$	$\widetilde{nlf}$	$\widetilde{nam}$	$\widetilde{no}$	
ndi	1				1				$\widetilde{ndi}$
nlf	-0.881	1			-0.872	1			$\widetilde{nlf}$
nam	-0.427	0.042	1		-0.377	0.032	1		$\widetilde{nam}$
no	0.137	-0.325	-0.113	1	-0.108	-0.156	-0.059	1	$\widetilde{no}$

Table C.2: Summary statistics of futures positions

This table reports the means, standard deviations, and autocorrelations of net positions, in thousands of contracts (unscaled, Panel A) and shares of open interest (Panel B), of intermediaries (ndi / ndis), hedge funds (nlf / nlfs), institutional investors (nam / nams), and others (no / nos). The bottom three rows of each panel report, respectively, the average time-series standard deviation (averaged across currencies), the cross-sectional standard-deviation of the within-currency means, and the total standard deviation across the panel.

Panel A	$\mu(ndi)$	$\mu(nlf)$	$\mu(nam)$	$\mu(no)$	$\sigma(ndi)$	$\sigma(nlf)$	$\sigma(nam)$	$\sigma(no)$	$\rho(ndi)$	$\rho(nlf)$	$\rho(nam)$	$\rho(no)$
AUD	-16.32	24.17	-7.29	-4.27	62.29	36.81	19.31	9.30	0.98	0.96	0.99	0.96
BRL	-7.64	5.00	2.44	0.02	13.91	7.32	2.79	11.15	0.96	0.94	0.98	0.96
CAD	-11.94	-2.95	2.58	6.17	48.37	35.72	12.97	6.80	0.97	0.97	0.98	0.91
CHF	5.74	-4.08	-0.63	0.47	22.91	15.50	1.10	2.86	0.94	0.91	0.87	0.91
EUR	27.16	-27.61	-2.25	9.48	82.10	62.41	25.45	17.69	0.98	0.98	0.98	0.93
GBP	25.93	8.42	-24.52	-7.08	72.64	48.09	25.48	10.50	0.98	0.97	0.99	0.94
JPY	19.44	-20.98	10.29	3.14	63.90	52.32	23.39	18.13	0.96	0.96	0.99	0.95
MXN	-32.75	14.71	18.78	-2.39	44.68	47.76	17.06	5.39	0.96	0.96	0.90	0.92
NZD	-5.93	7.93	-2.32	-0.30	11.75	9.37	5.90	1.35	0.97	0.96	0.99	0.89
Time-series	0.41	0.51	-0.32	0.58	46.95	35.03	14.83	9.24	0.97	0.96	0.96	0.93
Cross section					20.60	16.52	11.92	5.11				
Pooled					57.93	43.97	21.31	12.00				

Panel B	$\mu(\widetilde{ndi})$	$\mu(\widetilde{nlf})$	$\mu(\widetilde{nam})$	$\mu(\widetilde{no})$	$\sigma(\widetilde{ndi})$	$\sigma(\widetilde{nlf})$	$\sigma(\widetilde{nam})$	$\sigma(\widetilde{no})$	$\rho(\widetilde{ndi})$	$\rho(\widetilde{nlf})$	$\rho(\widetilde{nam})$	$\rho(\widetilde{no})$
AUD	-0.18	0.21	-0.05	-0.03	0.49	0.29	0.14	0.08	0.98	0.96	0.97	0.94
BRL	-0.19	0.16	0.09	-0.07	0.44	0.25	0.11	0.29	0.96	0.94	0.96	0.96
CAD	-0.12	-0.02	0.03	0.05	0.36	0.26	0.11	0.06	0.97	0.96	0.97	0.91
CHF	0.08	-0.05	-0.01	0.01	0.38	0.24	0.02	0.07	0.94	0.91	0.86	0.93
EUR	0.04	-0.06	-0.01	0.03	0.30	0.21	0.08	0.05	0.97	0.96	0.97	0.91
GBP	0.14	0.07	-0.14	-0.05	0.40	0.27	0.13	0.07	0.97	0.96	0.97	0.94
JPY	0.04	-0.06	0.08	0.01	0.35	0.26	0.15	0.10	0.96	0.96	0.98	0.94
MXN	-0.26	0.12	0.14	-0.02	0.34	0.37	0.11	0.05	0.95	0.95	0.88	0.91
NZD	-0.22	0.26	-0.06	-0.01	0.41	0.30	0.15	0.05	0.97	0.96	0.97	0.87
Time-series	-0.07	0.07	0.01	-0.01	0.38	0.27	0.11	0.09	0.97	0.95	0.95	0.92
Cross section					0.15	0.13	0.09	0.04				
Pooled					0.41	0.30	0.14	0.10				

Table C.3: The interaction of futures bets and currency betas

This table reports the results for time-series regressions of daily currency returns on S&P 500 returns and its interaction with the relative positioning of hedge funds $(\widetilde{nlf})$.

$$r_{i,t} = \alpha_i + \beta_i r_t^{S\&P} + \beta_i^* r_t^{S\&P} \cdot \widetilde{nlf}_{i,t} + \varepsilon_{i,t}, \tag{3.1}$$

where $r_{i,t}$ and $r_t^{S\&P}$ denote the currency return of currency i and the return on the S&P 500 from day $t-1$ to day t, respectively. The column on the far right shows the estimate of β^* obtained from the pooled panel regression.

$$r_{i,t} = \alpha_i + \beta_i r_t^{S\&P} + \beta^* r_t^{S\&P} \cdot \widetilde{nlf}_{i,t} + \varepsilon_{i,t}, \tag{3.2}$$

I report t-statistics based on robust standard errors in parentheses. The standard errors for the pooled regression (3.2) are clustered at the currency level.

	AUD	**BRL**	**CAD**	**CHF**	**EUR**	**GBP**	**JPY**	**MXN**	**NZD**	**pooled**
					Panel A: Full sample: 2006 - 2017					
α_i	-0.013	-0.056	-0.015	0.009	-0.007	-0.017	0.013	-0.028	-0.008	
	(-0.98)	(-2.18)	(-1.58)	(0.64)	(-0.64)	(-1.50)	(1.14)	(-2.31)	(-0.60)	
β_i	0.421	0.388	0.302	0.017	0.129	0.142	-0.244	0.389	0.371	
	(1.51)	(11.60)	(25.55)	(0.91)	(10.22)	(10.56)	(-16.86)	(20.36)	(23.50)	
β_i^*	0.028	0.009	0.268	0.211	-0.068	-0.074	0.300	-0.067	0.132	0.082
	(0.35)	(0.06)	(6.25)	(2.74)	(-0.96)	(-1.64)	(5.31)	(-1.74)	(3.31)	(1.73)
R^2 in %	35.76	11.21	33.56	0.61	7.23	8.83	18.34	35.56	29.27	21.50
Obs.	2,864	1,334	2,864	2,864	2,864	2,864	2,864	2,864	2,864	24,232
					Panel B: Pre-Crisis / Crisis: 2006 - 2009					
α_i	0.025		0.005	0.020	0.017	-0.012	0.028	-0.013	0.015	
	(0.83)		(0.23)	(0.84)	(0.78)	(-0.53)	(1.30)	(-0.60)	(0.51)	
β_i	0.430		0.270	0.001	0.108	0.125	-0.255	0.338	0.383	
	(11.56)		(15.45)	(0.03)	(6.21)	(7.03)	(-12.05)	(11.33)	(19.18)	
β_i^*	-0.140		0.190	0.143	-0.050	-0.047	0.125	-0.110	0.111	0.012
	(-1.18)		(2.84)	(1.87)	(-0.39)	(-0.70)	(1.41)	(-2.20)	(1.99)	(0.22)
R^2 in %	39.64		31.32	0.78	7.79	9.59	28.30	41.70	35.11	28.47
Obs.	923		923	923	923	923	923	923	923	7,370
					Panel C: Post-Crisis: 2010 - 2017					
α_i	-0.029	-0.056	-0.026	0.005	-0.021	-0.020	0.003	-0.039	-0.017	
	(-2.18)	(-2.17)	(-2.73)	(0.24)	(-1.57)	(-1.64)	(0.21)	(-2.75)	(-1.16)	
β_i	0.378	0.388	0.347	0.027	0.194	0.172	-0.195	0.473	0.317	
	(17.75)	(11.60)	(28.09)	(0.89)	(8.40)	(8.37)	(-10.48)	(25.18)	(10.23)	
β_i^*	0.274	0.009	0.319	0.266	0.123	-0.123	0.462	0.013	0.251	0.186
	(4.59)	(0.06)	(6.58)	(1.85)	(1.18)	(-2.23)	(6.35)	(0.27)	(3.54)	(2.79)
R^2 in %	32.41	11.21	37.09	0.63	7.66	8.74	12.39	33.85	23.83	17.72
Obs.	1,940	1,334	1,940	1,940	1,940	1,940	1,940	1,940	1,940	16,854

Table C.4: Futures positions and market risk

This table reports the results for pooled panel regressions of changes in the absolute size of futures positions by each trader group on contemporaneous shocks to "risk-on" and "risk-off" shocks, represented by the S&P 500 and the VIX. Negative (positive) changes in the size of futures positions ($\Delta|ndi|$, $\Delta|nlf|$, and $\Delta|nam|$) reflect contractions (expansions) of the outstanding bets of each trader group. To measure asymmetric shocks to the market risk-taking environment, I define $r_t^{S\&P^+} = [r_t^{S\&P}]^+$, $r_t^{S\&P^-} = [r_t^{S\&P}]^-$, $\Delta VIX^+ = [\Delta VIX_t]^+$, and $\Delta VIX^- = [\Delta VIX_t]^-$ (note that $r^{S\&P^-}$ and ΔVIX^+ represent the "risk-off" shock proxies). I then run the following contemporaneous regressions.

$$\Delta|ndi|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.3}$$
$$\Delta|ndi|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.4}$$
$$\Delta|nlf|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.5}$$
$$\Delta|nlf|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.6}$$
$$\Delta|nam|_{i,t} = \alpha_i + \eta r_t^{S\&P^-} + \gamma r_t^{S\&P^+} + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.7}$$
$$\Delta|nam|_{i,t} = \alpha_i + \eta \Delta VIX_t^+ + \gamma \Delta VIX_t^- + \delta\, r_{i,t-1} + \varepsilon_{i,t} \tag{3.8}$$

Again, $r_{i,t}$ denotes the currency return of currency i from week $t-1$ to week t. Standard errors are clustered at the currency level and t-statistics reported in parentheses.

	Pre-Crisis / Crisis: 2006 - 2009						*Post-Crisis: 2010 - 2017*					
	Intermediaries		Hedge Funds		Asset Managers		Intermediaries		Hedge Funds		Asset Managers	
Regression	(3.3)	(3.4)	(3.5)	(3.6)	(3.7)	(3.8)	(3.3)	(3.4)	(3.5)	(3.6)	(3.7)	(3.8)
$r^{S\&P^-}$	-13.67		-8.25		2.67		70.91		35.30		-16.57	
	(-0.81)		(-0.60)		(1.01)		(2.31)		(1.99)		(-1.33)	
$r^{S\&P^+}$	1.63		11.07		-12.15		38.02		15.84		1.60	
	(0.12)		(0.83)		(-1.45)		(2.24)		(0.85)		(0.13)	
ΔVIX^+		-4.86		-10.88		5.40		-56.41		-23.10		12.95
		(-0.38)		(-0.90)		(1.88)		(-2.28)		(-1.96)		(1.13)
ΔVIX^-		33.86		22.57		0.42		-4.13		11.08		-1.80
		(3.83)		(2.29)		(0.12)		(-0.30)		(1.01)		(-0.23)
$r_{i,t-1}$	111.60	117.55	73.07	76.88	-8.00	-7.92	-45.67	-39.23	-8.25	-0.47	-8.00	-8.17
	(2.53)	(2.49)	(1.69)	(1.75)	(-3.38)	(-3.01)	(-1.09)	(-0.92)	(-0.23)	(-0.01)	(-0.72)	(-0.75)
R^2 in (%)	2.27	2.59	1.42	1.63	0.94	0.78	0.94	0.94	0.28	0.20	0.36	0.46
Currencies	8	8	8	8	8	8	9	9	9	9	9	9
Obs.	1,476	1,476	1,476	1,476	1,476	1,476	3,354	3,354	3,354	3,354	3,354	3,354

Table C.5: Summary statistics of short-run currency-equity co-movement

This table reports the mean, standard deviation, and autocorrelation of currency betas, $\beta^{MKT}_{i,t\to t+1}$, and correlations with the S&P 500 and the VIX, $\rho^{SPX}_{i,t\to t+1}$ and $\rho^{VIX}_{i,t\to t+1}$, respectively. The correlations and betas are computed using closing prices for the 5 trading days following date t to result in weekly observations over the full sample from January 2006 to June 2017.

	ρ^{SPX}			β^{SPX}			ρ^{VIX}		
	Mean	St. dev.	Autocorr.	Mean	St. dev.	Autocorr.	Mean	St. dev.	Autocorr.
AUD	0.39	0.57	0.32	0.33	0.67	0.17	-0.32	0.55	0.29
BRL	0.39	0.48	0.12	0.38	0.75	0.03	-0.34	0.48	0.10
CAD	0.42	0.49	0.32	0.28	0.45	0.14	-0.34	0.50	0.19
CHF	0.01	0.56	0.23	-0.01	0.66	0.10	0.01	0.54	0.19
EUR	0.17	0.56	0.25	0.11	0.52	0.18	-0.14	0.55	0.18
GBP	0.19	0.52	0.16	0.13	0.53	0.08	-0.16	0.51	0.19
JPY	-0.32	0.52	0.21	-0.24	0.58	0.11	0.26	0.53	0.15
MXN	0.51	0.45	0.20	0.40	0.55	0.05	-0.46	0.46	0.14
NZD	0.33	0.52	0.34	0.28	0.69	0.19	-0.28	0.53	0.26
Mean	0.23	0.52	0.24	0.18	0.60	0.11	-0.20	0.52	0.19

Table C.6: Predicting betas using net futures positions

This table reports the results for pooled panel regressions of the realized correlations and betas of exchange rates with the S&P and the VIX on the scaled level of net positions of hedge funds ("Leveraged Funds", denoted by $\widetilde{nlf}$). The realized betas and correlations are computed using closing prices for the 5 trading days following date t. The data form an unbalanced panel for the two subsample periods 2006-2009 and 2010-2017. I run the following predictive regressions for $y_{i,t\to t+1} = \{\rho^{SPX}_{i,t\to t+1}, \beta^{MKT}_{i,t\to t+1}, \rho^{VIX}_{i,t\to t+1}\}$:

$$y_{i,t\to t+1} = \alpha_i + \eta\,\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, fd^w_{i,t} + \varepsilon_{i,t+1} \tag{3.9}$$

$$y_{i,t\to t+1} = \alpha + \eta\,\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, fd^w_{i,t} + \lambda\, y_{i,t-1} + \varepsilon_{i,t+1}. \tag{3.10}$$

$fd^w_{i,t}$ denotes the 1-week forward discount of currency i versus the dollar and $r_{i,t}$ denotes the currency return of currency i from week $t-1$ to week t. Standard errors are clustered at the currency level and t-statistics reported in parentheses.

	Panel A: with currency fixed effects					
	2006-2009			*2010-2017*		
	ρ^{SPX}	β^{SPX}	ρ^{VIX}	ρ^{SPX}	β^{SPX}	ρ^{VIX}
$\widetilde{nlf}$	-0.141	-0.025	0.155	0.315	0.215	-0.256
	(-1.32)	(-0.44)	(1.97)	(4.40)	(2.85)	(-4.21)
fd^w	0.838	0.540	-0.716	0.017	0.058	-0.085
	(1.52)	(1.59)	(-1.34)	(0.18)	(0.63)	(-0.92)
$r_{i,t-1}$	-1.235	-0.821	1.088	-1.053	0.503	1.890
	(-1.43)	(-1.34)	(1.97)	(-1.88)	(0.79)	(2.79)
R^2 in (%)	1.11	0.31	1.13	2.61	0.76	1.92
Obs.	1485	1485	1485	3371	3371	3371
	Panel B: without currency fixed effects					
$\widetilde{nlf}$	0.023	0.069	-0.009	0.267	0.263	-0.243
	(0.58)	(2.43)	(-0.20)	(4.75)	(3.72)	(-4.57)
fd^w	-2.147	-1.539	2.327	-0.132	-0.272	0.076
	(-4.80)	(-3.93)	(5.47)	(-0.57)	(-0.93)	(0.34)
$r_{i,t-1}$	-0.641	-0.535	0.291	-0.951	0.753	1.844
	(-0.67)	(-0.73)	(0.55)	(-1.70)	(1.59)	(2.73)
lag(dep. var.)	0.360	0.226	0.279	0.325	0.158	0.260
	(7.63)	(13.71)	(8.43)	(7.03)	(4.42)	(8.63)
R^2 in (%)	22.17	11.60	16.80	14.62	4.71	9.82
Obs.	1477	1477	1477	3371	3371	3371

Table C.7: Predicting betas using net futures positions

This table reports the results for pooled panel regressions of the realized correlations and betas of exchange rates with the S&P and the VIX on the scaled level of net positions of hedge funds ("Leveraged Funds", denoted by $\widetilde{nlf}$). The realized betas and correlations are computed using closing prices for the 5 trading days following date t. The data form an unbalanced panel for the two subsample periods 2006-2009 and 2010-2017. I run the following predictive regressions for $y_{i,t\to t+1} = \{\rho^{SPX}_{i,t\to t+1}, \beta^{MKT}_{i,t\to t+1}, \rho^{VIX}_{i,t\to t+1}\}$:

$$\Delta y_{i,t\to t+1} = \alpha_i + \eta\, \Delta\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, \Delta fd^w_{i,t} + \varepsilon_{i,t+1} \quad (C.1)$$

$$\Delta y_{i,t\to t+1} = \alpha + \eta\, \Delta\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, \Delta fd^w_{i,t} + \varepsilon_{i,t+1} \quad (C.2)$$

$fd^w_{i,t}$ denotes the 1-week forward discount of currency i versus the dollar and $r_{i,t}$ denotes the currency return of currency i from week $t-1$ to week t. Standard errors are clustered at the currency level and t-statistics reported in parentheses.

	2006-2009			*2010-2017*		
	ρ^{SPX}	β^{SPX}	ρ^{VIX}	ρ^{SPX}	β^{SPX}	ρ^{VIX}
Panel A: with currency fixed effects						
$\Delta\widetilde{nlf}$	-0.006	-0.045	-0.004	0.369	0.333	-0.352
	(-0.06)	(-0.39)	(-0.03)	(2.63)	(2.72)	(-1.70)
Δfd^w	-0.718	-1.172	1.031	0.077	0.084	-0.165
	(-0.79)	(-2.14)	(1.30)	(9.75)	(8.22)	(-18.99)
$r_{i,t-1}$	0.381	0.503	-1.687	-0.828	2.582	1.699
	(0.35)	(0.53)	(-2.45)	(-1.07)	(1.73)	(1.89)
R^2 in (%)	0.03	0.07	0.25	0.23	0.35	0.36
Panel B: without currency fixed effects						
$\Delta\widetilde{nlf}$	-0.006	-0.044	-0.005	0.369	0.333	-0.352
	(-0.06)	(-0.39)	(-0.03)	(2.63)	(2.72)	(-1.70)
Δfd^w	-0.718	-1.172	1.030	0.077	0.084	-0.165
	(-0.79)	(-2.14)	(1.30)	(9.74)	(8.25)	(-18.99)
$r_{i,t-1}$	0.383	0.503	-1.686	-0.829	2.582	1.698
	(0.35)	(0.53)	(-2.45)	(-1.07)	(1.73)	(1.89)
R^2 in (%)	0.03	0.07	0.25	0.23	0.35	0.36
Obs.	1476	1476	1476	3354	3354	3354

Table C.8: A contrarian trading strategy

This table reports the returns to the trading strategy described in subsection 3.2.3 for different conditioning thresholds of the S&P return. The strategy is designed to exploit temporary dislocations in FX markets following the unwinding of futures positions by hedge funds. Let $\Omega^{CW}_{i,t}$ denote the *number of futures contracts* in currency i against the dollar, included in the strategy at time t:

$$\Omega^{CW}_{i,t} = \frac{\omega^{CW}_{i,t}}{\sum_j |\omega^{CW}_{j,t}| e_{i,t} s_i}, \text{ where } \omega^{CW}_{i,t} = -\Delta nlf_{i,t} \ \mathbb{1}_{\{r^{S\&P}_t < x\}} \ \mathbb{1}_{\{\Delta |nlf|_{i,t} < 0\}} \ \mathbb{1}_{\{\text{sign}(\Delta nlf_{i,t}) = \text{sign}(r_{i,t})\}}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function which takes value 1 if $\cdot$ is true, and 0 otherwise, $e_{i,t}$ denotes the exchange rate, and s_i the contract size in units of foreign currency, such that $e_{i,t} s_i$ expresses the dollar notional of each contract. Denote by $\Omega^{EW}_{i,t}$ the *dollar notional amount* in futures contracts of currency i against the dollar at time t

$$\Omega^{EW}_{i,t} = \frac{\omega^{EW}_{i,t}}{\sum_j |\omega^{EW}_{j,t}|}, \text{ where } \omega^{EW}_{i,t} = -\text{sign}(\Delta nlf_{i,t}) \ \mathbb{1}_{\{r^{S\&P}_t < x\}} \ \mathbb{1}_{\{\Delta |nlf|_{i,t} < 0\}} \ \mathbb{1}_{\{\text{sign}(\Delta nlf_{i,t}) = \text{sign}(r_{i,t})\}}.$$

The *contract-weighted* (CW) strategy is described by $\Omega^{CW}_{i,t}$ and takes positions, which are proportional in the cross section to the amount of positions unwound by hedge funds over the previous week. The *equal-weighted* (EW) strategy is described by $\Omega^{EW}_{i,t}$ and fixes the dollar notional of each individual position, such that all non-zero positions taken at any point in time have the same absolute dollar exposure. Both strategies have a gross dollar notional of \$1 in each week, where the strategy is active. The below returns are based on 1-week forward exchange rates, with long (short) positions transacted at the ask (bid) price. Returns are unlevered and refer to positions formed weekly between January 5, 2010 and June 6, 2017. The strategy is inactive in week t, if $\Omega_{i,t} = 0 \, \forall \, i$.

Threshold (x)	$r^{S\&P} < 0$		$r^{S\&P} < -3\%$		None	
	CW	EW	CW	EW	CW	EW
Mean return p.w. (in %)	0.12	0.15	0.38	0.35	0.04	0.01
Std. deviation p.w. (in %)	1.19	1.14	1.33	1.27	1.16	1.07
Sharpe ratio p.a.	0.74	0.97	2.06	1.97	0.23	0.04
Total compound return (in %)	17.22	22.29	6.94	6.29	11.41	-0.02
Weeks total	389		389		389	
Weeks active	138		18		362	
Weeks active (in %)	35.48		4.63		93.06	

Table C.9: Predicting pairwise correlations using net futures positions

This table reports the results for within-currency pair regressions of pairwise currency correlations on differences in net positions of "Leveraged Funds" ($\widetilde{nlf}$) and "Asset Managers" ($\widetilde{nam}$), each scaled by open interest, and on interest differentials.

$$\rho_{i,j,t} = \alpha_{i,j} + \eta \mid \widetilde{nlf}_{i,t} - \widetilde{nlf}_{j,t} \mid + \gamma \mid \widetilde{nam}_{i,t} - \widetilde{nam}_{j,t} \mid + \phi \mid fd^w_{i,t} - fd^w_{j,t} \mid + \varepsilon_{i,j,t} \quad (3.11)$$

$$\rho_{i,j,t} = \alpha_{i,j} + \eta \mid \widetilde{nlf}_{i,t} - \widetilde{nlf}_{j,t} \mid + \varepsilon_{i,j,t} \quad (3.12)$$

$$\rho_{i,j,t} = \alpha_{i,j} + \gamma \mid \widetilde{nam}_{i,t} - \widetilde{nam}_{j,t} \mid + \varepsilon_{i,j,t} \quad (3.13)$$

$$\rho_{i,j,t} = \alpha_{i,j} + \phi \mid fd^w_{i,t} - fd^w_{j,t} \mid + \varepsilon_{i,j,t} \quad (3.14)$$

Correlations $\rho_{i,j,t}$ are computed using intraday exchange rates at 15-minute intervals during week t (i.e., the week following the observation of forward discounts and futures positions). Intraday data span the period from July 2010 to June 2017. Standard errors are clustered at the currency-pair level and t-statistics reported in parentheses.

Regression	(3.11)	(3.12)	(3.13)	(3.14)
$\widetilde{nlf}$ differential	-0.06	-0.05		
	(-1.70)	(-1.73)		
$\widetilde{nam}$ differential	-0.04		-0.04	
	(-0.69)		(-0.65)	
fd^w differential	32.08			30.12
	(0.58)			(0.54)
Average intercept	0.40	0.40	0.39	0.37
	(13.97)	(44.54)	(37.81)	(17.69)
Currency pair FE	Yes	Yes	Yes	Yes
Currency pairs	28	28	28	28
Observations	9,044	9,044	9,044	9,044
R^2	0.45%	0.33%	0.07%	0.04%

Table C.10: Regression of currency movements on contemporaneous net futures flows

This table reports the results for a contemporaneous regression of exchange rate movements on weekly net flows in futures positions. I regress weekly returns on the contemporaneous change in the net long position of "Leveraged Funds" (nlf) and "Asset Managers" (nam), first in absolute terms and then scaled by open interest ($\widetilde{nlf}$, $\widetilde{nam}$).

$$r_{i,t} = \alpha_i + \eta_i \Delta nlf_{i,t} + \gamma_i \Delta nam_{i,t} + \varepsilon_{i,t} \tag{3.15}$$

$$r_{i,t} = \alpha_i + \eta_i \Delta \widetilde{nlf}_{i,t} + \gamma_i \Delta \widetilde{nam}_{i,t} + \varepsilon_{i,t} \tag{3.16}$$

where $r_{i,t} = e_{i,t}/e_{i,t-1} - 1$ denotes the currency return on currency i versus the US dollar from week $t-1$ to week t. $\Delta nlf_{i,t}$ and $\Delta nam_{i,t}$, respectively, denote the change in the net positions of hedge funds and institutional investors in currency i versus the US dollar from week $t-1$ to week t, while $\Delta \widetilde{nlf}_{i,t}$ and $\Delta \widetilde{nam}_{i,t}$ refer analogously to the positions scaled by open interest. All futures positions are expressed in thousands of contracts. The estimated coefficients for regressions (3.15) and (3.16) are reported below with their respective robust t-statistics in parentheses. The weekly exchange rate movements and forward discounts are expressed in %.

(3.15)	*AUD*	*BRL*	*CAD*	*CHF*	*EUR*	*GBP*	*JPY*	*MXN*	*NZD*
Δnlf	0.07	0.15	0.05	0.07	0.04	0.04	0.04	0.03	0.26
	(11.96)	(3.00)	(8.10)	(5.63)	(11.22)	(9.35)	(12.38)	(6.53)	(8.61)
Δnam	0.01	0.29	0.07	0.65	0.05	0.01	0.08	0.00	0.36
	(0.21)	(1.53)	(1.88)	(2.37)	(3.97)	(0.55)	(4.97)	(0.36)	(4.39)
intercept	0.03	-0.19	-0.01	0.06	-0.01	-0.05	0.03	-0.07	0.06
	(0.40)	(-1.46)	(-0.20)	(0.94)	(-0.22)	(-0.85)	(0.57)	(-1.01)	(0.79)
Obs.	573	250	573	573	573	573	573	573	569
R^2 in %	19.08	4.05	11.05	11.42	19.17	13.14	25.36	6.60	15.61
(3.16)	*AUD*	*BRL*	*CAD*	*CHF*	*EUR*	*GBP*	*JPY*	*MXN*	*NZD*
$\Delta \widetilde{nlf}$	9.27	3.68	5.17	3.58	8.84	6.01	7.70	4.01	6.03
	(9.76)	(2.43)	(6.61)	(4.39)	(9.34)	(9.99)	(11.61)	(6.17)	(5.64)
$\Delta \widetilde{nam}$	-1.03	0.42	7.49	28.75	9.71	-2.56	7.38	0.67	10.91
	(-0.29)	(0.09)	(1.94)	(2.21)	(2.81)	(-0.86)	(3.27)	(0.35)	(4.45)
intercept	0.03	-0.19	-0.02	0.06	-0.01	-0.05	0.02	-0.07	0.06
	(0.41)	(-1.43)	(-0.27)	(0.91)	(-0.12)	(-0.88)	(0.46)	(-0.99)	(0.76)
Obs.	573	250	573	573	573	573	573	573	569
R^2 in %	16.19	2.72	8.17	7.33	14.02	12.51	19.53	6.82	11.53

Table C.11: Predicting betas using net futures positions – FOMC announcement weeks

This table reports the results for pooled panel regressions of the realized correlations and betas of exchange rates with the S&P and the VIX on the scaled level of net positions of hedge funds ("Leveraged Funds", denoted by $\widetilde{nlf}$). The realized betas and correlations are computed using closing prices for the 5 trading days following date t. The data form an unbalanced panel for the post-crisis 2010-2017. 58 of the 388 weeks in that period include a scheduled FOMC announcement, and the "Announcement" subsample contains 504 currency-week observations.I run the following predictive regressions for $y_{i,t\to t+1} = \{\rho^{SPX}_{i,t\to t+1}, \beta^{MKT}_{i,t\to t+1}, \rho^{VIX}_{i,t\to t+1}\}$:

$$y_{i,t\to t+1} = \alpha_i + \eta\,\widetilde{nlf}_{i,t} + \delta\, r_{i,t} + \phi\, fd^w_{i,t} + \varepsilon_{i,t+1} \tag{3.9}$$

$fd^w_{i,t}$ denotes the 1-week forward discount of currency i versus the dollar and $r_{i,t}$ denotes the currency return of currency i from week $t-1$ to week t. Standard errors are clustered at the currency level and t-statistics reported in parentheses. The panel entitled "Announcement" reports the results over all 58 weeks from 2010-2017 that contained a scheduled FOMC announcement, while the "Non-Announcement" panel reports the results for the remaining 330 weeks.

	Non-Announcement			*Announcement*		
	ρ^{SPX}	β^{SPX}	ρ^{VIX}	ρ^{SPX}	β^{SPX}	ρ^{VIX}
$\widetilde{nlf}$	0.334	0.247	-0.257	0.208	0.029	-0.247
	(4.50)	(3.28)	(-4.18)	(2.25)	(0.22)	(-2.31)
fd^w	0.029	0.090	-0.098	-0.247	-0.348	0.169
	(0.32)	(1.00)	(-1.08)	(-1.64)	(-2.33)	(1.40)
$r_{i,t-1}$	-1.153	0.710	1.965	-0.155	-0.577	1.737
	(-1.90)	(1.02)	(2.91)	(-0.09)	(-0.30)	(0.93)
R^2 in (%)	2.95	0.98	1.93	1.23	0.25	2.06
Currencies	9	9	9	9	9	9
Obs.	2867	2867	2867	504	504	504

C.2 Additional Figures

Figure C.1: Futures positions (in 000's of contracts) by currency

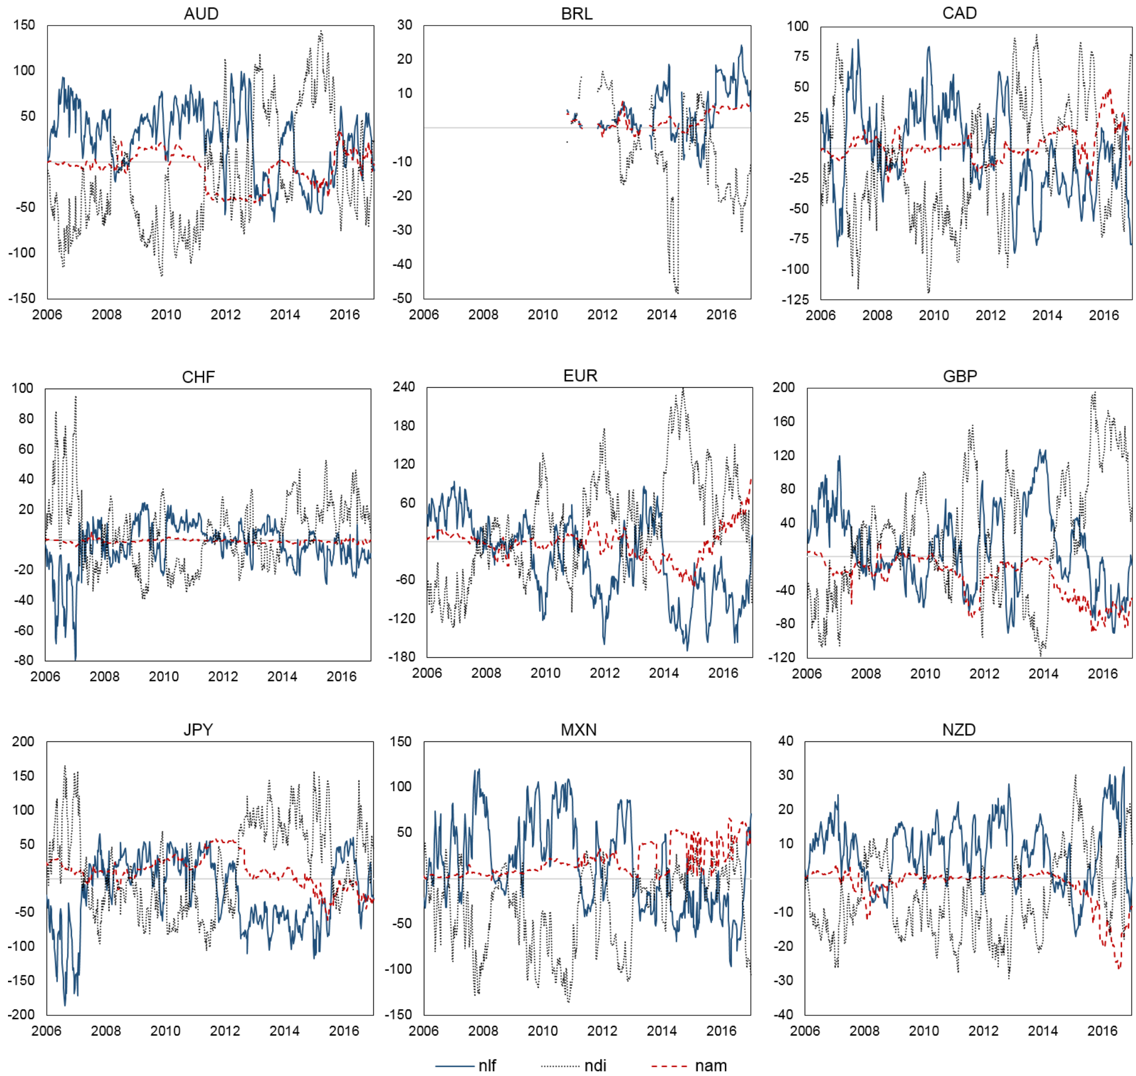

Hedge funds (*nlf*, solid, blue), intermediaries (*ndi*, dotted, black), and institutional investors (*nam*, dashed, red).

Figure C.2: Hedge fund positions and forward discounts

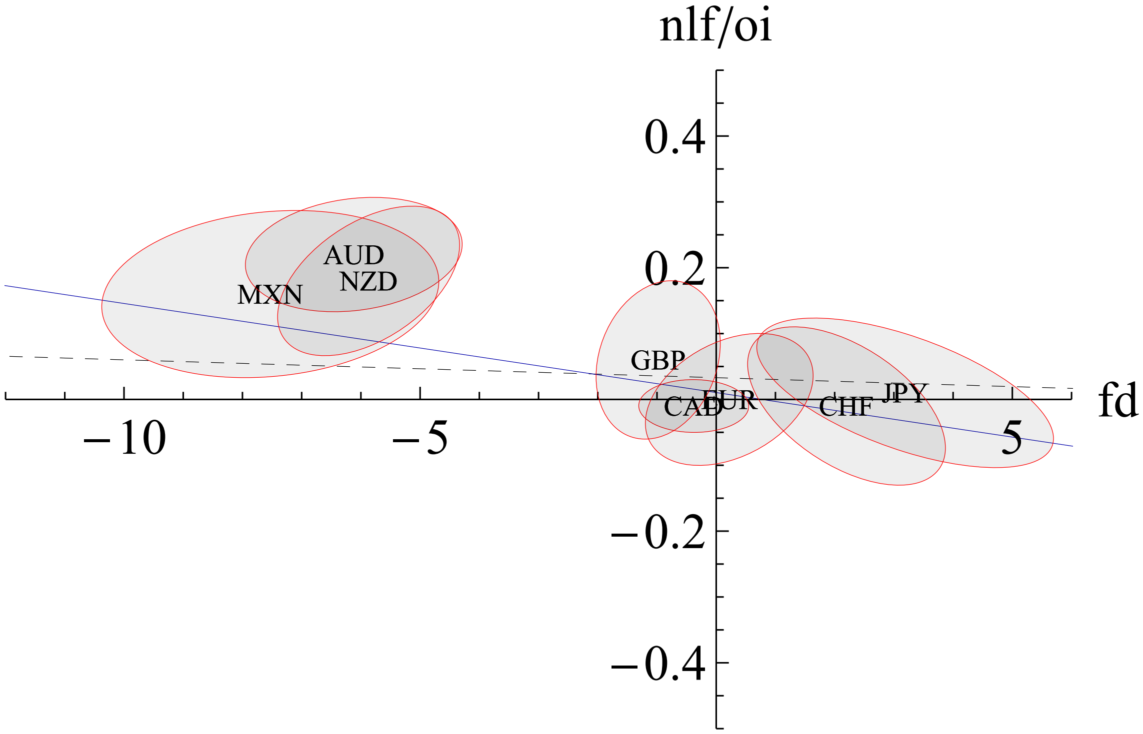

The centre of each confidence ellipse represents a currency's average forward discount and net hedge fund position. In population, each ellipse would contain 20% of its currency's data points under Normality. The orientation of each ellipse reflects the within-currency correlation between forward discount versus the US dollar and hedge fund positions, while the ellipse's size reflects their volatilities. BRL is omitted from the plot for readability, due to its large average forward discount of -25bps. The lines are derived from a univariate OLS regression of $\widetilde{nlf}$ on fd^w, including BRL (solid blue line) and excluding BRL (dashed black line), respectively.