BANK OF ENGLAND

Staff Working Paper No. 822

Attention to the tail(s): global financial conditions and exchange rate risks

Fernando Eguren-Martin and Andrej Sokol

September 2019

BANK OF ENGLAND

Staff Working Paper No. 822

Attention to the tail(s): global financial conditions and exchange rate risks

Fernando Eguren-Martin[(1)] and Andrej Sokol[(2)]

Abstract

We document how the entire distribution of exchange rate returns responds to changes in global financial conditions. We measure global financial conditions as the common component of country-specific financial condition indices, computed consistently across a large panel of developed and emerging economies. Based on quantile regression results, we provide a characterisation and ranking of the tail behaviour of a large sample of currencies in response to a tightening of global financial conditions, corroborating some of the prevailing narratives about safe haven and risky currencies. We then carry out a portfolio sorting exercise to identify the macroeconomic fundamentals associated with such different tail behaviour, and find that currency portfolios sorted on the basis of relative interest rates, current account balances and levels of international reserves display a higher likelihood of large losses in response to a tightening of global financial conditions.

Key words: Exchange rates, tail risks, financial conditions indices, global financial cycle, quantile regression.

JEL classification: F31, G15.

(1) Bank of England and University of Oxford. Email: fernando.egurenmartin@bankofengland.co.uk
(2) European Central Bank, Bank of England and CfM. Email: andrej.sokol@ecb.europa.eu

We would like to thank Gino Cenedese, Georgios Georgiadis, Simon Lloyd, Daniele Massacci, Angelo Ranaldo, Lucio Sarno, Livio Stracca, Giorgio Valente, seminar participants at Oxford University and the European Central Bank, conference participants at IAAE 2019 and one anonymous referee for useful comments and discussions. The views expressed in this paper are solely those of the authors and should not be taken to represent those of the Bank of England or the European Central Bank. Online appendix available at https://sites.google.com/site/fereguren/research

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Non-technical summary

Recent work in academia and policy institutions has emphasised the importance of the global financial cycle, a shorthand for the often strong positive co-movement of financial variables observed across countries. However, there is an asset class that stands out from the rest: exchange rates. Being relative prices, the scope for exchange rates to positively co-move at the global level is limited by construction. Against this backdrop, this paper uses the quantile regression methodology to characterise the response of exchange rate returns, with a particular emphasis on the tails of their distributions, to changes in global financial conditions, a proxy for the global financial cycle.

We start by developing a measure of global financial conditions based on capturing co-movement across a range of country-specific, multi-asset financial condition indices. Armed with this measure of global financial conditions, we use quantile regression to analyse the response of the entire distribution of an exchange rate returns series to a tightening in global financial conditions. By quantifying the shifts, in particular in the tails, of such conditional distributions, we are able to characterise currencies safe haven and risky-type dynamics in a richer way than was done in earlier studies, which focused overwhelmingly on average returns.

We find that the shifts in the distributions of a large set of currencies in the event of a tightening in global financial conditions mostly matches prevailing market narratives: for example, when financial conditions tighten, the risk of a sharp appreciation in the Japanese yen (typically regarded as a safe haven currency) increases significantly; in contrast, the Australian dollar (a risky currency) displays significantly larger crash risk in the same situation. The case of the euro is intermediate, as we find that both tails become fatter. An important feature of our approach is that we can quantify such dynamics, for example by assigning probabilities to the full range of exchange rate moves.

Our analysis also identifies a series of currency-specific risk factors associated with increased chances of a sharp depreciation in the event of a tightening in global financial conditions. We find that currencies of countries with (i) high interest rates, (ii) large current account deficits and (iii) low levels of international reserves are particularly at risk of suffering sharp depreciations in the event of a tightening in global financial conditions.

1 Introduction

Recent years have witnessed a heated debate about the extent and interpretation of the global co-movement of financial variables. Proponents of a so-called 'global financial cycle', beginning with Rey (2013), argue that the observed cross-country co-movement in asset prices cannot be fully explained by co-movement in real variables alone, and therefore must have a finance-specific component to it, such as some measure of global risk aversion. Others, such as Cerutti et al. (2017), argue against the very notion of a 'global financial cycle'.

Within that debate, there is an asset class that stands out from the rest: exchange rates. Being relative prices, the scope for them to co-move at the global level is limited by construction. Moreover, the relationship between exchange rate movements and overall financial conditions in a country is not a priori obvious: a given exchange rate move can 'tighten' access to finance for some agents in the economy, while 'loosening' access for others.

These considerations, and assuming the existence of a global component of financial conditions, motivate the main question addressed in this paper, namely, how different exchange rates co-move with that global component of financial conditions. Unlike most of the existing literature, we study the behaviour of the entire distribution of different currencies' returns in the face of changes in global financial conditions, with a particular focus on the tails, that is, on the likelihood of a sharp appreciation or depreciation. This way we can provide both a characterisation and quantification of the risks facing particular currencies under different scenarios for global financial conditions.

We exploit several novel quantification possibilities afforded by quantile regression to make two main contributions. First, we document the tail behaviour of exchange rate returns across a broad range of currencies. We show that simple quantile regressions can deliver marked improvements in fit in the tail regions even where standard R^2 measures are low, and then rank currencies according to how their left (depreciation) and right (appreciation) tails respond to a tightening of global financial conditions. Our exercise corroborates some of the prevailing narratives about safe haven and risky currencies, but also offers interesting new insights, including on the quantitative side. For example, our method allows to quantify the probabilities of a given exchange rate move in the face of observed or projected global financial conditions.

Second, we identify potential risk factors associated with the different tail behaviour of currencies. In order to do so, we conduct portfolio sorting exercises based on several macroeconomic fundamentals, and then study the responses of the resulting returns series to a tightening of global financial conditions. We find that the portfolios sorted on the basis of relative interest rates, current account balances and levels of international reserves display a higher likelihood of large losses in response to a tightening of global financial conditions. From a policy perspective, these results provide an empirical motivation for a close scrutiny of these variables when assessing countries' financial stability prospects.

1.1 Related literature

This paper is related to several literature strands. First, and most directly, it is related to papers that study the occurrence of tail events in exchange rate markets. On the negative returns side, there is a large literature that documents the existence of 'crash' or 'disaster' risk in popular FX strategies. Brunnermeier et al. (2009) find that carry trade strategies perform particularly poorly during periods of heightened risk aversion (as proxied by the VIX index), while Menkhoff et al. (2012) show similar results but focusing on periods of high FX volatility. Relatedly, Farhi and Gabaix (2016) and Farhi et al. (2009) study disaster risk embedded in option prices.

In principle, the poor performance of carry trades could be the result of both a sharp depreciation of high-interest-rate currencies and/or a sharp appreciation of low-interest-rate currencies. In that vein, some papers study the dynamics of particular currencies, namely those usually labelled as safe havens, which, according to market narratives, tend to appreciate sharply during periods of high risk aversion. Ranaldo and Soderlind (2010) and Habib and Stracca (2012) study the safe haven property of a series of currencies, and do indeed find robust evidence of substantial appreciation during periods of market stress. Fratzscher (2009) also looks at the dynamics of individual currencies under stress conditions in the context of the global financial crisis.

A common feature of these papers is that their empirical strategies focus on the mean returns of currencies or trading strategies. In contrast, our approach allows a detailed study of the entire distribution of exchange rate returns, including the tails, which are at the core of our analysis. Moreover, we propose a novel way of characterising periods of heightened (global) risk aversion, avoiding popular but imperfect proxies (e.g. the VIX index), or FX-based proxies which can become somewhat circular (e.g. FX volatility). Also, we look at a long list of individual exchange rates, facilitating a direct analysis of particular currencies.

The second literature strand the paper is related to is more methodological, and has to do with the recent surge in popularity of quantile regression, originally introduced by Koenker and Bassett (1978) in both macroeconomics and finance. Some recent contributions include Cenedese et al. (2014) for exchange rates, Gaglianone and Lima (2012) for unemployment, Korobilis (2017) for inflation and Crump et al. (2018) for equity returns. Most closely related to our study, Adrian et al. (2019) rely on quantile regression to characterise the tails of the GDP growth distribution conditional on domestic financial conditions.[1] We build on similar ideas, but focus instead on the distribution of exchange rate returns conditional on *global* financial conditions.

The last strand of literature we draw and build on has to do with measurement of financial conditions. We follow Arregui et al. (2018) in constructing country-specific financial condition indices that exploit a broad set of market-based indicators for a large panel of countries, which then allows us to extract a *global* financial conditions index. This measurement exercise is related to earlier attempts to characterise a 'global financial cycle', most notably by Miranda-Agrippino and Rey (2015),[2] but in the finance literature it also overlaps with various proposals to measure global risk aversion and other factors commonly used to price exchange rate rates (see e.g. Menkhoff et al. (2012) and Lustig et al. (2011)).

The rest of the paper is organised as follows: in Section 2, we describe our measure of global financial conditions. In Section 3 we discuss quantile regressions of effective exchange rate returns on global financial conditions. In Section 4 we introduce currency portfolio sorting based on macroeconomic fundamentals and identify potential factors associated with currencies' differential tail behaviour. In Section 5 we run a series of robustness checks on our results. In Section 6 we conclude, while the Appendix includes details about our data and methodology. An online appendix provides additional results and robustness checks.

2 Measuring global financial conditions

The existence of a global factor in financial conditions has been widely debated in economics over recent years.[3] Beginning with Rey (2013), a series of papers have emphasised (and measured) a strong co-movement in financial variables across countries (among others, see Bruno and Shin (2014), Cesa-Bianchi et al. (2018a,b), Ha et al. (2018)). These papers have suggested that this co-movement in financial conditions went beyond a reflection of co-movement in macroeconomic indicators, and hence was at least partly driven by a specific global factor in financial variables, such as risk appetite. The standard approach has been to measure common variation in a set of asset prices and/or credit quantities, interpreting

[1] Relatedly, Adrian et al. (2018) explore the term-structure of this relationship.

[2] See Cerutti et al. (2017) for a contrarian view on the existence of a global financial cycle. Also see Drehmann et al. (2012) for a characterisation of a more medium-term (domestic) financial cycle.

[3] This factor has typically been referred to as 'the global financial cycle'.

the result as an indicator of the ease at which finance could be accessed at a given time in a given country (see, for example, Miranda-Agrippino and Rey (2015)).

Existing measures of global financial conditions typically suffer from two shortcomings. First, the breadth of financial series considered tends to be limited, and usually skewed towards equity markets (as, for example, in Miranda-Agrippino and Rey (2015)). Second, the geographical coverage tends to be limited to advanced economies (e.g. Ha et al. (2018)) and, in some cases, a handful of emerging countries. Both of these limitations are due to data availability constraints: it is not straightforward to construct a panel dataset spanning a broad set of financial indicators for a large cross-section of countries.

In order to overcome these limitations we follow Arregui et al. (2018) and construct a panel dataset covering a broad set of monthly financial indicators for 43 countries from January 1991 to June 2018. The financial series included are as follows: term, sovereign, interbank and corporate spreads, long-term interest rates, equity returns and volatility, relative market capitalisation of the financial sector, house prices and credit growth.[4] To extract country-specific summary measures of financial conditions, we follow Koop and Korobilis (2014) and estimate factor models which allow for time variation in the parameters and attempt to 'clean' financial conditions from changes that reflect a response to macro-economic news (proxied by industrial production and CPI inflation).[5]

Armed with a set country-specific financial condition indices, we extract a global component by taking the cross-sectional mean.[6] The share of variance of individual country FCIs explained by this global component varies in the cross section, but averages around 30%. It is worth noting that this figure goes up to above 60% for several countries, including financial centres such as the US or the UK (all figures are reported in the Online Appendix). In what follows we take this series as our measure of global financial conditions.

Figure 1 shows the evolution of our measure over the last 30 years, which is broadly in line with prevailing narratives: for example, global financial conditions tighten sharply around the collapse of Lehman Brothers in 2008, and during the euro area crisis of 2010-2011. Importantly, our measure co-moves positively, but far from one-to-one, with other

[4] A detailed description of the variables used and corresponding data sources can be found in Appendix A.

[5] The objective is to obtain financial condition indices that seek to reflect 'pure' changes in financial conditions (e.g. shifts in risk aversion), in contrast to market changes which reflect news to the economic outlook. Note, however, that given the forward-looking nature of asset prices, it could still be the case that they respond to news to *expected* macro-economic developments, not captured in the contemporaneous series used in our approach.

[6] Taking instead the first principal component yields an almost indistinguishably similar series.

Figure 1 Global Financial Conditions Index, 1991-2017.

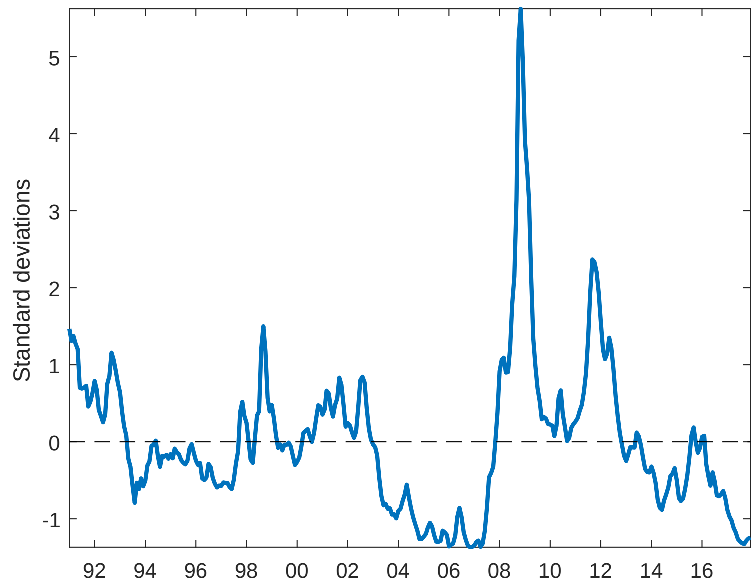

Note: Index in deviations from its historical mean. Higher values signal tighter financial conditions.

widely used US-centric measures such as the VIX index or the S&P index, and with the estimated factor in Miranda-Agrippino and Rey (2015).[7] [8]

In the context of this paper, it is also interesting to note that our measure of global financial conditions displays a positive correlation with factors widely used to price exchange rates. The correlation with the FX volatility factor from Menkhoff et al. (2012) is 0.7, while the correlation with (the negative of) the dollar and HML factors (originally proposed by Lustig et al. (2011)) is approximately 0.2. This is particularly interesting because these factors are computed using the very same exchange rate data that are then priced with them, while our measure does not directly contain any FX data at all.

[7] The correlation of our index of global financial conditions and (the negative of) the S&P index is approximately 0.3, while the correlation with the VIX index is approximately 0.8. The correlation with the global factor in Miranda-Agrippino and Rey (2015) is approximately 0.6.

[8] See Section 5 for a robustness check of our baseline exercise considering alternative measures of global financial conditions.

3 Quantifying exchange rate tail risks with quantile regression

As discussed in Section 2, asset prices tend to display a high degree of co-movement across countries. However, exchange rates are somewhat special. Being relative prices, the pattern and extent of their co-movement is more constrained than for other assets. This feature of exchange rates is the departing point of our analysis: we want to understand how different exchange rates co-move with changes in global financial conditions, and the underlying country-specific characteristics that are associated with such dynamics.

Our focus is on the whole distribution of exchange rate returns, and in particular on tail events. Specifically, we study how the likelihood of sharp exchange rate movements (in either direction) is affected by global financial conditions. To this end, we rely on quantile regression (Koenker and Bassett, 1978). Unlike standard regression, which provides an estimate of the conditional mean of a variable of interest given a set of explanatory variables, quantile regression allows to model the entire conditional distribution of a dependent variable given a set of covariates. This allows to capture features that are lost when only focussing on the average response.

3.1 Specification

Following the (limited) existing literature applying quantile regression to exchange rates (see, for example, Cenedese et al. (2014)), our baseline exercise studies the effect of global financial conditions on the distribution of exchange rate returns. We specify a linear model for their conditional quantiles as follows:

$$Q_{\Delta FX_{t+h}}(\tau|X_t) = \alpha_h(\tau) + \beta_h(\tau)GFC_t \tag{1}$$

where ΔFX_{t+h} is the monthly log change in the nominal effective exchange rate h months ahead and GFC_t is our measure of global financial conditions.[9] Function Q computes quantiles τ of the distribution of ΔFX_{t+h} given a set of covariates X_t (in this case, GFC_t and a constant). Appendix B discusses technical details.

We estimate this equation on a currency-by-currency basis for a panel of 61 countries from January 1991 to June 2018. The full list of currencies can be found in Appendix A. We

[9]The convention we adopt here is that positive FX changes represent an appreciation, and negative changes a depreciation. This is done in order to facilitate comparison with the next section, which looks at returns.

Figure 2 Impact of global financial conditions on the conditional quantiles of exchange rate returns.

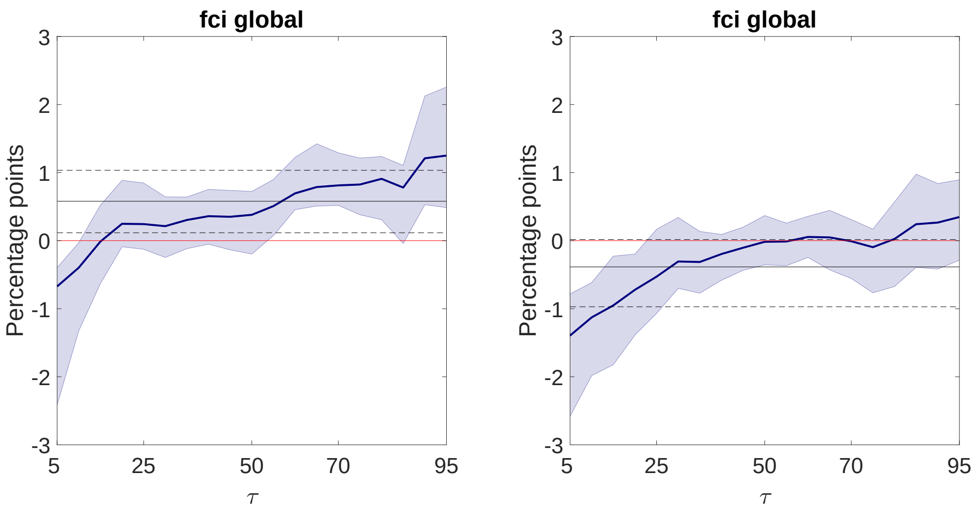

(a) Japan (b) Australia

Note: The blue lines plot the values of $\beta_h(\tau)$ across quantiles, while the black lines show OLS estimates of the same specification. 95% confidence intervals are computed from 1000 overlapping block bootstrap draws.

focus on nominal effective exchange rates in our baseline to identify idiosyncratic dynamics, avoiding potentially US-driven moves of US dollar bilaterals. [10] US dollar moves could still affect global financial conditions (especially given the prominence of the US dollar in global trade and financial markets), but we are interested on the effect of these moves on country-idiosyncratic nominal effective exchange rates.

Figure 2 shows the typical output from such regressions for two currencies, the Japanese yen (JPY) and the Australian dollar (AUD). The prevailing narrative in FX markets places these two currencies at the opposite ends of a spectrum: while the JPY is considered a safe haven (Ranaldo and Soderlind (2010), Habib and Stracca (2012)), which means that it tends to appreciate during periods of increased global risk aversion, the AUD is typically regarded as a risky currency that would instead depreciate in such circumstances.

The two panels show the impact of a one standard deviation change in global financial conditions on different quantiles of the conditional distribution of exchange rate returns over the same month, i.e. for $h = 0$. The Japanese yen exhibits positive coefficients approximately from the 25th quantile on, meaning that most of the conditional distribution shifts to the

[10] See Section 5 for a robustness exercise which re-estimates our baseline specification considering US dollar bilateral exchange rates, and another one that considers excess returns instead of plain exchange rate changes.

Table 1 Goodness of fit measures, selected currencies.

	$R^1(\tau)$					R^2
	0.05	0.25	0.5	0.75	0.95	
Australia	16.5	1.9	0.0	0.2	1.2	3.2
Euro area	1.7	0.4	0.1	0.4	7.5	0.0
Japan	2.8	0.7	1.7	6.5	12.3	5.8
Switzerland	0.1	0.1	2.2	3.7	9.0	2.1
United Kingdom	10.8	2.8	0.5	0.0	0.2	3.6
United States	0.3	0.1	1.9	2.4	9.8	4.7

right in the face of a tightening of global financial conditions (increasing the chances of an appreciation). On the other hand, the coefficients for the Australian dollar are for the most part not statistically different from 0, except for the first quartile, indicating an increased risk of a sharp depreciation. The black lines show coefficient values from simple OLS regressions as a benchmark. Figure 2 gives a sense of the advantages of our approach withe respect to OLS-based alternatives commonly used in the literature: while the OLS coefficients are broadly able to capture the different mean behaviour of the two currencies in the face of the same shock (albeit not significantly so for Australia), it is evident that much useful information for a richer chatacterisation of conditional distributions is discarded by only focussing on the conditional mean.

3.2 Goodness of fit

Another way of comparing insights from quantile regression and OLS is by looking at measures of goodness of fit. We follow Koenker and Machado (1999) and report quantile-specific $R^1(\tau)$ measures for all currencies. Unlike standard R^2 measures, which quantify the relative success of two models for the conditional mean function, and thus provide a global measure of goodness of fit over the entire conditional distribution, $R^1(\tau)$ measures provide information on the relative local success of two models of a conditional quantile function.

$R^1(\tau)$ is defined as

$$R^1(\tau) = 1 - \hat{V}(\tau)/\tilde{V}(\tau) \tag{2}$$

where $\hat{V}(\tau)$ denotes the sum of weighted absolute residuals of model (1) and $\tilde{V}(\tau)$ the sum of weighted absolute residuals of a model consisting only of a constant (which provides an estimate of the unconditional quantile τ).[11] The interpretation is thus analogous to that

[11] As explained in Appendix B, $\hat{V}(\tau)$ and $\tilde{V}(\tau)$ are simply the objective functions of the respective quantile regression problems, which take the form of weighted sums of absolute residuals, evaluated at the optimum.

of standard R^2: $R^1(\tau)$ expresses the improvement in fit, in terms of the relevant criterion function, obtained by adding covariates to the model.

Table 1 reports both R^2 and $R^1(\tau)$ measures for selected currencies. The first thing to note is that the overall improvement in fit from including our measure of global financial conditions as a covariate, proxied by the R^2 of a standard OLS regression, varies across countries, but is typically limited.

On the other hand, as far as $R^1(\tau)$ measures are concerned, a robust pattern seems to hold across countries (see Online Appendix for the full panel), namely, that the goodness of fit tends to generally improve in the tails, highlighting the information content of financial conditions that can be lost by exclusively focusing on the mean of exchange rates distribution. This is an important dimension along which our approach is superior in characterising exchange rate behaviour compared to mean-based approaches. It is also worth noting that the improvements tend to be concentrated in one tail, and the largest gains tend to accrue to the most extreme percentiles, so the 95^{th} for the Japanese yen, Swiss franc, US dollar and the euro, and the 5^{th} for the Australian dollar and UK pound.

3.3 Summary measures of tail behaviour

The information conveyed by the quantile-specific slope coefficients $\beta_h(\tau)$ (as shown in Figure 2) can be summarised visually by studying their effect on fitted probability density functions. In the same spirit as Adrian et al. (2019), who fit skew-t distributions to the predictive quantiles of GDP growth, we fit non-parametric density functions to the quantiles of exchange rate returns conditional on different values of global financial conditions.[12] The reason for using a non-parametric distribution is to allow for a more nuanced depiction of tail behaviour than is possible by only fitting a few parameters (four in the case of the skew-t) to the estimated conditional quantiles, which is particularly important in our application.

Specifically, we fit non-parametric distributions with Normal kernel ϕ and suitably chosen bandwidth h, whose density is given by

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{\tau=1}^{T} \phi\Big(\frac{x - \hat{q}(\tau)}{h}\Big), \tag{3}$$

to the fitted quantiles $\hat{q}(\tau)$ conditional on the average value of global financial conditions (which is 0), given by $\hat{\alpha}_h(\tau)$, and conditional on a one standard deviation increase in global

[12] A similar approach is followed by Gaglianone and Lima (2012) and Korobilis (2017).

Figure 3 Impact of a tightening of global financial conditions on the conditional distribution of exchange rate returns.

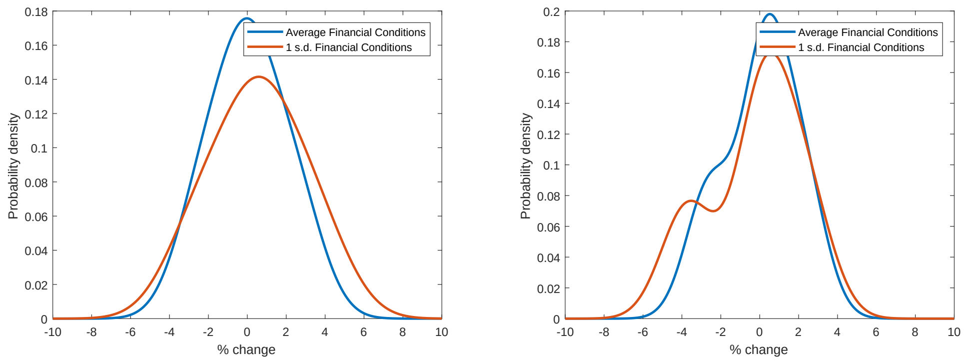

(a) Japan (b) Australia

financial conditions, given by $\hat{\alpha}_h(\tau) + \hat{\beta}_h(\tau)$.

Figure 3 illustrates the changes induced by a one standard deviation increase in global financial conditions on the conditional densities of the same two currencies analysed before. Our exercise does corroborate their usual characterisation: in the face of a tightening of global financial conditions, the right (appreciation) tail of the distribution of the Japanese yen shifts up significantly (increased chances of a sharp appreciation), while the left (depreciation) tail of the distribution of the Australian dollar shifts down (increased chances of a sharp depreciation). These 'fatter' tails not only confirm market narratives but, most importantly, also provide a quantification of the shift in risks. That is, our model can assign probabilities to each possible change in the value a particular currency, and make those probabilities a function of global financial conditions. This is particularly useful from a risk-monitoring perspective in the case of the tracking of countries' macroeconomic conditions, and from a risk-management perspective when thinking of exchange rates as asset prices underlying investment strategies.

To compare such heterogeneous tail behaviour across currencies we compute measures of divergence between the two distributions. In particular, we use a version of the Kullback-Leibler divergence, also known as relative entropy, to quantify the 'shifts' induced in the tail regions by a tightening of global financial conditions.[13] Given a fitted distribution $\hat{g}(x)$ conditional on average global financial conditions and another, $\hat{f}(x)$, conditional on a 1 standard deviation increase in global financial conditions, we compute downside and upside

[13]This is similar in spirit to the quantification of upside and downside risks in Adrian et al. (2019).

Figure 4 Downside and upside entropy measures of conditional exchange rate returns, selected currencies.

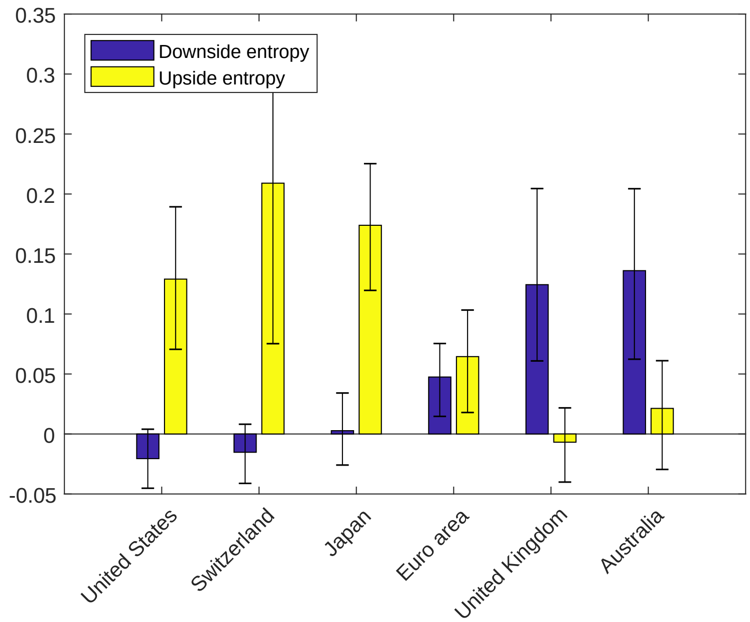

Note: 68% confidence intervals are computed from 1000 overlapping block bootstrap draws.

(relative) entropy outside of the interquartile range of $\hat{g}(x)$ as

$$\mathcal{L}^D = \int_{-\infty}^{\hat{G}^{-1}(0.25)} \log\left(\frac{\hat{f}(x)}{\hat{g}(x)}\right) \hat{f}(x)dx \tag{4}$$

$$\mathcal{L}^U = \int_{\hat{G}^{-1}(0.75)}^{\infty} \log\left(\frac{\hat{f}(x)}{\hat{g}(x)}\right) \hat{f}(x)dx. \tag{5}$$

Intuitively, downside and upside entropy measure the additional probability mass assigned to tail events when there is a tightening of global financial conditions. For safe haven currencies, upside entropy should be positive (denoting an increased probability of a large appreciation), whereas for risky currencies, downside entropy should be high.

Figure 4 shows the results for the same selection of currencies AS Table 1.[14] The ranking in terms of tail behaviour once again broadly confirms prevailing narratives: typical safe haven currencies such as the Japanese yen, the US dollar and the Swiss franc exhibit high upside entropy but hardly any downside entropy, whereas risky currencies such as the Australian dollar tend to exhibit a higher downside entropy. The case of the euro is somewhat interesting

[14] See Section 2 of the Online Appendix for the full sample of currencies.

Table 2 Changes in appreciation and depreciation probabilities due to a tightening of global financial conditions, selected currencies.

	Depreciation probability Δ			Appreciation probability Δ		
	5-7.5%	2.5-5%	0-2.5%	0-2.5%	2.5-5%	5-7.5%
Australia	3.6	3.9	-5.7	-4.1	1.8	0.4
Euro area	0.0	1.3	1.5	-8.2	5.5	0.0
Japan	1.1	0.3	-8.2	-3.2	6.1	3.7
Switzerland	0.0	0.4	-7.2	-1.7	8.4	0.0
United Kingdom	0.0	5.8	-1.5	-5.1	0.8	0.0
United States	0.0	0.4	-6.3	1.5	4.5	0.0

in that it exhibits similar degrees of both upside and downside entropy, meaning that both tails become fatter in response to a tightening of global financial conditions.

To provide a more tangible measure, Table 2 also reports changes in appreciation and depreciation probabilities induced by a tightening of global financial conditions, that is, the integral of $\hat{f}(x) - \hat{g}(x)$ over different ranges. Very large swings in returns in either direction (>7.5%) are never assigned very high probabilities, whereas appreciations or depreciations between 2.5% and 7.5% tend to be assigned higher chances, mostly in accord with usual currency characterisations. In the case of the euro, the most notable change is a reallocation of probability mass between the first two appreciation buckets. These results thus complement and qualify the information about local fit from $R^1(\tau)$ measures, and give a much more nuanced depiction of different currencies' tail behaviour.

In the next Section we turn to analysing the underlying country characteristics that are associated with the different responses of currencies' distributions to changes in global financial conditions.

4 Identifying risk factors: a portfolio sorting approach

What country characteristics are associated with the different exchange rate dynamics documented in the previous section? Or, in other words, are there any risk factors associated with specific tail behaviours that policymakers and investors should keep track of? To answer this question we rely on portfolio sorting exercises, popular in the FX and equity pricing literatures. In Subsection 4.1 we explain the rationale and mechanics behind our portfolio sorting exercises, then in Subsection 4.2 we identify risk factors by studying the returns of our portfolios following changes in global financial conditions.[15]

4.1 Portfolio sorting

Identifying country characteristics associated with the individual features of exchange rate returns distributions documented in Section 3 is challenging: for each country, conditional distributions are identified from the whole (time series) sample and offer a single summary statistic. However, it is likely that the risk factors associated with such dynamics change over time. For example, it would not be appropriate to try to associate a certain conditional exchange rate distribution to average fiscal deficits over 25 years, as this statistic is likely to hide significant variation over the sample. To address such concerns, we need to introduce a degree of time-variation in our analysis. To do so, we follow Cenedese (2015) and conduct portfolio sorting exercises, widely used in the equity and FX pricing literatures.

We start from a series of candidate variables that have been identified in the literature as being associated with particular reactions of exchange rates to changes in global financial conditions: interest rate differentials (with respect to the US), current account balances, fiscal balances and levels of international reserves.[16] A series of empirical papers have analysed the importance of these risk factors for average exchange rate dynamics: Brunnermeier et al. (2009), Lustig et al. (2011) and Menkhoff et al. (2012) study the risk features of high interest rate currencies, Della-Corte et al. (2016) that of currencies of countries with large external imbalances, while Fratzscher (2009) and Habib and Stracca (2012) assess the relevance of a wide range of variables, including fiscal balances and international reserves. In turn, these studies are grounded on a rich history of theoretical work linking these macroeconomic variables and exchange rate dynamics.[17]

We consider each of the candidate risk factor variables in turn and, at each point in time throughout our sample, begin by ranking countries according to the values they display for the variable under consideration. So for example, when working with current account balances,

[15]Throughout the current exercise we do not consider euro-zone currencies (for the entire sample) given the asymmetry between national/domestic risk factors and a zone-wide currency which value countries only influence partially. We drop the entire time-series to avoid sample selection issues.

[16]Exact definitions and data sources can be found in Appendix A. Interest rate differentials are implied from FX forward contracts. Note that recent CIP deviations mean that there is measurement error in this quantification of interest rate differentials. For simplicity, and given that forward discounts are computed based on contracts on US dollar bilateral exchange rates, we rely on interest rate differentials with respect to the US despite using nominal effective exchange rates.

[17]See, for example, Krugman (1979), Dornbusch and Fischer (1980), Wijnbergen (1991), Obstfeld and Rogoff (1995) and Gabaix and Maggiori (2015).

Figure 5 Portfolio ranking based on Current Account balance for selected countries.

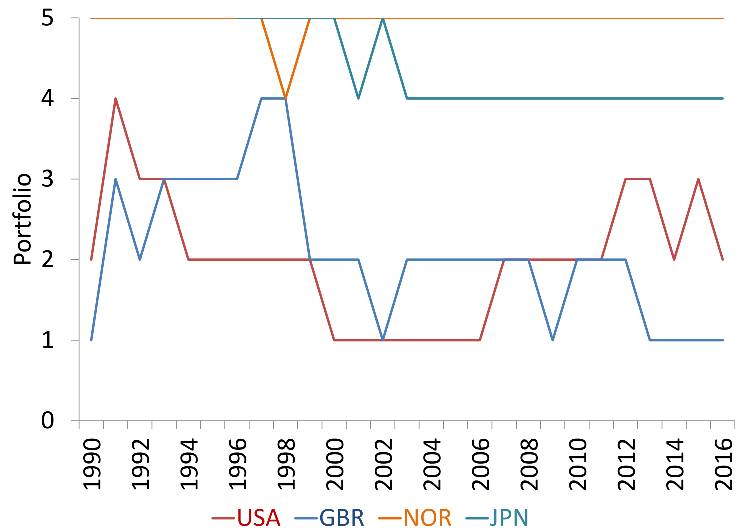

we rank countries from those displaying the highest current account surplus to those with the highest deficit.[18]

We then assign currencies to five portfolios according to this ranking. Continuing the previous example, the first portfolio contains the currencies with the largest current account deficits, while the fifth portfolio contains the currencies with the largest current account surpluses. As an example, Figure 5 plots such portfolio assignment for a selection of currencies over time. Finally, we compute the return of each portfolio over the month as the equally-weighted return of its component currencies, and compute the relative return between the first (most exposed to the risk factor) and fifth (least exposed to the risk factor) portfolios (i.e. of the high-minus-low portfolio).[19] This relative return is a proxy for the (FX) market compensation for exposure to the risk factor under consideration, and constitutes our variable of interest.[20]

The advantage of such a portfolio sorting approach is that it introduces time variation

[18]We rebalance portfolios annually given the limited availability of data for the sorting variables.

[19]We use pure FX-driven returns, i.e. log exchange rate changes. The convention matches that of the previous section, so that a positive return corresponds to an appreciation. See Section 5 for an alternative specification which considers excess returns.

[20]See Section 5 for a robustness exercise which considers an alternative version of the portfolio sorting in which the assignment of currencies to risk buckets is performed according to lagged values of the risk variables under consideration.

Table 3 Goodness of fit measures for the relative returns of sorted portfolios.

	$R^1(\tau)$					R^2
	0.05	0.25	0.5	0.75	0.95	
Carry	6.7	4.0	2.4	0.4	0.0	0.0
Current Account	7.8	4.0	3.9	1.2	1.4	0.0
Fiscal	0.5	0.1	0.0	0.7	1.2	0.0
Reserves	2.3	2.2	0.8	0.3	1.6	0.0

in the exposure to risk factors, which could be associated with particular exchange rate dynamics. This is achieved by allowing countries to have different levels of exposure at different points in time. For example, country A could exhibit a large current account surplus in period t and a large deficit in period $t+k$. In this situation, the return of country A's currency in period t will be assigned to the portfolio comprising surplus countries, while the return in period $t+k$ will be assigned to the portfolio comprising deficit countries. By doing this, our estimates do not depend on the whole time series of returns of a particular country or group of countries, but instead returns are computed dynamically depending on where countries lie in the ranking of risk factors.[21]

We conduct the exercise described above separately for each of our four risk factors, and then analyse how exposure to each of them is associated with differential responses of the tails of relative return distributions to changes in global financial conditions.

4.2 Risk factors and global financial conditions

After computing relative portfolio returns as described in Subsection 4.1, we proceed to analyse how their distributions are affected by shifts in global financial conditions, as in Section 3. In line with the previous section, we first first estimate conditional quantile functions for each relative portfolio returns series, and then fit two empirical distributions: one conditional on average global financial conditions, and another conditional on a one standard deviation tightening of global financial conditions. That is, we want to know the distribution of potential returns of each of the strategies considered (as reflected in the behaviour of the described relative returns) under both 'normal' and 'tight' global financial conditions. If the factors considered were true risk factors, we would expect the distribution of their potential relative returns to exhibit 'fatter' left tails under a tightening of global financial conditions

[21] In practice, these portfolios are moderately stable but not constant: currencies remain in their most common portfolio throughout 55% of the sample on average. If we consider the two most common portfolios, this number goes up to 80%.

Figure 6 Impact of a tightening of global financial conditions on the conditional distribution of relative portfolio returns.

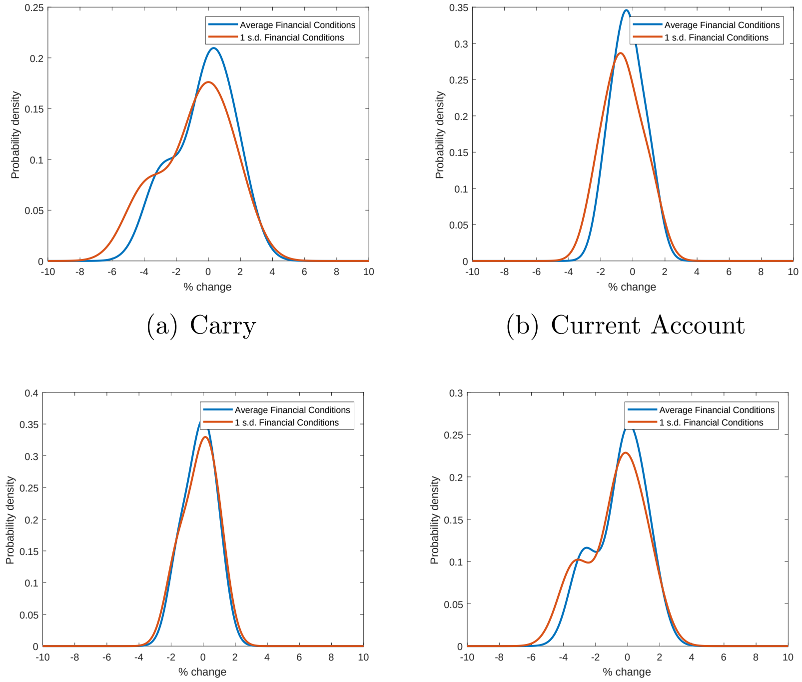

(a) Carry

(b) Current Account

(c) Fiscal

(d) Reserves

(that is, we would expect a larger likelihood of currencies exposed to that factor depreciating sharply). As before, we report goodness-of-fit measures for the quantile regressions, relative entropy measures as well as changes to probabilities of different outcomes.

Table 3 shows that the goodness of fit of quantile regressions of relative portfolio returns on global financial conditions improves in the tails, even when standard R^2 measures are very close to 0, in line with our findings for individual currencies. More specifically, it is the left tails (negative returns) that display the best fit, which suggests that global financial conditions are particularly useful for understanding the 'crash risk' of such strategies.[22] The gains are most pronounced for portfolios sorted by interest rate differentials (carry) and current account balances, whereas there's virtually no improvement in fit for portfolios sorted by fiscal balances. Portfolios sorted according to the level of international reserves are an intermediate case.

As for the shape of the distribution of conditional returns in the face of tighter global

[22]The fact that most of the action is concentrated in the left tails of the conditional distributions is to be expected if the factors considered are indeed risk factors associated with negative conditional returns of highly exposed currencies and conditional excess returns of currencies with low exposures.

Table 4 Changes in probabilities of relative returns in response to a tightening of global financial conditions, sorted portfolios.

	Loss probability change					Gain probability change		
	>10%	7.5-10%	5-7.5%	2.5-5%	0-2.5%	0-2.5%	2.5-5%	>5%
Carry	0.0	0.1	4.1	3.7	-0.4	-7.7	0.0	0.1
Current account	0.0	0.0	0.0	6.1	-0.1	-6.5	0.5	0.0
Fiscal	0.0	0.0	0.0	1.2	-3.1	1.7	0.2	0.0
Reserves	1.7	0.0	1.1	4.3	-0.6	-5.5	0.7	0.0

financial conditions, Figure 6 shows that currencies of countries with relatively high interest rates, large current account deficits and low levels of international reserves display a higher likelihood of experiencing a sharp depreciation.[23] This is also reflected in positive downside entropies (Figure 7) and increases in the chances of negative relative returns (Table 4), in line with our priors. The results for currencies of countries with high fiscal deficits are less clear-cut, as there is a very minor increase in both downside and upside entropies. [24] These findings are consistent with (mean-based) results in Brunnermeier et al. (2009), Lustig et al. (2011), Menkhoff et al. (2012), Della-Corte et al. (2016), Fratzscher (2009) and Habib and Stracca (2012).

In sum, in this section we show that global financial conditions contain useful information for characterising the relative returns distribution of currencies exposed to a series of intuitive risk factors. This is particularly true of (negative) tail outcomes. These insights can be of interest to policymakers assessing the financial stability outlook of countries, and to investors performing risk management calculations of their investment strategies.

5 Robustness

In this section we list a series of robustness checks on our baseline results along three dimensions. First, we explore alternative exchange rate measures, namely US dollar bilaterals and currencies' excess returns (net of interest rate differentials). Second, we modify our exercise in Section 4 by sorting portfolios using lagged values of the sorting variables in order to rule out the possibility of currency rankings being contemporaneously affected by changes in global financial conditions. Last, we explore a range of alternative proxy measures for

[23]To be precise, these conditional portfolio returns could also be the consequence of opposite dynamics for currencies that rank low in terms of exposure to the risk factors under consideration.

[24]However, this comes on the back of an overall worse fit across quantiles than for the other factors. (Table 3).

Figure 7 Downside and upside entropy measures of conditional relative portfolio returns.

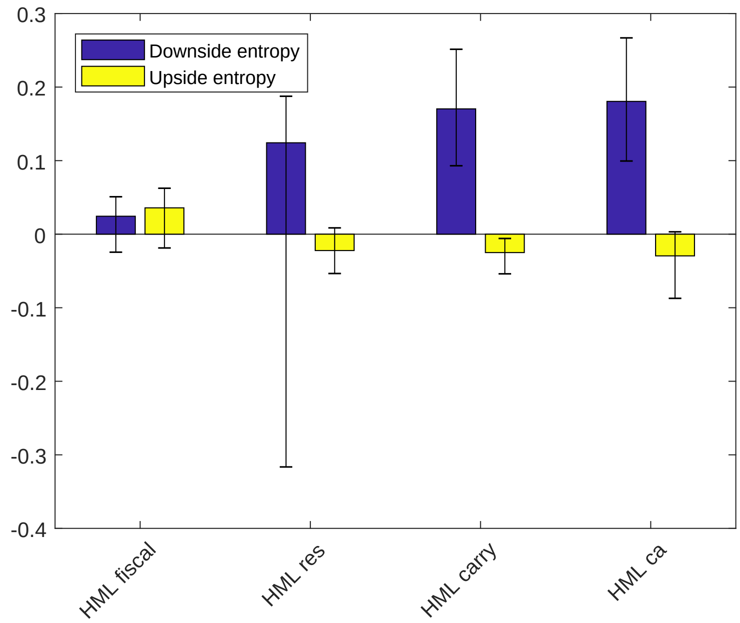

Note: 68% confidence intervals are computed from 1000 overlapping block bootstrap draws.

global financial conditions, including the VIX index. All results are available in the Online Appendix.

5.1 Exchange rate returns measures

Our baseline results in Section 3 are based on nominal effective exchange rates (NEERs). This choice is motivated by the desire to focus on plain exchange rate moves, abstracting from interest rate differentials, and to avoid US-driven, globally synchronised changes in bilateral dollar exchange rates. However, in order to facilitate comparisons with the existing literature, we also report results of exercises that consider alternative choices: NEERs-based excess returns and US dollar bilaterals. Charts and tables can be found in Section 3.1 of the Online Appendix.

Results are broadly unchanged when considering excess returns, which net out interest rate differentials of the currency under consideration vis-à-vis the rest of the world. This holds both for individual currency exercises, and for the portfolio exercises used to identify risk factors. For US dollar bilaterals the changes are also small. Tail behaviour rankings based on relative entropy are virtually unaltered, despite changes in their values in the expected direction: given the high conditional upside entropy of the US dollar NEER itself, the conditional upside entropies of other safe haven currencies become smaller when considering dollar bilaterals, while the downside entropies of risky currencies increase. The portfolio sorting exercise also yields results that are very close to our baseline specification. The only more notable change, also present when using NEER-based excess returns, is that the portfolio exercise based on fiscal balances does show some downside entropy; that is, there is potential evidence of fiscal deficits actually constituting another significant risk factor.

5.2 Portfolio sorting strategy

In our baseline results, the sorting of currencies into portfolios based on the values of risk factors is done contemporaneously. More specifically, the sorting is done at the annual frequency due to data availability, while the conditional returns of the resulting portfolios are measured at monthly frequency. One downside of such strategy is that it is liable to suffering from reverse causality, in that the risk factors could themselves change in response to changes in global financial conditions over the year, in turn affecting the composition of portfolios. A solution is to perform the portfolio sorting considering lagged values of the risk factors, which comes at the cost of potentially using out-of-date data, given the annual rebalancing of portfolios but monthly returns computation. With this caveat in mind, we check the robustness of the results reported in Section 4 to this alternative sorting strategy, but find that results, reported in Subsection 3.2 of the Online Appendix, are virtually unchanged.

5.3 Measurement of global financial conditions

Our global financial conditions index is one of many attempts in the literature to summarise moves in risky asset prices. Other widely used approaches include the global factor presented in Miranda-Agrippino and Rey (2015) (MAR) and the VIX index (see, for example, Habib and Stracca (2012), who identify their global risk shock based on VIX data).[25]

We repeat the exercises in Section 3 using both alternative measures. Quantile regressions based on these alternative indices typically yield conditional FX returns distributions which match qualitatively those of our baseline (especially for VIX-based results, less clearly so for MAR's index), but goodness of fit measures are typically worse. Also, in the case of MAR's index, the ranking of currencies based on downside and upside entropies is less intuitive, in the sense that it does not as clearly match prevailing market narratives. The results can be found in Subsection 3.3 of the Online Appendix.

[25] The correlations of our global financial conditions index with MAR's and the VIX are approximately 0.6 and 0.8 respectively.

6 Conclusion

We provide novel empirical evidence on the relationship between the entire distribution of currency returns and global financial conditions. Our results corroborate some of the prevailing narratives about safe haven and risky currencies, but also provide richer insights than existing studies focussing on mean returns, allowing for example to rank currencies according to their tail behaviour and to quantify the shifts in their distributions following changes in global financial conditions. We also document the role of commonly used macro-financial risk factors in explaining losses on FX portfolios in the face of tighter global financial conditions. These insights can be of interest to policymakers assessing the financial stability outlook of countries, and to investors performing risk management calculations of their investment strategies.

In ongoing research work, we are first of all exploring the usefulness of our approach for forecasting currency returns, that is, evaluating the out-of-sample performance of predictive densities obtained from our baseline specification 1 for $h > 0$. Furthermore, we are expanding the list of macroeconomic fundamentals underlying our portfolio sorting exercises to include other potentially relevant country characteristics, and also exploring alternative empirical strategies that could allow us to rank the relative importance of different risk factors in shifting the distribution of currency returns.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289.

Adrian, T., F. Grinberg, N. Liang, and S. Malik (2018): "The Term Structure of Growth-at-Risk," IMF Working Papers 18/180, International Monetary Fund.

Arregui, N., S. Elekdag, R. G. Gelos, R. Lafarguette, and D. Seneviratne (2018): "Can Countries Manage Their Financial Conditions Amid Globalization?" IMF Working Papers 18/15, International Monetary Fund.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen (2009): "Carry Trades and Currency Crashes," in *NBER Macroeconomics Annual 2008, Volume 23*, National Bureau of Economic Research, Inc, NBER Chapters, 313–347.

Bruno, V. and H. S. Shin (2014): "Cross-Border Banking and Global Liquidity," *The Review of Economic Studies*, 82, 535–564.

Cenedese, G. (2015): "Safe haven currencies: a portfolio perspective," Bank of England working papers 533, Bank of England.

Cenedese, G., L. Sarno, and I. Tsiakas (2014): "Foreign exchange risk and the predictability of carry trade returns," *Journal of Banking & Finance*, 42, 302–313.

Cerutti, E., S. Claessens, and A. K. Rose (2017): "How Important is the Global Financial Cycle? Evidence from Capital Flows," NBER Working Papers 23699, National Bureau of Economic Research, Inc.

Cesa-Bianchi, A., A. Ferrero, and A. Rebucci (2018a): "International credit supply shocks," *Journal of International Economics*, 112, 219–237.

Cesa-Bianchi, A., M. H. Pesaran, and A. Rebucci (2018b): "Uncertainty and Economic Activity: A Multi-Country Perspective," Working Paper 24325, National Bureau of Economic Research.

Crump, R. K., D. Giannone, and S. Hundtofte (2018): "Changing risk-return profiles," Staff Reports 850, Federal Reserve Bank of New York.

Della-Corte, P., S. J. Riddiough, and L. Sarno (2016): "Currency Premia and Global Imbalances," *Review of Financial Studies*, 29, 2161–2193.

Dornbusch, R. and S. Fischer (1980): "Exchange Rates and the Current Account," *American Economic Review*, 70, 960–971.

Drehmann, M., C. Borio, and K. Tsatsaronis (2012): "Characterising the financial cycle: don't lose sight of the medium term!" BIS Working Papers 380, Bank for International Settlements.

Farhi, E., S. P. Fraiberger, X. Gabaix, R. Ranciere, and A. Verdelhan (2009): "Crash Risk in Currency Markets," NBER Working Papers 15062, National Bureau of Economic Research, Inc.

Farhi, E. and X. Gabaix (2016): "Rare Disasters and Exchange Rates," *The Quarterly Journal of Economics*, 131, 1–52.

Fitzenberger, B. (1998): "The moving blocks bootstrap and robust inference for linear least squares and quantile regressions," *Journal of Econometrics*, 82, 235 – 287.

Fratzscher, M. (2009): "What explains global exchange rate movements during the financial crisis?" *Journal of International Money and Finance*, 28, 1390–1407.

Gabaix, X. and M. Maggiori (2015): "International Liquidity and Exchange Rate Dynamics," *The Quarterly Journal of Economics*, 130, 1369–1420.

Gaglianone, W. P. and L. R. Lima (2012): "Constructing Density Forecasts from Quantile Regressions," *Journal of Money, Credit and Banking*, 44, 1589–1607.

Ha, J., M. A. Kose, C. Otrok, and E. S. Prasad (2018): "Global Macro-Financial Cycles and Spillovers," Unpublished manuscript.

Habib, M. M. and L. Stracca (2012): "Getting beyond carry trade: What makes a safe haven currency?" *Journal of International Economics*, 87, 50–64.

Koenker, R. (2005): *Quantile Regression*, Econometric Society Monographs, Cambridge University Press.

Koenker, R. and G. Bassett (1978): "Regression Quantiles," *Econometrica*, 46, 33–50.

Koenker, R. and J. A. F. Machado (1999): "Goodness of Fit and Related Inference Processes for Quantile Regression," *Journal of the American Statistical Association*, 94, 1296–1310.

Koop, G. and D. Korobilis (2014): "A new index of financial conditions," *European Economic Review*, 71, 101–116.

Korobilis, D. (2017): "Quantile regression forecasts of inflation under model uncertainty," *International Journal of Forecasting*, 33, 11–20.

Krugman, P. (1979): "A Model of Balance-of-Payments Crises," *Journal of Money, Credit and Banking*, 11, 311–325.

Lustig, H., N. Roussanov, and A. Verdelhan (2011): "Common Risk Factors in Currency Markets," *Review of Financial Studies*, 24, 3731–3777.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012): "Carry Trades and Global Foreign Exchange Volatility," *Journal of Finance*, 67, 681–718.

Miranda-Agrippino, S. and H. Rey (2015): "US Monetary Policy and the Global Financial Cycle," NBER Working Papers 21722, National Bureau of Economic Research, Inc.

Mohammadi, S. (2009): "QUANTILEREG: MATLAB function to estimate quantile regression," Statistical Software Components, Boston College Department of Economics.

Obstfeld, M. and K. Rogoff (1995): "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 103, 624–660.

Ranaldo, A. and P. Soderlind (2010): "Safe Haven Currencies," *Review of Finance*, 14, 385–407.

Rey, H. (2013): "Dilemma not trilemma: the global cycle and monetary policy independence," *Proceedings - Economic Policy Symposium - Jackson Hole*, 1–2.

Wijnbergen, S. V. (1991): "Fiscal Deficits, Exchange Rate Crises and Inflation," *Review of Economic Studies*, 58, 81–92.

Xiao, Z. (2012): "9 - Time Series Quantile Regressions," in *Time Series Analysis: Methods and Applications*, ed. by T. S. Rao, S. S. Rao, and C. Rao, Elsevier, vol. 30 of *Handbook of Statistics*, 213 – 257.

Zhou, Z. and X. Shao (2013): "Inference for linear models with dependent errors," *Journal of the Royal Statistical Society Series B*, 75, 323–343.

A Appendix: Data

A.1 Exchange rates

The analysis in Section 3 is conducted using Nominal Effective Exchange Rates (NEERs) from the BIS from January 1994 to June 2018 for the following countries: Algeria, Argentina, Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Chile, China, Chinese Taipei, Colombia, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Euro area, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Korea, Latvia, Lithuania, Luxembourg, Malaysia, Malta, Mexico, Netherlands, New Zealand, Norway, Peru, Philippines, Poland, Portugal, Romania, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland, Thailand, Turkey, United Arab Emirates, United Kingdom, United States and Venezuela.

Exchange rate changes are computed as log differences on monthly averages; interest rate differentials (and the corresponding excess returns) are not considered in the baseline analysis.

A.2 Global financial conditions

As described in Section 2, our proxy for global financial conditions is constructed from a series of country-specific Financial Condition Indices (FCIs) following Arregui et al. (2018), which in turn base their method on Koop and Korobilis (2014). These country-specific FCIs consider the following variables:

- Long-term government interest rates: yield on nominal government bonds with maturity of 10 years. Source: Thomson Reuters Datastream.
- Sovereign spreads: for advanced economies, we calculate sovereign spreads as the difference between domestic long-term government interest rates and those of bonds of a benchmark country (Germany for Europe and US for rest of the world). For emerging market economies we use stripped spreads from JP Morgan's EMBI. Sources: Thomson Reuters Datastream and JP Morgan.
- Term spreads: difference between domestic long-term government interest rates and a domestic short term T-bill rate (with maturity of 3 months or closest). Sources: Thomson Reuters Datastream and Bank of America Merrill Lynch.

- Interbank spreads: difference between 3-month interbank rate (or closest) and 3-m T-Bill rate (or closest). Source: Thomson Reuters Datastream and national central banks.

- Corporate spreads: corporate spread indices. Sources: Bank of America Merril Lynch, Barclays, JPMorgan (CEMBI) and Standard & Poor's.

- Equity returns: monthly return of domestic stock index, measured in domestic currency. Source: Thomson Reuters Datastream.

- Equity volatility: realised monthly volatility computed using daily changes in equity index. Source: Thomson Reuters Datastream.

- Market capitalisation of financial sector: market capitalisation of MSCI Country Financials Index divided by MSCI Country Index. Source: MSCI Inc.

- Credit growth: monthly change of credit to households and non-profit institutions serving households, provided by all sectors. Source: BIS.

- House price returns: monthly returns based on residential property prices. Source: BIS.

The macroeconomic variables used to 'clean' financial condition indices are CPI inflation and industrial production (source: national sources via Thomson Reuters Datastream).

We compute FCIs at the monthly frequency from January 1995 to June 2018 for the following countries: Argentina, Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Chile, China, Colombia, Czech Republic, Denmark, Finland, France, Germany, Hungary, India, Ireland, Israel, Italy, Japan, Korea, Malaysia, Mexico, Netherlands, New Zealand, Norway, Peru, Philippines, Poland, Portugal, Russia, South Africa, Spain, Sweden, Switzerland, Thailand, Turkey, United Kingdom, United States and Venezuela. Where needed, we splice back the series (up to 1991) using the FCIs published in the IMF's April 2017 GFSR.

Armed with a set of country-specific FCIs we then compute our proxy of global financial conditions as a simple cross-sectional average of these.

A.3 Risk factors

We use a series of macro-financial variables as risk factors in the portfolio-sorting exercise conducted in Section 4. The variables considered are the following:

- Current account balance. Sources: IMF IFS and OECD databases.

- Interest rate differentials: relying on the CIP condition, we use FX-forward based forward discounts (vs. the US dollar) as a proxy for interest rate differentials.[26] Source: Thomson Reuters Datastream.

- International reserves: total international reserves. Source: IMF IFS database.

- Fiscal balance: fiscal position of the government after accounting for capital expenditures. Source: OECD.

- GDP: Gross Domestic Product, constant prices in domestic currency. Source: IMF IFS database.

B Appendix: Quantile regression

Given a linear model for the conditional quantile function

$$Q_y(\tau|X) = x\beta(\tau) \tag{B.1}$$

the quantile regression estimate $\hat{\beta}(\tau)$ is the minimiser of

$$\hat{V}(\tau) = \min_{\beta \in \mathbb{R}^p} \sum \rho_\tau \left(y_i - x_i'\beta\right) \tag{B.2}$$

where $\rho_\tau(u) = u[\tau - I(u < 0)]$ is the so-called check function.

As discussed in Koenker (2005), the solution of problem B.2 is amenable to linear programming techniques. However, in our MATLAB implementation, we have found it computationally more efficient to approximate the exact solution via an iteratively-reweighted-least-squares (IRLS) algorithm. This is motivated by the close relationship of B.2 to the problem of finding the least-absolute-deviations (LAD) estimator (which obtains for $\tau = 0.5$), and more generally of solving L^p-norm linear regression problems. Building on Mohammadi (2009), we proceed as follows: we start from an initial OLS estimate,

$$\hat{\beta}^{(0)}(\tau) = (x'x)^{-1} x'y.$$

We then take the residuals $\hat{u}_i^{(0)}(\tau) = y_i - x_i\hat{\beta}^{(0)}(\tau)$ and construct a diagonal matrix of

[26] We acknowledge the presence of measurement error due to deviations from the CIP condition after 2008.

weights $w^{(t)}, t > 0$, whose diagonal elements are given by

$$w_{ii}^{(t)}(\tau) = \frac{1}{\rho_{1-\tau}\left(u_i^{(t-1)}(\tau)\right)}$$

We then obtain an updated estimate $\hat{\beta}^{(t)}(\tau)$, residuals $\hat{u}^{(t)}(\tau)$ and weights $w^{(t+1)}(\tau)$ using weighted least squares:

$$\hat{\beta}^{(t)}(\tau) = \left(x'w^{(t)}(\tau)'x\right)^{-1} x'w^{(t)}(\tau)'y$$

and iterate until convergence. Essentially, the procedure approximates B.2 by a convergent sequence of weighted sums of square residuals, where the weights are chosen so as to approximate the check function ρ_τ with a quadratic one.

B.1 Bootstrapping

While there are several results available for inference in quantile regression with time-series data (see for example Xiao (2012), Zhou and Shao (2013)), we take a shortcut and deal with potential autocorrelation in the errors from B.2 by bootstrapping confidence intervals for all quantities of interest. Fitzenberger (1998) shows that a moving (or overlapping) block bootstrap procedure provides heteroskedasticity- and autocorrelation-consistent (HAC) standard errors for quantile regression coefficient estimators.

The procedure works as follows: letting $z_t = [y_t, x_t]$ denote the original data, T the sample size and b a suitably chosen block length, a resample z_{it}^* of length $T^* = b * round(T/b)$ is obtained by joining $round(T/b)$ draws (with replacement) of b consecutive elements of z_t (blocks), where the blocks are allowed to overlap. Each resample z_{it}^* yields an estimate of the quantile regression coefficients $\hat{\beta}_i^*(\tau)$ and can be used to compute all other statistics of interest, such as $\hat{V}_i(\tau)$ and thus $R^1(\tau)$ etc. Confidence intervals at level γ for $\hat{\beta}(\tau)$ and other quantities of interest are computed as

$$\left(2\hat{\beta}(\tau) - \hat{\beta}^*_{\frac{1-\gamma}{2}}(\tau), 2\hat{\beta}(\tau) - \hat{\beta}^*_{\frac{\gamma}{2}}(\tau)\right) \tag{B.3}$$

where $\hat{\beta}_p^*(\tau)$ denotes the $p-$th percentile of the bootstrapped draws $\hat{\beta}_i^*(\tau)$[27].

[27]In the computation of confidence intervals for $R^1(\tau)$ we instead compute directly percentiles from the bootstrapped draws to ensure non-negative values.