Global Financial Conditions and Exchange Rate Tail Risks ☆

Fernando Eguren-Martin[†] Andrej Sokol[‡]

January 18, 2019
PRELIMINARY AND INCOMPLETE
NOT FOR CIRCULATION

Abstract

We document how the 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 in 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 the currencies of countries with high interest rates, large current account deficits and low levels of international reserves display a higher likelihood of sharp depreciation in response to a tightening in global financial conditions.

Keywords: exchange rates, tail risks, financial conditions indices, global financial cycle, quantile regression.

JEL Codes: F31, G15.

☆We would like to thank Gino Cenedese, Giancarlo Corsetti, Simon Lloyd, Lucio Sarno and Giorgio Valente 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/.

[†]Bank of England and University of Oxford. Email: fernando.egurenmartin@bankofengland.co.uk.
[‡]Bank of England, European Central Bank and CfM. Email: andrej.sokol@ecb.europa.eu.

1 Introduction

Recent years have witnessed a heated debate about the extent and interpretation of global co-movement in 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 admitting a degree of global co-movement in financial conditions, motivate the main question addressed in this paper, namely, how different exchange rates co-move with that global component. Unlike most of the existing literature, we study the behaviour of the whole 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.

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 significant improvements in fit in the tail regions even where standard R^2 measures are low, and then rank currencies according to how their left (appreciation) and right (depreciation) tails respond to a tightening in global financial conditions. Our exercise corroborates some of the prevailing narratives about 'safe haven' and 'risky' currencies, but also offers interesting new insights.

Second, to identify potential risk factors associated with different tail behaviour of currencies, we conduct portfolio sorting exercises based on several macroeconomic fundamentals, and then study the responses of the resulting returns series to a tightening in global financial conditions. We find that the currencies of countries with high interest rates, large current account deficits and low levels of international reserves display a higher likelihood of sharp depreciation in response to a tightening in global financial conditions.

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 substantial 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.

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).

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 proposed by Koenker and Bassett (1978), in both macroeconomics and finance. Most notably, Adrian et al. (2016) 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

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

related to past 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 drivers of currencies' differential tail behaviour. In Section 5 we conclude, while the Appendix includes details about our data and methodology.

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. 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 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 respond to data availability constraints: it is not straightforward to construct a panel dataset spanning a wide 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 wide set of financial indicators for 43 countries over 1991-2017. The

[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 has typically been referred to as 'the global financial cycle'.

Figure 1 Global Financial Conditions Index, 1991-2017.

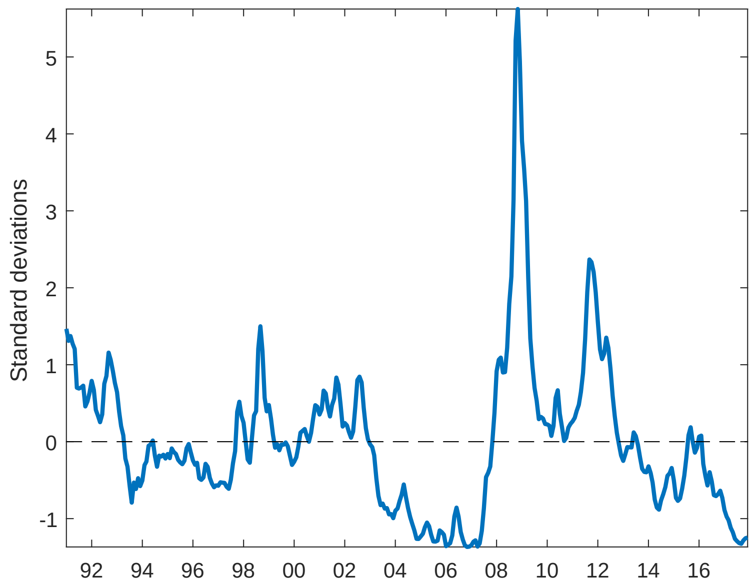

financial series included are as follows: term, sovereign, interbank and corporate spreads, long-term interest rates, equity returns and volatility, relative market capitalisation of the financial sector, house prices and credit growth.[4] To obtain country-specific summary measures of financial conditions out of these series we follow Koop and Korobilis (2014) and estimate factor models which allow for time variation in the parameters and 'clean' financial conditions from changes that reflect a response to macro-economic news (namely, industrial production and CPI inflation). That is, we 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.[5]

Armed with a set country-specific financial condition indices, we extract a global component by taking the cross-sectional mean. 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 (see Figure 7). In what follows we take this series as our measure of global financial conditions.

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

[5] Note that given the forward-looking nature of asset prices, it could still be the case that these responded to news to *expected* macro-economic developments, not captured in the contemporaneous series used in our approach.

Figure 1 shows the evolution of our measure over the last 30 years, which is broadly in line with the prevailing narrative: 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 widely used US-centric measures such as the VIX index or the S&P index.[6]

In the context of the current 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 contain any FX data at all.

3 Quantifying exchange rate tail risks using 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.

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

[6]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 0.7.

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

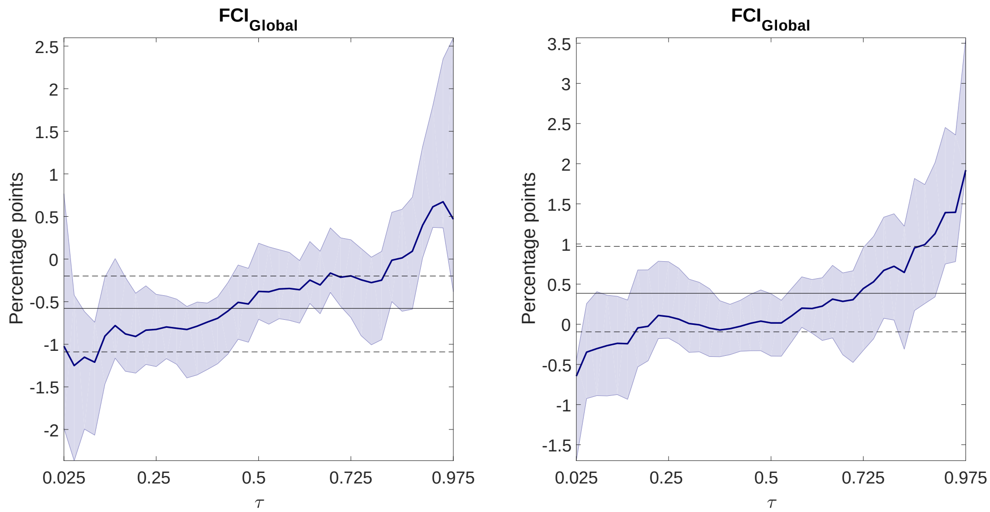

(a) Japan (b) Australia

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.

financial conditions on the distribution of exchange rate returns. We specify a linear model for their conditional quantiles as follows:

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

where ΔFX_{t+h} is the monthly log change in the nominal effective exchange rate[7] h months ahead and GFC_t is our measure of global financial conditions. Function Q computes quantiles τ of the distribution of ΔFX_{t+h} given X_t. Appendix B discusses technical details.

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 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.

[7]We consider nominal effective exchange rates to focus on country idiosyncratic dynamics, avoiding potentially US-driven moves of US dollar bilaterals (which can move in sync with global financial conditions).

Table 1 Goodness of fit measures, selected currencies.

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

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 $h = 0$.[8] The Japanese yen exhibits negative coefficients approximately up to the 75th quantile, meaning that most of the conditional distribution shifts to the left in the face of a tightening in global financial conditions. On the other hand, the coefficients for the Australian dollar are for the most part not statistically different from 0 below the last quartile or so, and positive thereafter, indicating an increased risk of a sharp depreciation. The black lines show coefficient values from simple OLS regressions as a benchmark. While in Figure 2 the OLS coefficients are still able to capture the different behaviour of the two currencies in the face of the same shock (this is not always as clear cut with other currencies, and even for Australia the coefficient is not statistically significant), it is clear that much interesting information is discarded by only focussing on the conditional mean.

Another interesting 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 $R^1(\tau)$ 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, $R^1(\tau)$ provide information on the local relative 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

[8]The convention we adopt is that negative FX changes represent an appreciation, and positive changes a depreciation.

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

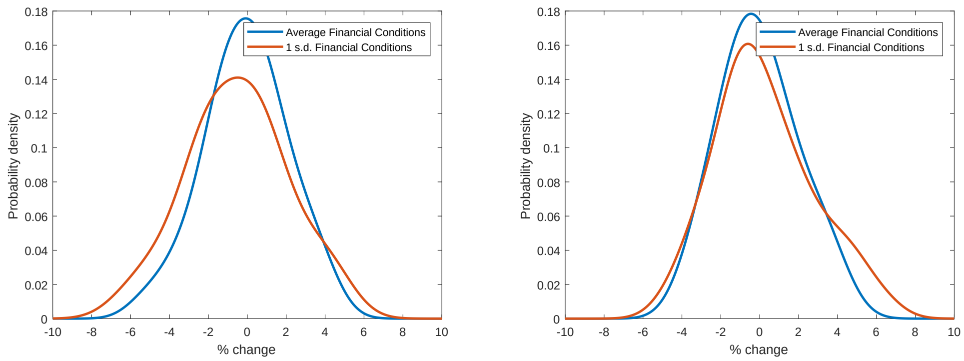

(a) Japan (b) Australia

estimate of the unconditional quantile τ).[9] The interpretation is thus analogous to that of standard R^2: $R^1(\tau)$ expresses the improvement in fit, in terms of appropriately weighted absolute residuals, obtained by adding covariates to the model.

Table 1 reports both R^2 and $R^1(\tau)$ measures for selected currencies. 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. However, as far as $R^1(\tau)$ measures are concerned, a robust pattern seems to hold across countries (see Table 3 for the full panel), namely, that the fit tends to generally improve in the tails.

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. (2016), 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. 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 1 standard deviation increase in global

[9] 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.

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

Figure 3 illustrates the changes induced by a 1 standard deviation increase in global financial conditions from their average level on the conditional densities of the same two currencies analysed before. Our exercise does corroborate their usual characterisation: in the face of a tightening in global financial conditions, the left (appreciation) tail of the distribution of the Japanese yen shifts down significantly (increased chances of a sharp appreciation), while the right (depreciation) tail of the distribution of the Australian dollar shifts up (increased chances of a sharp depreciation).

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 in global financial conditions.[10] 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 (relative) entropy outside of the interquartile range of $\hat{g}(x)$ as

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

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

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

Figure 4 shows the results for a selection of currencies, and Figure 8 for the whole sample. The ranking in terms of tail behaviour once again broadly confirms the prevailing narrative: typical 'safe haven' currencies such as the Japanese yen, the US dollar and the Swiss franc exhibit high downside entropy but hardly any upside entropy, whereas 'risky' currencies such as the Australian dollar exhibit a higher upside entropy. The case of the euro is somewhat interesting in that it exhibits similar degrees of both upside and downside entropy, meaning that both tails become fatter in response to a tightening in global financial conditions.

In the next Section we turn to analysing the underlying country characteristics that

[10] This is similar in spirit to the quantification of upside and downside risks in Adrian et al. (2016).

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

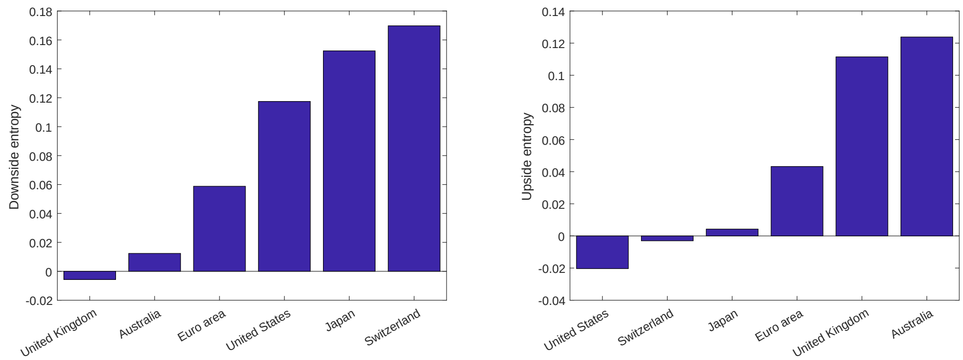

(a) Downside entropy

(b) Upside entropy

are associated with the different responses of currencies' distributions to changes in global financial conditions.

4 Identifying risk factors: a portfolio 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? 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.[11]

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 single statistic

[11]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.

is likely to hide significant within-period variation. 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 could plausibly be 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.[12][13] We consider each of the candidate 'risk factor' variables in turn and, at every point in time throughout our sample, begin by ranking countries according to the values they display for the variable under consideration. For example, when working with current account balances, we rank countries from those displaying the highest current account surplus to those with the highest deficit.[14]

Table 2 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	6.2
Current Account	7.8	4.0	3.9	1.2	1.4	5.8
Fiscal	0.5	0.1	0.0	0.7	1.2	0.3
Reserves	2.3	2.2	0.8	0.3	1.6	2.0

We then assign currencies to five portfolios according to this ranking. 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. 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 and fifth portfolios.[15] 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.

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

[12] Exact definitions and data sources can be found in Appendix A

[13] Interest rate differentials are implied from forward FX contracts. Note that recent CIP deviations mean that there is measurement error in this quantification of interest rate differentials. For simplicity, we rely on interest rate differentials with respect to the US despite using nominal effective exchange rates and not US dollar bilaterals.

[14] We rebalance portfolios annually given availability of underlying data for sorting variables.

[15] We use pure FX-driven returns (log exchange rate changes) and not consider interest rate differentials.

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.[16]

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 in global financial conditions. As before, we report goodness-of-fit measures for the quantile regressions and relative entropy measures to summarise the tail behaviour of the different relative portfolios.

Table 2 shows that the goodness of fit of quantile regressions of relative portfolio returns on global financial conditions improves in the tails, 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.

As for the shape of the distribution of conditional returns in the face of tighter global financial conditions, Figure 5 shows that currencies of countries with high interest rates, large current account deficits and low levels of international reserves display a higher likelihood of experiencing a sharp depreciation. This is also reflected in positive downside entropies (Figure 6), and is in line with our priors. The results for currencies of countries with high fiscal deficits are less clear-cut: there is a very minor increase in both downside and upside

[16]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 5 Impact of a tightening in global financial conditions on the conditional distribution of relative portfolio returns.

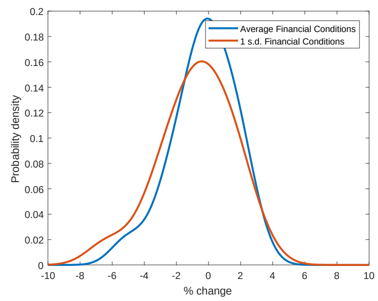

(a) Carry

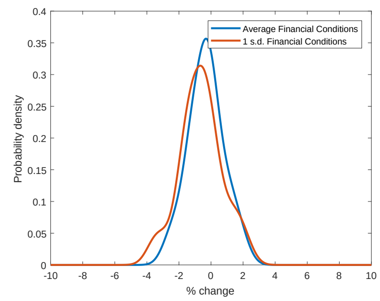

(b) Current Account

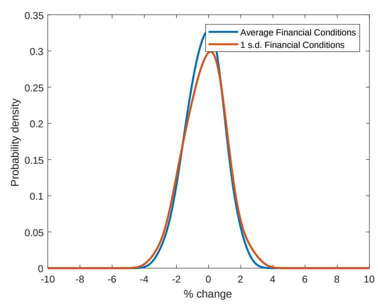

(c) Fiscal

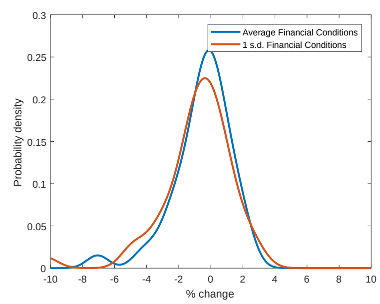

(d) Reserves

entropies, but the overall goodness of fit across quantiles is significantly lower than for the other risk factors (Table 2).

In sum, global financial conditions contain useful information for characterising the returns distribution of currencies exposed to a series of risk factors, namely high interest rates, large current account deficits and low levels of international reserves. This is particularly true of (negative) tail returns.

5 Conclusion

We provide novel evidence on the relation between 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. We also highlight the importance of commonly used 'risk factors' associated with an increased likelihood of a sharp depreciation in the face of tighter global financial conditions.

In ongoing research work, we are first of all exploring the usefulness of our approach for

Figure 6 Downside and upside entropy measures of conditional relative portfolio returns.

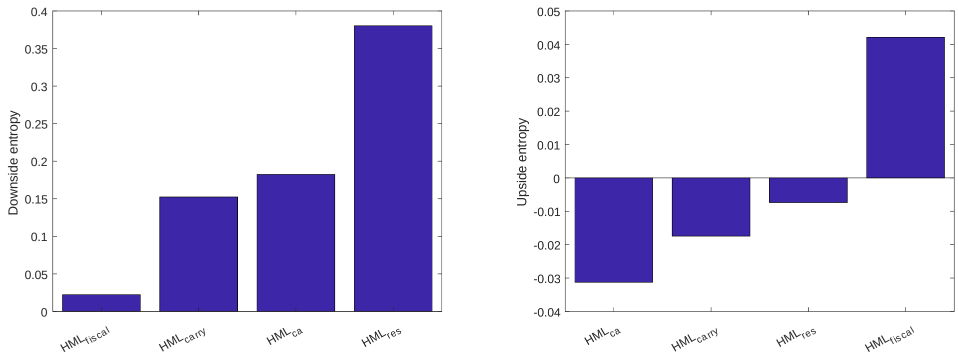

(a) Downside entropy

(b) Upside entropy

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.

6 Figures and Tables

Figure 7 Share of variance of country FCIs due to GFCI.

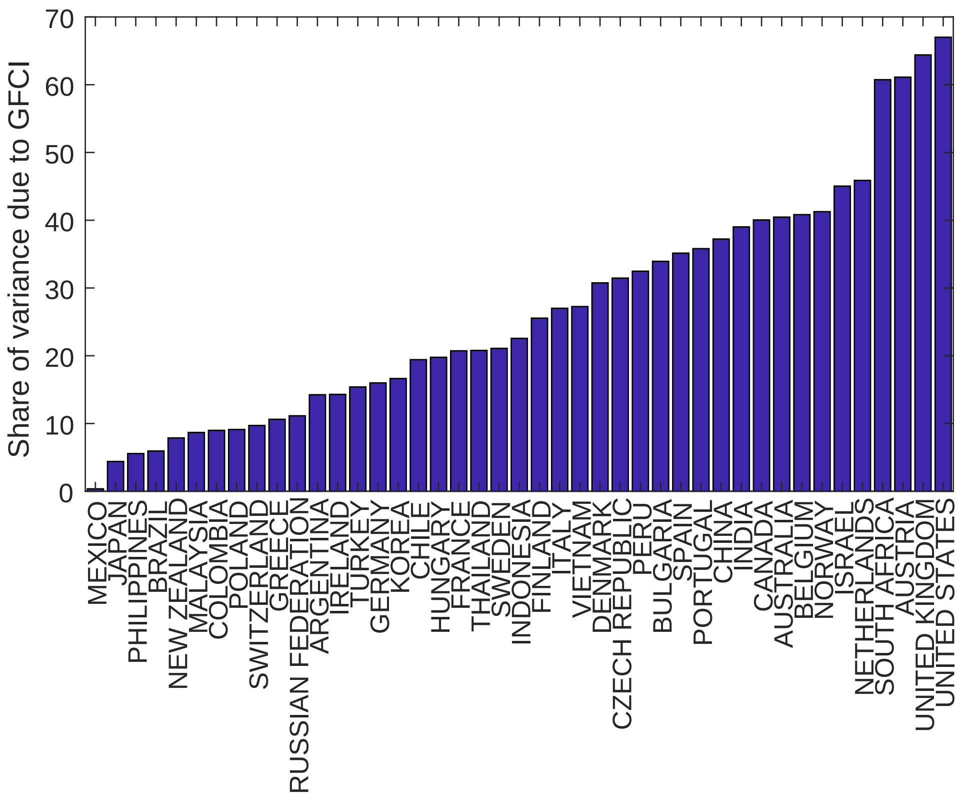

Table 3 Goodness of fit measures, full sample.

	$R^1(\tau)$					R^2
	0.05	0.25	0.5	0.75	0.95	
Algeria	6.0	0.0	0.5	0.3	2.8	0.0
Argentina	4.2	2.6	0.2	0.1	0.9	0.1
Australia	1.2	0.2	0.0	1.9	16.5	3.2
Austria	5.2	0.5	0.3	0.4	0.0	0.0
Belgium	5.6	0.1	0.1	0.5	1.8	0.0
Brazil	0.1	0.8	1.2	2.4	4.3	1.8
Bulgaria	2.1	0.7	0.0	0.1	1.0	0.4
Canada	3.3	1.2	1.6	2.0	2.8	5.7
Chile	0.9	0.3	0.8	1.1	13.5	2.5
China	10.5	2.7	1.6	0.1	1.8	2.6

Chinese Taipei	0.1	0.3	0.1	1.0	0.8	1.2
Colombia	0.3	0.0	0.6	1.9	9.9	2.4
Croatia	0.1	0.0	0.2	0.7	2.0	0.4
Cyprus	6.9	1.8	0.0	0.1	0.0	0.7
Czech Republic	9.2	0.3	1.3	4.9	10.2	1.7
Denmark	9.0	0.8	0.1	0.4	1.2	0.3
Estonia	4.8	1.2	0.0	0.2	0.7	1.1
Euro area	7.5	0.4	0.1	0.4	1.7	0.0
Finland	4.2	0.4	0.1	0.1	0.2	0.1
France	6.1	0.4	0.1	0.8	2.4	0.0
Germany	5.1	0.3	0.2	0.5	1.0	0.0
Greece	5.8	0.7	0.2	0.1	0.1	0.1
Hong Kong SAR	10.1	4.5	1.7	0.3	0.2	2.8
Hungary	3.0	0.1	1.2	2.9	14.7	3.3
Iceland	0.2	0.5	2.3	4.8	16.0	6.7
India	0.1	0.5	0.9	5.1	7.0	3.3
Indonesia	5.9	0.4	0.0	0.3	6.2	0.2
Ireland	6.1	0.1	0.3	2.2	3.5	0.3
Israel	4.4	0.0	0.3	1.8	2.4	1.5
Italy	3.4	0.2	0.2	0.4	0.1	0.0
Japan	12.3	6.5	1.7	0.7	2.8	5.8
Korea	1.2	0.3	1.4	5.2	19.1	4.5
Latvia	3.9	1.3	0.3	0.0	1.6	1.6
Lithuania	1.7	2.0	0.2	0.0	0.7	0.6
Luxembourg	4.8	0.1	0.2	0.3	0.0	0.0
Malaysia	0.0	0.3	0.2	0.0	1.6	0.1
Malta	4.5	0.0	0.5	0.2	3.7	0.0
Mexico	0.3	0.0	0.6	3.9	4.3	2.3
Netherlands	6.0	0.3	0.2	0.9	2.5	0.0
New Zealand	0.8	0.1	2.5	5.5	5.5	2.5
Norway	0.1	0.9	0.3	0.1	3.6	1.0
Peru	0.0	2.7	1.4	0.1	2.1	0.2
Philippines	0.2	0.4	0.1	0.0	1.0	0.3
Poland	2.1	0.0	2.6	9.4	20.4	8.9
Portugal	2.9	0.0	0.2	0.1	0.2	0.0

Romania	2.2	1.5	2.1	2.0	0.9	1.3
Russia	0.0	0.1	0.2	2.5	2.6	2.0
Saudi Arabia	10.9	2.6	1.5	0.2	0.7	2.8
Singapore	1.2	0.3	0.1	0.0	6.7	0.3
Slovakia	0.9	0.1	0.1	0.6	0.1	0.1
Slovenia	4.5	1.1	0.1	0.3	1.3	0.4
South Africa	0.6	0.6	0.5	0.4	4.3	1.8
Spain	4.2	0.4	0.1	0.5	0.0	0.0
Sweden	1.8	1.4	0.2	1.6	7.9	2.1
Switzerland	9.0	3.7	2.2	0.1	0.1	2.1
Thailand	0.1	0.7	1.0	1.1	1.7	0.7
Turkey	0.5	0.1	0.6	0.3	2.3	0.6
United Arab Emirates	12.6	3.5	2.2	0.1	0.2	3.8
United Kingdom	0.2	0.0	0.5	2.8	10.8	3.6
United States	9.8	2.4	1.9	0.1	0.3	4.7
Venezuela	3.8	1.0	0.2	0.0	0.5	0.5

Figure 8 Downside and upside entropy measures of conditional exchange rate returns.

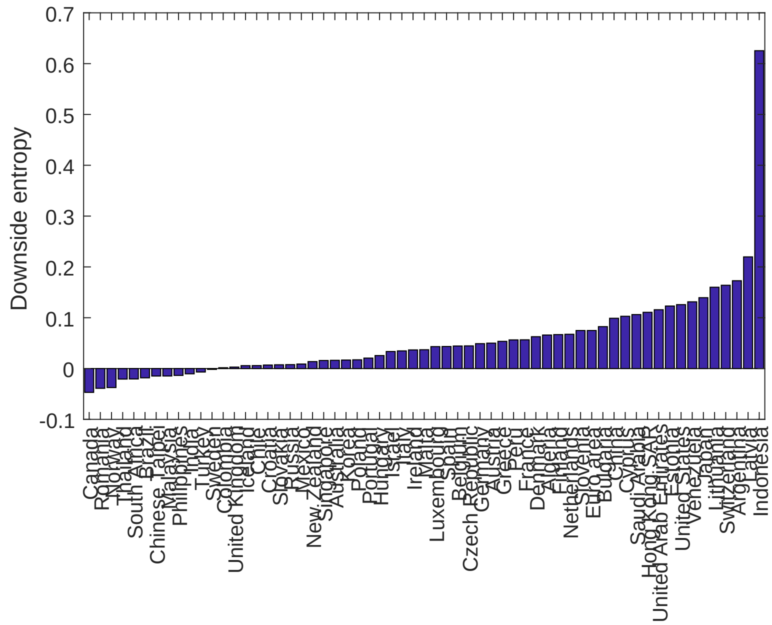

(a) Downside entropy

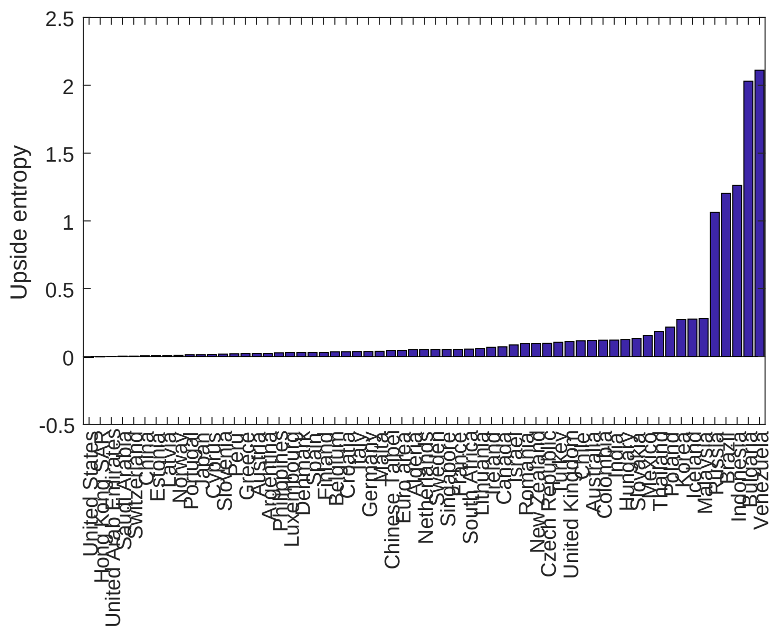

(b) Upside entropy

References

Adrian, T., Boyarchenko, N., Giannone, D., Sep. 2016. Vulnerable growth. Staff Reports 794, Federal Reserve Bank of New York.

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

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

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

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

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

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

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

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

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

Fitzenberger, B., 1998. The moving blocks bootstrap and robust inference for linear least squares and quantile regressions. Journal of Econometrics 82 (2), 235 – 287.

Ha, J., Kose, M. A., Otrok, C., Prasad, E. S., Aug. 2018. Global macro-financial cycles and spillovers. Unpublished manuscript.

Habib, M. M., Stracca, L., 2012. Getting beyond carry trade: What makes a safe haven currency? Journal of International Economics 87 (1), 50–64.

Koenker, R., 2005. Quantile Regression. Econometric Society Monographs. Cambridge University Press.

Koenker, R., Bassett, G., 1978. Regression quantiles. Econometrica 46 (1), 33–50.

Koenker, R., Machado, J. A. F., 1999. Goodness of fit and related inference processes for quantile regression. Journal of the American Statistical Association 94 (448), 1296–1310.

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

Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common Risk Factors in Currency Markets. Review of Financial Studies 24 (11), 3731–3777.

Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., April 2012. Carry Trades and Global Foreign Exchange Volatility. Journal of Finance 67 (2), 681–718.

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

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

Ranaldo, A., Soderlind, P., 2010. Safe Haven Currencies. Review of Finance 14 (3), 385–407.

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

Xiao, Z., 2012. 9 - time series quantile regressions. In: Rao, T. S., Rao, S. S., Rao, C. (Eds.), Time Series Analysis: Methods and Applications. Vol. 30 of Handbook of Statistics. Elsevier, pp. 213 – 257.

Zhou, Z., Shao, X., March 2013. Inference for linear models with dependent errors. Journal of the Royal Statistical Society Series B 75 (2), 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; interest rate differentials (and the consequent 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 bond 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.[17] 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,

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

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

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

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

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

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

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

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 such 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 $\hat{\beta}_i^*(\tau)$[18].

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