

# DISCUSSION PAPER SERIES

DP15817

## **How Puzzling Is the Forward Premium Puzzle? A Meta-Analysis**

Diana Zigrainova, Tomas Havranek, Zuzana Irsova  
and Jiri Novak

**INTERNATIONAL MACROECONOMICS AND FINANCE**

**CEPR**

# How Puzzling Is the Forward Premium Puzzle? A Meta-Analysis

*Diana Zigraiova, Tomas Havranek, Zuzana Irsova and Jiri Novak*

Discussion Paper DP15817  
Published 16 February 2021  
Submitted 15 February 2021

Centre for Economic Policy Research  
33 Great Sutton Street, London EC1V 0DX, UK  
Tel: +44 (0)20 7183 8801  
[www.cepr.org](http://www.cepr.org)

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Macroeconomics and Finance

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Diana Zigraiova, Tomas Havranek, Zuzana Irsova and Jiri Novak

# How Puzzling Is the Forward Premium Puzzle? A Meta-Analysis

## Abstract

A key theoretical prediction in financial economics is that under risk neutrality and rational expectations a currency's forward rates should form unbiased predictors of future spot rates. Yet scores of empirical studies report negative slope coefficients from regressions of spot rates on forward rates. We collect 3,643 estimates from 91 research articles and using recently developed techniques investigate the effect of publication and misspecification biases on the reported results. Correcting for these biases yields slope coefficients in the intervals (0.23,0.45) and (0.95,1.16) for the currencies of developed and emerging countries respectively, which implies that empirical evidence is in line with the theoretical prediction for emerging economies and less puzzling than commonly thought for developed economies. Our results also suggest that the coefficients are systematically influenced by the choice of data, numeraire currency, and estimation method.

JEL Classification: C83, F31, G14

Keywords: Forward Rate Bias, uncovered interest parity, meta-analysis, Publication bias, Model uncertainty

Diana Zigraiova - [diana.zigraiova@fsv.cuni.cz](mailto:diana.zigraiova@fsv.cuni.cz)  
*European Stability Mechanism, Luxembourg*

Tomas Havranek - [tomas.havranek@ies-prague.org](mailto:tomas.havranek@ies-prague.org)  
*Charles University, Prague and CEPR*

Zuzana Irsova - [zuzana.irsova@ies-prague.org](mailto:zuzana.irsova@ies-prague.org)  
*Charles University, Prague*

Jiri Novak - [jiri.novak@fsv.cuni.cz](mailto:jiri.novak@fsv.cuni.cz)  
*Charles University, Prague*

# How Puzzling Is the Forward Premium Puzzle? A Meta-Analysis\*

Diana Zigraiova<sup>a,b</sup>, Tomas Havranek<sup>b</sup>, Zuzana Irsova<sup>b</sup>, and Jiri Novak<sup>b</sup>

<sup>a</sup>European Stability Mechanism, Luxembourg

<sup>b</sup>Charles University, Prague

February 8, 2021

## Abstract

A key theoretical prediction in financial economics is that under risk neutrality and rational expectations a currency's forward rates should form unbiased predictors of future spot rates. Yet scores of empirical studies report negative slope coefficients from regressions of spot rates on forward rates. We collect 3,643 estimates from 91 research articles and using recently developed techniques investigate the effect of publication and misspecification biases on the reported results. Correcting for these biases yields slope coefficients in the intervals (0.23, 0.45) and (0.95, 1.16) for the currencies of developed and emerging countries respectively, which implies that empirical evidence is in line with the theoretical prediction for emerging economies and less puzzling than commonly thought for developed economies. Our results also suggest that the coefficients are systematically influenced by the choice of data, numeraire currency, and estimation method.

**Keywords:** Forward rate bias, uncovered interest parity, meta-analysis, publication bias, model uncertainty

**JEL Codes:** C83, F31, G14

## 1 Introduction

If forward exchange rates systematically differ from future spot rates, money can be made on the difference: a risk-neutral agent with rational expectations can exploit the inefficiency and hence, supposedly, the anomaly should disappear. It is therefore puzzling that the forward anomaly has been found again and again for dozens of different currencies, time periods, and identification designs. Yet the exact results in the literature vary, and the null hypothesis is

---

\*An online appendix with data and code is available at [meta-analysis.cz/forward](http://meta-analysis.cz/forward). Corresponding author: Jiri Novak, [jiri.novak@fsv.cuni.cz](mailto:jiri.novak@fsv.cuni.cz). We thank the participants of the MAER-Net Colloquium in Greenwich, 2019, and seminar participants at the European Stability Mechanism and the Crawford School of Public Policy, Australian National University, for their helpful comments. Diana Zigraiova acknowledges support from the European Union's 2020 Research and Innovation Staff Exchange programme under the Marie Skłodowska-Curie grant agreement #681228, Charles University (project Primus/17/HUM/16), and the Czech Science Foundation (grant #18-02513S). Tomas Havranek and Zuzana Irsova acknowledge support from the Czech Science Foundation (project 19-26812X). Jiri Novak acknowledges support from the Czech Science Foundation (project 21-09231S). The views expressed here are ours and not necessarily those of our employers.

not rejected universally. The anomaly is commonly labeled “forward premium puzzle,” because most studies do not estimate the relationship in levels (future spot rates on forward rates), but subtract current spot rates from both sides of the regression. Thus one obtains currency depreciation on the left-hand side and the forward discount on the right-hand side. The puzzle is that, according to many studies, depreciation is positively associated with forward *premium*, not a discount. So not only do researchers typically reject the hypothesis of the coefficient (we will call it  $\beta$ ) being equal to one, but often they find a statistically significant negative coefficient. If the covered interest parity holds, this result is equivalent to the finding that currencies with higher interest rates tend to appreciate. Thus tests of forward rate unbiasedness are closely related to tests of the uncovered interest parity.

The forward premium puzzle is a traditional problem in international economics and finance; as such, it has attracted the attention of dozens of researchers during the last four decades. Yet still no clear-cut consensus emerges on whether the puzzle really exists or whether it represents a statistical artifact, how large the departure from the null hypothesis is, and how material the implications are in practice. Important prospective solutions to the forward premium puzzle put forward in the last decade include infrequent portfolio decisions (Bacchetta & van Wincoop, 2010), investor overconfidence (Burnside *et al.*, 2011), omitted variables (Pippenger, 2011), sentiment (Yu, 2013), sovereign default risk (Coudert & Mignon, 2013), order flow (Breedon *et al.*, 2016), rare disasters (Farhi & Gabaix, 2016), and monetary policy changes (Kim *et al.*, 2017; Park & Park, 2017; Coulibaly & Kempf, 2019), a string of efforts that highlights persistent research activity in the field. What the literature lacks is a quantitative synthesis, or meta-analysis, that would take stock of the enormous body of work and shed light on potential biases and patterns that are impossible to detect in individual studies considered separately. That is what we attempt to achieve in this paper.

Meta-analysis brings value added in four key areas. First, it can quantify the results that are qualitatively well established in the literature but their size varies across studies. For example, most of the literature agrees that the  $\beta$  coefficient depends on country characteristics (Bansal & Dahlquist, 2000; Frankel & Poonawala, 2010) and monetary policy rules (McCallum, 1994; Park & Park, 2017). Because meta-analysis uses many results across a large number of datasets, it can provide more robust estimates for these patterns. Second, individual studies estimating  $\beta$  use different methodologies, and meta-analysis can quantify whether different methods have systematic effects on the results. Again, a large number of results from many different datasets is crucial for robust inference on the impact of methodology. Third, an important aspect of meta-analysis is the identification of a plausible range for  $\beta$  based on the entire corpus of empirical literature, because results from individual studies vary widely even for the same currencies. Finally, and perhaps most importantly, meta-analysis allows us to correct the literature for potential publication selection bias, which is impossible to do in individual studies.


Publication bias is the tendency of authors, editors, and referees to prefer results that are statistically significant and consistent with previous findings or underlying theory. The bias has been discussed, among others, by Stanley (2001), Stanley (2005), Stanley & Doucouliagos

(2010), Havranek (2015), Brodeur *et al.* (2016), Bruns & Ioannidis (2016), Stanley & Doucouliagos (2017), Christensen & Miguel (2018), Brodeur *et al.* (2020), and Blanco-Perez & Brodeur (2020). Ioannidis *et al.* (2017) show that publication bias looms large in economics and finance, exaggerating the mean reported coefficient twofold. That is not to say the bias always arises intentionally: for better or worse, many researchers use statistical significance as an implicit indicator of importance, and select the results for publication accordingly without an intention to inflate their results. A useful analogy appears in McCloskey & Ziliak (2019), who compare publication bias to the Lombard effect in psychoacoustics, when speakers involuntarily increase their effort with increasing noise. (For example, one would instinctively talk louder at a noisy party, without necessarily realizing so.) With large imprecision given by noisy data or inefficient estimation techniques, researchers may unintentionally try harder to search through potential specifications in order to find effects that are interesting, that produce sufficiently large  $t$ -statistics. Even if no intentional inflation of research results occurs, a correlation between estimates and the corresponding standard errors arises.

We test for publication bias using the recently developed techniques by Ioannidis *et al.* (2017, weighted average of adequately powered estimates), Andrews & Kasy (2019, selection model), Bom & Rachinger (2019, endogenous kink), and Furukawa (2020, stem-based technique), which are all, to some extent, based on the Lombard effect, but allow for a nonlinear relationship between the magnitude of publication bias and the size of the standard error. We also use a brand new p-uniform\* technique (van Aert & van Assen, 2021) that does not rely on the Lombard effect but uses the distribution of p-values. Our results based on these techniques suggest substantial publication bias. The corrected mean  $\beta$  estimates for the average currency vary between 0.2 and 0.8 depending on the method, far from the simple arithmetic average of  $-0.6$  computed from all the reported estimates. The pattern we observe in the literature is thus consistent with a breed of publication bias called confirmation bias, the tendency to publish results that corroborate the famous finding on the negative coefficient by Fama (1984) rather than estimates that point in the opposite direction. Thus, correcting the literature for publication bias makes the forward premium puzzle look much less puzzling than previously suggested.

We further explore how the published  $\beta$  estimates vary with the choice of data samples and estimation methodology. Various studies use data on different currency pairs sourced from different time periods, use different estimation techniques, and are published in journals of different reputation. For example, Froot & Thaler (1990) collected 75 such estimates published until the end of the 1980s and reported their mean to equal  $-0.88$ . Figure 1 shows that even without any correction for a potential publication bias the  $\beta$  estimates from the differences regression exhibit a tendency to increase over time, starting with values around  $-1$  in the 1980s and approaching values close to 0 at the end of the 2010s. To capture the context in which individual estimates are obtained, we collect 42 corresponding variables and then regress the reported estimates on these variables. Because of model uncertainty and collinearity inherent in such an exercise, we cannot place all the 42 variables into one regression, but have to use model

Figure 1: *Do estimates of  $\beta$  increase with time?*



*Notes:* The horizontal axis denotes the year when the first version of the paper appeared in Google Scholar. The solid line represents a linear trend and the surrounding shaded area shows the corresponding 95% confidence band. Only difference estimates (Equation 3) are included.

averaging techniques that run millions of different regressions with various combinations of the 42 variables and then weight these models according to fit and complexity. We employ both Bayesian (Eicher *et al.*, 2011; Steel, 2020) and frequentist (Hansen, 2007; Amini & Parmeter, 2012) model averaging.

Our results suggest that several data, method, and publication characteristics systematically affect the reported estimates. The most robust findings concern differences among individual currencies. The estimates for the currencies of emerging economies tend to be much larger than estimates for developed economies, even if we control for other features in which studies vary. This finding corroborates that of Frankel & Poonawala (2010), who also report that much less evidence exists for the forward premium puzzle in emerging economies compared to developed countries. Moreover, we also find substantially above-average estimates for the former French franc and Italian lira, while substantially below-average estimates for the euro, Japanese yen, and Swiss franc. Thus even among the currencies of developed countries the less risky ones (as measured by historical risk premia) are associated with more evidence of the puzzle. Taken together, our results support the conclusion of Frankel & Poonawala (2010) that a time-varying exchange risk premium does not represent a plausible explanation of the forward premium puzzle. Because larger risk premia are typically more volatile, the currencies of emerging countries and the riskier among developed countries should show more evidence for the puzzle if time-varying risk premia was an underlying explanation.

As the bottom line of our analysis we compute an implied  $\beta$  that uses all the available results reported in the literature but, aside from correcting for publication bias, gives more weight to estimates that are based on arguably more reliable and larger datasets, employ modern estimation techniques, and are published in the best journals. The implied  $\beta$  is constructed

using the parameters from model averaging and choosing values for each variable (for example, sample maximum for data size and the impact factor of the journal in which the study was published). We obtain an interval of (0.23, 0.45) for the currencies of developed economies and (0.95, 1.16) for emerging economies. Thus, exploiting the heterogeneity of published studies and correcting for the publication selection bias produces  $\beta$  estimates which suggest that for many currencies the forward premium puzzle is less puzzling than previously thought. For emerging economy currencies the estimated values are very close to the theoretical prediction of 1. For developed economy currencies the estimated values are well below 1; nevertheless, in contrast to the common interpretation of prior findings it is positive. Even after correcting for publication and misspecification biases we document negative  $\beta$  estimates the Swiss franc (-0.14, 0.08), the Japanese yen (-0.47, -0.25), and especially for the euro (-0.79, -0.57). Meta-analysis is thus no panacea, and there remains scope for other explanations to the puzzle than publication bias.

## 2 Testing Forward Rate Unbiasedness

In this section we briefly describe how the coefficient  $\beta$  is typically estimated in the literature; further details are provided in Section 5 and in the studies quoted in this section. We start with the straightforward theoretical relationship between forward and future spot rates. The forward rate should differ from the expected spot exchange rate by a premium  $rp_{t+k}$ , which is a compensation for the perceived risk of holding different currencies based on information available at time  $t$ . This can be written in logarithms as:

$$f_{t,t+k} = E_t s_{t+k} + rp_{t+k}, \quad (1)$$

where  $f_{t,t+k}$  is the forward value of the spot exchange rate  $s_t$  for a contract signed in period  $t$  that expires  $k$  periods in the future.

Since the expectation term in Equation 1 is not directly observable, researchers typically invoke rational expectations. Coupled with the assumption of risk neutrality we arrive at the following regression:

$$s_{t+k} = \alpha + \beta * f_{t,t+k} + v_t. \quad (2)$$

In practice, however, researchers often subtract  $s_t$  from both sides of Equation 2:

$$s_{t+k} - s_t = \alpha + \beta * (f_{t,t+k} - s_t) + v_t, \quad (3)$$

which has two benefits: i) both sides of the equation can now be typically considered stationary, ii) both sides also have an intuitive interpretation in percentage points, the left-hand side denoting depreciation, the right-hand side representing the forward discount. For the forward rate unbiasedness hypothesis to hold, and thus for the absence of the forward premium puzzle,  $\alpha$  should equal 0,  $\beta$  should equal 1, and  $v_t$  should be serially uncorrelated. Therefore, strictly speaking, testing of market efficiency involves testing a joint hypothesis. Nevertheless, the joint hypothesis is rarely tested in practice for the following three reasons. First, the slope coefficient



$\beta$  is what matters for the profitability of carry trades. Second, even from the academic point of view, most of the researchers are more interested in the slope coefficient than the intercept because the slope coefficient determines the forecasting power of forward rates. This forecasting power is important e.g. for the models used by central banks. Third, most researchers find a negative estimated coefficient for  $\beta$ , which allows them to reject the forward rate unbiasedness hypothesis without looking at the intercept. Because most researchers focus on the slope coefficient  $\beta$ , it is also our focus in this meta-analysis.

A large body of literature has had troubles confirming that  $\beta$  equals 1. What is more, researchers frequently find that  $\beta$  is zero or even negative (e.g., Backus *et al.*, 1993; Hai *et al.*, 1997; Bekaert, 1995; Byers & Peel, 1991; MacDonald & Taylor, 1990; McFarland *et al.*, 1994; Backus *et al.*, 2010). Froot & Thaler (1990), on the basis of 75 published regressions, compute that the average  $\beta$  is equal to  $-0.88$ . Moreover, under covered interest parity we have:

$$f_{t,t+k} - s_{t,k} = i_{t,t+k} - i_{t,t+k}^* \tag{4}$$

where  $i_{t,t+k}$  denotes the logarithm of one plus the interest rate paid on domestic assets for  $k$  periods while  $i_{t,t+k}^*$  applies to the rate paid on foreign assets. So Equation 3 can be also thought of as a test of uncovered interest parity, similarly rejected by many published studies.

The literature has attempted to explain the frequent rejection of the null hypothesis that  $\beta$  equals 1 in various ways. First, some authors attribute it solely to statistical considerations. It is only correct to regress the change in the spot exchange rate on the forward premium in Equation 3 if both variables are stationary. The forward premium thus needs to be integrated of order zero as well. Goodhart *et al.* (1997) attest that Equation 3 would be misspecified and the measured value  $\beta$  would be biased towards zero if forward premium was not  $I(0)$ . Crowder (1994) fails to reject the presence of a unit root in the forward premium series while Baillie & Bollerslev (1994) find forward premiums to be fractionally integrated processes. Second, Fama (1984) explains the frequent rejection of the null hypothesis by highly variable rational expectations risk premia in Equation 3. Numerous other studies support this result: for example, Domowitz & Hakkio (1985), Wolf (1987) and Baillie & Bollerslev (1989). On the other hand, Frankel (1982), Frankel (1986) and Frankel & Froot (1987) do not confirm the presence of significant risk premia and report instead that the empirical rejection of the unbiasedness hypothesis implies that expectations are generally not rational due to excessive speculations.

While the aforementioned explanations for the finding of the forward premium puzzle are the most prominent ones, others have been put forward in the literature. McCallum (1994) attributes the rejection of the hypothesis to the fact that monetary authorities aim to avoid sudden exchange rate changes and thus smooth interest rates. As a result, tests of forward unbiasedness suffer from the absence of an equation that would take into account the behavior of the monetary authority. Furthermore, the theory behind the forward rate unbiasedness does not indicate whether long-term or short-term interest rates should be used or whether using T-bill rates as opposed to rates on commercial papers should matter. To address this issue, Razzak (2002) performs the test using one-year forward exchange rates instead of one-month

rates and finds support for the null hypothesis when exchange rates are measured in US dollars. Nevertheless, no such support for the hypothesis emerges when other currencies are used for measuring exchange rates. In a similar vein, studies by Mussa (1979), Chinn & Meredith (2004), and Nadal De Simone & Razzak (1999) corroborate that long-term rates are more suitable for explaining the movements of spot exchange rates in tests of forward rate unbiasedness.

The exchange rate regime, time period, stage of a country’s development, and data contamination have also appeared to matter for the testing of the hypothesis. Flood & Rose (1996) offer evidence that negative estimates of  $\beta$  hold only for floating exchange rate regimes. Using data for the European Monetary System they show that a large part of the forward discount puzzle vanishes for fixed exchange rate regimes. Frankel & Poonawala (2010) show that for emerging market currencies there is a smaller bias in the estimated  $\beta$ . The coefficient is on average positive and never significantly less than zero. The study by Chiang (1988) indicates that empirical evidence against the forward rate unbiasedness hypothesis might disappear if the regression parameters are allowed to vary in time. As for data contamination, Cornell (1989) argues that estimates of the slope coefficient are biased towards  $\beta < 1$  due to data mismatch. He suspects that most studies do not find exactly the future spot exchange rate that corresponds to the forward rate in their data and proposes the use of a lagged forward discount as the right-hand-side variable in Equation 3 to deal with this problem. Using this technique Cornell (1989) cannot reject the forward rate unbiasedness for the Canadian dollar/US dollar rate. He, however, rejects the hypothesis for other currencies relative to the US dollar. Moreover, following the recommendation put forward by Cornell (1989), Bekaert & Hodrick (1993) investigate the unbiasedness hypothesis for the mark, pound and yen relative to the US dollar and find  $\beta$  significantly negative. Given the conflicting findings in the literature, it is surprising that no quantitative synthesis of the empirical evidence has been conducted. In the next section we describe the first step in such a meta-analysis, data collection.

### 3 Data

We use Google Scholar to search for studies testing the forward rate unbiasedness hypothesis. Google’s algorithm goes through the full text of studies, thus increasing the coverage of suitable published estimates, irrespective of the precise formulation of the study’s title, abstract, and keywords. This is the key advantage in contrast to other databases commonly used in research synthesis, such as the Web of Science. Our search query contains expressions “forward rate unbiasedness,” “forward premium puzzle,” “forward discount puzzle,” “forward premium anomaly,” “foreign exchange efficiency,” and “forward rate spot rate.” We terminate the search on July 31, 2018, and do not add any new studies beyond that date.<sup>1</sup> Our search strategy is

---

<sup>1</sup>As of February 1, 2021, we identify 6 new papers that have been published after July 31, 2018, and that conform to our search criteria. We do not add these studies to our dataset but briefly describe their results here. Cho & Chun (2019) argue that the puzzle can be caused by the nonstationarity of the forward premium. Their results suggest that if the premium is stationary, no puzzle appears. Haab & Nitschka (2020) find substantial forward premium puzzle for the Swiss franc, which is consistent with our results. Hassan & Mano (2019) study the related anomalies of dollar trade, carry trade, and forward premium puzzle. In contrast to most previous

described in more detail in Appendix A. We follow the reporting guidelines for meta-analysis in economics (Havranek *et al.*, 2020).


To be included in our dataset, a study must meet four criteria. First, the paper must be written in English. Second, at least one estimate in the study must originate from an equation regressing either  $s_{t+k}$  on  $f_{t,t+k}$  (Equation 2) or regressing  $s_{t+k} - s_t$  on  $f_{t,t+k} - s_t$  (Equation 3), as described in Section 2. That is, we do not collect estimates from studies that focus on uncovered interest parity and replace the forward discount with the interest rate differential. While such estimates are comparable to our dataset under the assumption of covered interest parity, the covered interest parity does not have to hold for all markets, especially after the financial crisis. Third, the study must be published. This criterion is mostly due to feasibility since even after restricting our efforts to published studies the dataset involves a manual collection of hundreds of thousands of data points; in any case, studies published in journals can be expected to contain fewer typos and be, on average, of higher quality due to peer review. Fourth, the study must report standard errors of the estimated  $\beta$  or other statistics from which the standard error can be computed. This requirement is necessary for tests of publication bias and for conducting weighted least squares. If the study assumes a non-linear relationship (for example, quadratic) between spot and forward rates, we use the marginal effect evaluated at the sample mean and compute the corresponding standard error using the delta method. If the study uses a regime-switching model, we collect estimates separately for each regime and control for the definition of the regime (for example, “high premium”) in our Bayesian model averaging analysis.

Using the search queries and the study inclusion criteria specified above, we obtain 3,643 estimates of the slope coefficient  $\beta$  from 91 published studies, which makes our paper one of the largest meta-analyses ever conducted in economics and finance. For the list of primary studies included in our meta-analysis, please see Appendix C. All data and codes are available in an online appendix at [meta-analysis.cz/forward](https://meta-analysis.cz/forward). To ensure that outliers do not drive our results, we winsorize the collected estimates and their standard errors at the 5% level. The main results, however, are not sensitive to the chosen level of winsorization. Figure 2 and Figure 3 show a box plot of the estimates for individual studies and currencies, respectively. We can observe the estimates vary greatly both within and across studies and countries, with most studies reporting both negative and positive estimates of  $\beta$ . The mean of all estimates is  $-0.34$ , which confirms that the finding of the forward premium puzzle is common in the literature. The histogram of all the estimated coefficients (the left panel of Figure 4) has two peaks: 0 and 1, while its left tail is much longer than the right tail, suggesting a relatively greater representation of negative than positive estimates in the literature. This seems to be in line with the prevalence of negative estimates of  $\beta$  reported in primary studies and could represent a type of confirmation bias, in which the findings consistent with the prevalent view

---


findings in the literature, they find that investors expect currencies with high interest rates to depreciate. Jetter *et al.* (2019) argue that shifts in monetary policy regimes explain the forward puzzle. Kumar (2020) also claims to find a solution to the puzzle, this time by controlling for liquidity shocks. Reed (2019) argues that the puzzle can be explained by using a Markov-switching model with adaptive learning. In general, these new studies corroborate the results of our meta-analysis: when properly estimated, the forward premium puzzle (and the associated opportunity for profitable carry trades) is smaller than commonly thought.

Figure 2: *Estimates vary both within and across studies*




*Notes:* The figure shows a box plot of beta estimates in individual studies. The length of each box represents the interquartile range (P25-P75), and the dividing line inside the box is the median value. The whiskers represent the highest and lowest data points within 1.5 times the range between the upper and lower quartiles. The vertical line denotes zero beta. For ease of exposition, outliers are excluded from the figure but included in all statistical tests.

Figure 3: *Estimates vary both within and across countries*



*Notes:* The figure shows a box plot of beta estimates for the currencies of individual countries. (If a euro area member country is shown in the figure, the estimate concerns its own currency prior to the adoption of the euro.) The length of each box represents the interquartile range (P25-P75), and the dividing line inside the box is the median value. The whiskers represent the highest and lowest data points within 1.5 times the range between the upper and lower quartiles. The vertical line denotes zero beta. For ease of exposition, outliers are excluded from the figure but included in all statistical tests.

Figure 4: *The distribution of the estimates*



*Notes:* The figure depicts histograms of the estimated beta for all estimates (on the left) and for the difference estimates (on the right, extracted from the differences equation (Equation 3)). The solid vertical lines indicate the mean of winsorized estimates. Extreme observations are cut from the figure but included (after winsorization for the most extreme ones) in all statistical tests.

are more likely to get selected for reporting and publication. In their survey, Jongen *et al.* (2008) observe that a negative association between currency depreciation and forward discount constitutes common wisdom in the literature.

In the preceding paragraph we discussed the distribution of all the estimates of  $\beta$  from the studies included in our dataset. Consequently, we did not differentiate between estimates originating from the two main tests of forward rate unbiasedness conducted in the literature, estimates computed in levels and differences. There are only 654 estimates obtained from Equation 2 in our dataset (level estimates), of which only 5.4% are negative and representing the puzzling result that forward rate is negatively related to the future spot rate. On the other hand, difference estimates extracted from Equation 3 are much more numerous: there are 2,989 of these estimates in our dataset (their distribution is depicted in the right-hand panel of Figure 4). Out of these 58.4% are negative, and the mean is  $-0.6$ , which is in line with the ongoing quest in the literature to explain the predominance of puzzling results. As we already discussed in the Introduction, the level estimates are problematic because of the likely unit root, and few modern studies use Equation 2. For this reason, in the remainder of the analysis we focus exclusively on the difference estimates.

Apart from the estimates of  $\beta$  and their standard errors, we collect 42 variables that capture different aspects of how the studies are designed. In consequence, we have to collect more than 150,000 data points from the primary studies. The data collection was performed by one of the co-authors while another double checked random portions of the data to minimize potential mistakes made during the data coding process.

Table 1 presents mean estimates of reported  $\beta$  for various currencies or groups of currencies. The reported mean overall coefficient is negative, in line with the forward rate unbiasedness puzzle. The most striking observation obtained from the table, however, is the apparent dif-

Table 1: *Results for different currencies vary widely*

	No. of obs.	Unweighted			Weighted		
		Mean	95% conf. int.		Mean	95% conf. int.	
Advanced currency	1,159	-0.507	-0.625	-0.388	-0.697	-0.792	-0.602
Emerging currency	407	0.364	0.247	0.482	0.759	0.613	0.905
Euro	102	-2.371	-2.891	-1.852	-2.261	-2.721	-1.800
German mark	181	-0.980	-1.293	-0.668	-0.937	-1.197	-0.676
British pound	269	-0.981	-1.244	-0.719	-1.003	-1.217	-0.789
French franc	133	-0.266	-0.639	0.108	-0.300	-0.587	-0.014
Italian lira	117	-0.272	-0.657	0.112	-0.169	-0.464	0.125
Swiss franc	168	-0.825	-1.149	-0.501	-1.048	-1.342	-0.754
Japanese yen	309	-1.587	-1.821	-1.352	-1.629	-1.813	-1.444
European currencies	1,434	-0.822	-0.940	-0.705	-0.834	-0.935	-0.734
Asian currencies	520	-0.922	-1.089	-0.755	-1.118	-1.272	-0.964
Other country currencies	416	-0.864	-1.051	-0.677	-0.572	-0.842	-0.303
All estimates	2,989	-0.602	-0.679	-0.526	-0.840	-0.905	-0.774

*Notes:* The table reports summary statistics of the reported estimates of beta from the differences specification (Equation 3). Weighted means are calculated using the inverse of the number of estimates reported per study as the weight. All variables are described in Table 3.

ference between the mean coefficients reported for the currencies of advanced and emerging economies, respectively. The mean estimate for the former is negative, while the mean estimate for the latter is positive, both being statistically significantly different from zero at the 5% level. Thus our data corroborate the results of Frankel & Poonawala (2010): the bias in the forward discount as a predictor of future changes in the spot exchange rate is less severe in emerging market currencies than in advanced country currencies. Frankel & Poonawala (2010) observe that the coefficient for emerging market currencies is on average slightly above zero, and even if negative, it is rarely significantly less than zero. Nevertheless, all of these results can be influenced by publication bias, an issue to which we turn next.

## 4 Publication Bias

Publication bias is the empirical observation in many sciences, including economics (Ioannidis *et al.*, 2017), that the reported results constitute a biased reflection of the universe of results obtained by researchers before they write up their papers. Why should reported results be biased? One reason is that statistical significance is sometimes perceived as evidence of scientific importance. This perception might represent a valid principal in some cases, but in general it means that the published results will exaggerate the true underlying effect unless the true effect is zero. Estimates that are, simply by chance, much larger than the true effect (in absolute value) will be statistically significant. Estimates that are, also by chance, much smaller than the true value will probably be insignificant. If statistical significance is taken for scientific importance, the former estimates will become overrepresented in the literature. Another reason for publication bias is a simple preference for a particular sign of the regression parameter, perhaps given by underlying theory or previous influential findings. A more precise label for the

problem is “selective reporting”, because there is no reason why the problem should be confined to journals: researchers write their working papers with the intention to publish. We use the term “publication bias” to keep consistency with previous studies on the topic.<sup>2</sup>

Our main identification assumption in testing for publication bias (we will relax the assumption later) is that in the absence of the bias the estimate and its standard error should be statistically independent quantities. This assumption follows from the properties of the econometric methods used to estimate  $\beta$  in the literature. In almost all the cases, the econometric techniques suggest that the ratio of the estimated  $\beta$  to its standard error has a symmetrical distribution (typically a  $t$ -distribution). In consequence, we should observe zero correlation between estimates and their standard errors. But if, intentionally or not, researchers prefer to report statistically significant results, they will, given a particular standard error, search for estimates that are large enough to bring the  $t$ -statistic above 2 in absolute value. This search can be conducted by choosing a subset of the entire dataset available, different estimation technique, or different control variables. A similar correlation between estimates and standard errors arises if estimates of a particular sign are discriminated against; the correlation follows from the observation that a regression of estimates on standard errors is heteroskedastic.

Figure 5 shows the traditional visual test, originating from medical research, of the correlation between estimates and standard errors (here precision, the inverse of standard error), and thus of publication bias. The figure is called a funnel plot, because in the absence of publication bias the observations should form a symmetrical inverted funnel. Intuitively, the most precise estimates should be close to the underlying mean value of the parameter in question, while less precise estimates should be more dispersed, giving rise to the funnel shape.<sup>3</sup> The funnel should be symmetrical because, in the traditional view, there is no reason for imprecise negative and positive estimates to have a different probability of publication. (But note that asymmetry can be also caused by small-sample bias or heterogeneity.) The first panel of Figure 5 shows two groups of estimates: red ones (lighter in grayscale) are derived from level regressions, spots on forwards. Blue ones (darker in grayscale) are derived from differences regressions, depreciation on the forward discount. The second panel uses only difference estimates. Two observations stand out. First, level estimates do not form a funnel, but are almost always very close to 1, with little dispersion irrespective to precision. The observation is consistent with level estimates being attracted to 1 via spurious regression. Second, the funnel for difference estimates is slightly asymmetrical, which is consistent with publication bias—but may be also caused by heterogeneity, which is suggested by the two peaks of the funnel. The most precise estimates are around zero or mildly positive, but many imprecise positive estimates seem to be missing from the literature, which leads to the overall observed mean of  $-0.6$ .

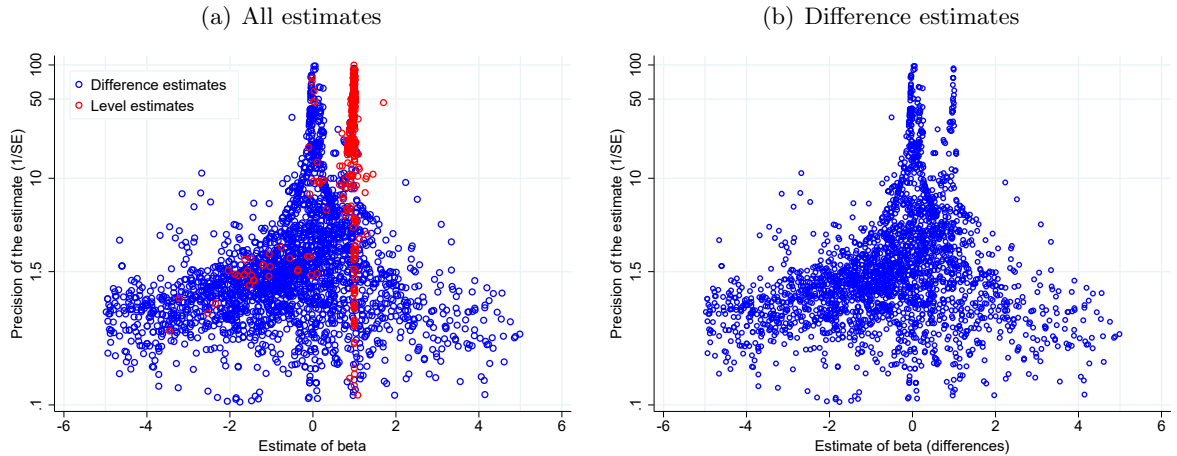
---

<sup>2</sup>Publication bias in economics has also been recently discussed, among others, by Astakhov *et al.* (2019), Bajzik *et al.* (2020), Cazachevici *et al.* (2020), Christensen & Miguel (2018), Havranek & Irsova (2017), Havranek *et al.* (2017), Imai *et al.* (2021), and Matousek *et al.* (2021).

<sup>3</sup>It is important to stress that all estimators we use are in fact weighted averages: all estimates of beta have a non-zero weight with the exception of the WAAP model by Ioannidis *et al.* (2017), which only assigns a non-zero weight to estimates that have power above 80%. But even this estimator does not focus on the single most precise estimate or study, which would surely be inefficient and produce volatile results.



Figure 5: *Funnel plots*



*Notes:* The figure depicts funnel plots for all estimates on the left and for the difference estimates on the right. In the absence of publication bias the scatter plot should resemble an inverted funnel that is symmetrical around the most precise estimates. Logarithmic scale is used to depict precision (the vertical axis). Red circles (lighter in grayscale) represent estimates extracted from the level equation (Equation 2); blue circles (darker in grayscale) show estimates from the differences equation (Equation 3). We use winsorized data, but because our winsorization is relatively conservative, for visual presentation we focus on estimates between  $-5$  and  $5$ . Outlying observations are included in all statistical tests.

In the next step we quantify the extent of publication bias numerically. As we have noted, if there is no bias, there should be no correlation between estimates and their standard errors. In the presence of bias, we will observe a correlation consistent with the Lombard effect, as researchers will (intentionally or not) increase their effort to find larger estimates of  $\beta$  in response to noise (Stanley, 2005):

$$\beta_{i,j} = \gamma_0 + \gamma_1 * SE(\beta_{i,j}) + \epsilon_{i,j}, \quad (5)$$

where  $\beta_{i,j}$  is the  $i$ -th slope coefficient estimate in study  $j$  collected from the differences specification, as detailed in Equation 3.  $SE(\beta_{i,j})$  is the standard error of this estimate.  $\gamma_1$  captures the severity of publication bias in the literature while  $\gamma_0$  measures the mean effect beyond bias. Nevertheless, the specification in Equation 5 is heteroskedastic since the right-hand-side variable (SE) captures the variance of the left-hand-side variable (coefficient estimate  $\beta$ ). To correct for this heteroskedasticity, we divide Equation 5 by the standard error of the estimate and obtain

$$t_{i,j} = \gamma_1 + \gamma_0 * \frac{1}{SE(\beta_{i,j})} + \omega_i, \quad (6)$$

where  $t_{i,j}$  is  $t$ -statistic of the  $i$ -th estimate of  $\beta$  from study  $j$ ,  $\gamma_1$  captures publication selectivity, and  $\gamma_0$  measures the corrected effect beyond bias (the mean  $\beta$  conditional on maximum precision, and therefore no publication bias—but note again that the estimated  $\gamma_0$  can be rewritten as a weighted average of all the reported estimates and does not rely merely on the most precise ones). The weighted specification has the additional allure of giving more precise estimates greater weight. Along with this inverse-variance weight traditionally used in meta-analysis, we also use a weight that assigns the same importance to each study. This latter issue is typically

not discussed in meta-analysis, because traditional meta-analyses (for example, in medicine or psychology) often use only one observation per study. In this paper we attempt to use all published information, and while some studies report only a couple of estimates, other studies report many dozen. We prefer to use all of this information but also do not give more weight in meta-analysis to studies that report many estimates. The solution is to add a weight based on the inverse of the number of estimates produced in each study. We also use bootstrapped standard errors instead of clustering at the level of individual studies. Because the number of estimates reported per study varies greatly, clustering does not work well in this case. Indeed, clustering produces larger standard errors than a regression using only median estimates per study, as noted by an anonymous referee who recomputed our results. Nevertheless, clustering standard errors would not change our results qualitatively.

Table 2: Linear and nonlinear techniques detect publication bias

<i>Panel A: FAT-PET</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.611 <sup>***</sup> (0.111)	0.605 <sup>***</sup> (0.0985)	0.332 <sup>*</sup> (0.182)	
Publication bias (Constant)		-2.057 <sup>***</sup> (0.332)	-2.034 <sup>***</sup> (0.256)	-1.040 <sup>**</sup> (0.509)	
Observations		2,989	2,989	2,989	
<i>Panel B: PEESE</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.622 <sup>***</sup> (0.105)	0.552 <sup>***</sup> (0.0768)	0.267 <sup>*</sup> (0.155)	
Publication bias (SE)		-0.362 <sup>***</sup> (0.0575)	-0.197 <sup>***</sup> (0.0236)	-0.510 <sup>*</sup> (0.295)	
Observations		2,989	2,989	2,989	
<i>Panel C: Advanced</i>	WAAP	Kinked model	Selection model	p-uniform <sup>*</sup>	Stem method
Mean beyond bias	0.288 <sup>***</sup> (0.029)	0.306 <sup>***</sup> (0.009)	0.222 <sup>***</sup> (0.025)	0.775 <sup>***</sup> (0.127)	0.791 <sup>**</sup> (0.313)
Observations	2,989	2,989	2,989	2,989	2,989

*Notes:* FE = study-level fixed effects, WLS = weighted least squares, IV = instrumental variables regression with the inverse of the square root of the number of observations used as an instrument of the standard error. Only difference estimates (Equation 3) are included. Panel A & B: weighted by both inverse variance and the inverse of the number of estimates reported per study; bootstrapped standard errors in parentheses. Panel C: WAAP (weighted average of adequately powered, Ioannidis *et al.*, 2017), kinked model (Bom & Rachinger, 2019), selection model (Andrews & Kasy, 2019), p-uniform<sup>\*</sup> (van Aert & van Assen, 2021), stem method (Furukawa, 2020). <sup>\*</sup>  $p < 0.10$ , <sup>\*\*</sup>  $p < 0.05$ , <sup>\*\*\*</sup>  $p < 0.01$ .

Panel A in Table 2 shows the results of running Equation 6. We prefer including dummy variables for each study, which controls for unobserved study-level characteristics that can be related to quality; the resulting specification is labeled “FE” for fixed effects. This specification also has disadvantages, because it throws away a lot of information and only exploits within-study variation. But on balance it is important to filter out variation that can yield biased estimates of meta-regression parameters, even if that means a decrease in efficiency. We observe that in this preferred specification publication bias is robustly negative, and the mean

$\beta$  corrected for the bias is around 0.6. The results are similar for a weighted least squares specification without study dummies. Nevertheless, some method choices in the primary studies could potentially influence the estimates and their standard errors in the same direction, which would make our estimates spurious. Moreover, the standard error is itself estimated with some error, which means that our baseline regression will suffer from attenuation bias. Therefore, as an important robustness check, we use the inverse of the square root of the number of observations as an instrument for the standard error, an estimator introduced by Stanley (2005). This quantity is correlated with the standard error by definition, but should not be much correlated with method choices (of course with exceptions because some methods are designed for data sets of particular size). Our results still show negative publication bias, but with much less precision, and the estimated mean effect beyond publication bias is only 0.3.

So far we have assumed that publication bias is linearly proportional to the size of the standard error. But in principle, this does not have to be the case. Intuitively, the linear relationship works well with intermediate values of the standard error. For example, given a true underlying effect of 0.5 and a standard error of 0.3, a researcher searching for a significant coefficient will need to find a specification in which the point estimate is 0.6 (simply by chance) in order to produce a t-statistic of 2 and statistical significance at the 5% level. If the standard error (which is given by noise in the data, for example) is 0.4, the researcher will have to produce a point estimate of 0.8, and so on, which means that the linear specification approximates publication selection perfectly. But if the standard error is small (say, 0.1), with a true underlying effect at 0.5 the researcher will find a significant effect in almost all specifications. A further decrease of the standard error to 0.05 will have little effect on the extent of publication bias, and the linear relationship breaks down. A similar problem arises with very large standard errors, for which it is impossible to plausibly generate large enough point estimates and report t-statistics above 2. More formally, the bias of linear estimators is shown by Stanley (2008) and Stanley & Doucouliagos (2014).

As a next step, we estimate a quadratic model of publication bias, the so-called PEESE (precision-effect estimate with standard error) developed by Stanley & Doucouliagos (2012) and Stanley & Doucouliagos (2014):

$$t_{i,j} = \gamma_1 * SE(\beta_{i,j}) + \gamma_0 * \frac{1}{SE(\beta_{i,j})} + \xi_i, \quad (7)$$

where  $\gamma_1$  again captures publication bias in the literature, while  $\gamma_0$  denotes the mean  $\beta$  corrected for publication bias, i.e. the slope coefficient when regressing the changes in the spot exchange rate on the forward discount. Panel B in Table 2 presents the results of this specification, which are similar to those of the linear specification. The corrected mean estimate for the fixed effects and weighted least squares specifications lies between 0.55 and 0.62. The IV estimation again brings coefficients which display the same signs, but are smaller and less precise.

Techniques more sophisticated than PEESE have recently been developed. In Panel C of Table 2 we use a battery of these advanced tests to evaluate the robustness of our results regarding the mean  $\beta$  corrected for publication bias. We employ five methods: the weighted

average of adequately powered estimates by Ioannidis *et al.* (2017), the stem-based method by Furukawa (2020), the selection model by Andrews & Kasy (2019), the p-uniform\* technique by van Aert & van Assen (2021), and the endogenous kink technique by Bom & Rachinger (2019). First, Ioannidis *et al.* (2017), using a survey of more than 60,000 estimates published in economics, find that the median statistical power among the published results in economics is 18%. They investigate how power is associated with publication bias and propose a correction technique that employs the estimates with power above 80%. Furthermore, using Monte Carlo simulations, Stanley *et al.* (2017) show that this technique outperforms the commonly used meta-analysis estimators. The results of their model are close to that of our IV estimation.

Second, Bom & Rachinger (2019) account for the case that estimates get reported only when they cross a particular precision threshold. In their method they estimate this threshold and introduce an “endogenous kink” to extend the linear test of publication bias. The technique gives us an estimate of 0.31, again close to the results of the IV estimation. Third, Andrews & Kasy (2019) apply the observation reported by many researchers (e.g. Brodeur *et al.*, 2016) that standard cut-offs for the p-value are associated with jumps in the distribution of reported estimates. Consequently, they build on Hedges (1992) in order to construct a selection model of publication probability for each estimate in the literature given its p-value. Using their technique, we obtain an estimate of 0.22, which is statistically significant at the 1% level.

Fourth, we use a new method developed in psychology by van Aert & van Assen (2021), p-uniform\*. The advantage of this method is that it does not need the assumption of zero correlation between estimates and standard errors in the absence of publication bias. The disadvantage is that the method can only use one estimate per study, and we use the medians. The technique works with the distribution of p-values and looks for an underlying effect for which the distribution is uniform (using the statistical principle that the distribution of p-values is uniform at the true effect). Here we obtain an estimate a bit larger than that of FE and WLS. Finally, the approach by Furukawa (2020) relies on the assumption that the most precise estimates suffer from little bias: authors do not find it difficult to produce statistically significant estimates when the standard errors are very small. Previous researchers in meta-analysis focused on a fraction of the most precise estimates in meta-analysis—such as the top ten method by Stanley *et al.* (2010). Furukawa (2020) finds a new way of estimating this fraction of estimates by exploiting the trade-off between bias and variance. His technique delivers a large estimate of the mean  $\beta$  corrected for publication bias, 0.8.

On balance, we find strong evidence of publication bias in the literature on the forward rate unbiasedness hypothesis. All 5 recently proposed techniques suggest that the mean  $\beta$  corrected for publication bias is positive, which contrasts with the naive mean of  $-0.6$  obtained by averaging the reported estimates of  $\beta$ . In the next part of the manuscript we consider the effects of differences in study design on the reported estimates. Explicitly controlling for heterogeneity is also important for a proper identification of publication bias. Nevertheless it does not mean that the estimate of publication bias in the next section is superior to the ones presented in this section. First, in this section we use study-level fixed effects, which filter out

all unobserved study-level characteristics. While in principle we can add fixed effects to the full heterogeneity model as well, in practice such a model would be of little use because many of the variables that we will consider are defined at the study level or have limited within-study variation. Second, in this section we use the instrumental variable estimator, which is unfeasible in weighted Bayesian model averaging with the dilution prior. Even though in the next section we explicitly control for dozens of variables, we might have missed some characteristics that affect the estimated parameter for publication bias, and thus the IV specification is important. Third, in Bayesian model averaging estimation in the next section we cannot consider non-linear estimators (with the potential exception of PEESE). So, in our view, the next section is important for the examination of various sources of heterogeneity in the literature, but regarding publication bias it provides merely a further robustness check.

## 5 Heterogeneity

We aim to identify the aspects of estimation context that systematically influence the reported estimates of the slope coefficient  $\beta$  from the differences specification of the forward rate unbiasedness hypothesis tests presented in Equation 3. To this end, we collect 42 variables that reflect country scope, data characteristics, estimation characteristics, regimes capturing different market conditions, databases used, and publication characteristics. In the following paragraphs we describe these variables.

### 5.1 Variables

**Country scope** Previous studies have hinted on potential differences between the reported  $\beta$ 's for countries in different stages of development. For instance, Bansal & Dahlquist (2000) observe that the puzzling finding of a negative  $\beta$  is systematically related to the use of data from advanced economies. Summary statistics of our data corroborate this finding: the mean  $\beta$  for advanced economy currencies equals  $-0.51$ , while the mean for developing economy currencies is  $0.36$ , both significantly different from zero. For this reason we include dummy variables for advanced economy currencies (38.8% of the estimates), emerging economy currencies (13.6%), estimates arising from both advanced and developing currency samples (4.8%), estimates specifically obtained for the former German mark (6.1%), French franc (4.4%), British pound (9%), Italian lira (4%), Japanese yen (10.3%), Swiss franc (5.6%), euro (3.4%), and for currencies from different geographical regions: Europe (48%), Asia (17.4%), non-European and non-Asian countries (14%), and for estimates from geographically mixed datasets (20.7%). In addition, we control for the numeraire currency against which other currencies are tested. Flood & Rose (1996) find a larger  $\beta$  for European currencies tested versus the German mark than for those tested versus the USD. Bansal & Dahlquist (2000) hypothesize that this finding could be due to the fact that the economies within the European Monetary System synchronized their monetary policy with Germany. The proportions of estimates in our dataset with different numeraire

currencies are the following: USD (70.7%), British pound (10.5%), Japanese yen (2%), Swiss franc (1.4%), euro (8.4%), German mark (7%).

**Data characteristics** This category comprises additional information regarding the data that were used to produce the estimates of  $\beta$ . In particular, we collect information on the forward rate horizon; i.e., the number of periods  $k$  after which settlement will occur, the frequency of the data used in the estimation, the size of the time series component of the data (the number of time periods used for estimation), the size of the cross-sectional component (the number of currencies included in a panel), and sources of data in primary studies. As for the horizon of the rates used to test the unbiasedness hypothesis, Razzak (2002) finds support for the unbiasedness hypothesis when one-year forward exchange rates are used instead of one-month contracts. His finding is consistent with the literature on uncovered interest rate parity, for which some authors also find support when using long-term interest rates (e.g., Chinn & Meredith, 2004; Nadal De Simone & Razzak, 1999). Next, we collect dummies for the frequencies of data as follows: daily (12.3%), weekly (13%), monthly (72%), quarterly (2.4%), and other frequency (0.2%). We further account for the so-called overlapping samples problem, where data frequency is higher than the maturity horizon of instruments, which introduces serial correlation in the error term of the regression. In our sample, 38% of the estimates originate from data that suffer from this problem, and many authors introduce corrections in the methods they apply in order to mitigate the issue. For instance, Goodhart (1988) applies an adjustment to the OLS covariance matrix proposed by Hansen (1982).

**Estimation** Researchers apply different estimation techniques to test the validity of the forward rate unbiasedness hypothesis. The most frequently used method is OLS regression (66.5% of estimates) followed by the regime switching model (13.8%), seemingly unrelated regressions (7.1%), and fixed effects regression (3.1%). Other methods, such as instrumental variables regression, error correction model and maximum likelihood, are also used, and generate 9.4% estimates of  $\beta$ . Some researchers advocate the use of the seemingly unrelated regressions technique which allows them to account for cross-correlation across currencies in their samples, arising from, among other things, the use of a common numeraire currency. For instance, Fama (1984) applies both OLS and seemingly unrelated regressions and finds that joint estimation of  $\beta$  improves the precision of the estimate. Moreover, he observes that the estimated slope coefficients from seemingly unrelated regressions are closer to zero compared to the estimates from OLS.

Other studies associate the existence of the forward premium anomaly with different market conditions or regimes. For this reason, they apply a variety of regime switching techniques to model this transition. For instance, Baillie & Chang (2011) apply logistic smooth transition dynamic regression models with interest differentials and foreign exchange market volatility as transition variables between two regimes. In one regime, they observe exchange rate movements that are characterized by persistent deviations from the uncovered interest rate parity, while in the other regime reversions to the parity occur. They show that the forward premium anomaly

ceases to manifest when foreign exchange market volatility is high. Moreover, we account for any additional variables that researchers may include in their regressions. For this reason, we control for the inclusion of lagged values of the forward rate (0.13%), interest rate differentials (0.13%), forward discount terms to the power of two and three (2.7%), and other controls (1.7%). Since these categories of controls do not comprise enough cases separately, we aggregate them into one variable “Controls”. Overall, 4.5% of the estimates in our dataset include additional control variables.

We also account for the units in which Equation 3 is specified. Typically, the forward discount and the change in the spot exchange rate are expressed as the difference between the natural logarithm of the forward rate and the natural logarithm of the spot rate at time  $t$ , and the difference between the natural logarithm of the spot rate at  $t+k$  and the logarithm of the spot rate at time  $t$ , respectively. This specification is the most frequent in the literature: over 95% of the estimates of  $\beta$  are obtained from this specification. Alternatively, about 5% of the estimates arise from the specification where the change in spot rates and the forward discount are expressed as a percentage of the spot rate at time  $t$ . More precisely, the percentage change in the spot exchange rate is expressed as  $\frac{(S_{t+k}-S_t)}{S_t}$  and the change in the forward discount is written as  $\frac{(F_t-S_t)}{S_t}$ , where  $S_t$  and  $F_t$  is the spot and the forward rate at  $t$ , respectively.

**Regimes** Researchers have investigated whether the apparent presence of the forward rate bias is subject to different market conditions; that is, if the unbiasedness hypothesis holds in some so-called regimes while it is violated in others. For instance, Baillie & Kiliç (2006) find using the logistic smooth transition dynamic regression model that the forward premium anomaly is more likely to occur during the periods of high volatility in US money growth while the periods of relative stability in terms of US money growth volatility are associated with forward rate unbiasedness. They also find that the growth of foreign money relative to US money leads to a higher likelihood of unbiasedness condition not being rejected. Therefore, money supply differentials serve as important transition variables between regimes in their study. Furthermore, Grossmann *et al.* (2014), using a sample of advanced economy currencies vis-a-vis the euro and the British pound, find that a significant forward premium anomaly exists for advanced country currencies during crisis periods when the numeraire currency sells at a premium or is overvalued. On the other hand, Zhou & Kutan (2005) do not find evidence for any forward premium asymmetry between the US dollar and the six currencies in their sample between 1977 and 1998. We include controls for large and positive forward premium (equals one for 5.7% of the estimates in our sample), negative forward premium (5.9%), overvalued currency (3.5%), undervalued currency (3.5%), large differential, which comprises controls for positive interest differential and other positive differentials (e.g., money growth differentials), for 3.3% of estimates, small differential, which includes controls for negative interest rate differential and other small/negative differentials, altogether for 3.5% of estimates, and a dummy for other regimes, which comprises both high and low foreign exchange volatility regimes among others, for 5.7% of collected estimates.

**Data sources** We control for the different sources of the data that researchers employ in primary studies. For example, some researchers advocate the use of survey data to address the issue whether the forward premium is due to systematic expectational errors or the risk premium (e.g., Froot & Frankel, 1989). Researchers use Datastream (51% of estimates), various bank data sources (14%), Data Resources, Inc. (3.2%), survey data (1.5%), and other minor sources (30%) to calculate the estimates of  $\beta$  reported in primary studies. We hypothesize that idiosyncratic private data sources (in contrast to standard public sources such as, for example, Datastream) can be subject to more random mistakes and measurement error, which would give rise to attenuation bias in the estimates of beta.

**Publication characteristics** To capture the potential aspects of study quality which are not reflected by the differences in data and methods across studies outlined above, we include three study-level variables: the year when the first draft of the paper appeared in Google Scholar (we opt for this measure instead of the nominal publication year due to increasing publication lags in economics and finance), the recursive discounted impact factor of the journal from RePEc, and the number of citations per year since the paper first appeared in Google Scholar. Table 3 presents the summary statistics of the aforementioned variables. Of course, these variables are rough proxies, but *ceteris paribus* we expect newer studies to be of higher quality (given by unobserved improvements in data and methods). We also expect that studies of higher quality are published on average in better outlets and receive more citations.

Table 3: Description and summary statistics of meta-regression variables

Variable	Description	Mean	SD	WM
Beta	The reported estimate of the coefficient beta.	-0.60	2.13	-0.84
Standard error	The reported standard error of the coefficient beta.	2.29	3.85	1.39
<i>Country scope</i>				
Advanced currencies	=1 if advanced-economy currencies are used.	0.39	0.49	0.32
Emerging currencies	=1 if emerging-economy currencies are used.	0.14	0.34	0.08
German mark	=1 if German mark is used.	0.06	0.24	0.09
French franc	=1 if French franc is used.	0.04	0.21	0.05
British pound	=1 if British pound is used.	0.09	0.29	0.15
Italian lira	=1 if Italian lira is used.	0.04	0.19	0.03
Japanese yen	=1 if Japanese yen is used.	0.10	0.30	0.15
Swiss franc	=1 if Swiss franc is used.	0.06	0.23	0.07
Euro	=1 if euro is used.	0.03	0.18	0.03
Mixed currencies	=1 if mixed currencies are used (reference category for different currencies).	0.05	0.21	0.02
European currencies	=1 if European currencies are used.	0.48	0.50	0.53
Other country currencies	=1 if non-European and non-Asian currencies are used.	0.14	0.35	0.21
Reference country currencies	=1 if Asian or multiple-country currencies are used (reference category for geographical location of currencies).	0.38	0.49	0.26
British pound base	=1 if British pound is used as the numeraire currency.	0.10	0.31	0.03
Euro base	=1 if euro is used as the numeraire currency.	0.08	0.28	0.01
German mark base	=1 if German mark is used as the numeraire currency.	0.07	0.26	0.05

Continued on next page



Table 3: Description and summary statistics of meta-regression variables (continued)

Variable	Description	Mean	SD	WM
Other base	=1 if other currency is used as the numeraire currency (reference category for currency bases).	0.74	0.44	0.91
<i>Data characteristics</i>				
Shorter horizon	=1 if forward contract maturity is less than one month.	0.06	0.23	0.03
1month horizon	=1 if forward contract maturity is one month.	0.65	0.48	0.70
1month-1year horizon	=1 if forward contract maturity is between one month and one year.	0.19	0.39	0.21
1year horizon	=1 if forward contract maturity is one year.	0.08	0.27	0.05
Longer horizon	=1 if forward contract maturity is more than one year (reference category for different horizons).	0.03	0.16	0.02
Daily frequency	=1 if data frequency is daily.	0.12	0.33	0.10
Weekly frequency	=1 if data frequency is weekly.	0.13	0.34	0.15
Monthly frequency	=1 if data frequency is monthly.	0.72	0.45	0.68
Other frequency	=1 if data frequency is quarterly or lower (reference category for different data frequencies).	0.03	0.16	0.08
Time difference	The logarithm of the number of observations in the forward contract maturity horizon.	1.01	1.30	0.76
Number of currencies	The logarithm of the number of currencies used in the estimation.	0.55	1.06	0.15
Overlapping problem	=1 if the overlapping samples problem is present.	0.38	0.49	0.31
<i>Estimation</i>				
OLS method	=1 if ordinary least squares method is used.	0.66	0.47	0.75
Fixed-effects method	=1 if fixed-effects regression is used.	0.03	0.17	0.02
SUR method	=1 if seemingly unrelated regression is used.	0.07	0.26	0.07
Regime-switching model	=1 if regime switching/transition model is used.	0.14	0.34	0.05
Other technique	=1 if other estimation techniques are used (reference category for different estimation techniques).	0.10	0.29	0.10
Controls	=1 if there are additional control variables included.	0.04	0.21	0.05
Spot rate percentage	=1 if spot rate change and forward premium are expressed in percentage of the spot rate.	0.05	0.22	0.07
<i>Regimes</i>				
Large differential	=1 if estimation period is characterized by large differentials (in interest rates, money growth, etc.).	0.03	0.18	0.01
Small differential	=1 if estimation period is characterized by small differentials (in interest rates, money growth, etc.).	0.03	0.18	0.01
Large positive premium	=1 if estimation period is characterized by a positive forward premium.	0.06	0.23	0.02
Low negative premium	=1 if estimation period is characterized by a negative forward premium.	0.06	0.24	0.02
Overvalued currency	=1 if estimation period is characterized by overvaluation of the currency.	0.03	0.18	0.00
Undervalued currency	=1 if estimation period is characterized by undervaluation of the currency.	0.03	0.18	0.00
<i>Data sources</i>				
Datastream source	=1 if data from Datastream are used in the estimation.	0.51	0.50	0.29
Bank data source	=1 if data from various bank sources are used in the estimation.	0.14	0.35	0.19
DataResources source	=1 if data from Data Resources, Inc., is used in the estimation.	0.03	0.18	0.13
Other source	=1 if other data sources are used in the estimation (reference category for different data sources).	0.32	0.47	0.40
<i>Publication characteristics</i>				
Impact factor	The recursive discounted impact factor from RePEc database.	0.46	0.61	0.51

Continued on next page

Table 3: Description and summary statistics of meta-regression variables (continued)

Variable	Description	Mean	SD	WM
Citations	The logarithm of the number of Google Scholar citations normalized by the number of years since the first draft of the paper appeared in Google Scholar.	1.70	1.80	1.72
First-draft year	The year when the first draft of the study appeared in Google Scholar normalized by the year of the earliest publication in our sample.	30.11	8.03	26.24

*Notes:* SD = standard deviation, WM = mean weighted by the inverse of the number of estimates reported per study. The impact factor was extracted from the RePEc database and the number of citations from Google Scholar. The remaining variables were collected from studies estimating  $\beta$ .

## 5.2 Methodology

To investigate which variables systematically explain the differences among the reported estimates of  $\beta$  extracted from Equation 3, the natural method would be to regress the reported  $\beta$ s on the variables capturing the context in which  $\beta$ s are calculated in primary studies. In other words, one wishes to estimate the following equation:

$$\beta_{ij} = \gamma_0 + \gamma_1 * SE(\beta_{ij}) + \sum_{l=1}^{42} \gamma_l * X_{l,ij} + \zeta_{ij}, \quad (8)$$

where  $\beta_{ij}$  denotes estimates of the slope coefficient  $i$  in study  $j$  obtained from regressing changes in spot foreign exchange rates on the forward premium, as detailed in Equation 3.  $X_{l,ij}$  represents the set of control variables that we introduced in Subsection 5.1,  $\gamma_1$  measures the severity of publication bias conditional on the inclusion of controls, and  $\gamma_0$  is the mean  $\beta$  estimate corrected for publication bias but also conditional on the variables included in  $X$ .

Nevertheless, including all the variables in  $X$  into one regression is problematic because of model uncertainty. While there is a strong rationale to include some of the variables, there are others which we would like to include as controls because they can also affect the slope coefficient  $\beta$  but for which little theory exists that would firmly justify their inclusion ex ante. Estimating Equation 8 as a single regression would negatively affect the precision of the coefficient estimates due to a large number of variables. There are several ways one can approach this problem, the most commonly traveled one being stepwise regression. Nevertheless, sequential  $t$ -testing does not properly account for the conditionality on the results of the previous  $t$ -test, and could thus accidentally exclude a useful variable at some stage. A more appropriate solution to model uncertainty is Bayesian Model Averaging (BMA). For more details on the technique, see Eicher *et al.* (2011) and Steel (2020). BMA has been used in meta-analysis, for example, by Havranek *et al.* (2015); Zigraiova & Havranek (2016); Havranek *et al.* (2018a,b,c).

BMA estimates many regressions using different subsets of the variables from the model space. In our case, since we consider 42 variables, this yields  $2^{42}$  possible models to estimate. To run all the models would be infeasible even with a modern computer. For this reason, we use Markov chain Monte Carlo (MCMC; Madigan *et al.*, 1995) algorithm that approximates

the model space and walks through the part of the model space that contains models with the highest posterior model probabilities (PMP). In frequentist terms, PMP is an analogue to information criteria, thus measuring how well the model performs compared to other models of similar complexity. BMA reports the posterior mean coefficient and posterior standard deviation of the coefficient, which are based on the weighted average of the coefficients from all the estimated models with weights being the PMP. Furthermore, for each variable BMA reports its posterior inclusion probability (PIP), which is equal to the sum of the PMPs of all the models in which this variable is included. In the baseline specification we apply the unit information g-prior (which gives the prior the same weight as one observation of the data), and the dilution prior (George, 2010), which re-weights each model by the determinant of the correlation matrix of the included variables. We apply alternative priors as well and report the results in Appendix B, and the main results are compared in Figure 7. Note that the hyper-g prior typically produces larger PIPs, and what matters for the comparison is the order of variables according to PIPs.


Regarding the thresholds for posterior inclusion probability, we build on Jeffreys (1961), who categorizes the values of PIP below 0.5 as irrelevant, between 0.5 and 0.75 as weak, values between 0.75 and 0.95 as positive, values between 0.95 and 0.99 as strong, and values above 0.99 as showing decisive evidence of an effect. In both OLS and fixed effects we cluster standard errors at the level of studies. Last but not least, we run frequentist model averaging (FMA). In our implementation of FMA we use Mallows' criteria as weights since they were shown to be asymptotically optimal (Hansen, 2007). Nevertheless, by using a frequentist approach there is no immediate alternative to the MCMC, and we find it infeasible to estimate all the  $2^{42}$  models. For this reason we follow Amini & Parmeter (2012) and orthogonalize the covariate space.

### 5.3 Results

The results of BMA are first summarized visually in Figure 6. The vertical axis shows the explanatory variables ordered by their PIP from the highest to the lowest value, and the horizontal axis shows individual models ordered by their posterior model probabilities; the best models are on the left. The blue color indicates that a variable is included in the model and its estimated coefficient sign is positive, while the red color stands for a negative coefficient sign. Blank cells indicate a variable is not included in the corresponding model.

Overall, there are 20 variables that can explain the variation in the reported estimates; 18 of them have PIP above 0.75, which means there is at least positive evidence for their effect on the estimated coefficient. Table 4 presents the numerical results of the BMA exercise as well as the results of the complementary frequentist model averaging. The posterior mean in Bayesian model averaging (or alternatively the estimated coefficient in frequentist model averaging) denotes the marginal effect of a study characteristic on the estimate of beta reported in the literature. For example, the value of the posterior mean of  $-0.66$  for Japanese yen means that, *ceteris paribus*, the reported betas for the Japanese yen are typically 0.66 smaller than the reported betas for the representative category in our estimation (a mixed sample of

Figure 6: *Model inclusion in BMA*



*Notes:* The figure depicts results of the baseline BMA model, where the response variable is an estimate of slope coefficient  $\beta$  from Equation 3. We employ the prior suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data), and the dilution prior suggested by George (2010), which accounts for collinearity. Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities for the 5,000 best models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model. All variables are described in Table 3.

developed and developing currencies). That is, the forward rate unbiasedness hypothesis is more likely to be violated for the Japanese yen than even for other currencies of developed countries (for example, the British pound, for which we find a marginal effect of zero). Note that the standard error (a proxy for publication bias) remains important, both in Bayesian and frequentist model averaging. The estimated magnitude of publication bias is similar to that reported by the instrumental variable specification from the previous section. In the next paragraphs we describe the results for individual variables.

**Country scope** The choice of the currency has strong implications for the estimated  $\beta$  coefficient. Estimates obtained from testing the unbiasedness hypothesis on emerging economy currencies are substantially larger than the estimates for advanced economies, and in BMA the effect is classified as *decisive* for explaining  $\beta$ . Larger estimates are also reported for the former French franc and Italian lira, but to a lesser degree. For some of the advanced economy currencies, especially the yen and euro, the reported  $\beta$  estimates are even smaller than the mean for advanced economies as a whole. The choice of the numeraire currency also has an impact on the reported  $\beta$ . Our results indicate that employing euro as the numeraire currency increases the reported coefficient by about 0.7 on average. The finding of a smaller forward bias in the estimates of  $\beta$  for emerging country currencies corroborates the results of Frankel & Poonawala (2010), who show that the coefficient for these currencies is on average positive and never significantly less than zero.

**Data characteristics** Our results suggest that if the spot rates and forward rates are sampled with weekly frequency the estimated  $\beta$  tends to be larger. Similarly, if one-month horizon rates (in contrast to longer horizons) are used to test the forward unbiasedness hypothesis the reported coefficients are smaller by around 0.26. In a similar vein, Nadal De Simone & Razzak (1999) find that long-term interest rates can explain a larger portion of the spot exchange rate movements.

**Estimation method** Our results suggest that estimates arising from regime switching models tend to be on average much smaller than those from seemingly unrelated regressions. In addition, simple OLS and panel fixed effects models yield estimates that are on average larger than those from the regime switching approaches but smaller than those from seemingly unrelated regressions. This is in line with Fama (1984), who applies both OLS and seemingly unrelated regressions to test for the unbiasedness hypothesis and finds that the estimated slope coefficients from the seemingly unrelated regressions are closer to zero (less negative) compared to the estimates from OLS. The omitted estimation category comprises other techniques, most prominently instrumental variables. Thus the results are consistent with the observation that methods that try to account for potential endogeneity tend to yield larger estimates.


**Regimes** We find that periods characterized by large differentials between (for example) interest rates or money growth coincide with less forward premium bias. The variable is decisive for explaining the reported  $\beta$ , and the estimated impact of this method choice is 2.8. Studies by Baillie & Kiliç (2006) and Baillie & Chang (2011) are consistent with our results as they observe that in periods characterized by large money supply differentials, large interest rate differentials, and foreign money growth in excess of US money growth, the uncovered interest rate parity is likely to hold. In addition, we find weak evidence for  $\beta$  being larger when the forward premium is negative. The existence of an asymmetric effect is confirmed also by Grossmann *et al.* (2014), who find that a significant forward premium anomaly exists for advanced country currencies when the numeraire currency sells at a premium.

Table 4: *Why do estimates of beta vary?*

Response variable: Estimated beta	Bayesian model averaging (baseline model)			Frequentist model averaging (robustness check)		
	P. mean	P. SD	PIP	Coef.	SE	p-value
Constant	0.00	0.00	0.01	0.00	0.01	0.76
Standard error	-1.07	0.34	0.95	-1.09	0.49	0.03
<i>Country scope</i>						
Advanced currencies	0.28	0.24	0.92	0.02	0.34	0.96
Emerging currencies	1.44	0.25	1.00	1.26	0.34	0.00
German mark	0.06	0.23	0.34	-0.20	0.37	0.58
French franc	0.78	0.26	0.94	0.42	0.38	0.26
British pound	-0.08	0.22	0.14	-0.48	0.36	0.18
Italian lira	0.85	0.30	0.94	0.52	0.39	0.19
Japanese yen	-0.66	0.24	1.00	-0.97	0.37	0.01
Swiss franc	-0.09	0.26	0.16	-0.62	0.38	0.10
Euro	-1.62	0.29	1.00	-2.01	0.40	0.00
European currencies	0.00	0.04	0.02	0.06	0.19	0.73
Other country currencies	0.00	0.02	0.02	-0.13	0.17	0.44
British pound base	0.00	0.03	0.01	0.31	0.22	0.16
Euro base	0.71	0.71	0.57	1.55	0.48	0.00
German mark base	0.50	0.18	0.95	0.52	0.15	0.00
<i>Data characteristics</i>						
Shorter horizon	0.01	0.10	0.03	0.82	0.50	0.10
1month horizon	-0.26	0.13	0.88	0.12	0.34	0.73
1month-1year horizon	0.01	0.06	0.06	0.29	0.29	0.32
1year horizon	0.00	0.02	0.01	0.13	0.27	0.62
Daily frequency	-0.07	0.17	0.18	-0.61	0.38	0.11
Weekly frequency	0.80	0.13	1.00	0.61	0.28	0.03
Monthly frequency	0.02	0.07	0.08	-0.03	0.17	0.86
Time difference	0.12	0.06	0.93	0.28	0.07	0.00
Number of currencies	-0.02	0.08	0.09	-0.18	0.11	0.11
Overlapping problem	0.00	0.02	0.02	-0.15	0.13	0.26
<i>Estimation</i>						
OLS method	-0.57	0.14	1.00	-0.76	0.11	0.00
Fixed-effects method	-0.78	0.37	0.88	-0.86	0.30	0.00
SUR method	-0.37	0.25	0.75	-0.60	0.17	0.00
Regime-switching model	-0.94	0.25	0.99	-1.02	0.24	0.00
Controls	0.00	0.02	0.01	-0.02	0.12	0.86
Spot rate percentage	0.00	0.03	0.02	0.36	0.15	0.02
<i>Regimes</i>						
Large differential	2.79	0.34	1.00	2.57	0.37	0.00
Small differential	-0.01	0.07	0.02	-0.36	0.37	0.33
Large positive premium	-0.36	0.36	0.55	-0.58	0.23	0.01
Low negative premium	0.12	0.25	0.21	0.45	0.24	0.06
Overvalued currency	0.00	0.07	0.01	0.41	0.61	0.50
Undervalued currency	0.02	0.16	0.02	0.75	0.63	0.24
<i>Data sources</i>						
Datastream source	-0.38	0.10	0.99	-0.51	0.10	0.00
Bank data source	0.00	0.02	0.02	0.00	0.09	0.99
DataResources source	0.00	0.03	0.02	-0.07	0.11	0.52
<i>Publication characteristics</i>						
Impact factor	0.00	0.01	0.01	0.08	0.06	0.19
Citations	0.00	0.00	0.03	0.04	0.02	0.03
First-draft year	0.03	0.00	1.00	0.03	0.01	0.00
Studies	74			74		
Observations	2,989			2,989		

*Notes:* Response variable = estimate of  $\beta$  from Equation 3. SE = standard error, P. mean = posterior mean, P. SD = posterior standard deviation, PIP = posterior inclusion probability. The posterior mean in Bayesian model averaging (or alternatively the estimated coefficient in frequentist model averaging) denotes the marginal effect of a study characteristic on the estimate of beta reported in the literature. For detailed description of all the variables see Table 3.

Figure 7: *Sensitivity of BMA to different priors*



Notes: UIP (unit information prior) and Dilution = priors recommended by Eicher *et al.* (2011) and George (2010), respectively. BRIC and Random = a g-prior by Fernandez *et al.* (2001) for parameters with the beta-binomial model prior (Ley & Steel, 2009) for model space; this ensures that each model size has equal prior probability. HQ and random = Hyper-g prior as suggested by Feldkircher & Zeugner (2012), HQ corresponds to Hannah-Quinn criterion. PIP stands for posterior inclusion probability.

**Data sources & Publication characteristics** The source of the data that is used to test the null hypothesis matters. We find decisive evidence that data extracted from Datastream are associated with 0.4 smaller estimates of  $\beta$ , which however contrasts with our hypothesis that public databases are less prone to mistakes and measurement error and therefore less attenuation bias. Furthermore, estimates of  $\beta$  have an increasing time trend, which corroborates the simple graph presented in the Introduction. We find decisive evidence that the reported  $\beta$  coefficients increase by about 0.03 per year. On the other hand, publication quality, as captured by the number of citations and impact factor of the outlet, do not affect the size of reported estimates of  $\beta$ . Consequently, there is not much evidence that estimates extracted from studies of higher quality are different from those originating from studies of lower quality after accounting for other aspects of the context in which the estimates were obtained.

## 5.4 Implied Estimates

In this subsection we calculate the mean  $\beta$  implied by the literature after we take into account the effect that currency, method, and data decisions have on the estimates of  $\beta$ . For this purpose, we create a hypothetical study in which we include all the 2,989 estimates of  $\beta$  extracted from the differences specification. We use three strategies how to define the values of individual variables (for example, the frequency, data horizon, or estimation technique). First, based on our reading of the literature we select the aspects of data and methodology that we believe

are preferred by the most prominent current studies. But of course such a choice is subjective, and other researchers would select other values. So, as a robustness check, we use the data and method choices of two prominent recent studies: Frankel & Poonawala (2010) and Breedon *et al.* (2016). Thus the implied beta in the second and third strategy can be viewed as results of a hypothetical study that would use all data available to all researchers in the literature but used the approach of Frankel & Poonawala (2010) and Breedon *et al.* (2016), respectively, all the while correcting for potential publication bias.

Table 5: Betas for different currencies implied by different estimation context

	Our preferred methodology	Frankel & Poonawala 2010	Breedon <i>et al.</i> 2016
Advanced currencies	0.309 (-0.298, 0.916)	0.448 (0.131, 0.764)	0.231 (-0.077, 0.539)
Emerging currencies	0.945 (0.326, 1.563)	1.164 (0.889, 1.438)	0.947 (0.683, 1.210)
Euro	-0.697 (-1.365, -0.030)	-0.570 (-0.866, -0.274)	-0.787 (-1.060, -0.514)
German mark	0.144 (-0.500, 0.789)	0.318 (0.024, 0.613)	0.101 (-0.173, 0.376)
French franc	0.461 (-0.205, 1.126)	0.626 (0.368, 0.884)	0.409 (0.189, 0.628)
British pound	0.069 (-0.570, 0.708)	0.157 (-0.127, 0.441)	-0.060 (-0.322, 0.202)
Italian lira	0.459 (-0.215, 1.132)	0.649 (0.395, 0.903)	0.432 (0.217, 0.647)
Japanese yen	-0.352 (-0.887, 0.183)	-0.250 (-0.548, 0.049)	-0.467 (-0.736, -0.198)
Swiss franc	0.058 (-0.579, 0.696)	0.078 (-0.218, 0.375)	-0.139 (-0.418, 0.141)
European currencies	0.115 (-0.624, 0.854)	0.474 (0.197, 0.750)	0.256 (0.039, 0.474)
Other country currencies	0.111 (-0.626, 0.849)	0.344 (0.059, 0.629)	0.127 (-0.108, 0.361)

*Notes:* The table presents mean estimates of  $\beta$  conditional on selected data, estimation, and publication characteristics. The exercise is akin to constructing a hypothetical study that uses all estimates in the literature but puts more weight on selected aspects of study design. The first column shows the betas conditional on the methodology that we prefer after reading the literature (see the text for details). In the next two columns we also show betas conditional on the methodology of two prominent recent papers: Frankel & Poonawala (2010) and Breedon *et al.* (2016).

Effectively, we calculate an implied estimate of  $\beta$  by using the results of BMA in Subsection 5.3 and calculating a linear combination of BMA coefficients and the chosen values for each variable. Regarding the first strategy, our subjective definition of preferred methodology, in the linear combination of the coefficients we plug in sample maxima for variables that are in line with the preferred methodology in the literature, sample minima for variables that depart from preferred methodology, and sample means for variables where preferred methodology is not obvious. To be specific, we prefer precise estimates, estimates using foreign exchange spot and forward rates over longer horizons, controlling for overlapping samples problem in the data,



and using more complex estimation techniques than OLS; so we plug sample minima for the standard error (here we actually plug zero because we want to control for the entirety of publication bias), less-than-one-month dummy, overlapping-problem dummy, and OLS dummy. We also prefer the use of foreign exchange rates of longer maturities, estimation methods allowing for different regimes or cross-sectional dependence across currencies, estimates extracted from more recent studies and studies of higher publication quality. In line with these best practice preferences in the literature we plug sample maxima one-year horizon dummy, time difference, regime-switching and SUR dummies, impact factor, citation count, and publication year. For all the remaining variables included in the BMA we cannot reliably discern preferred methodology and thus plug in their sample means.


The results of the exercise for different currencies are shown in Table 5. Aggregating across developed countries we observe an implied  $\beta$  estimate between 0.23 and 0.45, which is in line with our baseline result obtained after correcting for publication bias. This estimate is below the theoretically predicted value of 1; however, in contrast to the conclusion in numerous prior studies, it is positive rather than negative. Furthermore, we obtain an implied  $\beta$  estimate of 0.95–1.16 for emerging economy currencies, which is consistent with no forward premium puzzle. While the wide confidence intervals of the implied estimates do not allow us to conclude for some of the estimates that they are statistically different from zero at the 5% level, all these estimates are clearly larger than the simple mean estimates of  $\beta$ , which are reported in Table 1. Regarding the individual currencies, our bottom-line estimates of the beta  $\beta$  parameter is between  $-0.79$  and  $-0.57$  for the euro, 0.10 and 0.32 for German mark, 0.41 and 0.63 for French franc,  $-0.1$  and 0.16 for British pound, 0.43 and 0.65 for Italian lira,  $-0.47$  and  $-0.25$  for Japanese yen, and  $-0.14$  and 0.08 for Swiss franc.

## 6 Extensions and Robustness Checks

This section presents a battery of new specifications and robustness checks that strengthen and provide nuances for our baseline results, and we thank four anonymous referees of the European Economic Review for pointing us at these important issues. We start by focusing on estimates that stem from a regression of spot rates on forward rates in levels, not differences. As we have noted earlier, we excluded the level estimates from the baseline analysis because they show a very different pattern from the difference estimates and are likely to represent a distinct population of studies, one that is best analyzed separately. These estimates are closely concentrated around 1, as shown in Figure 8. There are two possible interpretations of this fact. First, the level estimates can be better suited than difference estimates to identify the underlying beta coefficients, and their results are consistent with the forward rate unbiasedness hypothesis. Second, the concentration of level estimates in the vicinity of 1 can be given by serious problems in measurement. In the presence of nonstationarity and absence of cointegration, a regression of spot rates on forward rates in levels can lead to a spurious slope coefficient. In our opinion the second explanation is more likely, which is the principal reason why we excluded the level

estimates from the baseline analysis reported earlier. With few exceptions, the level estimates are reported in older studies and are not pursued by the modern literature.

Figure 8: *A funnel plot and histogram for level estimates*



*Notes:* The figure represents funnel plot (on the right) and histogram (on the left) of the estimates extracted from the level equation (Equation 2). In the absence of publication bias the scatter plot should resemble an inverted funnel that is symmetrical around the most precise estimates. Logarithmic scale is used to depict precision (the vertical axis of the funnel plot). Outlying observations are cut from the figure for ease of exposition but included in all statistical tests. The solid vertical line in the histogram shows the mean from the winsorized estimates of beta from the level specification.

The funnel plot in Figure 8 suggests some asymmetry, which can be explained by small sample bias, publication bias, or heterogeneity. It is important to note that the level estimates not concentrated around 1 come from only a couple of studies, most of which test for cointegration. These studies show a pattern similar to the one displayed by difference estimates. The results of publication bias tests for the level estimates are reported in Table 6. They show remarkable consistency with little evidence for publication bias and the estimated underlying beta around 1, which is also close to the simple mean of the reported level estimates. The only outlier is the instrumental variable estimation, where the inverse of the square root of the number of observations is used as the instrument for the standard error (Stanley, 2005). Nevertheless, in this case the instrument is weak, and the instrumental variable estimate is thus uninformative. We conclude that the level estimates are on average close to 1 according to all major techniques that we employ, there is little evidence for publication bias, and we find some evidence for heterogeneity with a few studies that potentially do not suffer from spurious regression (and hence show results relatively similar to the difference estimates). The interpretation of our baseline results is not affected by these new findings.

Another issue that warrants attention is our treatment of outliers. Some of the reported estimates are clearly implausible, larger than 50 in the absolute value. We have made sure these estimates are coded properly from the primary studies, and in several cases we have contacted the authors of these studies and asked them for details on these estimates. Most of these estimates come from data on the currencies of developing countries, where much more noise is

Table 6: *Tests of publication bias (level estimates)*

<i>Panel A: FAT-PET</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.934 <sup>***</sup> (0.0573)	0.988 <sup>***</sup> (0.0109)	1.027 <sup>***</sup> (0.159)	
Publication bias (Constant)		0.875 (1.791)	-1.627 <sup>***</sup> (0.405)	-3.438 (7.343)	
Observations		654	654	654	
<i>Panel B: PEESE</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.933 <sup>***</sup> (0.0313)	0.967 <sup>***</sup> (0.009)	0.779 (0.995)	
Publication bias (SE)		0.298 (0.291)	-0.0116 <sup>***</sup> (0.00421)	5.838 (38.61)	
Observations		654	654	654	
<i>Panel C: Advanced</i>	WAAP	Kinked model	Selection model	p-uniform*	Stem method
Mean beyond bias	0.979 <sup>***</sup> (0.005)	0.979 <sup>***</sup> (0.004)	0.993 <sup>***</sup> (0.001)	0.966 <sup>***</sup> (0.010)	0.989 <sup>***</sup> (0.008)
Observations	654	654	654	654	654

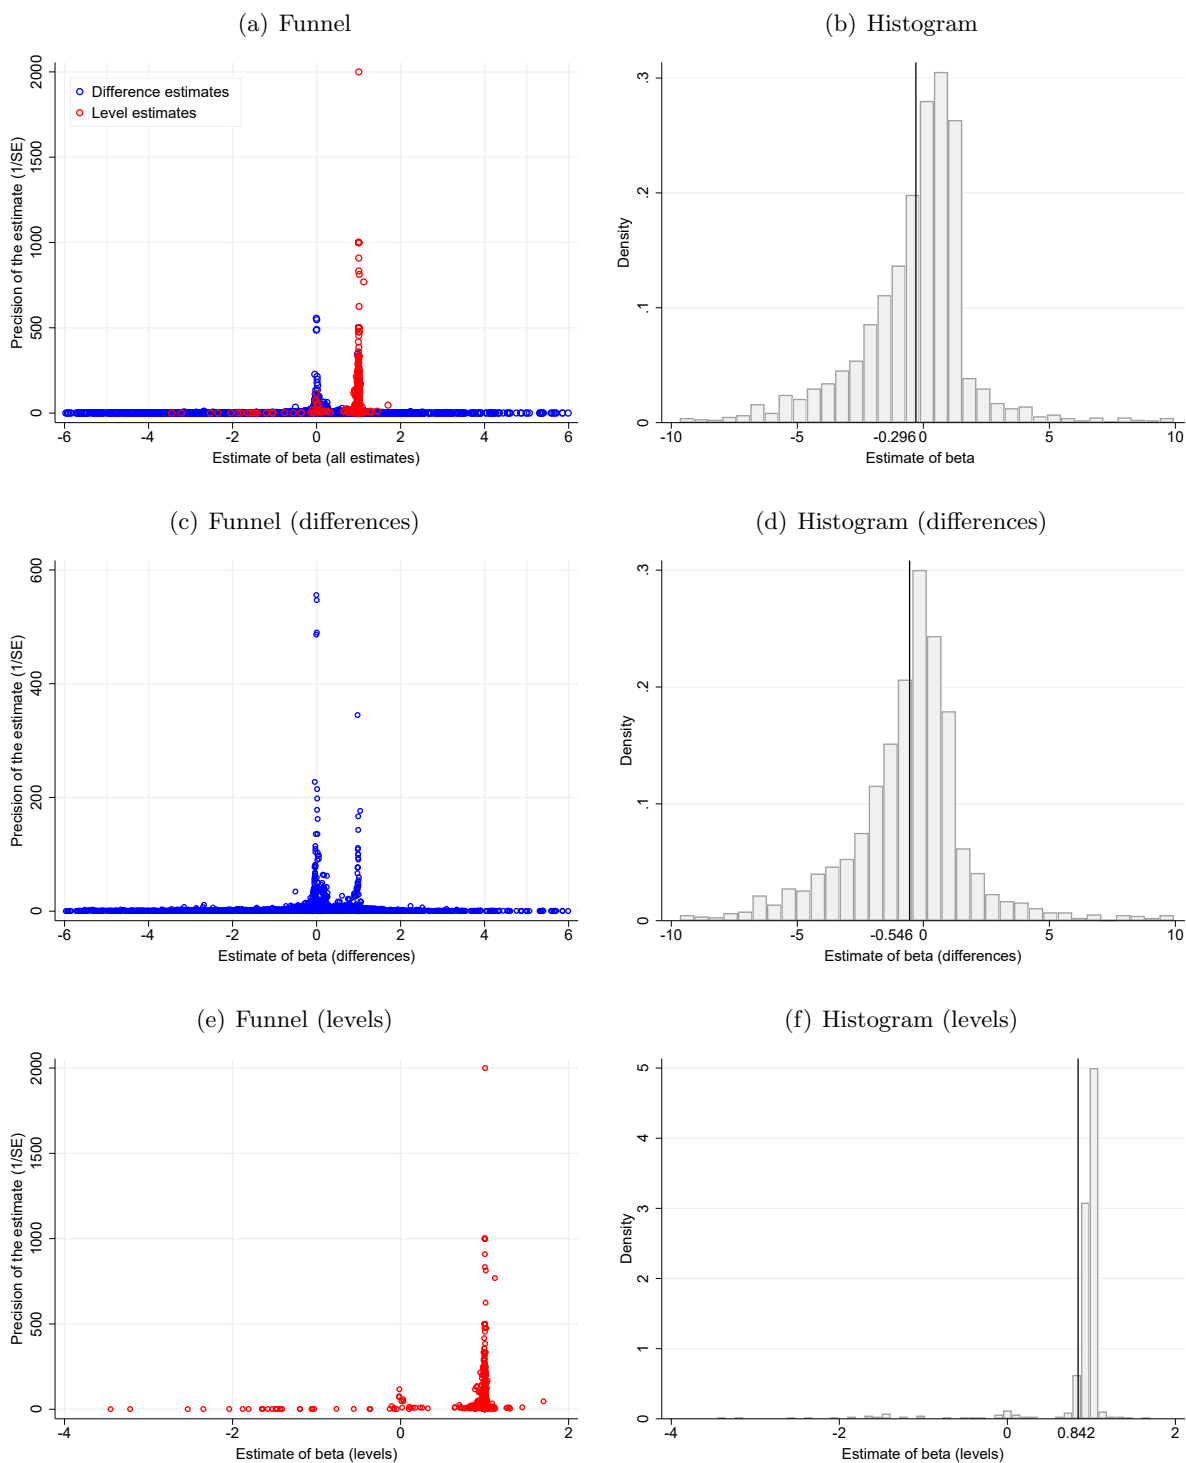
*Notes:* FE = study-level fixed effects, WLS = weighted least squares, IV = instrumental variables regression with the inverse of the square root of the number of observations used as an instrument of the standard error (the instrument is weak in this case). Only level estimates (Equation 2) are included. Panel A & B: weighted by both inverse variance and the inverse of the number of estimates reported per study; bootstrapped standard errors in parentheses. Panel C: WAAP (weighted average of adequately powered, Ioannidis *et al.*, 2017), kinked model (Bom & Rächinger, 2019), selection model (Andrews & Kasy, 2019), p-uniform\* (van Aert & van Assen, 2021), stem method (Furukawa, 2020).

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

present than in the data on developed countries. Because we attempt to include all estimates, we collect the extreme ones as well. In our baseline estimation we prefer to winsorize the extreme values of the reported standard errors and the magnitude of the beta. We believe that winsorization has two main benefits relative to data trimming. First, because it is symmetrical, winsorization preserves the median and in many cases also the mean of the underlying dataset. Data trimming at specific values often leads to significant changes in both statistics. Second, winsorization uses all the observations. For example, an estimate of beta that equals 50 cannot plausibly be used in a meta-analysis at face value, because it would hugely distort basic statistics such as the mean and standard deviation and also any regression result. But we still find it useful to preserve the information that the estimate exists and is large. Such a preservation is achieved by winsorizing, though of course the winsorization level is subjective. We choose our winsorization level (5%) as the smallest adjustment of outliers after which the results stabilize (that is, a 6% or 10% level would bring very similar outcomes). In other words, we prefer not to put much weight on results that would crucially rely on a few influential estimates.

But of course such influential estimates can be important and useful. Moreover, winsorizing stacks data at pre-defined percentiles, which might create other problems for the statistical tests we rely on, including Bayesian model averaging. So it is important to show that our results do not depend on winsorizing and are robust to a plausible and conservative trimming of outliers. First, in Figure 9 we show histograms and funnel plots that include all outliers in the standard

Figure 9: *Funnel plots and histograms including more extreme observations*



*Notes:* The figure depicts funnels and histograms with outlying values liberally included. To be specific, all reported values of precision are shown. We still cut the most extreme estimates of beta (which sometimes surpass 50 in absolute value), because their inclusion would make the figures illegible. The horizontal lines in histograms depict mean reported betas (without winsorization).

errors and almost all outliers in the magnitude of beta (with the exception of those that are larger than 10 in absolute value). The first row of the figure shows the plots for all estimates, the second row shows the plots for difference estimates, and the third row shows the plots for level estimates. The distribution of level estimates, as we have noted, is concentrated around 1 with a couple of (typically negative) outliers. The distribution of difference estimates is well-behaved, with two main features distinguishing it from a normal distribution: asymmetry and fat tails. Regarding funnel plots (where we now use a linear instead of a logarithmic scale for precision), the level estimates display a single clear peak at one.

In contrast, the difference estimates display two peaks, one at zero and the other at one. While the first peak is formed by estimates from many studies, the second peak only includes estimates from 4 studies (out of a total of 74 studies for difference estimates). These 4 studies are not connected by a single data or method aspect that could be controlled for in a meta-regression analysis. Instead, they all offer idiosyncratic solutions to the forward puzzle. For example, Hall *et al.* (2013) offer an explanation based on a time-varying estimator robust to measurement error and misspecifications, while Pippenger (2011) puts forward an explanation based on omitted variables. The bottom line of these studies is that when estimated correctly, the beta coefficient is precisely estimated at being close to 1. While we fail to code a variable that would capture all the idiosyncratic explanations of these 4 studies, we believe it would be inappropriate to discount them and we therefore include them in the analysis with the rest of the sample.

In Table 7 we examine how our results regarding publication bias change when we abandon winsorizing and rely instead on conservative trimming of the most extreme outliers. We still observe the principle that guided us regarding the choice of the winsorization threshold: we want to trim as few observations as possible in order for our results to stabilize. That is, for our final threshold it holds that if we trim more outliers, our results will not change qualitatively (or will be stronger). We proceed in two steps. First, we focus on the threshold for the magnitude of the reported beta. The threshold at which our results stabilize is 6, which means that we omit all estimates larger than 6 in the absolute value. Such estimates are considered implausible by the bulk of the literature, so we see little harm in doing so. Second, we focus on the reported standard errors. Here the problem is much more delicate, because outliers in precision (observations with extremely small standard errors) are very influential in meta-analysis: in all estimations, we use inverse variance as a weight. Nominally, these estimates are the most informative. But in rare cases, extremely small standard errors can also arise due to undetectable mistakes in reporting or inappropriate estimation of standard errors themselves, which could cause a bias in the meta-analysis. We thus use  $DF_{beta}$  to identify the values of the standard error that differ the most from the estimated regression line in the funnel asymmetry test. We find that our results stabilize after trimming merely 3 of the 2,989 observations for difference estimates. Moreover, the results in Table 7 are similar to those of our baseline model, with the exception of the stem-based estimator, which has troubles converging and yields an insignificant estimate of the underlying beta.

Table 7: *Tests of publication bias (trimmed dataset, no winsorization)*


<i>Panel A: FAT-PET</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.698 <sup>***</sup> (0.183)	0.677 <sup>***</sup> (0.158)	0.449 <sup>**</sup> (0.201)	
Publication bias (Constant)		-2.459 <sup>***</sup> (0.643)	-2.367 <sup>***</sup> (0.488)	-1.385 <sup>**</sup> (0.579)	
Observations		2,756	2,756	2,756	
<i>Panel B: PEESE</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.699 <sup>***</sup> (0.191)	0.645 <sup>***</sup> (0.140)	0.469 <sup>**</sup> (0.224)	
Publication bias (SE)		-0.342 <sup>***</sup> (0.097)	-0.00270 (0.00337)	-0.582 <sup>*</sup> (0.329)	
Observations		2,756	2,756	2,756	
<i>Panel C: Advanced</i>	WAAP	Kinked model	Selection model	p-uniform*	Stem method
Mean beyond bias	0.295 <sup>***</sup> (0.028)	0.305 <sup>***</sup> (0.009)	0.205 <sup>***</sup> (0.025)	0.906 <sup>***</sup> (0.188)	0.312 (0.277)
Observations	2,756	2,756	2,756	2,756	2,756

*Notes:* Instead of winsorization here we trim estimates of beta larger than 6 in absolute value. We also drop 3 most influential values of the standard error based on DFbeta. FE = study-level fixed effects, WLS = weighted least squares, IV = instrumental variables regression with the inverse of the square root of the number of observations used as an instrument of the standard error. Only difference estimates (Equation 3) are included. Panel A & B: weighted by both inverse variance and the inverse of the number of estimates reported per study; bootstrapped standard errors in parentheses. Panel C: WAAP (weighted average of adequately powered, Ioannidis *et al.*, 2017), kinked model (Bom & Rachinger, 2019), selection model (Andrews & Kasy, 2019), p-uniform\* (van Aert & van Assen, 2021), stem method (Furukawa, 2020). The stem method has troubles converging in this case. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Figure 10 shows the results of Bayesian model averaging when we use the aforementioned trimming strategy (omitting estimates larger than 6 in absolute value and also the 3 most influential values of the standard error) instead of winsorizing. The remaining aspects and setup of the Bayesian model averaging exercise are the same as in the baseline case. We observe that our main results hold. In particular, the standard error has a posterior inclusion probability of almost one, which suggests publication bias even after we control for the context in which the estimates were obtained in primary studies (that is, the apparent asymmetry of the funnel plot is most likely not caused by heterogeneity). The reported estimates of beta tend to be larger for the currencies of emerging countries. Among advanced currencies, small estimates of beta are typically reported for the euro and yen, while estimates for the French franc and Italian lira tend to be larger. Large differentials are associated with less forward premium puzzle. To conclude, this extension suggests that our results are not driven by winsorizing and are robust to a plausible and conservative omitting of extreme outliers.

Another important extension is a separate analysis of publication bias and the underlying corrected beta for a sub-sample of the currencies of developed countries. The corresponding results are shown in Table 8. We do not conduct a separate analysis for emerging countries, because many of the advanced techniques do not converge given the relatively small number of estimates available for such an exercise. In general, we prefer to examine heterogeneity (in-

Figure 10: *Model inclusion in BMA (trimmed dataset, no winsorization)*



*Notes:* Instead of winsorization here we trim estimates of beta larger than 6 in absolute value. We also drop 3 most influential values of the standard error based on  $DF_{beta}$ . The response variable in BMA is an estimate of slope coefficient  $\beta$  from Equation 3. We employ the prior suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data), and the dilution prior suggested by George (2010), which accounts for collinearity. Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities for the 5,000 best models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model. All variables are described in Table 3.

cluding the differences between developed and emerging countries) within the Bayesian model averaging framework, which is tractable and allows us to control for many other aspects of the context in which the estimate was obtained. Nevertheless, the new results for developed countries presented in Table 8 also suggest the presence of publication bias that prefers negative estimates (that is, a confirmation bias in the direction of the earliest influential published estimates). The corrected mean effects beyond publication bias are between 0.3 and 0.9 with the exception of the instrumental variable estimator (where the instrument is weak for this subsample) and the selection model (which fails to converge). Thus, as in the baseline results

Table 8: *Tests of publication bias (currencies of developed countries)*

<i>Panel A: FAT-PET</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.657 <sup>***</sup> (0.117)	0.639 <sup>***</sup> (0.108)	0.359 <sup>*</sup> (0.196)	
Publication bias (Constant)		-2.331 <sup>***</sup> (0.340)	-2.264 <sup>***</sup> (0.262)	-1.258 <sup>**</sup> (0.523)	
Observations		2,582	2,582	2,582	
<i>Panel B: PEESE</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.666 <sup>***</sup> (0.109)	0.582 <sup>***</sup> (0.0883)	0.275 (0.182)	
Publication bias (SE)		-0.390 <sup>***</sup> (0.0537)	-0.219 <sup>***</sup> (0.0226)	-0.597 <sup>**</sup> (0.301)	
Observations		2,582	2,582	2,582	
<i>Panel C: Advanced</i>	WAAP	Kinked model	Selection model	p-uniform*	Stem method
Mean beyond bias	0.401 <sup>***</sup> (0.0326)	0.424 <sup>***</sup> (0.0105)	0.084 (0.101)	0.851 <sup>***</sup> (0.127)	0.791 <sup>***</sup> (0.295)
Observations	2,582	2,582	2,582	2,582	2,582

*Notes:* FE = study-level fixed effects, WLS = weighted least squares, IV = instrumental variables regression with the inverse of the square root of the number of observations used as an instrument of the standard error. Only difference estimates (Equation 3) for the currencies of advanced countries are included. Panel A & B: weighted by both inverse variance and the inverse of the number of estimates reported per study; bootstrapped standard errors in parentheses. Panel C: WAAP (weighted average of adequately powered, Ioannidis *et al.*, 2017), kinked model (Bom & Rachinger, 2019), selection model (Andrews & Kasy, 2019), p-uniform\* (van Aert & van Assen, 2021), stem method (Furukawa, 2020). The selection model does not converge in this sub-sample. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

presented in the previous section, we find that for developed countries the forward rate unbiasedness hypothesis does not hold but that the underlying beta coefficient is on average positive in contrast to the findings reported previously in much of the literature.

As we have noted earlier, the forward rate unbiasedness hypothesis is equivalent to the uncovered interest rate parity hypothesis if the covered interest rate parity also holds. Therefore, one could, in principle, additionally include estimates from studies that replace forward rate discounts with interest rate differentials. In the baseline analysis we choose to focus on studies that use forwards, because such studies are plentiful, consistent with each other, and we do not need to invoke the covered interest parity (which might not hold universally after the wave of regulations enacted as a response to the financial crisis of 2008/2009). For some countries, however, data on forwards rates are not readily available, and interest rate differentials are thus used as substitutes, which means that studies based on the latter may be used as a robustness check to our results.

Because data collection in meta-analysis is extremely laborious, we have decided to complement our main dataset of 3,643 estimates from studies using forwards with a randomly selected sample of 200 estimates using interest rate differentials. We exclude one outlier and conduct a battery of tests regarding publication bias and the underlying corrected mean beta. The results, reported in Table 9, show that some methods have troubles converging (p-uniform\*, selection model, and the stem method). The instrument for the standard error is weak in this case. On



Table 9: *Tests of publication bias (interest rate differentials)*


<i>Panel A: FAT-PET</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.00387 (0.0167)	0.0345** (0.0138)	-0.329 (2.439)	
Publication bias (Constant)		0.780*** (0.176)	0.629*** (0.194)	2.423 (10.68)	
Observations		199	199	199	
<i>Panel B: PEESE</i>		FE	WLS	IV	
Mean beyond bias (1/SE)		0.00322 (0.0179)	0.0493* (0.0273)	-6.092 (12.32)	
Publication bias (SE)		-0.101 (0.0627)	-0.0590* (0.0341)	33.30 (44.05)	
Observations		199	199	199	
<i>Panel C: Advanced</i>	WAAP	Kinked model	Selection model	p-uniform*	Stem method
Mean beyond bias	0.0498*** (0.0148)	0.0323** (0.0129)	0.721 (0.987)	0.603 (0.721)	NA (NA)
Observations	199	199	199	199	199

*Notes:* Here we include studies that use data on interest rate differentials instead of forward rates. FE = study-level fixed effects, WLS = weighted least squares, IV = instrumental variables regression with the inverse of the square root of the number of observations used as an instrument of the standard error (the instrument is weak in this case). Only difference estimates (Equation 3) are included. Panel A & B: weighted by both inverse variance and the inverse of the number of estimates reported per study; bootstrapped standard errors in parentheses. Panel C: WAAP (weighted average of adequately powered, Ioannidis *et al.*, 2017), kinked model (Bom & Rachinger, 2019), selection model (Andrews & Kasy, 2019), p-uniform\* (van Aert & van Assen, 2021), stem method (Furukawa, 2020). In this case the selection model and p-uniform\* have troubles converging. The stem method does not converge at all. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

balance, we find little evidence for publication bias. The corrected mean beta is in most models positive but small. Thus the results of this extension are not entirely in line with the main results. The beta coefficient is not negative, which implies (in line with our baseline results and in contrast to much of the literature) that carry trades are not profitable. Nevertheless, the coefficient is far from 1, indicating a violation of the forward rate unbiasedness hypothesis. A potential explanation of the difference between this extension and our baseline results is a violation of the covered interest rate parity. Another explanation is that for a proper analysis of interest rate differentials we need more data, and thus that a separate meta-analysis on this topic is warranted.

In the baseline Bayesian model averaging analysis we include dummy variables for the most important individual currencies and groups of currencies. Nevertheless, a plausible alternative specification involves replacing the dummy variables with country-level variables that allow for a more parsimonious and structured representation of the underlying relationships. In the next extension we thus replace all the currency-specific dummy variables (with the exception of variables specifying the base currency used in computations) with two variables: GDP per capita (based on data from the World Bank) and exchange rate regime (a dummy based on data from the IMF; the variable equals 0 for a fixed regime, 1 for an intermediate regime, and 2 for a floating regime). We are thus able to explicitly test the exchange rate regime hypothesis and


Figure 11: *Model inclusion in BMA (country variables)*



*Notes:* Here we replace the various dummies for individual currencies by two variables: GDP per capita (from the World Bank and used in logs) and Exchange rate regime (from the IMF; the variable equals 0 for a fixed regime, 1 for an intermediate regime, and 2 for a floating regime). The response variable in BMA is an estimate of slope coefficient  $\beta$  from Equation 3. We employ the prior suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data), and the dilution prior suggested by George (2010), which accounts for collinearity. Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities for the 5,000 best models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model.

also the hypothesis that more developed countries display more forward premium puzzle. The results are shown in Figure 11. Although the structure of the model differs a lot (the individual currency variables are omitted), we still find that the main results are consistent with the baseline BMA exercise. Namely, even in this extension we find evidence for publication bias, the importance of using OLS, and the role of a large differential. Regarding the two new country-level variables, we confirm that the reported beta coefficients tend to be smaller for countries with higher GDP per capita. In contrast, we find no evidence that the reported beta differs across countries with different exchange rate regimes.

Figure 12: *Model inclusion in BMA (no survey data)*



Notes: Here we omit all estimates that use survey data. The response variable in BMA is an estimate of slope coefficient  $\beta$  from Equation 3. We employ the prior suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data), and the dilution prior suggested by George (2010), which accounts for collinearity. Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities for the 5,000 best models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model.

As a final robustness check, we omit from our dataset estimates constructed using survey data. Those studies are in fact aimed at investigating whether a time-varying risk premium rather than a violation of the rational expectations hypothesis explains the forward bias. In consequence, these estimates might be inconsistent with the rest of our dataset. Nevertheless, there are very few estimates based on survey data in our sample, and the results in Figure 12 are similar to those presented in the baseline BMA exercise in the previous section.

## 7 Conclusion

We present the first quantitative synthesis of the extensive empirical literature on the forward premium puzzle. We collect 3,643 estimates of  $\beta$  from 91 studies, which makes this synthesis

one of the largest meta-analyses ever conducted in economics and finance. Our results suggest that, after correction for various biases, the average slope coefficient  $\beta$  in the literature is positive but smaller than 1. Furthermore, we exploit the heterogeneity in data samples and estimation methodologies and examine the impact of 42 study and estimation characteristics on the reported  $\beta$  estimates. To address the problem of model uncertainty arising from the presence of many potential explanatory variables we use the Bayesian and frequentist model averaging techniques. We observe systematic differences between currencies of more and less advanced economies with higher  $\beta$  estimates for emerging country currencies, and to a lesser extent the former French franc and Italian lira. We also find that the  $\beta$  estimates tend to be larger when focusing on weekly observations and sophisticated estimation methods that account for potential endogeneity. Furthermore, we find evidence that  $\beta$  estimates are regime-dependent as they differ across different time periods and they tend to increase over time. On the contrary, we document no systematic effect of publication quality proxies on the results.

As the bottom line of our analysis we use 2,989 estimates of  $\beta$  from the differences specification to construct a hypothetical study based on weighted study characteristics. We obtain an implied  $\beta$  estimate for developed countries between 0.23 and 0.45, which is close to our estimates that correct for the publication bias without any judgment on the relative desirability of data samples and methodology. Nevertheless, our  $\beta$  estimates based on the hypothetical study exercise exhibit wide confidence intervals, which imply that some of the estimates are not statistically different from 0 at the conventional 5% significance level.

The meta-analysis has implications for the various explanations of the forward puzzle that have been raised in the literature. Our basic result is that, when the literature is evaluated as a whole and corrected for various biases (especially publication bias), we find little evidence of the forward premium puzzle for the currencies of emerging countries and only modest evidence of the puzzle in developing countries. These results are in line with Frankel & Poonawala (2010) and cast doubt on the most common explanation of the puzzle: time-varying risk premia. Because larger risk premia are typically more volatile, the currencies of emerging countries should show more evidence for the puzzle if time-varying risk premia was an underlying explanation. We also find that there is less forward bias if there are large interest rate differentials (that is, large opportunities to make money), which is consistent with the view that transaction costs (or, potentially, attention) play a role in the puzzle. Moreover, we corroborate the results of several studies, such as Pippenger (2011) and Hall *et al.* (2013), that the apparent forward puzzle can be due to misspecifications in estimation. Finally, we propose an explanation that has not appeared in the literature: the puzzle can be driven by publication bias, specifically confirmation bias in the direction of the early influential estimates by Fama (1984).

Three qualifications of our results are in order. First, we focus on studies published in peer-reviewed journals and (mainly) on studies using forward rates. One could also include unpublished papers and collect more data from studies using the interest rate differential. We prefer published studies from unpublished ones because the former are likely to be of a higher quality due to the peer-review process. Studies using the interest rate differential produce

estimates comparable to our  $\beta$  only if the covered interest parity holds, which does not have to be the case for all markets, especially after the financial crisis. Second, estimates reported within one study are unlikely to be independent. We attempt to tackle this issue by bootstrapping the standard errors in our regressions at the study level, but we admit that bootstrapping is not an ultimate solution to the problem of sample overlap. Third, our analysis of implied beta is inevitably subjective, because judgment is required on various aspects of study design. Different authors may choose a somewhat different set of characteristics. Nevertheless, plausible changes to the definition of preferred methodology in line with two recent prominent papers (Frankel & Poonawala, 2010; Breedon *et al.*, 2016) keep our results qualitatively unchanged.

## References

- VAN AERT, R. C. & M. VAN ASSEN (2021): “Correcting for publication bias in a meta-analysis with the p-uniform\* method.” *Working paper*, Tilburg University & Utrecht University.
- AMINI, S. M. & C. PARMETER (2012): “Comparison of model averaging techniques: Assessing growth determinants.” *Journal of Applied Econometrics* **27**(5): pp. 870–876.
- ANDREWS, I. & M. KASY (2019): “Identification of and correction for publication bias.” *American Economic Review* **109**(8): pp. 2766–94.
- ASTAKHOV, A., T. HAVRANEK, & J. NOVAK (2019): “Firm Size And Stock Returns: A Quantitative Survey.” *Journal of Economic Surveys* **33**(5): pp. 1463–1492.
- BACCHETTA, P. & E. VAN WINCOOP (2010): “Infrequent Portfolio Decisions: A Solution to the Forward Discount Puzzle.” *American Economic Review* **100**(3): pp. 870–904.
- BACKUS, D. K., F. GAVAZZONI, C. TELMER, & S. E. ZIN (2010): “Monetary Policy and the Uncovered Interest Parity Puzzle.” *NBER Working Papers 16218*, National Bureau of Economic Research, Inc.
- BACKUS, D. K., A. W. GREGORY, & C. I. TELMER (1993): “Accounting for Forward Rates in Markets for Foreign Currency.” *Journal of Finance* **48**(5): pp. 1887–1908.
- BAILLIE, R. T. & T. BOLLERSLEV (1989): “Common Stochastic Trends in a System of Exchange Rates.” *Journal of Finance* **44**(1): pp. 167–81.
- BAILLIE, R. T. & T. BOLLERSLEV (1994): “Cointegration, Fractional Cointegration, and Exchange Rate Dynamics.” *Journal of Finance* **49**(2): pp. 737–45.
- BAILLIE, R. T. & S. S. CHANG (2011): “Carry trades, momentum trading and the forward premium anomaly.” *Journal of Financial Markets* **14**(3): pp. 441 – 464.
- BAILLIE, R. T. & R. KILIÇ (2006): “Do asymmetric and nonlinear adjustments explain the forward premium anomaly?” *Journal of International Money and Finance* **25**(1): pp. 22 – 47.
- BAJZIK, J., T. HAVRANEK, Z. IRSOVA, & J. SCHWARZ (2020): “Estimating the Armington elasticity: The importance of study design and publication bias.” *Journal of International Economics* **127**(C): p. 103383.
- BANSAL, R. & M. DAHLQUIST (2000): “The forward premium puzzle: different tales from developed and emerging economies.” *Journal of International Economics* **51**(1): pp. 115–144.
- BEKAERT, G. (1995): “The Time Variation of Expected Returns and Volatility in Foreign-Exchange Markets.” *Journal of Business & Economic Statistics* **13**(4): pp. 397–408.
- BEKAERT, G. & R. J. HODRICK (1993): “On biases in the measurement of foreign exchange risk premiums.” *Journal of International Money and Finance* **12**(2): pp. 115–138.
- BLANCO-PEREZ, C. & A. BRODEUR (2020): “Publication Bias and Editorial Statement on Negative Findings.” *Economic Journal* **130**(629): pp. 1226–1247.
- BOM, P. R. D. & H. RACHINGER (2019): “A kinked meta-regression model for publication bias correction.” *Research Synthesis Methods* **10**(4): pp. 497–514.
- BREEDON, F., D. RIME, & P. VITALE (2016): “Carry Trades, Order Flow, and the Forward Bias Puzzle.” *Journal of Money, Credit and Banking* **48**(6): pp. 1113–1134.
- BRODEUR, A., N. COOK, & A. HEYES (2020): “Methods Matter: P-Hacking and Causal Inference in Economics.” *American Economic Review* **110**(11): pp. 3634–3660.
- BRODEUR, A., M. LE, M. SANGNIER, & Y. ZYLBERBERG (2016): “Star wars: The empirics strike back.” *American Economic Journal: Applied Economics* **8**(1): pp. 1–32.
- BRUNS, S. B. & J. P. A. IOANNIDIS (2016): “p-Curve and p-Hacking in Observational Research.” *PloS ONE* **11**(2): p. e0149144.
- BURNSIDE, C., B. HAN, D. HIRSHLEIFER, & T. Y. WANG (2011): “Investor Overconfidence and the Forward Premium Puzzle.” *Review of Economic Studies* **78**(2): pp. 523–558.

- BYERS, J. D. & D. A. PEEL (1991): “Some evidence on the efficiency of the sterling-dollar and sterling-franc forward exchange rates in the interwar period.” *Economics Letters* **35(3)**: pp. 317–322.
- CAZACHEVICI, A., T. HAVRANEK, & R. HORVATH (2020): “Remittances and Economic Growth: A Meta-Analysis.” *World Development* **134**: p. 105021.
- CHIANG, T. C. (1988): “The Forward Rate as a Predictor of the Future Spot Rate—A Stochastic Coefficient Approach.” *Journal of Money, Credit and Banking* **20(2)**: pp. 212–32.
- CHINN, M. & G. MEREDITH (2004): “Monetary policy and long-horizon uncovered interest parity.” *IMF Staff Papers* **51(3)**: pp. 409–430.
- CHO, D. & S. CHUN (2019): “Can structural changes in the persistence of the forward premium explain the forward premium anomaly?” *Journal of International Financial Markets, Institutions and Money* **58(C)**: pp. 225–235.
- CHRISTENSEN, G. & E. MIGUEL (2018): “Transparency, Reproducibility, and the Credibility of Economics Research.” *Journal of Economic Literature* **56(3)**: pp. 920–980.
- CORNELL, B. (1989): “The impact of data errors on measurement of the foreign exchange risk premium.” *Journal of International Money and Finance* **8(1)**: pp. 147–157.
- COUDERT, V. & V. MIGNON (2013): “The “forward premium puzzle” and the sovereign default risk.” *Journal of International Money and Finance* **32(C)**: pp. 491–511.
- COULIBALY, D. & H. KEMPF (2019): “Inflation targeting and the forward bias puzzle in emerging countries.” *Journal of International Money and Finance* **90(C)**: pp. 19–33.
- CROWDER, W. J. (1994): “Foreign exchange market efficiency and common stochastic trends.” *Journal of International Money and Finance* **13(5)**: pp. 551–564.
- DOMOWITZ, I. & C. S. HAKKIO (1985): “Conditional variance and the risk premium in the foreign exchange market.” *Journal of International Economics* **19(1-2)**: pp. 47–66.
- EICHER, T. S., C. PAPAGEORGIOU, & A. E. RAFTERY (2011): “Default priors and predictive performance in Bayesian model averaging, with application to growth determinants.” *Journal of Applied Econometrics* **26(1)**: pp. 30–55.
- FAMA, E. F. (1984): “Forward and spot exchange rates.” *Journal of Monetary Economics* **14(3)**: pp. 319–338.
- FARHI, E. & X. GABAIX (2016): “Rare Disasters and Exchange Rates.” *The Quarterly Journal of Economics* **131(1)**: pp. 1–52.
- FELDKIRCHER, M. & S. ZEUGNER (2012): “The impact of data revisions on the robustness of growth determinants—a note on determinants of economic growth: Will data tell?” *Journal of Applied Econometrics* **27(4)**: pp. 686–694.
- FERNANDEZ, C., E. LEY, & M. F. J. STEEL (2001): “Benchmark priors for Bayesian Model Averaging.” *Journal of Econometrics* **100(2)**: pp. 381–427.
- FLOOD, R. P. & A. K. ROSE (1996): “Fixes: Of the Forward Discount Puzzle.” *The Review of Economics and Statistics* **78(4)**: pp. 748–52.
- FRANKEL, J. & K. FROOT (1987): “Using survey data to test standard propositions regarding exchange rate expectations.” *American Economic Review* **77(1)**: pp. 133–53.
- FRANKEL, J. & J. POONAWALA (2010): “The forward market in emerging currencies: Less biased than in major currencies.” *Journal of International Money and Finance* **29(3)**: pp. 585–598.
- FRANKEL, J. A. (1982): “In search of the exchange risk premium: A six-currency test assuming mean-variance optimization.” *Journal of International Money and Finance* **1(1)**: pp. 255–274.
- FRANKEL, J. A. (1986): “The implications of mean-variance optimization for four questions in international macroeconomics.” *Journal of International Money and Finance* **5(1)**: pp. S53–S75.
- FROOT, K. A. & J. A. FRANKEL (1989): “Forward Discount Bias: Is it an Exchange Risk Premium?” *The Quarterly Journal of Economics* **104(1)**: pp. 139–161.
- FROOT, K. A. & R. H. THALER (1990): “Foreign Exchange.” *Journal of Economic Perspectives* **4(3)**: pp. 179–92.
- FURUKAWA, C. (2020): “Publication bias under aggregation frictions: Theory, evidence, and a new correction method.” *Working paper*, MIT.
- GEORGE, E. I. (2010): “Dilution priors: Compensating for model space redundancy.” In “IMS Collections Borrowing Strength: Theory Powering Applications – A Festschrift for Lawrence D. Brown,” volume 6, p. 158–165. Institute of Mathematical Statistics.
- GOODHART, C. (1988): “The foreign exchange market: A random walk with a dragging anchor.” *Economica* **55(220)**: pp. 437–460.
- GOODHART, C. A. E., P. C. MCMAHON, & Y. L. NGAMA (1997): “Why Does the Spot-Forward Discount Fail to Predict Changes in Future Spot Rates.” *International Journal of Finance & Economics* **2(2)**: pp. 121–29.
- GROSSMANN, A., A. A. LEE, & M. W. SIMPSON (2014): “Forward premium anomaly of the British pound and the euro.” *International Review of Financial Analysis* **34(C)**: pp. 140–156.
- HAAB, D. R. & T. NITSCHKA (2020): “Carry trade and forward premium puzzle from the perspective of a safe haven currency.” *Review of International Economics* **28(2)**: pp. 376–394.
- HAI, W., N. C. MARK, & Y. WU (1997): “Understand-

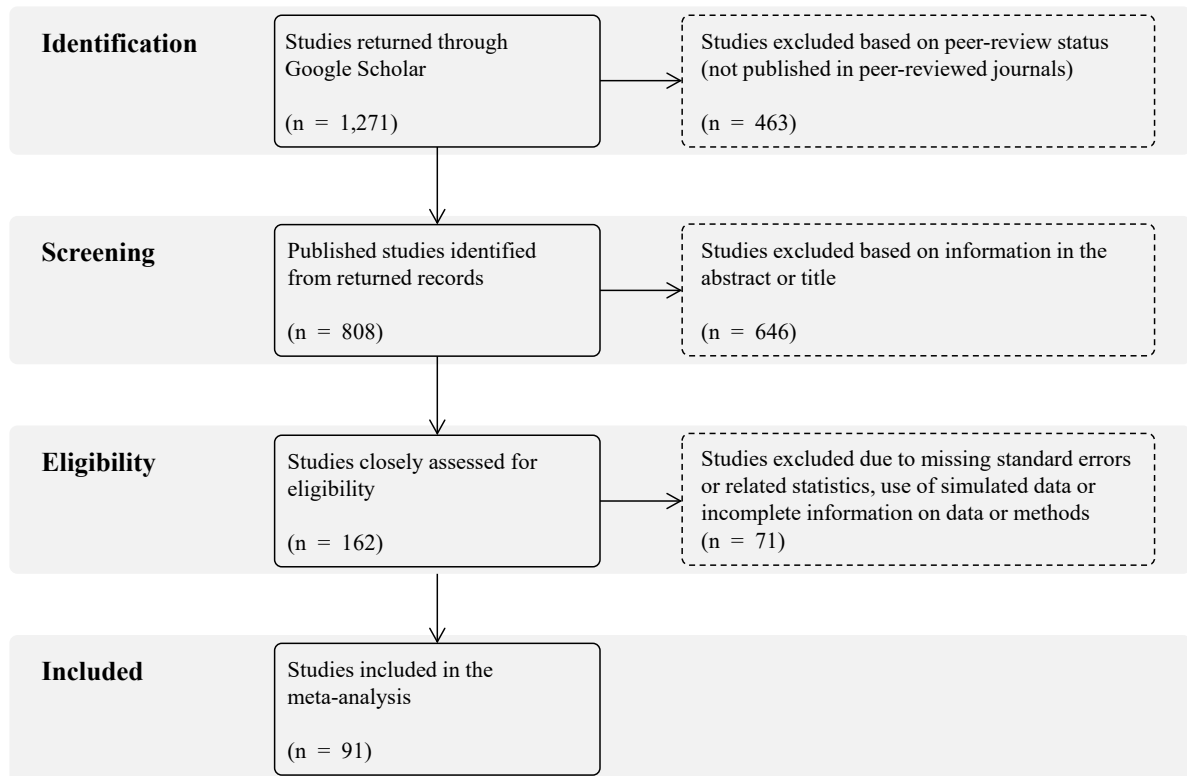
- ing Spot and Forward Exchange Rate Regressions.” *Journal of Applied Econometrics* **12(6)**: pp. 715–34.
- HALL, S. G., A. KENJEGALIEV, P. A. V. B. SWAMY, & G. S. TAVLAS (2013): “The forward rate premium puzzle: a case of misspecification?” *Studies in Non-linear Dynamics & Econometrics* **17(3)**: pp. 265–279.
- HANSEN, B. E. (2007): “Least Squares Model Averaging.” *Econometrica* **75(4)**: pp. 1175–1189.
- HANSEN, L. P. (1982): “Large sample properties of generalized method of moments estimators.” *Econometrica* **50(4)**: pp. 1029–1054.
- HASSAN, T. A. & R. C. MANO (2019): “Forward and Spot Exchange Rates in a Multi-Currency World.” *The Quarterly Journal of Economics* **134(1)**: pp. 397–450.
- HAVRANEK, T. (2015): “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting.” *Journal of the European Economic Association* **13(6)**: pp. 1180–1204.
- HAVRANEK, T., D. HERMAN, , & Z. IRSOVA (2018a): “Does Daylight Saving Save Electricity? A Meta-Analysis.” *The Energy Journal* **39(2)**: pp. 35–61.
- HAVRANEK, T., R. HORVATH, Z. IRSOVA, & M. RUSNAK (2015): “Cross-Country Heterogeneity in Intertemporal Substitution.” *Journal of International Economics* **96(1)**: pp. 100–118.
- HAVRANEK, T. & Z. IRSOVA (2017): “Do Borders Really Slash Trade? A Meta-Analysis.” *IMF Economic Review* **65(2)**: pp. 365–396.
- HAVRANEK, T., Z. IRSOVA, & T. VLACH (2018b): “Measuring the Income Elasticity of Water Demand: The Importance of Publication and Endogeneity Biases.” *Land Economics* **94(2)**: pp. 259–283.
- HAVRANEK, T., Z. IRSOVA, & O. ZEYNALOVA (2018c): “Tuition Fees and University Enrolment: A Meta-Regression Analysis.” *Oxford Bulletin of Economics and Statistics* **80(6)**: pp. 1145–1184.
- HAVRANEK, T., M. RUSNAK, & A. SOKOLOVA (2017): “Habit formation in consumption: A meta-analysis.” *European Economic Review* **95**: pp. 142–167.
- HAVRANEK, T., T. D. STANLEY, H. DOUCOULIAGOS, P. BOM, J. GEYER-KLINGEBERG, I. IWASAKI, W. R. REED, K. ROST, & R. C. M. VAN AERT (2020): “Reporting Guidelines for Meta-Analysis in Economics.” *Journal of Economic Surveys* **34(3)**: pp. 469–475.
- HEDGES, L. V. (1992): “Modeling publication selection effects in meta-analysis.” *Statistical Science* **7(2)**: pp. 246–255.
- IMAI, T., T. A. RUTTER, & C. F. CAMERER (2021): “Meta-Analysis of Present-Bias Estimation using Convex Time Budgets.” *Economic Journal* (**forthcoming**).
- IOANNIDIS, J. P. A., T. D. STANLEY, & H. DOUCOULIAGOS (2017): “The Power of Bias in Economics Research.” *The Economic Journal* **127(605)**: pp. 236–265.
- JEFFREYS, H. (1961): *Theory of Probability (3rd edn)*. Oxford University Press.
- JETTER, M., A. NIKOLSKO-RZHEVSKYY, & O. OGROKHINA (2019): “Can policy shifts explain the forward discount puzzle?” *Empirical Economics* **57(6)**: pp. 1891–1909.
- JONGEN, R., W. F. VERSCHOOR, & C. C. WOLFF (2008): “Foreign Exchange Rate Expectations: Survey And Synthesis.” *Journal of Economic Surveys* **22(1)**: pp. 140–165.
- KIM, S.-H., S. MOON, & C. VELASCO (2017): “Delayed Overshooting: Is It an ’80s Puzzle?” *Journal of Political Economy* **125(5)**: pp. 1570–1598.
- KUMAR, V. (2020): “Liquidity shocks: A new solution to the forward premium puzzle.” *Economic Modelling* **91(C)**: pp. 445–454.
- LEY, E. & M. F. STEEL (2009): “On the Effect of Prior Assumptions in Bayesian Model Averaging with Applications to Growth Regression.” *Applied Econometrics* **24**: pp. 651–674.
- MACDONALD, R. & M. P. TAYLOR (1990): “The Term Structure of Forward Foreign Exchange Premia: The Inter-war Experience.” *The Manchester School of Economic & Social Studies* **58(1)**: pp. 54–65.
- MADIGAN, D., J. YORK, & D. ALLARD (1995): “Bayesian graphical models for discrete data.” *International Statistical Review* **63(2)**: pp. 215–232.
- MATOUSEK, J., T. HAVRANEK, & Z. IRSOVA (2021): “Individual Discount Rates: A Meta-Analysis of Experimental Evidence.” *Experimental Economics* (**forthcoming**).
- MCCALLUM, B. T. (1994): “A reconsideration of the uncovered interest parity relationship.” *Journal of Monetary Economics* **33(1)**: pp. 105–132.
- MCCLOSKEY, D. N. & S. T. ZILIAK (2019): “What Quantitative Methods Should We Teach to Graduate Students? A Comment on Swann’s Is Precise Econometrics an Illusion?” *Journal of Economic Education* **50(4)**: pp. 356–361.
- McFARLAND, J. W., P. C. McMAHON, & Y. NGAMA (1994): “Forward exchange rates and expectations during the 1920s: A re-examination of the evidence.” *Journal of International Money and Finance* **13(6)**: pp. 627–636.
- MUSSA, M. (1979): “Empirical regularities in the behavior of exchange rates and theories of the foreign exchange market.” *Carnegie-Rochester Conference Series on Public Policy* **11(1)**: pp. 9–57.
- NADAL DE SIMONE, F. & W. RAZZAK (1999): “Nominal Exchange Rates and Nominal Interest Rate Differentials.” *IMF Working Papers 99/141*, International Monetary Fund.
- PARK, C. & S. PARK (2017): “Can monetary policy cause the uncovered interest parity puzzle?” *Japan and the World Economy* **41(C)**: pp. 34–44.
- PIPPENGER, J. (2011): “The solution to the forward-bias puzzle.” *Journal of International Financial*

- Markets, Institutions and Money* **21(2)**: pp. 296–304.
- RAZZAK, W. A. (2002): “The Forward Rate Unbiasedness Hypothesis Revisited.” *International Journal of Finance & Economics* **7(4)**: pp. 293–308.
- REED, J. R. (2019): “The forward premium puzzle and Markov-switching adaptive learning.” *Journal of Macroeconomics* **59(C)**: pp. 1–17.
- STANLEY, T. & H. DOUCOULIAGOS (2012): *Meta-Regression Analysis in Economics and Business*. London: Routledge.
- STANLEY, T. & H. DOUCOULIAGOS (2014): “Meta-regression approximations to reduce publication selection bias.” *Research Synthesis Methods* **5(1)**: pp. 60–78.
- STANLEY, T. D. (2001): “Wheat from Chaff: Meta-Analysis as Quantitative Literature Review.” *Journal of Economic Perspectives* **15(3)**: pp. 131–150.
- STANLEY, T. D. (2005): “Beyond Publication Bias.” *Journal of Economic Surveys* **19(3)**: pp. 309–345.
- STANLEY, T. D. (2008): “Meta-Regression Methods for Detecting and Estimating Empirical Effects in the Presence of Publication Selection.” *Oxford Bulletin of Economics and Statistics* **70(1)**: pp. 103–127.
- STANLEY, T. D. & H. DOUCOULIAGOS (2010): “Picture This: A Simple Graph that Reveals Much Ado about Research.” *Journal of Economic Surveys* **24(1)**: pp. 170–191.
- STANLEY, T. D. & H. DOUCOULIAGOS (2017): “Neither Fixed nor Random: Weighted Least Squares Meta-Regression.” *Research Synthesis Methods* **8(1)**: pp. 19–42.
- STANLEY, T. D., H. DOUCOULIAGOS, & J. P. IOANNIDIS (2017): “Finding the Power to Reduce Publication Bias.” *Statistics in Medicine* **36(10)**: pp. 1580–1598.
- STANLEY, T. D., S. B. JARRELL, & H. DOUCOULIAGOS (2010): “Could It Be Better to Discard 90% of the Data? A Statistical Paradox.” *The American Statistician* **64(1)**: pp. 70–77.
- STEEL, M. F. J. (2020): “Model Averaging and its Use in Economics.” *Journal of Economic Literature* **58(3)**: pp. 644–719.
- WOLF, C. C. P. (1987): “Forward Foreign Exchange Rates, Expected Spot Rates, and Premia: A Signal-Extraction Approach.” *Journal of Finance* **42(2)**: pp. 395–406.
- YU, J. (2013): “A sentiment-based explanation of the forward premium puzzle.” *Journal of Monetary Economics* **60(4)**: pp. 474–491.
- ZHOU, S. & A. M. KUTAN (2005): “Does the forward premium anomaly depend on the sample period used or on the sign of the premium?” *International Review of Economics & Finance* **14(1)**: pp. 17 – 25.
- ZIGRAIOVA, D. & T. HAVRANEK (2016): “Bank Competition And Financial Stability: Much Ado About Nothing?” *Journal of Economic Surveys* **30(5)**: pp. 944–981.



## A Details on the Selection of Studies

Figure A1: A PRISMA flow diagram




## B BMA Diagnostics and Robustness Checks (for online publication)

Table B1: *Diagnostics of the baseline BMA estimation*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
20.0264	$5 \cdot 10^6$	$1 \cdot 10^6$	7.664059 mins	751,865
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$1.8 \cdot 10^{13}$	0.0000043%	78%	0.9993	2,989
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random-dilution / 22	UIP	$A_V = 0.9997$		


*Notes:* In the baseline model, we employ the priors suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data) and using the dilution prior suggested by George (2010), which accounts for collinearity. The results of this BMA exercise are reported in Table 4.

Figure B1: *Model size and convergence of the baseline BMA estimation*



*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of the BMA exercise reported in Table 4.

Figure B2: *Model inclusion in BMA (BRIC and random priors)*




*Notes:* The figure depicts results of the BMA model, where the response variable is an estimate of slope coefficient  $\beta$  from Equation 3. In this specification, we employ BRIC g-prior suggested by Fernandez *et al.* (2001) and the beta-binomial (random) model prior according to Ley & Steel (2009). Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities from the 5,000 best models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model. All the variables are reported in Table 3.

Table B2: *Diagnostics of the BMA estimation (BRIC and random priors)*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
20.0179	$5 \cdot 10^6$	$1 \cdot 10^6$	8.213935 mins	750,145
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$1.8 \cdot 10^{13}$	0.0000043%	78%	0.9994	2,989
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random / 22	BRIC	$A_v = 0.9997$		


*Notes:* The specification uses a BRIC g-prior suggested by Fernandez *et al.* (2001) and the beta-binomial model prior according to Ley & Steel (2009).

Figure B3: *Model size and convergence of the BMA estimation (BRIC and random priors)*



*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of the BMA exercise using the BRIC prior for estimated parameters and the random prior for estimated models.

Figure B4: *Model inclusion in BMA (hyper-g and random priors)*



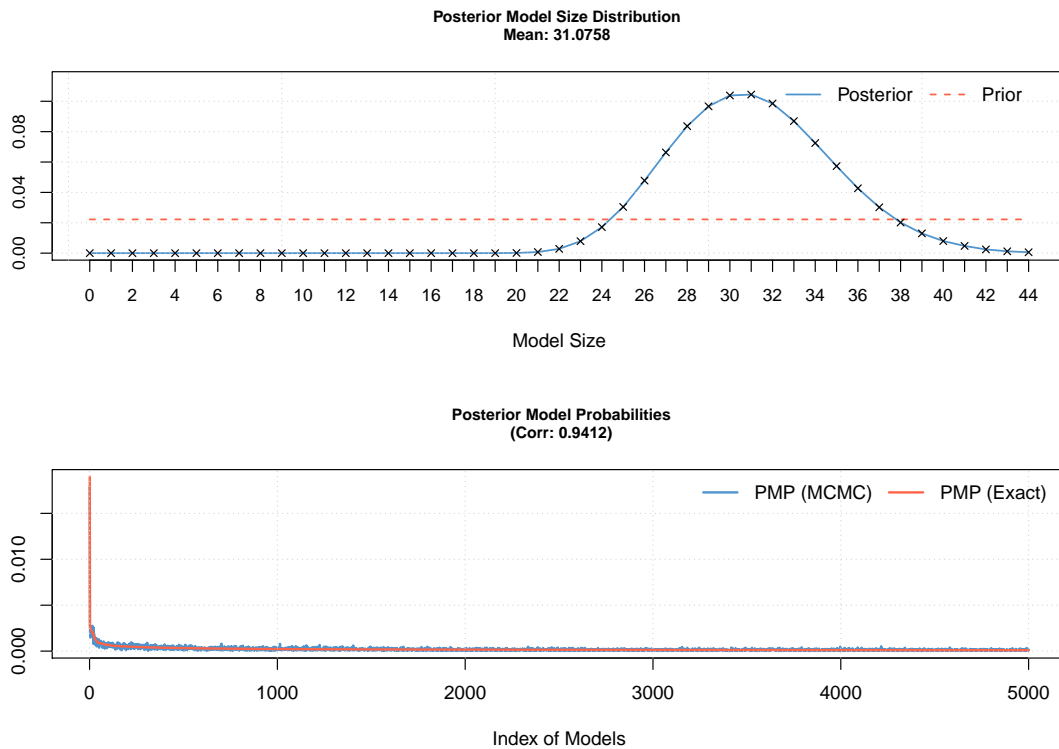
*Notes:* The figure depicts results of the BMA model, where the response variable is an estimate of slope coefficient  $\beta$  from Equation 3. In this specification, we employ the data-dependent hyper-g prior suggested by Feldkircher & Zeugner (2012) and the beta-binomial (random) model prior according to Ley & Steel (2009). Columns show individual models, and variables are listed in descending order by their posterior inclusion probabilities. The horizontal axis shows cumulative posterior model probabilities from the 5,000 best models. To ensure convergence of the Markov Chain Monte Carlo sampler, we use 5,000,000 iterations with 1,000,000 burn-ins to allow the sampler to converge to the part of the model space with high posterior probability models. Blue color (darker in grayscale) = the variable is included in the model with a positive sign. Red color (lighter in grayscale) = the variable is included in the model with a negative sign. No color = the variable is missing from the model. All the variables are reported in Table 3.

Table B3: *Diagnostics of the BMA estimation (hyper-g and random priors)*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
31.0758	$5 \cdot 10^6$	$1 \cdot 10^6$	20.05682 mins	2,118,785
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$1.8 \cdot 10^{13}$	0.000012%	3.3%	0.9412	2,989
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random / 22	hyper (a=2.001)	Av = 0.9746 St. dev. = 0.007		

*Notes:* The specification reported in Figure B4 uses the data-dependent hyper-g prior suggested by Feldkircher & Zeugner (2012) and a random model prior advocated by Ley & Steel (2009).

Figure B5: *Model size and convergence of the BMA estimation (hyper-g and random priors)*




*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of the BMA exercise using the hyper-g prior for estimated parameters and random prior for estimated models (Figure B4).

Table B4: *Diagnostics of the BMA estimation (trimmed data)*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
14.8389	$5 \cdot 10^6$	$1 \cdot 10^6$	9.700674 mins	694,939
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$1.8 \cdot 10^{13}$	0.000004%	79%	0.9997	2,989
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random-dilution / 22	UIP	$A_v = 0.9996$		

*Notes:* In the trimming-data model reported in Figure 10, we employ the priors suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data) and using the dilution prior suggested by George (2010), which accounts for collinearity.

Figure B6: *Model size and convergence of the BMA estimation (trimmed data)*



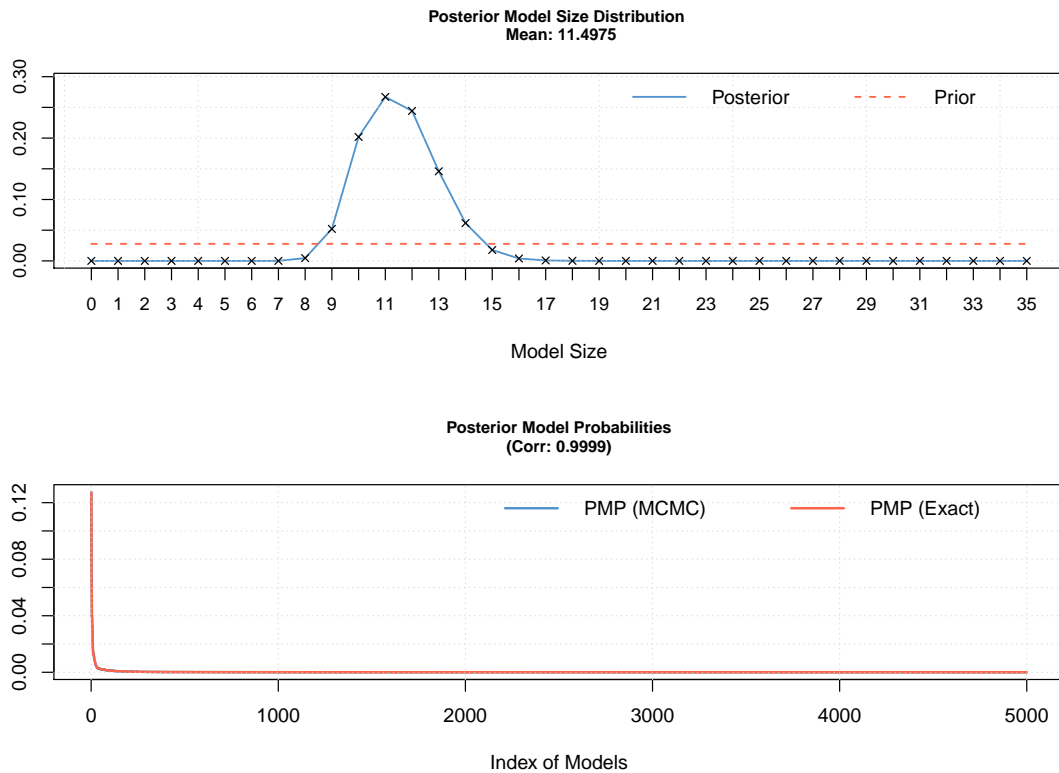
*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of the trimmed-data BMA exercise (Figure 10) using unit information prior and dilution prior.

Table B5: *Diagnostics of the BMA estimation (country variables)*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
11.4975	$5 \cdot 10^6$	$1 \cdot 10^6$	4.995831 mins	597,272
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$3.4 \cdot 10^{10}$	0.0017%	97%	0.9999	2,989
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random-dilution / 22	UIP	$A_v = 0.9997$		

*Notes:* In this specification reported in Figure 11, we employ the priors suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data) and using the dilution prior suggested by George (2010), which accounts for collinearity.

Figure B7: *Model size and convergence of the BMA estimation (country variables)*



*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of the cross-country BMA exercise (Figure 11) employing the unit information prior and dilution prior.




Table B6: *Diagnostics of the BMA estimation (no survey data)*

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>	<i>No. models visited</i>
20.4277	$5 \cdot 10^6$	$1 \cdot 10^6$	6.198806 mins	727,647
<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>	<i>Corr PMP</i>	<i>No. obs.</i>
$1.8 \cdot 10^{13}$	0.0000041%	91%	0.9994	2,943
<i>Model prior</i>	<i>g-prior</i>	<i>Shrinkage-stats</i>		
Random-dilution / 22	UIP	Av = 0.9997		

*Notes:* In this specification reported in Figure 12, we employ the priors suggested by Eicher *et al.* (2011), who recommend using the unit information prior (the prior has the same weight as one observation of data) and using the dilution prior suggested by George (2010), which accounts for collinearity.

Figure B8: *Model size and convergence the BMA estimation (no survey data)*



*Notes:* The figure depicts the posterior model size distribution and the posterior model probabilities of BMA exercise using data WO survey datasets (Figure 12) employing the unit information prior and dilution prior.

## C Studies Included in the Meta-analysis (for online publication)

- AGGARWAL, R., W. T. LIN, & S. K. MOHANTY (2008): "Are Forward Exchange Rates Rational Forecasts of Future Spot Rates? An Improved Econometric Analysis for the Major Currencies." *Multinational Finance Journal* **12(1-2)**: pp. 1–20.
- AGGARWAL, R., B. M. LUCEY, & S. K. MOHANTY (2009): "The Forward Exchange Rate Bias Puzzle Is Persistent: Evidence from Stochastic and Nonparametric Cointegration Tests." *The Financial Review* **44(4)**: pp. 625–645.
- AHMAD, R., S. G. RHEE, & Y. M. WONG (2012): "Foreign exchange market efficiency under recent crises: Asia-Pacific focus." *Journal of International Money and Finance* **31(6)**: pp. 1574–1592.
- AL-ZOUBI, H. A. (2011): "A new look at the forward premium 'puzzle'." *Journal of Futures Markets* **31(7)**: pp. 599–628.
- ALEXIUS, A. (2001): "Uncovered Interest Parity Revisited." *Review of International Economics* **9(3)**: pp. 505–517.
- AMRI, S. (2008): "Analysing the forward premium anomaly using a Logistic Smooth Transition Regression model." *Economics Bulletin* **6(26)**: pp. 1–18.
- AZOUZI, D., R. V. KUMAR, & C. ALOUI (2011): "Forward Rate Unbiasedness Hypothesis in the Tunisian Exchange Rate Market." *International Journal of Academic Research in Business and Social Sciences* **1(2)**: pp. 17–44.
- BAI, S. & A. V. MOLLIK (2010): "Currency crisis and the forward discount bias: Evidence from emerging economies under breaks." *Journal of International Financial Markets, Institutions and Money* **20(5)**: pp. 556–574.
- BAILLIE, R. T. & S. S. CHANG (2011): "Carry trades, momentum trading and the forward premium anomaly." *Journal of Financial Markets* **14(3)**: pp. 441–464.
- BAILLIE, R. T. & D. CHO (2014): "Time variation in the standard forward premium regression: Some new models and tests." *Journal of Empirical Finance* **29(C)**: pp. 52–63.
- BAILLIE, R. T. & R. KILIC (2006): "Do asymmetric and nonlinear adjustments explain the forward premium anomaly?" *Journal of International Money and Finance* **25(1)**: pp. 22–47.
- BAILLIE, R. T., R. E. LIPPENS, & P. C. MCMAHON (1983): "Testing Rational Expectations and Efficiency in the Foreign Exchange Market." *Econometrica* **51(3)**: pp. 553–563.
- BAKSHI, G. S. & A. NAKA (1997): "Unbiasedness of the Forward Exchange Rates." *The Financial Review* **32(1)**: pp. 145–162.
- BANSAL, R. & M. DAHLQUIST (2000): "The forward premium puzzle: different tales from developed and emerging economies." *Journal of International Economics* **51(1)**: pp. 115–144.
- BARNHART, S., R. MCNOWN, & M. WALLACE (2002): "Some answers to puzzles in testing unbiasedness in the foreign exchange market." *Applied Financial Economics* **12(10)**: pp. 687–696.
- BARNHART, S. W., R. MCNOWN, & M. S. WALLACE (1999): "Non-Informative Tests of the Unbiased Forward Exchange Rate." *Journal of Financial and Quantitative Analysis* **34(2)**: pp. 265–291.
- BARNHART, S. W. & A. C. SZAKMARY (1991): "Testing the Unbiased Forward Rate Hypothesis: Evidence on Unit Roots, Co-Integration, and Stochastic Coefficients." *Journal of Financial and Quantitative Analysis* **26(2)**: pp. 245–267.
- BEKAERT, G. & R. J. HODRICK (1993): "On biases in the measurement of foreign exchange risk premiums." *Journal of International Money and Finance* **12(2)**: pp. 115–138.
- BHATTI, R. & I. MOOSA (1995): "An alternative approach to testing uncovered interest parity." *Applied Economics Letters* **2(12)**: pp. 478–481.
- BHATTI, R. H. (2014): "The existence of uncovered parity in the CIS countries." *Economic Modelling* **40(C)**: pp. 227–241.
- BOOTHE, P. & D. LONGWORTH (1986): "Foreign exchange market efficiency tests: Implications of recent empirical findings." *Journal of International Money and Finance* **5(2)**: pp. 135–152.
- BREDON, F., D. RIME, & P. VITALE (2016): "Carry Trades, Order Flow, and the Forward Bias Puzzle." *Journal of Money, Credit and Banking* **48(6)**: pp. 1113–1134.
- BREUER, J. B. & M. E. WOHR (1996): "The Road Less Travelled: Institutional Aspects of Data and Their Influence on Empirical Estimates with an Application to Tests of Forward Rate Unbiasedness." *Economic Journal* **106(434)**: pp. 26–38.
- BURNSIDE, C., B. HAN, D. HIRSHLEIFER, & T. Y. WANG (2011): "Investor Overconfidence and the Forward Premium Puzzle." *Review of Economic Studies* **78(2)**: pp. 523–558.
- CAVAGLIA, S. M. F. G., W. F. C. VERSCHOOR, & C. C. P. WOLFF (1994): "On the Biasedness of Forward Foreign Exchange Rates: Irrationality or Risk Premia?" *The Journal of Business* **67(3)**: pp. 321–343.
- CHAKRABORTY, A. (2007): "Learning, Forward Premium Puzzle and Exchange Rate Fundamentals under Sticky Prices." *Economics Bulletin* **6(34)**: pp. 1–13.
- CHAKRABORTY, A. & S. E. HAYNES (2008): "Econo-

- metrics of the Forward Premium Puzzle." *Economics Bulletin* **6(42)**: pp. 1–17.
- CHIANG, T. C. (1986): "Empirical Analysis On The Predictors Of Future Spot Rates." *Journal of Financial Research* **9(2)**: pp. 153–162.
- CHIANG, T. C. (1988): "The Forward Rate as a Predictor of the Future Spot Rate—A Stochastic Coefficient Approach." *Journal of Money, Credit and Banking* **20(2)**: pp. 212–232.
- CHO, D. (2018): "On the persistence of the forward premium in the joint presence of nonlinearity, asymmetry, and structural changes." *Economic Modelling* **70(C)**: pp. 310–319.
- CHRISTENSEN, M. (2000): "Uncovered interest parity and policy behavior: new evidence." *Economics Letters* **69(1)**: pp. 81–87.
- CLARIDA, R., J. DAVIS, & N. PEDERSEN (2009): "Currency carry trade regimes: Beyond the Fama regression." *Journal of International Money and Finance* **28(8)**: pp. 1375–1389.
- DA COSTA, C. E., J. DE JESUS FILHO, & P. MATOS (2016): "Forward-premium puzzle: is it time to abandon the usual regression?" *Applied Economics* **48(30)**: pp. 2852–2867.
- DIAMANDIS, P. F., D. A. GEORGOUTSOS, & G. P. KOURETAS (2008): "Testing the forward rate unbiasedness hypothesis during the 1920s." *Journal of International Financial Markets, Institutions and Money* **18(4)**: pp. 358–373.
- DING, L. (2012): "The Thursday effect of the forward premium puzzle." *International Review of Economics & Finance* **21(1)**: pp. 302–318.
- DING, L. & L. TO (2010): "The Forward Premium Puzzle Across Maturities." *Economics Bulletin* **30(2)**: pp. 1113–1119.
- ENGEL, C. (1996): "The forward discount anomaly and the risk premium: A survey of recent evidence." *Journal of Empirical Finance* **3(2)**: pp. 123–192.
- FAMA, E. F. (1984): "Forward and spot exchange rates." *Journal of Monetary Economics* **14(3)**: pp. 319–338.
- FRANKEL, J. & J. POONAWALA (2010): "The forward market in emerging currencies: Less biased than in major currencies." *Journal of International Money and Finance* **29(3)**: pp. 585–598.
- FRANKEL, J. A. (1980): "Tests of rational expectations in the forward exchange market." *Southern Economic Journal* **46(4)**: pp. 1083–1101.
- FRANKEL, J. A. & M. D. CHINN (1993): "Exchange Rate Expectations and the Risk Premium: Tests for a Cross Section of 17 Currencies." *Review of International Economics* **1(2)**: pp. 136–144.
- FRENKEL, J. A. (1976): "A monetary approach to the exchange rate: Doctrinal aspects and empirical evidence." *The Scandinavian Journal of Economics* **78(2)**: pp. 200–224.
- FRENKEL, J. A. (1981): "Flexible Exchange Rates, Prices, and the Role of 'News': Lessons from the 1970s." *Journal of Political Economy* **89(4)**: pp. 665–705.
- FROOT, K. A. & J. A. FRANKEL (1989): "Forward Discount Bias: Is it an Exchange Risk Premium?" *The Quarterly Journal of Economics* **104(1)**: pp. 139–161.
- FUKUTA, Y. & M. SAITO (2002): "Forward Discount Puzzle and Liquidity Effects: Some Evidence from Exchange Rates among the United States, Canada, and Japan." *Journal of Money, Credit and Banking* **34(4)**: pp. 1014–1033.
- GOODHART, C. (1988): "The foreign exchange market: A random walk with a dragging anchor." *Economica* **55(220)**: pp. 437–460.
- GOSPODINOV, N. (2009): "A New Look at the Forward Premium Puzzle." *Journal of Financial Econometrics* **7(3)**: pp. 312–338.
- GREGORY, A. W. & T. H. MCCURDY (1984): "Testing the unbiasedness hypothesis in the forward foreign exchange market: A specification analysis." *Journal of International Money and Finance* **3(3)**: pp. 357–368.
- GREGORY, A. W. & T. H. MCCURDY (1986): "The unbiasedness hypothesis in the forward foreign exchange market: A specification analysis with application to France, Italy, Japan, the United Kingdom and West Germany." *European Economic Review* **30(2)**: pp. 365–381.
- GROSSMANN, A., A. A. LEE, & M. W. SIMPSON (2014): "Forward premium anomaly of the British pound and the euro." *International Review of Financial Analysis* **34(C)**: pp. 140–156.
- HAI, W., N. C. MARK, & Y. WU (1997): "Understanding Spot and Forward Exchange Rate Regressions." *Journal of Applied Econometrics* **12(6)**: pp. 715–734.
- HALL, S. G., A. KENJEGALIEV, P. A. V. B. SWAMY, & G. S. TAVLAS (2013): "The forward rate premium puzzle: a case of misspecification?" *Studies in Nonlinear Dynamics & Econometrics* **17(3)**: pp. 265–279.
- HO, T.-W. (2002): "The Forward Rate Unbiasedness Hypothesis revisited." *Applied Financial Economics* **12(11)**: pp. 799–804.
- HO, T.-W. & W.-S. MO (2016): "Testing the Persistence of the Forward Premium: Structural Changes or Misspecification?" *Open Economies Review* **27(1)**: pp. 119–138.
- HODGSON, D. J., O. LINTON, & K. VORKINK (2004): "Testing forward exchange rate unbiasedness efficiently: a semiparametric approach." *Journal of Applied Economics* **7(2)**: pp. 325–353.
- HUANG, R. D. (1984): "Some alternative tests of forward exchange rates as predictors of future spot rates." *Journal of International Money and Finance* **3(2)**: pp. 153–167.

- HUISMAN, R., K. KOEDIJK, C. KOOL, & F. NISSEN (1998): "Extreme support for uncovered interest parity." *Journal of International Money and Finance* **17(1)**: pp. 211–228.
- KUMAR, R. V. & D. AZOUZI (2011): "Tunisian and Indian Forex Markets: A Comparison on Forward Rate Unbiased Hypothesis." *Romanian Economic Journal* **14(40)**: pp. 81–98.
- KUMAR, S. & S. TRÁĽCK (2014): "Unbiasedness and risk premiums in the Indian currency futures market." *Journal of International Financial Markets, Institutions and Money* **29(C)**: pp. 13–32.
- LIU, F. & P. SERCU (2009): "The Forex Forward Puzzle: The Career Risk Hypothesis." *The Financial Review* **44(3)**: pp. 371–404.
- LONGWORTH, D. (1981): "Testing the Efficiency of the Canadian-U.S. Exchange Market under the Assumption of no Risk Premium." *Journal of Finance* **36(1)**: pp. 43–49.
- LORING, G. & B. LUCEY (2013): "An analysis of forward exchange rate biasedness across developed and developing country currencies: Do observed patterns persist out of sample?" *Emerging Markets Review* **17(C)**: pp. 14–28.
- LUNTEL, K. B. & K. PAUDYAL (1998): "Common stochastic trends between forward and spot exchange rates." *Journal of International Money and Finance* **17(2)**: pp. 279–297.
- MACKLEM, R. T. (1991): "Forward exchange rates and risk premiums in artificial economies." *Journal of International Money and Finance* **10(3)**: pp. 365–391.
- MARK, N. C. & Y.-K. MOH (2007): "Official interventions and the forward premium anomaly." *Journal of Empirical Finance* **14(4)**: pp. 499–522.
- MARK, N. C. & Y. WU (1998): "Rethinking Deviations from Uncovered Interest Parity: The Role of Covariance Risk and Noise." *Economic Journal* **108(451)**: pp. 1686–1706.
- MAYNARD, A. (2003): "Testing for Forward-Rate Unbiasedness: On Regression in Levels and in Returns." *The Review of Economics and Statistics* **85(2)**: pp. 313–327.
- MAYNARD, A. & P. C. B. PHILLIPS (2001): "Rethinking an old empirical puzzle: econometric evidence on the forward discount anomaly." *Journal of Applied Econometrics* **16(6)**: pp. 671–708.
- MCCALLUM, B. T. (1994): "A reconsideration of the uncovered interest parity relationship." *Journal of Monetary Economics* **33(1)**: pp. 105–132.
- McFARLAND, J. W., P. C. McMAHON, & Y. NGAMA (1994): "Forward exchange rates and expectations during the 1920s: A re-examination of the evidence." *Journal of International Money and Finance* **13(6)**: pp. 627–636.
- McFARLAND, J. W., P. C. B. PHILLIPS, & P. C. McMAHON (1996): "Robust Tests of Forward Exchange Market Efficiency with Empirical Evidence from the 1920s." *Journal of Applied Econometrics* **11(1)**: pp. 1–22.
- MEHL, A. & L. CAPPIELLO (2009): "Uncovered Interest Parity at Long Horizons: Evidence on Emerging Economies." *Review of International Economics* **17(5)**: pp. 1019–1037.
- MOOSA, I. A. (2004): "An empirical examination of the Post Keynesian view of forward exchange rates." *Journal of Post Keynesian Economics* **26(3)**: pp. 395–418.
- NAKA, A. & G. WHITNEY (1995): "The unbiased forward rate hypothesis re-examined." *Journal of International Money and Finance* **14(6)**: pp. 857–867.
- NEWBOLD, P., M. E. WOHR, T. RAYNER, N. KELLARD, & C. ENNEW (1998): "Two puzzles in the analysis of foreign exchange market efficiency." *International Review of Financial Analysis* **7(2)**: pp. 95–111.
- NIKOLAOU, K. & L. SARNO (2006): "New evidence on the forward unbiasedness hypothesis in the foreign-exchange market." *Journal of Futures Markets* **26(7)**: pp. 627–656.
- OLMO, J. & K. PILBEAM (2009): "Uncovered Interest Parity: Are Empirical Rejections of It Valid?" *Journal of Economic Integration* **24(2)**: pp. 369–384.
- PHILLIPS, P. C. B. & J. W. McFARLAND (1997): "Forward exchange market unbiasedness: the case of the Australian dollar since 1984." *Journal of International Money and Finance* **16(6)**: pp. 885–907.
- PIPPENGER, J. (2011): "The solution to the forward-bias puzzle." *Journal of International Financial Markets, Institutions and Money* **21(2)**: pp. 296–304.
- RAZZAK, W. A. (2002): "The Forward Rate Unbiasedness Hypothesis Revisited." *International Journal of Finance & Economics* **7(4)**: pp. 293–308.
- ROLL, R. & S. YAN (2000): "An explanation of the forward premium 'puzzle'." *European Financial Management* **6(2)**: pp. 121–148.
- SAKOULIS, G., E. ZIVOT, & K. CHOI (2010): "Structural change in the forward discount: Implications for the forward rate unbiasedness hypothesis." *Journal of Empirical Finance* **17(5)**: pp. 957–966.
- SARNO, L., G. VALENTE, & H. LEON (2006): "Nonlinearity in Deviations from Uncovered Interest Parity: An Explanation of the Forward Bias Puzzle." *Review of Finance* **10(3)**: pp. 443–482.
- SNAITH, S., J. COAKLEY, & N. KELLARD (2013): "Does the forward premium puzzle disappear over the horizon?" *Journal of Banking & Finance* **37(9)**: pp. 3681–3693.
- SOSVILLA-RIVERO, S. & Y. B. PARK (1992): "Further tests on the forward exchange rate unbiasedness hypothesis." *Economics Letters* **40(3)**: pp. 325–331.
- VASILYEV, D., V. BUSYGIN, & S. BUSYGIN (2017): "Testing and interpreting uncovered interest parity

- in Russia.” *Russian Journal of Economics* **3(2)**: pp. 158 – 173.
- VERSCHOOR, W. F. C. & C. C. P. WOLFF (2001): “Scandinavian forward discount bias risk premia.” *Economics Letters* **73(1)**: pp. 65–72.
- VILLANUEVA, O. M. (2007): “Spot-forward cointegration, structural breaks and FX market unbiasedness.” *Journal of International Financial Markets, Institutions and Money* **17(1)**: pp. 58–78.
- WANG, P. & P. WANG (2009): “Does a “correct” parameter estimate tell a better story about foreign exchange market efficiency?” *Journal of International Money and Finance* **28(2)**: pp. 183 – 197.
- WESTERLUND, J. (2007): “Estimating Cointegrated Panels with Common Factors and the Forward Rate Unbiasedness Hypothesis.” *Journal of Financial Econometrics* **5(3)**: pp. 491–522.
- WU, Y. & H. ZHANG (1997): “Forward premiums as unbiased predictors of future currency depreciation: a non-parametric analysis.” *Journal of International Money and Finance* **16(4)**: pp. 609–623.
- ZHOU, S. (2002): “The forward premium anomaly and the trend behavior of the exchange rates.” *Economics Letters* **76(2)**: pp. 273–279.
- ZHOU, S. & A. M. KUTAN (2005): “Does the forward premium anomaly depend on the sample period used or on the sign of the premium?” *International Review of Economics & Finance* **14(1)**: pp. 17–25.
- ZIVOT, E. (2000): “Cointegration and forward and spot exchange rate regressions.” *Journal of International Money and Finance* **19(6)**: pp. 785–812.