Commodity Prices and Currencies*

Alexandre Jeanneret
UNSW Business School

Valeri Sokolovski
University of Alberta

September 10, 2023

Abstract

We introduce an empirical approach to identify *commodity currencies* as those with significant commodity price exposure. This categorization aligns with the importance of a country's commodity sector across multiple dimensions. Studying these currencies, we find that monthly changes in a country's commodity export prices predict its exchange rate, especially when uncertainty is high. This predictability extends to the carry trade and is driven exclusively by investments in commodity-exposed currencies. These results hold out-of-sample, surpassing the random walk benchmark, particularly for emerging currencies. We explain our findings using a model incorporating heterogeneous beliefs among agents regarding the informativeness of news.

Keywords: Exchange rates, commodities, predictability, carry trade, FX volatility.
JEL codes: C32, F31, G15.

*We are grateful for comments and suggestions from Patrick Augustin, Geert Bekaert, Philippe Cote, Pasquale Della Corte, Christian Dorion, Mathieu Fournier, Mark Huson, Òscar Jordá, Richard Levich, Ľuboš Pástor, Ella Patelli, Steven Riddiough (discussant), Nick Roussanov, Ivan Shaliastovich (discussant), Kirsten Smart, Elvira Sojli, Fabricius Somogyi, Simon Van Norden, Pascale Valery, Adrien Verdelhan, Robert Vigfusson (discussant), Colin Ward, Masahiro Watanabe, and seminar participants at the 2019 JPMCC International Symposium, Fulcrum Asset Management, 2023 Frontiers in Finance Conference, 2021 Vienna Symposium on Foreign Exchange Markets, 2021 FIRN, and HEC Montréal. Alexandre Jeanneret (corresponding author) is with the School of Banking and Finance, UNSW Business School. Email: a.jeanneret@unsw.edu.au. Website: www.alexandrejeanneret.net. Valeri Sokolovski is with the Alberta Business School, University of Alberta. E-mail: vsokolov@ualberta.ca; Website: www.valerisokolovski.com.

Commodity Prices and Currencies

Abstract

We introduce an empirical approach to identify *commodity currencies* as those with significant commodity price exposure. This categorization aligns with the importance of a country's commodity sector across multiple dimensions. Studying these currencies, we find that monthly changes in a country's commodity export prices predict its exchange rate, especially when uncertainty is high. This predictability extends to the carry trade and is driven exclusively by investments in commodity-exposed currencies. These results hold out-of-sample, surpassing the random walk benchmark, particularly for emerging currencies. We explain our findings using a model incorporating heterogeneous beliefs among agents regarding the informativeness of news.

Keywords: Carry trade, exchange rates, commodities, FX volatility, predictability.
JEL codes: C32, F31, G15.

1 Introduction

The exchange rate is arguably the most important price in an open economy. Yet the literature has often struggled to empirically connect exchange rates with economic fundamentals. This phenomenon is broadly known as the "exchange rate disconnect", and it remains one of the most persistent puzzles in international finance (Obstfeld and Rogoff, 2000; Itskhoki and Mukhin, 2021). In the short term, exchange rate moves are hard to explain and even harder to predict (Meese and Rogoff, 1983). In this paper, we reexamine short-term exchange rate predictability by focusing on a distinct set of currencies for which the link between exchange rates and fundamentals is empirically strong and theoretically unambiguous: the "commodity currencies".[1]

To illustrate this connection, Figure 1 shows how the Norwegian krone (NOK) and the Russian ruble (RUB) perform in relation to the price of oil (their main commodity export). Despite the NOK being a developed G10 currency and the RUB an emerging currency, both exchange rates closely track the oil price over the long term (2004-2020) as well as around exogenous short-term shocks to oil prices, like during the Russia-Saudi Arabia oil price war.[2] In contrast, it is difficult to think of another variable that would be as closely related to, for example, the Swiss franc or the Japanese yen. The primary reason for this tight link is that commodities play a vital role in multiple sectors of these economies, naturally driving their exchange rates.[3]

Figure 1 [about here]

Although commodity prices significantly influence these exchange rates, participants in the foreign exchange (FX) market are strongly heterogeneous, face asymmetric information, and are likely to trade for a variety of idiosyncratic reasons (Ranaldo and Somogyi, 2021). Their trades could stem from shifts in other fundamentals, monetary policies, market sentiment, corporate decisions, or even noise. Consequently, it is reasonable to expect that fluctuations in commodity prices would only be gradually incorporated into the exchange rates of commodity-producing countries. This gradual information diffusion may result in short-term predictability for these countries' exchange rates, especially for the less frequently traded currencies.

[1] "Commodity currencies" are typically defined as currencies of countries in which primary commodities constitute a significant share of production and exports (we provide our formal definition later in the paper). Chen and Rogoff (2003) introduced this term in their paper titled "Commodity Currencies", finding that commodity prices strongly influence the real exchange rates of Australia, Canada, and New Zealand.

[2] This conflict stemmed from a breakdown in negotiations between the Organization of the Petroleum Exporting Countries (OPEC) and Russia over proposed oil production cuts amid the COVID-19 crisis. On 8 March 2020, Saudi Arabia unexpectedly announced discounts of $6 to $8 per barrel to international customers. This announcement led to a 30% drop in oil prices and a depreciation of the NOK and the RUB. On 2 April 2020, US President Trump threatened to withdraw military support unless OPEC and its allies reduced production. Oil prices surged by about 25% that day, and both currencies subsequently appreciated.

[3] For example, commodity exports constitute 67% of Norway and Russia's exports; commodity-related revenue accounts for 19% and 33% of their total fiscal revenue, respectively; and commodity-linked companies constitute 47% and 69% of their total stock market capitalization, respectively.

Commodity producers' currencies are interesting for another critical reason. These currencies typically offer high-interest rates (Ready, Roussanov, and Ward, 2017), and thus tend to play a crucial role in the currency carry trade - a highly popular FX strategy involving borrowing in low-interest rate countries and investing in high-interest rate countries. The profitability of this strategy has puzzled financial researchers and spawned an extensive literature (see, e.g., Daniel, Hodrick, and Lu, 2017). While substantial progress has been made to explain unconditional carry trade returns and the cross-sectional differences in currency returns (see Hassan and Zhang (2021) for a survey), there has been limited exploration of carry trade predictability to date.

This paper makes four key contributions to the FX literature. First, we develop a simple model showing how changes in commodity prices could impact the contemporaneous and future exchange rates of commodity-exporting countries, especially in times of elevated uncertainty. Second, we introduce a new empirical approach to identifying currencies with significant exposure to their countries' exported commodities. Third, we exploit changes in country-level commodity export prices to provide evidence of unconditional and conditional exchange rate predictability for commodity currencies, both in and out-of-sample. Fourth, we show that commodity price fluctuations are valuable for predicting the performance the carry trade, but that this predictability is driven exclusively by a small set of less-traded commodity currencies, mostly from emerging markets. This evidence of currency predictability has significant implications for investment and policy decisions, which depend highly on the ability to forecast exchange rates.

We start with a motivating, information-based model to study predictability in the FX market. Building on existing models that consider differences of opinion among traders, we explore how new information could drive exchange rate predictability in an economy with heterogenous agents.[4] We find that, when there is disagreement about the informativeness of public news (e.g., commodity price movements), agents trade based on their differing beliefs. Consequently, this news is only gradually incorporated into exchange rates, resulting in short-term predictability. The model's key insight is that variations in commodity prices (i.e., fundamental shocks for commodity producers) affect both current and future exchange rate changes. This information-based model provides key predictions about short-term exchange rate dynamics, complementing the macro-finance literature on currency risk premiums.

Guided by these theoretical insights, we empirically investigate the role of commodity export prices in exchange rate predictability. We start by identifying a set of commodity-exposed currencies, using a sample of 41 (developed and emerging market) currencies and country-specific commodity export price indexes spanning from January 1985 to April 2020.[5] Thus far, the definition of a commodity currency in existing literature has often lacked consistency, with studies frequently analyzing small

[4]Related models with heterogenous investors include Bacchetta and Van Wincoop (2006) and Cespa, Gargano, Riddiough, and Sarno (2022).

[5]The country-specific commodity export price indexes are constructed as the export-weighted changes in international market prices of up to 45 individual commodities. The weights are time-varying to ensure that changes in the price indexes reflect variations in the relevant commodity prices for each country at any given point in time.

sets of reasonable but arbitrarily chosen candidates from major commodity producers (e.g., Chen and Rogoff, 2003; Chen, Rogoff, and Rossi, 2010; Ferraro, Rogoff, and Rossi, 2015; Ready et al., 2017). In contrast, we propose a formal definition of a *commodity currency* as a country's currency with a positive and statistically significant covariance (beta) with its commodity export prices. In other words, a currency that, on average, tends to appreciate when commodity export prices rise and depreciate when they fall.

Upon examining our sample of 41 currencies, we find nine countries whose currencies display a positive and statistically significant commodity price beta: Australia, Brazil, Canada, Mexico, New Zealand, Norway, Peru, Russia, and South Africa. We show that our market-based categorization helps capture the importance of a country's commodity sector across multiple dimensions, including exports, GDP, financial markets, and fiscal revenue. And it does so using a single measure. In contrast, none of these economic measures alone are sufficient for adequate identification of the commodity currencies. For instance, while Colombia has the highest share of commodity exports (68%) in our sample, it does not display a statistically or economically significant commodity price beta once we account for the US dollar effect.[6] In comparison, Mexico displays a strongly significant commodity price beta, yet commodities reflect a relatively modest fraction of its total exports (18%). However, commodity-related revenue constitutes a substantial portion of Mexico's government's income: 55% in 2007 and 28% in 2017 (OECD, 2020). We thus offer the first formal identification of "commodity currencies", which can differ from those of the largest commodity exporters.

Next, we examine how variations in country-level commodity export prices help predict exchange rates for this set of currencies. Our analysis focuses on one-month-ahead predictability using non-overlapping data. We find that these currencies appreciate, both statistically and economically, following an increase in commodity export prices. A one-standard-deviation rise in a country's commodity export prices predicts a 0.37% currency appreciation over the next month (equivalent to 4.4% per annum). Notably, this predictability is short-lived, disappearing after four months. This supports an information-based mechanism where news is gradually reflected in exchange rates. Furthermore, considering the full cross-section of commodity price betas, we find that predictability increases with the currency's commodity beta. This is in line with our model. Intuitively, if commodity shocks have no contemporaneous effect on a currency (zero beta), we should not see a predictive effect either.

We also provide evidence of superior out-of-sample exchange rate predictability relative to the traditional benchmark random walk model. The FX market is considered as one of the most active, liquid, and efficient markets in the world, and predicting short-term exchange rates is known to be notoriously difficult.[7] Consistent with this view, we find limited out-of-sample predictability for

[6] As commodity prices are typically denominated in US dollars, a dollar depreciation tends to mechanically increase commodity prices and appreciate other currencies vis-a-vis the US dollar. Controlling for the dollar factor helps isolate the portion of exchange rate changes unaffected by these effects. Unsurprisingly, ignoring the dollar factor generates unreasonable predictions. For example, one would erroneously categorize the Singapore dollar and the Swedish krona, currencies from countries with minimal commodity export shares, as commodity currencies.

[7] While there is ample evidence of exchange rate predictability at medium to long-term horizons (see, e.g., Mark (1995a), Engel, Mark, and West (2007), Balduzzi and Chiang (2019), and Eichenbaum, Johannsen, and Rebelo

highly-traded currencies, such as the Australian and Canadian dollar. However, we do find robust predictability for the less-liquid commodity currencies, especially those from emerging markets. For instance, predictability is particularly strong for the Brazilian real and Russian ruble, two currencies whose average trading volume is around seven times smaller than that of the Australian dollar. More generally, we uncover a significant inverse relationship between a currency's out-of-sample predictability and its average daily trading volume, in line with a delayed reaction channel.

Our empirical approach mitigates endogeneity concerns by ensuring that exchange rate and commodity export price changes are not driven by variations in financial market conditions. Specifically, we control for a number of exchange rate predictors, including each country's interest rate differential (Fama, 1984), aggregate FX volatility (Bakshi and Panayotov, 2013; Menkhoff, Sarno, Schmeling, and Schrimpf, 2012; Karnaukh, Ranaldo, and Söderlind, 2015), funding liquidity (Mancini, Ranaldo, and Wrampelmeyer, 2013), aggregate market uncertainty (Brunnermeier, Nagel, and Pedersen, 2008; Lustig, Roussanov, and Verdelhan, 2011), and a US recession indicator to account for the aggregate commodity declines during global economic slowdowns. Hence, the commodity export prices we exploit contain unique information unspanned by the factors considered in the extant FX literature.[8]

Our findings remain robust regardless of the chosen base currency. While we primarily focus on exchange rates from the US investors' perspective, commodity export price changes also predict the future performance of commodity-exposed currencies relative to the euro, Swiss franc, and Japanese yen. This dismisses the concern that the exchange rate predictability we document is merely driven by US dollar effects. Furthermore, through a counterfactual analysis, we confirm that the predictive relation between commodity export prices and exchange rates is not present for the set of currencies unrelated to commodity prices. This strengthens our argument against omitted variables (e.g., reflecting global economic conditions) potentially driving our findings. In sum, we find that commodity export prices hold valuable predictive information for only the commodity-exposed currencies' exchange rates, regardless of the base currency.

Next, we explore conditional exchange rate predictability during normal and stressed FX market conditions. Our model predicts that higher FX uncertainty reduces trading among risk-averse agents. This suggests that newly available information takes more time to be incorporated into the exchange rate, leading to stronger predictability. Following this theoretical prediction, we assess the conditional impact of commodity export prices on future exchange rates in a regime-switching environment, using Jordà (2005)'s local projection method. We find that exchange rate predictability is concentrated in times of elevated FX uncertainty, as measured by either realized volatility or dispersion in professional FX forecasts. Thus, the level FX uncertainty plays a key role in conditional exchange rate predictability.

(2021)), most economic variables fail to predict exchange rates at short horizons (i.e., monthly). See Rossi (2013) for a comprehensive literature review. Forward-looking financial measures, however, sometimes have more success in predicting short-term exchange rates (see, e.g., Londono and Zhou, 2017; Della Corte, Jeanneret, and Patelli, 2023) Nevertheless, a common assumption is that exchange rates follow a random walk at short horizons.

[8]Our results are thus unlikely to be influenced by investors jointly trading commodities and currencies, adjusting their positions in both assets as global financial conditions change.

One may be concerned that time variations in FX uncertainty are linked to broader global market changes rather than being specific to the FX market itself. For example, aggregate liquidity tends to evaporate when FX volatility increases (e.g., Karnaukh et al., 2015). Similarly, FX volatility tends to surge when investor fears (VIX) increase, such as during periods of financial turmoil (e.g., Menkhoff et al., 2012). Currencies can also become riskier when FX dealers face tighter funding constraints and money-market premiums increase (e.g., Brunnermeier et al., 2008; Ranaldo and Söderlind, 2010), as indicated by a higher TED spread. However, we find that exchange rate predictability of commodity export prices remains statistically significant and concentrated in times of elevated FX volatility, after orthogonalizing the latter to FX illiquidity, the VIX, and the TED spread.

Lastly, we explore our results' implications for the carry trade. We expand upon previous research on the predictability of carry trade returns using commodity prices in several ways, all of which support our economic narrative.[9] First, we find that the investment component of our strategy is heavily concentrated in emerging currencies, complementing previous work based on G10 currencies (Bakshi and Panayotov, 2013). For example, in contrast to common belief, we find that the investment portfolio rarely contains the Australian dollar, but frequently includes the Brazilian real, the Russian ruble, and the South African rand. Effectively, our carry trade strategy invests in G10 currencies only 6.4% of the time, indicating little overlap with Bakshi and Panayotov (2013). Second, we provide evidence that carry trade predictability arises largely from the consideration of *country-specific* commodity export prices. Even when accounting for global commodity price indices (from CRB, Goldman Sachs, or the oil price), commodity export prices maintain their predictive power, while these global predictors do not. Third, we delve into the origin of carry trade predictability (both theoretically and empirically), a question largely left unaddressed by the existing literature. We form interest-rate-sorted currency portfolios and find that their return predictability significantly increases with their average commodity currency membership. For example, commodity currencies represent an average of 36.8% of the investment leg (top quintile), but only 5.8% of the short leg (bottom quintile). Consequently, commodity prices' predictive power for carry trade returns is driven solely by the currencies with significant commodity price betas that are part of the investment portfolio. Additionally, we find that predictability is concentrated in times of elevated FX volatility, mirroring our findings for individual exchange rates. Furthermore, we exploit our large cross-section of currencies and consider a counterfactual carry trade strategy that excludes commodity currencies. While this alternative strategy remains unconditionally almost as profitable as the unconstrained carry trade, it no longer exhibits commodity export price predictability. In sum, we find that the predictability of the carry trade with commodity prices is purely driven by the exchange rate predictability of a few commodity currencies, which we study in this paper, not because commodity prices capture a global risk factor or because commodity investing coincides with an appetite for risk-taking.

[9] Bakshi and Panayotov (2013) study the predictability of carry trade returns using global predictors such as FX volatility, funding liquidity, and the Commodity Research Bureau (CRB) commodity price index. Relatedly, Opie and Riddiough (2020) study international portfolio hedging using FX factors and document the predictability of Lustig et al. (2011)'s carry trade factor with the CRB index, before using it for portfolio optimization.

Our paper also contributes to the literature on the relationship between commodity prices and exchange rates. Amano and Van Norden (1998) finds that oil prices Granger-cause the real US dollar exchange rate, while Chen and Rogoff (2003) provide evidence that commodity prices are in-sample predictors of quarterly exchange rates for Australia, Canada, and New Zealand. Chen et al. (2010) add two more countries (Chile and South Africa) to their analysis and find that exchange rates predict global commodity prices, but that the reverse does not hold out-of-sample. In contrast, Ferraro et al. (2015), analyzing five commodity-producing countries (Australia, Canada, Norway, Chile and South Africa), find significant out-of-sample predictability (with oil, copper, or gold prices), but only at the daily frequency. We expand on these earlier studies on multiple key dimensions. First, we consider a large cross-section of currencies, which allows us to (i) analyze previously ignored currencies, (ii) identify commodity currencies in a systematic way, (iii) relate commodity exposures to a broad set of economic fundamentals, and (iv) to conduct counterfactual exercises utilizing the full cross-section. Second, in contrast to Chen et al. (2010), we find evidence of significant monthly predictability, especially for emerging market currencies. This predictability, however, is short-lived, aligning with the earlier findings of no quarterly predictability (Chen et al., 2010) and significant daily predictability at (Ferraro et al., 2015). Third, we offer a tractable model that explains why short-term exchange rate predictability with commodity export prices is concentrated in periods of high FX uncertainty. Lastly, we demonstrate our results' critical implications for the carry trade.

Our work is closely related to Ready et al. (2017). Their general-equilibrium model shows that (i) commodity-exporting countries have lower aggregate risk and thus higher interest rates, compared to countries producing final goods; and (ii) commodity currencies appreciate in "good" times and depreciate in "bad" times, thereby earning a risk premium. Their model thus rationalizes why commodity currencies have relatively higher interest rates and offer higher returns, particularly when goods markets are more segmented due to higher trade costs. The authors empirically validate their model's cross-sectional implications and the theoretical prediction that shipping costs positively forecast carry trade returns. Our paper differs from this influential work in several ways. First, we consider an information-based explanation for exchange rate predictability, where exchange rates slowly adjust to commodity price changes. This contrasts with Ready et al. (2017)'s risk-based model, which explicitly predicts no short-term predictability. Our results support our model, showing strong predictability at short horizons, particularly during high FX uncertainty and among emerging currencies. Second, their complete markets model applies elegantly to developed economies, where aggregate consumption is expected to be relatively stable. Our framework, on the other hand, is well-suited to emerging currencies which appear to be less liquid and thus more likely to incorporate information with a delay. This is important because, when considering a large cross-section of currencies, emerging currencies constitute the bulk of the investment side of the carry trade. Overall, the differences between our work and that of Ready et al. (2017) provide complementary insights into the role of commodity prices in carry trade predictability.

More broadly, we contribute to the understanding of the carry trade performance. Existing literature has identified various common risk factors that help explain the cross-section of currency returns and thus the *unconditional* carry trade profitability.[10] Our work complements this strand of the literature by providing evidence that changes in country-specific commodity export prices can explain the time variation in exchange rate changes and, in turn, the *conditional* carry trade performance. Additionally, our findings revisit the connection between the carry trade and individual currency returns. For example, Verdelhan (2018) suggests that the carry trade exposes investors to global risk factors, such that a currency that is more exposed to the carry trade factor is viewed as riskier and earns a higher expected return. Our findings emphasize that individual exchange rates are subject to currency-specific shocks, which in turn could affect the carry trade performance. This is because the carry trade's long portfolio is heavily concentrated in a small set of commodity currencies, whose fluctuations cannot be fully diversified away.[11] Our work therefore highlights a new "individual FX returns to carry trade" channel that complements the existing "carry trade to individual FX returns" channel.

The remainder of the paper is organized as follows. Section 2 presents a simple model that provides guidance on exchange rate predictability. Section 3 describes the data and identifies the set of commodity currencies. Sections 4 and 5 discuss our main empirical findings on the unconditional and conditional exchange rate predictability, respectively. Section 6 extends the analysis to the carry trade, while Section 7 provides an out-of-sample analysis. Section 8 concludes.

2 Motivating theory

We present a stylized model of exchange rate determination with heterogeneous agents to explore how variation in commodity export prices can generate exchange rate predictability. The setting builds on existing equilibrium models with differences of opinion among traders.[12] In our model, agents "agree to disagree" on the relevance of using commodity export prices for predicting exchange rates, even if they have access to the same publicly available information. When agents trade based on their different beliefs, we find that new information is slow to reflect in exchange rates. This delay then results in future exchange rates being predictable by changes in commodity export prices.

Our framework departs from the existing macro-finance literature, which has developed models to uncover sources of currency risk premiums. This literature focuses primarily on the cross-sectional analysis of currency excess returns or on long-term predictability.[13] In contrast to existing risk

[10]Unconditional currency (excess) returns reflect compensation for investors' exposure to global factors, such as consumption growth risk (Lustig and Verdelhan, 2007), consumption habits (Verdelhan, 2010), average excess returns (Lustig et al., 2011), systematic FX volatility (Menkhoff et al., 2012), systematic liquidity (Mancini et al., 2013), global imbalance risk (Della Corte, Riddiough, and Sarno, 2016), crash risk (Chernov, Graveline, and Zviadadze, 2018), the business cycle (Colacito, Riddiough, and Sarno, 2020), and FX liquidity risk (Söderlind and Somogyi, 2023).

[11]Fluctuations in country-specific commodity export prices, similar to firm-specific shocks in a granular economy (Gabaix, 2011), have significant implications for aggregate asset pricing.

[12]See, for example, Harrison and Kreps (1978), Harris and Raviv (1993), Kandel and Pearson (1995), Cao and Ou-Yang (2008), Banerjee, Kaniel, and Kremer (2009), Banerjee and Kremer (2010), Bhamra and Uppal (2014), Dumas, Lewis, and Osambela (2017), and Atmaz and Basak (2018) for theoretical models of stock prices.

[13]See the references cited in footnote 7. Notable exceptions are the early literature on the profitability of technical

premium explanations, we consider an information-based model to study how public news becomes incorporated into exchange rates, both contemporaneously and with a delay, thereby generating short-term predictability. The proposed framework is particularly well-suited for understanding how current and future exchange rates of commodity exporters should vary with the prices of commodity exports, which are arguably an important source of revenue for such countries. The model provides new testable predictions, which we use to guide our empirical analysis.

2.1 Environment

Consider a three-date, two-period economy with dates indexed by $t = 0, 1, 2$. We define the (log) exchange rate s_t as the date-t price in US dollars of a unit of foreign currency. At date 2, the exchange rate is given by

$$s_2 = \bar{s} + \Phi, \tag{1}$$

where $\bar{s}$ determines the initial exchange rate level, which is known at date 0. Φ is a normally distributed variable with mean 0 and volatility σ. The component Φ reflects fundamental information on the date-2 exchange rate level, such that $\Phi > 0$ ($\Phi < 0$) represents an appreciation (depreciation) of the foreign currency. The distribution of s_2, including its parameters, is common knowledge to all agents. Risk-free rates are set to 0 for convenience, i.e., we abstract from the role of UIP. Additionally, money supply plays no role in the model, so we need not specify a two-country economy.

2.2 Heterogenous beliefs

It is well established that the FX market involves different categories of market participants such as corporates, commercial banks, or asset managers.[14] Each participant has a distinct objective depending on (i) the extent to which the agent exploits available information, and (ii) whether the agent is a liquidity maker or taker.

Building on this insight, we consider three types of agents in the market. First, there is a research-intensive informed agent (hereafter the "Informed trader"), who learns about the fundamental exchange rate component Φ using public information available at date 1. Second, there is an uninformed agent ("Uninformed trader", thereafter), who offers liquidity in the market, akin to a market maker. This agent views the exchange rate as a random walk in the spirit of Meese and Rogoff (1983) and thus does not attempt to learn about Φ. Third, there is a "Noise trader" buying/selling

analysis in the currency market (see, e.g., Levich and Thomas III, 1993) and the recent work by Cespa et al. (2022), which finds that trading volume helps predict one-day-ahead exchange rate changes.

[14] Heterogeneity in agents' information is a strong feature of the FX market due to its opaque OTC nature characterized by a decentralized network and dealership structure. The rise of electronic trading and settlement in recent years has also amplified market fragmentation and asymmetric information across market participants. See Ranaldo and Somogyi (2021) for recent empirical evidence, and King, Osler, and Rime (2012) for a comprehensive review of the FX market structure.

currencies for exogenous reasons (e.g., a corporate), which reflects any non-informational trading in the FX market.

All agents are ex-ante identical, trade competitively, and have common knowledge about each other's views. Additionally, all agents have the same initial prior of $\bar{s}$ for the future exchange rate level, therefore $s_0 = \bar{s}$. Heterogeneity across agents arises due to differences in beliefs about the usefulness of public information released at date 1, which we describe below. That is, only a fraction of agents have the ability or willingness to process new information and trade on it, consistent with Cespa et al. (2022), among others.[15]

2.3 News and expectations

At date 1, the Informed trader (identified by the subscript I) learns about the fundamental component Φ from the public news

$$p \equiv \Phi + \epsilon, \tag{2}$$

where p is an unbiased, albeit noisy, signal for the fundamental component Φ, and ϵ is a normally distributed noise term with mean zero and variance σ_ϵ^2.

In the case of a commodity-exporting country like Australia, for example, the spot exchange rate s reflects the number of US dollars per Australian dollar. For a country like this, an important piece of public news is the price of its exported commodities. This price tells us much about the country's terms of trade and gives us insights into its exchange rate.[16]

The Informed trader processes the public news p and uses Bayesian updating to form new beliefs about s_2:

$$\mathbb{E}_{I,1}[s_2] = \bar{s} + \eta p \tag{3}$$

$$\mathbb{V}_{I,1}[s_2] = (1-\eta)\,\sigma^2, \tag{4}$$

where $\mathbb{E}_{i,1} \equiv \mathbb{E}_i[\cdot|\mathcal{F}_{i,1}]$ and $\mathbb{V}_{i,1} \equiv \mathbb{V}_i[\cdot|\mathcal{F}_{i,1}]$ denote the conditional expectation and variance given the agent i's information set $\mathcal{F}_{i,t}$ at time t, while η is the informativeness (or signal-to-noise ratio) of the public news p:

$$\eta = \frac{\mathbb{COV}_{I,1}[p,\Phi]}{\mathbb{V}_{I,1}[p]} = \frac{\sigma^2}{\sigma^2+\sigma_\epsilon^2} \in (0,1), \tag{5}$$

[15] See Menkhoff, Sarno, Schmeling, and Schrimpf (2016) for evidence that different groups of FX market participants differ markedly in their predictive ability.

[16] There exists a long-standing literature on the terms in trade's role in explaining exchange rates, particularly for commodity exporters. See Neary (1988) for an early discussion. See also Chen and Rogoff (2003) and Chen et al. (2010) and the references therein.

where $\mathbb{COV}_{I,1}[p,\Phi]$ denotes the covariance between the public news p and the fundamental Φ, as measured by the Informed agent at date 1. The Informed trader thus learns about the fundamental level of the exchange rate and takes a position in the market based on the new publicly available information.

The Uninformed trader (identified by the subscript U), however, either does not believe that news p contains any valuable information or is unable to process it. The expected exchange rate for the Uninformed trader at date 1 is

$$\mathbb{E}_{U,1}[s_2] = \bar{s} \neq \mathbb{E}_{I,1}[s_2] = \bar{s} + \eta p \tag{6}$$

$$\mathbb{V}_{U,1}[s_2] = \sigma^2 > \mathbb{V}_{I,1}[s_2] = (1-\eta)\sigma^2. \tag{7}$$

Both agents I and U "agree to disagree" on the relevant information set and, thus, on the expected exchange rate level. Each agent believes that no other agent holds information of any additional value to his or her information set, following classic models based on differences of opinions (e.g., Harrison and Kreps, 1978). The fact that agents have heterogeneous beliefs has long been accepted as a key feature in financial and FX markets, as sophisticated investors, analysts, and economists often publicly disagree about their forecasts.

Note that the difference in expectations across agents given by $\mathbb{E}_{I,1}[s_2] - \mathbb{E}_{U,1}[s_2] = \frac{\sigma^2}{\sigma^2+\sigma_\epsilon^2}p$ increases with the level of exchange rate uncertainty σ^2. A higher σ means the Informed trader has a stronger informational advantage of using the public news (the signal-to-noise ratio increases, see Equation 5). However, as the public news becomes pure noise, $\frac{\sigma^2}{\sigma^2+\sigma_\epsilon^2} \to 0$, this informational advantage vanishes.

2.4 Optimal demand and equilibrium exchange rate

All agents maximize CARA utility over terminal wealth, with risk-aversion set to one for notational simplicity, as in Banerjee and Kremer (2010) and Cespa et al. (2022), among others. Optimal demand for agent $i = I, U$ at date 1 is

$$x_{i,1} = \frac{\mathbb{E}_{i,1}[s_2] - s_1}{\mathbb{V}_{i,1}[s_2]}, \tag{8}$$

while the aggregate demand/supply of the noise trader, denoted by $x_{N,t}$, is normally distributed with mean 0 and volatility σ_N. In our model, the role of the noise trader's shocks is to allow exchange rates to also vary for non-fundamental reasons, as one would expect in the data.

Imposing market clearing conditions, the equilibrium exchange rate at date 1 (i.e., after the public news p is revealed) is equal to (see Internet Appendix A.1)

$$s_1 \quad = \quad \bar{\mu}_s + \bar{\sigma}_s^2 x_{N,1} \tag{9}$$

with

$$\bar{\mu}_s \quad = \quad \omega_I \mathbb{E}_{I,1}[s_2] + \omega_U \mathbb{E}_{U,1}[s_2] = \bar{s} + \underbrace{\omega_I \eta}_{<1} p \tag{10}$$

$$\bar{\sigma}_s^2 \quad = \quad \omega_I \mathbb{V}_{I,1}[s_2] + \omega_U \mathbb{V}_{U,1}[s_2] = \underbrace{(1 - \omega_I \eta)}_{<1} \sigma^2, \tag{11}$$

where ω_I and ω_U reflect the relative weights of the Informed and Uninformed traders, respectively, while $\bar{\sigma}_s^2$ is the aggregate degree of uncertainty about exchange rate s_2. Note that $\bar{\sigma}_s^2$ reflects the uncertainty perceived by the "average" agent, which differs from the true level of exchange rate uncertainty, σ^2.

From Equation (11), the equilibrium exchange rate corresponds to the average valuation across agents and, thus, only partially reflects the available public information about Φ. When there is disagreement across agents, the "average" agent puts a weight on the public news that is lower than the informativeness of p, given that $\omega_I \eta < 1$. Hence, the equilibrium exchange rate at date 1 underreacts to new public information.

2.5 Impact on contemporaneous and future exchange rate changes

Let $\Delta s_1 \equiv s_1 - s_0$ be the first-period (log) exchange rate change. From Equation (9) and $s_0 = \bar{s}$, it follows that:

$$\Delta s_1 \quad = \quad \underbrace{\left(\bar{\mu}_s + \bar{\sigma}_s^2 x_{N,1}\right)}_{s_1} - \underbrace{\bar{s}}_{s_0} \tag{12}$$

$$= \quad \omega_I \eta p + \underbrace{\bar{\sigma}_s^2 x_{N,1}}_{\text{noise}} \tag{13}$$

given that $\bar{\mu}_s = \bar{s} + \omega_I \eta p$ from Equation (10). The contemporaneous impact of the public news p on Δs_1 can be expressed as $\frac{\delta \Delta s_1}{\delta p} = \omega_I \eta > 0$, which increases with the fraction of Informed traders in the market (ω_I). It also increases with the informativeness of the news (η). The price impact of trade is thus positively related to the asymmetric use of public information across FX traders, in line with the empirical findings of Ranaldo and Somogyi (2021).

We now discuss the implication for exchange rate predictability. The second-period (log) exchange rate change, $\Delta s_2 \equiv s_2 - s_1$, is given by:

$$\Delta s_2 = \underbrace{\bar{s} + \Phi}_{s_2} - \underbrace{\left(\bar{\mu}_s + \bar{\sigma}_s^2 x_{N,1}\right)}_{s_1} \tag{14}$$

$$= \Phi - \omega_I \eta p - \bar{\sigma}_s^2 x_{N,1} \tag{15}$$

$$= [1 - \omega_I \eta]\, p - \underbrace{\epsilon - \bar{\sigma}_s^2 x_{N,1}}_{\text{noise}}, \tag{16}$$

as the date-2 exchange rate is $s_2 = \bar{s} + \Phi$ once the fundamental information is revealed. Given that information is gradually incorporated into prices, the public news p released at date 1 becomes useful for predicting the future exchange rate. In other words, $\frac{\delta \Delta s_2}{\delta p} = 1 - \omega_I \eta > 0$.[17]

Despite being as parsimonious as possible, our model generates two insightful predictions: (i) an increase in a country's commodity export prices p can only be informative about the future exchange rate if it generates a contemporaneous currency appreciation ($\frac{\delta \Delta s_1}{\delta p} > 0$). So, a "commodity currency" should be a currency that has a positive and significant exposure to current changes in commodity export prices. (ii) For such currencies, equilibrium exchange rates slowly reflect newly available information when the FX market consists of participants with heterogenous beliefs. This gradual diffusion of new public information about commodity export prices into the exchange rate generates short-term predictability ($\frac{\delta \Delta s_2}{\delta p} > 0$).

2.6 Model discussion

The model developed above contains several simplifying assumptions that help provide a good balance between tractability and realism. We focus on three periods only so that the news exploited by the more informed agent is short-lived. This is arguably a reasonable assumption for three reasons. First, it is equivalent to assuming a model with additional periods but with news that is identically and independently distributed over time, as in Llorente, Michaely, Saar, and Wang (2002). Second, Cespa et al. (2022) finds that short-term exchange rate predictability also arises in an overlapping-generations (OLG) framework. Third, introducing persistence in the news would amplify rather than weaken exchange rate predictability. Another assumption is that investors have CARA preferences and exchange rates are lognormally distributed, which allows for a closed-form solution in the model. This assumption precludes any income effect, as investors' positions are independent of wealth. It could be a fruitful avenue for future research to consider more general agent preferences while studying exchange rate predictability in a richer environment. Relaxing these assumptions is possible,

[17] The model implies the possibility of negative serial correlation ($Cov(\Delta s_1, \Delta s_2) < 0$) due to noise trader shocks, for example if $x_{N,1}$ is large. However, the positive exposure of Δs_2 to the signal p is independent of the level of noise trading, given that $\frac{\delta \Delta s_2}{\delta p} = 1 - \omega_I \eta$. Nevertheless, while the level of noise trading has no impact on the economic relation of interest, i.e., the effect of p on Δs_2, noise trading increases the variance of Δs_2 and could thus reduce the statistical significance of this exposure once we estimate the empirical counterpart of Equation (16).

but we believe it adds little to the main message of the paper.

Our theoretical analysis is expected to be particularly relevant in the context of commodity exporters' (e.g., Australia or Russia) currencies. The predictions suggest that fluctuations in commodity export prices, a public and informative source of news for these countries, should impact their contemporaneous *and* future exchange rate changes. In contrast, we should not expect to observe any of these relations for countries with negligible commodity exports (e.g., Switzerland). This is because commodity export prices should not be viewed as informative for their exchange rates. Guided by these insights, we provide a comprehensive analysis of how changes in commodity export prices predict exchange rate changes for a set of meaningful commodity currencies.

3 Identifying commodity currencies

In this section, we ask what constitutes a "commodity currency". Countries that specialize in exporting basic commodities are typically labeled as commodity countries, and their respective currencies are often regarded as commodity currencies. However, there is some degree of arbitrariness in the definition, with many studies analyzing just small sets of candidates.[18] For example, what should be the threshold to categorize a currency as a commodity currency? When a country's exports are composed of more than 20% of commodities, or rather 30%, or even 50%? Additionally, shouldn't the type of commodities matter? For example, is exporting dairy products comparable to exporting oil and gas? Clearly, there is a lot of latitude as to what defines a commodity currency. However, a reasonable classification approach should be clear, theoretically motivated, statistically sound, and subject to minimal discretion in the criteria.

In light of this, we propose a formal identification of commodity currencies based on a market-based approach. We first provide a definition based on a currency's commodity price beta. Next, we describe our data, after which we discuss the identified set of commodity currencies.

3.1 Definition

Guided by our theory, we define a commodity currency as a currency that varies positively with its country's commodity export prices (i.e., the currency appreciates when commodity export prices increase and depreciates when they fall). This is in line with Ready et al. (2017)'s equilibrium model showing that commodity currencies' exchange rates are positively correlated with commodity prices. Our definition is economically intuitive; if a country's exports are a key factor (valuable public

[18] Chen and Rogoff (2003) only consider Australia, Canada, and New Zealand; Chen et al. (2010) consider the three countries in the sample of Chen and Rogoff (2003) plus Chile and South Africa; Ferraro et al. (2015) consider Australia, Canada, Chile, Norway, and South Africa; Ready et al. (2017) examine 21 developed countries and do not formally categorize currencies as commodity currencies. However they do refer to a familiar group of commodity exporters (Australia, Canada, New Zealand, and Norway) in their discussions.

information) for the traders of this currency, commodity export prices must contemporaneously affect its exchange rate, i.e., the currency has a positive commodity price beta. A key advantage of this approach is that the beta embeds all relevant information regarding the importance and type of a country's commodity exports in one single metric.

3.2 Data

We now describe our primary data, which consists of individual foreign exchange rates and country-specific commodity export price indexes. We discuss the auxiliary data when introducing it in our analysis as well as in the Internet Appendix. The sample period runs from January 1985 to April 2020.

FX data We collect daily spot and one-month forward exchange rates relative to the US dollar from WM/Reuters via Datastream. Exchange rates are defined as units of US dollars per unit of foreign currency, so that an increase in the exchange rate indicates an appreciation of the foreign currency. Monthly data are obtained by sampling end-of-month exchange rates. Our sample includes 41 developed and emerging market currencies. Namely, the currencies of Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Chile, Colombia, Croatia, Czechia, Denmark, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Italy, Japan, Malaysia, Mexico, the Netherlands, New Zealand, Norway, Peru, the Philippines, Poland, Portugal, Russia, Singapore, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Thailand, the United Kingdom, and the euro area. The euro series starts in January 1999. After this date, euro area countries are excluded and only the euro series remains. We filter these data following Lustig et al. (2011) and Dahlquist and Hasseltoft (2020).

Our sample of currencies is similar to that of Lustig et al. (2011), but includes additional commodity exporters such as Colombia, Chile, and Peru. Our sample differs, however, from the work of Ready et al. (2017), who also study commodity currencies but focus exclusively on developed countries (this is because the equilibrium model they test requires sample countries to be financially integrated). We, on the other hand, do not need to restrict our sample, as our empirical analysis is guided by a differences of opinion theoretical framework which is applicable to all currencies.[19]

Commodity export prices We use commodity price data from the International Monetary Fund (IMF) Commodity Term of Trade database, which provides country-specific commodity price indexes for many countries. These indexes are constructed for each country as trade-weighted changes in the international market prices of up to 45 individual commodities (including agricultural raw materials, energy, food and beverages, and metals). Given our focus on commodity export prices, we use the

[19] Our framework may be especially relevant for emerging market currencies, as they are relatively less efficient, potentially making differences of opinion more prevalent (e.g., Pukthuanthong-Le and Thomas, 2008).

export-weighted indexes for each of the 41 countries in our sample. The weights are each country's individual commodity exports, scaled by its overall commodity trade. To account for variations in commodity trade over time, the weights are time-varying (specifically, lagged three-year rolling averages).[20] The index methodology ensures that changes in the price indexes reflect variations in the relevant commodity export prices for each country at each point in time. It is thus well-suited for our analysis.

3.3 Commodity price beta

To identify the set of commodity currencies, we estimate the following benchmark regression at a monthly frequency for each individual currency:

$$\Delta s_{i,t} = \alpha_i + \beta_i \Delta \mathsf{CEP}_{i,t} + \gamma_i \mathsf{DOL}_t + \varepsilon_{i,t}, \tag{17}$$

where $\Delta s_{i,t}$ denotes the log change in nominal bilateral exchange rate in US dollar per unit of currency i in month t (i.e., an increase corresponds to an appreciation of currency i), while $\Delta\mathsf{CEP}_{i,t}$ denotes the log change in the commodity export price index of country i in month t. The dollar factor $\mathsf{DOL}_{i,t}$ is computed as the average change in exchange rates against the US dollar in month t, following Verdelhan (2018).[21] The coefficient, β_i, is the currency i's sensitivity to its country's commodity export prices, which we call the "commodity price beta".[22]

Figure 2 [about here]

Results Figure 2 displays the country-level commodity price beta, β_i, and its 90% confidence interval, based on White (1980)'s standard errors. To be categorized as a commodity currency, we require that a currency's commodity price beta β_i is positive and statistically significant at the 10% level. Our procedure identifies nine countries' currencies that display a positive and statistically significant commodity price beta: Australia, Brazil, Canada, Mexico, New Zealand, Norway, Peru, Russia, and South Africa. This categorization is in line with the priors established by the existing literature, where currencies of Australia, Canada, New Zealand, and Norway are typically considered commodity currencies. Compared to previous works, however, we consider a larger sample of currencies, including several emerging market currencies. We are thus able to identify additional, less-studied commodity currencies, such as the Brazilian real, the Mexico peso, the Peruvian sol, and the Russian ruble.

Although our benchmark analysis focuses on the core commodity currencies, our procedure identifies a spectrum of betas across the cross-section of countries. For example, we are able

[20] See Gruss and Kebhaj (2019) for additional details on the data.

[21] The relation between exchange rates and commodity prices may also be correlated with global financial conditions. Figure A.3 in the Internet Appendix considers a specification that controls for the dollar factor and the US stock market return. The inclusion of this additional risk factor does not materially affect the results. As there is no clear consensus on the null model for bilateral exchange rate changes, we choose the most parsimonious specification.

[22] Table A.1 in the Internet Appendix reports the regression results by country.

to identify commodity importers like Japan with a negative and statistically significant commodity price beta (we exploit the full cross-section of commodity betas in some of our tests in the next section). It may also be interesting to note that the UK, a country that is never classified as a commodity currency in the literature, has a positive, albeit statistically insignificant, commodity price beta. This makes economic sense given that the UK historically exported a non-trivial amount of oil, while natural resource firms make up a substantial share of its public stock market (we explore how commodity beta relates to economic fundamentals below).

Role of the dollar factor The inclusion of the dollar factor in regression (17) is fundamental to adequately identifying commodity currencies. To see this, we estimate univariate regressions without controlling for the average dollar effect, which produces nonsensical results. As illustrated by Figure A.2 in the Internet Appendix, one would conclude that 26 out of the 41 currencies should be categorized as commodity currencies. It would include, for example, the Singapore dollar and the Swedish krona, two currencies of countries producing a tiny share of commodity goods. The positive commodity price beta is largely due to their commodity export prices being denominated in the US dollar. Thus a dollar depreciation often mechanically leads to an increase in commodity prices and an appreciation of other currencies, such as the Swedish krona, vis-a-vis the US dollar. Controlling for the dollar effect is thus of paramount importance for the correct identification of commodity currencies.

Commodity beta vs. economic fundamentals We now verify that our categorization lines up well with the importance of the commodity sector in a country's exports, GDP, financial markets, and fiscal revenue.[23] Panel A of Figure 3 displays each country's commodity price beta against the average primary commodity share of its total exports. On average, the greater the primary commodity share of a country's exports, the more sensitive a country's currency is to its commodity export prices. A cross-country regression of the commodity price beta on the average commodity share of exports yields a positive and statistically significant coefficient, with an R^2 of 49%. Similar relationships are obtained using total commodity rents as a fraction of GDP (Panel B), a commodity sector share in a country's stock market (Panel C), and the share of commodity revenue of total government revenue (Panel D). Table 1 shows that, based on these measures, a one standard deviation difference

[23] We source additional data for the analysis in this sub-section. Data on primary countries' commodity share of total exports are from the United Nations (UN) Comtrade Database (the data are annual). Data on total natural resources rents (as a percentage of GDP) are from the World Bank (these data are available for the years 1990, 2000, 2013, and annually thereafter). We source data on the market capitalization of listed firms in each country from Datastream. We calculate the total commodity firm market capitalization for each country by aggregating across Datastream industrial groupings (INDM). We consider firms to be commodity firms if their industrial sector is, for example, Aluminum, Coal, Forestry, General Mining, Oil, Crude Producers, Sugar, or Tobacco. The market capitalization data are monthly, and the time period matches that of our FX data. There is no centralized source for international data on the fraction of commodity-linked fiscal revenue. Moreover, such data are typically not reported for countries that are not large commodity producers. Therefore, the revenue share data are estimates stemming from a variety of sources, including the Norwegian Petroleum Directorate, the Organization for Economic Co-operation and Development (OECD), and the United Kingdom Office for Budget Responsibility, among others.

between countries explains 0.66 to 0.77 standard deviation difference in their commodity price betas. Notice that each of these measures provides a complementary explanation of why a currency ends up being categorized as a commodity currency. For example, Mexico's commodities (mostly oil) represent a modest 18% of its exports, but the average share in the government revenue amounts to a staggering 42%. We observe the opposite for Australia, which has export and revenue shares of 67% and 2%, respectively.

This analysis uncovers two main findings. First, considering a single metric to capture the importance of the commodity sector in a country is typically insufficient for the adequate identification of commodity currencies. For example, Chile and Colombia are often referred to as commodity countries, as they are major commodity producers of copper and coffee. The commodity share of their exports is sizeable (52% and 68%, respectively), but their commodity betas are not statistically significant.[24] In contrast, we find that Canada displays a significant commodity beta, although commodities reflect a more modest fraction of its total exports (36%). Second, the four metrics combined explain only 68% of the cross-sectional variation in the commodity beta (see Table 1), which underscores the additional and unique information captured by our measure.

Table 1 and Figure 3 [about here]

In sum, we propose a new market-based approach to determine the currencies that should reasonably be classified as commodity currencies. We identify nine currencies that are statistically exposed to their country's commodity export price fluctuations. We hereafter focus on this set of commodity currencies and provide new evidence for exchange rate predictability.

4 Unconditional exchange rate predictability

In this section, we show that, in our set of commodity currencies, changes in country-level commodity export prices predict these countries' exchange rates. We then confirm that this result holds after controlling for variation in FX and global financial conditions. We first present the empirical approach, describe the control variables, and then discuss the results.

[24] The case of Saudi Arabia provides an extreme example. Although oil represents 90% of government revenue and 87% of exports, the Saudi riyal is almost uncorrelated to oil prices, as illustrated by Figure A.1 in the Internet Appendix. So merely being the currency of a large commodity exporter is not enough to qualify as a commodity currency.

4.1 Baseline specification

We assess exchange rate predictability with commodity export prices by running panel regressions at the monthly frequency based on the following specification:

$$\frac{1}{k}\Delta s_{i,t+k} \quad = \quad \alpha_{i,k} + \beta_k \Delta \mathsf{CEP}_{i,t} + \gamma_k \mathsf{IRD}_{i,t} + \boldsymbol{\theta}_k \mathbf{x}_t + u_{i,t+k}, \tag{18}$$

where $s_{i,t}$ is the log of the nominal exchange rate between currency i and the US dollar in month t, $\Delta s_{i,t+k} = s_{i,t+k} - s_{i,t}$ is the exchange rate change between months t and $t+k$, $\Delta\mathsf{CEP}_{i,t}$ is the (standardized) log change of the country i's commodity export price index observed in month t, $\mathsf{IRD}_{i,t}$ is the interest rate differential between currency i and the US dollar in month t, $\mathbf{x}_t$ comprises a set of control variables observed in month t, and κ is the forecast horizon. We consider exchange rate predictability for up to 12-month horizons, but focus our analysis on the one-month-ahead predictability to exploit non-overlapping data. This specification can be viewed as the empirical counterpart of Equation (16) in the model. We include currency fixed effects, denoted by $\alpha_{i,k}$, to control for time-invariant differences across exchange rate movements. Standard errors are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with optimally-selected bandwidth. Observations are monthly and the sample period ranges between January 1985 and April 2020.

We control for a number of exchange rate predictors. First, we include each country's interest rate differential ($\mathsf{IRD}_{i,t}$) to account for the UIP (see, e.g., Fama, 1984; Bansal and Dahlquist, 2000), which we derive from one-month forward exchange rates following Verdelhan (2018). Second, we account for the dollar factor following Lustig et al. (2011) and Verdelhan (2018). Third, we consider fluctuations in FX volatility, constructed as in Bakshi and Panayotov (2013).[25] This choice builds on the empirical evidence that uncertainty in the FX market helps explain the cross-section of currency returns (Menkhoff et al., 2012; Karnaukh et al., 2015) and predict future exchange rate changes (Bakshi and Panayotov, 2013). Fourth, we control for changes in funding liquidity as measured by the TED spread (the interest rate difference between the 3-month interbank deposits, LIBOR, and 3-month US Treasury bills). An elevated TED spread is associated with tighter funding conditions in the inter-bank market and, typically, with less liquidity (Mancini et al., 2013). Fifth, we control for aggregate market uncertainty, measured by the CBOE equity-option implied volatility index (VIX). Both the TED and the VIX have been shown to predict exchange rate changes (see, e.g., Brunnermeier et al., 2008). Finally, we control for a recession indicator based on the National Bureau of Economic Research (NBER) business-cycle dates to address the potential concern that variations in commodity prices reflect changes in global economic conditions (for example, the declines in commodity prices observed during the Great Recession and the COVID-19 crisis).[26]

[25] For each currency, we construct monthly volatility as the square root of the sum of squares of daily log currency changes against the US dollar over a month. We then average each volatility across our sample of 41 currencies to obtain the aggregate FX volatility measure.

[26] The data for the VIX, TED, and NBER recession indicators are from the Federal Reserve Bank of St. Louis. The

4.2 Unconditional results

Table 2 reports the results for the one-month-ahead exchange rate predictability. We find that variation in commodity export prices positively predicts future exchange rate changes for our set of commodity currencies. In a univariate specification, reported in Column (1), we observe an estimate $\beta_1 = 0.374$ that is statistically significant at the 1% level. This result is robust to including various control predictors such as the one-month interest rate differential (Column 2), the dollar factor (Column 3), changes in FX volatility (Column 4), changes in the TED spread (Column 5), changes in the VIX (Column 6), and the NBER recession indicator (Column 7). It is worth noting that commodity export prices appear to be the only statistically significant source of information for predicting commodity countries' exchange rates beyond the interest rate differential.[27]

Accounting for all controls leaves the coefficient estimate almost unchanged ($\beta_1 = 0.367$). This estimate implies that one-standard-deviation increase (decrease) in a country's commodity prices predicts a future appreciation (depreciation) for that country's currency of about $12\times 0.367 = 4.4\%$. This effect is economically meaningful. Given the small set of commodity currencies, it is important to verify that our results are not driven by one particular currency or time period. As reported in the Internet Appendix, we find that the coefficient estimates remain stable when dropping one currency at a time (see Figure A.4), and hold both in the first (1985-2003) and second (2003-2020) halves of our sample (see Tables A.2 and A.3).

Table 2 [about here]

Figure 4 reports the least-squares estimates of β_k, for horizon k up to 12 months. This analysis illustrates the dynamic exchange rate predictive ability of commodity export prices, accounting for the full set of controls. The predictability is positive and statistically significant in the short term, but dissipates gradually, becoming negligible after four months. There are two direct implications of this result. First, a long position in a commodity currency could be profitable for several months following an increase in the prices of the commodities they export. The exchange rate predictability for commodity currencies can thus be practically applicable. Second, the short persistence in predictability is in line with the underreaction to new information that we expect to disappear rather quickly: new public information (i.e., a change in commodity export prices) is gradually reflected in exchange rates. This channel contrasts the return predictability arising from a change in risk premium, which is persistent and typically increases with the forecast horizon (see, e.g., Lettau and Ludvigson, 2001). The pattern of Figure 4 is thus consistent with our information-based explanation.

Figure 4 [about here]

sample spans the 1986.01–2020.04 period when we include the TED spread or the VIX as controls. Although we refer to the VIX for convenience, we use the VXO (the old version of the VIX) to benefit from a longer sample period.

[27] The interest rate differential has a positive sign (i.e., a higher interest rate relates to a future foreign currency appreciation), contradicting the forward unbiasedness hypothesis, but in line with the existing literature.

Overall, our findings indicate that fluctuations in commodity export prices help predict exchange rates over the next month. This result is in contrast to the general wisdom that exchange rates are well approximated by a naive random walk model (e.g., Meese and Rogoff, 1983), especially at short horizons.[28] Additionally, such predictability is strong, both statistically and economically, and is not explained by alternative exchange rate predictors. With the inclusion of various key control variables, our empirical approach ensures that exchange rate changes and commodity export price fluctuations are uncorrelated with changes in financial market conditions, thereby addressing potential endogeneity concerns. Specifically, the coefficient β_k is unlikely to be capturing the potential effects stemming from investors who jointly trade commodities and currencies, and who may be simultaneously adjusting their positions in both assets as global financial conditions vary. We now discuss various additional robustness tests.

4.2.1 Alternative base currencies

We first ensure that our inferences are robust to the choice of the base currency. In our benchmark analysis, we computed exchange rates from the US investor's perspective (i.e., the US dollar as the base currency). However, this choice may, for example, introduce US dollar-specific effects into the analysis because most commodities are traded in US dollar. This can make it difficult to isolate the effect of commodity prices on commodity currencies. To address this potential concern, we now consider exchange rates with respect to the Swiss franc, the euro, and the Japanese yen. Columns (1), (2), and (3) of Table 3 report the results. When using these alternative base currencies, we continue to find that commodity export prices predict future exchange rates of commodity currencies. Not only are the coefficient β_k estimates of a similar magnitude to the baseline case, but the explanatory power (R^2) even increases for such alternative base currencies. This analysis alleviates the concern that the exchange rate predictability that we document for commodity currencies is purely driven by a US dollar effect.

Table 3 [about here]

4.2.2 Counterfactual exercise

We now consider a counterfactual exercise, which aims to verify the mechanism underpinning the exchange rate predictability for commodity currencies. We expect that an increase (decrease) in commodity export prices is good (bad) news for the economy of a commodity country. Trading on expectations, FX participants would react to the news by increasing (decreasing) their demand for the country's currency. Therefore, following the change in commodity export prices, a commodity country's currency would appreciate (depreciate) with this increase in buying (selling) pressure.

[28]Section 7 verifies that the predictability also holds in an out-of-sample analysis.

However, there is no obvious economic reason why this mechanism should affect the currencies of countries with limited commodity exports.

Guided by this intuition, we repeat our analysis (specification 18) on a sub-sample of non-commodity currencies and test the hypothesis of no predictability. For this exercise, we determine the set of non-commodity currencies using two distinct approaches. First, we simply exclude from our full sample the currencies that exhibit a statistically significant and positive commodity price beta. We report the result for this sub-sample in Column (4) of Table 3. Second, we only include the currencies of commodity importers in Column (5), i.e., those exhibiting a statistically significant negative commodity beta using the methodology described in Section 3. In both cases, the results affirm our hypothesis that there is no exchange rate predictability associated with commodity export prices in the non-commodity currency group. Irrespective of how we identify the non-commodity currencies, the estimate of the coefficient of interest, β_k, is consistently at least half the size of the baseline case (Column 4 of Table 2) and is never statistically significant. This finding also mitigates the concern that omitted variables that are correlated with commodity prices and exchange rates (e.g., those reflecting global economic conditions) are skewing our results.

Departing from traditional sample splits, we consider an interaction between $\Delta\text{CEP}_{i,t}$ and β_i, where β_i is the commodity price beta estimated using the contemporaneous regression (17). This augmented specification allows us to explore a more continuous analysis of the role of commodity currencies. Specifically, we expect this interaction term to be positive and statistically significant, thereby suggesting that changes in commodity export prices matter more for the exchange rate predictability of currencies with relatively higher commodity price beta. This hypothesis is empirically verified in Column (6) of Table 3.

Overall, we can conclude from these analyses that commodity export prices contain useful information for predicting exchange rates. In line with the model's intuition, this result holds only in the short-term and exclusively for commodity currencies, i.e., those having a positive and significant commodity price beta.

5 Conditional exchange rate predictability

In this section, we explore how exchange rate predictability with commodity export prices varies across FX market conditions. We first use our stylized model developed in Section 2 to provide guidance on conditional exchange rate predictability. The theory predicts that the impact of public news on the future exchange rate increases with the level of FX uncertainty. To test this prediction, we use a smooth-transition model to estimate regimes and find strong evidence that exchange rate predictability is concentrated in times of elevated FX uncertainty.

5.1 Theoretical prediction

Our simple model provides valuable insights into conditional exchange rate predictability. Recall from Equation (16) that the second period's exchange rate change Δs_2 varies positively with the public news p, the change in commodity export prices in our case, as $\frac{\delta \Delta s_2}{\delta p} > 0$. Taking the perspective of the econometrician, we can express the exposure of Δs_2 to the change in commodity export prices p as follows (see Internet Appendix A.2):

$$\beta_1 \quad = \quad \frac{\mathbb{COV}\left[\Delta s_2, p\right]}{\mathbb{V}\left[p\right]} = \frac{\sigma^2}{\sigma^2 + \sigma_\epsilon^2}(1 - \omega_I), \tag{19}$$

where β_1 is akin to the slope coefficient of our one-month-ahead predictive regression (18).

Our model suggests that the impact of commodity export prices on the next-period exchange rate increases with the level of FX uncertainty σ, as $\frac{\partial \beta_1}{\partial \sigma} > 0$. To better understand this prediction, remember that both informed and uninformed traders are risk-averse, meaning they are less inclined to trade with each other when the exchange rate is more volatile. The decrease in trading activity implies that the exchange rate incorporates less news during the initial period, enhancing the predictability of the future exchange rate. Conversely, lower volatility would be associated with higher trade volume, causing a more rapid reflection of news in the exchange rate, and subsequently reducing predictability (i.e., $\beta_1 \to 0$ when $\sigma \to 0$). In sum, exchange rate predictability with commodity export prices should strengthen in times of higher FX uncertainty. We now propose an empirical approach to test this prediction.

5.2 Identifying FX uncertainty regimes

We assess the conditional impact of commodity export prices on future exchange rate changes in a regime-switching environment. Specifically, we use Jordà (2005)'s local projection method in a system that admits a smooth transition across two regimes, namely a high (H) and a low (L) FX uncertainty regime.[29] Our baseline measure of FX uncertainty is hereafter the average of the realized exchange rate volatility computed with daily returns across each of the nine commodity currencies.[30]

The transition function F_t from the low FX volatility regime (L) to the high FX volatility regime (H) is given by:

$$F_t = \frac{\exp\left(\theta \frac{\sigma_t - c}{\mathsf{std}(\sigma)}\right)}{1 + \exp\left(\theta \frac{\sigma_t - c}{\mathsf{std}(\sigma)}\right)}, \tag{20}$$

[29] This approach is similar to the smooth-transition-local projection model used in Tenreyro and Thwaites (2016) and Ramey and Zubairy (2018), which analyze monetary and fiscal policies. Studies using alternative forms of smooth transition models to study exchange rate or carry trade predictability include Taylor and Peel (2000), Taylor, Peel, and Sarno (2001), Kilian and Taylor (2003), and Christiansen, Ranaldo, and Söderlind (2011).

[30] We also use the dispersion in professional forecasts as an alternative, forward-looking measure of FX uncertainty. We discuss the results in Section 5.7.

where the state variable σ_t corresponds to the level of FX volatility in month t, and std(σ) is the standard deviation of σ_t. The parameter θ determines the speed of transition across regimes, while the parameter c fixes the threshold between the two regimes.

Following the literature (e.g., Granger and Terasvirta, 1993), we fix the parameters of the transition function (20). We calibrate the speed of transition across regimes, θ, and the threshold, c, to obtain a proper interpretation of the FX regimes. In particular, we set $\theta = 3$ as in Tenreyro and Thwaites (2016) and determine c such that the probability that $F_t > 0.8$ is 0.2, thus ensuring that the FX market is in the high FX volatility regime only 20% of the time.

Figure 5 displays the probability of being in the high FX volatility regime, which appears to be elevated in multiple instances of market stress. This includes during the Russian debt crisis in 1998, the severe oil price crash and carry trade reversal in 2008, the European debt crisis in 2012, when the Fed entered a monetary tightening phase (and increased US interest rates 9 consecutive times) between 2015 and 2018, or at the peak of the Covid-19 crisis (in around March 2020) when commodity currencies and commodity prices had crashed. So, our transition function F_t adequately captures most of the key events with a large impact on the commodity and FX markets.

Figure 5 [about here]

5.3 Non-linear specification

In a panel setting, we estimate how a country i's exchange rate responds to changes in its commodity export prices, contingent on whether it is in a high (H) or low (L) FX volatility regime. The specification extends the unconditional case (18) to capture conditional exchange rate predictability as follows:

$$\begin{aligned}\frac{1}{k}\Delta s_{i,t+k} &= \alpha_{i,k} + \beta_{k,H}\underbrace{\Delta\text{CEP}_{i,t} \times F_{t-1}}_{\text{High FX volatility}} \\ &\quad +\beta_{k,L}\underbrace{\Delta\text{CEP}_{i,t} \times (1 - F_{t-1})}_{\text{Low FX volatility}} + \boldsymbol{\gamma}\mathbf{x}_t + v_{i,t+k}, \qquad (21)\end{aligned}$$

where F_{t-1} is the smooth transition function reflecting the probability of being in a, FX volatility regime in month $t-1$.

The coefficients $\beta_{k,H}$ and $\beta_{k,L}$ capture whether the changes in a country's export commodity prices at time t predict the k-month-ahead exchange rate, conditional on FX volatility being high and low, respectively. It is important to note that we only use the information available at time $t-1$ to categorize the level of FX volatility, i.e., the transition function F_{t-1} is lagged by one month. This adjustment serves two purposes. First, the lag structure avoids inadvertently capturing potential predictive effects of changes in FX volatility. Second, it alleviates the concern of reverse causality, as exchange rate changes could also affect FX volatility.

5.4 Results of the conditional case

The results reported in Table 4 indicate that exchange rate predictability varies substantially with the level of FX volatility. In times of high FX volatility, the estimate of $\beta_{k,H}$ is positive and statistically significant, while there is no apparent predictability when the FX market is less volatile (i.e., $\beta_{k,L}$ is much smaller in magnitude and not statistically significant from zero). Accounting for all controls, we obtain $\beta_{k,H} = 0.515$ (in contrast to $\beta_{k,L} = 0.124$) at the one-month horizon ($k = 1$), increasing from 0.367 in the unconditional case. That is, a one-standard-deviation increase in a country's commodity export prices predicts a one-month-ahead currency appreciation of about 6.2% per annum when FX volatility is high.

Table 4 [about here]

Figure 6 reports the estimates of $\beta_{k,H}$ and $\beta_{k,L}$ for horizons k up to 12 months, accounting for the full set of control predictors. We observe that, in the most volatile periods, predictability remains positive and statistically significant for up to three months ahead, while we do not see significant predictability when FX volatility is low.

Figure 6 [about here]

In sum, these results provide evidence that exchange rate predictability for commodity currencies is concentrated in times of high FX volatility, consistent with our model prediction. We now consider various tests to assess the robustness of this finding.

5.5 Stripping out global market conditions

One may be concerned that time variation in FX volatility reflects changes in the global market environment and, thus, conditions that are not specifically tied to the FX market. For example, the FX market tends to display more volatile fluctuations and become less liquid when investors' fears (VIX) increase, such as during periods of financial turmoil (e.g., Menkhoff et al., 2012). Currencies also become more volatile when FX dealers face tighter funding constraints and money-market premiums increase (e.g., Brunnermeier et al., 2008; Ranaldo and Söderlind, 2010), as reflected by a higher TED spread. Over our sample period, the correlation between FX volatility and the VIX and TED spread is 0.56 and 0.40, respectively. Hence, FX volatility indeed varies positively with these measures of global market conditions.

To ensure that FX volatility reflects primarily FX market stress, we now strip out the effect of such global risk conditions. Specifically, we orthogonalize the monthly measure of FX volatility to the VIX or the TED spread and re-estimate the transition functions. Columns (1) and (2) of Table 5 report the results when we eliminate the part of FX volatility that covaries with aggregate uncertainty (VIX) and funding liquidity conditions (TED spread), respectively. Column (3) uses FX volatility orthogonalized to both the VIX and the TED spread. In all three cases, the predictability of exchange rates with commodity export prices remains statistically significant and concentrated in times of elevated FX volatility.[31] This analysis confirms the primary role of FX market conditions in driving the asymmetric predictability of exchange rates with commodity export prices.

Table 5 [about here]

5.6 Disentangling FX volatility from liquidity

Next, we verify that FX volatility is not capturing the level of (il)liquidity in the FX market, as a substantial body of research suggests that volatility and illiquidity are highly interlinked. The lack of liquidity, as reflected in bid-ask spreads, can be positively affected by volatility due to higher adverse selection and inventory risk (e.g., Stoll, 1978). Empirically, Karnaukh et al. (2015) find that the liquidity of currencies tends to evaporate when their volatility increases, while Ranaldo and de Magistris (2022) show that higher volatility and illiquidity arise jointly when there is more disagreement among FX traders.

Given the potentially tight link between periods of volatile exchange rates and market dry-outs, we now orthogonalize FX volatility with respect to FX illiquidity to verify the validity of the baseline results. We use the systematic FX illiquidity index proposed by Karnaukh et al. (2015), which is constructed as the average level of illiquidity of 33 currency pairs, mostly based on bid-ask spreads.[32] The results presented in Column (4) of Table 5 show that exchange rate predictability remains concentrated when FX volatility is high, even after stripping out the effect of currency liquidity dry-outs.

5.7 Alternative measure of FX uncertainty

Finally, we exploit professional forecasts to compute an alternative exchange rate uncertainty measure. We construct the level of FX uncertainty as the average of the forecast dispersion (measured as the standard deviation of individual forecasts) across our set of commodity currencies. The monthly exchange rate forecast data are from the Foreign Exchange Consensus Forecasts (the series is only available from January 2003). This approach complements our baseline analysis in two ways: first,

[31] Results are qualitatively and quantitatively similar when orthogonalizing FX volatility to aggregate stock market excess returns.

[32] The index, which is available from January 1991, largely reflects the illiquidity of developed and emerging currencies. The emerging countries present in the index are Hungary, India, Mexico, Poland, South Africa, and Turkey. We use an updated series kindly shared by Angelo Ranaldo (the series is available for the period January 1991 to July 2019.

it is a forward-looking measure of uncertainty, while realized volatility uses historical data. Second, it is not based on currency returns. Column (5) of Table 5 reports results when we condition exchange rate predictability on the level of FX forecast dispersion. The results confirm that the predictive ability of commodity price changes is concentrated when there is higher uncertainty in the FX market.

6 Implications for the carry trade

We now explore the implications of our predictability results for the carry trade, one of the most popular zero-cost and dollar-neutral investment strategies in the FX market. A fundamental aspect of the carry trade strategy, which consists of borrowing in low-yield currencies and investing in high-yield currencies, is the composition of the currency portfolios. The common perception is that high-yield currencies often belong to major commodity producers like Australia and New Zealand, whereas low-yield currencies are typically associated with commodity-importing countries like Japan and Switzerland.[33] While the prior literature takes this classification for granted, the reality is more nuanced in a large cross-section of currencies. To see this, Figure 7 plots each country's estimated commodity price beta against its average monthly forward premium. While we observe, in line with the extant literature, a positive and statistically significant relationship between the two measures, the explained variation (R^2) is only 30%. That is, high-yield currencies are not necessarily obvious commodity currencies (e.g., consider Portugal before the euro), while some more established commodity currencies may not consistently offer high yields (e.g., Canada).

Since the carry trade involves investing in both commodity and non-commodity currencies, it's uncertain how much of the predictability we find for commodity currencies extends to the carry trade. If a substantial portion of the long leg of the carry trade portfolio consists of commodity currencies, commodity export prices might significantly influence the carry trade's future performance. Indeed, the existing literature finds that carry trade returns seem to be predictable with commodity price changes (Bakshi and Panayotov, 2013). In this section, we aim to reexamine these results with the aim of uncovering the source of this predictability.

Figure 7 [about here]

6.1 Methodology and statistics

For this analysis, we construct an aggregate index of commodity export prices. This series is constructed as the equal-weighted average of the commodity export price returns corresponding to each

[33] This premise is at the core of the model in Ready et al. (2017), predicting that currencies that deliver high (low) yields belong to countries that specialize in exporting (importing) commodities in equilibrium.

of the nine commodity currencies identified in Section 3. We refer to this new index as the average commodity export price and denote its return by $\Delta\overline{\text{CEP}}$.

We explore the carry trade strategy from the perspective of the US investor. Following Lustig et al. (2011) and Menkhoff et al. (2012), at the end of each month t, we allocate all the available currencies in our sample into five portfolios based on their forward discounts $f_t - s_t$, where f_t and s_t denote the logarithm of the spot and one-month forward exchange rates, respectively. Exchange rates are in units of foreign currency per US dollar.[34] Portfolio 1 (P1) contains currencies with the smallest forward discounts (or lowest interest rates), and Portfolio 5 (P5) contains currencies with the largest forward discounts (or highest interest rates). We rebalance portfolios at the end of each month. We compute portfolio returns by taking the (equally weighted) average of the returns of each currency in the portfolio. The return of the carry trade is the return difference between the high-yield portfolio (P5) and the low-yield portfolio (P1).

Panel A of Table 6 reports the descriptive statistics for the currency portfolios and for the carry trade (P5-P1). Consistent with the literature, the carry trade is a very profitable strategy, with a long-short return of 7.60% per annum and a Sharpe ratio of 0.85 over the 1985-2020 period. Panel B presents the fraction of time each of the commodity currencies is a member of a specific portfolio. It is interesting to notice that Canada is never part of the investment portfolio (P5) and sometimes even becomes a funding currency (P1). This finding suggests that the positive link we observe between the Canadian dollar and the oil price becomes irrelevant for explaining and predicting the performance of the carry trade. Additionally, once we consider a large sample of currencies, the Australian dollar is regularly in P3 and P4, given its relatively high yield, but is rarely an investment currency. In contrast, the Brazilian real, the Russian ruble, and the South African rand are primary constituents of the investment portfolio. This portfolio composition analysis reveals that our carry trade strategy differs substantially from that in Bakshi and Panayotov (2013), which focuses on G10 currencies.[35] Our investment portfolio only contains 6.4% of G10 currencies, on average, which indicates little overlap between the two carry trade strategies.

Table 6 [about here]

6.2 Carry trade predictability

Panel A of Table 7 reports the results for the one-month-ahead carry trade return predictive regressions. We find that changes in aggregate commodity export prices positively predict future carry

[34] Sorting on forward discounts is equivalent to sorting on interest rate differentials under the covered interest rate parity: $i_t^* - i_t - \Delta s_{t+1} \cong f_t - s_{t+1}$, where i_t^* and i_t denote the 1-month foreign and US dollar nominal risk-free rates, respectively.

[35] Over the 1985-2011 period, these authors find that the Australian dollar, the New Zealand dollar, and the British pound were among the highest-yielding currencies. In contrast, once we consider a larger cross section, the carry trade rarely requires buying these currencies, as various non-G10 currencies have higher yields and thus constitute the largest part of the investment portfolio. So one should view both works as complementary.

trade returns in a univariate specification, as reported in Column (1). The results remain similar when we control for the dollar factor (Column 2), changes in aggregate FX volatility (Column 3), changes in the TED spread (Column 4), changes in the VIX (Column 5), and the NBER recession indicator (Column 6). When including all controls, we observe a predictive regression estimate that is statistically significant at the 1% level.

The predictability of carry trade returns is also economically meaningful: the coefficient estimate implies that a one-standard-deviation increase in aggregate commodity export prices leads to a carry trade return of 0.105% in the following month (or about 1.3% on an annualized basis). It is worth noting that, among the control predictors, only FX volatility has a statistically significant predictive ability for the carry trade returns.[36]

Table 7 [about here]

Having established that commodity export prices help forecast carry trade returns, we now examine the source of the predictability. Specifically, we verify that the predictive power of commodity export prices for carry trade returns stems from the predictability of individual commodity countries' exchange rates, and not, for example, due to the influence of omitted risk factors. We provide several direct tests of this intuition. First, we compare the predictability of each of the five carry trade portfolios against the frequency of having commodity currencies in each of the portfolios. Figure 8 shows that the estimated slope coefficient from regressing the one-month-ahead portfolio return on $\Delta\overline{\text{CEP}}$ strongly and significantly increases with the average commodity currency membership.[37] That is, the more commodity currencies a portfolio contains, the better we can predict the future performance of that portfolio.

Figure 8 [about here]

Second, we construct a counterfactual carry trade portfolio with a sub-sample of currencies that excludes our set of commodity currencies. We then re-estimate the predictive regression on the new carry trade returns. The results, reported in Column (1) of Table 8, show that commodity export prices have no predictive ability on carry trade returns if we exclude commodity currencies from the carry trade portfolio. Similarly, we reevaluate the predictive regression using the return on P1 (low-interest currencies) as the dependent variable. The regression results reported in Column (2) confirm that commodity export prices have no predictability for this leg of the carry trade ($\beta_1 =$

[36] The estimated coefficient is positive, in line with the findings of Menkhoff et al. (2012) that FX volatility changes are a priced risk factor in the cross-section of interest rate sorted currency portfolios.

[37] Table A.4 in the Internet Appendix presents the regression results for each portfolio. The results indicate that the predictive power, as measured by the R^2, also increases monotonically with a portfolio's average commodity currency membership.

0.022 and is statistically insignificant). In contrast, commodity export prices strongly predict the returns on P5 (β_1 = 0.125 and is statistically significant at the 1% level), as P5 regularly contains commodity currencies (see Panel B of Table 6).

Table 8 [about here]

It is important to stress what these findings do and do not imply. Our results demonstrate that fluctuations in commodity export prices help predict the profitability of the carry trade, but only when the carry trade strategy invests in commodity currencies. Once we exclude commodity currencies, this predictability vanishes. However, this does not mean the carry trade becomes unprofitable.

To verify this, we examine the performance of our counterfactual carry trade strategy that excludes the nine commodity currencies (we report the performance statistics in Table A.5 in the Internet Appendix). We observe that the long-short return is 6.3% per annum while the Sharpe ratio is 0.75, both of which are similar to the unrestricted carry trade. Therefore, commodity currencies play a critical role in the carry trade in terms of predictability, but only a minor role in explaining the average carry trade performance.[38]

We then provide evidence that the return predictability with commodity prices arises largely from the consideration of country-specific export prices. To show this, we consider other global commodity price indexes as alternative predictors. We use the percentage changes of CRB Raw Industrials subindex of the CRB commodity index, as it has been used in the exchange rate literature (e.g., Ready et al., 2017) and is shown to be relevant for carry trade predictability (e.g., Bakshi and Panayotov, 2013), the generic Goldman Sachs commodity index (a well-known benchmark in practice), and the West Texas Intermediate Crude Oil price (e.g., Ferraro et al., 2015).[39] Columns (4) and (5) of Table 8 report the results when we control for the CRB and Goldman Sachs indexes, respectively, while Column (6) controls for oil price changes. Our aggregate commodity export price index continues to have predictive power, while the other global commodity measures do not.[40] This analysis confirms the relevance of using country-specific commodity export prices for forecasting the exchange rates of commodity currencies and, in turn, the carry trade's performance. Finally, we show that the conditional predictability results uncovered for individual commodity currencies echo those for the carry trade. Columns (7) and (8) of Table 8 suggest that the return predictability is concentrated in times of elevated FX volatility and FX forecast dispersion, respectively.

Overall, our findings enhance our understanding of the relationship between individual currency returns and the carry trade. Conventionally, it is believed that the carry trade exposes investors to

[38]Consistent with this view, Bekaert and Panayotov (2020) propose an alternative carry trade strategy that almost never includes the Australian dollar and the Norwegian Krone, two key commodity currencies.

[39]The two indexes are sourced from Datastream and Bloomberg, respectively, and the oil price from FRED.

[40]Table A.6 shows that, on their own, each of these commodity indices helps predict carry trade performance. However, their explanatory power (R^2) is on average about half of that obtained with $\Delta\overline{\text{CEP}}$. In addition, when combining all predictors, only $\Delta\overline{\text{CEP}}$ becomes significant, both economically and statistically.

global sources of risk, leading to higher returns for currencies more exposed to this factor (Verdelhan, 2018). Complementing this view, we show that individual exchange rates are subject to currency-specific fluctuations stemming from the commodity market, influencing the time variation in carry trade performance. This occurs because the carry trade's long portfolio is concentrated in a small set of commodity currencies, the shocks of which cannot be truly diversified away. Consequently, our results suggest a bottom-up channel, whereby news and frictions impacting individual currencies drive carry trade returns. This complements the more traditional top-down perspective, which posits that carry trade shocks affect individual currency returns.

7 Out-of-sample analysis

This section investigates the out-of-sample predictive ability of commodity export prices for exchange rate and carry trade returns. Since the influential work of Meese and Rogoff (1983), numerous studies find that even economically meaningful variables often fail to yield accurate out-of-sample exchange rate forecasts. Considering the prevailing belief that exchange rates, particularly over a short horizon (e.g., Mark, 1995b), are inherently unpredictable, the random walk (RW) model has become the default benchmark model for assessing the predictive performance of exchange rates. However, we demonstrate that leveraging the information embedded in commodity export prices outperforms the RW model out-of-sample, particularly for the less-traded and emerging currencies. We also show strong implications for predicting the performance of the carry trade. We first describe the forecast accuracy's out-of-sample tests and then present the empirical results.

7.1 Tests of forecast accuracy

For each month t and currency i, we regress the exchange rate changes measured between time $t+1$ and t on the changes in CEP through the following predictive regression:

$$\Delta s_{i,t+1} \quad = \quad \alpha_i + \beta_{1,i} \Delta \text{CEP}_{i,t} + u_{i,t+1}. \tag{22}$$

We then produce the one-month-ahead out-of-sample forecasts given the information available at time t as $\mathbb{E}_t[\Delta s_{i,t+1}] = \widehat{\alpha_i} + \widehat{\beta_{1,i}}\,\text{CEP}_{i,t}$, where $\widehat{\alpha_i}$ and $\widehat{\beta_{1,i}}$ denote the least-squares estimates based on an expanding window of monthly data (starting with 24 months). Following Campbell and Thompson (2008), we also impose an economic sign restriction by setting $\beta_{1,i}$ equal to zero when its estimate is negative. Such restriction is consistent with the prediction of our theory and mitigates the parameter instability arising from using a shorter time series.

The benchmark case, which is the RW model with drift, generates the one-month-ahead forecast as $\mathbb{E}_t[\Delta s_{i,t+1}] = \overline{\alpha_i}$, where $\overline{\alpha_i}$ is the conditional average exchange rate change, based on an expanding window. We will compare the performance of a parsimonious restricted null model (with $\beta_1 = 0$) to an unrestricted model that nests the parsimonious model (with $\beta_1 \neq 0$). We now describe a set of statistical criteria based on out-of-sample forecasts and then summarize our empirical findings.

We first compute, for each currency, the out-of-sample R^2 statistic of Campbell and Thompson (2008) as $R^2_{oos} = 1 - (MSE^{CEP}/MSE^{RW})$, where MSE is a given model's mean-squared error (MSE). A related statistic is Welch and Goyal (2008)'s out-of-sample root mean-squared error difference, which is computed as $\Delta RMSE = \sqrt{MSE^{RW}} - \sqrt{MSE^{CEP}}$. For both statistics, a positive value would imply that using CEP outperforms the benchmark RW model. We also assess whether using CEP delivers a lower MSE than the RW model using the statistic of Clark and West (2007) for the null of equal predictive ability for nested models. Clark and West (2007)'s statistic is defined as $CW = MSE^{RW} - (MSE^{CEP} - adj)$, where the adjustment term adj captures the average squared difference between the RW-based forecasts and the CEP-based forecasts.

For all these statistics, we compute bootstrapped critical values by generating 1,000 artificial samples under the null of no predictability as in Mark (1995b) and Kilian (1999). This procedure preserves the autocorrelation structure of the predictive variable and maintains the cross-correlation structure of the residual. The bootstrap algorithm is summarized in Internet Appendix B.

7.2 Empirical evidence

Table 9 reports the test statistics discussed above. In Panel A, we present the findings at the currency level. These findings reveal compelling evidence of enhanced predictive ability compared to the benchmark model for four of the commodity currencies. Notably, the Brazilian real and the Russia ruble, characterized by the highest commodity price beta (see Figure 2), demonstrate particularly strong predictive power as they feature the largest test statistics and the highest significance levels. Conversely, the Australian dollar and Canadian dollar, which rank among the world's most widely traded and liquid currencies (see Column 8, Table 9), unsurprisingly exhibit weak out-of-sample predictability.

Confirming this intuition, Figure 9 demonstrates the inverse relationship between a currency's out-of-sample predictability (measured by the R^2_{oos}) and its liquidity (approximated by the average daily trading volume).[41] Notably, all currencies exhibiting statistically significant out-of-sample predictability have daily average trading volumes under 100 billion USD. For instance, the Brazilian real and Russian ruble exhibit volumes of around 38 billion and 48 billion USD, respectively, which is considerably lower than the Australian dollar's volume of 316 billion USD. This highlights a stark liquidity contrast between large developed and emerging currencies. Hence, while the FX market is recognized as the world's largest and most liquid market, with trillions of dollars of daily turnover, this fact obscures a crucial detail: the bulk of this volume is concentrated in a few developed currencies (such as the euro and the Japanese yen).[42] Our analysis confirms that discovering predictability within

[41] The average daily FX turnover (in billion USD) for each country's currency is calculated using the BIS triennial surveys for the years 2004-2019.

[42] The average daily FX turnover for the euro and yen are 1502 and 861 billion USD, respectively (untabulated),

heavily traded currencies is inherently challenging. In contrast, we uncover robust predictability in currencies with lower liquidity, especially those from emerging markets.

Table 9 and Figure 9 [about here]

In Panel B, we expand our analysis by constructing various indexes to reduce noise and assess their predictability using our commodity export price index $\Delta\overline{\text{CEP}}$. First, we consider the performance of an aggregate commodity currency index, computed as the equal-weighted average monthly change of the nine commodity currencies. We find robust out-of-sample predictability of commodity export prices at the 1% significance level across the three statistical criteria. Moreover, we observe heightened predictability when we focus exclusively on emerging countries.[43] This finding reinforces our earlier conclusion that predictability is concentrated in currencies with slower adjustments to new commodity price information, as these currencies typically have lower trading volumes and higher volatility.

Lastly, we explore the out-of-sample predictive ability of commodity export prices for the carry trade profitability. The results reveal robust predictability at the 5% significance level, which improves both economically and statistically when focusing on the investment leg of the carry trade (P5). This is consistent with our evidence that this portfolio is largely exposed to commodity currencies. In contrast, the predictability of the carry trade vanishes once we exclude commodity currencies from the strategy, confirming that it is primarily attributed to the small set of currencies with significant commodity price beta, and is not a systematic phenomenon.

Overall, we find that commodity export prices contain important information for commodity currencies' out-of-sample exchange rate forecasting. However, our analysis also suggests that this predictability is driven by slow information diffusion in the relatively less-liquid currencies, while the most liquid currencies remain difficult to predict. This finding has profound implications for the performance of FX investors' trading strategies as well as for central banks and policymakers, providing them with valuable insights for their decision-making processes.

8 Conclusion

This paper investigates how fluctuations in commodity export prices impact the future performance of individual commodity currencies. Using a cross-section of 41 developed and emerging currencies, we deploy an empirical approach to identify a set of nine currencies as "commodity currencies" based on their positive and statistically significant exposure (beta) to their country's commodity export prices.

which is over 10 times greater than, for example, the turnover of the New Zealand dollar (a G10 currency) and over 20 times greater than Brazilian real's turnover.

[43] Our set of emerging market commodity currencies includes the Brazilian real, Mexican peso, Russian ruble, South African rand, and the Peruvian sol.

Our categorization aligns with the importance of the commodity sector in a country's exports, GDP, financial markets, and fiscal revenue. Within this set of currencies, we demonstrate that changes in a country's commodity export prices positively predict its one-month-ahead exchange rate, even after controlling for standard currency and financial predictors. Conditionally, this predictability is concentrated in times of elevated FX uncertainty and remains robust out-of-sample, especially for emerging market and less-traded currencies.

These results provide a fresh perspective on the conventional view that the FX market is highly active, liquid, and efficient. While it is arguably challenging to find predictability in the most-traded currencies like the Australian and Canadian dollar, we uncover strong and robust predictability in less-liquid currencies, particularly those from emerging markets. These currencies seem to slowly incorporate news into exchange rates, leading to short-term predictability. These results are consistent with an information-based mechanism, which we explore with a simple, stylized model.

Our findings also hold significant implications for the carry trade, a widely adopted FX investment strategy. The carry trade involves borrowing in low-yield currencies (e.g., the Japanese yen and Swiss franc), which we find typically exhibit low or negative commodity price beta. It also involves investing in high-yield currencies (e.g., the Brazilian real and Russian ruble), which we find to have high and positive commodity price beta. Our empirical analysis reveals that commodity export prices are powerful predictors of carry trade returns, outperforming the predictive ability of traditional global commodity price indices. Notably, this predictive power is primarily attributed to the currencies that are part of the investment portfolio and have significant commodity price beta. These turn out to be mostly emerging market currencies. Consequently, we find that the predictive ability of commodity export prices is the result of the exchange rate predictability of a few less-liquid commodity currencies - and not a reflection of a global systematic risk factor or an alignment of commodity investments with an appetite for risk.

Future research could explore the links between commodity options and FX options on commodity currencies and test whether these options are consistently priced. For example, it would be interesting to examine the role of crash risk in commodities, as experienced during the 2020 oil market turmoil, in explaining tail risk premiums in commodity currencies. This line of research could significantly advance our understanding of the FX options market.

References

Amano, Robert A., and Simon Van Norden, 1998, Oil prices and the rise and fall of the US real exchange rate, *Journal of International Money and Finance* 17, 299–316.

Atmaz, Adem, and Suleyman Basak, 2018, Belief dispersion in the stock market, *The Journal of Finance* 73, 1225–1279.

Aumann, Robert J., 1976, Agreeing to disagree, *The Annals of Statistics* 1236–1239.

Bacchetta, Philippe, and Eric Van Wincoop, 2006, Can information heterogeneity explain the exchange rate determination puzzle?, *American Economic Review* 96, 552–576.

Bakshi, Gurdip, and George Panayotov, 2013, Predictability of currency carry trades and asset pricing implications, *Journal of Financial Economics* 110, 139–163.

Balduzzi, Pierluigi, and I-Hsuan Ethan Chiang, 2019, Real Exchange Rates and Currency Risk Premiums, *The Review of Asset Pricing Studies* 10, 94–121.

Banerjee, Snehal, Ron Kaniel, and Ilan Kremer, 2009, Price drift as an outcome of differences in higher-order beliefs, *The Review of Financial Studies* 22, 3707–3734.

Banerjee, Snehal, and Ilan Kremer, 2010, Disagreement and learning: Dynamic patterns of trade, *The Journal of Finance* 65, 1269–1302.

Bansal, Ravi, and Magnus Dahlquist, 2000, The forward premium puzzle: different tales from developed and emerging economies, *Journal of International Economics* 51, 115–144.

Bekaert, Geert, and George Panayotov, 2020, Good carry, bad carry, *Journal of Financial and Quantitative Analysis* 55, 1063–1094.

Bhamra, Harjoat S., and Raman Uppal, 2014, Asset prices with heterogeneity in preferences and beliefs, *The Review of Financial Studies* 27, 519–580.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2008, Carry trades and currency crashes, *NBER Macroeconomics Annual* 23, 313–348.

Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *The Review of Financial Studies* 21, 1509–1531.

Cao, H. Henry, and Hui Ou-Yang, 2008, Differences of opinion of public information and speculative trading in stocks and options, *The Review of Financial Studies* 22, 299–335.

Cespa, Giovanni, Antonio Gargano, Steven J. Riddiough, and Lucio Sarno, 2022, Foreign exchange volume, *The Review of Financial Studies* 35, 2386–2427.

Chen, Yu-Chin, and Kenneth Rogoff, 2003, Commodity currencies, *Journal of International Economics* 60, 133–160.

Chen, Yu-Chin, Kenneth S. Rogoff, and Barbara Rossi, 2010, Can exchange rates forecast commodity prices?, *Quarterly Journal of Economics* 125, 1145–1194.

Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze, 2018, Crash risk in currency returns, *Journal of Financial and Quantitative Analysis* 1–34.

Christiansen, Charlotte, Angelo Ranaldo, and Paul Söderlind, 2011, The time-varying systematic risk of carry trade strategies, *Journal of Financial and Quantitative Analysis* 46, 1107–1125.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Colacito, Riccardo, Steven J. Riddiough, and Lucio Sarno, 2020, Business cycles and currency returns, *Journal of Financial Economics* 137, 659–678.

Dahlquist, Magnus, and Henrik Hasseltoft, 2020, Economic momentum and currency returns, *Journal of Financial Economics* 136, 152–167.

Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The carry trade: Risks and drawdowns, *Critical Finance Review* 6, 211–262.

Della Corte, Pasquale, Alexandre Jeanneret, and Ella D.S. Patelli, 2023, A credit-based theory of the currency risk premium, *Journal of Financial Economics* 149, 473–496.

Della Corte, Pasquale, Steven J. Riddiough, and Lucio Sarno, 2016, Currency premia and global imbalances, *The Review of Financial Studies* 29, 2161–2193.

Driscoll, John C., and Aart C. Kraay, 1998, Consistent covariance matrix estimation with spatially dependent panel data, *Review of Economics and Statistics* 80, 549–560.

Dumas, Bernard, Karen K. Lewis, and Emilio Osambela, 2017, Differences of opinion and international equity markets, *The Review of Financial Studies* 30, 750–800.

Eichenbaum, M. S., B. K. Johannsen, and S. T. Rebelo, 2021, Monetary Policy and the Predictability of Nominal Exchange Rates, *The Review of Economic Studies* 88, 192–228.

Engel, Charles, Nelson C. Mark, and Kenneth D. West, 2007, Exchange rate models are not as bad as you think [with comments and discussion], *NBER macroeconomics annual* 22, 381–473.

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

Ferraro, Domenico, Kenneth Rogoff, and Barbara Rossi, 2015, Can oil prices forecast exchange rates? An empirical analysis of the relationship between commodity prices and exchange rates, *Journal of International Money and Finance* 54, 116–141.

Gabaix, Xavier, 2011, The granular origins of aggregate fluctuations, *Econometrica* 79, 733–772.

Granger, C. W., and T. Terasvirta, 1993, Modelling nonlinear economic relationships, *Oxford University Press, New York* .

Gruss, Bertrand, and Suhaib Kebhaj, 2019, Commodity terms of trade: A new database, *IMF Working Paper* 19/21.

Harris, Milton, and Artur Raviv, 1993, Differences of opinion make a horse race, *The Review of Financial Studies* 6, 473–506.

Harrison, J. Michael, and David M. Kreps, 1978, Speculative investor behavior in a stock market with heterogeneous expectations, *The Quarterly Journal of Economics* 92, 323–336.

Hassan, Tarek A., and Tony Zhang, 2021, The economics of currency risk, *Annual Reviews of Economics* 13, 281–307.

Itskhoki, Oleg, and Dmitry Mukhin, 2021, Exchange rate disconnect in general equilibrium, *Journal of Political Economy* 129, 2183–2232.

Jordà, Òscar, 2005, Estimation and inference of impulse responses by local projections, *American Economic Review* 95, 161–182.

Kandel, Eugene, and Neil D. Pearson, 1995, Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy* 103, 831–872.

Karnaukh, Nina, Angelo Ranaldo, and Paul Söderlind, 2015, Understanding FX liquidity, *Review of Financial Studies* 28, 3073–3108.

Kilian, Lutz, 1999, Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions?, *Journal of Applied Econometrics* 14, 491–510.

Kilian, Lutz, and Mark P Taylor, 2003, Why is it so difficult to beat the random walk forecast of exchange rates?, *Journal of International Economics* 60, 85–107.

King, Michael R., Carol Osler, and Dagfinn Rime, 2012, Foreign exchange market structure, players, and evolution, *Handbook of Exchange Rates* 1–44.

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.

Levich, Richard M., and Lee R. Thomas III, 1993, The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach, *Journal of international Money and Finance* 12, 451–474.

Llorente, Guillermo, Roni Michaely, Gideon Saar, and Jiang Wang, 2002, Dynamic volume-return relation of individual stocks, *The Review of Financial Studies* 15, 1005–1047.

Londono, Juan M., and Hao Zhou, 2017, Variance risk premiums and the forward premium puzzle, *Journal of Financial Economics* 124, 415–440.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *The Review of Financial Studies* 24, 3731–3777.

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Mancini, Loriano, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums, *The Journal of Finance* 68, 1805–1841.

Mark, Nelson C., 1995a, Exchange rates and fundamentals: Evidence on long-horizon predictability, *The American Economic Review* 201–218.

Mark, Nelson C., 1995b, Exchange rates and fundamentals: Evidence on long-horizon predictability, *The American Economic Review* 85, 201–218.

Meese, Richard A., and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of International Economics* 14, 3–24.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012, Carry trades and global foreign exchange volatility, *The Journal of Finance* 67, 681–718.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2016, Information flows in foreign exchange markets: Dissecting customer currency trades, *The Journal of Finance* 71, 601–634.

Neary, Peter, 1988, Determinants of the equilibrium real exchange rate, *The American Economic Review* 78, 210–215.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–08.

Obstfeld, Maurice, and Kenneth Rogoff, 2000, The six major puzzles in international macroeconomics: is there a common cause?, *NBER Macroeconomics Annual* 15, 339–390.

OECD, 2020, *Government at a Glance: Latin America and the Caribbean 2020*.

Opie, Wei, and Steven J. Riddiough, 2020, Global currency hedging with common risk factors, *Journal of Financial Economics* 136, 780–805.

Pukthuanthong-Le, Kuntara, and Lee R. Thomas, III, 2008, Weak-form efficiency in currency markets, *Financial Analysts Journal* 64, 31–52.

Ramey, Valerie A., and Sarah Zubairy, 2018, Government spending multipliers in good times and in bad: Evidence from US historical data, *Journal of Political Economy* 126, 850–901.

Ranaldo, Angelo, and Paolo Santucci de Magistris, 2022, Liquidity in the global currency market, *Journal of Financial Economics* 146, 859–883.

Ranaldo, Angelo, and Paul Söderlind, 2010, Safe haven currencies, *Review of Finance* 14, 385–407.

Ranaldo, Angelo, and Fabricius Somogyi, 2021, Asymmetric information risk in fx markets, *Journal of Financial Economics* 140, 391–411.

Ready, Robert, Nikolai Roussanov, and Colin Ward, 2017, Commodity trade and the carry trade: A tale of two countries, *The Journal of Finance* 72, 2629–2684.

Rossi, Barbara, 2013, Exchange rate predictability, *Journal of Economic Literature* 51, 1063–1119.

Söderlind, Paul, and Fabricius Somogyi, 2023, Liquidity risk and currency premia, *forthcoming, Management Science* .

Stoll, Hans R., 1978, The supply of dealer services in securities markets, *The Journal of Finance* 33, 1133–1151.

Taylor, Mark P, and David A Peel, 2000, Nonlinear adjustment, long-run equilibrium and exchange rate fundamentals, *Journal of international Money and Finance* 19, 33–53.

Taylor, Mark P., David A Peel, and Lucio Sarno, 2001, Nonlinear mean-reversion in real exchange rates: toward a solution to the purchasing power parity puzzles, *International Economic Review* 42, 1015–1042.

Tenreyro, Silvana, and Gregory Thwaites, 2016, Pushing on a string: US monetary policy is less powerful in recessions, *American Economic Journal: Macroeconomics* 8, 43–74.

Verdelhan, Adrien, 2010, A habit-based explanation of the exchange rate risk premium, *The Journal of Finance* 65, 123–146.

Verdelhan, Adrien, 2018, The share of systematic variation in bilateral exchange rates, *The Journal of Finance* 73, 375–418.

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *The Review of Financial Studies* 21, 1455–1508.

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.

Table 1: **Drivers of the commodity price beta**

This table reports cross-sectional regression results of each country's commodity price beta on the commodity share of exports, total commodity rents (as % of GDP), commodity sector share in each country's stock market, and share of commodity revenue of total government revenue. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity price index changes and the dollar factor over a sample period of January 1985 and April 2020. Explanatory variables are averages calculated over the same (or shorter due to data availability) sample period as the commodity price betas. All the variables are standardized (a constant term is included in all regressions, but not reported as it is zero in all specifications). White (1980) standard errors are reported in parentheses. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Comm. export share	0.683^{***}				0.296^{*}
	(0.161)				(0.159)
Comm. rents share		0.768^{***}			-0.008
		(0.143)			(0.249)
Comm. mkt. cap. share			0.681^{***}		0.380^{***}
			(0.163)		(0.135)
Comm. gov. revenue share				0.663^{***}	0.388^{**}
				(0.190)	(0.144)
R^2 %	45.27	57.92	45.00	42.48	67.96
Observations	41	41	41	41	41

Table 2: **Predictability for commodity currencies**

This table reports panel regression results on the exchange rate predictive ability of commodity export prices. The dependent variable is the one-month-ahead exchange rate change of each commodity currency against the US dollar. The variable ΔCEP is the change of the corresponding country's commodity price index, which is export-weighted and rebalanced monthly. Model (1) is the univariate regression, Model (2) controls for the one-month interest rate differential (IRD) of each currency relative to the US dollar, Model (3) includes the dollar factor (DOL), Model (4) includes changes in aggregate currency volatility (FX Vol), while Models (5) and (6) include changes in funding liquidity (TED) and in aggregate uncertainty in the US market (VIX), respectively. Model (7) includes an indicator variable equal to one during NBER recessions and zero otherwise. All specifications include currency fixed effects. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with an optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the econometric specification and controls. The sample consists of monthly observations between January 1985 and April 2020.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔCEP	0.374^{***}	0.380^{***}	0.367^{***}	0.372^{***}	0.376^{***}	0.377^{***}	0.367^{***}
	(0.140)	(0.142)	(0.130)	(0.125)	(0.126)	(0.126)	(0.116)
IRD		0.249^{**}	0.248^{**}	0.247^{**}	0.245^{**}	0.243^{**}	0.245^{**}
		(0.113)	(0.114)	(0.113)	(0.118)	(0.117)	(0.114)
DOL			0.020	0.024	0.029	0.017	0.015
			(0.061)	(0.057)	(0.057)	(0.060)	(0.060)
ΔFX Vol				0.068	0.047	0.075	0.083
				(0.206)	(0.196)	(0.212)	(0.212)
ΔTED					−0.240	−0.154	−0.162
					(1.198)	(1.193)	(1.193)
ΔVIX						−0.021	−0.022
						(0.045)	(0.045)
NBER							−0.323
							(0.841)
Observations	2967	2966	2966	2966	2904	2904	2904
R^2 %	0.83	1.10	1.08	1.07	1.19	1.24	1.28
Currencies	9	9	9	9	9	9	9
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 3: **Predictability for commodity currencies – additional results**

This table reports panel regression results on the exchange rate predictive ability of commodity export prices under various specifications. The dependent variable is each commodity currency's one-month-ahead exchange rate change. The variable ΔCEP is the change of the corresponding country's commodity price index, which is export-weighted and rebalanced monthly. Models (1), (2), and (3) report results for individual commodity currencies against the Swiss franc, the euro (spliced with the Deutsche mark), and the Japanese yen, respectively. Models (4) to (5) report results for non-commodity currencies against the US dollar. Model (4) only includes the currencies that do not exhibit a statistically significant positive commodity price beta, using the methodology described in Section 3. Model (5) only includes commodity importers' currencies, i.e., currencies exhibiting a statistically significant negative commodity price beta. Model (6) uses all currencies and includes an interaction between ΔCEP and the commodity price beta β. All specifications include currency fixed effects and controls. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with an optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the econometric specification and controls. The sample consists of monthly observations between January 1985 and April 2020.

	Commodity currencies			**Non-commodity currencies**		**All**
	(1)	(2)	(3)	(4)	(5)	(6)
ΔCEP	0.317^{**}	0.293^{**}	0.397^{*}	0.164	0.032	0.179^{*}
	(0.141)	(0.121)	(0.236)	(0.106)	(0.125)	(0.099)
ΔCEP × β						0.471^{***}
						(0.166)
R^2 %	1.91	1.95	1.93	0.32	0.55	0.65
Observations	2,904	2,898	2,904	7179	1983	10083
Currencies	9	9	9	32	7	41
Base currency	CHF	EUR	JPY	USD	USD	USD
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Table 4: **Conditional predictability for commodity currencies**

This table reports panel regression results on the conditional exchange rate predictive ability of commodity export prices. The dependent variable is each commodity currency's one-month-ahead exchange rate change against the US dollar. The variable ΔCEP is the change of the corresponding country's commodity price index, which is export-weighted and rebalanced monthly. The conditioning variable F reflects the probability of being in a high FX volatility regime, as defined in Section 5.2. Model (1) is the univariate regression, Model (2) controls for the one-month interest rate differential (IRD) relative to the US dollar, Model (3) includes the dollar factor (DOL), Model (4) includes changes in aggregate currency volatility (FX Vol), while Models (5) and (6) include changes in funding liquidity (TED) and aggregate uncertainty in the US market (VIX), respectively. Model (7) includes an indicator variable equal to one during NBER recessions and zero otherwise. All specifications include currency fixed effects. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with an optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the econometric specification and controls. The sample consists of monthly observations between January 1985 and April 2020.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔCEP × F	0.511^{***}	0.521^{***}	0.510^{***}	0.516^{***}	0.532^{***}	0.525^{***}	0.515^{***}
	(0.188)	(0.190)	(0.175)	(0.169)	(0.182)	(0.182)	(0.163)
ΔCEP × $(1-F)$	0.186	0.185	0.172	0.177	0.153	0.165	0.156
	(0.200)	(0.200)	(0.204)	(0.207)	(0.209)	(0.213)	(0.215)
IRD		0.256^{**}	0.254^{**}	0.254^{**}	0.249^{**}	0.246^{**}	0.249^{**}
		(0.110)	(0.111)	(0.110)	(0.115)	(0.114)	(0.112)
DOL			0.018	0.022	0.029	0.018	0.017
			(0.060)	(0.056)	(0.056)	(0.060)	(0.060)
ΔFX Vol				0.076	0.051	0.076	0.085
				(0.201)	(0.189)	(0.207)	(0.208)
ΔTED					−0.287	−0.207	−0.215
					(1.209)	(1.214)	(1.211)
ΔVIX						−0.019	−0.020
						(0.046)	(0.046)
NBER							−0.322
							(0.811)
R^2 %	0.93	1.21	1.19	1.19	1.31	1.34	1.38
Observations	2,962	2,961	2,961	2,961	2,904	2,904	2,904
Currencies	9	9	9	9	9	9	9
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: **Conditional predictability for commodity currencies – robustness**

This table reports panel regression results on the conditional exchange rate predictive ability of commodity export prices under various specifications. The dependent variable is the one-month-ahead exchange rate change of each commodity currency. The variable ΔCEP is the change of the corresponding country's commodity prices, which is export-weighted and rebalanced monthly. The conditioning variable F reflects the probability of being in a high FX volatility regime, as defined in Section 5.2. Models (1) and (2) report the predictability results conditional on the level of FX volatility orthogonalized with respect to the VIX and the TED spread, respectively. Model (3) uses the level of FX volatility orthogonalized with respect to the VIX and the TED spread, while Model (4) uses the level of FX volatility orthogonalized with respect to Karnaukh et al. (2015)'s level of FX illiquidity. Model (5) reports the predictability results conditional on the dispersion in professional FX forecasts as an alternative measure of FX uncertainty. All specifications include currency fixed effects and controls. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with an optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 3 presents the commodity price indices, while Section 4 describes the econometric specification and controls. The sample consists of monthly observations between January 1985 and April 2020.

	Orthogonalized FX volatility				**Alternative FX uncertainty**
	⊥ VIX	⊥ TED	⊥ VIX & TED	⊥ FX Liq	FX Forecast Disp.
	(1)	(2)	(3)	(4)	(5)
ΔCEP × F	0.507^{***}	0.471^{***}	0.508^{***}	0.579^{***}	0.632^{***}
	(0.174)	(0.159)	(0.172)	(0.195)	(0.237)
ΔCEP × (1−F)	0.200	0.216	0.190	0.093	0.155
	(0.174)	(0.202)	(0.178)	(0.198)	(0.254)
R^2 %	1.37	1.32	1.38	1.61	1.96
Observations	2,904	2,904	2,904	2,533	1,811
Currencies	9	9	9	9	9
Currency FE	Yes	Yes	Yes	Yes	Yes

Table 6: **Carry trade descriptive statistics**

This table presents key characteristics of the carry trade portfolios. Five equally weighted portfolios are formed every month, with Portfolio 1 (P1) containing the currencies with the lowest interest rates and Portfolio 5 (P5) containing those with the highest. The carry trade is the strategy that is long in P5 and short in P1. Panel A presents summary statistics (mean, standard deviation, skewness, and Sharpe ratio) of the monthly excess returns of the five portfolios and the carry trade. The mean and standard deviations are annualized. IRD is the average annualized one-month interest rate differential relative to the US dollar. ΔS denotes the average annualized monthly appreciation of portfolio member currencies relative to the US dollar. Portfolios are rebalanced each month and a total of 41 currencies are considered, however, the number of available currencies for portfolio construction varies between periods depending on data availability. Panel B presents the fraction of time during which each of the commodity currencies is a member of a specific portfolio. Average membership is the average frequency that commodity currencies belong to a specific currency portfolio. Row G10 reports the average portfolio membership of the G10 currencies. The sample consists of monthly observations between January 1985 and April 2020.

	P1	P2	P3	P4	P5	P5-P1
Panel A: **Summary statistics**						
Mean (%)	−0.447	1.277	2.439	2.288	7.148	7.595
Standard deviation (%)	8.025	8.847	8.653	8.829	10.773	8.970
Sharpe ratio	−0.056	0.144	0.282	0.259	0.664	0.847
Skewness	0.042	−0.329	−0.261	−0.653	−0.655	−0.737
IRD (%)	−1.821	−0.002	1.512	3.742	9.497	11.319
ΔS (%)	1.375	1.279	0.928	−1.454	−2.349	−3.724
Panel B: **Commodity currency membership (%)**						
Australia	4.7	9.2	31.6	40.3	14.2	
Brazil	1.0	0.0	0.0	16.6	82.4	
Canada	19.6	38.7	31.6	10.1	0.0	
Mexico	0.0	0.0	3.9	43.9	52.1	
New Zealand	4.2	9.7	25.5	35.1	25.5	
Norway	8.5	23.3	40.8	19.3	8.0	
Peru	9.8	8.8	26.9	45.6	8.8	
Russia	4.1	4.7	11.9	14.5	64.8	
South Africa	0.2	3.3	5.0	16.1	75.4	
Average membership	5.8	10.9	19.7	26.9	36.8	
G10	31.6	24.3	20.9	16.7	6.4	

Table 7: **Carry trade return predictability**

This table presents results on the predictability of carry trade returns with commodity export prices. Five equally weighted portfolios are formed every month, with Portfolio 1 (P1) containing the currencies with the lowest interest rates and Portfolio 5 (P5) containing those with the highest. The carry trade is the strategy that is long in P5 and short in P1. We report the results of a regression of one-month-ahead carry trade returns on $\Delta\overline{\text{CEP}}$, which is the average (equally weighted) change across each commodity country's commodity export price index. Model (1) is the univariate regression, Model (2) controls for the dollar factor (DOL), and Model (3) controls for changes in aggregate currency volatility (FX Vol). Models (4) and (5) include changes in funding liquidity (TED) and in aggregate uncertainty in the US market (VIX), respectively. Model (6) additionally controls for an indicator variable equal to one during NBER recessions and zero otherwise. The standard errors in parentheses are adjusted using the Newey and West (1987) kernel, where the bandwidth is selected optimally. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the controls. The sample consists of monthly observations between January 1985 and April 2020.

	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta\overline{\text{CEP}}$	0.107^{***}	0.095^{***}	0.105^{***}	0.102^{***}	0.103^{***}	0.105^{***}
	(0.035)	(0.034)	(0.035)	(0.035)	(0.035)	(0.036)
DOL		0.063	0.085	0.088	0.076	0.077
		(0.054)	(0.056)	(0.059)	(0.062)	(0.062)
ΔFX Vol			0.449^{***}	0.343^{**}	0.382^{**}	0.376^{**}
			(0.150)	(0.158)	(0.173)	(0.172)
ΔTED				−0.450	−0.334	−0.331
				(0.761)	(0.814)	(0.819)
ΔVIX					−0.029	−0.028
					(0.040)	(0.040)
NBER						0.244
						(0.462)
Constant	0.633^{***}	0.633^{***}	0.629^{***}	0.631^{***}	0.632^{***}	0.611^{***}
	(0.134)	(0.133)	(0.135)	(0.136)	(0.135)	(0.142)
R^2 %	1.92	1.99	3.48	2.73	2.76	2.58
Observations	423	423	423	410	410	410

Table 8: **Carry trade return predictability - additional results**

This table presents results on the predictability of carry trade and portfolio returns with commodity export prices under alternative specifications. Five equally weighted portfolios are formed every month, with Portfolio 1 (P1) containing the currencies with the lowest interest rates and Portfolio 5 (P5) containing those with the highest. The carry trade is the strategy that is long in P5 and short in P1. The benchmark predictor is $\Delta\overline{\text{CEP}}$, which is the average (equally-weighted) change across each commodity country's commodity export price index. Model (1) reports results on the one-month-ahead returns to a carry trade strategy constructed from non-commodity currencies (Without CC), i.e., those that do not exhibit a statistically significant positive commodity price beta, using the methodology described in Section 3. Model (2) reports results using the one-month-ahead change in P1, which rarely contains commodity currencies. Models (3)-(8) report results using the one-month-ahead change in P5, which frequently contains commodity currencies. Models (4), (5), and (6) control for ΔCRB, ΔGSCI, and ΔOil, which are the percentage change in the Commodity Research Bureau (CRB) index, Goldman Sachs commodity index (GSCI), and the West Texas Intermediate Crude Oil price, respectively. Model (7) reports results conditional on the probability of being in a high FX volatility regime, as defined in Section 5.2. Model (8) reports results conditional on the dispersion in professional FX forecasts as an alternative measure of FX uncertainty. All specifications include controls, which are presented in Section 4. The standard errors in parentheses are adjusted using the Newey and West (1987) kernel where the bandwidth is selected optimally. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample consists of monthly observations between January 1985 and April 2020.

	Counterfactual cases		**Base case**	**Controlling for alt. commodity indices**			**Conditional predictability**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta\overline{\text{CEP}}$	0.049	0.023	0.142^{***}	0.119^{***}	0.137^{***}	0.167^{***}		
	(0.030)	(0.032)	(0.039)	(0.041)	(0.053)	(0.054)		
$\Delta\overline{\text{CEP}} \times F$							0.158^{***}	0.200^{***}
							(0.051)	(0.069)
$\Delta\overline{\text{CEP}} \times (1{-}F)$							0.150^{*}	0.086
							(0.082)	(0.090)
ΔCRB				0.073				
				(0.059)				
ΔGSCI					0.005			
					(0.036)			
ΔOil						−0.014		
						(0.019)		
Constant	0.477^{***}	0.080	−0.237	-0.254^{*}	−0.238	−0.235	−0.238	−0.188
	(0.140)	(0.115)	(0.151)	(0.153)	(0.152)	(0.151)	(0.158)	(0.186)
R^2 %	1.00	0.51	2.69	2.75	2.45	2.58	2.38	3.61
Observations	410	410	410	410	410	410	410	206
Portfolio	Without CC	P1	P5	P5	P5	P5	P5	P5

Table 9: **Out-of-sample predictability**

This table reports the out-of-sample predictive ability of the commodity export prices (CEP) against the random walk (RW) for commodity currencies' exchange rate changes. Panel A presents results for each commodity currency, while Panel B considers various aggregate series. These include an (equally weighted) commodity currencies index (using the nine commodity currencies or emerging market (EM) currencies only); the carry trade return, which is long in currencies with the highest interest rates (P5) and short currencies with the lowest interest rates (P1); the return on P5; and the return on the counterfactual carry trade strategy constructed from non-commodity currencies (Without CC). Columns (1-2) report Campbell and Thompson (2008)'s out-of-sample R^2 statistic, given by $R^2_{oos} = 1 - (MSE^{CEP}/MSE^{RW})$, where MSE is the mean-squared error (MSE) of a given model. Columns (3-4) report Welch and Goyal (2008)'s out-of-sample root mean-squared error difference, which is computed as $\Delta RMSE = \sqrt{MSE^{RW}} - \sqrt{MSE^{CEP}}$. Columns (5-6) report the statistic of Clark and West (2007), defined as $CW = MSE^{RW} - (MSE^{CEP} - adj)$, where the adjustment term, adj, captures the average squared difference between the RW-based forecasts and the CEP-based forecasts. For all statistics, a positive value would imply that using CEP outperforms the benchmark RW model. The p-values are computed based on 1,000 bootstrap replications, using the methodology described in the Internet Appendix B. Column (7) reports the average daily FX turnover (in billion USD) for each country's currency calculated using the BIS triennial surveys for the years 2004-2019. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample consists of monthly observations between January 1985 and April 2020.

	R^2_{oos}		$\Delta RMSE$		CW		**Daily turnover**
	statistics	p-value	statistics	p-value	statistics	p-value	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: **Currency level**							
Australia	-0.81	0.542	-0.013	0.514	-0.0004	0.605	315.88
Brazil	3.48^{**}	0.021	0.083^{**}	0.021	0.0199^{**}	0.014	37.67
Canada	-0.23^{*}	0.094	-0.003^{*}	0.094	0.0002	0.202	211.70
Mexico	0.35	0.103	0.005	0.103	0.0013	0.179	76.35
New Zealand	1.03^{**}	0.026	0.017^{**}	0.028	0.0069^{***}	0.005	82.14
Norway	0.22^{*}	0.083	0.003^{*}	0.083	0.0010	0.156	71.60
Russia	3.00^{**}	0.023	0.071^{**}	0.021	0.0237^{***}	0.005	48.13
South Africa	-0.27	0.135	-0.006	0.129	0.0006	0.255	42.31
Peru	1.67^{**}	0.018	0.014^{**}	0.017	0.0009^{*}	0.052	2.38
Panel B: **Aggregate level**							
Comm. currency index (all)	2.56^{***}	0.001	0.030^{***}	0.001	0.0026^{***}	0.004	
Comm. currency index (EM)	3.66^{***}	0.001	0.050^{***}	0.001	0.0044^{***}	0.010	
Carry trade	1.45^{**}	0.014	0.019^{**}	0.014	0.0025^{**}	0.016	
P5	2.31^{***}	0.005	0.035^{***}	0.006	0.0043^{***}	0.008	
Counterfactual:							
Carry trade (Without CC)	-0.35	0.216	-0.004	0.222	0.0000	0.385	

Figure 1: **Norwegian krone and Russian ruble performance relative to the oil price**

This figure illustrates the relationship between selected currencies and the price of their main commodity exports. Panels (a) and (b) respectively show the Norwegian krone (NOK) and the Russian ruble (RUB) monthly exchange rate indexes alongside the West Texas Intermediate (WTI) crude oil price (in USD per barrel) from January 2004 to April 2020. Panels (c) and (d) show the daily time series of the NOK and RUB indexes and the WTI oil price over a shorter period, around the Russia - Saudi Arabia oil price war in early 2020. The two vertical dashed lines denote the start of the price war (8 March 2020, when Saudi Arabia announced a significant oil price reduction to its customers) and its unofficial end (3 April 2020, when Russian President Putin made a public announcement that global production could be cut). The indexes are constructed using the average changes of the FCU/USD, FCU/CHF, FCU/EUR, and FCU/JPY exchange rates. The exchange rates are expressed as base currency (e.g., USD) per foreign currency unit (FCU), i.e., a decrease in the rate implies a depreciation of the foreign currency. The index value is set to 100 on 31 March 2004 in the top panels and on 1 February 2020 in the bottom panels. The economic ratios reported in the panels are averages and are discussed in Section 3.

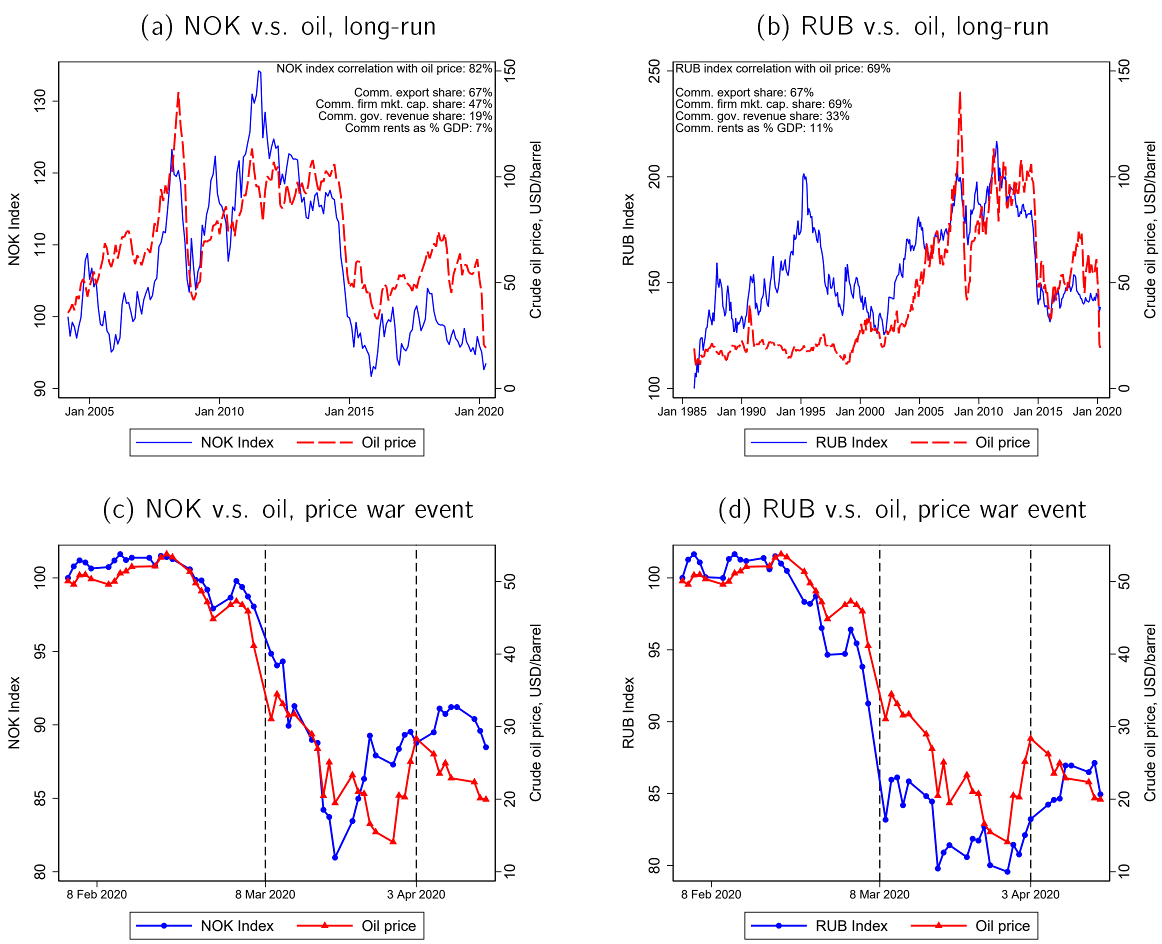

Figure 2: **Commodity price beta by country**

This figure illustrates each country's commodity price beta, which reflects the country's exchange rate exposure to its commodity export prices. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity price index changes and the dollar factor (the average change in exchange rates against the US dollar). Country-level commodity price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Countries are ranked according to their commodity price beta. Red diamonds denote commodity price betas that are positive and statistically significant at the 10% level. Plot whiskers represent the 90% confidence intervals. The sample consists of monthly observations between January 1985 and April 2020.

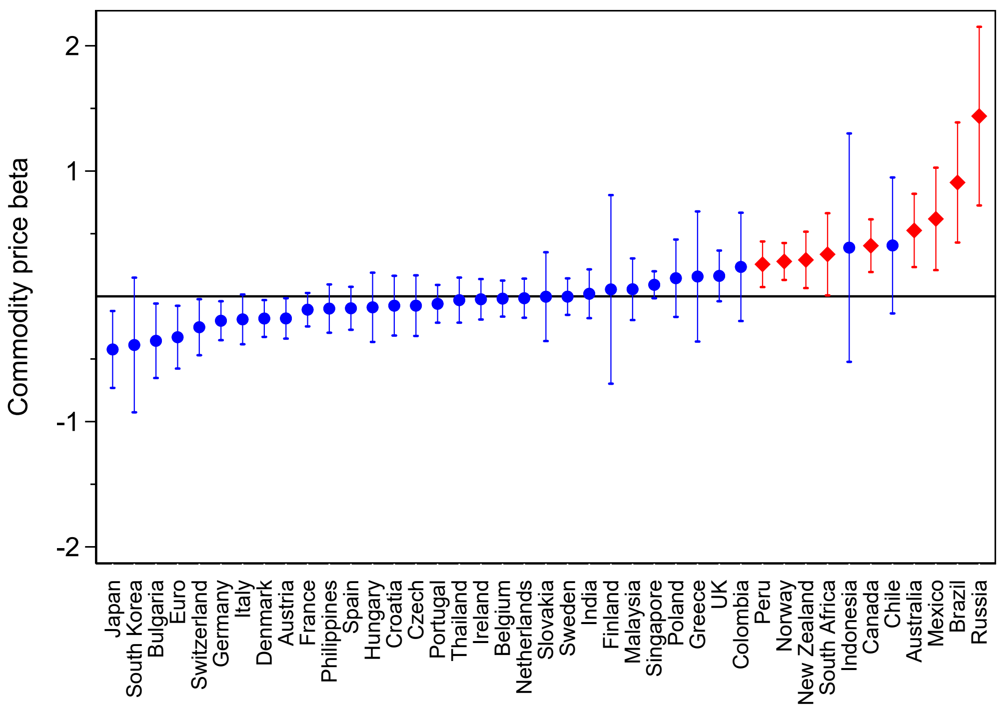

Figure 3: **Commodity price beta and the importance of the commodity sector**

This figure reports the relationships between country-level commodity price betas and proxies for the importance of the commodity sector in a country. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity price index changes and the dollar factor. Country-level commodity price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Red diamonds denote commodity price betas that are positive and statistically significant at the 10% level. Country-level commodity price betas are plotted against (a) the average share (in %) of raw commodity exports of each country's total exports, (b) the average total natural resource rents as % of GDP, (c) the average share (in %) of the commodity sector of each country's stock market capitalization, and (d) the average share (in %) of commodity-based revenue of total government revenue. The scatter plots include fitted regression lines and 95% confidence intervals. Commodity price betas are estimated over the sample period between January 1985 and April 2020, and the commodity share variables are averages calculated over the same (or shorter due to data availability) sample period.

(a) Commodity export share

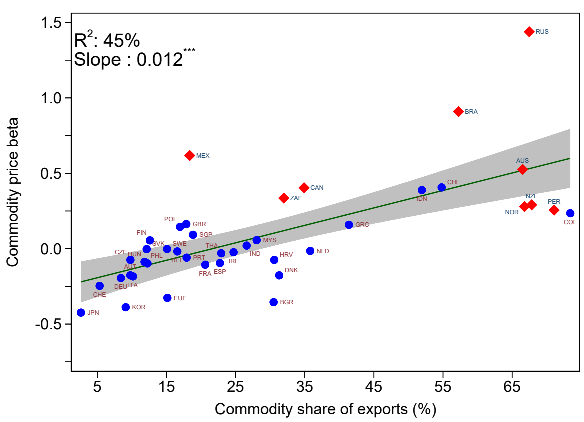

(b) Commodity rents GDP share

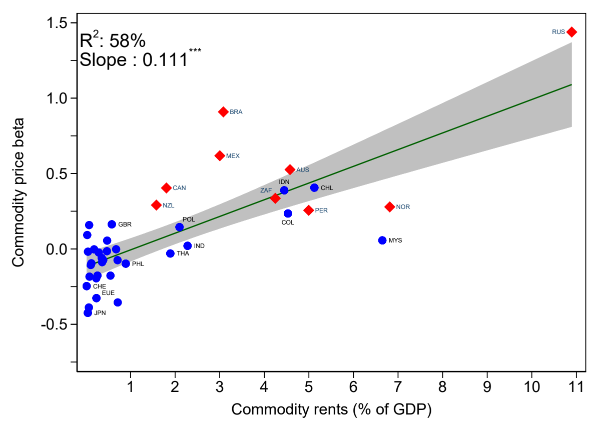

(c) Commodity market cap. share

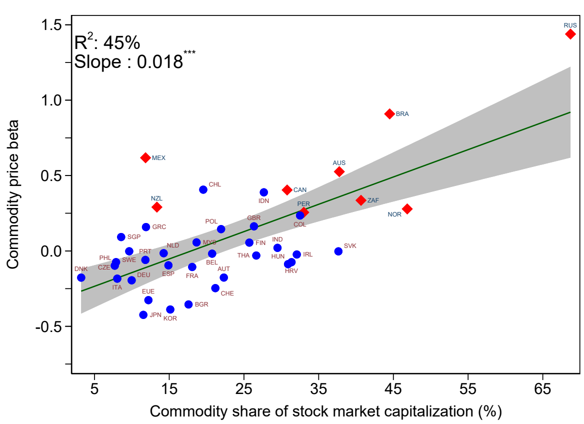

(d) Commodity gov. revenue share

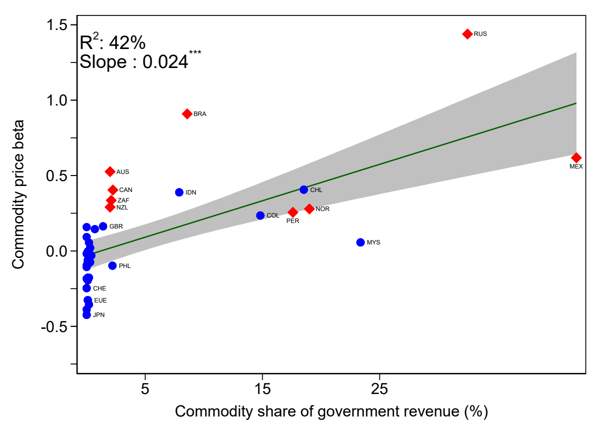

Figure 4: **Unconditional exchange rate predictability over different horizons**

This figure presents the exchange rate predictive ability of commodity export prices over different horizons. The reported slope coefficient is estimated by regressing the exchange rate changes of each commodity currency, computed up to 12 months ahead, on the changes of the corresponding country's commodity price index, which is export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. The panel regression specification includes the control variables, namely the one-month interest rate differential, the dollar factor, changes in aggregate FX volatility, changes in funding liquidity (TED spread), changes in aggregate uncertainty in the US market (VIX), and an NBER recession indicator. Reported 95% confidence intervals use standard errors based on Driscoll and Kraay (1998), which are adjusted for serial correlation using the Newey and West (1987) kernel with a bandwidth equal to the forecasting horizon. Section 4 describes the econometric specification and controls. The sample consists of monthly observations between January 1985 and April 2020.

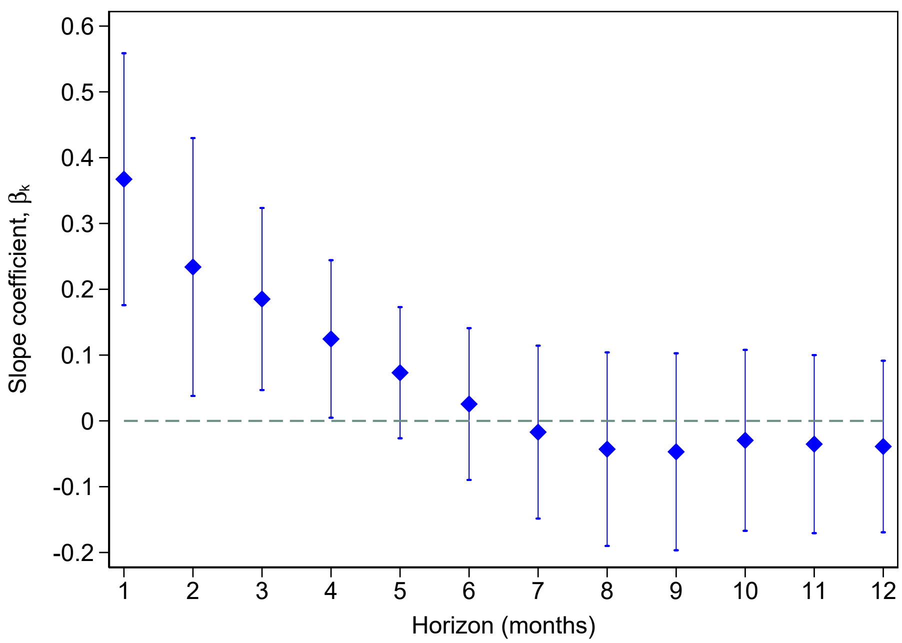

Figure 5: **Regime probabilities and FX volatility**

The figure shows the probability of being in times of high FX volatility. The reported series are the smoothed transition functions of changing from a low to a high FX volatility regime, using the methodology described in Section 5.3. Panel (a) uses the raw measure of FX volatility, constructed as the average realized volatility across 41 currencies against the US dollar. Panel (b) uses a version orthogonalized with respect to global financial conditions, as measured by aggregate uncertainty in the US market (VIX) and funding liquidity (TED spread). Grey areas indicate NBER recession periods. The sample consists of monthly observations between January 1985 and April 2020.

(a) FX volatility

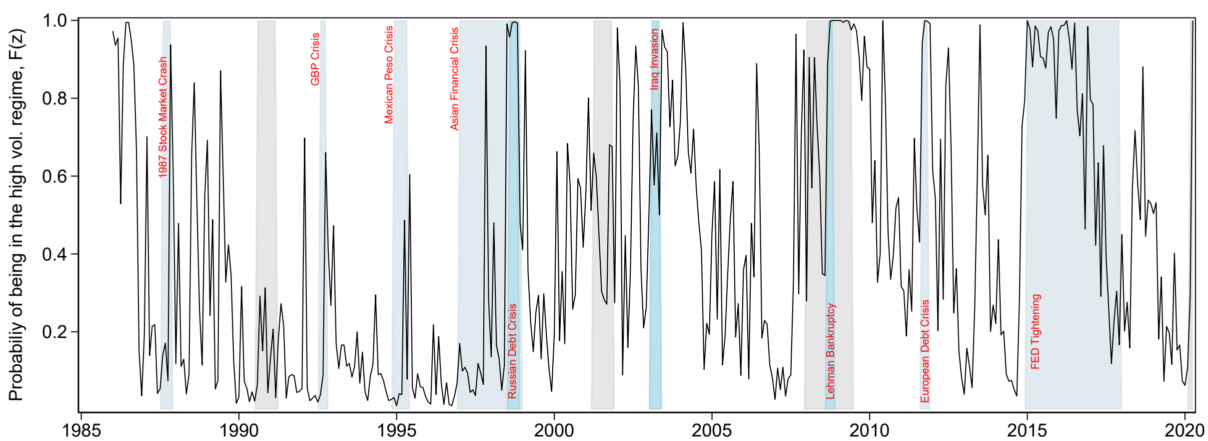

(b) FX volatility orthogonalized to VIX and TED

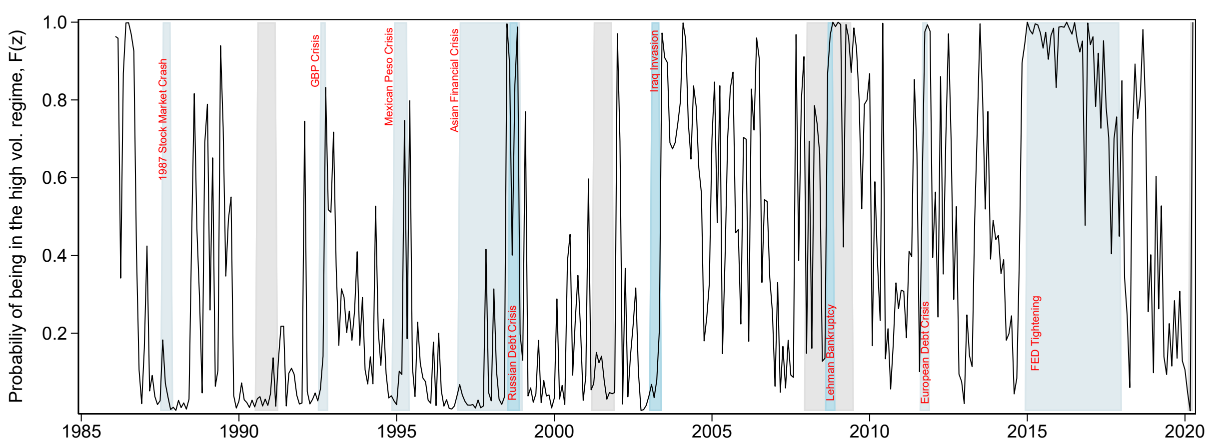

Figure 6: **Conditional exchange rate predictability over different horizons**

This figure presents the exchange rate predictive ability of commodity export prices over different horizons, conditional on the level of FX volatility. The reported conditional slope coefficients are estimated by regressing the exchange rate changes of each commodity currency, computed up to 12 months ahead, on the changes of the corresponding country's commodity price index, which is export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Panels (a) and (b) present results in times of high and low FX volatility, respectively. The panel regression specification includes the control variables, namely the one-month interest rate differential, changes in aggregate FX volatility, changes in funding liquidity (TED spread), changes in aggregate uncertainty in the US market (VIX), and an NBER recession indicator. Reported 95% confidence intervals use standard errors based on Driscoll and Kraay (1998), which are adjusted for serial correlation using the Newey and West (1987) kernel with a bandwidth equal to the forecasting horizon. Section 5.3 presents the econometric specification, and Section 4 describes the controls. The sample consists of monthly observations between January 1985 and April 2020.

(a) High FX volatility regime

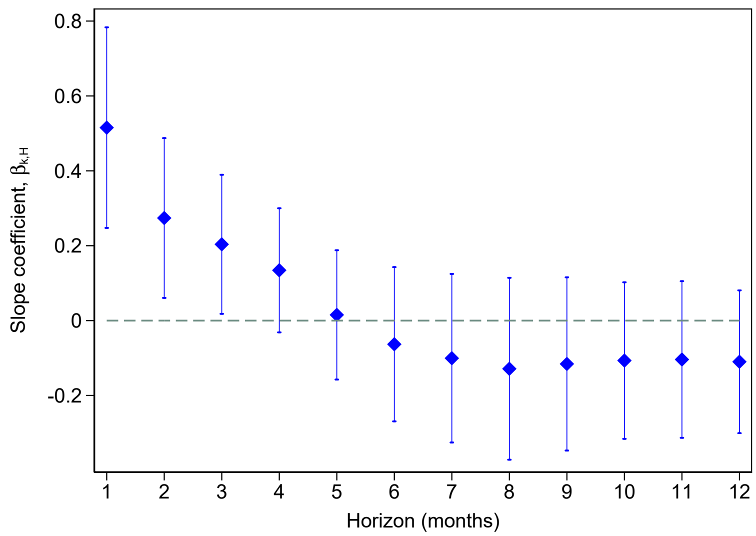

(b) Low FX volatility regime

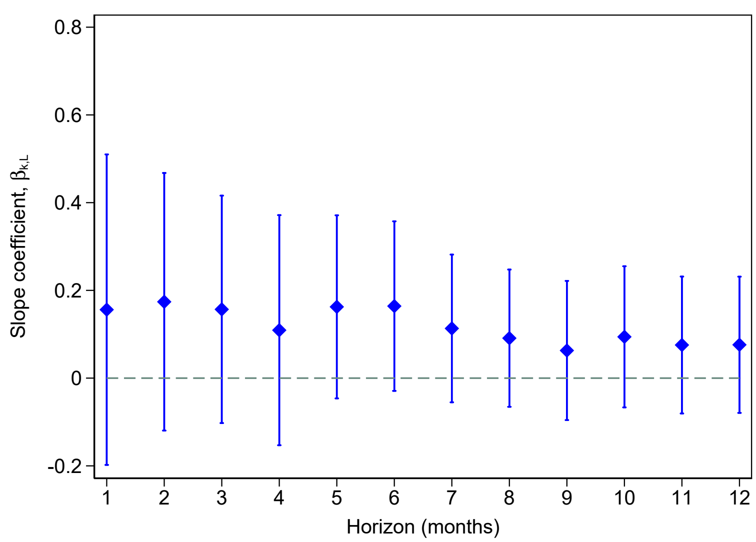

Figure 7: **Commodity price beta and the forward premium**

This figure plots the country-level commodity price beta against its currency's forward premium. Each country's commodity price beta is estimated by regressing the spot exchange rate changes on its commodity price index changes and the dollar factor. Country-specific commodity price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Red diamonds denote commodity price betas that are positive and statistically significant at the 10% level. Country-level commodity price betas are plotted against the average monthly forward premium (in % and annualized). The scatter plot includes a fitted regression line and a 95% confidence interval. The sample consists of monthly observations between January 1985 and April 2020.

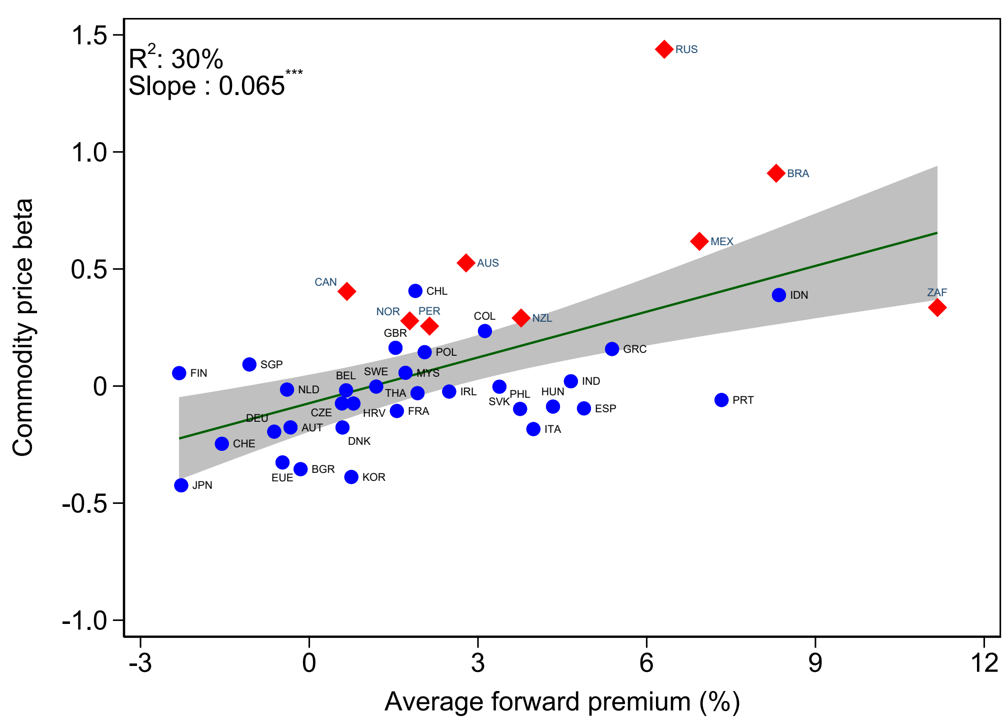

Figure 8: **Predictability by carry trade portfolios**

This figure plots, for each carry trade portfolio, the exchange rate predictability with commodity export prices against the prevalence of commodity currencies in that portfolio. Five equally weighted portfolios are formed every month, with Portfolio 1 (P1) containing the currencies with the lowest interest rates and Portfolio 5 (P5) containing those with the highest. The carry trade is the strategy that is long in P5 and short in P1. Exchange rate predictability is measured with the estimated slope coefficient from the regression of the one-month-ahead spot exchange rate changes for each of the five interest-rate sorted portfolios on the changes in commodity export prices. Commodity currency membership (in %) is the frequency the nine commodity currencies are members of a specific currency portfolio. Membership statistics for each commodity currency are reported in Table 6. Predictive regressions include a full set of controls, which are presented in Section 4. Red diamonds denote estimated slope coefficients that are statistically significant at the 5% level. The scatter plot includes a fitted regression line and a 95% confidence interval. The sample that is used for the predictive regressions and membership calculation consists of monthly observations between January 1985 and April 2020.

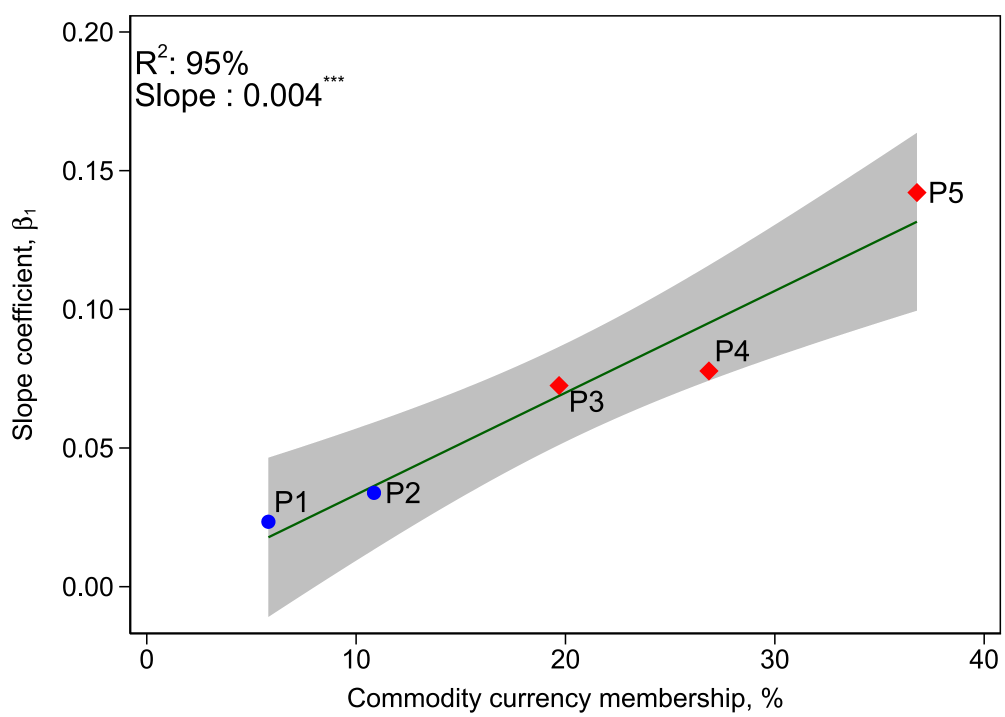

Figure 9: **Out-of-sample predictability relative to daily FX turnover**

This figure reports the relationship between a currency's out-of-sample predictability and its liquidity. We consider the predictive ability of the commodity export prices (CEP) against the random walk (RW) for commodity currencies' exchange rate changes. The degree of out-of-sample predictability is measured by the out-of-sample R^2 statistic of Campbell and Thompson (2008), given by $R^2_{oos} = 1 - (MSE^{CEP}/MSE^{RW})$, where MSE is the mean-squared error (MSE) of a given model. Liquidity is measured by the average daily FX turnover (in billion USD) of each country's currency (calculated using the BIS triennial surveys for the years 2004-2019). Red diamonds denote cases where the CEP model significantly, based on bootstrapped p-values, outperforms the benchmark RW model (see Table 9 for details). The scatter plot includes a fitted regression line and a 95% confidence interval. The sample used for the out-of-sample predictive regressions consists of monthly observations between January 1985 and April 2020.

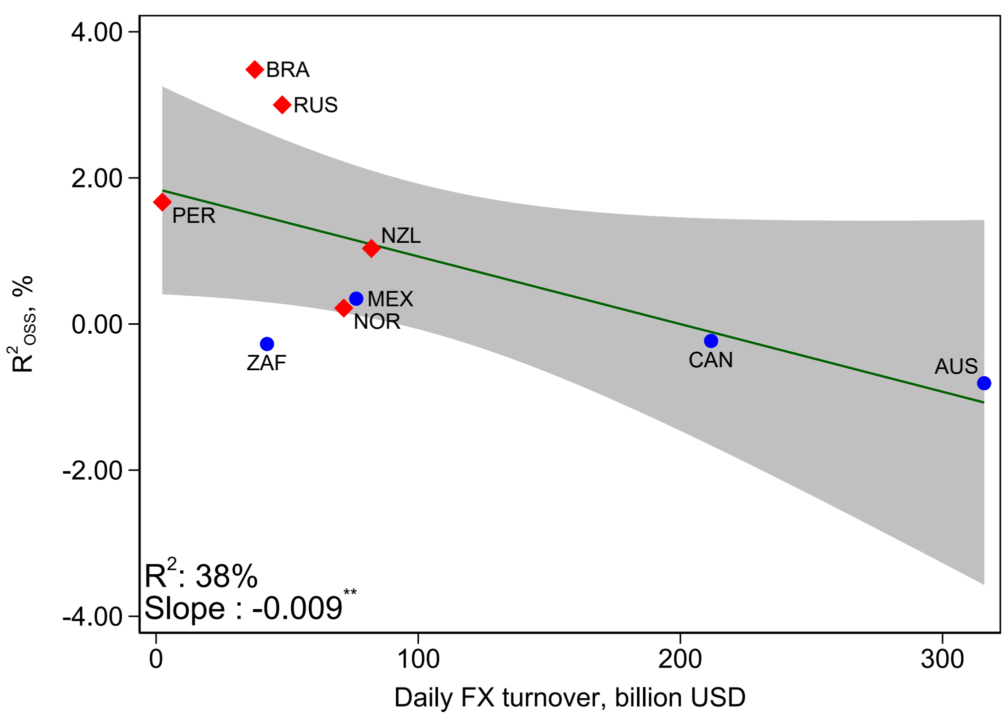

Internet Appendix to

“Commodity Prices and Currencies”

(not for publication)

Abstract

This Internet Appendix presents supplementary material and results not included in the main body of the paper.

September 10, 2023

A Theoretical derivation

A.1 Equilibrium exchange rate

This section derives the equilibrium exchange rate at date 1, i.e., s_1. Note that the optimal demand for agent $i = I, U$ at date 1 is given by

$$x_{i,1} = \frac{\mathbb{E}_{i,1}[s_2] - s_1}{\mathbb{V}_{i,1}[s_2]}, \tag{A.1}$$

while the aggregate demand/supply of the noise trader $x_{N,t}$ is normally distributed with mean zero and volatility σ_N.

Market clearing at each date t imposes the following condition:

$$w_I x_{I,t} + w_U x_{U,t} + x_{N,t} = 0, \tag{A.2}$$

where agents I and U have relative weights given by w_I and w_U, respectively, where $w_I + w_U = 1$.

By market clearing condition at date 1, we have

$$-x_{N,1} = w_I \underbrace{\frac{\mathbb{E}_{I,1}[s_2] - s_1}{\mathbb{V}_{I,1}[s_2]}}_{x_{I,1}} + w_U \underbrace{\frac{\mathbb{E}_{U,1}[s_2] - s_1}{\mathbb{V}_{U,1}[s_2]}}_{x_{U,1}} \tag{A.3}$$

$$= \frac{w_I}{\mathbb{V}_{I,1}[s_2]}\mathbb{E}_{I,1}[s_2] + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\mathbb{E}_{U,1}[s_2] - \left(\frac{w_I}{\mathbb{V}_{I,1}[s_2]} + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\right) s_1, \tag{A.4}$$

and the (log) exchange rate at date 1 then satisfies:

$$s_1 = \underbrace{\left(\frac{w_I}{\mathbb{V}_{I,1}[s_2]} + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\right)^{-1}}_{\equiv \bar{\sigma}_s^2} \left(\frac{w_I}{\mathbb{V}_{I,1}[s_2]}\mathbb{E}_{I,1}[s_2] + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\mathbb{E}_{U,1}[s_2] + x_{N,1}\right) \tag{A.5}$$

$$= \underbrace{\frac{w_I \bar{\sigma}_s^2}{\mathbb{V}_{I,1}[s_2]}}_{\equiv \omega_I}\mathbb{E}_{I,1}[s_2] + \underbrace{\frac{w_U \bar{\sigma}_s^2}{\mathbb{V}_{U,1}[s_2]}}_{\equiv \omega_U}\mathbb{E}_{U,1}[s_2] + \bar{\sigma}_s^2 x_{N,1}, \tag{A.6}$$

where ω_I and $\omega_U = 1 - \omega_L$ are weights reflecting how the belief of agent $i = I, U$ affects the equilibrium exchange rate s_1. These weights are proportional to the agents respective precision $\frac{1}{\mathbb{V}_{I,1}[s_2]}$ and $\frac{1}{\mathbb{V}_{U,1}[s_2]}$. The aggregate degree of uncertainty about exchange rate s_2, denoted by $\bar{\sigma}_s^2$, can be conveniently expressed as follows:

$$\bar{\sigma}_s^2 \quad = \quad \omega_I \mathbb{V}_{I,1}[s_2] + \omega_U \mathbb{V}_{U,1}[s_2] \tag{A.7}$$

and thus reflects the uncertainty about exchange rate s_2, as perceived by the "average" agent.

To verify that $\left(\frac{w_I}{\mathbb{V}_{I,1}[s_2]} + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\right)^{-1} = \bar{\sigma}_s^2$ in Equation (A.5), replace w_I by $\frac{\omega_I \mathbb{V}_{I,1}}{\bar{\sigma}_s^2}$ and w_U by $\frac{\omega_U \mathbb{V}_{U,1}}{\bar{\sigma}_s^2}$, as defined in Equation (A.6). We have

$$\left(\frac{w_I}{\mathbb{V}_{I,1}[s_2]} + \frac{w_U}{\mathbb{V}_{U,1}[s_2]}\right)^{-1} = \left(\frac{\omega_I}{\bar{\sigma}_s^2} + \frac{\omega_U}{\bar{\sigma}_s^2}\right)^{-1} = \left(\frac{\omega_I + \omega_U}{\bar{\sigma}_s^2}\right)^{-1} = \bar{\sigma}_s^2, \tag{A.8}$$

given that $\omega_I + \omega_U = 1$.

A.2 Commodity beta

Here, we derive the commodity price beta discussed in Section 5. The exposure of the second-period (log) exchange rate change Δs_2 to the public news p, denoted by the commodity price beta β, can be expressed as follows:

$$\beta_1 \quad = \quad \frac{\mathbb{COV}[\Delta s_2, p]}{\mathbb{V}[p]} = \frac{\sigma^2}{\sigma^2 + \sigma_\epsilon^2}(1 - \omega_I), \tag{A.9}$$

where the covariance between Δs_2 and p satisfies

$$\begin{aligned}
\mathbb{COV}[\Delta s_2, p] &= \mathbb{COV}\left[\Phi - \omega_I \eta p - \bar{\sigma}_s^2 x_{N,1}, p\right] && \text{(A.10)}\\
&= \mathbb{COV}\left[\Phi - \omega_I \eta(\Phi + \epsilon) - \bar{\sigma}_s^2 x_{N,1}, \Phi + \epsilon\right] && \text{(A.11)}\\
&= \sigma^2 - \omega_I \eta \left(\sigma^2 + \sigma_\epsilon^2\right) && \text{(A.12)}\\
&= \sigma^2 - \omega_I \sigma^2 && \text{(A.13)}\\
&= \sigma^2(1 - \omega_I) && \text{(A.14)}
\end{aligned}$$

using Equation (15) for Δs_2 and $\sigma^2 + \sigma_\epsilon^2 = \frac{\sigma^2}{\eta}$ from Equation (5), while the variance of p equals

$$\begin{aligned}
\mathbb{V}[p] &= \mathbb{V}[\Phi + \epsilon] && \text{(A.15)}\\
&= \sigma^2 + \sigma_\epsilon^2. && \text{(A.16)}
\end{aligned}$$

B Bootstrap algorithm

Our bootstrap algorithm follows Mark (1995b) and Kilian (1999) and imposes the null of no predictability to generate the critical values for our out-of-sample test statistics. This procedure consists of the following steps:

1. Given the sequence of observations for $\{\Delta s_t, \Delta \text{CEP}_t\}$, define the out-of-sample window and generate M out-of-sample forecasts by running the predictive regressions

$$\Delta s_t = \alpha + \beta \Delta \text{CEP}_{t-1} + \varepsilon_{t+1}$$

   both under the null (i.e., $\beta = 0$) and under the alternative, based on an expanding window. Compute the statistic of interest $\widehat{\tau}$.

2. The data generating process for $\{\Delta s_t, \Delta \text{CEP}_t\}$ under the null of a RW model with drift, is assumed to be

$$\begin{aligned} \Delta s_t &= \alpha + u_{1,t} \\ \Delta \text{CEP}_t &= \phi_0 + \phi_1 \Delta \text{CEP}_{t-1} + \ldots + \phi_p \Delta \text{CEP}_{t-p} + u_{2,t}, \end{aligned}$$

   where the lag order p is determined by the Bayesian information criterion. Estimate this specification using the full sample of observations via least-squares, and store the estimates $\widehat{a}$, $\widehat{\phi}_0, \ldots, \widehat{\phi}_p$, and the residual residuals $\widehat{u}_t = (\widehat{u}_{1,t}, \widehat{u}_{2,t})'$.

3. Generate a sequence of pseudo-observations $\{\Delta s_t^\star, \Delta \text{CEP}_t^\star\}$ of the same length as the original data series $\{\Delta s_t, \Delta \text{CEP}_t\}$ as follows:

$$\begin{aligned} \Delta s_t^\star &= \widehat{\alpha} + u_{1,t}^\star \\ \Delta \text{CEP}_t^\star &= \widehat{\phi}_0 + \widehat{\phi}_1 \Delta \text{CEP}_{t-1}^\star + \ldots + \widehat{\phi}_p \Delta \text{CEP}_{t-p}^\star + u_{2,t}^\star, \end{aligned}$$

   where the pseudo-innovation term $u_t^\star = (u_{1,t}^\star, u_{2,t}^\star)'$ is randomly drawn with replacement from the set of observed residuals $\widehat{u}_t = (\widehat{u}_{1,t}, \widehat{u}_{2,t})'$. The initial observations $\left(\Delta \text{CEP}_{t-1}^\star, \ldots, \Delta \text{CEP}_{t-p}^\star\right)'$ are randomly drawn from the actual data. Repeat this step $N_{\text{step}} = 1,000$ times.

4. For each of the N_{step} bootstrap replications, generate M out-of-sample forecasts by running the predictive regressions

$$\Delta s_t^\star = \alpha^\star + \beta^\star \Delta \text{CEP}_{t-1}^\star + u_{1,t}^\star$$

   both under the null and the alternative. Construct the test statistic of interest $\widehat{\tau}^\star$.

5. Compute the one-sided p-value as follows

$$p\text{-value} = \frac{1}{N_{\mathsf{step}}} \sum_{j=1}^{N_{\mathsf{step}}} I(\widehat{\tau}^{\star} > \widehat{\tau}),$$

where $I(\cdot)$ denotes an indicator function, which is equal to 1 when its argument is true and 0 otherwise.

Table A.1: **Commodity price beta and country-level regressions**

This table reports the regression results of estimating each country's commodity price beta, which reflects the country's exchange rate exposure to its commodity export prices. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity price index changes (ΔCEP) and the dollar factor (DOL). Country-level commodity price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Section 3 describes the econometric specification and the controls. The t-statistics are based on White (1980) standard errors. The sample consists of monthly observations between January 1985 and April 2020.

Country	Constant	t-stat	ΔCEP	t-stat	DOL	t-stat	R^2	N
Australia	−0.05	−0.40	0.53	2.96	0.73	8.79	31.63	424
Austria	0.25	2.72	−0.18	−1.80	1.22	29.59	87.28	168
Belgium	0.23	2.62	−0.02	−0.21	1.25	31.54	88.69	168
Brazil	−0.19	−0.78	0.91	3.12	1.16	8.29	45.71	193
Bulgaria	0.07	0.72	−0.35	−1.96	1.09	18.18	73.48	193
Canada	−0.01	−0.12	0.40	3.16	0.39	8.14	26.67	424
Chile	−0.05	−0.25	0.41	1.23	0.90	7.60	41.39	193
Colombia	−0.05	−0.26	0.24	0.89	1.19	12.26	47.73	193
Croatia	0.07	0.62	−0.07	−0.51	1.06	20.56	72.44	193
Czechia	0.23	1.83	−0.07	−0.51	1.28	19.71	62.85	280
Denmark	0.13	1.96	−0.18	−1.98	1.16	35.00	78.95	424
Euro	0.08	0.82	−0.33	−2.14	1.13	19.68	71.35	255
Finland	0.32	0.82	0.06	0.12	0.88	5.05	54.74	24
France	0.20	2.24	−0.11	−1.31	1.20	30.02	87.90	168
Germany	0.25	2.80	−0.19	−2.06	1.26	34.05	88.52	168
Greece	−0.27	−0.72	0.16	0.50	0.82	6.04	30.32	42
Hungary	−0.04	−0.33	−0.09	−0.52	1.45	14.31	65.97	270
India	−0.22	−2.12	0.02	0.17	0.51	9.12	31.20	270
Indonesia	−0.61	−1.16	0.39	0.70	1.03	5.33	9.70	202
Ireland	0.11	1.07	−0.02	−0.24	1.17	30.24	82.78	168
Italy	−0.03	−0.24	−0.18	−1.52	1.07	18.27	68.12	168
Japan	0.21	1.51	−0.42	−2.27	0.61	8.29	18.32	424
Malaysia	−0.13	−0.88	0.06	0.38	0.72	11.02	38.51	197
Mexico	−0.31	−1.98	0.62	2.49	0.61	6.38	27.90	280
Netherlands	0.25	2.80	−0.01	−0.16	1.25	31.89	88.24	168
New Zealand	0.07	0.48	0.29	2.12	0.83	10.87	34.65	424
Norway	−0.02	−0.30	0.28	3.11	1.14	25.71	73.36	424
Peru	0.03	0.31	0.26	2.31	0.32	6.65	30.90	191
Philippines	−0.17	−1.35	−0.10	−0.83	0.46	6.68	17.42	280
Poland	0.01	0.05	0.14	0.77	1.49	14.95	72.17	218
Portugal	−0.14	−1.34	−0.06	−0.65	1.10	23.36	80.83	168
Russia	−0.39	−1.70	1.44	3.32	0.99	6.19	48.01	193
Singapore	0.11	2.08	0.09	1.41	0.48	15.22	55.12	424
Slovakia	−0.03	−0.11	−0.00	−0.01	1.57	10.60	72.17	45
South Africa	−0.50	−2.89	0.34	1.68	0.96	10.43	29.51	422
South Korea	0.04	0.25	−0.39	−1.19	1.05	10.62	50.56	218
Spain	−0.01	−0.10	−0.10	−0.92	1.14	21.82	74.36	168
Sweden	−0.02	−0.19	−0.00	−0.03	1.13	28.19	68.66	424
Switzerland	0.23	2.53	−0.25	−1.81	1.15	22.54	66.09	422
Thailand	0.01	0.08	−0.03	−0.28	0.66	6.93	21.92	280
UK	0.02	0.24	0.16	1.33	0.84	14.57	48.39	424

Table A.2: **Exchange rate predictability for commodity currencies – early period**

This table reports panel regression results on the exchange rate predictive ability of commodity export prices, using the first half of the sample period. The dependent variable is the one-month-ahead exchange rate change of each commodity currency against the US dollar. The variable ΔCEP is the change of the corresponding country's commodity price index, which is export-weighted, rebalanced monthly, and standardized. Model (1) is the univariate regression, Model (2) controls for the one-month interest rate differential (IRD) of each currency relative to the US dollar, Model (3) includes the dollar factor (DOL), Model (4) includes changes in aggregate currency volatility (FX Vol), while Models (5) and (6) include changes in funding liquidity (TED) and in aggregate uncertainty in the US market (VIX), respectively. Model (7) includes an indicator variable equal to one during NBER recessions and zero otherwise. All specifications include currency fixed effects. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the econometric specification and the controls. The sample consists of monthly observations between January 1985 and March 2003.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔCEP	0.301^{**}	0.289^{**}	0.277^{**}	0.263^{**}	0.296^{**}	0.286^{**}	0.264^{**}
	(0.123)	(0.120)	(0.117)	(0.117)	(0.118)	(0.119)	(0.118)
IRD		0.382^{***}	0.375^{***}	0.373^{***}	0.380^{***}	0.385^{***}	0.382^{***}
		(0.134)	(0.136)	(0.138)	(0.138)	(0.141)	(0.141)
DOL			0.031	0.036	0.053	0.049	0.049
			(0.057)	(0.056)	(0.057)	(0.057)	(0.058)
ΔFX Vol				0.208	0.145	0.099	0.102
				(0.143)	(0.170)	(0.169)	(0.170)
ΔTED					0.654	0.517	0.514
					(0.561)	(0.559)	(0.554)
ΔVIX						0.044^{**}	0.044^{**}
						(0.018)	(0.018)
NBER							−0.414
							(0.293)
Observations	1161	1161	1161	1161	1099	1099	1099
R^2 %	0.00	1.01	0.99	1.21	1.90	2.39	2.46
Currencies	6	6	6	6	6	6	6
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A.3: **Exchange rate predictability for commodity currencies – recent period**

This table reports panel regression results on the exchange rate predictive ability of commodity export prices, using the second half of the sample period. The dependent variable is the one-month-ahead exchange rate change of each commodity currency against the US dollar. The variable ΔCEP is the change of the corresponding country's commodity price index, which is export-weighted, rebalanced monthly, and standardized. Model (1) is the univariate regression, Model (2) controls for the one-month interest rate differential (IRD) of each currency relative to the US dollar, Model (3) includes the dollar factor (DOL), Model (4) includes changes in aggregate currency volatility (FX Vol), while Models (5) and (6) include changes in funding liquidity (TED) and in aggregate uncertainty in the US market (VIX), respectively. Model (7) includes an indicator variable equal to one during NBER recessions and zero otherwise. All specifications include currency fixed effects. Standard errors, in parentheses, are based on Driscoll and Kraay (1998) and are adjusted for serial correlation using the Newey and West (1987) kernel with optimally-selected bandwidth. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. Section 4 describes the econometric specification and the controls. The sample consists of monthly observations between April 2003 and April 2020.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔCEP	0.394**	0.400**	0.395**	0.390**	0.394**	0.417**	0.412***
	(0.176)	(0.176)	(0.162)	(0.154)	(0.162)	(0.165)	(0.155)
IRD		0.142	0.142	0.142	0.128	0.107	0.114
		(0.169)	(0.169)	(0.172)	(0.185)	(0.185)	(0.174)
DOL			0.006	0.002	−0.005	−0.087	−0.090
			(0.099)	(0.090)	(0.086)	(0.108)	(0.109)
ΔFX Vol				−0.051	0.008	0.102	0.115
				(0.332)	(0.302)	(0.342)	(0.345)
ΔTED					−0.741	−0.435	−0.447
					(1.691)	(1.635)	(1.647)
ΔVIX						−0.080	−0.081
						(0.073)	(0.073)
NBER							−0.293
							(1.237)
R^2 %	0.87	0.89	0.83	0.79	0.95	1.72	1.72
Observations	1806	1805	1805	1805	1805	1805	1805
Currencies	9	9	9	9	9	9	9
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A.4: **Predictability of the carry trade portfolios**

This table presents results on the predictability of the one-month-ahead carry trade portfolios with commodity export prices. Five equally weighted portfolios are formed every month, with portfolio 1 (P1) containing the currencies with the lowest interest rates and portfolio 5 (P5) containing those with the highest. The benchmark predictor is $\Delta\overline{\text{CEP}}$, which is the average (equally-weighted) return across each commodity country's commodity export price index. All specifications include controls, which are presented in Section 4. The table also reports the average frequency commodity currencies (CC) are members of a specific currency portfolio (CC membership %). The standard errors in parentheses are adjusted using the Newey and West (1987) kernel where the bandwidth is selected optimally. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample consists of monthly observations between January 1985 and April 2020.

	(1)	(2)	(3)	(4)	(5)
$\Delta\overline{\text{CEP}}$	0.023	0.034	0.073^{**}	0.078^{**}	0.142^{***}
	(0.032)	(0.035)	(0.032)	(0.034)	(0.039)
Constant	0.080	0.084	0.031	−0.121	−0.237
	(0.115)	(0.122)	(0.111)	(0.119)	(0.151)
R^2 %	0.51	-0.45	0.99	1.28	2.69
Observations	410	410	410	410	410
Portfolio	P1	P2	P3	P4	P5
CC membership %	5.81	10.85	19.69	26.85	36.79

Table A.5: **Descriptive statistics of the carry trade without commodity currencies**

This table presents key characteristics of the carry trade portfolios, excluding commodity currencies. Five equally weighted portfolios are formed every month, with portfolio 1 (P1) containing the currencies with the lowest interest rates and portfolio 5 (P5) containing those with the highest. The carry trade is the strategy that is long in P5 and short in P1. We present summary statistics (mean, standard deviation, skewness, and Sharpe ratio) of the monthly excess returns of the five portfolios and the carry trade. The mean and standard deviation are annualized. IRD is the average annualized one-month interest rate differential relative to the US dollar. ΔS denotes the average annualized monthly appreciation of portfolio member currencies relative to the US dollar. All portfolios exclude the nine commodity currencies as per the definition in Section 3. Portfolios are rebalanced each month and a total of 32 currencies are considered, however, the number of available currencies for portfolio construction varies between periods depending on data availability. The sample consists of monthly observations between January 1985 and April 2020.

	P1	P2	P3	P4	P5	P5-P1
Mean (%)	−0.525	0.978	1.805	2.708	5.808	6.333
Standard deviation (%)	8.685	9.036	9.387	9.065	11.215	8.459
Sharpe ratio	−0.060	0.108	0.192	0.299	0.518	0.749
Skewwess	0.100	−0.230	−0.300	−0.965	−0.301	−0.394
IRD	−2.119	−0.425	0.863	2.829	6.773	8.893
ΔS	1.595	1.402	0.941	−0.122	−0.965	−2.560

Table A.6: **Carry trade return predictability – alternative commodity price indexes**

This table presents results on the predictability of the one-month-ahead carry trade return with different global indexes of commodity prices. Five equally weighted portfolios are formed every month, with portfolio 1 (P1) containing the currencies with the lowest interest rates and portfolio 5 (P5) containing those with the highest. The analysis focuses here on the performance of the investment leg of the carry trade, P5. Model (1) reports the benchmark results using $\Delta\overline{\text{CEP}}$ (the average, equally-weighted, return across each country's commodity export price index). Models (2)-(4) report results using ΔCRB, ΔGSCI, and ΔOil, respectively, to forecast spot exchange rate changes. ΔCRB, ΔGSCI, and ΔOil are the monthly percentage change in the Commodity Research Bureau (CRB) index, Goldman Sachs commodity index (GSCI), and West Texas Intermediary crude oil price respectively. Models (5) reports predictability results using $\Delta\overline{\text{CEP}}$ and the alternative commodity indexes jointly. All specifications include controls, which are presented in Section 4. The standard errors in parentheses are adjusted using the Newey and West (1987) kernel where the bandwidth is selected optimally. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively. The sample consists of monthly observations between January 1985 and April 2020.

	(1)	(2)	(3)	(4)	(5)
$\Delta\overline{\text{CEP}}$	0.142^{***}				0.136^{**}
	(0.039)				(0.057)
ΔCRB		0.141^{***}			0.058
		(0.055)			(0.062)
ΔGSCI			0.058^{**}		0.052
			(0.027)		(0.074)
ΔOil				0.022	−0.036
				(0.014)	(0.039)
Constant	−0.237	-0.255^{*}	−0.237	−0.226	-0.255^{*}
	(0.151)	(0.155)	(0.154)	(0.153)	(0.153)
R^2 %	2.69	1.51	1.21	0.60	2.58
Observations	410	410	410	410	410
Portfolio	P5	P5	P5	P5	P5

Figure A.1: **Saudi Arabia riyal performance over time**

The figure shows the performance of the Saudi Arabia riyal (SAD) over the long run. The exchange rate is expressed as USD per unit of the SAD, i.e., a decrease in the exchange rate implies a depreciation of SAD. Reported data statistics on commodity share of total exports are from the United Nations Comtrade database. Commodity firm market capitalization share is calculated using market capitalization data from Datastream. Data on natural resources rents as a percentage of GDP are from the World Bank. Data on commodity share of total government revenue are from the Federal Reserve Bank of St. Louis. The reported economic shares are time series averages. The correlation with the WTI Intermediary oil price is calculated, using monthly data, over the period between January 1993 (first available data point) and April 2020.

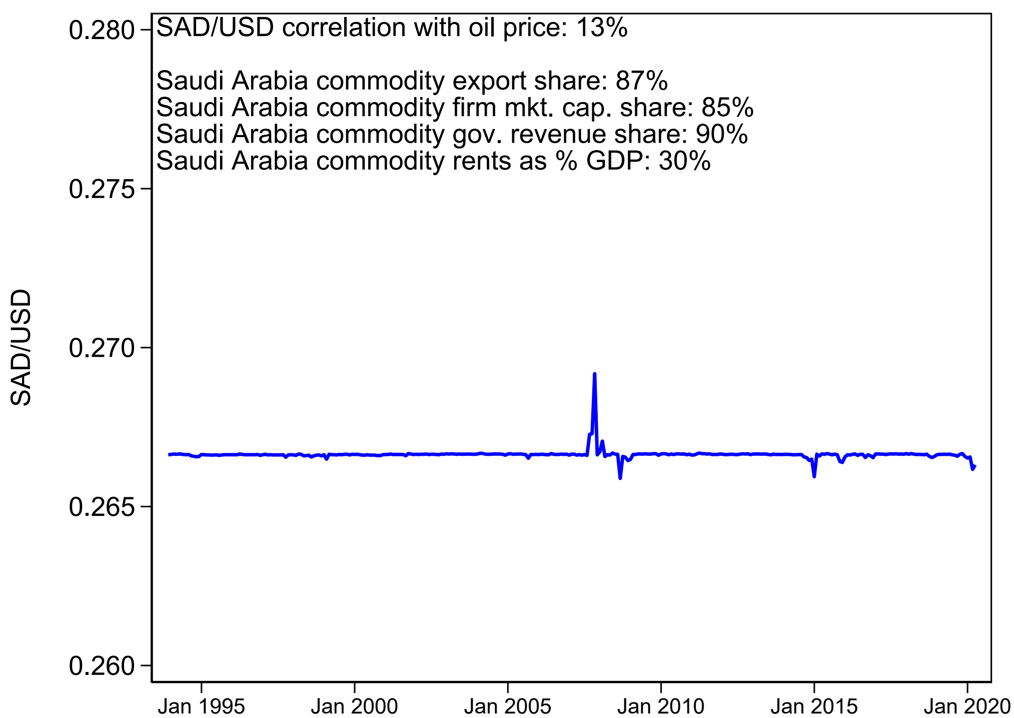

Figure A.2: **Commodity price beta – without the dollar factor**

This figure illustrates each country's commodity price beta, which captures the exchange rate exposure to commodity export prices obtained without controlling for the dollar factor. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity export price index changes. Country-level commodity export price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Countries are ranked according to their commodity price beta. Red diamonds denote commodity price betas that are positive and statistically significant at the 10% level. Plot whiskers represent the 90% confidence intervals. The sample consists of monthly observations between January 1985 and April 2020.

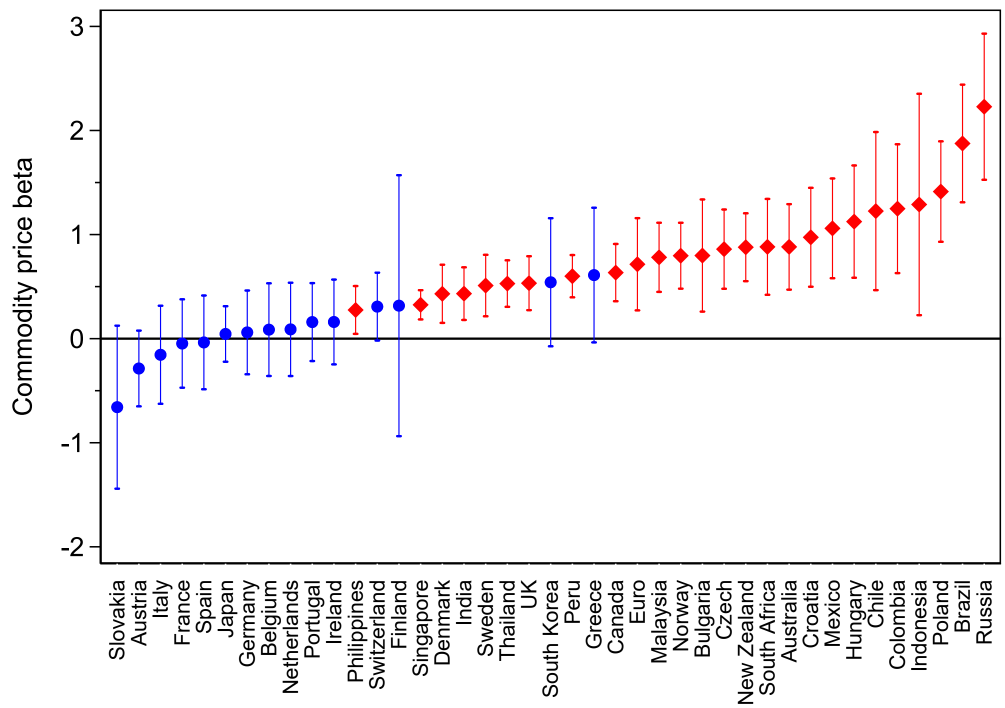

Figure A.3: **Commodity price beta – with dollar and market factors**

This figure illustrates each country's commodity price beta, which captures the exchange rate exposure to commodity export prices obtained by controlling for the dollar factor and the US equity market return. Commodity price betas are estimated by regressing a country's spot exchange rate changes on the corresponding commodity price index changes, the dollar factor, and the Fama/French market return. Country-level commodity export price indexes are export-weighted, rebalanced monthly, and standardized. All exchange rates are against the US dollar. Countries are ranked according to their commodity price beta. Red diamonds denote commodity price betas that are positive and statistically significant at the 10% level. Plot whiskers represent the 90% confidence intervals. The sample consists of monthly observations between January 1985 and April 2020.

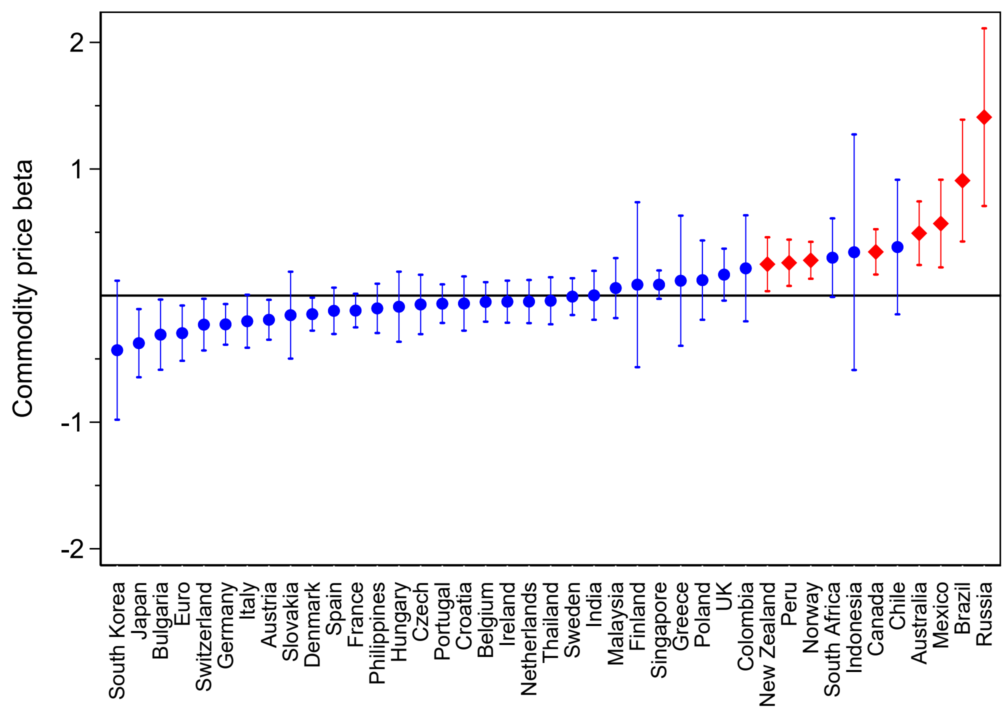

Figure A.4: **Unconditional exchange rate predictability – excluding currencies**

This figure presents the exchange rate predictive ability of commodity export prices when sequentially excluding currencies from the original data sample. The reported slope coefficient is estimated in a panel by regressing the exchange rate changes of each commodity currency, dropping one country at a time, on the changes of the corresponding country's commodity export price index, which is export-weighted and rebalanced monthly. All exchange rates are against the US dollar. The panel regression specification includes the control variables, namely the one-month interest rate differential, the dollar factor, changes in aggregate FX volatility, changes in funding liquidity (TED spread), changes in aggregate uncertainty in the US market (VIX), and an NBER recession indicator. Reported 95% confidence intervals use standard errors based on Driscoll and Kraay (1998), which are adjusted for serial correlation using the Newey and West (1987) kernel with optimal bandwidth. The gray dashed line indicates the slope coefficient obtained with the full sample (Table 2, Column 7). The sample consists of monthly observations between January 1985 and April 2020.