

Business Cycles and Currency Returns*

Steven J. Riddiough[†]
University of Melbourne

Lucio Sarno[‡]
Cass Business School and CEPR

This Version: December 2017

Abstract

We find a strong link between currency returns and the relative strength of the business cycle. Buying currencies of strong economies and selling currencies of weak economies generates high returns in both the cross-section and time-series of countries. These returns stem primarily from spot exchange rate predictability, are uncorrelated with common currency strategies, and cannot be understood using traditional risk factors. We also show that a business cycle factor implied by our results is priced in a broad cross-section of currency excess returns. These results contrast with a vast literature that detects no linkages between currency fluctuations and macroeconomic variables.

Keywords: exchange rates; currency risk premium; business cycles.

JEL Classification: F31, G12, G15.

*We are grateful for comments and suggestions to Antje Berndt, Ric Colacito, Magnus Dahlquist, Pasquale Della Corte, Antonio Gargano, Federico Gavazzoni, Michael Halling, Jianlei Han, Kai Li, Matteo Maggiori, Nelson Mark, Thomas Maurer, Michael Moore, Anella Munro, Dagfinn Rime, Maik Schmeling, Thomas Stolper, Andrea Tamoni, Raman Uppal, Geoff Warren, Mungo Wilson, Josef Zechner, Qi Zeng and to other participants at presentations held at: 2017 China International Conference in Finance; 2017 Wellington Finance Summit; 6th Workshop on ‘Financial Determinants of Exchange Rates,’ Bank of England; 2016 Australian National University Annual Summer Camp; the 9th INQUIRE Business School Workshop, Cass Business School; BI Norwegian Business School, Oslo; University of Geneva; University of Melbourne. We also thank Sidharth Dube, Harvey Huang and Rachel Quan for excellent research assistance. Steven Riddiough gratefully acknowledges financial support from the Faculty of Business and Economics, University of Melbourne. The paper was the recipient of the CFA Institute Best Paper Prize 2017, FIRN Annual Conference, and the Runner-up WINNER (Wealth Innovation, Neuro- and Entrepreneurial Results) Best Paper Award (funded by ZZ Vermögensverwaltung and POK Phringer Privatstiftung in collaboration with WU Vienna University of Economics and Business). All errors are ours.

[†]Steven Riddiough is with the Faculty of Business and Economics, University of Melbourne. Email: steven.riddiough@unimelb.edu.au.

[‡]Lucio Sarno is with Cass Business School, City, University of London and CEPR, Email: luccio.sarno@city.ac.uk.

1 Introduction

It is a common perception that fluctuations in exchange rates cannot be satisfactorily explained, never mind predicted, by macroeconomic fundamentals either in- or out-of-sample. This is especially evident at short and intermediate horizons from one month to a year that are relevant to most investors (Meese and Rogoff, 1983; Engel, Mark, and West, 2007; Rossi, 2013). This macroeconomic ‘disconnect’ is puzzling since exchange rate fluctuations should be driven by an underlying macroeconomic process. In the same way that the return on a company’s stock is theoretically linked to the company’s fundamentals, the return on a country’s currency should be a function of its underpinning economic fundamentals. Yet we have a limited understanding of whether or how currency returns are related to macroeconomic conditions across countries.

In this paper, we focus on the broadest measure of aggregate macroeconomic conditions constituting a key building block in theoretical models of exchange rates – the business cycle – and provide new evidence suggesting that business cycles are a key driver and powerful predictor of currency excess returns and spot exchange rate fluctuations. The business cycle is a crucial component of present value models of exchange rate determination using Taylor rule reaction functions, in which the currencies of economies where current output is above potential are expected to appreciate (Engel and West, 2005). Moreover, in the macro-finance literature the existence of a relationship between business cycles and currency returns is a necessary condition for risk-based models to have empirical validity (Cochrane, 2017).¹

Our empirical approach moves away from traditional forecasting of *bilateral* exchange rate movements using time-series regression analysis, which has been the common focus in most of the earlier literature.² By contrast, we focus on *relative* business cycles using a portfolio approach that has proven successful in recent studies shedding new light on exchange rate determination (Lustig and Verdelhan, 2007; Lustig, Roussanov, and Verdelhan, 2011; Menkhoff et al., 2012b, 2017; ?; Verdelhan, 2017; Lustig and Richmond, 2017). This approach allows us to gather information in both the cross-section and time-series of countries’ business cycles to view predictability through the lens of an international investor trying to exploit a potential link between business cycles and exchange rates. Such an investor cares principally about currency excess returns and evaluates predictability in a multi-currency portfolio setting on the basis of economic metrics. In our empirical analysis we address the following set of questions: do business cycles predict currency returns in a portfolio setting? If so, could an investor earn positive returns from an investment strategy capitalizing on this predictability? And can the returns be understood as compensation for risk?

We split our empirical analysis into two parts. In the first part we explore predictability using the full sample of data to measure business cycles. We do so in light of the earlier evidence in the literature finding macro-fundamentals cannot predict exchange rates even when the econometrician has perfect foresight of macroeconomic outcomes (Meese and Rogoff, 1983; Cheung, Chinn, and Garcia Pascual, 2005). In the second part, we move entirely out of sample to avoid revised or forward looking data, taking the perspective of a currency investor to consider the investment performance, diversification benefits, and risk characteristics associated with currency return predictability.

¹The business cycle is a common feature of macro-finance models of asset pricing with endogenous risk premia, which argues in various ways that the spillovers of business cycles from one country to another make some currencies safer or riskier than others (see, e.g., Martin, 2013; Hassan, 2013; Tran, 2013; Ready, Roussanov, and Ward, 2016; Richmond, 2016). The earlier work of Lustig and Verdelhan (2007) considers a link between consumption growth and currency excess returns. However, our contribution in this paper is entirely empirical and we do not attempt to test directly a specific theoretical model, rather our objective is to examine the existence of an empirical link between business cycles and currency returns.

²See, e.g. Mark (1995), Molodtsova, Nikolsko-Rzhevskyy, and Papell (2008), Rossi (2013).

Throughout the paper we use the output gap as our measure of business cycle conditions. However, the output gap is a common macroeconomic measure of the aggregate state of the economy, defined as the percentage deviation in output from its long-run trend. The output gap is not directly observable, and so we measure it using several commonly adopted methods in the literature, including the Hodrick-Prescott (1980, HP) filter, Baxter-King (1999) filter, the quadratic time trend used by Clarida, Gali, and Gertler (1998), and the linear projection method recently introduced by Hamilton (2017).³ Using monthly data from 1983 to 2016 for a cross-section of 27 countries, we find that sorting currencies into portfolios on the basis of output gaps generates a monotonic increase in excess returns as we move from the portfolios of the weakest economies’ currencies – i.e. those with current output most below potential – to the portfolios of the strongest economies’ currencies – i.e. those with current output most above potential. Furthermore, we find that a 1/N time-series portfolio strategy that takes long positions in countries with output gaps above the United States and short positions in countries with output gaps below the United States also generates impressive predictability. In both cases, the predictive power for currency excess returns is driven almost entirely by the spot exchange rate component rather than interest rate differentials. This result stands in contrast to a large literature in empirical exchange rate modelling, which for decades has struggled to find evidence of a connection between macro variables and exchange rates.

Importantly, our approach is not a regurgitation of existing findings in the literature. Sorting currencies by output gaps is not equivalent, for example, to the currency carry trade that involves sorting currencies by their nominal interest rates, as implemented by, e.g., Lustig, Roussanov, and Verdelhan (2011). We highlight the intuition for this point in Figure 1, using two common carry trade currencies – the Australian dollar and Japanese yen. The interest rate differential is highly persistent in the data and has been consistently positive between Australia and Japan in recent decades. A carry trade investor would thus have always been long the Australian dollar and short the Japanese yen. In contrast the output gap differential varies substantially over time, and thus an output-gap investor would have taken long and short positions in both the Australian dollar and Japanese yen as their relative business cycles fluctuated. Moreover, while we document a strong factor structure in currency portfolio returns sorted by output gaps, we find the second principal component is orthogonal to the analogous ‘*Slope*’ factor documented by Lustig, Roussanov, and Verdelhan (2011).

Our findings suggest a strong predictive link from business cycles to currency returns, and raise questions as to why our results differ from those in the long-standing international macroeconomics literature. The answers lie in both the time-series and cross-section. Our portfolio results from the time-series indicate that high output gap currencies *usually* appreciate – the output-gap provides a better-than-chance prediction of future exchange rate movements and the size of the return increases with the output gap. But the high volatility of exchange rates

³Hamilton (2017) provides a quantitative analysis of the main drawbacks of the HP filter and suggests an alternative procedure for detrending output and measuring the output gap to achieve the objectives of the HP filter without the drawbacks. Although the focus is on improving over the HP filter out-of-sample, Hamilton’s analysis and criticisms are relevant for all other filters commonly used in this literature. Therefore, we use the Hamilton procedure in our out-of-sample analysis, implementing the procedure recursively conditioning only on data available at the time of sorting. While the output gap is a common measure of business cycle conditions in the macroeconomics literature, it has received comparatively little attention in financial economics. Cooper and Priestley (2009) provide a notable exception, finding that the output gap can help predict future stock returns for the United States and other G7 countries both in-sample and out-of-sample. In international macroeconomics, Molodtsova, Nikolsko-Rzhevskyy, and Papell (2008) show that ‘Taylor rule’ models that incorporate output gap and inflation information display predictive power for spot exchange rate changes in time series regressions for three major exchange rates, although this result was challenged by Rogoff and Stavrageva (2008) who argue against the robustness of the predictability across different subsample periods.




Figure 1: Australia and Japan: Interest Rate and Output Gap Spreads

The figure plots the interest-rate- and output-gap-spread between Australia and Japan. When a series is above the origin, it indicates the Australian value was higher (i.e., either a higher interest rate or output gap).

relative to macro-fundamentals makes it difficult to outperform a random walk on statistical grounds. The second reason relates to a cross-sectional ‘level’ factor (the average output gap relative to the United States), which leads to coefficient instability. In periods when the United States has a relatively low output gap, all currencies tend to appreciate against the U.S. dollar – even those with output gaps *below* the United States. In contrast, currencies of weak economies depreciate substantially against the U.S. dollar when the United States is relatively strong compared to all other countries. The time-series literature typically fails to account for these state-dependent coefficients, while the cross-section is immune to them.

The in-sample analysis uses the entire sample to calculate output gaps. While this is useful for studying the relationship between relative economic conditions and exchange rates, the use of forward-looking information raises questions as to whether the relationship is economically valuable. Moving out-of-sample and using real-time output data from 1999 to 2016, we indeed find that our results are qualitatively identical. The data we use mimics the information set available to investors during this period and thus sorting is conditional only on information available at that time. The out-of-sample Sharpe ratios for the high-minus-low cross-sectional strategy, which we denote as GAP_{CS} , reaches 0.72 (without transaction costs) and 0.60 (net of costs). In the time-series, the strategy which goes long (short) currencies issued by countries with output gaps above (below) the United States, denoted as GAP_{TS} , generates a Sharpe ratio of 0.65 (0.57 after costs). The two strategies exhibit a correlation of around 40% and

thus the investment performance increases further once these strategies are combined.⁴ The results imply that currency excess returns are higher for stronger economies, i.e. those in a more favorable state of the business cycle, and are robust to various ways of constructing our portfolio strategy. For example, the results hold when assigning linear or rank weights (and thus trading all currencies simultaneously), which reassures us that the results are not driven by a few outlier currencies but apply generally to the broader cross section.

The out-of-sample performance, as with the in-sample results, stems mainly from the predictability of spot exchange rates rather than from interest rate differentials. That is, currencies of strong economies tend to appreciate and those of weak economies tend to depreciate over the subsequent month. This feature makes the returns from exploiting business cycle information different from the returns delivered by most canonical currency investment strategies, and most notably distinct from carry (which generates a negative exchange rate return). Indeed, we find the time-series correlations between each of GAP_{CS} and GAP_{TS} strategies and the currency carry trade are essentially zero, and the correlations with other canonical currency investment strategies – such as dollar carry, momentum, and value strategies – are also close to zero.⁵ The observed predictability of spot exchange rates is a rare finding in this literature and accounts for the lack of correlation with standard currency strategies. This apparent lack of correlation also implies that the output gap strategies provide useful diversification gains to an investor who adds one or more of the strategies to a conventional menu of currency portfolios, and we quantify these gains in the empirical analysis.

We then ask whether the returns of output-gap-sorted portfolios reflect compensation for risk. We address this question by testing the pricing power of conventional risk factors using standard linear asset pricing models, and find no evidence that the returns can be explained using various factors that have been used to explain variation in the cross section of currency returns. Finally, we discuss our empirical findings in the context of the existing theoretical literature and consider the possibility that business cycles proxy for a priced state variable as implied by many macro-finance models of currency premia. To do so, we consider the pricing power of a business cycle factor, taken to equal the returns on the GAP_{CS} strategy, and test whether it is priced in the cross section of currencies. We find that the pricing power of the factor is strong and not confined to portfolios sorted on output gaps, extending to other popular currency cross-sections, including portfolios sorted on carry (interest rate differentials), momentum and value. We thus provide the first tentative evidence in support of a large theoretical macro-finance literature linking business cycles to currency premia.

The remainder of the paper is as follows. Section 2 describes the data and defines the currency portfolios studied in the empirical analysis. Section 3 reports in-sample results on the predictive information content of business cycles for currency excess returns, whereas Section 4 presents out-of-sample results on the performance and diversification gains from incorporating information on relative business cycles. Section 5 reports the results for asset pricing tests designed to explore whether the returns to output-gap-sorted portfolios can be understood as compensation for risk, and whether a business cycle risk factor implied by our results is priced

⁴Moskowitz, Ooi, and Pedersen (2012) provide evidence of time-series strategy performance, while Baz et al. (2015) consider combinations of time-series and cross-sectional strategies across asset classes using carry, value and momentum signals.

⁵Lustig, Roussanov, and Verdelhan (2014) propose a dollar carry trade strategy which trades a basket of currencies against the U.S. dollar on the basis of the average forward discount relative to the United States. Their strategy is uncorrelated with the standard carry trade, and the returns compensate U.S. investors for taking on aggregate risk by shorting the dollar in bad times, when the U.S. price of risk is high. Our GAP_{CS} strategy is distinct conceptually – in that it directly sorts on relative business cycles across countries rather than on interest rate (forward discount) information – and empirically we document that the returns of the dollar carry trade are uncorrelated with the returns of our GAP_{CS} strategy.

in the cross section of currencies. Section 6 concludes.

2 Data and Currency Portfolios

This section describes the main data employed in the empirical analysis as well as the construction of output gaps both in- and out-of-sample.

2.1 Data on Spot and Forward Exchange Rates

We collect daily spot and 1-month forward exchange rates vis-à-vis the U.S. dollar from Barclays and Reuters via Datastream. The empirical analysis uses monthly data obtained by sampling end-of-month rates from October 1983 to January 2016. Our sample comprises 27 countries: Australia, Austria, Belgium, Brazil, Canada, Chile, Czech Republic, Germany, Finland, France, Iceland, Ireland, Italy, Japan, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, South Korea, Spain, Sweden, Switzerland, Turkey, United Kingdom, and the United States.⁶

2.2 Currency Excess Returns

We define spot and forward exchange rates at time t as $Spot_t$ and Fwd_t . Exchange rates are defined as units of U.S. dollars per unit of foreign currency such that an increase in $Spot_t$ indicates an appreciation of the foreign currency. The excess return on buying a foreign currency in the forward market at time t and selling in the spot market at time $t + 1$ is computed as

$$RX_{t+1} = \frac{(Spot_{t+1} - Fwd_t)}{Spot_t}, \quad (1)$$

which is equivalent to the spot exchange rate return minus the forward premium

$$RX_{t+1} = \frac{Spot_{t+1} - Spot_t}{Spot_t} - \frac{Fwd_t - Spot_t}{Spot_t}. \quad (2)$$

According to the Covered Interest Parity (CIP) condition, the forward premium approximately equals the interest rate differential $(Fwd_t - Spot_t) / Spot_t \simeq i_t - i_t^*$, where i_t and i_t^* represent the U.S. and the foreign riskless rates respectively, over the maturity of the forward contract. Since CIP generally holds closely in the data at low frequency (e.g., Akram, Rime, and Sarno, 2008), the currency excess return is approximately equal to the exchange rate return (i.e., $(Spot_{t+1} - Spot_t) / Spot_t$) plus the interest rate differential relative to the United States (i.e., $i_t^* - i_t$). As a matter of convenience, throughout this paper we refer to $fd_t = (Spot_t - Fwd_t) / Spot_t = i_t^* - i_t$ as the forward discount or interest rate differential relative to the United States.

2.3 The Output Gap and Data on Economic Activity

The output gap is defined as the logarithm of the difference between actual (y_t) and ‘potential’ (\bar{y}_t) output ($gap_t = y_t - \bar{y}_t$). A country’s potential output is not directly observable, and it therefore needs to be estimated. Numerous statistical methods have been proposed to measure potential output \bar{y}_t , with the principal aim being to decompose a country’s output into the trend and cyclical components. The trend component can be viewed as the economy’s natural

⁶Full details of the Datastream mnemonics we use and sample period available for each currency pair is provided in the Internet Appendix.

or potential growth path, from which growth cyclically deviates. The cyclical component is thus a measure of short-term deviations and serves as our empirical proxy for the output gap.

To measure economic activity we use data on (log) industrial production obtained from the OECD’s *Original Release Data and Revisions Database*. The database provides monthly ‘vintages’ of data, allowing researchers to know exactly what information was available to market participants on particularly macroeconomic variables at the end of each given month and is thus free of revisions and forward looking information. For example, in December 1999, the available U.S. industrial production data ran from January 1960 until October 1999. In our out-of-sample analysis, we thus assume that an investor at the end of December 1999 would have no knowledge of either the November or December 1999 values and is thus able to condition investment decisions on information available up until October 1999. This assumption inherently biases our out-of-sample results downwards, given the availability of other relevant information at the time. Nonetheless, our out-of-sample results can provide a lower bound for assessing the usefulness of relative business cycle information across countries.

2.3.1 In-sample

For the in-sample analysis we use the April 2016 vintage of data. The full series of monthly industrial production data begin at various dates across countries. The earliest start date is January 1960, and we end the sample in January 2016 to coincide with the last industrial production data point available for most countries in our sample. We estimate output gaps using various statistical techniques to extract a cyclical component from macroeconomic data: (i) the linear projection method of Hamilton (2017), (ii) the Hodrick-Prescott (HP) (1980, 1997) filtered output gap, (iii) the Baxter-King (1999) filter, and the quadratic trend specification used by Clarida, Gali, and Gertler (1998). We provide further details of the parameters and functional forms of these statistical techniques in the Internet Appendix.

2.3.2 Out-of-sample

For the out-of-sample analysis, we use the full set of monthly industrial production vintages from December 1999 until January 2016.⁷ Each monthly vintage records the industrial production data available to an investor in that particular month. In the out-of-sample analysis, we construct output gap estimates using each monthly vintage in turn, applying the linear projection method of Hamilton (2017), and therefore the resulting estimate at time t is conditioned only on information available at t .

The linear projection methodology requires the estimation of the following time-series regression:

$$y_{i,t} = \alpha_i + \beta_{i,1}y_{i,t-24} + \beta_{i,2}y_{i,t-25} + \beta_{i,3}y_{i,t-26} + \beta_{i,4}y_{i,t-27} + \varepsilon_{i,t} \quad (3)$$

where y_t is the (log) value of industrial production for country i available at time t . We regress time- t values on their corresponding values from two-years (24-months) prior, and include three further lags following the suggestion of Hamilton (2017). We measure the cyclical component as $c_t = y_t - \hat{y}_t$, where \hat{y}_t is the fitted value from the regression in (3). Our measure is therefore purely backward-looking, making no use of either revised data or forward-looking information.

⁷The data becomes available from February 1999 onwards, but the early months have unusually short samples and missing observations. We therefore choose to begin the analysis with the most complete dataset, which begins with the December 1999 vintage.

2.4 Output Gap Portfolios

2.4.1 Cross-Section

At the end of each month t , we sort currencies into five portfolios on the basis of their output gap. Portfolio 1 corresponds to the weakest countries with the lowest output gaps (output most below potential), whereas Portfolio 5 comprises the strongest countries with the highest output gap (output most above potential). We compute the excess return for each portfolio as an equally weighted average of individual currency excess returns within the portfolio. In our out-of-sample analysis, we refer to the difference between Portfolio 5 (P_5) and Portfolio 1 (P_1) as the GAP_{CS} strategy, which is a tradeable investment strategy that exploits the relative cross-sectional spread in business cycle conditions around the world. This approach is equivalent to a strong-minus-weak strategy that buys the currencies of strong economies (characterized by relatively high output gaps) and sells the currencies of weak economies (characterized by relatively low output gaps).

2.5 Time-Series

At the end of each month t , we form a $1/N$ (equally weighted) strategy that takes long positions in the currencies of countries with output gaps above the United States and short positions in the currencies of countries with output gaps below the United States. The strategy thus invests in all currencies available at each point in time, under the expectation that countries with higher (lower) output gaps than the United States should subsequently offer higher (lower) currency excess returns. In our out-of-sample analysis, we refer to this portfolio strategy as GAP_{TS} .

3 Business Cycles and Currency Returns: In-Sample

In this section, we explore the business cycles' predictive power for currency excess returns in-sample, using the full sample of industrial production data. This approach follows a vast literature studying whether macro fundamentals can predict exchange rates when the econometrician has perfect foresight of macroeconomic outcomes (Meese and Rogoff, 1983; Mark, 1995). Our benchmark is a cross-sectional portfolio sort, in which currencies are sorted into five bins (P_1, P_2, P_3, P_4, P_5) based on quintiles of the cross-sectional distribution of output gaps. Within each bin, currencies are equally weighted. We report results for the cross-sectional portfolio that goes long in P_5 (strong economy currencies, i.e., the highest output gaps) and short in P_1 (weak economy currencies, i.e., the lowest output gaps). In addition to this cross-sectional strategy, we also implement a time-series strategy that takes an equal weight in all currencies, with the direction of the trade depending on whether the output gap of a country is higher (long position) or lower (short position) than the U.S. output gap.

This analysis uses the longest possible data sample for output gaps applied to revised data and allows us address the general question of whether there is a meaningful link between economic conditions across countries and currency returns, both in the cross section and the time series. We examine the out-of-sample investment performance in the following section.

3.1 The Link Between Business Cycles and Currency Returns

In Table 1, we present the average excess returns on the five output-gap-sorted portfolios, which monotonically increase from P_1 to P_5 for three of the four output gap measures employed; the only exception is the Baxter-King filter, which displays a slight non-monotonicity from portfolio

P_2 to P_3 . The spread in returns between P_1 to P_5 is sizeable, from 4.13% to 6.66% per annum, and is statistically different from zero at the 1% level for each measure of the output gap employed. Interestingly, the time-series portfolio also delivers a statistically significant spread in returns, which ranges from 2.14% to 3.83% per annum.

Further scrutiny of the results in Table 1 reveal that the predictability is mainly driven by predicting spot exchange rates (in the row denoted fx), whereas the interest rate differential contributes very little (in the row denoted ir). This finding is quite different from that observed with carry trade strategies in which returns are entirely driven by exploiting interest rate differentials across countries, and typically the exchange rate component of the excess return is negative.⁸ The last three rows in Table 1 report a measure of turnover and the spread in both interest rate differentials and output gaps in each of the five output-gap-sorted portfolios. The turnover measure is slightly higher than reported in the literature for carry trade strategies but lower than momentum strategies (see, e.g., Menkhoff et al., 2012a,b). We note a tendency exists for interest rate differentials to increase as we move from P_1 to P_5 , albeit non-monotonically; however, the spread is rather low, consistent with the fact that the returns from sorting on output gaps are not driven by interest rate differentials.

The results in Table 1 are qualitatively identical for all four measures of the output gap considered, indicating that they lead to similar portfolio sorts each month (i.e., similar ranking of countries by the state of the business cycle). In Table 2, we report evidence on the correlation across the portfolio sorts obtained by the different output gap measures. Panel A of Table 2 reports both linear correlation coefficients (below the diagonal) and Spearman rank correlations (above the diagonal) for each pair of output gap measures. To calculate these measures we first collect the cross-sectional correlations (either linear or rank) at each month t . We then calculate the mean of these correlations by averaging across the full sample (the values do not, therefore, reflect a *single* time-series correlation). While the correlations are not perfect they are sizeable, and in a range between 0.41 to 0.65. In Panel B of Table 2 we also report the results from a principal component analysis applied to the output gap estimates for the four different measures. To do so, each month we collect the four cross-sectional vectors of output gaps and extract the proportion of variance explained by the four associated principal components. The average percentage of cross-sectional variation explained by the first principal component is a hefty 86%, indicating that the output gap measures have a very strong common component.

In Table 3, we present the results from a principal component decomposition of the returns of the five portfolios sorted on output gaps, across each of the four measures of output gap considered. The results indicate a strong factor structure in currency portfolio returns sorted by output gaps. The first principal component accounts for most of the variation in portfolio returns, but the loadings appear to be almost identical across the five portfolios, suggesting that this component is essentially a ‘level’ factor. The second principal component appears instead to be a ‘slope’ factor and the loadings of the five portfolios on this component display a tendency to increase (monotonically for two of the output gap measures) from negative values for P_1 to positive values for P_5 . Therefore, it is the second principal component that is key to explaining the cross-sectional difference in excess returns. These features of the factor structure resembles the features displayed by carry portfolios sorted on interest rate differentials, where the ‘slope’ factor is key to understand carry trade excess returns. However, we also find that the second principal component for portfolios sorted on output gaps is orthogonal to the analogous ‘*Slope*’ factor documented by Lustig, Roussanov, and Verdelhan (2011), confirming that sorting currencies on output gaps is very different from sorting currencies on interest rates.

⁸For comparison, we present the equivalent descriptive statistics for forward-premia-sorted (carry) portfolios in Table A.1 of the Internet Appendix.

Overall, the results in this section suggest a strong predictive link between the relative state of the business cycles and future currency excess returns, which is mainly driven by spot exchange rate predictability. Currencies of strong economies tend to appreciate, while currencies of weak economies tend to depreciate over the subsequent month. While this result is supportive of much theory of exchange rate determination that predicts the existence of a relationship between exchange rates and business cycles, it stands in sharp contrast with much literature – based generally on time series analysis – that has struggled to find evidence in favor of such a relationship. This sharp difference in results is due to our method of analysis, based on a portfolio setting. Our portfolio results from the time-series indicate that high output gap currencies *usually* appreciate – the output-gap provides a better-than-chance prediction of future exchange rate movements and the size of the return increases with the output gap.

To further investigate this finding, we also calculate whether a market timing strategy that goes long (short) currencies with output gaps above the United States is *statistically* significant. To do so, we implement the z-score test of Henriksson and Merton (1981) across all observations in our sample. Specifically, we calculate p/r , for which r is the number of correct forecasts and n is the total number of forecasts. If the output gap can predict the direction of exchange rate movements, we should reject the null hypothesis ($p = 0.5$) in favour of the alternative ($p > 0.5$). We evaluate this hypothesis by calculating the test statistic

$$Z = \frac{E[p] - 0.5}{\sqrt{Var[p]}} \sim \mathcal{N}(0, 1) \quad (4)$$

where $E[p] = r/n$ and $Var[p] = p(1 - p)/n$. We report results in Table 4. Across the four methods for constructing output gaps, we find each offers predictability in the direction of both currency excess returns as well as exchange rate returns. The value of $E[p]$ ranges from 0.53 to 0.54, while the test statistics are large and highly statistically significant at the 1% level in all cases. These findings are consistent with the positive returns generated using a time-series portfolio strategy. When we run the same test across the full-sample of 26 bilateral pairs we find, however, that only between 7 and 9 pairs appear to generate predictability of currency excess returns and only between 3 and 6 offer exchange rate predictability that is statistically significant at the 5% level. In contrast, between 15 and 24 (17 and 23) of currency pairs generate an $E[p]$ greater than 0.5 for excess returns (exchange rate returns), indicating the difficulty in overcoming statistical hurdles without sufficient statistical power. Therefore it is perhaps of little surprise, given the high volatility of exchange rates relative to macro fundamentals, that it is difficult to outperform a random walk on purely statistical grounds.

We provide further evidence of the time-series effect in Figure 2. In the top-left hand plot, we show the relationship between the distribution of output gaps (across all observations) and the subsequent average currency excess return. Following a monotonically increasing pattern, currencies with the highest output gaps (relative to the United States) tend to subsequently offer the highest currency excess returns. Focusing only on the time series, however, fails to consider whether there is also a ‘level’ effect – Baz et al. (2015) show that cross-sectional strategies are, under certain conditions, equal to time-series strategies hedged for the ‘level’ effect. In this case, the ‘level’ represents the average conditional variable, i.e., the average output gap of each currency against the United States.

In the top-middle and top-right plots in Figure 2, we consider the level effect by presenting the same information for periods in which the average output gap is above that in the United States (‘strong foreign growth’) and for periods in which the average output gap is below that in the United States (‘weak foreign growth’). It can be seen that during relatively weak periods in the United States all foreign currencies tend to appreciate against the U.S. dollar – even those with *weaker* economies (i.e., with lower output gaps).

These results are consistent with parameter instability in time series studies. Similarly, negative output gap currencies depreciate substantially against the U.S. dollar when the United States is relatively strong compared to other countries. The time-series literature cannot account for these state-dependent coefficients, while the cross-sectional approach is immune to this issue and exploits the fact that the currencies of *relatively* strong economies tend to appreciate relative to the currencies of relatively weaker economies – ignoring any base (level) effect.

The time-series and cross-sectional relationships are highlighted further in the bottom three plots in Figure 2. They show the relationship between output gap differentials and average currency excess returns within three portfolios: weak economies (Portfolio 1), average growth economies (Portfolio 3), and strong economies (Portfolio 5). Within each portfolio, we see the monotonic time-series effect – stronger economies tend to offer higher subsequent returns. The cross-sectional effect can be seen by viewing the average relative output gap across portfolios. In Portfolio 5, even negative output gaps relative to the United States tend to be associated with positive returns, but in Portfolio 1 these same negative output gaps are associated with strongly negative returns, which a time series approach does not uncover.

4 Business Cycles and Currency Returns: Out-of-Sample

The results in the previous section are obtained in sample in the sense that the output gap is constructed using revised industrial production data across the full sample for each country, which is clearly unavailable in real time. In this section we analyze the investment performance of currency strategies that sort on output gaps by moving the analysis out of sample. To do so, we employ the procedure of Hamilton (2017) to construct the output gap recursively as we move through the sample. This procedure is applied to vintage data, which form the information set for industrial production data an investor could have used if implementing the strategy in real time, accounting for both delays in data releases and revisions. These vintage data are available from December 1999. In addition, for the cross-sectional strategy, GAP_{CS} we also construct portfolios based on linear weights and rank weights as implemented by Asness, Moskowitz, and Pedersen (2013), which allows us to trade all currencies available rather than just the corner portfolios. These portfolio construction schemes serve as useful checks that the performance is not generated purely by some outlier currencies. Furthermore, we analyze the impact of transaction costs.

In addition to analyzing the performance of currency strategies sorting on output gaps, we also compare the out-of-sample returns arising from these strategies to a number of other portfolio strategies. This is useful to assess whether sorting on output gaps simply recovers returns that can be obtained in other or simpler ways, or whether they constitute a novel source of exchange rate predictability which can offer diversification gains to investors. Therefore, before presenting the out-of-sample evidence on the strategies performance, we describe other currency portfolio strategies which we use for comparison.

4.1 Alternative Currency Portfolios, Factors and Strategies

Carry Trade Portfolios. At the end of each month t , we allocate currencies to five portfolios on the basis of their forward discounts (or interest rate differential relative to the United States). This exercise implies that currencies with the lowest forward discounts (or lowest interest rate differential relative to the United States) are assigned to Portfolio 1, whereas currencies with the highest forward discounts (or highest interest rate differential relative to the United States) are assigned to Portfolio 5. We compute the excess return for each portfolio as an equally

weighted average of the currency excess returns within that portfolio. The strategy that is long Portfolio 5 and short Portfolio 1 is referred to as *CAR*.

Currency Momentum Portfolios. At the end of each period t , we form five portfolios based on exchange rate returns over the previous month. We assign the 20% of all currencies with the lowest lagged exchange rate returns to Portfolio 1 and the 20% of all currencies with the highest lagged exchange rate returns to Portfolio 5. We then compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio. A strategy long in Portfolio 5 (*winner currencies*) and short in Portfolio 1 (*loser currencies*) is denoted as *MOM*.

Value Portfolios. At the end of each period t , we form five portfolios based on the lagged five-year real exchange rate return as in Asness, Moskowitz, and Pedersen (2013). This measure of currency value is based on calculating a deviation from relative purchasing power parity. Specifically, relative inflation over a 5-year window vis-à-vis the United States is compared with the foreign exchange rate appreciation over the same period versus the U.S. dollar. To provide a more stable measure of the foreign exchange rate appreciation, Asness, Moskowitz, and Pedersen (2013) calculate the appreciation as today's FX rate minus the average FX rate observed 4.5 to 5.5 years earlier. If inflation growth in the foreign economy outpaced that in the U.S. but the U.S. dollar did not appreciate against the foreign currency by an offsetting amount, then the foreign currency is considered 'overvalued'.

To construct currency value portfolios, we collect monthly data on consumer price indices from the IMF's *International Financial Statistics* database beginning in October 1978 and also collect additional foreign exchange spot rate data from *Global Financial Data* beginning in April 1978, such that the first currency value signals are obtained in October 1983. We assign the 20% of all currencies with the highest lagged real exchange rate return to Portfolio 1 and the 20% of all currencies with the lowest lagged real exchange rate return to Portfolio 5. We compute the excess return for each portfolio as an equally weighted average of the currency excess returns within that portfolio. A strategy long in Portfolio 5 (*undervalued currencies*) and short in Portfolio 1 (*overvalued currencies*) is denoted as *VAL*.

Other Factors and Portfolios. In addition to the portfolios described above, we also compare the properties of the output-gap portfolios against other popular strategies and factors in the literature. These include: (i) **Dollar** factor, proposed by Lustig, Roussanov, and Verdelhan (2011), which is essentially a market factor in currency space, equal to the average return of a large basket of foreign currencies against the U.S. dollar; (ii) the **Dollar-Carry** strategy as proposed by Lustig, Roussanov, and Verdelhan (2014), which conditions the Dollar factor on the average forward premia of currencies against the U.S. and thus goes long (short) the U.S. dollar whenever interest rates are relatively high (low) in the United States; (iii) the **Global Imbalance** factor of Della Corte, Riddiough, and Sarno (2016), which is factor that compensates investors for financing risky economies with large stocks of liabilities that issue the majority of those in foreign currency; (iv) the **Trend-Following** risk factors proposed by Fung and Hsieh (2001), which reflect the option-like returns typically generated by hedge funds. We use the foreign exchange and interest-rate trend-following returns; (v) the Pástor and Stambaugh (2003) measure of **Aggregate Market Liquidity** and (vi) the **Market Risk Premium** collected from Kenneth French's website.⁹

⁹The hedge fund risk factors returns are available on David A. Hsieh's website at <http://faculty.fuqua.duke.edu/~dah7/DataLibrary/TF-FAC.xls>. We collect liquidity data from Lubos Pastor's website at <http://>

4.2 The Out-of-Sample Investment Performance

In Panel A of Table 5 we report the gross returns from implementing the GAP_{CS} strategy out of sample using high-minus-low portfolios, linear weights and rank portfolios. We also report the returns from implementing the time-series variant of this strategy, GAP_{TS} , and combinations of the GAP_{CS} strategy with GAP_{CS} . We observe that all investment strategies presented generate statistically significant returns at the 1% level. We also find that using linear weights or rank weights GAP_{CS} improves its Sharpe ratio from 0.6 to 0.72 (for linear weights) and 0.69 (for rank weights). GAP_{TS} generates a Sharpe ratio of 0.65 out of sample, and the Sharpe ratio always increases when GAP_{CS} is combined with the GAP_{CS} strategy, peaking at 0.8 for the 50/50 combination with GAP_{CS} with linear weights. This is due to the fact that the correlation between GAP_{CS} and GAP_{TS} is positive but far from perfect.

In Panel B of Table 5 we report results in the same format as Panel A for returns net of transaction costs, i.e. accounting for bid-ask spreads.¹⁰ These results suggest that transaction costs do not have a material effect on the investment performance of GAP_{CS} and GAP_{TS} , or any of their combinations. The Sharpe ratio for GAP_{CS} goes down from 0.60 to 0.5 when adjusting for bid-ask spreads, and the Sharpe ratio for GAP_{TS} reduces from 0.65 to 0.57, while the highest Sharpe ratio arises again from the combination of GAP_{CS} with linear weights and GAP_{TS} . In short, transaction costs do not wipe out the performance of strategies that sort on output gaps out of sample, and the Sharpe ratios remain attractive even after accounting for transaction costs, so that qualitatively our out-of-sample results are identical with and without accounting for costs.

The results in Table 5 also confirm that the predictive power stems mainly from spot rate predictability rather than interest rate differentials: approximately 90% of the total return is delivered from the FX component across all portfolios considered. Therefore the basic features of exchange rate predictability recorded in sample appear to hold out of sample.

4.3 Relationship with Other Strategies

Table 6 reports a battery of correlation coefficients between the returns from the strategies sorting on output gap and the various combinations examined in Table 5 with the returns from a variety of currency strategies and equity-based strategies. The results are reported for returns obtained out of sample. The main point arising from this table is that the returns of the each variant of the GAP_{CS} strategy, the GAP_{TS} strategy and their combinations are generally uncorrelated with any of the strategies and factors considered. The GAP_{CS} strategy does not show any statistically significant correlation. While some cases of statistically significant correlations exist in the table for some variants of output-gap-sorted strategies, the correlations are generally low, with the largest correlation being a mere 20%. This result suggests the strategies sorting on output gap contain novel economic information and seem unlikely to be a mechanical relabelling of an existing currency strategy or factor.

[//faculty.chicagobooth.edu/lubos.pastor/research/](http://faculty.chicagobooth.edu/lubos.pastor/research/) and market data from Kenneth French's website at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. We thank each author for making their data publicly available.

¹⁰The bid-ask spread data available are for quoted spreads and not effective spreads. Because it is known that quoted spreads are much higher than effective spreads, we follow earlier work (e.g., Goyal and Saretto, 2009; Menkhoff et al., 2012a, 2017), and employ 50% of the quoted bid-ask spread as the actual spread. Even this number seems conservative: Gilmore and Hayashi (2011) find transaction costs due to bid-ask spreads are likely to be much lower than our 50% rule.

4.4 Diversification Gains

Taken together, the previous results suggest that the GAP_{CS} strategy has creditable excess returns overall, low correlation with conventional currency strategies, and the appealing characteristic of strong predictive power for spot exchange rate returns. The importance of these features is twofold. First, a currency investor would likely gain substantial diversification benefits from adding GAP_{CS} strategy to a currency portfolio to enhance risk-adjusted returns. Second, a spot currency trader interested in forecasting exchange rate fluctuations (as opposed to currency excess returns) might value the signals provided by output gaps.

To better understand the value of the GAP_{CS} strategy for a currency investor, we combine it with various canonical currency strategies and assess its value added in terms of performance. In Panel A of Table 7 we first show the out-of-sample returns from carry, dollar carry, momentum, and value strategies during our sample period. We can see that the value strategy performs worse during this sample, with a Sharpe ratio of basically zero, and carry performs best with a Sharpe ratio of 0.58. We also consider a strategy that combines the above four canonical strategies with equal weights (EW). This generates a higher Sharpe ratio of 0.74 relative to each individual strategy by exploiting (albeit simplistically with equal weights) the imperfect correlation of returns across the individual strategies. In essence, the results in Panel A of Table 7 provide us with a benchmark on performance of standard currency strategies, and we ask whether combining them with GAP_{CS} strategy improves performance and to which extent. We report results in Panel B of Table 7, both when we combine each individual strategy with GAP_{CS} and when we add GAP_{CS} to the equally-weighted strategy alongside carry, dollar carry, momentum, and value. The results indicate that adding the GAP_{CS} strategy to this menu of strategies delivers substantially higher Sharpe ratios. For example, the Sharpe ratio of the carry trade improves from 0.58 to 0.80, and the equal-weighted strategy which includes all four benchmark strategies and GAP_{CS} delivers a Sharpe ratio of 0.85, in contrast to 0.74 that is obtained when GAP_{CS} is excluded.

Overall, we view these findings as a confirmation of the value the GAP strategy adds when included in a currency portfolio, driven by its desirable return and correlation properties with existing currency-based strategies.

5 Asset Pricing and Implications

In this section, we begin by investigating if a range of alternative pricing models can explain the returns generated by output-gap sorted portfolios. The purpose of this analysis is to evaluate whether the relationship between currency returns and business cycles can be understood from a risk-return perspective. We go on to consider the role business cycles may also play as a novel source of risk, motivated by the broad macro-finance literature and inherent link between the stochastic discount factor (SDF) and aggregate macroeconomic conditions, before discussing the implications and challenges for future theoretical work.

Methodology. We denote the discrete excess returns on portfolio j in period t as RX_t^j . In the absence of arbitrage opportunities, risk-adjusted excess returns have a price of zero and satisfy the following Euler equation:

$$E_t[M_{t+1}RX_{t+1}^j] = 0 \tag{5}$$

with an SDF linear in the pricing factors f_{t+1} , given by

$$M_{t+1} = 1 - b' (f_{t+1} - \mu) \quad (6)$$

where b is the vector of factor loadings, and μ denotes the factor means. This specification implies a beta pricing model in which the expected excess return on portfolio j is equal to the factor risk price λ times the risk quantities β^j . The beta pricing model is defined as

$$E[RX^j] = \lambda' \beta^j \quad (7)$$

where the market price of risk $\lambda = \Sigma_f b$ can be obtained via the factor loadings b . $\Sigma_f = E[(f_t - \mu)(f_t - \mu)']$, is the variance-covariance matrix of the risk factors, and β^j are the regression coefficients of each portfolio's excess return RX_{t+1}^j on the risk factors f_{t+1} .

5.1 Pricing Output-Gap Portfolios

Risk Factors and Pricing Kernel. The recent literature on cross-sectional asset pricing in currency markets has considered a two-factor SDF. The first risk factor is the expected market excess return, approximated by the average excess return on a portfolio strategy that is long in all foreign currencies with equal weights and short in the domestic currency – the *DOL* factor. For the second risk factor, the literature has employed several return-based factors such as the slope factor (essentially *CAR*) of Lustig, Roussanov, and Verdelhan (2011) or the global volatility risk factor of Menkhoff et al. (2012a).

Following this literature, we start from a two-factor SDF with *DOL* as the first factor, and then consider various second factors, beginning with standard risk factors such as the slope factor (*CAR*) proposed by Lustig, Roussanov, and Verdelhan (2011); the global imbalance factor (*IMB*) of Della Corte, Riddiough, and Sarno (2016), the volatility factor (*VOL*) of Menkhoff et al. (2012a). We also consider alternative currency strategies as factors, such as value (*VAL*) and 1-month momentum (*MOM*) before moving to considering the U.S. equity market (*MKT*) and hedge-fund ‘trend following’ risk factors from Fung and Hsieh (2001); we include the foreign exchange (*PTFSFX*) and interest-rate trend-following factors (*PTFSIR*) but find similar results for all other trend-following returns.

Test Assets. Our test assets include five output-gap sorted currency portfolios using out-of-sample conditioning information as described in Section 4. The portfolio returns therefore begin in January 2000 and end in January 2016. We also follow the suggestion of Lewellen, Nagel, and Shanken (2010) and include the second risk factor as a test asset, e.g, when evaluating the SDF that is constructed using *DOL* and *CAR* factors, we price five output-gap-sorted portfolios *plus* the *CAR* factor. Adding the factor has the dual effect of weakening the strong factor structure and ensuring that the price of risk associated with the factor is approximately equal to the average return of the factor – a necessary condition, since the factor must price itself with $\beta = 1$.¹¹

Cross-Sectional Regressions. Table 8 presents the cross-sectional asset pricing results, including estimates of factor loadings b and the market prices of risk λ . The factor loadings b are estimated via the Generalized Method of Moments (*GMM*) of Hansen (1982). To implement *GMM*, we use the pricing errors as a set of moments and a prespecified weighting matrix. Since the objective is to test whether the model can explain the cross-section of expected currency

¹¹To maintain this approach we construct a tradeable version of the *VOL* factor as the fitted values in a regression of global FX volatility on currency returns following the procedure implemented by Menkhoff et al. (2012a).

excess returns, we only rely on unconditional moments and do not employ instruments other than a constant and a vector of ones. The first stage *GMM* estimation used here employs an identity-weighting matrix, which tells us how much attention to pay to each moment condition. With an identity matrix, *GMM* attempts to price all currency portfolios equally well.

We report estimates of b and λ , and standard errors based on Newey and West (1987). The model’s performance is evaluated using the cross-sectional R^2 and the *HJ* distance measure of Hansen and Jagannathan (1997), which quantifies the mean-squared distance between the SDF of a proposed model and the set of admissible SDFs. To test whether the *HJ* distance is statistically significant, we simulate p -values using a weighted sum of χ_1^2 distributed random variables (see, Jagannathan and Wang, 1996; Ren and Shimotsu, 2009). The p -values of the *HJ* distance measure are reported in brackets.

In Panel A of Table 8 we report results for two-factor SDF models that include *DOL* and, in turn, the carry factor *CAR*, the volatility risk factor *VOL* and the global imbalance risk factor *IMB*. The results suggest a rejection of all three pricing models. In each case the cross-sectional R^2 is low (never greater than 41%), which is surprising, considering the relative ease in achieving *high* R^2 statistics when test assets are characterized by a strong factor structure (Lewellen, Nagel, and Shanken, 2010). Furthermore, the p -values from the *HJ* distance measure are always below 5%, indicating that the model cannot account adequately account for the average variation in the test asset returns.

The results also confirm that *DOL* offers no pricing power (beyond acting as a constant in the model), as evidenced by not contributing significantly to explaining variation in the SDF. This finding is consistent with the findings of Lustig, Roussanov, and Verdelhan (2011) when pricing interest-rate-sorted portfolios. Only the global imbalance factor has a statistically significant coefficient in the SDF regression, although the model’s ability to explain the returns is weak – the cross-sectional R^2 is just 17% and we can reject the null hypothesis that the SDF is correctly specified.

In Panel B of Table 8 we report asset pricing tests for two-factor models which include as second factor the excess return to FX momentum (*MOM*) and value (*VAL*) strategies, constructed as described in Section 2. In Panel C, we experiment further with the SDF specification by considering two-factor models where the second factor is the equity market factor (*MKT*) or one of the two trend-following factors described in Section 2 (*PTFSFX* and *PTFSIR*), which prior literature has found useful in understanding hedge fund returns. All of the asset pricing models are rejected on the basis of the *HJ* distance measure, displaying lack of statistical significance in their estimated coefficients and low explanatory power on the basis of the R^2 .

In sum, the results uniformly reject the ability of standard asset pricing models plus a range of alternatives, to be able to account for returns to the output-gap portfolios. This finding does not rule out a risk-based explanation but highlights the novelty of the returns and the need for alternative risk factors to account for the cross section.

5.2 A Business Cycle Factor?

Next, we consider the possibility that the returns from portfolios sorted on output gap are compensation for risk linked to the relative state of business cycle conditions. The theoretical link between aggregate macroeconomic conditions and asset prices is fundamental to the study of asset pricing, and most classes of risk-based models require the SDF to be a function of the business cycle (Cochrane, 2017). To explore this possibility, we construct a *GAP* factor as the excess return from the GAP_{CS} strategy and ask whether it is priced in the cross section of currencies. Specifically, we carry out asset pricing tests from two SDF specification: a two-

factor model including *DOL* and *CAR*, which is the most common benchmark in the literature since its introduction by Lustig, Roussanov, and Verdelhan (2011), and a three-factor model which also includes the *GAP* factor. This allows to gauge the additional pricing power of a business cycle factor beyond the benchmark. The *GAP* factor essentially measures the excess returns generated by sorting currencies on the output gap information (i.e., the return to the *GAP_{CS}* strategy), and is increasing in the spread of output gaps across the world: it is therefore a measure of the return arising from divergence in business cycles, so that the more business cycles diverge across countries the more the currencies of fast-growing countries appreciate.

Test Portfolios. We consider three sets of test portfolios, increasing in the number of portfolios. We first consider the five output-gap-sorted portfolios, which constitute a small set of test assets for the purpose of asset pricing tests. Again, Lewellen, Nagel, and Shanken (2010) show that a strong factor structure in test asset returns can give rise to misleading results in empirical work, and this outcome is especially the case in small cross sections. Therefore, we also conduct asset pricing tests on: 10 portfolios sorted on currency value and momentum (i.e., out-of-sample test assets where the sorting variable is neither carry nor the output gap); and a larger cross-section of 20 portfolios which comprises the 5 portfolios sorted on output gap, plus 5 portfolios sorted on forward premia (carry), 5 portfolios sorted on momentum, and 5 portfolios sorted on value. We also add the *CAR* factor to each set of test assets and the *GAP* factor when it forms a part of the SDF.

Cross-Sectional Regressions. Starting from Panel A of Table 9, we ask whether a two-factor model including *DOL* and *CAR* portfolios can price the three sets of test assets described above. We focus our interest on the sign and the statistical significance of the market price of risk λ attached to the *CAR* factor and of the associated factor loading b .¹² We know from Table 8 that this SDF specification cannot price the returns from output gap-sorted portfolios. We find that this SDF, which is known to be powerful at pricing carry portfolios, does not explain satisfactorily also the other cross sections considered. Specifically, the factor loading on *CAR* is statistically insignificant from zero, and the R^2 is low. The *HJ* distance test does not indicate a rejection of the model but, with an insignificant factor loading and low R^2 , the *HJ* distance result cannot be considered as supportive of the SDF.

When augmenting the SDF specification with the *GAP* factor, we find that both the loading and the price of risk for the *GAP* factor enter with positive and statistically significant coefficients. Moreover, the factor loading on *CAR* continues to be statistically insignificant. The R^2 for the three-factor model including the *DOL*, *CAR* and *GAP* factors is substantially higher (in the range between 54% and 63%) than the two-factor specification that excludes *GAP* (where the range is between 37% and 42%). Further support in favor of the pricing power of the SDF specification that includes the *GAP* factor comes from the fact that the *HJ* distance is statistically insignificant in each of the asset pricing tests carried out.

5.3 Theoretical Implications

The asset pricing results suggest that standard risk factor used in the literature cannot explain the returns from currency portfolios that sort on output gap. However, the business cycle risk factor we have employed does so, and also appears to be priced in other cross sections of currency returns which are notoriously difficult to price (e.g. momentum). This result is

¹²In general, throughout the literature on currency asset pricing, *DOL* does not display a significant price of risk in cross-sectional tests, and the factor loadings of different portfolio returns do not show a significant spread. This finding was confirmed in Table 8 and occurs with our results in Table 9 as well.

supportive of the broad macro-finance literature that predicts a link between business cycle risk and asset returns, including long-run risk models, rare disaster models, and habit models (Cochrane, 2017). On the other hand, this literature typically predicts a tight link between business cycles and interest rates, and aims at explaining returns to the carry trade (Colacito and Croce, 2013; Farhi and Gabaix, 2016; Ready, Roussanov, and Ward, 2016; Richmond, 2016).

The lack of correlation between the returns to the carry trade strategy (and also momentum and value) and the returns from the GAP_{CS} strategy pose a challenge to these existing theories. Accounting for this lack of correlation requires richer macro-finance models which can account for persistent differentials in interest rates across countries (to explain carry trades) while at the same time allowing for currency excess returns to offer compensation for divergence in business cycles.¹³ Most likely, achieving this task requires moving beyond models with a single shock, which generate a tight connection between carry and business cycles. Moreover, most international macro-finance models are endowment economies, where the output gap plays no role (e.g., Verdelhan, 2010), whereas most international models with production economies and a meaningful account of the output gap (like a new Keynesian model) either do not have risk premia or cannot account very well for the failure of uncovered interest rate parity (Obstfeld and Rogoff, 1995; Cavallo and Ghironi, 2002; Clarida, Gali, and Gertler, 2002). In essence, understanding deeply the facts established in our paper requires an international macro-finance theory of exchange rate determination with endogenous risk premia and, most likely, two sources of shocks. This constitutes a serious and important challenge for theory in this area.

6 Conclusions

The results in this paper are supportive of a strong link between relative business cycles and currency returns. The relationship is evident in-sample and out-of-sample, in both the time-series and cross-section, and is driven almost entirely by the foreign exchange return rather than nominal interest rate differentials. Any theoretical explanation for the relationship between currency excess returns and business cycles must therefore be different from a carry trade based explanation.

The results stand in contrast to the well-known difficulty of predicting foreign exchange returns using macroeconomic variables. Our approach is different and uses relative information and portfolio-based economic evaluation. We show the results improve for two reasons. First, in the time-series, while there is statistical predictability of foreign exchange *direction*, it is only evident once all observations are pooled – bilateral predictability still appears elusive. Critically, the economic evaluation via portfolios reduces volatility, increases the power of the test and thus magnifies the pooled statistical predictability. Second, in the cross-section, our relative sorting procedure is immune from level effects, which determine the relationship between business cycles and currency returns over time and across states. But these level effects matter: when the United States is relatively weak, all currencies tend to appreciate, even those with comparatively weaker economies than the United States. These effects are not captured using standard time-series predictive regressions.

The return predictability we document is useful for global investors seeking novel sources of portfolio diversification. The output-gap based portfolio has zero correlation with the carry trade and close to zero correlations with all other currency trading strategies we consider. Furthermore, output gap portfolios do not appear to offer compensation for risks documented

¹³Colacito et al. (2016) provide promising theoretical developments in accounting for highly persistent interest-rate differentials that may provide a foundation for incorporating additional important empirical facts.

in the literature to date. The macro-finance theoretical literature has, however, the central feature that business cycles are inherently related to stochastic discount factors and thus risk loadings. And indeed, we find that a factor created using our approach is not only priced in the cross section of output-gap-sorted portfolios, but also in other cross sections including currency carry, value and momentum portfolios. We thus provide the first tentative evidence supportive of this link in the broad international asset pricing literature.

The evidence in this paper points towards several directions for future theoretical work. First, we highlight the need to break the link between the aggregate state of the macro-economy and carry-based theories – carry-based theories need to incorporate extremely persistent interest-rate differentials that mean-revert over much longer periods than do relative business cycles. Second, to account for both carry and output-gap returns requires multiple-shock models and thus multiple sources of return premia, possibly across bonds and currency markets. Finally, the results indicate the need for a meaningful output gap to be modelled in international finance theory and thus for deviations from aggregate trend growth to play a role in asset price determination.

References

- Akram, Q. F., D. Rime, and L. Sarno. 2008. Arbitrage in the foreign exchange market: Turning on the microscope. *Journal of International Economics* 76:237–53.
- Asness, C. S., T. J. Moskowitz, and L. H. Pedersen. 2013. Value and momentum everywhere. *Journal of Finance* 68:929–85.
- Baxter, M., and R. G. King. 1999. Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81:575–93.
- Baz, J., N. M. Granger, C. R. Harvey, N. Le Roux, and S. Rattray. 2015. Dissecting investment strategies in the cross section and time series. Working paper: Man Group.
- Cavallo, M., and F. Ghironi. 2002. Net foreign assets and the exchange rate. *Journal of Monetary Economics* 49:1057–97.
- Cheung, Y.-W., M. D. Chinn, and A. Garcia Pascual. 2005. What do we know about recent exchange rate models? in-sample and out-of-sample performance evaluated. In P. De Grauwe, ed., *Exchange Rate Modelling: Where Do We Stand?* MIT Press.
- Clarida, R., J. Gali, and M. Gertler. 2002. A simple framework for international monetary policy analysis. *Journal of Monetary Economics* 49:879–904.
- . 1998. Monetary policy rules in practice: Some international evidence. *European Economic Review* 42:1033–67.
- Cochrane, J. H. 2017. Macro-finance. *Review of Finance, forthcoming* .
- Colacito, R., and M. M. Croce. 2013. International asset pricing with recursive preferences. *Journal of Finance* 68:2651–86.
- Colacito, R., M. M. Croce, F. Gavazzoni, and R. C. Ready. 2016. Currency risk factors in a recursive multi-country economy. Working paper: UNC-Chapel Hill, INSEAD, University of Rochester.
- Cooper, I., and R. Priestley. 2009. Time-varying risk premiums and the output gap. *Review of Financial Studies* 22:2801–33.
- Della Corte, P., S. J. Riddiough, and L. Sarno. 2016. Currency premia and global imbalances. *Review of Financial Studies* 29:2161–93.
- Engel, C., N. C. Mark, and K. D. West. 2007. Exchange rate models are not as bad as you think. *NBER Macroeconomics Annual* 22:381–441.
- Engel, C., and K. D. West. 2005. Exchange rates and fundamentals. *Journal of Political Economy* 113:485–517.
- Farhi, E., and X. Gabaix. 2016. Rare disasters and exchange rates. *Quarterly Journal of Economics forthcoming* 131:1–52.
- Fung, W., and D. A. Hsieh. 2001. The risk in hedge fund strategies: Theory and evidence from trend followers. *Review of Financial Studies* 14:313–41.

- Gilmore, S., and F. Hayashi. 2011. Emerging market currency excess returns. *American Economic Journal: Macroeconomics* 3:85–111.
- Goyal, A., and A. Saretto. 2009. Cross-section of option returns and volatility. *Journal of Financial Economics* 94:310–26.
- Hamilton, J. D. 2017. Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics*, forthcoming .
- Hansen, L. P. 1982. Large sample properties of generalized method of moments. *Econometrica* 50:1029–54.
- Hansen, L. P., and R. Jagannathan. 1997. Assessing specification errors in stochastic discount factor models. *Journal of Finance* 52:557–90.
- Hassan, T. A. 2013. Country size, currency unions, and international asset returns. *Journal of Finance* 68:2269–308.
- Henriksson, R. D., and R. C. Merton. 1981. On market timing and investment performance. ii. statistical procedures for evaluating forecasting skills. *Journal of Business* 54:513–33.
- Hodrick, R. J., and E. C. Prescott. 1980. Postwar u.s. business cycles: An empirical investigation. Working Paper: Carnegie-Mellon University.
- . 1997. Postwar u.s. business cycles: An empirical investigation. *Journal of Money, Credit and Banking* 29:1–16.
- Jagannathan, R., and Z. Wang. 1996. The conditional capm and the cross-section of expected returns. *Journal of Finance* 51:3–53.
- Lewellen, J., S. Nagel, and J. Shanken. 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* 96:175–94.
- Lustig, H., and R. J. Richmond. 2017. Gravity in fx r^2 : Understanding the factor structure in exchange rates. Working paper: Stanford Graduate School of Business and New York University.
- Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *Review of Financial Studies* 24:3731–77.
- . 2014. Countercyclical currency risk premia. *Journal of Financial Economics* 111:527–53.
- Lustig, H., and A. Verdelhan. 2007. The cross-section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97:89–117.
- Mark, N. C. 1995. Exchange rates and fundamentals: Evidence on long horizon predictability. *American Economic Review* 85:201–18.
- Martin, I. 2013. The forward premium puzzle in a two-country world. Working paper: London School of Economics & Political Science.
- Meese, R. A., and K. Rogoff. 1983. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics* 14:3–24.

- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012a. Carry trades and global foreign exchange volatility. *Journal of Finance* 67:681–718.
- . 2012b. Currency momentum strategies. *Journal of Financial Economics* 106:660–84.
- . 2017. Currency value. *Review of Financial Studies* 30:416–41.
- Molodtsova, T., A. Nikolsko-Rzhevskyy, and D. H. Papell. 2008. Taylor rules with real-time data: A tale of two countries and one exchange rate. *Journal of Monetary Economics* 55:167–80.
- Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen. 2012. Time-series momentum. *Journal of Financial Economics* 104:228–50.
- Newey, W. K., and K. D. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–8.
- Obstfeld, M., and K. Rogoff. 1995. Exchange rate dynamics redux. *Journal of Political Economy* 103:624–60.
- Pástor, L., and R. F. Stambaugh. 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111:642–85.
- Ready, R., N. Roussanov, and C. Ward. 2016. Commodity trade and the carry trade. *Journal of Finance*, forthcoming .
- Ren, Y., and K. Shimotsu. 2009. Improvement in finite sample properties of the hansen-jagannathan distance test. *Journal of Empirical Finance* 16:483–506.
- Richmond, R. J. 2016. Trade network centrality and currency risk premia. Working Paper: New York University.
- Rogoff, K. S., and V. Stavrakeva. 2008. The continuing puzzle of short horizon exchange rate forecasting. Working paper: Harvard University and London Business School.
- Rossi, B. 2013. Exchange rate predictability. *Journal of Economic Literature* 51:1063–119.
- Tran, N.-K. 2013. Growth risk of nontraded industries and asset pricing. Working paper: Washington University in St. Louis.
- Verdelhan, A. 2010. A habit-based explanation of the exchange rate risk premium. *Journal of Finance* 65:123–46.
- . 2017. The share of systematic variation in bilateral exchange rates. *Journal of Finance*, forthcoming .

Table 1: Descriptive Statistics of Output-Gap-Sorted Portfolios

The table presents descriptive statistics for five currency portfolios sorted by output gaps. The output gap at time- t is estimated as (log) industrial production *minus* the (log) trend in industrial production. The trend is estimated in four ways using a (i) Hodrick-Prescott filter; (ii) Baxter-King Filter, (iii) linear projection and (iv) quadratic time trend. Portfolios are rebalanced monthly with strong (weak) economy currencies entering P_5 (P_1). We report summary statistics for the annualized excess mean return and its decomposition between the exchange rate (fx) and interest rate (ir) components. We also report the Sharpe ratio ($Sharpe$), standard deviation (std), skewness ($skew$), kurtosis ($kurt$), maximum drawdown (mdd), average turnover (t/o), average forward premium (fp) and average output gap (gap) for each portfolio. The Cross-Section is a *high-minus-low* portfolio long P_5 and short P_1 . The Time-Series is a 1/N portfolio long (short) currencies issued by countries with an output gap above (below) the United States' output gap. The superscripts *, **, *** represent significance of the Cross-Section and Time-Series portfolios at the 10%, 5% and 1% level using Newey and West (1987) standard errors. The sample is from October 1983 to January 2016.

	<i>Hodrick-Prescott Filter</i>							<i>Baxter-King Filter</i>						
	P_1	P_2	P_3	P_4	P_5	Cross-Section	Time-Series	P_1	P_2	P_3	P_4	P_5	Cross-Section	Time-Series
<i>mean</i> (%)	-0.25	0.96	2.77	4.00	6.41	6.66***	2.45***	-0.44	2.45	2.39	3.72	5.97	6.41***	3.83***
<i>fx</i> (%)	-2.34	-1.03	0.88	1.58	2.72	5.06	2.03	-2.34	0.65	0.49	1.23	1.92	4.26	3.38
<i>ir</i> (%)	2.09	1.99	1.89	2.41	3.69	1.60	0.41	1.90	1.80	1.90	2.49	4.06	2.15	0.44
<i>Sharpe</i>	-0.02	0.11	0.27	0.43	0.71	0.82	0.54	-0.04	0.26	0.26	0.39	0.68	0.77	0.74
<i>std</i>	10.18	9.09	10.12	9.32	9.05	8.14	4.57	10.15	9.50	9.33	9.51	8.82	8.31	5.21
<i>skew</i>	-0.06	-0.47	-0.28	-0.27	-0.28	0.01	-0.92	-0.08	0.06	-0.21	-0.29	-0.61	0.07	-0.39
<i>kurt</i>	4.49	4.72	4.75	4.39	3.97	4.32	10.89	3.94	3.83	4.32	4.22	5.00	4.25	5.70
<i>mdd</i> (%)	42.5	34.2	23.9	23.6	24.4	9.0	8.6	49.0	28.8	26.1	24.6	21.8	23.0	9.2
<i>t/o</i> (%)	44.8	58.2	67.2	60.6	44.8			10.0	21.8	29.8	23.1	11.6		
<i>fp</i> (t, %)	2.23	2.03	1.80	2.45	4.15			1.91	1.87	1.75	2.38	4.81		
<i>gap</i> (t, %)	-3.08	-0.96	0.11	1.17	3.01			-2.75	-0.84	0.18	1.30	3.01		

	<i>Linear projection</i>							<i>Quadratic time-trend</i>						
	P_1	P_2	P_3	P_4	P_5	Cross-Section	Time-Series	P_1	P_2	P_3	P_4	P_5	Cross-Section	Time-Series
<i>mean</i> (%)	1.10	2.01	2.59	2.95	5.23	4.13***	3.12***	0.27	1.99	3.08	4.21	4.83	4.56***	2.14**
<i>fx</i> (%)	-1.63	0.09	0.41	0.88	2.05	3.68	2.80	-1.04	-0.16	0.02	1.54	1.75	2.80	1.95
<i>ir</i> (%)	2.73	1.92	2.19	2.08	3.18	0.45	0.32	1.31	2.15	3.06	2.67	3.08	1.76	0.19
<i>Sharpe</i>	0.11	0.23	0.27	0.31	0.55	0.54	0.61	0.03	0.21	0.32	0.49	0.51	0.60	0.41
<i>std</i>	9.95	8.87	9.59	9.56	9.52	7.61	5.15	10.05	9.58	9.76	8.66	9.53	7.56	5.27
<i>skew</i>	-0.07	-0.39	-0.13	-0.35	-0.38	-0.29	-1.17	-0.23	-0.21	-0.07	-0.51	-0.15	-0.68	-1.24
<i>kurt</i>	4.70	5.29	3.96	4.09	4.76	5.50	11.10	4.23	4.58	4.29	4.50	5.02	6.29	10.93
<i>mdd</i> (%)	34.1	29.4	29.8	32.7	23.2	37.7	12.2	38.9	28.5	31.1	24.8	21.5	18.3	16.5
<i>t/o</i> (%)	26.7	42.7	53.0	44.1	27.7			20.0	32.9	44.3	33.8	19.7		
<i>fp</i> (t, %)	2.70	1.99	2.10	2.32	3.58			1.17	2.04	3.15	2.87	3.35		
<i>gap</i> (t, %)	-1.30	-0.28	0.32	0.92	2.02			-8.41	-3.35	-0.22	2.59	7.78		

Table 2: The Correlation and Factor Structure of Output Gap Measures

The table presents the average cross-sectional correlation and factor structure in output gaps. The output gap at time- t is estimated as (log) industrial production *minus* the (log) trend in industrial production. The trend is estimated in four ways using a (i) Hodrick-Prescott filter; (ii) Baxter-King Filter, (iii) linear projection and (iv) quadratic time trend. In Panel A, the entries below the diagonal are linear (Pearson) correlations, which are calculated by taking the time-series average of monthly cross-sectional correlations for all available currencies. The entries above the diagonal are Spearman rank correlations, also calculated by taking the time-series average of monthly cross-sectional correlations for all available currencies. In Panel B, we present the average percentage of the cross-sectional variation accounted for by each principal component (PC). To calculate, we estimate the variation explained by each PC each month and then take a time-series average across all months in the sample. The sample is from October 1983 to January 2016.

Panel A: Output-Gap Correlations				
	HP	BK	LP	QT
<i>Hodrick-Prescott Filter (HP)</i>		0.63	0.55	0.45
<i>Baxter-King Filter (BK)</i>	0.65		0.57	0.58
<i>Linear Projection (LP)</i>	0.53	0.56		0.47
<i>Quadratic Time-trend (QT)</i>	0.41	0.53	0.48	

Panel B: Output-Gap Factor Structure				
	PC1	PC2	PC3	PC4
<i>var explained</i>	86%	10%	3%	1%

Table 3: Principal Components of Output-Gap-Sorted Portfolios

The table presents results from a principal component decomposition of the returns to five currency portfolios sorted by output gaps. The output gap at time- t is estimated as (log) industrial production *minus* the (log) trend in industrial production. The trend is estimated in four ways using a (i) Hodrick-Prescott filter; (ii) Baxter-King Filter, (iii) linear projection and (iv) quadratic time trend. Portfolios are rebalanced monthly with strong (weak) economy currencies entering P_5 (P_1). We report the loading of each portfolio on all five principal components (PCs) and the percentage of total return variation explained by each PC. The sample is from October 1983 to January 2016.

<i>Hodrick-Prescott Filter</i>						<i>Baxter-King Filter</i>					
	PC1	PC2	PC3	PC4	PC5		PC1	PC2	PC3	PC4	PC5
P_1	0.47	-0.71	0.46	-0.08	-0.23	P_1	0.48	-0.51	0.69	0.11	-0.16
P_2	0.43	-0.19	-0.33	0.00	0.82	P_2	0.46	-0.18	-0.24	-0.26	0.79
P_3	0.49	0.18	-0.25	0.75	-0.32	P_3	0.45	-0.18	-0.52	-0.39	-0.59
P_4	0.44	0.19	-0.47	-0.64	-0.38	P_4	0.46	0.25	-0.27	0.81	-0.02
P_5	0.40	0.63	0.63	-0.11	0.19	P_5	0.38	0.78	0.36	-0.34	-0.05
<i>var explained</i>	78%	8%	5%	5%	4%	<i>var explained</i>	79%	7%	6%	4%	4%
<i>Linear Projection</i>						<i>Quadratic Time Trend</i>					
	PC1	PC2	PC3	PC4	PC5		PC1	PC2	PC3	PC4	PC5
P_1	0.47	-0.79	0.14	-0.38	-0.05	P_1	0.48	-0.08	0.70	-0.34	-0.40
P_2	0.41	-0.16	-0.36	0.75	-0.33	P_2	0.47	-0.10	0.24	0.21	0.82
P_3	0.45	0.39	-0.64	-0.49	-0.04	P_3	0.45	-0.52	-0.60	-0.40	-0.04
P_4	0.46	0.14	0.19	0.23	0.82	P_4	0.40	-0.10	-0.13	0.80	-0.41
P_5	0.44	0.43	0.64	-0.05	-0.45	P_5	0.43	0.84	-0.29	-0.17	-0.02
<i>var explained</i>	78%	7%	6%	5%	4%	<i>var explained</i>	78%	7%	6%	5%	4%

Table 4: Can Output Gaps Predict the Direction of Currency and Foreign Exchange Returns?

The table presents evidence of the directional predictability of currency excess returns and foreign exchange rate returns using the approach of Henriksson and Merton (1981). At each month-end- t we predict the return at $t + 1$ to be positive if the output gap (relative to the United States) is positive. The output gap at time- t is estimated as (log) industrial production *minus* the (log) trend in industrial production. The trend is estimated in four ways using a (i) Hodrick-Prescott filter; (ii) Baxter-King Filter, (iii) linear projection and (iv) quadratic time trend. Across all observations in the sample, we calculate $E[p]$, where $p = r/n$ in which r is the number of correct predictions and n is the total number of predictions. The test statistics (*test-stat*) is given by $Z = \frac{E[p]-0.5}{\sqrt{Var[p]}} \sim \mathcal{N}(0, 1)$, where $Var[p] = p(1 - p)/n$. A value $> \approx 1.64$ indicates the null hypothesis (no directional predictability) can be rejected at the 5% confidence level. We also report the bilateral success rate. This number reflects the number of bilateral currency pairs for which the value $p > 0.5$. We also report the bilateral significance, which indicates the number of bilateral currency pairs for which the null hypothesis (no directional predictability) could be rejected at the 5% confidence level. The sample is from October 1983 to January 2016.

	Hodrick- Prescott	Baxter- King	Linear Projection	Quadratic Time Trend
<i>currency excess returns</i>				
$E[p]$	0.54***	0.53***	0.53***	0.53***
<i>test-stat</i>	6.10	4.04	5.08	4.83
<i>bilateral success</i>	24/26	17/26	15/26	19/26
<i>bilateral significance</i>	7/26	9/26	7/26	9/26
<i>foreign exchange returns</i>				
$E[p]$	0.53***	0.53***	0.53***	0.53***
<i>test-stat</i>	5.05	4.15	4.97	4.57
<i>bilateral success</i>	23/26	19/26	17/26	20/26
<i>bilateral significance</i>	6/26	3/26	6/26	5/26

Table 5: Out-of-Sample Currency Trading Strategies' Investment Performance

The table presents investment performance for output-gap currency trading strategies. The output gap is estimated using monthly ‘vintages’ of real-time industrial production data from the OECD’s *Real-Time Data and Revisions Database*. To estimate the output-gap we follow the linear projection procedure in Hamilton (2017) by running the regression, $y_t = \alpha + \beta_1 y_{t-24} + \beta_2 y_{t-25} + \beta_3 y_{t-26} + \beta_4 y_{t-27} + \varepsilon_t$ each month, in which y is (log) industrial production. The output gap is constructed as the difference between the most recently available data point at time- t (y_t) and the fitted value from the regression. *GAP* is a *high-minus-low* portfolio formed as $P_5 - P_1$, after sorting currencies into five portfolios. *LIN* and *RNK* take a position in all currencies with the weight determined by either the magnitude or relative size of the output gap. *TS* is a 1/ N time-series strategy long (short) currencies issued by countries with an output gap above (below) the United States’ output gap. The three *COM* portfolios take 50 – 50 weights in the time-series and either the *GAP*, *LIN* or *RNK* strategy. We report summary statistics for the annualized mean, which is then further split between the exchange rate (*fx*) and interest rate (*ir*) components, we also report the Sharpe ratio (*Sharpe*), skewness (*skew*), kurtosis (*kurt*) and maximum drawdown (*mdd*). The superscripts *, **, *** represent significance of the strategy mean excess returns at the 10%, 5% and 1% confidence levels using Newey and West (1987) corrected standard errors. The sample runs from December 1999 to January 2016.

Panel A: Out-of-Sample Linear Projection (<i>excluding b-a spread</i>)							
	GAP	LIN	RNK	TS	COM _{GAP}	COM _{LIN}	COM _{RNK}
<i>mean</i> (%)	4.16***	2.11***	3.71***	2.59**	3.36***	2.34***	3.14***
<i>fx</i> (%)	3.67	1.84	3.38	2.65	3.14	2.22	2.99
<i>ir</i> (%)	0.49	0.26	0.33	-0.06	0.22	0.11	0.14
<i>Sharpe</i>	0.60	0.72	0.69	0.65	0.72	0.80	0.78
<i>skew</i>	0.18	0.27	0.05	-0.53	-0.05	-0.42	-0.33
<i>kurt</i>	2.81	3.27	3.53	4.56	3.03	3.70	3.77
<i>mdd</i> (%)	6.90	2.92	5.41	4.00	4.66	2.92	4.02

Panel B: Out-of-Sample Linear Projection (<i>including b-a spread</i>)							
	GAP	LIN	RNK	TS	COM _{GAP}	COM _{LIN}	COM _{RNK}
<i>mean</i> (%)	3.44**	1.75***	3.02***	2.28**	2.85**	2.01***	2.64***
<i>fx</i> (%)	3.14	1.58	2.87	2.42	2.75	1.97	2.62
<i>ir</i> (%)	0.31	0.18	0.15	-0.14	0.10	0.03	0.02
<i>Sharpe</i>	0.50	0.60	0.56	0.57	0.61	0.69	0.66
<i>skew</i>	0.18	0.27	0.04	-0.53	-0.05	-0.42	-0.33
<i>kurt</i>	2.80	3.26	3.52	4.58	3.03	3.71	3.79
<i>mdd</i> (%)	6.89	2.92	5.41	4.01	4.66	2.92	4.03

Table 6: Correlations Between Trading Strategies

The table presents linear correlation coefficients between trading strategies. In the upper panel we report correlations between output-gap currency trading strategies. The output gap is estimated using monthly ‘vintages’ of real-time industrial production data from the OECD’s *Real-Time Data and Revisions Database*. To estimate the output-gap we follow the linear projection procedure in Hamilton (2017) by running the regression, $y_t = \alpha + \beta_1 y_{t-24} + \beta_1 y_{t-25} + \beta_1 y_{t-26} + \beta_1 y_{t-27} + \varepsilon_t$ each month, in which y is (log) industrial production. The output gap is constructed as the difference between the most recently available data point at time- t (y_t) and the fitted value from the regression. *GAP* is a *high-minus-low* portfolio formed as $P_5 - P_1$, after sorting currencies into five portfolios. *LIN* and *RNK* take a position in all currencies with the weight determined by either the magnitude or relative size of the output gap. *TS* is a 1/ N time-series strategy long (short) currencies issued by countries with an output gap above (below) the United States’ output gap. The three *COM* portfolios take 50 – 50 weights in the time-series and either the *GAP*, *LIN* or *RNK* strategy. In the lower panel we present correlations between the output-gap currency trading strategies and various currency and equity-based strategies. We include full details of these strategies in Section XY. The superscripts *, ** represent significance of the correlation coefficients at the 5% and 1% confidence levels. The sample runs from December 1999 to January 2016.

<i>Output-gap Currency Trading Strategies</i>							
	GAP	LIN	RNK	TS	COM _{GAP}	COM _{LIN}	COM _{RNK}
<i>High-minus-Low (GAP)</i>							
<i>Linear Weights (LIN)</i>	0.86**						
<i>Rank Weights (RNK)</i>	0.88**	0.94**					
<i>Time Series (TS)</i>	0.42**	0.40**	0.45**				
<i>GAP Model Combo (COM_{GAP})</i>	0.92**	0.81**	0.84**	0.74**			
<i>LIN Model Combo (COM_{LIN})</i>	0.72**	0.78**	0.78**	0.89**	0.92**		
<i>RNK Model Combo (COM_{RNK})</i>	0.80**	0.84**	0.90**	0.80**	0.94**	0.97**	
<i>Alternative Trading Strategies in Currency and Equity Markets</i>							
	GAP	LIN	RNK	TS	COM _{GAP}	COM _{LIN}	COM _{RNK}
<i>HML_{fx}</i>	0.04	-0.01	-0.04	-0.07	-0.01	-0.06	-0.06
<i>Dollar</i>	0.13	0.16*	0.18*	0.16*	0.17*	0.19**	0.20**
<i>Dollar Carry</i>	0.13	0.16*	0.17*	0.16*	0.16*	0.19**	0.19**
<i>Value</i>	0.06	0.07	0.13	0.13	0.10	0.12	0.15*
<i>Momentum</i>	0.13	0.16*	0.09	0.00	0.10	0.08	0.06
<i>Global Imbalance</i>	0.13	0.11	0.11	0.19*	0.17*	0.18*	0.16*
<i>Foreign Exchange Trend Strategy</i>	0.07	0.07	0.07	0.14	0.11	0.13	0.11
<i>Interest Rate Trend Strategy</i>	0.09	0.12	0.12	-0.17*	-0.01	-0.06	-0.01
<i>Illiquidity</i>	0.08	0.07	0.02	0.12	0.11	0.12	0.07
<i>U.S. Equity</i>	-0.11	-0.07	-0.05	0.03	-0.07	-0.01	-0.02

Table 7: Diversification Benefits of Adding GAP to Common Currency Strategies

The table presents the investment performance of common currency trading strategies and the impact on performance from adding an output-gap trading strategy. The output gap is estimated using monthly ‘vintages’ of real-time industrial production data from the OECD’s *Real-Time Data and Revisions Database*. To estimate the output-gap we follow the linear projection procedure in Hamilton (2017) by running the regression, $y_t = \alpha + \beta_1 y_{t-24} + \beta_1 y_{t-25} + \beta_1 y_{t-26} + \beta_1 y_{t-27} + \varepsilon_t$ each month, in which y is (log) industrial production. The output gap is constructed as the difference between the most recently available data point at time- t (y_t) and the fitted value from the regression. *GAP* is a *high-minus-low* portfolio formed as $P_5 - P_1$, after sorting currencies into five portfolios. In Panel A, we report the investment performance of popular currency investment strategies, in which *CAR* is the currency carry trade; *DCAR* is a ‘dollar’ carry trade; *MOM* is a momentum trade; *VAL* is a value trade and *EW* is a 1/ N portfolio that takes an equal position in each individual currency strategy. We include full details of these strategies in Section XY. In Panel B, we add the *GAP* strategy to each individual strategy and to the broader equally-weighted (*EW*) portfolio. We report summary statistics for the annualized excess mean return, the Sharpe ratio (*Sharpe*), standard deviation (*std*), skewness (*skew*), kurtosis (*kurt*), maximum drawdown, percentage increase in Sharpe ratio ($\% \Delta$ Sharpe) and weight in the *GAP* portfolio (w_{GAP}). The superscripts *, **, *** represent significance of the strategy mean excess returns at the 10%, 5% and 1% confidence levels using Newey and West (1987) corrected standard errors. The sample runs from December 1999 to January 2016.

Panel A: Excluding Output Gap Strategy					
	CAR	DCAR	MOM	VAL	EW
<i>mean (%)</i>	6.34**	2.60	1.41	0.05	2.60***
<i>Sharpe</i>	0.58	0.31	0.16	0.01	0.74
<i>std</i>	10.91	8.51	9.01	8.65	3.53
<i>skew</i>	-0.72	-0.49	0.28	0.47	-0.24
<i>kurt</i>	5.23	4.78	3.31	4.42	4.13
<i>mdd (%)</i>	0.16	0.17	0.23	0.39	0.07
Panel B: Including Output Gap Strategy					
	GAP+ CAR	GAP+ DCAR	GAP+ MOM	GAP+ VAL	EW
<i>mean (%)</i>	5.24***	3.37**	2.77**	2.09*	2.91***
<i>Sharpe</i>	0.80	0.58	0.46	0.37	0.85
<i>std</i>	6.56	5.82	6.04	5.69	3.41
<i>skew</i>	-0.40	-0.46	0.27	0.39	-0.10
<i>kurt</i>	4.47	3.30	3.19	3.63	3.85
<i>mdd (%)</i>	0.09	0.09	0.10	0.16	0.06
$\% \Delta$ <i>Sharpe</i>	37.5	89.5	194	6714	15.6
w_{GAP} (%)	50.0	50.0	50.0	50.0	20.0

Table 8: Cross-Sectional Asset Pricing Tests of Output Gap Portfolios

The table presents cross-sectional asset pricing results. We construct various two-factor linear SDF's that include the *DOL* factor plus a second pricing factor. In each model, we price six test asset portfolios including five output-gap sorted currency portfolios plus the second pricing factor. We report Generalized Method of Moments (GMM) one-step estimates of factor loadings on the pricing kernel (*b*'s) and prices of factor risk (λ 's). The superscripts *, **, *** represent significance of the coefficients at the 10%, 5% and 1% confidence levels using Newey and West (1987) corrected standard errors. We also report goodness-of-fit statistics for each model including the R^2 statistic and the Hansen-Jagannathan distance statistic (*HJ*) with simulated *p*-values in brackets. The *HJ* statistic measures the distance between the estimated pricing kernel and the efficient set of permissible pricing kernels. A *p*-value less than 5% indicates the null hypothesis that the pricing kernel is efficient can be rejected at the 95% confidence level. We provide full details of the pricing factors in Section XY. The sample runs from December 1999 to January 2016.

Panel A: Standard Currency Factors						
	SDF		Risk		Model Fit	
	Loadings (<i>b</i>)		Prices (λ)		R^2	HJ_{dist}
	DOL	FAC	DOL	FAC		
DOL + CAR	0.25	0.43	0.03	0.06**	0.37	0.28 [0.03]
DOL + IMB	0.79	1.20***	0.04*	0.05***	0.17	0.31 [0.04]
DOL + VOL	0.24	-1.28	0.03	-0.00	0.41	0.38 [0.00]
Panel B: Alternative Currency Factors						
	SDF		Risk		Model Fit	
	Loadings (<i>b</i>)		Prices (λ)		R^2	HJ_{dist}
	DOL	FAC	DOL	FAC		
DOL + VAL	0.29	0.03	0.03	0.00	0.28	0.27 [0.02]
DOL + MOM	0.32	0.21	0.03	0.16	0.19	0.26 [0.02]
Panel C: Equity and Trend-Following Hedge Fund Strategies						
	SDF		Risk		Model Fit	
	Loadings (<i>b</i>)		Prices (λ)		R^2	HJ_{dist}
	DOL	FAC	DOL	FAC		
DOL + MKT	0.29	0.03	0.03	0.02	0.09	0.28 [0.02]
DOL + PTFSFX	0.29	0.03	0.03	0.00	0.31	0.27 [0.02]
DOL + PTFSIR	0.28	-0.06	0.03	-0.01	0.40	0.28 [0.01]

Table 9: Cross-Sectional Asset Pricing Tests including GAP as a Pricing Factor

The table presents cross-sectional asset pricing results for three sets of test portfolios. The SDF is constructed as a linear combination of *DOL* and *CAR* (2 Pricing Factors, left-hand-side) and *DOL*, *CAR* and *GAP* (3 Pricing Factors, right-hand-side). On the left-hand side we always add *CAR* as a test asset, while on the right-hand side we always add *CAR* and *GAP* as test assets. We report Generalized Method of Moments (GMM) one-step estimates of factor loadings on the pricing kernel (*b*'s) and prices of factor risk (λ 's). The superscripts *, **, *** represent significance of the coefficients at the 10%, 5% and 1% confidence levels using Newey and West (1987) corrected standard errors. In addition, we report goodness-of-fit statistics for each model including the R^2 statistic and the Hansen-Jagannathan distance statistic (*HJ*) with simulated *p*-values in brackets. The *HJ* statistic measures the distance between the estimated pricing kernel and the efficient set of permissible pricing kernels. A *p*-value less than 5% indicates the null hypothesis that the pricing kernel is efficient can be rejected at the 95% confidence level. The sample runs from December 1999 to January 2016.

	2 Pricing Factors (<i>DOL</i> + <i>CAR</i>)						3 Pricing Factors (<i>DOL</i> + <i>CAR</i> + <i>GAP</i>)							
	Factor Loadings (<i>b</i>)		Risk Prices (λ)		Model Fit		Factor Loadings (<i>b</i>)			Risk Prices (λ)			Model Fit	
	DOL	CAR	DOL	CAR	R^2	HJ_{dist}	DOL	CAR	GAP	DOL	CAR	GAP	R^2	HJ_{dist}
5 Test Portfolios (output gap)	0.25	0.43	0.026	0.064**	0.37	0.28 [0.03]	0.18	0.42	0.72**	0.026	0.064**	0.044**	0.63	0.39 [0.94]
10 Test Portfolios (value, momentum)	0.21	0.41	0.022	0.061**	0.38	0.23 [0.81]	0.13	0.41	0.77**	0.021	0.062**	0.046**	0.54	0.21 [0.87]
20 Test Portfolios (output gap, carry, value, momentum)	0.23	0.41	0.023	0.061**	0.42	0.70 [0.96]	0.14	0.40	0.77**	0.023	0.061**	0.047***	0.61	0.70 [1.00]

Figure 2: Excess Currency Returns and the Distribution of Output Gaps

The figure plots the average excess currency return at time- t , conditional on the output gap observed at time- $t-1$. For each plot, we sort observed output gaps and subsequent currency excess returns across a set of observations from the lowest (weakest) output gap to the highest (strongest) output gap. We split the data into five quintiles and report the average excess currency return (orange bar) and average output gap (blue line). The ‘All Periods’ figure includes every observation in the sample. The ‘Strong’ (‘Weak’) Foreign Growth sample includes all observations in which the average output gap of all countries is above (below) that in the United States. In the lower-half of the figure, we plot the figures in which the observations are restricted to particular portfolio allocations. For example, the observations in Portfolio 1 reflect the lowest output gaps at any given point that would have been allocated to Portfolio 1, when sorting currencies into five equally weighted portfolios as in Table 1.

