

School of Finance



University of St.Gallen

MONETARY POLICY AND CURRENCY RETURNS: THE FORESIGHT SAGA

DMITRY BORISENKO

IGOR POZDEEV

WORKING PAPERS ON FINANCE NO. 2017/08

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

MAY 2017

THIS VERSION: OCTOBER 2017



Monetary Policy and Currency Returns: the Foresight Saga^{*}

Dmitry Borisenko[†], Igor Pozdeev[‡]

This version: October 2017

Abstract

We document a drift in exchange rates before monetary policy changes across major economies. Currencies tend to depreciate by 0.8 percent over ten days before policy rate cuts and appreciate by 0.5 percent before policy rate increases. We show that available fixed income instruments allow to accurately forecast monetary policy decisions and thus that the drift is foreseeable and exploitable by investors. Our baseline specification of a trading strategy constructed by going long in currencies against USD before predicted local interest rate hikes and short in currencies before predicted cuts earns on average a statistically significant excess return of 38 basis points per ten-day period after trading costs. We further demonstrate that this return is robust to the choice of holding horizon and monetary policy forecast rule. Our results thus pose a major challenge for the risk-based explanations of the exchange rate dynamics and highlight an important side effect of monetary policy decisions.

Keywords: Monetary Policy, Policy Expectations, Predictability, Overnight Index Swap, Foreign Exchange.

JEL Classification: E43, E52, E58, F31, G12

*We would like to thank (in alphabetical order) Michael Bauer, Anastasia Berezinskaya, Angelo Ranaldo, Paul Söderlind and Annette Vissing-Jørgensen for their valuable comments and insightful feedback.

[†]E-mail: dmitry.borisenko@student.unisg.ch

[‡]E-mail: igor.pozdeev@unisg.ch

I. Introduction

By 2001, the central banks of most developed countries had adopted the practice of a fixed number of scheduled meetings per year, each culminating in a decision about the policy rate – most often a short-term rate. The decision would then by virtue of the expectation hypothesis propagate through the yield curve and eventually find its way into the domestic currency return, as documented in the vast literature on the uncovered interest parity and carry trade strategies. This indirect effect of monetary policy on exchange rates has been extensively studied; the direct, or immediate, relation has overgrown with anecdotal evidence – such as the January 2015 rate cut by the Swiss national bank trying to prevent a rapid appreciation of the franc – but been far less researched.

Mueller et al. (2017) were the first to document the abnormal positive return of being long in foreign (from the perspective of an American investor) currencies in the hours around FOMC announcements. They find that the effect is more pronounced for the high interest rate currencies, and that a simple *ex post* conditioning on the sentiment of policy decisions allows to improve the strategy performance. Karnaugh (2016) conducts similar research in the low-frequency dimension. She reports that the US dollar (synthetic exchange rate of the USD vs. a basket of currencies) tends to depreciate days before Federal funds rate cuts and appreciate before rate hikes. Using the rates implied in the Federal funds futures' prices to bet on the direction of the upcoming target rate change several days in advance, she constructs a strategy of high profitability between 1994 and 2015.

In this paper, we bring the currencies and policy announcements of other developed countries into the picture. Our primary contribution is to document an economically and statistically significant drift in exchange rates several days in advance of changes in target policy rates across major economies. We show that a randomly selected currency is expected to depreciate against the USD by 80 basis points over the 10 days before a randomly selected rate cut, and appreciate by 50 basis points in the opposite case, which is statistically significant at the 5% level. We find that the multiperiod appreciation before rate hikes and depreciation before rate cuts is a phenomenon common to most currencies – not only to the US dollar, as shown by Karnaugh (2016). We further demonstrate that this drift can be exploited by investors as a trading strategy. Using overnight index swaps, we forecast upcoming rate changes and go long in currencies ten days in advance of an expected rate hike and short in those with an expected rate cut. The strategy features a statistically and economically significant excess return of 38 basis points per event after transaction costs, a cumulative of 165

percent since late 2000.

Our second contribution is to point out that forecasting monetary policy direction is a classification problem and thus subject to the discretionary choice of the classification rule. The holding period is another “tweaker” for the trader or researcher to adjust. Thus, there exist many possible forecast-based trading strategies: one for each element of the Cartesian product of the set of possible classification rules and the set of possible holding periods. Since it is never clear *ex ante* which strategy specification will result in a significant return, the backward-looking bias discussed i.a. in Bailey and Lopez de Prado (2014) might contaminate the inference. We construct a plethora of specifications and show that our findings are very robust: the average return across all specifications amounts to 85 percent since late 2000 while several specifications lead to as much as a 200 percent return and only a handful – to money loss. Interestingly, when treated in the same way, the dollar-FOMC pre-announcement drift of Karnaukh (2016) is found to be less significant, averaging to zero across all specifications.

Our third contribution is to the literature on forecasting future policy rates. While evidence on the predictive power of the Fed funds futures is abundant (Krueger and Kuttner (1996) and Piazzesi and Swanson (2008), to name a few), this paper is to our knowledge the first extensive treatment of how overnight index swaps (OIS) can be used with the same purpose. We find that policy rate forecasts extracted from OIS rates have been most accurate since mid-2000s. For example, out of 20 rate increases and 13 rate cuts which happened in the USA in the bespoke period, 19 and 10 respectively could be correctly predicted by the OIS-implied rates twelve days in advance, which is on par with the Federal funds futures scoring 19 and 11 respectively.

Still, as we use only the market prices of one certain instrument, our forecasts are based on an information set that is surely narrower than that of the real-world market participants.¹ Thus our results are likely to be conservative.

Our paper thus extends the strand of literature on responses of asset prices to macroeconomic announcements. For the stock market, Lucca and Moench (2015) find strong positive returns of the S&P500 index around FOMC announcements. In contrast to the main finding of their paper, we show that exchange rates do not respond to the upcoming rate hikes and cuts in the same manner. Cieslak et al. (2016) report that the stock returns in the US are cyclical and centered on the FOMC meetings. For bonds, Hördahl et al. (2015) investigate the movements of the yield curve after the release of major U.S. macroeconomic announcements, and Kontonikas et al. (2016) study the

¹Additional sources of information are the prices of other fixed income derivatives, analysts’ surveys and often the regulators’ own words. For example, Norges Bank adds monetary policy projections with an own view on the future policy rates in its quarterly reports.

dynamics of the corporate bond returns after monetary policy shocks. For the FX market, the above mentioned papers by Mueller et al. (2017) and Karnaugh (2016) are the major references. In contrast to many of these papers however, ours covers only target rate decisions among all monetary policy announcements and is silent about unconventional policy tools. Although the latter have enjoyed elevated popularity since 2008 and although the FX patterns that we document might actually be shaped by announcements of policy easing and tightening, which include but are not limited to target rate changes, predicting the sentiment of a generic announcement is much more difficult and subject to data manipulation than predicting a rate hike or cut. In a way, we concentrate on a subset of monetary policy announcements, arguably the most important one, while leaving the rest for future research.

What could explain our results? Standard asset pricing theory links excess returns to systematic risks which can not be diversified away thus commanding a risk premium. Policy announcements provide markets with information about authorities' future actions. Recent theoretical models of Ai and Bansal (2016) and Pástor and Veronesi (2013) tie these information releases to the risk premium compensating investors for uncertainty regarding the path of the future policy.

It is difficult to reconcile our findings with these risk-based explanations: first, we show that excess returns earned before the announcement day dwarfs the announcement day returns documented by Mueller et al. (2017); second, the pre-announcement drift in exchange rates does not appear before the announcements at which no policy rate change was implemented; third, our finding of monetary policy shifts being highly predictable leaves little room for the uncertainty resolution argument.

Alternative theories feature inattentive investors, infrequent rebalancing decisions and other impediments to perfect markets. Duffie (2010) develops a limited participation model with heterogeneous agents where the "inattentive" investors trade less frequently than "professional intermediaries". In this setup, the aggregate level of risk does not change before scheduled events, unlike its distribution among the investor types, with intermediaries bearing a larger share, thus demanding compensation for the risk. As Lucca and Moench (2015) point out, it is not clear in the setup of Duffie (2010), why it would be optimal for inattentive investors to sell their positions out to intermediaries instead of maintaining their holdings and reaping the premium.

Bacchetta and Van Wincoop (2010) present an overlapping generations model where infrequent rebalancing decisions stem from the costs of active portfolio management. In their setup, agents optimally stick to passive currency management if costs of active management are prohibitively high. The infrequent rebalancings in turn lead to the

delayed exchange rate overshooting with depreciation of foreign currency over several periods after an interest rate cut implemented by the foreign central bank. Although the setup of Bacchetta and Van Wincoop (2010) helps to rationalize the persistence in currency returns, it does not explain why the drift appears before changes in interest rates.

The rest of the paper is organized as follows. Section II outlines the methodology of event studies, policy expectations recovery and trading strategy construction; Section III summarizes the data used; Section IV presents our findings; Sections V concludes.

II. Methodology

This section describes the empirical design of our study. First, we outline the methodology of event studies in a multicurrency framework. Then, we discuss the payoff structure of overnight index swaps and federal funds futures, and describe the techniques to extract the implied future interest rates. The section concludes with the description of the spot and excess returns of a trading strategy and related costs.

A. Event Study

To detect the pre-announcemnt drift in the currency markets, we use an event study framework.

Event studies in finance have not changed much since Fama et al. (1969). In our case the test assets are exchange rates, and the events are monetary policy announcements of respective regulators, such that each test asset is associated with multiple events. Two choices are important in the design of any event study: of the event window span, and of the model for what is considered “normal” as opposed to “abnormal”.

The former choice is dictated by the possible duration of the exercised effect and by the necessity to retain an “uncontaminated” portion of the sample for inference purposes. Mostly interested in the pre-announcement dynamics of the assets, we choose the period of 10 days before and 5 days after each announcement as the event window, using the rest of the sample for estimation. We also exclude the event day from both samples, thus differentiating between the pre-event and post-event windows.

We use the constant mean model discussed i.a. in Brown and Warner (1980) as the model for the “normal” currency returns, the mean being zero. This way, abnormal

returns are the same as returns. We will briefly discuss the quality of this model towards the end of this subsection.

Define $d_{i,k}$ to be the date of announcement $k \in \{1, \dots, K\}$ relating to currency $n \in \{1, \dots, N\}$. As discussed above, the event window spans w_b days before and w_a days after $d_{i,k}$. We cut the series of (dollar) returns of currency i into k subsamples of length $w_a - w_b$). We reindex these subsamples to have incremental ordinal indexes

$$\{s\} = \{w_b, \dots, -1, 0, +1, \dots, w_a\},$$

understood to denote s days after an event: for example, the day of event will have index 0, and the day corresponding to two days before it will have index -2 .

A cumulative abnormal return (CAR) is defined as:

$$R_{i,k,s}^{ca} = \begin{cases} \sum_{t=s}^{-1} R_{i,k,t} & s < 0, \\ \sum_{t=+1}^s R_{i,k,t} & s > 0, \end{cases} \quad (1)$$

such that the s -period CAR before an event is understood to be realized by buying the currency in period $-s$ and selling it in period -1 ; the return after an event is realized by buying the currency in period 1 and selling it in period s after the event. In what follows we will concentrate our attention on the pre-announcement returns.

The average cumulative abnormal return is defined as the average over events and over currencies of the cumulative return in equation (1):

$$\bar{R}_s^{ca} = \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K R_{i,k,s}^{ca} \quad (2)$$

Appendix A shows that \bar{R}_s^{ca} is approximately normally distributed with mean zero and variance defined therein.

The assumption that the log-returns of exchange rates are a zero-mean process is necessary because their true mean cannot be precisely estimated on the sample of 16 years that we have, let alone on the shorter subsamples between consecutive events. Still, even when looking at longer datasets, spot returns appear close to driftless, indistinguishable from such at the standard significance levels. Additionally, we can in part account for the possible misspecification by incorporating the zero-mean assumption into the variance formula in Appendix A.

B. Recovering Implied Rates

The literature on assessing the expectations about future monetary policy actions from observable asset prices is vast: for example, Krueger and Kuttner (1996), Kuttner (2001) and Karnaukh (2016) use the federal funds futures, Cochrane and Piazzesi (2002) employ the one-month eurodollar deposit rate. Gürkaynak et al. (2007) compare the predictive power of rates implied by a variety of traded assets in forecasting future monetary policy actions in the US. Our contribution to this strand of literature is two-fold. First, the empirical evidence on predictability of target rate changes primarily considers the United States. We find that the changes in policy rates are also predictable in the major economies outside the US. Second, Gürkaynak et al. (2007) report the federal funds futures to provide the best market-based measure of near-term monetary policy expectations. Since the federal fund futures contracts are unique to the United States, we recover expected policy rates from the overnight index swaps (OIS) which so far did not receive much attention in the literature on policy rates prediction, despite they and their underlying rates have been gaining popularity in derivative pricing and monetary policy practice.² We show the OIS-implied rates to be accurate predictors of the future monetary policy actions in the other countries, performing on par with the federal funds futures in the US. In the rest of this section we describe the payoff structure and extraction of the expected future policy rates from the federal funds futures and OIS contracts.

Overnight index swaps (OIS) are fixed/floating interest rate swaps where the floating leg pays the cumulative return on an underlying rate, e.g. the effective federal funds rate in the US or the SONIA in the UK. At the settlement day T the payoff of the floating leg of an OIS with notional amount of \$1 and start date tomorrow (day 1) is:

$$\pi^T = \prod_{t=1}^T (1 + r_t) - 1, \quad (3)$$

where t is the first day of the swap, r_s is the annualized underlying rate. The buyer will pay a fixed rate called the swap rate w_t , which is known at the inception of the swap, so the net payoff at maturity equals π^T .

In the absence of arbitrage opportunities, the price of the swap³ today (day 0) with

²For example in April 2017 the Bank of England recommended SONIA as the sterling near risk-free reference rate benchmark, furthermore Hull and White (2013) argue that for derivatives pricing OIS rates are superior to the traditional LIBOR rates.

³The actual prices are quoted in annualized terms, but we use rates per period equal to the maturity of the contract (e.g. monthly) to avoid cumbersome formulas.

start date tomorrow (day 1) is equal to the risk-neutral expectation of (3):

$$w_0 = E_0 [\pi^T] = E_0 \left[\prod_{t=1}^T (1 + r_t) - 1 \right], \quad (4)$$

where the expectation is taken under the risk-neutral measure. Let us assume a policy meeting takes place at date t^* , and the rate r^* announced at the meeting becomes effective at $t^* + 1$. We also assume the current rate stays constant until the announcement, and the rate then set prevails from the effective date until the expiration of the contract. That said, equation (4) can be rewritten as:

$$\begin{aligned} w_0 &= E_0 \left[\prod_{s=1}^{t^*} (1 + r_0) \prod_{t=t^*+1}^T (1 + r^*) - 1 \right] \\ &= (1 + r_0)^{t^*} E_0 [(1 + r^*)^{T-t^*}] - 1, \end{aligned} \quad (5)$$

Neglecting the Jensen's inequality, we arrive at the expected rate at the announcement date:

$$E_0 [r^*] = \left((w_0 + 1)(1 + r_0)^{-t^*} \right)^{\frac{1}{T-t^*}} - 1 \quad (6)$$

Federal funds futures are traded on the Chicago Mercantile Exchange (CME) and pay the average effective federal funds rate over the month at the corresponding month's end with the rate being carried forward over weekends. The payoff from holding a futures for delivery in month m is thus:

$$\pi^m = \frac{1}{T_m} \sum_{s \in m} r_s, \quad (7)$$

where T_m is the number of calendar days in month m . Two major advantages of these contracts are their high liquidity and zero counterparty risk because of the daily marking-to-market. Krueger and Kuttner (1996) and Gürkaynak et al. (2007) find the futures-implied rate to be an accurate predictor of the near-term monetary policy shifts in the US.

Similarly to OIS we start with the time t risk-neutral price of the federal funds futures contract with delivery in month m :

$$f_t^m = \frac{1}{T_m} E_t \left[\sum_{s \in m} r_s \right], \quad (8)$$

Assuming that the Fed funds rate on average remains at the same level between consecutive FOMC meetings, it is straightforward to extract the expectation of the rate set

at the next meeting. Since there are 8 meetings in a year, two scenarios are possible before any meeting k taking place in month m : either the next calendar month will witness another meeting $k + 1$, or the next month is “free” of meetings. In the second case the expected rate set at meeting k is the price of the futures contract expiring in the month $m + 1$. Otherwise the expected rate is a combination of the settlement price of the previous contract and today’s price of this month’s contract:

$$E_t [r^k] = \begin{cases} 100 - f_t^{m+1}, & (k+1) \notin (m+1) \\ \frac{T_m}{T_m-t} \left(f_t^m - \frac{t}{T_m} f_{T_{m-1}}^{m-1} \right), & (k+1) \in (m+1) \end{cases} \quad (9)$$

C. Trading Strategy

We construct a simple trading strategy based on expected shifts in policy rates. Assuming a US investor’s perspective, for a foreign central bank’s target rate decision announced on day T we forecast the new policy rate on day $T - h - 2$, and establish a position in the corresponding currency at the end of the next day $T - h - 1$ to avoid any potential overlap between interest rate derivatives and currencies. The position is then held for h days and liquidated one day before the announcement at $T - 1$. Should a rate hike be expected, we go long in the foreign currency vs. USD, should a rate cut be expected, we go short in the foreign currency vs. USD, and open no position otherwise. The log spot return over h periods realized at time $T - 1$ is therefore:

$$R_{T-1}(h) = d_{T-h-2} \sum_{t=T-h}^{T-1} (r_t) = d(h)r(h), \quad (10)$$

where r_t is the daily currency log spot return and d_{T-h-2} is a categorical variable, capturing the $T - h - 2$ expectation of the policy rate change on the announcement day and is equal to 1 if a hike is expected, -1 if a cut is expected and 0 otherwise. Conversely for the FOMC announcements we buy (sell) USD against an equally-weighted portfolio of currencies – the dollar index – if increase (decrease) in the federal funds rate is expected.

We recover the expected policy rates from the OIS contracts and federal funds futures using equations (6) and (9). With an exception of the US, the underlying rates for OIS differ from the policy rates set by central banks, hence the derivatives-implied expectations of the latter can be in addition to time-varying risk premia⁴ contaminated

⁴Although given our short policy rate forecast horizons the risk premium is of a lesser concern, Piazzesi and Swanson (2008) document the predictable time-varying risk premium in the federal funds futures of maturities higher than one month.

with noise. To address this issue, we define the expected change in the target rate $E_{T-h-2} [\Delta i_T]$ as the difference between the derivatives-implied rate expected to prevail after the announcement and the corresponding underlying rate with both rates averaged over the five preceding days.⁵ We further employ a simple rule to evaluate the expected shift in the target rate by defining the categorical variable d_{T-h-2} as:

$$d(h, \tau) = \begin{cases} 1, & \text{if } E_{T-h-2} [\Delta i_T] > \tau, \text{ rate hike expected} \\ 0, & \text{if } |E_{T-h-2} [\Delta i_T]| \leq \tau, \text{ no change expected} \\ -1, & \text{if } E_{T-h-2} [\Delta i_T] < -\tau, \text{ rate cut expected} \end{cases}$$

where τ is a threshold level. Denote $a = 1, \dots, A$ to be a chronological sequence of all policy rate announcements for every currency, the cumulative US dollar return on the aggregate strategy as of announcement a can be written as:

$$R_a(h, \tau) = \sum_{a=1}^A [d_a(h, \tau) r_a(h)], \quad (11)$$

Throughout this paper we employ the holding period and threshold of 10 days and 10 basis points as the baseline values. We further demonstrate that our results are robust to the variation in these parameters.

The strategies constructed this way admit cross-sectional leverage: if signals in different countries are separated by a period shorter than the holding period, we do not split the invested capital, but multiply it. This is a computationally convenient and realistic setup given the preponderance of leveraged transactions on the FX markets. The US Commodity Futures Trading Commission allows for a 50:1 leverage in the off-exchange retail FX forex trading,⁶ which corresponds to a possibility of opening 50 positions in our setup. The average leverage for the baseline strategy that we construct is 1.38, and it is less than or equal to 2 on 94% of all days. In Appendix D we provide a detailed description of how deleveraged strategies are constructed, and show that our findings are robust to restricting leverage.

A zero-cost foreign exchange trading requires investors to pay (or receive) the interest rate differential between the base currency and the counter currency. A common way in the academic literature to calculate the h -period excess return is to take the difference between the (log of) spot price in period $t + h$ and the price of a forward contract

⁵The choice of the smoothing window is inconsequential for our results.

⁶Or 30:1 leverage on-exchange currency futures trading, 50:1 in the commercial bank forex trading and 200:1 in the offshore off-exchange retail forex trading.

with maturity h opened in period t :

$$rx_{t+h} = \log S_{t+h} - \log F_{t \rightarrow h}, \quad (12)$$

$$= \underbrace{\log S_{t+h} - \log S_t}_{\Delta s_{t+h}} + \underbrace{\log S_t - \log F_{t \rightarrow h}}_{d_t}, \quad (13)$$

where Δs_t is the spot return, and d_t should under the Covered interest parity be equal to the interest rate differential. However, as forward prices are readily obtainable for a limited number of maturities only (e.g. one week, two weeks, etc.), but our task is to construct a strategy with the holding period of *several days*, we turn to the foreign exchange swaps which allow to earn the interest rate differential on the daily basis. In fact, foreign exchange swaps are the most traded instrument on the FX market, with the turnover in short-term swaps (maturity of under seven days) and spot transactions approaching USD 1.6 trillion for each of the instruments, exceeding the daily turnover of forward contracts of *any* maturity by a factor of two (BIS (2016)). Most of the FX positions are usually opened out of speculative interest and eventually reversed before the actual delivery of the transacted currency takes place. Until then every position kept open at 5pm New York time is being rolled over: the delivery is then postponed by one day, and the price of the contract is adjusted by adding the tom/next swap points. The tom/next swap points are closely linked to the interest rate differential between the two legs of the FX position and are positive (negative) if the interest rate in the base currency is lower (higher) than that in the counter currency, in which case the holding period return on the position rolled over falls (rises) *ceteris paribus*.

Now, imagine postponing the delivery for h periods: in this case the end-of period log-return is:

$$rx_{t+h} = \log S_{t+h} - \log(S_t + \sum_{\tau=1}^h w_\tau), \quad (14)$$

where w_τ is the tom/next swap points. Seen at time t , the same return is expected to be achieved by selling short an h -day forward and closing the position at its expiration. Hence, the change in the opening price by the time the position is closed can be *ex ante* thought of as the forward premium or discount, which brings us back to equation (13). In Appendix C we discuss the plausibility of this approach and compare it on the monthly frequency to the more common technique in (13).

III. Data

In this section we describe our dataset. First we provide a brief overview of monetary policy implementation procedures across the major economies, then we describe our currency and fixed income data.

A. *Announcements of Central Banks*

In the 1990s central banks started to adopt the policy of announcing target interest rate changes on pre-scheduled dates. We collect data on policy rate announcements for the following countries: United States, United Kingdom, Australia, Canada, New Zealand, Switzerland, Sweden, Norway and the Eurozone. Our sample spans the period from November 2000 to March 2017. By November 2000, all countries in the sample adopted the practice of interest rate announcements on pre-scheduled dates. We do not include Japan since the Bank of Japan has been switching between various monetary policy tools over the past 20 years.⁷ The targets and announcement schedules, however, have been different across the central banks:

Australia. The Reserve Bank of Australia began to announce the target rate decisions on pre-scheduled dates in 1981. The monetary policy meetings usually occur eleven times a year. Between 1990 and 1996 the Bank changed the Cash rate on 21 occasions from which ten cuts and two hikes were implemented outside the scheduled Board's meetings. There were two further unscheduled cuts in 1997. Until 1998, from time to time the Board gave the Governor discretion to implement a change in the cash rate in an agreed manner. From 1998 onwards, the Bank sticks to its schedule of announcing decisions on the first Tuesday of each Month except January. Before 2008 RBA announced the interest rate decision on the day following the meeting day simultaneously with the new policy coming into effect. Starting from 2008, the decision is announced on the meeting day and becomes effective on the following day.

Canada. The Bank of Canada introduced pre-scheduled interest rate announcements

⁷On March 19th 2001 the Bank of Japan abandoned targeting of the uncollateralized overnight call rate (MUTAN), leaving the rate to be determined by the market. The MUTAN was expected to be capped from above by *the official discount rate* on the Lombard-type lending facility where eligible financial institutions could receive loans posting eligible collateral. Simultaneously the main operating target for monetary policy was changed to current accounts at the Bank of Japan. Subsequently the Bank resumed targeting the average call rate on March 9th 2006, switching to a band on October 5th 2010, and abandoning once again the interest rate targeting in favor of the monetary base targeting on April 4th 2013. Finally, the Bank introduced negative interest rates on the current accounts on January 29th 2016 (effective from February 16th) and "yield curve control" on September 21st 2016 as additional policy measures. See also Kuttner (2014) for a comprehensive overview of Japan's monetary policy from 1980 to 2012.

in November 2000. The announcements take place eight times a year with decision becoming effective on the announcement day.

Eurozone. The European Central Bank (ECB) held a monetary policy meeting twice a month from 1999 to 2001, then once a month from 2002 to 2015, switching to a six-week cycle in 2015. The ECB targets three rates: (i) the deposit facility which allows banks to place deposits at the ECB; (ii) the marginal lending facility which offers overnight loans to the Eurozone's banking system; (iii) the main refinancing operations (or MRO) rate at is the rate at the ECB injects and withdraws liquidity using repo operations, normally, with a maturity of one week. The Bank announces its interest rate decisions on the meeting day, the changes in policy become effective on the day set at the meeting, usually from the next day to a week.

New Zealand. The Reserve Bank of New Zealand announces its Official Cash Rate on pre-scheduled meetings since April 1999. The bank holds around eight policy meetings a year, with the interest rate decisions becoming effective on the announcement day.

Norway. Norges Bank started to announce interest rate decisions on pre-scheduled meetings on June 16th 1999. The meetings took place once a month until June 2000 when the monetary policy meetings began to occur once every six weeks. The decision is normally announced on the day of the meeting and becomes effective on the next day.

Sweden. The Riksbank adopted the policy rate announcements on pre-scheduled meetings on October 6th 1999, with the first meeting in the February 2000. Since then and until 2008 the Bank held monetary policy meetings once every six to eight weeks. From 2008 onwards the Riksbank holds six ordinary monetary policy meetings per year. The decision is normally announced on the day following the day of the meeting and becomes effective in a week.

Switzerland. In contrast to other central banks mentioned here which target overnight rates, the Swiss National Bank operates on the higher maturity region of the yield curve, targeting the 3-month Swiss Franc Libor. Since 2000 the Bank abandoned money supply targeting in favor of interest rate targeting. Policy meetings take place four times a year with decision becoming effective immediately. From September 2011 to January 2015 the SNB focused its monetary policy on sustaining the exchange rate cap to the euro.

United Kingdom. In June 1998 the Bank of England received autonomy over the monetary policy. The Bank's Monetary Policy Committee (MPC) held meetings every month

until September 2016, since then the official interest rate is reviewed eight times a year. The interest rate decision is announced on the day following the MPC meeting day and comes into effect on the next day.

United States. Since February 1994, the Federal Open Markets Committee (FOMC), a part of the Federal Reserve System overseeing the monetary policy in the United States, announces its decisions on eight pre-scheduled meetings a year. The target range for the Federal funds rate is announced on the second day of the meeting and becomes effective on the following day. For a detailed description of the FOMC meetings and statement releases see e.g. Lucca and Moench (2015) and references therein.

Table I reports summary of the scheduled policy announcements for the central banks discussed above. The second and third columns show the fixed announcement schedule adoption date and the key policy rates respectively. The last three columns report the total number of announcements and the numbers of hikes and cuts in policy rates of each central bank. The joint sample is from November 2001, when the Bank of Canada adopted the fixed schedule, to March 2017. The period of Swiss franc – euro cap (from September 2011 to January 2015) is excluded for Switzerland. The total numbers of hikes and cuts are 155 and 180 respectively, resulting in the sample size well above the total number of all events for the FOMC announcements considered in the previous literature.

We further consider the scheduled monetary policy meetings only, although some extraordinary meetings became known to market participants well in advance (e.g. the meeting of Norges Bank on October 15, 2008 was announced on October 8th). First, the policy actions undertaken during unscheduled meetings constitute a small fraction of all target rate changes.⁸ Second, we aim to keep our results conservative and robust to outliers by ruling out extreme events like the September 2001 terrorist attacks and the coordinated interest rate cut by a number of central banks on October 8th 2008.

[Table 1 about here.]

B. Exchange Rates and Currency Returns

We use Bloomberg daily spot exchange rates against USD for the following countries: Australia, Canada, Japan, New Zealand, Norway, Sweden, Switzerland, United Kingdom and the Eurozone. We collect the quotes for different fixing times to ensure that

⁸With a notable exception of Switzerland, where roughly three quarters of the target rate changes from 2000 to 2017 were implemented during unscheduled meetings.

the announcement day is not overlapped for any of the currencies. Thus we use 5pm London fixing time for the Eurozone, Norway, Sweden, Switzerland, and the United Kingdom; 5pm New York time for Canada and the US; 8pm Tokyo time for Australia and New Zealand. Respective bid and ask prices are used to adjust for the trading costs. The long and short tom/next swap points are also from Bloomberg: except for AUD, EUR, GBP and NZD, these are quoted as units of foreign currency per unit of USD, such that we have to convert them first to conform with the perspective of a US investor.

For the FOMC announcements we construct the dollar index – an equally weighted portfolio of currency returns against USD, with each currency, including JPY, fixed at 5pm New York time.

C. *Overnight Index Swaps and Federal Funds Futures*

We collect 1-month swap rates from Bloomberg and use rates from Datastream where Bloomberg quotes are unavailable. The availability of the OIS data is as follows: the Eurozone since January 1999; United Kingdom and Switzerland since late 2000; Australia, Canada, and the US since August 2001; New Zealand and Sweden since September 2002.

The overnight rates underlying the OIS are the federal funds effective rate for the US, SONIA for the UK, RBA Cash Rate for Australia, Official Cash Rate for New Zealand, CORRA for Canada, TOIS fixing for Switzerland, STIBOR for Sweden, and EONIA for the Eurozone.⁹ The data on these rates is from Bloomberg.

In order to assess predictive power of the OIS-implied rates we also collect the data on the federal funds futures contracts considered to be staple in the literature. This data comes from the Chicago Mercantile Exchange.

IV. Results

In this section we present the empirical results of the paper. We begin with documenting a pre-announcement drift in currency returns preceding shifts in monetary policy around the world. We then demonstrate that this drift is exploitable by investors first, by showing that monetary policy actions are predictable and second, that a trading

⁹There are no overnight interest rates data available for Norway.

strategy aiming to forecast future monetary policy action and then buy (sell) currencies whose monetary authorities are expected to raise (cut) their policy rates earns substantial returns.

A. *Drift in Spot Exchange Rates Before Announcements*

Figures 1 and 2 depict the results of the event study, with events being announcements of the local central banks to raise and cut the target interest rate respectively, and the test assets being spot returns of the currencies of the corresponding countries. As shown in the bottom panel of Figure 1, a randomly selected currency before a randomly selected rate hike is expected to appreciate by 50 bps over ten days, 30 bps over five days, and 10 bps on the day before the event day. The individual cumulative spot returns are presented in the upper panel, making it evident that of all currencies, only the Swedish krona and the Swiss frank slightly depreciate on average over the ten-day period ahead of rate hikes. The pattern is reversed before rate cuts, as can be seen in the lower panel of Figure 2: in this case, a currency is expected to depreciate by about 70 bps over ten days, half that over five days, and 10 bps on the pre-announcement day. The average return is significant at the 5% level for periods of all lengths.

[Figure 1 about here.]

[Figure 2 about here.]

As seen in the figures, currencies experience a statistically significant and economically large drift in the direction of the policy rate changes, more pronounced in the case of rate cuts. The spot exchange rates begin to move at least ten days in advance of the central banks' announcements. Interestingly, the drift mostly dissipates, and the abnormal returns evaporate in the post-announcement period. As a robustness check, in Appendix B we redo the same exercise using two different counter currencies – GBP and JPY – rather than USD, and confirm an equally strong significant downward trend before interest rate cuts, and an upward yet insignificant trend before hikes.

In Figures 3 and 4, we contrast the observed patterns in the spot returns to those around the FOMC announcements. In general, the effect of the Fed policy rate changes on the foreign currencies is opposite in sign to that of the local rate changes: an average foreign currency tends to depreciate against the US dollar before the Fed funds rate is increased, and appreciate in the opposite scenario. However, this effect only

manifests itself over a short period of time, and no significant cumulative appreciation or depreciation can be detected earlier than four days ahead of events. The effect is also weaker economically: the average depreciation before the Fed funds rate hikes is lower in magnitude than that before the local rate cuts at any considered horizon, and the average appreciation in the opposite case is lower for 7 out of 10 horizons.

[Figure 3 about here.]

[Figure 4 about here.]

Given the magnitude of the abnormal returns and the horizon over which the drift manifests itself, the natural question is to what extent the market participants are able to exploit it in a trading strategy. To exploit the pre-announcement drift, investors need to be able to accurately predict upcoming monetary policy actions and earn significant return after accounting for transaction costs. We address this issue in the rest of the section.

B. Recovering Monetary Policy Expectations

Using the 1-month OIS and the forecast horizon of 12 days (for the 10-day holding period to be possible), we estimate the reference rates expected to prevail after each announcement. Figure 5 shows the error plots constructed thereof. The post-announcement rates can be forecast with a mean absolute error below 10 bps, the highest differences occurring for Switzerland and the Eurozone. The mean error (not reported here) rarely exceeds 1 bps and reaches the maximum of 4 bps in the case of Switzerland. As a comparison, the lower right panel depicts the forecasts of the federal funds rate calculated using the Fed funds futures: these exhibit a slightly higher mean absolute error, but overall are as strong a predictor of the policy shifts.

[Figure 5 about here.]

Being interested not in the level of implied rates per se, but rather in the direction which the implied rates imply (no pun implied), in Figure 6 we show the confusion matrices corresponding to each error plot above. We use the threshold of 10 bps to separate expected cuts from hikes, the same 12-day forecasting horizon, and 5 days to average the implied and the underlying rates. The “worst” cases of forecasting a direction opposite to the announced are almost absent in the sample: these are located

in the southwest and northeast corners of the matrices and never exceed 1. The ratio of correctly predicted directions is high, the worst being the one for rate cuts in Sweden.

Interestingly, rate cuts appear to be predictable with a lower accuracy than rate hikes. This is partly because they tend to happen in times of economic distress, when both the prices of OIS and the underlying rates become volatile and subject to large risk premia, such that the forecasts get distorted.

As in Figure 5, the bottom right panel refers to the Fed funds futures-based predictions. Since 2001, just one more cut was correctly predicted by the Fed funds futures.

[Figure 6 about here.]

Overall, using the information implied in the OIS rates to predict the upcoming monetary policy decisions is justified *ex post* by low absolute errors and a high percentage of correctly captured change directions. Not reported here are the outcomes of the forecasting exercise with different values of the forecasting horizon and threshold. In general, the prediction accuracy increases as the horizon shrinks (and vice versa).

C. *Is the Pre-Announcement Drift Exploitable by Investors?*

We start with spot returns thus recasting the results of the event study in the beginning of this section as a trading strategy. Figure 7 plots the cumulative performance of a strategy in which the investor goes long in currencies whose monetary authorities are expected to raise the policy rate, and short in currencies with expected interest rate cuts. Panel 7a shows the return plotted against time and Panel 7b shows the return plotted event-by-event. The solid line represents cumulative return of the forecast-based strategy, while the dashed line represents cumulative return of an investor with perfect foresight. The investor makes a decision whether to open a position twelve days ahead of the announcement. The rate change is forecast as the difference between the implied post-announcement rate extracted from the OIS and the underlying rate with both rates averaged over the five previous days. The investor establishes a position only if this difference exceeds a threshold of ten basis points in absolute value. For each predicted target rate change the FX position is held for ten days and liquidated one day ahead of the corresponding announcement. For the FOMC announcements the position in USD is established against the dollar index. The sample is from November 2001 to March 2017. The numbers in the upper panel report the mean return, its standard error (both in basis points) and the Sharpe ratio per one holding period. The

standard error is Newey and West (1987) HAC with optimal number of lags according to Newey and West (1994). The upper panel of figure 7 also plots returns of the perfect foresight strategy for the same sample (that is, the dashed line basically plots the return earned if the OIS-based forecasts were 100% accurate).

Over 16.5 years, the simple strategy based on the expected monetary policy shifts generated a total spot return of 160% with the average per-event return of 45.58 basis points (with t-statistic of over 3) and a ten-day Sharpe ratio of 0.23, outperforming its perfect foresight counterpart by approximately 10% over the course of the sample. Consistent with the event study results, the spot exchange rates tend to front-run impending rate changes, and the high predictability of monetary policy allows to profitably exploit it.

To check that the strategy performance is not shaped by a handful of extreme events, in the bottom panel of Figure 7 we plot the performance on the event line instead of the timeline. As can be seen, the strategy also delivers stable and positive returns event-by-event. Given the economic and statistical significance of our results we go on to investigate their robustness to the choice of the holding period and threshold.

[Figure 7 about here.]

In order to control for the uncertainty in the choice of the threshold and holding period and address the data snooping problem, we generate a universe of 375 trading strategies with holding periods ranging from 1 to 15 days and threshold levels ranging from 1 to 25 basis points. Figure 8 plots the results of this exercise. The fact that none of the gray lines depicting the cumulative return of different trading strategies culminates in a negative return, indicates that our finding is robust to the choice of the trading strategy parameters.

[Figure 8 about here.]

We further check if our results can be explained by the FOMC pre-announcement drift of the dollar factor documented by Karnaukh (2016). Figure 9 repeats the analysis in Figure 8 for the FOMC announcements and the dollar index only. Over the whole universe of 375 strategies buying and selling the dollar index around the US interest rates hikes and cuts, the average performance is almost exactly zero, indicating that the FOMC pre-announcement drift does not drive our results and making our evidence qualitatively different from that in previous studies.¹⁰

¹⁰Similar to Karnaukh (2016) we observe economically and statistically significant pre-announcement

[Figure 9 about here.]

Now, we bring the trading strategy closer to a real-life application. First, we account for the bid-ask spread by opening long positions at the ask and short ones at the bid price. Second, we make all open positions subject to rollovers at the end of the trading day.

Figure 10 plots the performance of the baseline strategy with the holding period of ten days and the threshold level of ten basis points. The cumulative return of the forecast-based strategy drops by approximately 20 percentage points to 140%, and per-event return falls to 38 basis points, remaining statistically significant at the 1% level. A similar reduction is observed for the strategy based on the perfect target rate change predictions. The bottom panel plots performance event by event. Similarly to the results reported for the spot rate, the performance is not driven by a small number of outliers: the surge in returns during the acute stage of the 2007–2009 financial crisis merely reflects the correctly predicted worldwide target rate cuts accompanied by depreciation of the corresponding currencies. Simply discarding the period of the apparent (in the time domain) spike, that is from June 2008 to June 2009, reduces the per-event return and its standard error to 24.19 and 10.74 basis points respectively, and the ten-days Sharpe ratio to 0.14.

Figure 11 shows the performance of 375 trading strategies across various holding periods and thresholds similar to those in figure 8. After accounting for rollovers and transaction costs, trading strategies deliver on average 20 percent lower cumulative return in comparison with their frictionless spot-based counterparts. In fact, the five strategies that have generated a negative cumulative return over the course of the sample are confined to the five lowest threshold values with holding period of one day, implying high occurrence of false positive interest rate change predictions and, hence, higher impact of transaction costs due to increased turnover. Similarly to the results in Figure 9, we do not find significant pre-FOMC-announcement drift in bid-ask spread-adjusted excess returns when controlling for uncertainty in choice of holding horizon and threshold level.

[Figure 10 about here.]

[Figure 11 about here.]

drift for a number of strategies trading the dollar index around FOMC announcements, primarily with short holding periods, it is unclear however whether investors could have learned the corresponding holding period and threshold values.

Overall, the trading strategy exercise provides evidence of the short-horizon predictability of currency returns. This drift can not be attributed to the behavior of the dollar index before the target rate announcements documented in the previous literature. Furthermore, we demonstrate that the drift is perfectly exploitable by investors who face transaction costs and have to roll their spot positions overnight. It is important to point out, that some of the strategies admit leverage due to overlapping holding periods among predicted target rate changes for different currencies. In Appendix D we further demonstrate that our results also robust to restricting leverage.

V. Conclusion

We describe a persistent pattern in the dynamics of exchange rates before policy rate announcements of respective central banks: currencies start to moderately appreciate days before declared interest rate hikes, and significantly depreciate before rate cuts. Given that a transparent monetary policy favored by most regulators since 2000s begets a high predictability of policy rate changes, we show that the pattern is profitably exploitable on the FX market. We document that policy rate decisions can be accurately forecast with information embedded in overnight index swaps, more so when the best classification rule is known in advance. However, the multitude of possible classification rules makes it difficult to accurately backtest trading strategies. We show that the final payoff of the strategies can be sensitive to the choice of the rule. Still, the payoff of the trading strategy that we construct using the OIS-implied information and a cross-section of currencies remains positive and large whatever the specification.


Our findings are difficult to reconcile with the existing theories of the determinants of exchange rates. Robust returns of the pre-announcement trading that we see might be a consequence of a gradual resolution of uncertainty about the approaching policy change and heterogeneous agents entering the currency market one by one as soon as their risk aversion allows to place a bet. The more risk averse investors would in this case enter the market last, when the monetary policy uncertainty is low, and the less risk averse ones would enter earlier, thus constantly buoying the demand for the currency. An attack at modeling the mechanism behind our findings would be a logical continuation of the research on the dependency between monetary policy and exchange rates.

References

- Ai, Hengjie, and Ravi Bansal, 2016, Risk preferences and the macro announcement premium, Technical report, National Bureau of Economic Research.
- Akram, Q Farooq, Dagfinn Rime, and Lucio Sarno, 2008, Arbitrage in the foreign exchange market: Turning on the microscope, *Journal of International Economics* 76, 237–253.
- Bacchetta, Philippe, and Eric Van Wincoop, 2010, Infrequent portfolio decisions: A solution to the forward discount puzzle, *The American Economic Review* 100, 870–904.
- Backus, David K., Federico Gavazzeni, Christopher Telmer, and Stanley E. Zin, 2010, Monetary policy and the uncovered interest parity puzzle, Working Paper 16218, National Bureau of Economic Research.
- Bailey, David H., and Marcos Lopez de Prado, 2014, The deflated sharpe ratio: Correcting for selection bias, backtest overfitting and non-normality, *Journal of Portfolio Management* 40, 94–107.
- BIS, 2016, *Triennial Central Bank Survey of Foreign Exchange and Derivative Market Activity in 2016*, Bank for International Settlements, Basel.
- Brown, Stephen, and Jerold B. Warner, 1980, Measuring security price performance, *Journal of Financial Economics* 8, 205–258.
- Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, 2016, Stock returns over the fomc cycle .
- Cochrane, John H, and Monika Piazzesi, 2002, The fed and interest rates: A high-frequency identification, *The American Economic Review* 92, 90–95.
- Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2017, Deviations from covered interest rate parity, Technical report, National Bureau of Economic Research.
- Duffie, Darrel, 2010, Asset price dynamics with slow-moving capital,(american finance association presidential address), *Journal of Finance* 65, 1238–68.
- Fama, Eugene F., Lawrence Fisher, Michael C. Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information, *International Economic Review* 10, 1–21.
- Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson, 2007, Market-based measures of monetary policy expectations, *Journal of Business & Economic Statistics* 25, 201–212.


- Hördahl, Peter, Eli M Remolona, and Giorgio Valente, 2015, Expectations and risk premia at 8:30am: Macroeconomic announcements and the yield curve .
- Hull, John C, and Alan White, 2013, Libor vs. ois: The derivatives discounting dilemma, *Journal of Investment Management* 11, 14–27.
- Karnaukh, Nina, 2016, The dollar ahead of fomc target rate changes, *Job Market Paper* .
- Kontonikas, Alexandros, Paulo Maio, and Zivile Zekaite, 2016, Monetary policy and corporate bond returns, Working papers.
- Krueger, Joel T, and Kenneth N Kuttner, 1996, The fed funds futures rate as a predictor of federal reserve policy, *Journal of Futures Markets* 16, 865–879.
- Kuttner, Kenneth N, 2001, Monetary policy surprises and interest rates: Evidence from the fed funds futures market, *Journal of monetary economics* 47, 523–544.
- Kuttner, Kenneth N, 2014, Monetary policy during Japan’s great recession: From self-induced paralysis to rooseveltian resolve, *PIIE Briefing* 14, 66–79.
- Lucca, David O, and Emanuel Moench, 2015, The pre-fomc announcement drift, *The Journal of Finance* 70, 329–371.
- Mueller, Philippe, Alireza Tahbaz-Salehi, and Andrea Vedolin, 2017, Exchange rates and monetary policy uncertainty, *The Journal of Finance* .
- Newey, Whitney K, and Kenneth D West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Newey, Whitney K, and Kenneth D West, 1994, Automatic lag selection in covariance matrix estimation, *The Review of Economic Studies* 61, 631–653.
- Pástor, Lúboš, and Pietro Veronesi, 2013, Political uncertainty and risk premia, *Journal of Financial Economics* 110, 520–545.
- Piazzesi, Monika, and Eric T Swanson, 2008, Futures prices as risk-adjusted forecasts of monetary policy, *Journal of Monetary Economics* 55, 677–691.

Figure 1: Exchange rates around local interest rate hikes.




This figure depicts cumulative currency returns around interest rate hikes announced by the local central banks. Panel 1a shows returns on individual currencies and Panel 1b shows the average return over all currencies weighted in proportion to the number of hikes each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in USD. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from November 2000 to March 2017, thus covering a total of 135 hikes.

Figure 2: Exchange rates around local interest rate cuts.




This figure depicts cumulative currency returns around interest rate cuts announced by the local central banks. Panel 2a shows returns on individual currencies and Panel 2b shows the average return over all currencies weighted in proportion to the number of cuts each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in USD. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from November 2000 to March 2017, thus covering a total of 162 cuts.

Figure 3: Exchange rates around the Fed funds rate hikes.




This figure depicts cumulative currency returns around the Fed funds rate hikes announced by the FOMC. Panel 1a shows returns on individual currencies and Panel 1b shows the average return over all currencies weighted in proportion to the number of hikes each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in USD. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from November 2000 to March 2017, thus covering a total of 18 hikes.

Figure 4: Exchange rates around the Fed funds rate cuts.




This figure depicts cumulative currency returns around the Fed funds rate cuts announced by the FOMC. Panel 2a shows returns on individual currencies and Panel 2b shows the average return over all currencies weighted in proportion to the number of cuts each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in USD. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from November 2000 to March 2017, thus covering a total of 18 cuts.

Figure 5: Forecasting interest rates.




This figure compares the expected reference rates recovered before announcements to the actual post-announcement rates. We use 1-month OIS rates and the forecast horizon of 12 days to recover the implied rates. They are compared to a 12-period average of the post-announcement rates. The x -axis keeps the expected, and the y -axis – the realized rates, in percent p.a. The value reported in the lower right corner of each subplot is the mean absolute forecast error, in basis points. In the lower right panel the OIS as the material for recovering the expectations are substituted with the Fed funds futures. The sample period is different for each currency.

Figure 6: Forecasting policy rate decisions.




This figure shows the confusion matrices of the policy rate forecasts. A rate hike (cut) is expected 12 days before announcements if the 5-day moving average of the implied rate at that day is by 10 bps higher (lower) than the similarly smoothed reference rate. Entry (x, y) (x denotes rows) in any such matrix contains the number of cases when direction x was predicted, and direction y announced. In each matrix, the column sum is the total number of decisions to decrease the policy rate, keep it unchanged and raise respectively. Higher numbers are highlighted with a warmer color. In the lower right panel the OIS as the material for recovering the expectations are substituted with the Fed funds futures. The sample period is different for each currency.

Figure 7: Pre-announcement trading (spot returns): cumulative performance.




This figure depicts the cumulative return on a trading strategy buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts). The position is established 11 days in advance of each announcement day if the forecast interest rate change exceeds 10 basis points in absolute value. The position is then held for 10 days and liquidated on the day preceding the announcement day. The rate change is forecast 12 days before the announcement day as the difference between the OIS-implied rate averaged over the five previous days and the corresponding underlying rate averaged over the same horizon. Panel 7a shows the return plotted against time and Panel 7b shows the return plotted event-by-event. The numbers in Panel 7a are mean return, its standard error (both in basis points) and the Sharpe ratio per one holding period. The standard error is Newey and West (1987) HAC with optimal number of lags according to Newey and West (1994). The returns are spot returns in USD on the following currencies AUD, CAD, CHF, EUR, GBP, NZD, SEK, and the dollar index. The sample is from November 2000 to March 2017.

Figure 8: Pre-announcement trading (spot returns): robustness to the choice of holding period and threshold.




This figure plots cumulative returns on 375 trading strategies buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts) for various holding horizons and expected policy rate cutoff levels. In the case of expected policy rate hike, the strategy $S^k(h, \tau)$ buys currency k against USD (or buys the dollar index for the FOMC announcements). The position is established $h + 1$ days in advance of the announcement day, only if the difference between the average OIS-implied post-announcement rate over days $h + 2, \dots, h + 6$ exceeds the average corresponding underlying rate over the same horizon by τ or more. Similarly, the currency is sold if an interest rate cut is expected and the implied rate is below the underlying rate by at least τ basis points. The position is then held for h days and liquidated on the day preceding the announcement day. The set of trading strategies (plotted in gray) is generated for $h \in [1, 15]$ and $\tau \in [1, 25] bps$, the solid black line depicts the cross-sectional mean across all trading strategies and the dashed black lines represent the 1st and 9th empirical deciles of the distribution of the cumulative returns at each point of time. The returns are spot returns in USD on the following currencies AUD, CAD, CHF, EUR, GBP, NZD, SEK, and the dollar index. The sample is from November 2000 to March 2017.

Figure 9: Pre-announcement trading (spot returns, FOMC and the dollar index only): robustness to the choice of holding period and threshold.




This figure plots cumulative returns on 375 trading strategies buying (selling) the dollar index in anticipation of interest rate hikes (cuts) in the US for various holding horizons and expected policy rate cutoff levels. In the case of expected policy rate hike, the strategy $S(h, \tau)$ buys the dollar index. The position is established $h + 1$ days in advance of the FOMC announcement day, only if the difference between the average OIS-implied post-announcement rate over the days $h + 2, \dots, h + 6$ exceeds the average effective federal funds rate over the same horizon by τ or more. Similarly, the currency is sold if an interest rate cut is expected and the implied rate is below the underlying rate by at least τ basis points. The position is then held for h days and liquidated on the day preceding the announcement day. The set of trading strategies (plotted in gray) is generated for $h \in [1, 15]$ and $\tau \in [1, 25] bps$, the solid black line depicts the cross-sectional mean across all trading strategies and the dashed black lines represent the 1st and 9th empirical deciles of the distribution of the cumulative returns at each point of time. The returns are spot returns in USD and the dollar index includes the following currencies: AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD, SEK. The sample is from November 2000 to March 2017.

Figure 10: Pre-announcement trading (bid-ask adjusted excess returns): strategy performance.



This figure depicts the cumulative return on a trading strategy buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts). The position in the spot rate is established 11 days in advance of each announcement day, only if the forecast interest rate change exceeds 10 basis points in absolute value. The position is then rolled over using tom/next swaps for 10 days and liquidated at the spot rate on the day preceding the announcement day. The rate change is forecast 12 days before the announcement day as the difference between the OIS-implied rate averaged over the five previous days and the corresponding underlying rate averaged over the same horizon. Panel 10a shows the return plotted against time and Panel 10b shows the return plotted event-by-event. The numbers in Panel 10a are mean return, its standard error (both in basis points) and the Sharpe ratio per one holding period. The standard error is Newey and West (1987) HAC with optimal number of lags according to Newey and West (1994). The returns are bid-ask spread-adjusted excess returns in USD on the following currencies AUD, CAD, CHF, EUR, GBP, NZD, SEK, and the dollar index. The sample is from November 2000 to March 2017.

Figure 11: Pre-announcement trading (bid-ask adjusted excess returns): robustness to the choice of holding period and threshold.



This figure plots cumulative returns on 375 trading strategies buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts) for various holding horizons and expected policy rate cutoff levels. In the case of expected policy rate hike, the strategy $S^k(h, \tau)$ buys currency k against USD (or buys the dollar index for the FOMC announcements). The position in the spot rate is established $h + 1$ days in advance of the announcement day, only if the difference between the average OIS-implied post-announcement rate over the days $h + 2, \dots, h + 6$ exceeds the average corresponding underlying rate over the same horizon by τ or more. Similarly, the currency is sold if an interest rate cut is expected and the implied rate is below the underlying rate by at least τ basis points. The position is rolled over for h days using tom/next swaps and liquidated at the spot rate on the day preceding the announcement day. The set of trading strategies (plotted in gray) is generated for $h \in [1, 15]$ and $\tau \in [1, 25] bps$, the solid black line depicts the cross-sectional mean across all trading strategies and the dashed black lines represent the 1st and 9th empirical deciles of the distribution of the cumulative returns at each point of time. The returns are bid-ask spread-adjusted excess returns in USD on the following currencies AUD, CAD, CHF, EUR, GBP, NZD, SEK, and the dollar index. The sample is from November 2000 to March 2017.

Table I: Central Banks' Policy Meetings Summary

Country	Announcements since	Target Rate	Events	Hikes	Cuts
Australia	1980s	Cash Rate	180	19	24
Canada	Nov-2000	Target for the Overnight Rate	131	18	25
Eurozone	Jan-1999	Rate on the Deposit Facility	202	11	20
New Zealand	Apr-1999	Official Cash Rate	130	23	22
Norway	Jun-1999	Sight Deposit Rate	128	21	23
Sweden	Oct-1999	Repo Rate	110	24	23
Switzerland	Jan-2000	3-month CHF LIBOR	52	9	5
United Kingdom	Jun-1998	Bank Rate	196	10	20
United States	Feb-1994	Federal Funds Rate	132	20	18
Total Events			1254	155	180

This table summarizes the policy announcements across countries. The first three columns contain countries, date of adoption of interest rate target announcements on prescheduled dates by the countries' central banks, and the corresponding interest rates respectively. The last three columns contain the total number of meetings, and the numbers of hikes and cuts for each country. The sample spans period from November 2001 when all countries adopted target rate announcements on fixed dates to March 2017, and considers shceduled announcements only. The period of Swiss franc – euro peg (September 2011 to January 2015) is omitted.

Appendix A Event study

Cutting and pivoting the sample of abnormal returns of currency i results in the following matrix:

$$R_i = \begin{bmatrix} R_{i,1,w_b} & R_{i,2,w_b} & \dots & R_{i,K,w_b} \\ R_{i,1,(w_b+1)} & R_{i,2,(w_b+1)} & \dots & R_{i,K,(w_b+1)} \\ \vdots & \vdots & \ddots & \vdots \\ R_{i,1,w_a} & R_{i,2,w_a} & \dots & R_{i,K,w_a} \end{bmatrix} \quad (\text{A.1})$$

where each row corresponds to a cross-section of returns a certain number of days after a generic event. Return $R_{i,k,s}$ is thus read as “return of currency i in period s after event k ”.

As already stated in Section II, a cumulative abnormal return (CAR) is defined as:

$$R_{i,k,s}^{ca} = \begin{cases} \sum_{t=s}^{-1} R_{i,k,t} & s < 0, \\ \sum_{t=+1}^s R_{i,k,t} & s > 0, \end{cases} \quad (\text{A.2})$$

The average-across-events CAR is defined as:

$$\bar{R}_{i,s}^{ca} = \frac{1}{K} \sum_{k=1}^K R_{i,k,s}^{ca} \quad (\text{A.3})$$

which corresponds to the average across columns of matrix (A.1). Finally, the average-across-assets CAR is the average-across-events CARs averaged across the assets:

$$\begin{aligned} \bar{R}_s^{ca} &= \frac{1}{N} \sum_{i=1}^N \bar{R}_{i,s}^{ca} \\ &= \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K R_{i,k,s}^{ca} \\ &= \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K \sum_{t=s}^{-1} R_{i,k,t} \end{aligned} \quad (\text{A.4})$$

in the pre-announcement case.

Distributional properties of \bar{R}_s^{ca} are derived from equation (A.4). For the mean:

$$\mathbb{E} [\bar{R}_s^{ca}] = \frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K \sum_{t=s}^{-1} \mathbb{E} [R_{i,k,t}] = 0 \quad (\text{A.5})$$

under the constant mean zero model for the abnormal return. For the variance, note that:

$$\text{var} \left[\sum_{m=1}^M x_m \right] = \sum_{m=1}^M \text{var} [x_m] + 2 \sum_{i \neq j} \text{cov} [x_i, x_j], \quad (\text{A.6})$$


which is simply equal to the first addend on the right-hand side if the cross-covariances are all zero. Given that policy announcements in any particular country are widely dispersed through time, the covariances stemming from the sum over k in equation (A.4) vanish. So do those stemming from the sum over i since the announcements made by the regulators of different countries are not synchronized and only rarely coincide. The inter-temporal covariances stemming from the sum over t are minuscule on the FX markets at the daily frequency, so we treat them as being zero. With that in mind:

$$\text{var} [\bar{R}_s^{ca}] = \frac{1}{(NK)^2} \sum_{i=1}^N \sum_{k=1}^K \sum_{t=-s}^{-1} \text{var} [R_{i,k,t}] \quad (\text{A.7})$$

We estimate $\text{var} [R_{i,k,t}]$ as the variance of returns of currency i in the period between two consecutive event windows, which given the zero-mean assumption discussed earlier amounts to:

$$\begin{aligned} \text{var} [R_{i,k,t}] &= \frac{1}{T} \sum_{t=s_0}^{s_1} R_{i,t}^2 \\ s_0 &= d_{i,k-1} + w_a + 1 \\ s_1 &= d_{i,k} - w_b - 1 \end{aligned} \quad (\text{A.8})$$

where T is the number of periods between the two event windows. This rather cumbersome formula in reality represents a very simple concept depicted below:




Here three events related to currency i are dated with $d_{i,1} = 04/15$, $d_{i,2} = 05/29$, and $d_{i,3} = 06/13$. The gray shaded area around each corresponds to (w_b, w_a) days around each event, and a hatched area before each event window is used for estimation of the variance of abnormal returns around that event.

Appendix B Event study, different counter currencies


We repeat the event study using GBP and JPY instead of USD as the counter currency. With a different counter currency, two things change when compared to Figures 1 and 2: first, the events associated with the new currency "disappear", and second, the USD appears as an additional currency. The results are presented in Figures B.1 and B.2 (with GBP as the counter currency) and B.3-B.4 (with JPY as the counter currency), and are supportive of our previous findings: the appreciation before interest rate hikes is not pronounced, unlike the depreciation before cuts. Over the 10-day pre-event period the currencies depreciate vs. GBP (JPY) by 60 (100) basis points, which is statistically significant at the 5% level.

Figure B.1: Exchange rates around local interest rate hikes (counter currency GBP).




This figure depicts cumulative currency returns around interest rate hikes announced by the local central banks, with the counter currency being GBP. Panel B.1a shows returns on individual currencies and Panel B.1b shows the average return over all currencies weighted in proportion to the number of hikes each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in GBP. The sample includes AUD, CAD, CHF, EUR, NOK, NZD, SEK and USD for the period from November 2000 to March 2017.

Figure B.2: Exchange rates around local interest rate cuts (counter currency GBP).




This figure depicts cumulative currency returns around interest rate cuts announced by the local central banks, with the counter currency being GBP. Panel B.2a shows returns on individual currencies and Panel B.2b shows the average return over all currencies weighted in proportion to the number of cuts each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in GBP. The sample includes AUD, CAD, CHF, EUR, NOK, NZD, SEK and USD for the period from November 2000 to March 2017.

Figure B.3: Exchange rates around local interest rate hikes (counter currency JPY).



This figure depicts cumulative currency returns around interest rate hikes announced by the local central banks, with the counter currency being JPY. Panel B.3a shows returns on individual currencies and Panel B.3b shows the average return over all currencies weighted in proportion to the number of hikes each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in JPY. The sample includes AUD, CAD, CHF, EUR, jpy, NOK, NZD, SEK and USD for the period from November 2000 to March 2017.

Figure B.4: Exchange rates around local interest rate cuts (counter currency JPY).



This figure depicts cumulative currency returns around interest rate cuts announced by the local central banks, with the counter currency being JPY. Panel B.4a shows returns on individual currencies and Panel B.4b shows the average return over all currencies weighted in proportion to the number of cuts each currency experienced. The announcement day is marked by zero. A pre-announcement spot return is realized by opening a long position in the currency x days and reversing it one day before the announcement; the post-announcement returns are realized by opening a long position in the currency on the first day following the announcement and holding it for x days, whereby x is read off the abscissa. The shaded area in the bottom panel represents the 95% confidence interval for the average value around zero. All returns are spot returns in JPY. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD, SEK and USD for the period from November 2000 to March 2017.

Appendix C Swap points vs. forward discounts

Recall that the monthly excess log-return on a currency pair is:

$$\begin{aligned} rx_{t+1} &= s_{t+1} - f_t \\ &= s_{t+1} - s_t + s_t - f_t \\ &= \Delta s_{t+1} + d_t, \end{aligned} \tag{C.1}$$

where Δs_{t+1} is the log spot return, d_t is the forward discount, approximately equal to the interest rate differential, and t is understood to index months. On the other hand, the same return is expected to be achieved by rolling over a spot position from t to $t + 1$, assuming a total of h days in the month:


$$\hat{rx}_{t+1} = \log S_{t+1} - \log(S_t + \sum_{\tau=1}^h w_\tau), \tag{C.2}$$

We could use the Taylor expansion of $\log(S_{t+1} + \sum_{\tau=1}^h w_\tau)$ around S_t (since the second addend is usually very small on the frequencies higher than the monthly) to rewrite equation (C.2) as follows:

$$\begin{aligned} \hat{rx}_{t+1} &= \log S_{t+1} + \frac{1}{S_t} \sum_{\tau=1}^h w_\tau - \log S_t \\ &= \Delta s_{t+1} + \hat{d}_t, \\ \hat{d}_t &= \frac{1}{S_t} \sum_{\tau=1}^h w_\tau \end{aligned} \tag{C.3}$$

Obviously, asking how close \hat{rx}_{t+1} is to rx_{t+1} is tantamount to asking if the previous month's forward discounts are accurate predictors of their next month's cumulative daily counterparts. The wedge – if any – should be driven by both the failure of the expectation hypothesis and omnipresent market frictions. Without claim at a rigorous study of this wedge, which would be beyond the scope of our work, and rather as a quick check that it is small, in Figure C.1 we compare d_t from equation (C.1) \hat{d}_t from equation (C.3). Though the cumulative daily rollovers are more volatile and would thus introduce additional noise to the excess return series, they closely follow the monthly forward discounts. The maximum mean absolute difference between the two series occurs for AUD and reaches 0.14% p.a., which is negligible compared to the magnitudes of returns of the strategies that we construct.

Figure C.1: Daily tom/next swap points vs. 1-month forward discounts.




This figure shows the part of the FX monthly excess return, in percent p.a., attributed to the interest rate differential between the respective currency and the US Dollar. The colder-colored line depicts the case of opening a spot position at the beginning of each month and rolling it over daily until the end of the month. Because of missing data, the average over the non-missing observations within each month is taken and multiplied by 30 to arrive at the monthly figure. The warmer-colored line depicts the case of an investor entering a short forward contract at the end of the previous month and closing it at month's end. The number in the lower left corner stands for the mean absolute difference between the two series, in percent p.a. All quotes are mid quotes from November 2000 to March 2017.

Appendix D Restricting leverage

Although funding constraints for foreign exchange speculation are rather lax, we demonstrate that our results are not driven by higher amount of leverage before profitable trades and lower leverage before the unprofitable ones. We repeat the trading strategy exercise for a leverage-constrained investor.

Figure D.1 plots the cumulative unrealized profit and loss (i.e. the liquidation value of all open positions in excess of initial equity) of the baseline strategy. At the beginning of the sample we endow the investor with one US dollar of equity, and require her to fully collateralize all open positions on the *net* basis, that is for a dollar of equity she is allowed to hold one dollar long and one dollar short in different currencies. At each portfolio rebalancing, first, we calculate the margin closeout value which is the sum of the portfolio balance and the unrealized profits and losses; then we allocate an equal (in absolute terms) share of the closeout value to all currencies that are required to be held according to signals. A leveraged portfolio would have opened positions sum up to a multiple of the margin closeout value; we restrict leverage to 1 to make the strategy comparable with the popular long-short strategies on the FX market. We conduct every transaction at the London fixing time to avoid any overlaps in the positions. Furthermore, when local predicted signals conflict with the FOMC signals, the former are given priority over the latter. For example, if an interest rate hike is predicted both in the US and Australia, the portfolio is long AUD and short every other currency. With leverage excluded, the cumulative performance exceeds 100 percent over the whole sample which corresponds to about 6.5 percent per year. The average ten-day return of 31.44 basis points is statistically significant at the 1% level, having the *t*-statistic of 2.47.

Figure D.1: Pre-announcement trading with restricted leverage (bid-ask adjusted excess returns): policy rate expectations.



This figure depicts the cumulative unrealized profit and loss of a trading strategy buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts). The position in the spot rate is established 11 days in advance of each announcement day, only if the forecast interest rate change exceeds 10 basis points in absolute value. The position is then rolled over using tom/next swaps for 10 days and liquidated at the spot rate on the day preceding the announcement day. The rate change is predicted 12 days before the announcement day as the difference between the OIS-implied rate averaged over the five previous days and the corresponding underlying rate averaged over the same horizon. The numbers refer to the mean, standard error of the mean (both in basis points) and the Sharpe ratio of daily log changes in market value of the portfolio scaled to 10 days to represent the average holding period. The standard error is Newey and West (1987) HAC with optimal number of lags according to Newey and West (1994). The unrealized profit and loss is in USD and accounts for bid-ask spread. The sample includes AUD, CAD, CHF, EUR, GBP, NZD, SEK, JPY, NOK, with the last two being traded around FOMC announcements only. The sample is from November 2000 to March 2017.