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

Monetary Policy and Currency Returns: the Foresight Saga*

Dmitry Borisenko†, Igor Pozdeev‡

This draft: May 2017

Abstract

We document a drift in exchange rates before monetary policy changes across major economies. Currencies tend to depreciate by 0.7 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. A simple trading strategy buying currencies against USD ten days ahead of predicted local interest rate hikes and selling currencies before predicted cuts earns on average a statistically significant return of 42 basis points per ten-day period. 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.

Keywords: Monetary Policy, Policy Expectations, Predictability, Overnight Index Swap, Foreign Exchange.
JEL Classification: E43, E52, E58, F31, G12

*We would like to thank Paul Söderlind, Angelo Ranaldo and Annette Vissing-Jørgensen for their valuable comments and insightful feedback.
†*E-mail*: dmitry.borisenko@student.unisg.ch
‡*E-mail*: igor.pozdeev@unisg.ch

1 Introduction

By 2001 the regulators of most developed countries have adopted the practice of a fixed number of meetings per year each culminating in a decision about the policy rate – most often a short-term rate. The shock would then by virtue of the expectation hypothesis propagate through the yield curve and eventually find its way into the exchange rate of the domestic currency, as documented in the vast literature on the relation between a currency return and the level of the risk-free interest that the currency can buy. The notorious uncovered interest puzzle and the carry trade strategies are all part of this literature. Thus, the indirect effect of monetary policy on exchange rates has been studied well; the direct 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 abnormal positive high-frequency returns of a simple trading strategy that goes long foreign (from the perspective of an American investor) currencies on the days of 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. Karnaukh (2016) takes their research to 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 reports a puzzling high profitability of such trading strategy between 1994 an 2015.

Our major contribution to the research on the monetary policy – exchange rates relation is to (a) cast a critical view on some previously established results and (b) bring the currencies and the policy announcements of other developed countries into the picture. We find that the multiperiod appreciation before rate raises and depreciation before rate cuts is a phenomenon common for most currencies – not only for the US dollar, as shown in Karnaukh (2016) – and that it totally dwarfs their post-announcement exchange rate dynamics. Similar to the latter work, we construct a trading strategy by forecasting upcoming rate changes and dynamically rebalancing the portfolio of currencies. We stress however, that this forecasting exercise is a classification problem and thus subject to the discretionary choice of the classification rule. When controlled for the backward-looking bias, as discussed i.a. in Bailey and Lopez de Prado (2014), the dollar-FOMC pre-announcement drift of Karnaukh (2016) is found to be less significant and rather an exception from the highly significant drifts of other currencies.

In contrast, we show that the return of the trading strategy based on predicted policy rate changes in a cross-section of countries is large and robust to the choice of the classification rule. It exhibits a total return of 80% when averaged across all rules, with several specifications resulting in as much as a 150% return and neither one leading to a money loss.

Still, forecasts of policy rate changes based on the market-implied information and a classification rule are likely to be too conservative due to various premia – e.g. liquidity and term – and the microstructure noise imminently present in the market prices. One would most probably expect the information set of market participants in the real world to be larger than that of the market participants in our paper. We thus complement our findings with studying the perfect foresight strategy, which is based on a perfect knowledge of future policy rate changes and constitutes the ultimate unattainable benchmark for currency traders. First, we show that a randomly selected currency is expected to depreciate against the USD by 70 bps over the 10 days before rate cuts, and appreciate by 50 bps in the opposite case, which is statistically significant at the 95% level. No pronounced pattern is observable on the days following the announcements. Second, we construct a strategy by going long currencies about to witness a rate hike and selling short those with an approaching rate cut (over the same time window and cross-section of currencies as for the forecast-based strategy), and document a profit of about 140% since late 2000, tantamount to 8.5% p.a.

What could explain these 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, with investors being compensated for uncertainty about 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, however the distribution of risk among the investor types does, 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 maintain 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, and 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 infrequent rebalancings 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.

Our minor contribution is to the literature on measuring expectations about future policy rates. While evidence on the forecasting 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 to gauge policy expectations. We find that policy rate forecasts extracted from OIS rates have been most accurate since mid-2000. 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.

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 in the same manner to the upcoming rate hikes and cuts. 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 dynamics of the corporate bond returns after monetary policy shocks. For the FX market, the above mentioned papers by Mueller et al. (2017) and Karnaukh (2016) are the major references.

The rest of the paper is organized as follows. Section 2 outlines the way we perform event studies, recover policy expectations and construct trading strategies; Section 3 summarizes the data we use; Section 4 presents our findings; Sections 5 concludes.

2 Methodology

This section describes empirical design of our study. First we outline the methodology of event study in a multicurrency framework. Then we discuss payoff structure of overnight index swaps and federal fund futures and how one can extract expected interest rates implied by these contracts. The section concludes with description of a trading strategy aiming to buy currencies before expected hikes and sell currencies before expected cuts.

2.1 Event Study

To detect the pre-announcemnt drift on 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, 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 $k \in \{1, \ldots, K\}$ relating to currency $n \in \{1, \ldots, N\}$. 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

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

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:

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

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 eq. (1):

$$\overline{R^{ca}_s} = \frac{1}{NK} \sum_{i=1}^{N} \sum_{k=1}^{K} R^{ca}_{i,k,s} \tag{2}$$

Appendix A shows that $\overline{R^{ca}_s}$ 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 the Appendix.

2.2 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[1]. 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:

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

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 swap[2] today (day 0) with start date tomorrow (day 1) is equal to the risk-neutral expectation of (3):

$$w_0 = \mathrm{E}_0\left[\pi^T\right] = \mathrm{E}_0\left[\prod_{t=1}^{T}(1 + r_t) - 1\right], \tag{4}$$

where the expectation is taken under the risk-neutral measure. Let us assume a pol-

[1]For 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.

[2]The 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.

icy 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:

$$\begin{aligned} w_0 &= \mathrm{E}_0 \left[\prod_{s=1}^{t^*} (1 + r_0) \prod_{t=t^*+1}^{T} (1 + r^*) - 1 \right] \\ &= (1 + r_0)^{t^*} \mathrm{E}_0 \left[(1 + r^*)^{T - t^*} \right] - 1, \end{aligned} \tag{5}$$

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

$$\mathrm{E}_0 \left[r^* \right] = \left((w_0 + 1)(1 + r_0)^{-t^*} \right)^{\frac{1}{T - t^*}} - 1 \tag{6}$$

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:

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

where T_m is the number of calendar days in month m. Two major advantages of these contracts is that they enjoy high liquidity and are marked to market on the daily basis which mitigates counterparty risk. Krueger and Kuttner (1996) and Gürkaynak et al. (2007) find the futures-implied rate to be a sufficiently 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:

$$f_t^m = \frac{1}{T_m} \mathrm{E}_t \left[\sum_{s \in m} r_s \right], \tag{8}$$

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:

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

2.3 Trading Strategy

We construct a simple trading strategy based on expected shifts in policy rates. Assuming a US investor 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 buy the foreign currency, should a rate cut be expected, we sell the foreign currency and buy USD, otherwise no position is established. The log (excess) return over h periods realized at time $T-1$ is therefore:

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

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 as given by (6) and (9). With an exception of the US, the underlying rates for OIS differ from the policy rates set by central banks and thus can be noisy, furthermore derivatives-implied measures of the future path of monetary policy can be noisy as well containing for instance a time-varying risk premium[3]. To address this issue, on the day of the policy rate forecast we define expected change in the target rate $\mathrm{E}_{T-h-2}\left[\Delta i_T\right]$ 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[4]. We further employ a simple rule to evaluate

[3]Although 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.

[4]The choice of the smoothing window is inconsequential for our results.

the expected shift in the target rate by defining the categorical variable d_{T-h-2} as:

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

where τ is a threshold level. Denote $a = 1, ..., A$ 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:

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

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.

3 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.

3.1 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 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[5]. The targets and announcement schedules, how-

[5]On March 19th 2001 the Bank of Japan abandoned targeting of the uncollaterallized 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

ever, 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 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

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.

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 in this study, 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 peg 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, has been announcing 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 for example Lucca and Moench (2015) and references therein.

Table 1 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 peg (from September 2011 to January 2015) is excluded for Switzerland. The total numbers of hikes and cuts are 155 and 180 respectively, giving 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[6]. 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.]

3.2 Exchange Rates and Currency Returns

We use 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 from Bloomberg for different fixing times to ensure that the annoncement day is not overlapped for any of the currencies. Thus we use 17:00 London fixing time for the Eurozone, Norway, Sweeden, Switzerland, and the United Kingdom; 17:00 New York time for Canada and the US; 20:00 Tokyo time for Australia and New Zealand.

We take the perspective of a US investor such that foreign currencies can be thought of as assets with dollar prices S_t^i, i denoting a currency. The one-period spot log-return equals:

$$\Delta s_t^i = s_t^i - s_{t-1}^i = \log S_t^i - \log S_{t-1}^i \tag{12}$$

where S_t^i is the exchange rate of the foreign currency. For the FOMC announcements we construct the dollar index – an equally weighted portfolio of currency returns against USD, whith each currency, including JPY, fixed at 17:00 New York time. We perform all the analysis on spot returns. Whereas returns to a trading strategy should in fact be excess returns, we do not find significant differences when redoing the calculations on excess returns. These results are available on request.

3.3 Overnight Index Swaps and Federal Funds Futures

We collect 1-month swap rates from Datastream (provided by ICAP and by Thomson Reuters where ICAP data are unavailable). The availability of the OIS data is as fol-

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

lows: Australia, Canada, Switzerland and the US since late 2001; the Eurozone since January 1999; Sweden since September 2002; United Kingdom since August 2007, and New Zealand since March 2009.

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.[7] The data on these rates are 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.

4 Results

In this section we present the empirical results of the paper. First, in an event study, we document a pre-announcement drift in currency returns preceding shifts in monetary policy around the world. Then we analyze these findings in the context of a trading strategy aiming to forecast future monetary policy action and then buy (sell) currencies whose monetary authorities are about to raise (cut) their policy rates. Finally we take a closer look at the predictive power of the monetary policy action forecasts.

4.1 Event study

Figures 1-2 show the results of an event study, with events being announcements of the decisions of the local central banks to cut and raise the target interest rate respectively, and the test assets being spot returns of the currencies of the corresponding countries. In each figure the top panel plots cumulative returns of individual currencies while the bottom panel plots the average return over all currencies weighted proportionally to the number of hikes and cuts each currency experienced over the period from November 2000 to March 2017. The currencies are AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK, and the sample includes a total of 135 hikes and 162 cuts. 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 pattern is reversed before the rate cuts: a currency is expected to depreciate by about 70 bps over ten days, half that over five days, and 10 bps on the

[7] There are no overnight interest rates data available for Norway.

pre-anouncement day. The shaded area in the bottom panel of each figure shows the 95% confidence interval assuming zero currency return outside the event window. To construct it, the set of days complementary to the set of event windows is used: in Figure 1 we take the days without announcements as well as the days of announced cuts and no-change decisions to construct the confidence interval; in Figure 2 we take the days without announcements as well as the days of announced raises and no-change decisions;

[Figure 1 about here.]

[Figure 2 about here.]

As seen in the figures, the currencies experience a statistically significant and economically large drift in the direction of the policy rate changes. The exchange rates begin to move at least ten days in advance of the central banks' announcements. Furthermore, after the announcment day the drift completely dissipates and the abnormal returns evaporate in the post-announcement period.

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 predict shifts in monetary policy around the world, and whether the pre-announcement drift in exchange rates can be exploited as a trading strategy. We address this issue in the rest of the section.

4.2 Trading Strategy

Figure 3 plots the cumulative performance of a strategy where an investor is going long currencies whose monetary authorities are expected to raise the policy rates and shorting currencies with expected interest rate cuts. The investor makes a decision whether to trade a currency 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 upper panel charts the cumulative return against time and the lower panel plots the performance event-by-event, to stress that the results are not dominated by

a handful of announcements. 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 3 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 total return 120% with average per-event return of 42.52 basis points (with t-statistic of over 3) and a ten-day Sharpe ratio of 0.2, underperforming its perfect foresight counterpart by 20% over the course of the sample. Furthermore the bottom panel of Figure 3 demonstrates that the performance of the strategy is not driven by a few rare events that delivered high returns over the course of the sample.Given the economic and statistical significance of this result we further investigate whether its robustness to the choice of the holding period and threshold.

[Figure 3 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 4 plots the results of this exercise. The gray lines chart returns on individual strategies, the solid black line shows the average cumulative return over all strategies at each point of time, and the dashed black lines depict the first and ninth deciles of the cumulative returns' distribution at each point of time. Strikingly, the entire support of the end-date empirical distribution of the cumulative returns is positive, indicating that our result is robust to the choice of the trading strategy parameters.

[Figure 4 about here.]

Finally, we check whether these results can be attributed to the FOMC pre-announcement drift of the dollar factor documented by Karnaukh (2016). Figure 5 repeats the analysis of Figure 4 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, thus making our evidence qualitatively different from that in previous studies[8].

[Figure 5 about here.]

Overall, the trading strategy exercise provides evidence of the short-horizon predictability of shifts in key policy rates around the world and hence of the corresponding pre-announcement drift in exchange rates we documented earlier in this section. Furthermore this drift can not be attributed to behavior of the dollar index before the target rate announcements documented in the previous literature.

4.3 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 be set at each announcement. Figure 6 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 implied Fed funds rates extracted from the Fed funds futures: these exhibit a slightly higher mean absolute error, but at the first glance are as fine a predictor of the policy shifts.

[Figure 6 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 7 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. Entry (x, y) (x denotes rows) in any such matrix contains the number of cases when direction x was predicted, and direction y announced. 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

[8]Similar to Karnaukh (2016) we observe economically and statistically significant pre-announcement 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.

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 6, 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 7 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).

5 Conclusion

We describe a strong 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 raises, and significantly depreciate before rate decreases. 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 is sensitive to the choice of the rule: for example, when averaged over many valid specifications, the profitability of the pre-FOMC forecast-based trading reported in the literature disappears. In contrast to that, the trading strategy based on exploiting the OIS-implied information over a cross-section of currencies remains positive and large whatever the specification.

The non-negligible gap between the returns of the forecast-based and perfect foresight strategies supports the observation that exchange rates incorporate information about

future monetary policy decisions over and above that dissolved in the fixed income market. Further research is needed to understand this gap. Its roots could be found in a risk premium (or several premia) "contaminating" the OIS prices that we use. However, we show that at least for the FOMC announcements the OIS are as strong a predictor as the Federal funds futures, which were found virtually free of risk premia by Piazzesi and Swanson (2008). This leaves hope that the same holds for the other OIS rates in our sample. In any case, the gap is a non-negligible unrealized return for a currency trader, a (probably surprising) consequence of the transparent monetary policy for a central bank, and a foresight saga for both.

Our findings are difficult to reconcile with the existing theories for 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 Gavazzoni, 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.

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.

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, L'uboš, 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.

Appendix A Event study

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

$$R_i = \begin{bmatrix} R_{i,1,w_b} & R_{i,2,w_b} & \dots & R_{i,K,w_b} \\ R_{i,1,(w_b+1)} & R_{i,2,(w_b+1)} & \dots & R_{i,K,(w_b+1)} \\ \vdots & \vdots & \ddots & \vdots \\ R_{i,1,w_a} & R_{i,2,w_a} & \dots & R_{i,K,w_a} \end{bmatrix} \tag{13}$$

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 2, a cumulative abnormal return (CAR) is defined as:

$$R^{ca}_{i,k,s} = \begin{cases} \sum_{t=s}^{-1} R_{i,k,t} & s < 0, \\ \sum_{t=+1}^{s} R_{i,k,t} & s > 0, \end{cases} \tag{14}$$

The average-across-events CAR is defined as:

$$\overline{R^{ca}_{i,s}} = \frac{1}{K} \sum_{k=1}^{K} R^{ca}_{i,k,s} \tag{15}$$

which corresponds to the average across columns of matrix (13). Finally, the average-across-assets CAR is the average-across-events CARs averaged across the assets:

$$\begin{aligned} \overline{R^{ca}_{s}} &= \frac{1}{N} \sum_{i=1}^{N} \overline{R^{ca}_{i,s}} \\ &= \frac{1}{NK} \sum_{i=1}^{N} \sum_{k=1}^{K} R^{ca}_{i,k,s} \\ &= \frac{1}{NK} \sum_{i=1}^{N} \sum_{k=1}^{K} \sum_{t=s}^{-1} R_{i,k,t} \end{aligned} \tag{16}$$

in the pre-announcement case.

Distributional properties of $\overline{R^{ca}_{s}}$ are derived from eq. (16). For the mean:

$$\mathrm{E}\left[\overline{R^{ca}_{s}}\right] = \frac{1}{NK} \sum_{i=1}^{N} \sum_{k=1}^{K} \sum_{t=s}^{-1} \mathrm{E}\left[R_{i,k,t}\right] = 0 \tag{17}$$

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

$$\text{var}\left[\sum_{m=1}^{M} x_m\right] = \sum_{m=1}^{M} \text{var}\left[x_m\right] + 2\sum_{i\neq j} \text{cov}\left[x_i, x_j\right], \tag{18}$$

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 eq. (16) 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:

$$\text{var}\left[\overline{R_s^{ca}}\right] = \frac{1}{(NK)^2}\sum_{i=1}^{N}\sum_{k=1}^{K}\sum_{t=-s}^{-1} \text{var}\left[R_{i,k,t}\right] \tag{19}$$

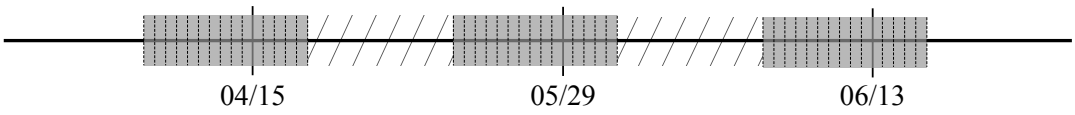
We estimate var $[R_{i,k,t}]$ as the variance of returns of currency i in the period between two consecutive event windows, which given the zero-mean assumption discussed earlier amounts to:

$$\text{var}\left[R_{i,k,t}\right] = \frac{1}{T}\sum_{t=s_0}^{s_1} R_{i,t}^2 \tag{20}$$

$$s_0 = d_{i,k-1} + w_a + 1$$

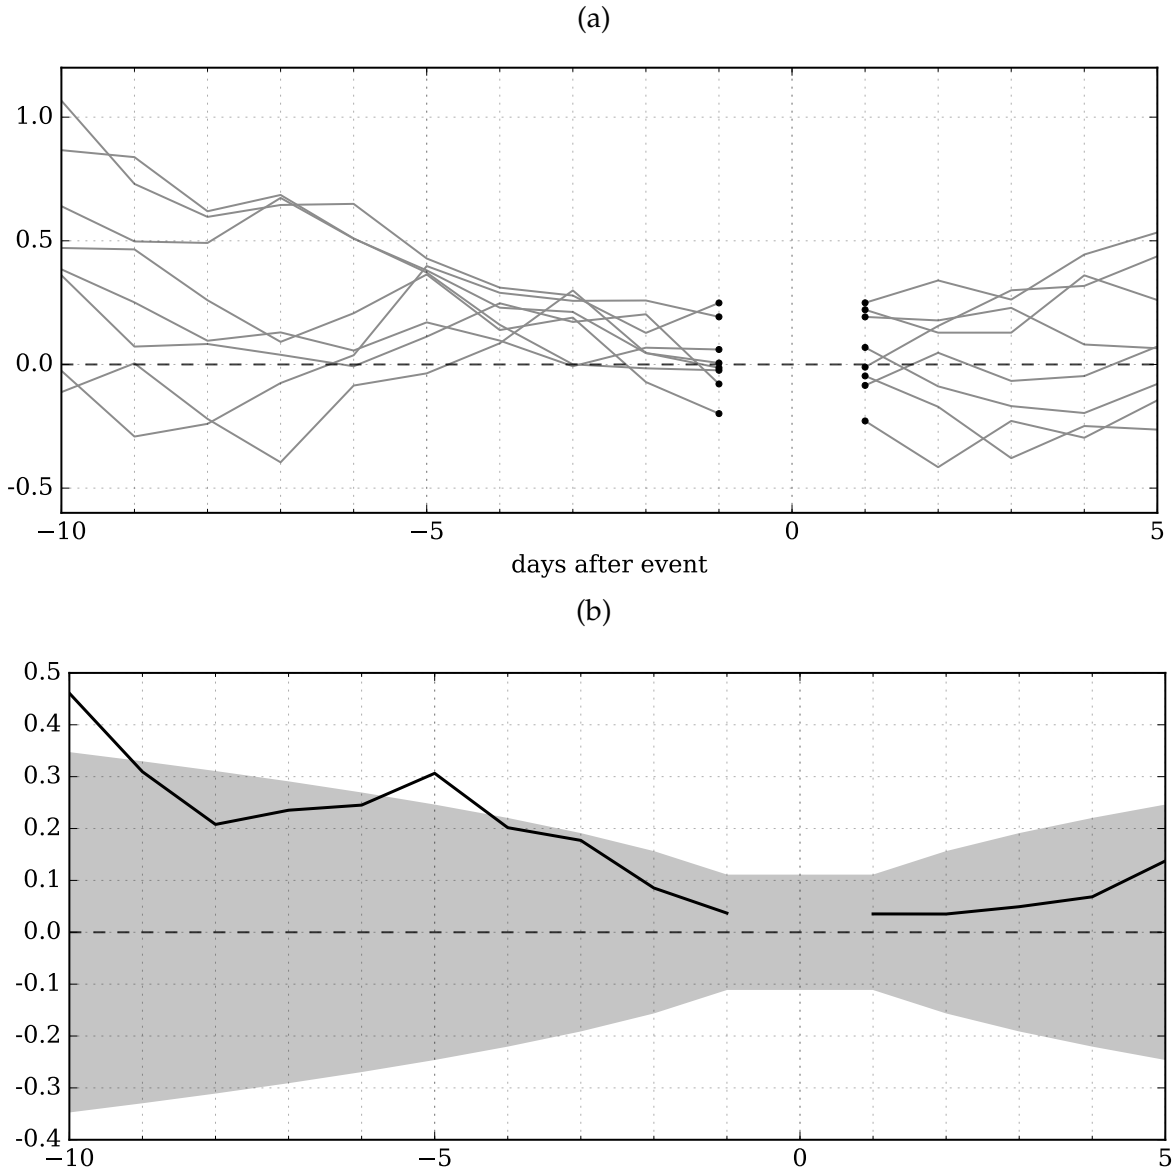
$$s_1 = d_{i,k} - w_b - 1$$

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:



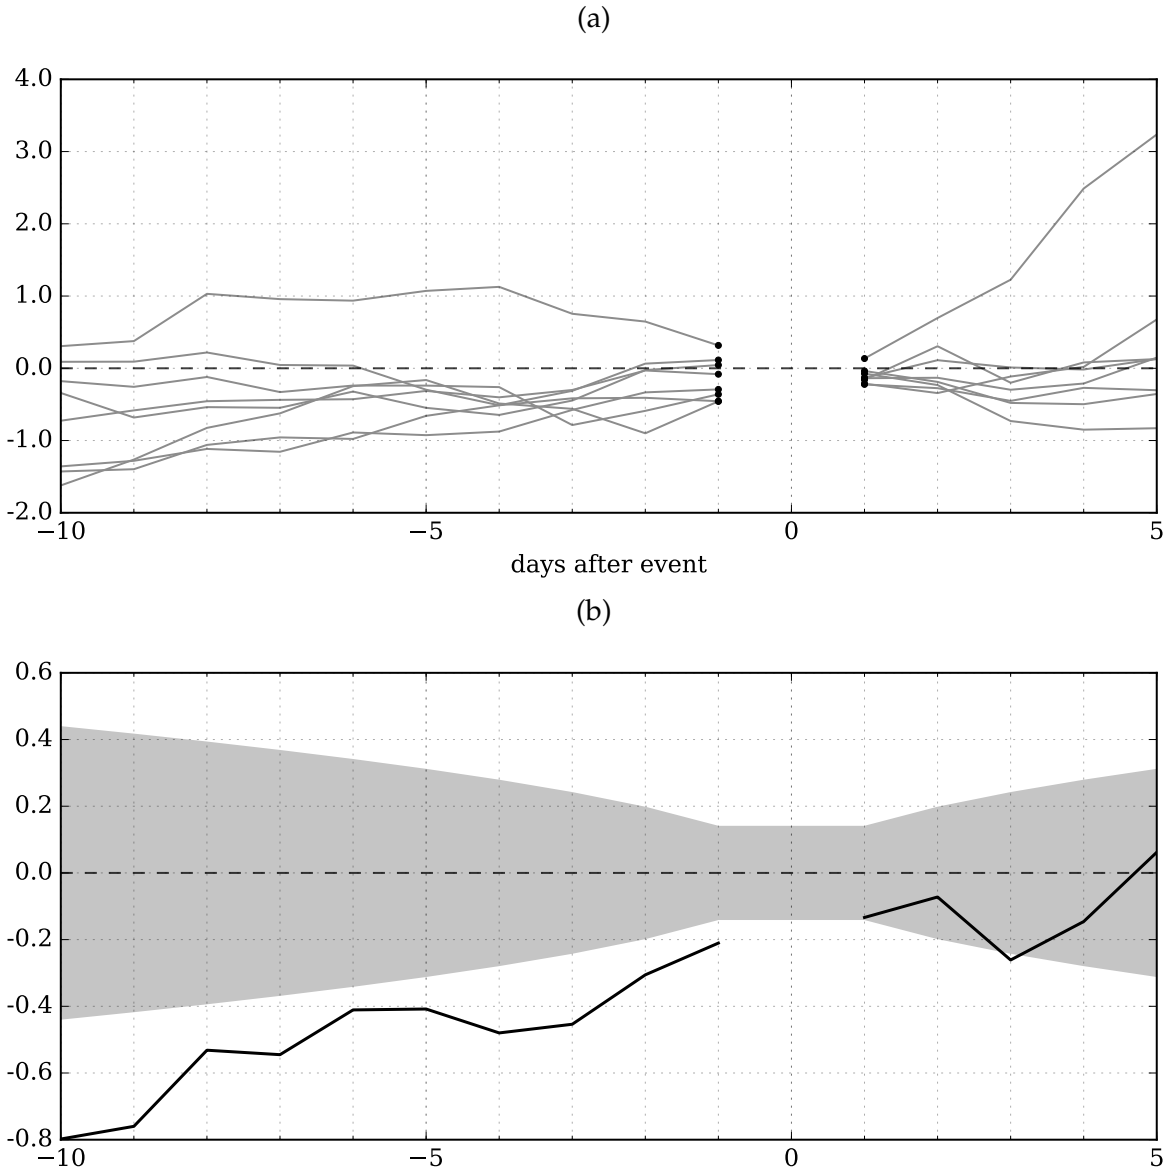
Here three events related to currency i are dated with $d_{i,1} = 04/15$, $d_{i,2} = 05/29$, and $d_{i,3} = 06/13$. The gray shaded area around each corresponds to (w_b, w_a) days around each event, and a hatched area before each event window is used for estimation of the variance of abnormal returns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

Figure 1: **Exchange rates around 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. The pre-announcement returns are realized by buying the currency x days and reversing the position one day before the announcement; the post-announcement returns are realized by buying 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.

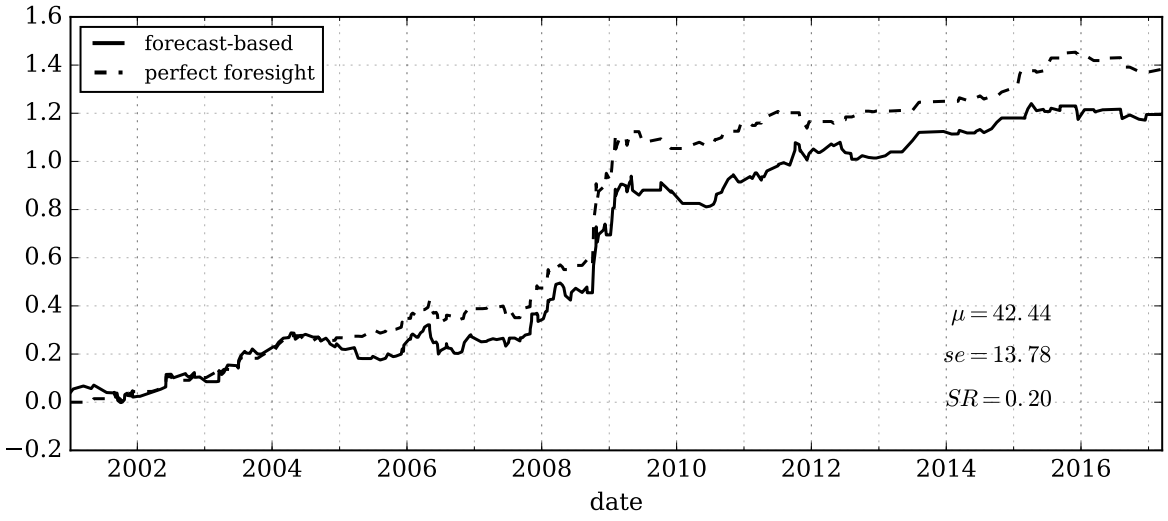
Figure 2: **Exchange rates around interest rate cuts.**



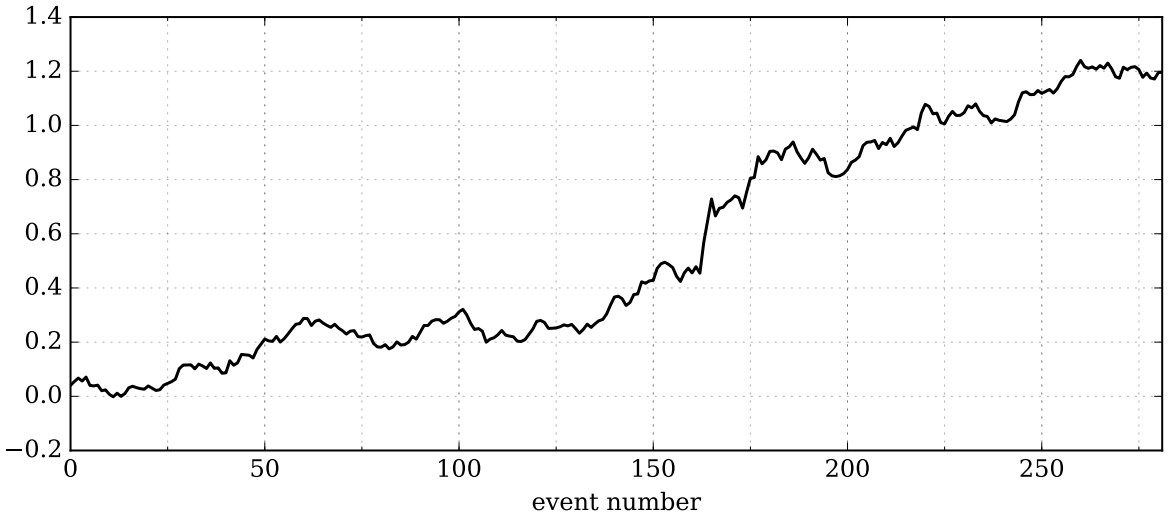
This figure depicts cumulative currency returns around interest rate cuts implemented by the local central banks. Panel 2a shows returns on individual currencies and Panel 2b shows the average return over all currencies weighted in proportion to the number of cuts each currency experienced. The announcement day is marked by zero. The pre-announcement returns are realized by buying the currency x days and reversing the position one day before the announcement; the post-announcement returns are realized by buying 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.

Figure 3: **Pre-announcement trading: policy rate expectations.**

(a)

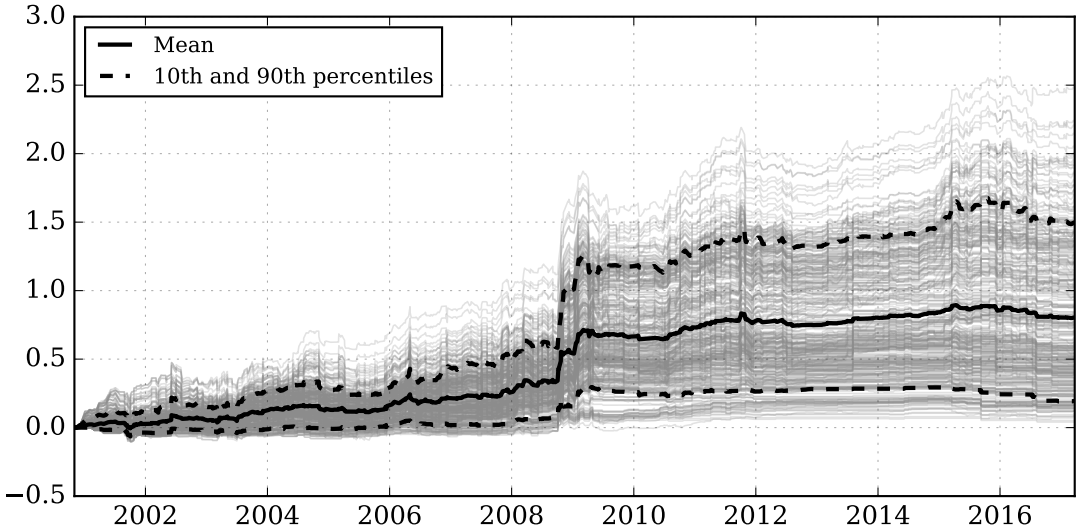


(b)



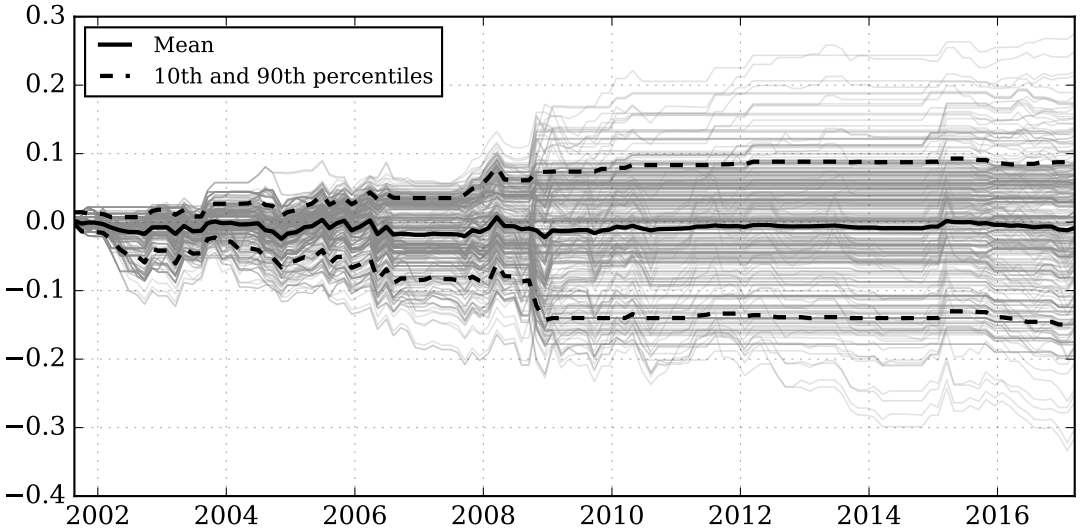
This figure depicts cumulative return on a trading strategy buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts). The position is established 11 days in advance of each announcement day, only if the forecasted interest rate change exceeds 10 basis points in absolute value. The position is then held for 10 days and liquidated on the day preceding the announcement day. The rate change is forecasted 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 3a shows the return plotted against time and Panel 3b shows the return plotted event-by-event. The numbers in Panel 3a 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 4: **Pre-announcement trading: robustness to the choice of holding period and threshold.**



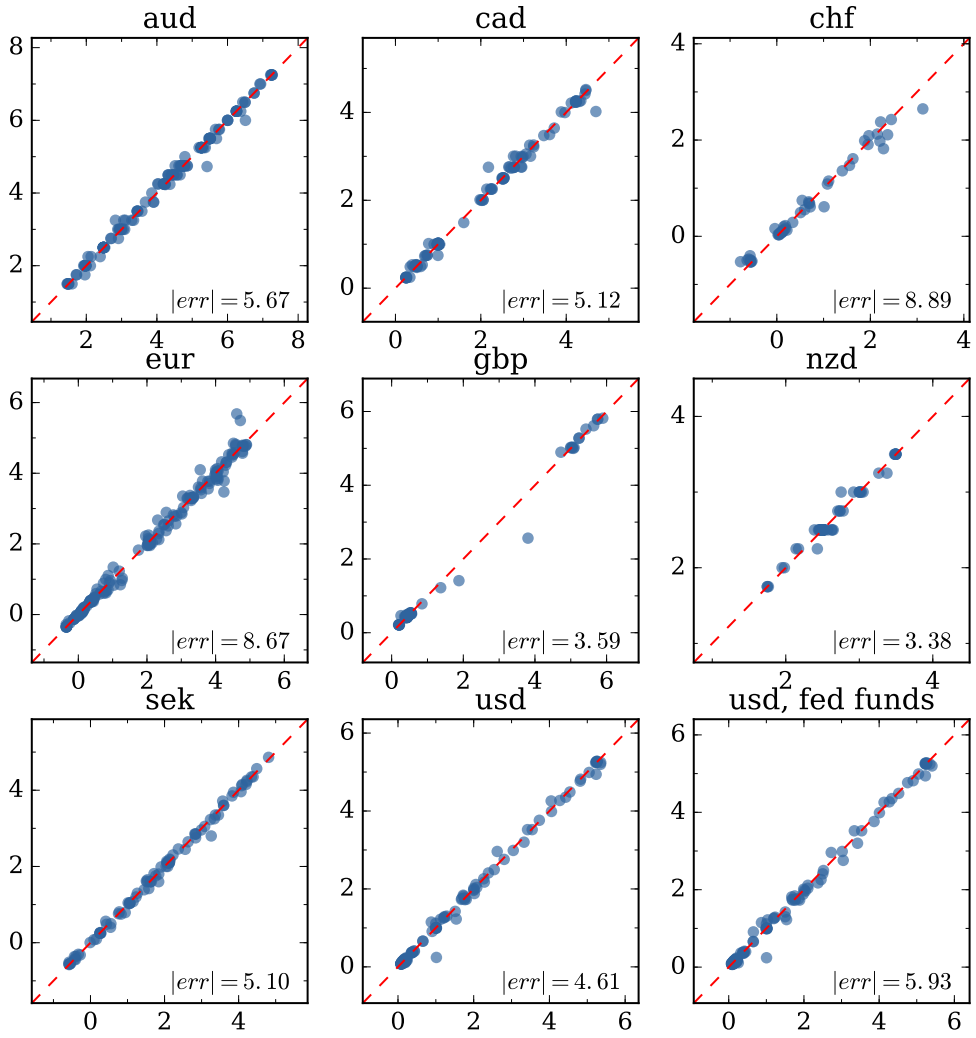
This figure plots cumulative returns on 375 trading strategies buying (selling) currencies against USD in anticipation of local interest rate hikes (cuts) for various holding horizons and expected policy rate cutoff levels. In the case of expected policy rate hike, the strategy $S^k(h,\tau)$ buys currency 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 the days $h+2,...,h+6$ 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 5: **Pre-announcement trading FOMC and the dollar index: robustness to the choice of holding period and threshold.**



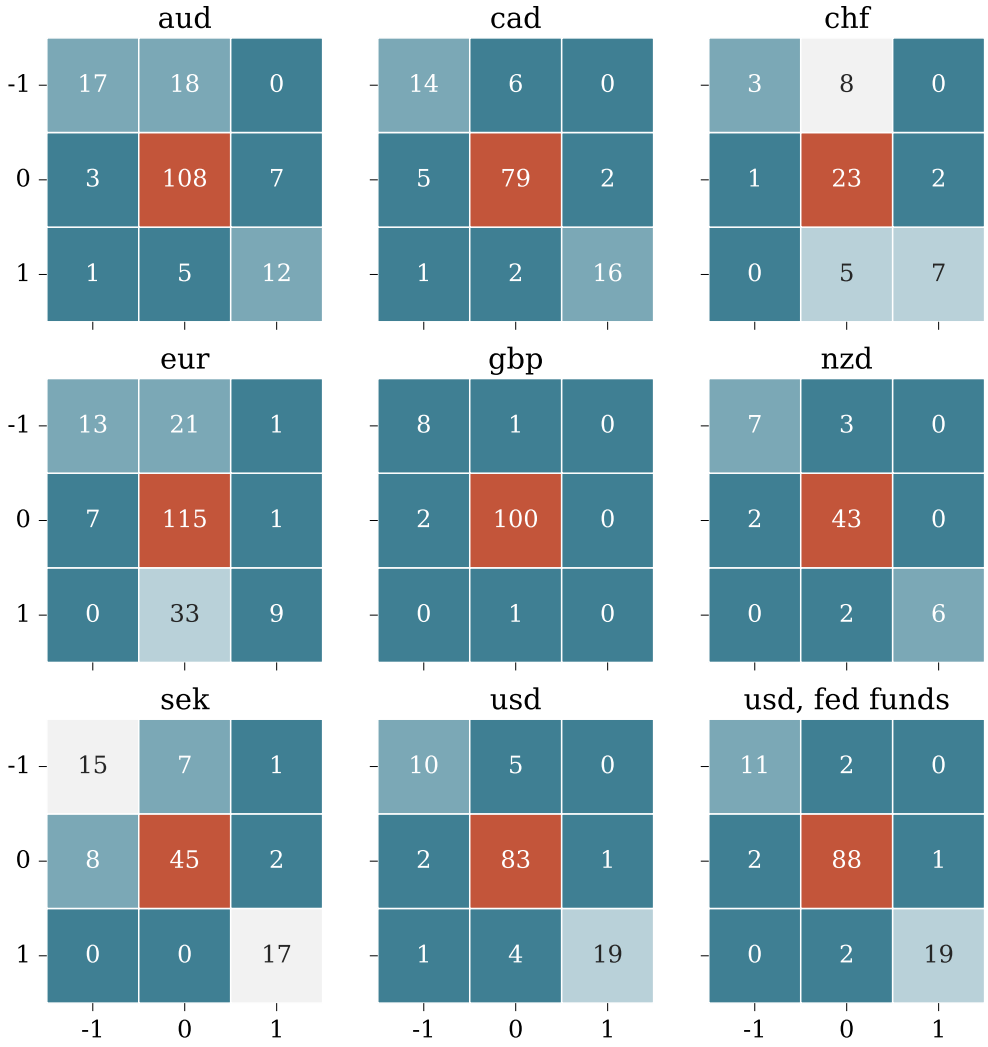
This figure plots cumulative returns on 375 trading strategies buying (selling) the dollar index in anticipation of interest rate hikes (cuts) in the US for various holding horizons and expected policy rate cutoff levels. In the case of expected policy rate hike, the strategy $S(h, \tau)$ buys the dollar 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 $h + 2, ..., h + 6$ 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 6: **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 7: **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. The x-axis keeps the actually announced direction of the rate change, the y-axis – the predicted direction. 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.

Table 1: 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.