Essays on FX Variance Risk Premium, Monetary Policy and Currency Returns

DISSERTATION
of the University of St. Gallen,
School of Management,
Economics, Law, Social Sciences
and International Affairs
to obtain the title of
Doctor of Philosophy in Economics and Finance

submitted by
Igor Pozdeev
from
Russia

Approved on the Application of

Prof. Paul Söderlind, PhD

and

Prof. Semyon Malamud, PhD

Dissertation no. 4978

D-Druck Spescha, St. Gallen 2020

Essays on FX Variance Risk Premium, Monetary Policy and Currency Returns

DISSERTATION
of the University of St. Gallen,
School of Management,
Economics, Law, Social Sciences
and International Affairs
to obtain the title of
Doctor of Philosophy in Economics and Finance

submitted by
Igor Pozdeev
from
Russia

Approved on the Application of

Prof. Paul Söderlind, PhD

and

Prof. Semyon Malamud, PhD
Prof. Francesco Audrino, PhD

Dissertation no. 4978

D-Druck Spescha, St. Gallen 2020

The University of St. Gallen, School of Management, Economics, Law, Social Sciences and International Affairs hereby consents to the printing of the present dissertation, without hereby expressing any opinion on the views herein expressed.

St. Gallen, November 15, 2019

The President:

Prof. Dr. Thomas Bieger

to my parents

Contents

Summary

The first chapter of this thesis deals with variance risk on the FX market. Therein, I recover risk-neutralized covariance matrices of currency returns and combine them with ex post realized covariance matrices to document an overall negative FX variance risk premium, but also possibility of strategies with a significantly positive one. Among portfolios with the most negative premium estimates, the US dollar index and Carry trade familiarly emerge. I report that hedging shocks to news about future FX variance is costly with a horizon-decaying pattern, suggesting a strong preference of FX market investors to hedge such shocks.

The second chapter studies the dynamics of currency spot and excess returns before policy rate announcements of central banks in developed economies. Dmitry Borisenko and I show that currencies depreciate by 8 basis points per day over ten days before target rate cuts and appreciate by 5 basis points per day before rate hikes. We use fixed income derivatives to forecast monetary policy decisions to make the above drift exploitable: our baseline specification of the trading strategy using these forecasts earns on average significant 4.5% per year, and would earn only a percentage point more if the forecast quality were perfect. This return is robust to the choice of the trading strategy specification and cannot be attributed to exposure to established FX, equity and alternative risk premia.

In the third chapter, Nikola Mirkov, Paul Söderlind and I explore whether verbal interventions by the Swiss National Bank (SNB) affected market beliefs in the desired direction during the period from 2011 to 2015, when the SNB imposed a cap on the Swiss franc against the euro. We find that the speeches including the passionate "unlimited quantities" were followed by decreased forward-looking measures of uncertainty regarding the future value of the franc and steered market beliefs toward franc depreciation, therefore reinforcing the credibility of the Swiss franc cap.

In the fourth and final chapter, Dmitry Borisenko and I document no significant risk premium embedded in overnight index swap (OIS) rates for maturities up to one year, which allows us to extract the implied future policy target rates and predict the outcome of monetary policy meetings around the world: in the US, a randomly selected triplet of a target rate hike, cut and no-change is correctly classified using the OIS-implied rates in 99.9 and 98 percent of cases five and ten days before a FOMC announcement respectively, which exceeds the prediction accuracy of the federal funds futures-implied rates. We report similarly high prediction accuracy for other developed countries.

Zusammenfassung

Die Varianz-Risikoprämie ist eine der wichtigsten und robustesten Risikoprämien in der Investitionswelt. Im ersten Kapitel dieser Arbeit untersuche ich anhand impliziter risikoneutraler und geschätzter realisierter Kovarianzmatrizen die FX Varianz-Risikoprämie. Ich zeige, dass es Portfolios gibt, die eine positive Varianz-Risikoprämie vorweisen, obwohl die Prämie im Grossen und Ganzen als negativ auftritt. Weiterhin finde ich, dass der US Dollar-Index und das bekannte Carry Trade Portfolio als die zwei Portfolios mit der negativsten Prämie vorkommen, was deren prominente Rolle im FX Asset Pricing festigt. Unter der Faktoren, die die Varianz-Risikoprämie beeinflussen, dokumentiere ich das kurzfristige Momentum (positiv wirkend) und die kürzlich realisierte Varianz (negativ wirkend). Die Varianz-Risikoprämie von dem Carry Trade Portfolio dominiert die des US Dollar Indexes in der Querschnittsanalyse der FX Rendite.

Das zweite Kapitel hat mit der Dynamik der FX Rendite vor Leitzinsentscheidungen der Zentralbanken in entwickelten Ländern zu tun. Dmitry Borisenko und ich finden eine persistente Abwertung der Währungen vor Zinssenkungen und eine Aufwertung vor Zinserhöhungen. Darüber hinaus lassen sich die Entscheidungen gut anhand von Derivaten wie Overnight Index Swaps prognostizieren, so dass eine Handelsstrategie möglich ist: wenn man lange (kurze) Positionen in Währungen der Länder in denen eine Zinssenkung (Zinserhöhung) prognostiziert ist, eröffnet, kann man mit einer jährlichen Rendite von 4.5 Prozent ($t > 2$, das Sharpe Ratio von 0.65) rechnen.

In Kapitel drei untersuche ich mit Nikola Mirkov und Paul Söderlind, ob die Worte der Vorsitzenden der Schweizerischen Nationalbank (SNB) die Erwartunggen der Investoren bezüglich des Franken/Euro Wechselkurses in die gewünschte Richtung steuerten, als die Mindestkursregime gültig war. Wir stellen fest, dass die Vorträge, in denen die Wortkombinazion "unlimited quantities" und "utmost determination" erwähnt werden, die forward-looking Maße der Unsicherheit negativ beeinflussten und damit die Erwartungen in die Richtung der Frankenabwertung steuerten, was die Glaubwürdigkeit des Mindestkurses verstärken konnte.

Letzlich, im vierten Kapitel – der Zusammenarbeit mit Dmitry Borisenko – dokumentiere ich keine signifikante Risikoprämie in den OIS Zinsen, so dass sie als unverzerrte Prognosen der eigenen Bezugszinsen vorkommen. Wir extragieren die in den Swapzinsen impliziten zukünftigen gelpolitischen Leitzinsen, um die Zinsentscheidungen in entwickelten Ländern zu prognostizieren, und zeigen, dass diese über eine ausgezeichnete Prognosefähigkeit verfügen.

Chapter 1

Variance Risk on the FX Market*

Igor Pozdeev†

Abstract

I recover risk-neutralized covariance matrices of currency returns and combine them with ex post realized covariance matrices to study FX variance risk. Eigendecomposition of daily matrix differences conveys three findings: evidence for an overall negative FX variance risk premium; a special place for the US dollar index and Carry trade which exhibit the most negative such premium; existence of strategies with a significantly positive such premium. A more negative variance risk premium of currency portfolios is associated with negative spot return momentum and high recently realized variance. The Carry trade variance risk dominates the US dollar variance risk as a priced factor, contributing to resolution of the differential pricing of "good and bad" carry portfolios. Hedging shocks to news about future FX variance is costly with a horizon-decaying pattern, similar to what is observed on the equity market; however, the costs do not become statistically or economically insignificant even at the longest horizons, suggesting a strong preference of FX market investors to hedge such shocks.

*I would like to thank (in alphabetical order) Anastasia Berezinskaya, Dmitry Borisenko, Peter Carr, Robert Engle, Ralph Koijen, Chiara Legnazzi, Semyon Malamud, Paul Söderlind and Dare Wale, as well as participants of the SFI Research Days, Tandon School of Engineering brown bag seminar and University of St. Gallen seminar for their helpful comments and insightful feedback.
†University of St. Gallen, *E-mail*: igor.pozdeev@unisg.ch

1.1 Introduction

Stochastic variance is considered with little disagreement an inherent feature of financial asset dynamics and a major source of risk for investors. The associated variance risk premium – the average payoff of an asset that is perfectly correlated with a variance process – is a phenomenon intuitive, backed by some of influential asset pricing models and empirically observed. Research on this type of risk was pioneered and has been shaped by its equity market strand, such that several important findings remain endemic to the latter. The FX market strand, which is where my paper belongs, has not only received less attention, but also failed to incorporate staple asset pricing factors and investment vehicles such as the Carry trade into the set of research objects, concentrating instead on USD currency pairs. That said, my work aims to align the two strands by replicating several results from the equity market in the FX market setting, and studying FX market variance risk from the perspective of currency portfolios.

Similar to Carr and Wu (2009), I perform my analysis through the lens of synthetic variance swaps. A variance swap is a contract that pays its buyer an amount equal to the realized variance of a return series in exchange for a predetermined swap rate; the variance risk premium is conveniently defined as the expectation of this net payoff. The notion of "synthetic" means the swaps I am talking about are not actually traded contracts as in Aït-Sahalia, Karaman, and Mancini (2018) and Dew-Becker et al. (2017), but rather no-arbitrage replicas thereof: under certain assumptions, the swap rate can be approximated by the model-free implied variance of the swap underlying, which in turn can be recovered from a cross-section of option prices. I follow Mueller, Stathopoulos, and Vedolin (2017) to recover such variances for all cross-rates of currencies in my sample and construct option-implied *covariance matrices* of currency returns, from which implied variances of generic portfolios naturally follow.

That said, I present my findings as follows. First, I show that the difference between ex post realized and option-implied covariance matrices is most of the time an indefinite, occasionally a negative definite and never a positive definite matrix. This means that in theory, strategies of both negative and positive variance swap payoffs are possible to construct, and hence that it is hard to make a definite statement about the sign of "the" FX market variance risk premium, although there is a stronger case for an overall negative one. Interestingly, the two eigenvectors corresponding to the most negative eigenvalues in the decomposition of the above matrix difference are aligned with the US dollar index and the Carry trade portfolio respectively, implying that these two portfolios are associated with the most negative variance risk premium of all and thus cementing the special place they have in the FX market fabric.

Second, I document that past spot return information influences conditional variance risk premium in the cross-section and over time. To do so, I take a set of portfolios and regress payoffs of swaps written on their variance onto a set of variables observed at swap inception. I find that portfolios with negative momentum exhibit a more negative price of variance risk, and hence are *ceteris paribus* relatively more expensive to insure against variance increase. Portfolios with a high recently realized variance of returns behave in a similar way, meaning that after periods of high realized variance, investors become more variance risk-averse.

Third, I turn attention to staple FX portfolios and show that protection against higher variance is priced for the Carry trade, Momentum, Value and the US dollar index – in the sense that a buyer of a swap written on the variance of any of them expects to lose money on average. These portfolios hold a special place in the FX market research. The former three have been discovered as profitable trading strategies likely exposed to non-diversifiable risks (Lustig, Roussanov, and Verdelhan (2011), Asness, Moskowitz, and Pedersen (2013), Menkhoff et al. (2011))[1]; the US dollar index has been shown as an important driver of the time-series dynamics of currency returns (Verdelhan (2018)), and is clearly equivalent to the Dollar carry strategy (Lustig, Roussanov, and Verdelhan (2014)) in terms of conditional variance of returns. The variance risk premium is negative and similar in magnitude across strategies; for the US dollar index it is comparable to the US equity market variance risk premium. Curiously, the Carry trade strategy, much more volatile than the other FX strategies and notorious for crashes and variance spikes, is far from being associated with the most negative or significant price of own variance risk, dominated in that respect by the US dollar index and the Value strategy in terms of statistical and economic significance respectively. That said, my findings add to the puzzle reported by Caballero and Doyle (2012) and Jurek (2014) who note that it has been strikingly cheap to hedge the Carry trade portfolio with FX options.

Fourth, motivated by my first set of results, I use the US dollar index and Carry trade variance swap payoff series as risk factors in the linear asset pricing framework to show that information in the former is subsumed by that in the latter. With the "good and bad" carry trade portfolios of Bekaert and Panayotov (2018) for test assets, I find only the Carry trade variance risk to be priced, the estimated price of risk being negative and close to the sample average value of the factor, whereas the US dollar index variance risk price is positive in some specifications and never significant.

Fifth and finally, I provide FX market-specific support for the equity market evidence that shocks to *news about future variance* are priced the weaker the more distant the future is, as if investors cared more about the next-period variance than about the variance expected farther ahead. In asset pricing models with preference for early resolution of uncertainty, and in macroeconomic models where shocks to variance expected tomorrow can lead to higher variance today, investors are willing to pay for protection against shocks to expected variance. Surprisingly, Dew-Becker et al. (2017) report that hedging shocks

[1] These portfolio are also on the list of tracked FX indexes published by Deutsche Bank.

to expected long-term US stock market variance has been remarkably cheap, the Sharpe ratios of hedging strategies quickly going to zero and disappearing beyond the horizon of one quarter. To replicate their result in the FX market setting, I construct synthetic forward variance claims and calculate payoffs from rolling these claims over for various FX portfolios. I find that both payoffs and Sharpe ratios of rolling over forward claims on future quarterly variance tend to diminish with horizon. However, the decline is not as steep as on the equity market, and even for hedging shocks to the variance 4 quarters from now, the costs are statistically different from zero, and Sharpe ratios – comparable to those of the portfolio excess returns themselves. I also zoom in on the shortest segment of 1-4 months and discover a more or less flat term structure of shocks to future variance prices and respective Sharpe ratios. Overall, my findings point to a stronger desire of FX market investors to hedge shocks to long-term variance.

My work belongs to and draws from the relatively young literature on the FX market variance risk as a subset of the variance risk research in general. Negative price of variance risk has been previously documented for the US aggregate stock market (Carr and Wu (2009), Bollerslev, Tauchen, and Zhou (2009)), non-US aggregate stock markets (Bollerslev, Marrone, et al. (2014)), US Treasury bonds (Choi, Mueller, and Vedolin (2017)) and commodities (Tee and Ting (2017)). Della Corte, Ramadorai, and Sarno (2016) present a first evidence of the variance risk premium on the FX market by constructing a dollar-neutral portfolio of currencies sorted on a crude measure of conditional variance risk premium. Over-weighting (under-weighting) currencies with high (low) premium, the authors construct a strategy that enhances the mean-variance investment opportunities on the FX market. Ammann and Buesser (2013) find further support for negative variance risk premium in individual exchange rates. However, variance risk of FX portfolios such as the Carry trade and the US dollar index has not received similar attention.

Since variance of a portfolio is a quadratic form of the covariance matrix of portfolio constituents and portfolio weights, I first develop vector representation of currency trading strategies. Then, I recover risk-neutralized covariance matrices of currency returns by making use of the concept of model-free implied variance developed by Britten-Jones and Neuberger (2000) and Carr and Wu (2009) and the assumption of no triangular arbitrage. A similar exercise has been undertaken before. Walter and Lopez (2000) and Mueller, Stathopoulos, and Vedolin (2017) construct option-implied correlations between appreciation rates of currencies against the US dollar: the former paper criticizes the minuscule information content thereof, while the latter explores the properties of a trading strategy that is long (short) currencies with the highest (lowest) loading on the measure of FX correlation risk. My work differs from these studies both in research questions and methodology: namely, the covariances and correlations are never the center of my research, but rather a tool to calculate portfolio variances.[2] Relatedly, Jurek (2014)

[2] One clear advantage of this is while the payoff of a variance swap can at least in theory be replicated

constructs the analogue of VIX for the carry trade portfolio, but only uses it to highlight the GARCH-in-mean effect in carry trade returns.

The rest of the paper is structured as follows. Section 1.2 describes the data I use and outlines the vector representation of currency trading strategies, discusses recovery of the option-implied covariances and construction of the variance risk premium estimates. Sections 1.3 and 1.4 present the findings. Section 1.5 concludes.

1.2 Data and methodology

1.2.1 Notation

By default, I treat currencies as assets from the point of view of an American investor, such that exchange rate S_x is the dollar "price" of currency x and could be referred to by quote XXXUSD if the three-letter ISO code of currency x is XXX. All τ-period "returns" defined as $R_x(t, t+\tau) = S_x(t+\tau)/S_x(t)$ are thus by default appreciation rates against USD.

When explicitly written with a double subscript S_{xy}, exchange rate against currency y is meant rather than against the US dollar, which can be thought of as "price" of currency x expressed in units of currency y, and "return" is then the appreciation rate of currency x against y.

In what follows, $s_x = \log S_x$, such that the log-return of currency x is defined as:

$$r_x(t, t+\tau) = \Delta s_x(t+\tau) = s_x(t+\tau) - s_x(t). \tag{1.1}$$

1.2.2 Data

I use FX option data from Bloomberg. In the next few paragraphs, I only present the basic facts about FX option data; a comprehensive introduction into the FX option market conventions is given in Wystup (2007) and Malz (2014).

Prices of FX options tend to be expressed in terms of the option Black-Scholes implied volatility (IV). This does not assume that the Black and Scholes (1973) model is understood to hold, but rather represents a one-to-one continuous mapping from the space of currency-denominated option prices to the space of unitless volatility, which allows for easier comparison between options of different strikes and maturities. Bloomberg provides implied volatility quotes against forward deltas (δ) rather than against strike prices,

with a portfolio of options (see Carr and Madan (1998)), no such replication is possible for individual correlations, which might impact their information content and can serve as the resolution of the critique of Walter and Lopez (2000). I thank Peter Carr for this observation.

the forward delta[3] being the derivative of the Black-Scholes pricing function with respect to the forward rate of the underlying. Henceforth, delta is understood to be the forward delta. Just as implied volatility is a mapping from the option price, delta is a mapping from the strike price. As deltas are bounded between 0 and 1 in magnitude, they are usually multiplied by 100 for quotation purposes to become numbers such as 25, 15 etc.

The most liquid part of the FX option market is concentrated in at-the-money options (ATM) and option contracts, of which Bloomberg provides risk reversals (RR) and butterfly spreads (BF). The notion of at-the-money in Bloomberg is the so-called "delta parity", implying that a call option is at-the-money if it and an otherwise identical put option have the same absolute delta. The other two instruments are linear combinations of plain vanilla call and put options: risk reversals give the holder exposure to the skewness, and butterfly spreads – to the volatility of the underlying exchange rates. That said, for any given day, there are implied volatilities of 10-, 15-, 25- and 35-delta contracts of both types as well as of one ATM option provided by Bloomberg, a total of nine quotes[4]. These are indicative quotes, collected by the data vendor from its suppliers at a particular time of day (usually right after closing of exchanges in the ET time zone) and are snapped at 17:00 New York time. The contracts are conveniently structured to be of constant maturity of 1, 2, 3, 4, 6, 9, 12 months.

As the risk reversals and butterfly spreads are linear combinations of put (P) and call (C) options, it is possible to solve for two call option IVs given the IVs of both contracts and of the ATM. The only necessary assumption is that of no arbitrage (including the put-call parity relation). For example, as shown in Malz (2014):

$$\sigma(C(\delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) + 0.5\sigma(RR(\delta)), \tag{1.2}$$
$$\sigma(C(1-\delta)) = \sigma(ATM(\delta)) + \sigma(BF(\delta)) - 0.5\sigma(RR(\delta)) \tag{1.3}$$

where $\sigma(Y(\delta))$ is the IV of contract of type $Y \in \{ATM, BF, RR, C, P\}$ having delta δ. Omitted are the maturity of the contracts and the time subscripts required to be the same for all contracts in the above equations. As discussed in Reiswich and Wystup (2009), (1.2)-(1.3) is only valid for small risk reversal implied volatilities. Still, it offers a handy, widely used relation, does not rely on a parametric form proposed in that paper and thus does not conflict with the cubic spline volatility smile interpolation I use later.

That said, it is possible to obtain $2N+1$ distinct (δ, σ) points from N (RR, BF) quotes and one ATM quote, from where it is straightforward to get to price-strike pairs (C, K) as shown in Wystup (2007). The data other than option quotes needed for the calculations is also collected from Bloomberg. Specifically, for every currency pair I collect the spot exchange rate, and – for each considered maturity – the forward rate rate and two OIS

[3] As provided by Bloomberg, delta is net of currency premium.

[4] Sometimes certain quotes would be missing on a given day; most often these are the less liquid 10- and 35-delta ones.

rates as proxy for the risk-free rates. I substitute the risk-free rate of the less traded currency with a synthetic risk-free rate obtained from the covered interest parity (CIP), whereby the ranking is by turnover of OTC foreign exchange instruments reported in BIS (2016a). For instance, in the case of AUDCHF, the true Australian dollar OIS rate will be taken and used to infer the Swiss franc risk-free rate, although the OIS rate for the latter currency is also available.

For construction of portfolio spot and excess returns, I use the same spot and forward rates from Bloomberg. The PPP data used for construction of the Value signal is from OECD.

My sample of currencies includes the Australian and Canadian dollar, Swiss franc, euro, British pound, Japanese yen, New Zealand and US dollar; the sample period is from 06/2009 to 12/2017.

Local-economy stock indexes used for construction of equity variance swap rates are obtained from websites of local stock exchanges: for Australia *S&P/ASX 300*, for Canada *S&P/TSX Composite*, for Switzerland *SMI*, for the Eurozone *Euro STOXX*, for the UK *FTSE 100*, for Japan *NIKKEI 225*, and for the US *S&P500*. Respective VIX-like indexes (usually referring to the same basket of stocks that the index is comprised by) are from Bloomberg. As there is no VIX-like index for the New Zealand market, I exclude this country from calculations where both the stock and FX variance risk premium is considered.

1.2.3 Variance swap returns and variance risk premium

Imagine an investor at time t wishing to receive a payoff equal to the variance $Var_{t+\tau}(r(t, t+\tau))$ of a return over some time interval $(t, t+\tau)$. Two problems arise: first, since the variance is a latent characteristic of the return process, it is not observed and has to be estimated from data at time $t+\tau$ to be paid out; second, the fair price of this payoff at time t has to be determined.

As a solution to the first problem, Corollary 1 in Andersen et al. (2003) equates[5] the conditional variance of an arbitrage-free process to the conditional expectation of its quadratic variation. Assume that exchange rate S follows an arbitrage-free process $\{S(t)\}$, and let $s(t) = \log S(t)$ as before, such that the continuously compounded appreciation rate is

$$r(t) = ds(t) = \lim_{\Delta t \to 0} s(t + \Delta t) - s(t). \tag{1.4}$$

[5]Under certain yet not implausible assumptions discussed therein.

The quadratic variation accumulated from time t to time $t+\tau$ is defined as:

$$[r, r]_{t,t+\tau} = \int_t^{t+\tau} (ds(k))^2, \tag{1.5}$$

and is closely related to the variance of the process:

$$Var_t(r(t, t+\tau)) = E_t([r, r]_{t,t+\tau}) = E_t\left(\int_t^{t+\tau} (ds(k))^2\right), \tag{1.6}$$

The same Corollary suggests that a natural *ex post* estimator of the variance on the left-hand side of eq. (1.5) – the quantity the investor expects to receive – is the discretized version of the right-hand side[6]:

$$Var_{t+\tau}(r(t, t+\tau)) = \frac{254}{D_\tau}\sum_{d=1}^{D_\tau} r\left(t + \frac{d-1}{D_\tau}, t + \frac{d}{D_\tau}\right)^2 = RV(t, t+\tau), \tag{1.7}$$

where $t + d/D_\tau$ denotes day d of time interval $(t, t+\tau)$, D_τ is the total number of days in that interval, and 254 is the annualisation factor approximately equal to the number of trading days in a year. This quantity is also called the realized variance, hence mnemonic RV. Eq. (1.7) assumes that daily returns have zero mean: although this seems restrictive and arguably more suitable at frequencies higher than daily, average currency spot returns are notoriously indistinguishable from zero (cf. Lustig, Roussanov, and Verdelhan (2011)), such that slightly biased but lower-variance estimators of type (1.7) have been favored in related econometric literature.

The fair price to swap the quantity in eq. (1.5) for (such that no money changes hands at time t) is by the standard argument its risk-neutralized expectation:

$$IV_t(\tau) = E_t^Q\left(\int_t^{t+\tau} (ds(k))^2\right), \tag{1.8}$$

also called the variance swap rate. Mnemonic IV has to do with the fact that this rate is essentially the model-free implied variance of the log-price process $\{s(t)\}$. Britten-Jones and Neuberger (2000), building on the results of Carr and Madan (1998) and Breeden and Litzenberger (1978), equate[7] the conditional $\mathbb{Q}$-expectation of the accumulated quadratic variation to the price of a continuous portfolio of options expiring in $(t+\tau)$, weighted by strikes:

$$E_t^Q\left(\int_t^{t+1} (ds(k)^2\right) = 2\int_0^\infty \frac{C_t(\tau, X) - (S(t) - X)^+}{X^2} dX, \tag{1.9}$$

where $C_t(\tau, X)$ is the time-t price of a call option on S with strike X and maturity of τ, and $(v)^+ = \max(v, 0)$. The integral in eq. (1.9) can be evaluated numerically, but a careful

[6] Estimators of this type are commonly used, see for instance Moreira and Muir (2017), Trolle and Schwartz (2010) and for daily observed FX returns Della Corte, Ramadorai, and Sarno (2016).

[7] Under assumptions that were relaxed by Jiang and Tian (2005) to include jump-diffusion processes.

inspection reveals two potential sources of errors in doing so: first, while the integration in eq. (1.9) is from zero to infinity, options are traded over a much narrower range, and second, even in that range, strikes are far from being sampled continuously. Addressing these issues, I follow the literature and take the usual steps[8], which are graphically represented in Figure 1.1. First, as shown by the hollow dots, the observed option prices are transformed into the Black-Scholes implied volatilities to obtain a volatility smile. Then, as shown by the solid line, the smile is interpolated within the available strike range using a cubic spline. Third, as shown by the dashed line, the smile is extrapolated by keeping it constant at the level of the endpoints. Finally, the Black-Scholes volatilities are converted back into prices. For the estimations, I use a grid of 2000 points over the moneyness range between 2/3 and 5/3, and the Simpson's rule to perform the integration.[9]

Figure 1.1: Preparing options data for integration

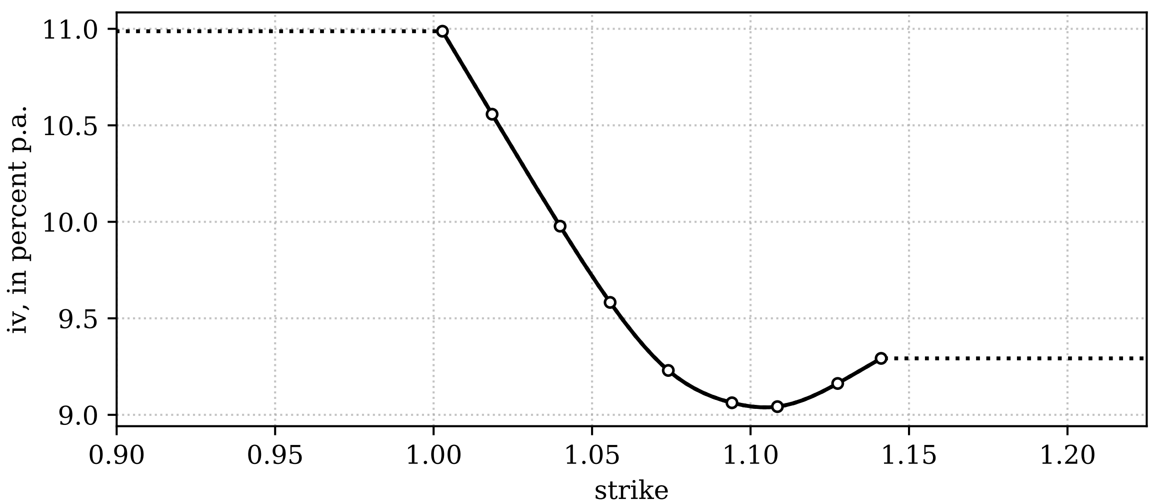

This figure shows the basic steps of inter- and extrapolating the observed option prices, before attempting the numerical integration of eq. (1.9). The hollow points correspond to the actually observed Black-Scholes implied volatilities of plain vanilla call options. The solid line depicts the cubic spline fitted to these volatilities to produce a smoothly interpolated smile. The smile is then extrapolated to the left and right of the observed strike range by keeping the respective endpoint volatilities constant, as shown by the dashed line. The example features call option prices extracted from the 10-, 15-, 25- and 35-delta risk reversals and butterfly spreads as well as the at-the-money option on EURUSD on 01/17/2017.

That said, in form of a variance swap the investor purchases protection against rising variance, the time-$(t+\tau)$ return amounting to:

$$vs(t, t+\tau) = RV(t, t+\tau) - IV_t(\tau), \tag{1.10}$$

[8]The same approach, if not without slight variations, has been used by Jiang and Tian (2005), Driessen, Maenhout, and Vilkov (2009), Buraschi, Kosowski, and Trojani (2014), Della Corte, Ramadorai, and Sarno (2016) and many others.

[9]Another way to arrive at an estimate of a risk-neutral moment would be through calibration of a parametric density to the set of observed option prices, from which calculation of moments is straightforward (see Mirkov, Pozdeev, and Söderlind (2016) for an example). I have ascertained that deviating from the model-free approach towards a parametric one does not lead to much different variance estimates.

where $RV(t, t+\tau)$ is the realized variance over time interval $(t, t+\tau)$, and $IV_t(\tau)$ is the model-free implied variance observed from a cross-section of options at time t.

The variance risk premium is defined as the expected value of the variance swap return in (1.10):

$$vrp_t(\tau) = E_t\left(vs(t, t+\tau)\right) = E_t\left(RV(t, t+\tau)\right) - IV_t(\tau), \tag{1.11}$$

Thus, it is the difference between the objective and risk-neutral expectation of the future realized variance of a stochastic return: if the difference is negative for some asset, investors are ready to pay for hendging the return variance of that asset.

1.2.4 Vector representation of currency trading strategies

Absence of triangular arbitrage allows to represent any currency position or trading strategy in a vector form, using only exchange rates against one common currency. A vector representation is necessary for computation of moments of trading strategy returns.

I define an N-currency portfolio as a zero-leverage[10] dynamic trading strategy, rebalanced monthly and represented with a vector of weights $(w_{1,t}, w_{2,t}, \ldots, w_{N,t})'$ known at the end of the previous month. For instance, any position involving the euro, the Australian and US dollar can be represented with a 2×1 vector $(w_{aud}, w_{eur})'$ of weights of AUDUSD and EURUSD. The strategy of having 40% of the portfolio in long AUDUSD and 60% in long EURUSD is defined as $(0.4, 0.6)'$, and the weights have to sum up to 1 in absolute value for the zero leverage constraint to bind. The strategy of having 40% in short AUD and 60% in short EUR is represented as $(-0.4, -0.6)'$, the weights again summing up to 1 in absolute value. Now, having 100% in long AUD and 100% in short EUR, represented as $(1.0, -1.0)'$ would also be a valid zero-leverage portfolio, as it is tantamount to be 100% long AUDEUR. In the latter case, the weights do not sum up to 1, but rather the short leg and the long leg do so separately. With that in mind, a currency portfolio must be a process $\{v_{t+1}\}$ of the form:

$$v_{t+1} = (w_{1,t}, w_{2,t}, \ldots, w_{n,t})', \tag{1.12}$$

$$E_t[v_{t+1}] = v_{t+1}, \tag{1.13}$$

$$\sum_{n=1}^{N} |w_{i,t}| + \left|\sum_{n=1}^{N} w_{i,t}\right| = 2, \tag{1.14}$$

where eq. (1.13) says that the composition of the portfolio at time *t+1* is known at time t. The constraint in (1.14) can be graphically summarized in Figure 1.2 for the example with AUDUSD and EURUSD: all combinations on the solid rhombus are valid (w_{aud}, w_{eur}) portfolios. Section 1.2.6 contains examples of popular trading strategies in

[10] On the ForEx, it would mean that having \$1 in the margin account it is only possible to open \$1 worth of positions.

vector representation. To stress that the time-*(t+1)* portfolio composition is known at

Figure 1.2: Portfolio constraints with AUDUSD and EURUSD.

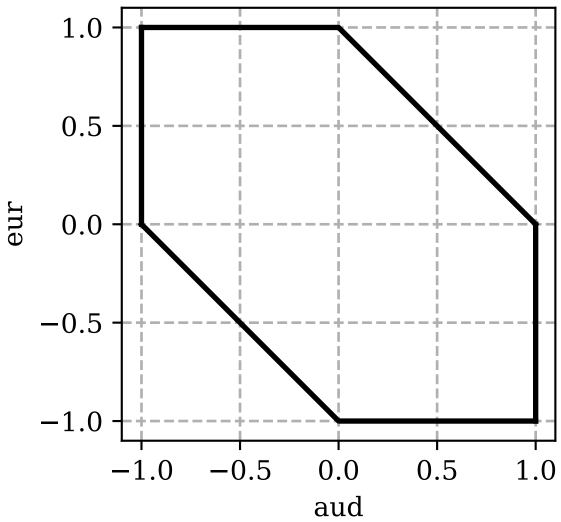

time t, I introduce $w_t = E_t[v_{t+1}] = (w_{1,t+1}, w_{2,t+1}, \ldots, w_{N,t+1})'$. Also, I denote the time-*(t+1)* return of any strategy as f_{t+1} to differentiate it from individual currency returns. This return is calculated in the usual way:

$$f_{t+1} = w_t' r_{t+1}, \tag{1.15}$$

where $r = (r_1, r_2, \ldots, r_n)'$ is the vector of individual currency returns. The conditional variance of the strategy return is:

$$Var_t(f_{t+1}) = w_t' \Omega_t w_t, \tag{1.16}$$

where Ω_t is the conditional covariance matrix of currency returns against the US dollar or any other currency.

1.2.5 Covariance matrices of currency returns

The cornerstone of this paper is the conditional covariance matrix of time-$(t+\tau)$ currency appreciation rates against a common counter currency, whereby the conditioning is w.r.t. the time-t information. Indexing the rows and columns of any such matrix with the base currencies, its (x, y) element reads:

$$\Omega_t[x, y] = \begin{cases} Var_t(r_x(t, t+\tau)), & x = y, \\ Cov_t\left(r_x(t, t+\tau), r_y(t, t+\tau)\right), & x \neq y \end{cases}$$

Absence of triangular arbitrage implies that the log-appreciation rate of currency x against y can be expressed in terms of their log-appreciation rates against a common currency (time subscripts can be dropped as long as returns are contemporaneous):

$$r_{xy} = r_x - r_y. \tag{1.17}$$

Applying the variance operator to both sides of eq. (1.17) results in:

$$Var(r_{xy}) = Var(r_x) + Var(r_y) - 2Cov(r_x, r_y),$$

which can be rearranged to isolate the covariance as follows:

$$Cov(r_x, r_y) = \frac{1}{2}(Var(r_x) + Var(r_y) - Var(r_{xy})). \tag{1.18}$$

The latter equation obviously holds irrespective of the time subscripts, of the set of information used for conditioning the moments, and of the measure under which the moments are taken. Thus, to obtain the covariance between returns of x and y one needs the variance of the appreciation rate of x against y as well as of each of them against a common currency.

1.2.6 Currency portfolios

US dollar index, *USD*

A return of a currency index is defined as an equally weighted average appreciation rate of foreign currencies against that particular currency: when the index goes up, the currency *depreciates* to the basket of other currencies. For each currency, the composition of this portfolio is the same after each rebalancing, e.g. for the US dollar:

aud	cad	chf	eur	gbp	jpy	nzd
1/7	1/7	1/7	1/7	1/7	1/7	1/7

Carry trade, *CAR*

As in Lustig, Roussanov, and Verdelhan (2011), currencies are sorted by how much risk-free interest rates in respective economies exceed that in the US. Long positions are opened in the two currencies with the highest, and short position – in the two with the lowest such difference. Instead of taking any particular type of risk-free rate, I rely on the covered interest parity and proxy the above interest rate difference as the negative of the currency's forward discount, smoothing the series over 3 months to avoid FX market microstructure effects.[11] A common representation of the carry trade strategy is as follows:

aud	cad	chf	eur	gbp	jpy	nzd
0.5	0	−0.5	0	0	−0.5	0.5

Momentum, *MOM*

As in Asness, Moskowitz, and Pedersen (2013), currencies are sorted by their cumulative appreciation against the US dollar over the past 12 months excluding the most recent month. Long positions are opened in the two currencies that have appreciated the most, and short positions – in those that have appreciated the least.

Value, *VAL*

As in Della Corte, Ramadorai, and Sarno (2016), currencies are sorted by their real exchange rate against the US dollar. Long positions are opened in the two currencies with the lowest, and short positions – in those with the highest real exchange rate. The real exchange rate is defined as $REER_t = PPP_t/S_t$, where PPP_t is the purchasing power parity rate from OECD data, and S_t is the nominal exchange rate.

Variance risk premium, *VRP*

As in Della Corte, Ramadorai, and Sarno (2016), currencies are sorted by a proxy of the time-t variance risk premium for holding that currency against USD. The premium is proxied as the difference between the realized in the previous year and the 1-year model-free implied variance of the currency appreciation rate against the US dollar. Long positions are opened in two currencies with the largest, and short – in those with the lowest such difference.

1.3 General results about FX variance risk

1.3.1 Extreme attainable variance risk premia

Access to both realized and implied covariance matrices makes it possible to determine which portfolios are associated with a consistently negative variance risk premium, and

[11] See Baba and Packer (2009) and Du, Tepper, and Verdelhan (2018) for a recent discussion on the parity deviations during the financial crisis.

which – if any at all – with a positive one. The exercise can be restated as an eigenvalue problem.

As mentioned in Section 1.2.3 (specifically, eq. (1.11)), variance risk premium of a portfolio is equal to the expectation of the payoff of a swap written on the portfolio's variance. The payoff equation trivially involves quadratic forms defined by the objective and risk-neutralized covariance matrices:

$$vs(t, t+\tau) = w_i' \Omega^{\mathbb{P}}(t, t+\tau) w_i - w_i' \Omega_t^{\mathbb{Q}}(\tau) w_i \tag{1.19}$$

$$= w_i' \underbrace{\left(\Omega^{\mathbb{P}}(t, t+\tau) - \Omega_t^{\mathbb{Q}}(\tau)\right)}_{\Theta(t, t+\tau)} w_i, \tag{1.20}$$

where w_i is a vector of portfolio weights, and Θ, which I call *variance swap payoff structure matrix*, is introduced to denote the difference in question. Period-by-period, a condition necessary and sufficient to make *any* quadratic form defined by this matrix negative (positive), Θ must be negative (positive) definite, which it would if all of its eigenvalues are less (greater) than zero. Hence, I calculate the time series of eigenvalues of the realized $\Theta(t, t+\tau)$. I find that the proportion of days when only a negative variance swap payoff could have been realized, rises from 0.01 at the 1-month to 0.15 at the 4-month horizon, whereas the proportion of days with only a positive payoff being possible is exactly zero irrespective of the horizon. This speaks in favor of an overall negative FX variance risk premium, but is also clearly indicative of a theoretical possibility to construct portfolios with a positive one.

By construction, eigenvalues of Θ are equivalent to variance swap payoffs for portfolios formed with corresponding eigenvectors as weights. Since Θ is Hermitian, the lowest (highest) value is paired with the vector that minimizes (maximizes) the quadratic form in eq. (1.20) under the condition that vector elements sum up to one when squared.[12] For the 3-month horizon only, Figure 1.3 presents the time series average of the eigenvalues ranked from the lowest to highest, as well as the corresponding eigenvectors. There is strong evidence that a broad range of average payoffs have been possible, from very negative for portfolio 1 to positive for portfolio 7. Interestingly, portfolio 1 – associated with the lowest eigenvalues – is essentially the US Dollar index, as the weights of different currencies are all similar in magnitude and sign and identified precisely. Portfolio 2 is reminiscent of a typical carry trade, separating the Australian, New Zealand and Canadian dollar from the Japanese yen, Swiss franc and euro, which are the high and low interest rate currencies respectively. This further stresses the special role the US dollar index and the Carry portfolio play on the FX market. At the same time, portfolio 7 – the one associated with the highest eigenvalues – turns out to exhibit a consistently positive variance risk premium, suggesting a possibility of constructing such a .

[12] That said, the portfolios obtained from such eigendecomposition are leveraged.

Figure 1.3: Eigendecomposition of variance swap payoff structure matrices

This figure shows the times series average of eigenvectors (left heatmap) and corresponding eigenvalues (right heatmap) calculated daily by performing eigendecomposition of the difference between two covariance matrices: the one realized over a 3-month period and the model-free option-implied one, observed at the start of the period (the observations of these matrix difference are thus daily and overlapping). The eigenvectors can be thought of as weights of USD exchange rates in portfolios, and eigenvalues – as portfolio variance swap payoffs. Values exceeding in magnitude $2\times$ own standard error are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment with the automatic lag selection of Newey and West (1994). The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2009 to 06/2018.

Mostly identical results emerge when realized covariance entering the computation of variance swap payoff structure matrices is substituted with expected covariance, proxied by past realized covariance in the spirit of many previous variance risk premium studies.[13] The *ex ante* variance risk premium in this case is similarly negative (positive) for portfolios corresponding to the lowest (highest) eigenvalues, the time-series average of the first two eigenvectors again mimicking the US dollar index and a typical carry trade portfolio.

1.3.2 Explaining variance risk premium of FX strategies

Given the observed heterogeneity in variance risk premium estimates, a naturally arising question is whether the conditional variance risk premium is explainable by characteristics. This question translates to estimation of the coefficients β in the following panel relation:

$$w_i'\Theta(t,t+\tau)w_i = \beta_i' X_{i,t} + \epsilon_{i,t}, \tag{1.21}$$

[13]See, for example, Bollerslev, Tauchen, and Zhou (2009), Della Corte, Ramadorai, and Sarno (2016), Zhou (2017) and others.

where the left-hand side is the variance swap payoff of portfolio i reaped at time $t + \tau$, Θ and w_i are as defined in eq. (1.20), and X_i is a vector of portfolio characteristics. I choose X_i to comprise the portfolio average interest rate difference (weighted by w_i), momentum signal (return of portfolio i in the previous year without the most recent month), value signal (weighted average of the logarithms of individual value signals as in Section 1.2.6) and the recently realized variance (average squared return over the 5 days leading to t). Since there are infinitely many portfolios characterized by weights w_i to estimate eq. (1.21), I only concentrate on the subset of seven portfolios constructed from rescaled eigenvectors in Figure 1.3. The choice is admittedly arbitrary, but as I describe earlier, the portfolios therein include a version of the US dollar index, a static carry trade portfolio etc., thus offering a representative range of test assets. For this exercise, I use the 3-month variance risk premium estimates as the most statistically significant ones. As for the estimator, I choose panel OLS with time fixed effects and the Fama-MacBeth procedure. Additionally, I control for entity fixed effects in a separate specification. Within the panel OLS framework, time fixed effects are needed to pinpoint the sources of cross-sectional variation in risk premium estimates, and the entity fixed effects would allow to zoom in on the sources of time series variation over and above a common trend. The Fama-MacBeth estimator is used as a robustness check.

The results are presented in Table 1.1. Overall, it is the backward-looking spot return information that is a significant determinant of cross-sectional differences in FX variance risk premium estimates and a driver of conditional risk premium dynamics. Portfolios that have enjoyed positive spot return over the previous incomplete year tend to be cheaper to buy variance protection for, as evidenced by the positive and significant coefficient, robust across all estimators. Each extra percentage point of positive return translates *ceteris paribus* to 50-100%% of extra positive difference between the realized and implied variance. Equally significant and robust is the effect of higher recently realized variance: every extra squared percentage point thereof is associated with the variance risk premium becoming more negative by 1-2%%.

As to the other variables, the relation is either not significant, or not robust, or both. Portfolios with a relatively higher forward discount tend to exhibit slightly more positive variance risk premium, although the effect is reversed in the time dimension and barely detectable with the Fama-MacBeth estimator. Portfolios of currencies that are "cheap" in real terms are expected to have a more negative variance risk premium when judged by the panel OLS coefficients, but this is at stark odds with the Fama-MacBeth estimation results.

Table 1.1: Explaining differences in variance risk premium estimates

	(1)	(2)	(3)
estimator	panel OLS	panel OLS	Fama-MacBeth
N obs.	17136	17136	17136
R^2	0.08	0.07	0.09
R^2 (within)	0.03	0.10	0.01
R^2 (between)	−0.39	0.31	0.97
R^2 (overall)	0.00	0.11	0.09
forward discount	3.44	−3.65	1.05
	(1.73)	(−0.73)	(0.47)
momentum signal	70.62	55.67	100.56
	(3.72)	(2.71)	(2.92)
value signal	−11.10	−41.69	44.63
	(−0.45)	(−1.46)	(3.10)
recent RV	−1.13	−1.12	−2.10
	(−3.28)	(−3.21)	(−2.21)
time f.e.	yes	yes	n/a
entity f.e.	no	yes	n/a

This table shows estimates of coefficients in eq. (1.21) which relates subsequent 3-month portfolio variance swap payoffs to portfolio characteristics observed at swap inception. Each column features a different estimator: in (1) panel OLS with time fixed effects is used, in (2) – panel OLS with time and entity fixed effects, and in (3) – Fama-MacBeth methodology. All variance-related variables are in percent-squared per year; returns are in percent per year, and value signal is dimensionless. In parentheses below coefficient estimates are respective t-statistics, calculated using Driscoll and Kraay (1998) standard errors. The sample period is from 01/2009 to 06/2018.

1.4 Variance risk of popular FX trading strategies

1.4.1 Time series of FX strategies' variance swap payoffs

Next, I turn to documenting the variance risk premium estimates for a number of popular FX portfolios. As discussed in Section 1.2.3 (specifically, see eq. (1.11)), the unconditional premium can be expressed as the unconditional mean of variance swap payoffs, or as the difference between the average realized and implied variance.

Figure 1.4a shows (the square root of) the average realized and implied variance of portfolio returns for horizons of 1 to 12 months. Not surprisingly, the carry trade strategy is the most volatile of all, both in terms of the average realized and average implied variance, at all horizons, while the US dollar index is the least volatile. Together with the Value strategy, the Carry is also the However, at least at the shorter horizons (1-4 months), the Carry is not the A restatement of the previously reported result from

Section 1.3.1, the Carry trade is however not the most expensive strategy to hedge the variance risk of, as Figure 1.4b suggests: its unconditional variance risk premium is either exceeded or matched in magnitude by that of the Value strategy, and both fall short of the US dollar index in terms of the Sharpe ratio of variance swap payoffs. This appears puzzling, given the reputation of the Carry trade as a strategy prone to crashes and volatility spikes much more than the other considered portfolios, and resembles the findings of Jurek (2014) who documents that hedging its downside with FX options increases the performance dramatically, as if the options on the carry-related cross-rates such as AUDJPY were "too cheap".

Table 1.2 suggests a possible venue for reconciliation of the puzzle: the time series average values of variance swap payoffs are pushed upwards by infrequent large positive observations. The *median* values on the other hand are uniformly twice as negative for the Carry as for the other strategies. Admittedly, these facts distort the analysis somewhat, but not dramatically: even at twice as higher a Sharpe ratio, the carry trade variance swap payoffs would only come close to resemble the US dollar index in terms of variance-riskiness.

FX variance risk premium is dwarfed by that of local stock market returns across the economies in my sample, as Figure 1.5 indicates. Therein, I plot average local stock market variance swap *returns* and Sharpe ratios against the corresponding currency index values. A currency index is defined for different counter currencies similar to the US dollar index. A variance swap return is defined by scaling the variance swap payoff by the inverse of the swap rate, not dissimilar from return on a fully collateralized swap position:

$$vs(t, t+\tau) = \frac{RV(t, t+\tau) - IV_t(\tau)}{IV_t(\tau)} \times 100 \times \frac{12}{\tau}, \tag{1.22}$$

For this exercise, swap returns are preferred to payoffs because stock markets are on average much more volatile than the FX market, which would distort the cross-market comparison. At the 1-month horizon depicted in the Figure, FX variance swap returns and respective Sharpe ratios are twice as small as their equity market counterparts.

1.4.2 Asset pricing with variance risk

Given that the US dollar index and Carry trade strategy have been documented to play an important role in FX asset pricing (see e.g. Lustig, Roussanov, and Verdelhan (2011) and Verdelhan (2018)), and taking into account my previously reported findings, namely coincidence of their composition with principal component loadings, as well as significance of respective variance risk premium estimates – it is of interest to test the corresponding sources of variance risk for cross-sectional asset pricing properties. To do so, I use a linear asset pricing framework and a range of "good and bad" carry trade portfolios constructed as in Bekaert and Panayotov (2018). Each such portfolio represents a carry trade set up on a narrower subsample of G10 currencies[14]. Bekaert and Panayotov (2018) show that these portfolios differ a lot in terms of profitability and distribution characteristics, and as such they provide a fruitful soil for testing asset pricing models.[15]

The model I estimate assumes that the stochastic discount factor of a representative investor linearly depends – among other things – on the US dollar and Carry trade variance swap payoffs, which leads to the following representation of the expected return of portfolio i:[16]

$$E[rx_i(t, t+\tau)] = \beta_i(\tau)'\lambda(\tau), \tag{1.23}$$

where rx_i is the excess return of the portfolio, λ is a vector of risk premia ("prices of risk"), and β_i is the vector of portfolio loadings on the risk factor ("quantities of risk"), which I allow to differ by horizon τ. Hence, the model implies that the cross-sectional differences in expected portfolio returns arise due to differences in their loadings on the risk factors, since the prices of risk are the same for each portfolio. For the risk factors, I take the US dollar index and Carry trade *returns* as in Lustig, Roussanov, and Verdelhan (2011), together referred to as benchmark factors, and the two corresponding *variance swap payoffs*, together referred to as variance risk factors. I estimate the model in the two-step way as has become staple in asset pricing with non-traded factors: the loadings are estimated in the first step, and then used to pin down the risk premia. Although the above four factors are theoretically traded (the benchmark ones clearly so, while the variance risk ones are such by analogy with the stock market variance swaps and a no-arbitrage condition), Table 1.2 shows that the time-series average payoffs as the estimates of the risk premium are rather imprecise, especially at the one-month horizon. I will demonstrate however that both the sign and the order of magnitude of the variance risk premium obtained in the second step are the same as those of the sample averages.

Table 1.3 shows the estimates of the four prices of risk from linear regressions of 1- to 4-month "good and bad" carry portfolios' returns on the risk factors. The first model is a standard specification with the returns of the broad US dollar index and the baseline Carry trade strategy on the right-hand side. A repetition of an old result, expected returns of the test assets are aligned with the exposures to the Carry, but not to the Dollar factor. The second model describes returns as a function of the two variance risk factors, namely the US dollar and Carry trade variance swap payoffs. There is clear evidence that only the Carry trade variance risk is priced: the corresponding lambdas are significant, and

[14]Here, in test assets construction I use the whole G10 sample, including the Danish and Norwegian krone and the Swedish krona for the portfolio construction. The US dollar index and Carry variance swap payoffs are constructed without these currencies, as before.

[15]Clearly, many of these portfolios have returns that are either not linearly independent of return of the other portfolios, or almost perfectly correlated with any other series, a nuisance exacerbated by the small sample size. After ensuring that the matrix of test asset returns has full column rank, I am left with 32 portfolios with the average pairwise return correlation of 0.6.

[16]For details, see Munk (2013)

the price of risk estimates negative and of a similar magnitude as the sample values. The US dollar index variance risk exposures are not meaningfully aligned with expected returns, as the lambdas are economically minuscule and statistically indistinguishable from zero. In the third model, which embraces both the benchmark and variance risk factors, the price of the Carry trade variance risk becomes insignificant, dominated by the ever significant Carry trade return factor, but the numerical value is barely unchanged; the price of the US dollar index variance risk remains dwarfed by it. Noteworthy is the fact that the drop in significance is less pronounced for horizon of 3 months, where the variance risk premium is more precisely estimated in the first place.

1.4.3 Hedging shocks to future variance

Previously, Figure 1.4b presented evidence of a negative significant variance risk premium in FX portfolio returns all the way up to the 4-month horizon, and suggested that it has been historically costly for investors to hedge variance of FX trading strategies. In this subsection, I answer a subtler question: do investors only care about hedging the shortest-term (next month or next quarter) variance, or do they also price shocks to news about future variance? In doing so, I closely follow the methodology and notation of Dew-Becker et al. (2017) who conduct a similar analysis for the US equity market. I deviate from their study in that I do not rely on interpolation of the term structure of forward variances, but rather report the results for monthly variance at the horizon of up to 4 month and for quarterly variance at longer horizons. Additionally – as everywhere throughout my work – I do not rescale variance swap payoffs to arrive at variance swap returns, but rather draw inference from time-series of the former.

To visualize the concept of shocks to future variance, consider a currency investor with a three-period investment horizon. There are two types of variance shocks that she will be exposed to in period 1: the first type are shocks to the variance in that period only, while the second type are shocks to the *expectation* of variance in periods 2 and 3. As noted by Dew-Becker et al. (2017), many established long-run risk models in finance (e.g. Bansal and Yaron (2004), Drechsler and Yaron (2011)) and recent works on the macroeconomic consequences of shocks to news about future uncertainty (e.g. Bloom (2009), Fernández-Villaverde et al. (2011)) predict that investors are willing to hedge such news.

To test this implication, I calculate effective costs of hedging shocks to future τ-period variance, $\tau = 2, 3, 4$ months or quarters, whereby the hedging is performed by rolling over forward variance claims. Given definition of variance in eqs. (1.6) and (1.8) and of the discretized estimator thereof in eq. (1.7), it is straightforward to express the variance realized solely in period τ as:

$$RV(\tau, \tau) = RV(t, t + \tau) - RV(t, t + \tau - 1). \tag{1.24}$$

The price of this variance, or the period-τ forward variance claim price, follows trivially from the risk-neutral pricing argument:

$$F_t(\tau) = IV_t(t+\tau) - IV_t(t+\tau-1). \tag{1.25}$$

In other words, in period t it costs the investor $F_t(\tau)$ to buy forward the variance that will be realized in period $t+\tau$. Come period $t+1$, she can sell forward that variance (which would now be a contract expiring in $\tau - 1$ months), and realize a gain if there has been a positive shock to the expectation thereof. In order to maintain exposure to variance shocks at the τ-period horizon, she would buy forward the $t+\tau+1$ variance and then repeat the same steps, a succession familiar to futures market participants. A significantly large time-series average of payoffs from this succession would indicate that investors price shocks to variance at that particular horizon. By definition, period-1 forward variance is just the implied variance at that horizon, and the corresponding rollover payoff is simply the period-1 variance swap payoff.

Figure 1.6 shows the term structure of average forward monthly variance claims prices. For all portfolios, the curves are strictly upward sloping suggesting negative returns from the rollover strategy described above. Beyond the horizon of 1 month, the Carry trade exhibits a larger curve slope than the other strategies (the curves of those appear of about the same steepness), although at that horizon, the Value strategy and the Dollar index are the steepest, a repetition of the previously stated fact about the 1-month variance swap payoffs.

Figures 1.7a-1.7b depict the payoffs and respective Sharpe ratios of rolling over forward monthly variance claims at different horizons. All payoffs are negative, consistent with the upward-sloping term structure of forward variance prices; most are statistically significant, and Sharpe ratios are large for the US dollar index, Carry trade and Value strategy. More importantly, except for the US Dollar index, rolling over FX forward variance claims is as costly at the 4 months as it is at the 1 month horizon. Hence, it can be concluded that investors are willing to pay a non-negligible price to hedge shocks to news about variance at least within a popular rebalancing frequency[17].

Figures 1.8a-1.8b show the outcome of the same exercise conducted on quarterly forward variances at thresholds of 3, 6, 9 and 12 months: therein, $F(\tau)$ is the average payoff of rolling over forward claims on the variance realized over $(\tau - 3 < t < \tau)$, such that $F(3m)$ is again just the 3-month variance swap payoff. A somewhat different pattern emerges: shocks to the nearest and second nearest quarter variance are priced much stronger than shocks to more distant variance, as evidenced by declining payoffs and Sharpe ratios. For all portfolios, the fourth nearest quarter values are twice as low in magnitude as the second nearest quarter ones. However, most values are still statistically significant even

[17] For instance, used in Deutsche Bank FX indexes.

at the longest horizons, and Sharpe ratios are as high as the those of the strategy excess returns.

Overall, there is evidence that investors care about shocks to news about future variance of FX portfolios, and despite the fact that these shocks are "cheaper" the farther the hedged variance, investors are still ready to enter transactions of return profile comparable to trading strategies themselves to hedge event the shocks to news about the most distant variance away. This finding stands in contrast with what Dew-Becker et al. (2017) find for the US equity market: for a sample of US market variance swaps, the authors report that shocks to uncertainty at horizons larger than one quarter could have been hedged virtually for free.

1.5 Conclusions

FX portfolios, such as the US dollar index and Carry trade, are at the heart of FX market academic research and investment management, yet have been ignored by studies on FX variance risk premium. In this paper, I aim to redress the balance by using option-implied covariance matrices of currency returns and representing trading strategies in vector form to construct synthetic swaps on the variance of these strategies and study FX variance risk.

I find evidence of an overall negative FX variance risk premium, as the eigenvalues of the matrix difference between the implied and subsequently realized covariances are predominantly negative. The US dollar index and a typical carry trade strategy appear as the two portfolios with the most negative variance risk premium estimates of *all* possible portfolios, further stressing the own special place on the FX market. At the same time, strategies with a significantly positive ex post variance risk premium are achievable, up to the limits of predictability of the realized covariance.

Cross-sectional and time-series differences in conditional variance risk premium estimates can be explained by the past time-series information such as the spot return and realized variance: portfolios of negative momentum and high recently realized variance are on average the most expensive to insure against rising variance.

The Carry trade strategy, despite being the most volatile and crash-prone, does not stand out as the one with the highest negative variance risk price, superseded in that respect by the US Dollar index. Nevertheless, the Carry trade variance risk is an important factor for explaining the cross-section of FX returns, just as the Carry trade excess returns are a priced risk factor. The Carry trade variance risk is priced in the cross-section of "good and bad" carry portfolios, contributing to the resolution of their differential pricing, whereas the US dollar variance risk is not.

Shocks to news about future variance on the FX market are priced, the price being the smaller the more distant the expectation horizon. However, this decline is not as pronounced as for the equity market, where Sharpe ratios of strategies providing insurance against such shocks drop to zero beyond the horizon of 3 months. For FX portfolios, hedging shocks to news about future variance has been statistically and economically costly even for longer horizons of 6, 9 and 12 months.

Figure 1.4: Variance and variance risk of FX trading strategies

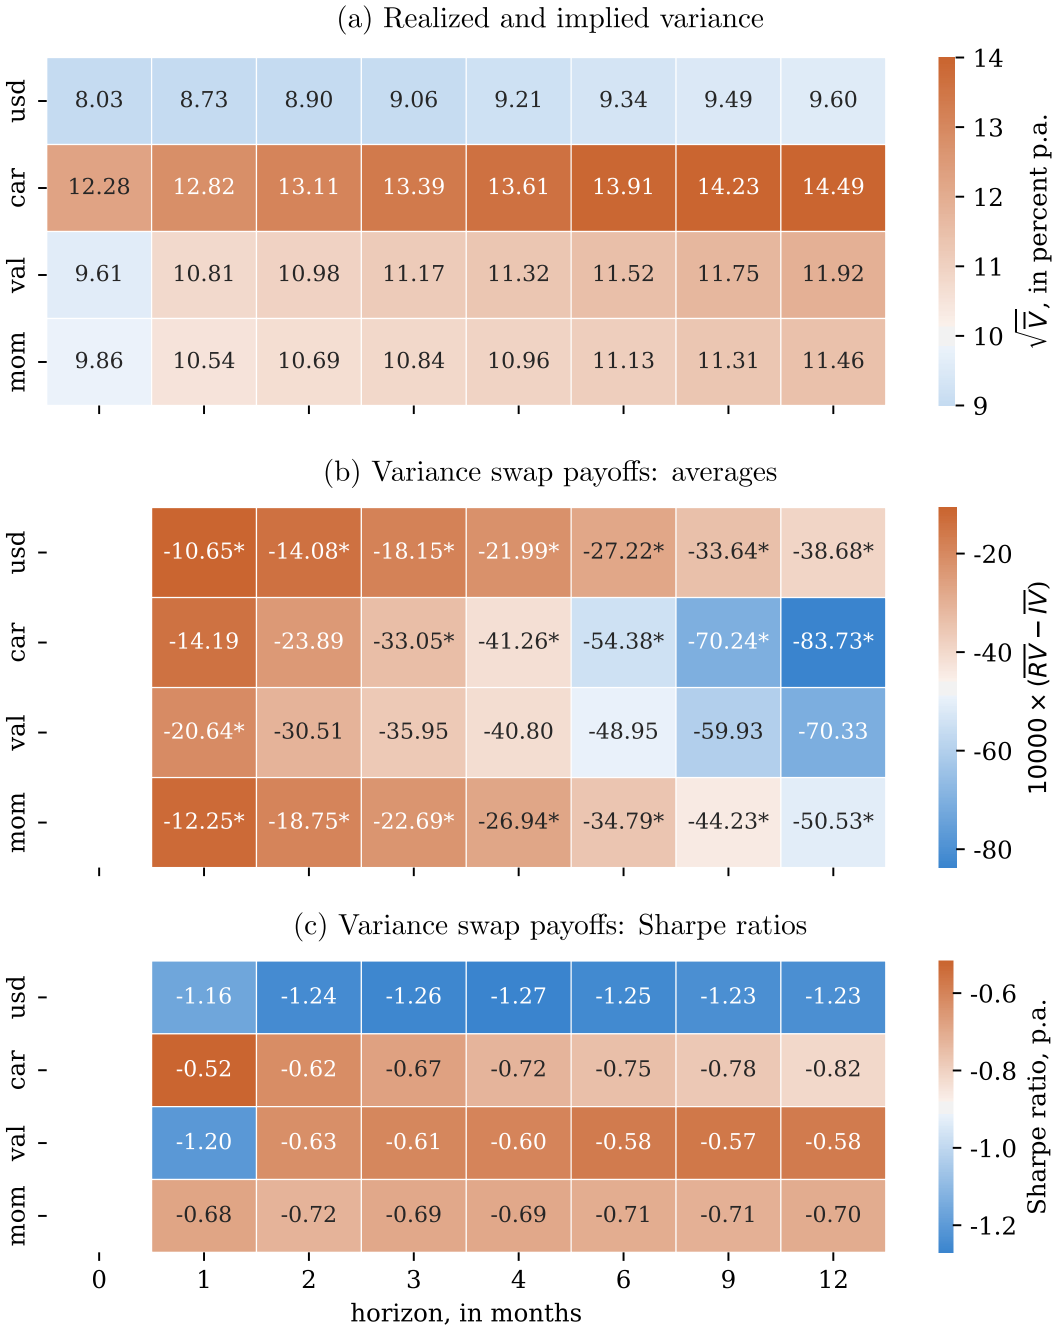

In this table, panel A shows the square root of the average realized (as horizon-0) and implied (for horizons of 1 to 12 months) variance of currency portfolios, in % p.a. The former are calculated as in eq. (1.7), the latter – as in eq. (1.9). Panel B shows the time-series averages of synthetic variance swap returns, calculated as in eq. (1.10), in squared percent p.a., whereby the estimates with a t-statistic exceeding 2.0 in absolute value are marked with an asterisk. Panel C shows the Sharpe ratios of the payoffs from panel B, annualized. Observations are daily and overlapping; t-statistics are calculated with the Newey and West (1987) adjustment with the number of lags set to 22×horizon. The colorbar to the right maps heatmap colors to numeric values. The sample period is from 01/2009 to 06/2018.

Table 1.2: Variance swap payoffs of for FX trading strategies: sample statistics

	usd	car	val	mom	vrp
1 month					
mean	−10.65	−14.19	−20.64	−12.25	−8.00
	(2.39)	(9.18)	(6.00)	(5.37)	(4.98)
sharpe	−1.16	−0.52	−1.20	−0.68	−0.43
std	31.79	95.13	59.37	62.07	65.16
median	−11.21	−18.82	−15.53	−11.82	−11.15
skewness	1.02	1.21	−2.14	1.05	3.11
2 months					
mean	−14.08	−23.89	−30.51	−18.75	−12.88
	(3.41)	(12.51)	(19.11)	(7.83)	(5.64)
sharpe	−1.24	−0.62	−0.63	−0.72	−0.53
std	27.70	94.35	119.13	63.58	59.13
median	−12.67	−24.04	−12.36	−12.84	−14.80
skewness	−0.23	0.42	−3.07	−1.33	1.93
3 months					
mean	−18.15	−33.05	−35.95	−22.69	−17.39
	(4.51)	(15.45)	(22.73)	(10.06)	(6.34)
sharpe	−1.26	−0.67	−0.61	−0.69	−0.63
std	28.83	98.05	118.72	65.56	54.95
median	−15.12	−28.03	−14.18	−15.63	−16.48
skewness	−0.49	−0.08	−2.94	−1.73	0.76
4 months					
mean	−21.99	−41.26	−40.80	−26.94	−22.87
	(5.51)	(17.65)	(25.48)	(11.99)	(7.06)
sharpe	−1.27	−0.72	−0.60	−0.69	−0.75
std	29.99	98.86	117.57	67.51	52.98
median	−18.05	−35.73	−17.43	−18.80	−18.70
skewness	−0.70	−0.38	−2.79	−1.81	−0.22

This table shows descriptive statistics of variance swap payoffs for a number of FX trading strategies and holding horizons τ (in months). A payoff is defined as in eq. (1.10), in percent-squared per year. Observations are daily and overlapping. Standard errors of the mean, calculated with the Newey and West (1987) adjustment with the number of lags set to $22 \times \tau$ are reported in parentheses below respective values. Sharpe ratios are annualized. Definitions of the strategies as in Section 1.2.6. The sample period is from 01/2009 to 06/2018.

Figure 1.5: Variance swap returns of currency and local stock market indexes

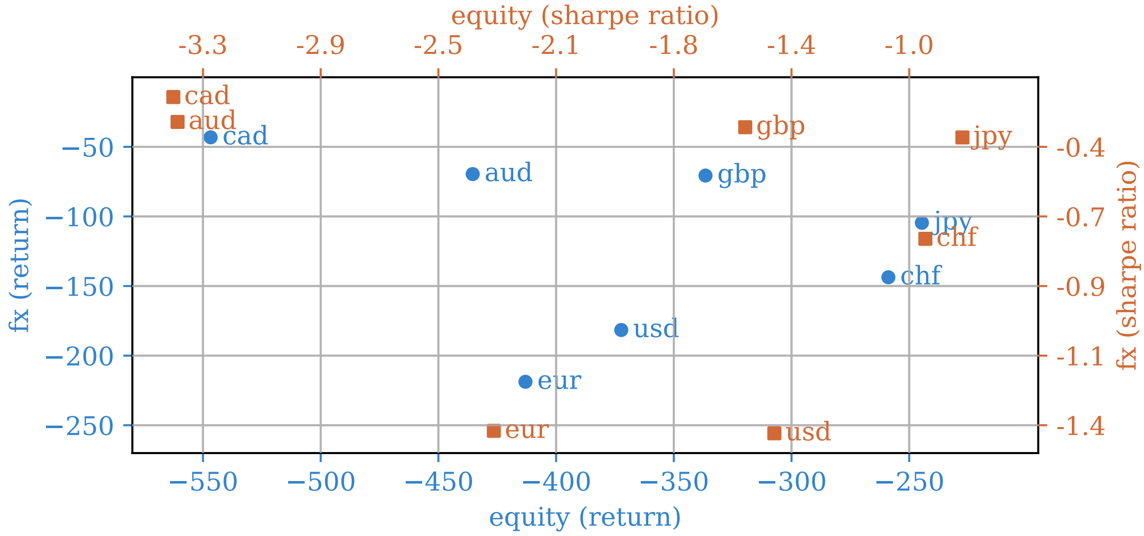

This figure depicts average 1-month variance swap returns (bottom x- and left y-axis, in blue) and respective Sharpe ratios (top x- and righ y-axis) for several local equity markets (on the x-axes) and corresponding currencies (on the y-axes). Returns are in percent, calculated as in eq. (1.22), whereby the FX variance swap rates are as in eq. (1.9), and the equity market swap rates are taken to be squared values of respective VIX indexes (details in Section 1.2.2. Both returns and Sharpe ratios are annualized. The sample period is from 01/2009 to 06/2018.

Figure 1.6: Forward variance claims prices: average values

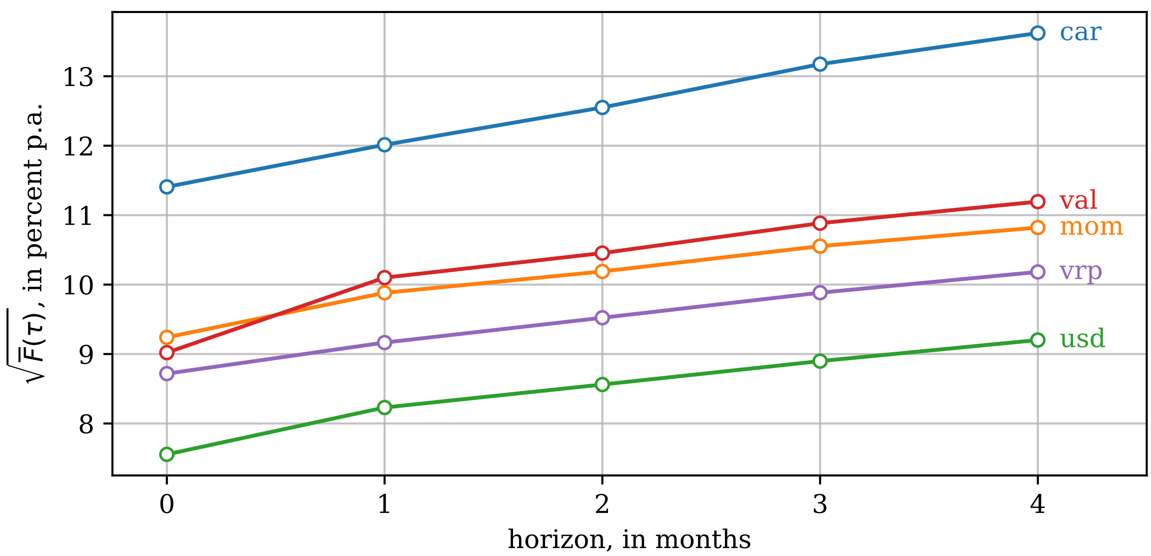

This figure shows the square root of the average τ-month forward variance prices for a number of currency portfolios (along the y-axis) and horizons (along the x-axis). Each value is thus a one-to-one increasing transformation of the average τ-month forward price of monthly variance. $F(0)$ is the average realized variance, and by definition, $F(1) = IV(1)$. Observations are daily and overlapping; one month is taken to be 22 days long. The sample period is from 01/2009 to 06/2018.

Table 1.3: Risk premium estimates in the cross-section of "good and bad" carry trades

	(1)	(2)	(3)	(1)	(2)	(3)
		(A) 1 month			(B) 2 months	
rx_{CAR}	7.33		7.37	6.66		6.79
	(2.31)		(2.37)	(2.22)		(2.33)
rx_{USD}	−2.03		−1.87	2.12		−0.49
	(−0.45)		(−0.38)	(0.37)		(−0.09)
vsp_{CAR}		−38.74	−36.00		−45.30	−39.09
		(−2.09)	(−0.78)		(−2.39)	(−1.29)
vsp_{USD}		9.44	5.82		−0.50	0.29
		(0.92)	(0.67)		(−0.04)	(0.03)
		(C) 3 months			(D) 4 months	
rx_{CAR}	6.30		6.43	6.07		6.06
	(2.14)		(2.35)	(2.10)		(2.23)
rx_{USD}	1.51		−1.28	1.40		−0.14
	(0.25)		(−0.19)	(0.23)		(−0.02)
vsp_{CAR}		−53.17	−48.29		−53.87	−37.89
		(−2.27)	(−1.58)		(−2.07)	(−1.15)
vsp_{USD}		−7.97	−6.21		−12.22	−7.97
		(−0.46)	(−0.46)		(−0.65)	(−0.55)

Each panel of this table shows risk premium estimates and respective t-statistics from a set of linear asset pricing model. The models differ in terms of explanatory variables: in column (1), these are the US dollar index and Carry trade returns, in column (2) – the corresponding variance swap payoffs, and in column (3) – the four variables jointly. On the left-hand side are 1-, 2-, 3- and 4- month excess returns of the "good and bad" carry trade portfolios as in Bekaert and Panayotov (2018), the horizon increasing from panel A to D. Standard errors are calculated with the Newey and West (1987) adjustment (the number of lags set to $\times 22\times$horizon). Returns are in percent p.a., variance swap payoffs are in currency per \$10000 notional. The sample period is from 01/2009 to 06/2018.

Figure 1.7: Rolling over forward variance claims: monthly variance

(a) Payoffs

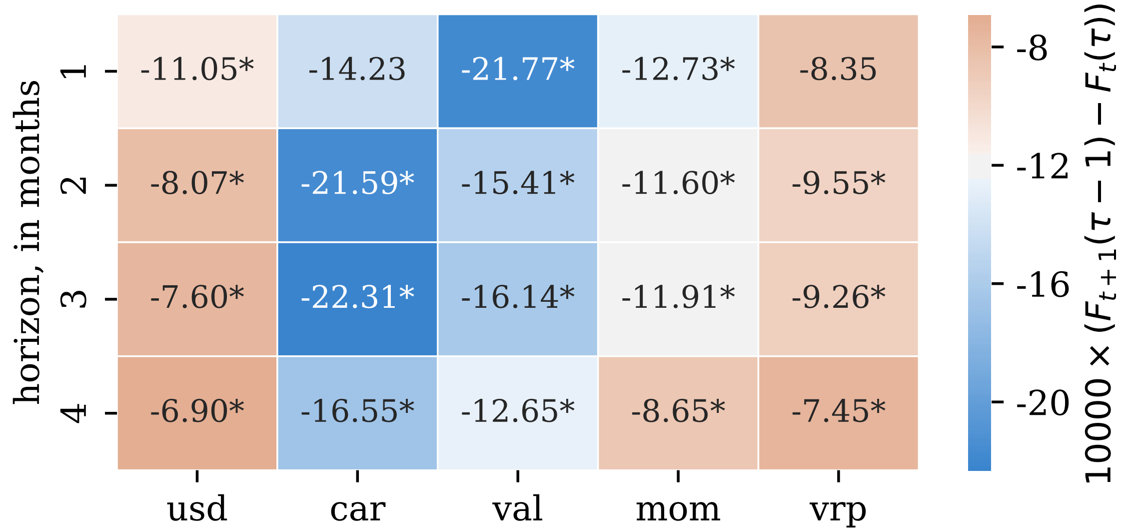

(b) Sharpe ratios

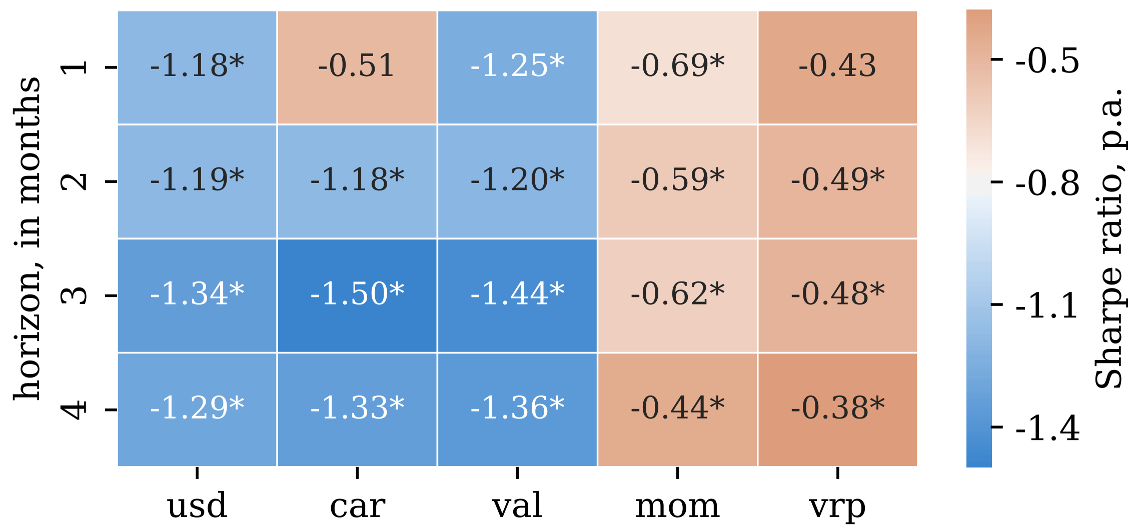

Panel A shows average monthly payoffs of rolling over forward monthly variance claims, and Panel B – respective Sharpe ratios. Horizon-1 claims are equivalent to 1-month variance swaps. Values exceeding in magnitude 2× own standard error are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment (the number of lags set to 2 × 22×months). The colorbar to the right maps heatmap colors to numeric values. Observations are daily and overlapping; one month is taken to be 22 days long. The sample period is from 01/2009 to 06/2018.

Figure 1.8: Rolling over forward variance claims: quarterly variance

(a) Payoffs

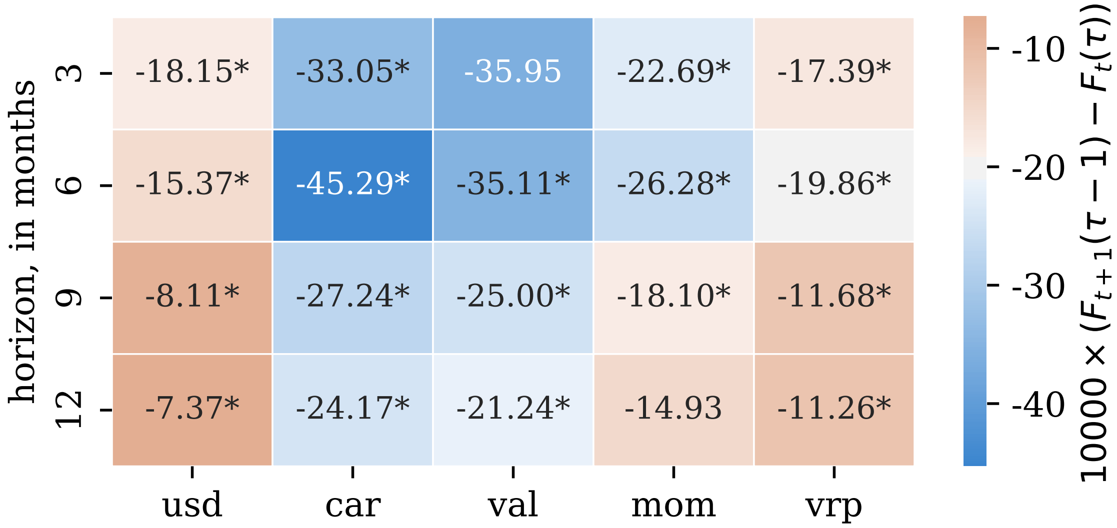

(b) Sharpe ratios

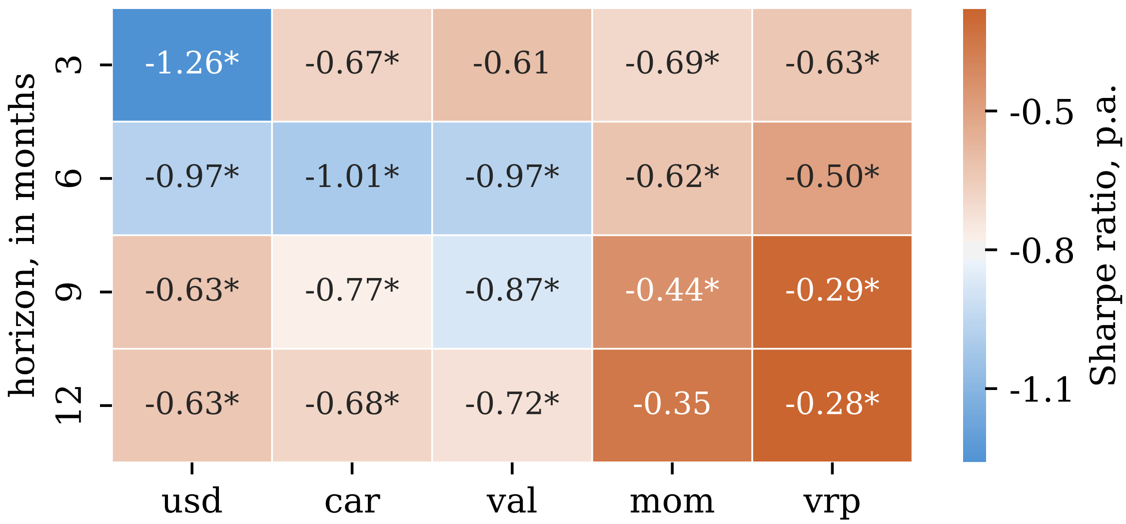

Panel A shows average quarterly payoffs of rolling over forward quarterly variance claims, and Panel B – respective Sharpe ratios. Horizon-3 claims are equivalent to 3-month variance swaps. Values exceeding in magnitude two own standard errors are marked with an asterisk (*), whereby the errors are calculated with the Newey and West (1987) adjustment (the number of lags set to $2 \times 22 \times$months). The colorbar to the right maps heatmap colors to numeric values. Observations are daily and overlapping; one quarter is taken to be 66 days long. The sample period is from 01/2009 to 06/2018.

Chapter 2

Monetary Policy and Currency Returns: the Foresight Saga*

Dmitry Borisenko† and Igor Pozdeev‡

Abstract

We document a conditional drift in currency returns before target rate announcements in developed economies. Currencies depreciate by 8 basis points per day over ten days before target rate cuts and appreciate by 5 basis points per day before rate hikes. Overnight index swaps allow to forecast monetary policy decisions accurately enough to make the drift exploitable by investors: our baseline specification of the trading strategy constructed by going long and short currencies before predicted local rate hikes and cuts earns on average 4.5% per year ($t > 2$, Sharpe ratio of 0.65), and would earn a percentage point more if the forecast quality were perfect. This return is robust to the choice of the trading strategy specification and cannot be attributed to exposure to established FX, equity and alternative risk premia. Our paper contributes to research on asset price dynamics around monetary policy decisions, the FX market anomaly literature, and to monetary policy predictability literature.

*We would like to thank (in alphabetical order) Michael Bauer, Anastasia Berezinskaya, Julien Huggonier, Marcial Messmer, Angelo Ranaldo, Gregor von Schweinitz, Paul Söderlind, Annette Vissing-Jørgensen and Jie Zhang, as well as participants of the 10th FIW Research Conference "International Economics", 30th Australasian Finance and Banking Conference, 2017 Auckland Finance Meeting and 2017 Paris Financial Management Conference for their valuable comments and insightful feedback.

†University of St. Gallen, *E-mail*: dmitry.borisenko@student.unisg.ch

‡Swiss Institute of Banking and Finance, University of St. Gallen, *E-mail*: igor.pozdeev@unisg.ch

2.1 Introduction

By 2001, central banks of most developed economies had adopted the practice of announcing monetary policy decisions at or immediately after regular pre-scheduled meetings. Recently, several studies have highlighted predictable patterns in the dynamics of asset prices around the meeting dates, challenging the existing asset pricing theories. In this paper we contribute to this literature by providing international evidence on profitably exploitable pre-announcement drift in currency returns, distinct from that documented previously for other asset classes.

In an event study framework, we find that over the 10-day period preceding an interest rate cut in an economy, its currency tends to depreciate against the US dollar by 8 basis points per day. The depreciation turns to an appreciation of 5 basis points per day before an interest rate hike, and virtually disappears before decisions to keep the target rate unchanged. These values are statistically different from zero at the 5% level and robust to exclusion of the recent Financial crisis period. Over the 5-day period, the numbers are approximately halved and still significant at the 5% level.

We support this result with a counter currency-agnostic panel regression and for the majority of currencies report depreciation and appreciation of a similar magnitude, but now against a generic counter currency. There is some heterogeneity in drift strength: for instance, the depreciation of the US dollar ahead of rate cuts initiated by the Federal Open Market Committee (FOMC) is estimated at 3-4 basis points per day over the 10-day period, whereas the Australian dollar tends to fall by 7 basis points per day before similar actions by the Reserve bank of Australia. In general, typical carry trade investment currencies exhibit larger magnitudes of the drift, at 5-11 basis points per day. The pre-hike appreciation, while evident for most currencies in our sample, is missing for the US dollar which weakly depreciates over the 10-day period irrespective of the policy decision type. With small differences, we find the same patterns over the 5- and 15-day periods.

Taking a look at the pre-announcement drift from a different angle, we show that it can be exploited by means of a trading strategy. From the perspective of a US investor, holding long (short) positions in currencies of economies about to witness an interest rate hike (cut) and closing them immediately before the announcement results in an average annualized return of 5% and a Sharpe ratio of 0.83 for the baseline strategy specification. Changing the account currency to the Japanese yen, thus adding the US policy events into the mix, does not meaningfully alter the significance of strategy returns.

Our findings would seem less surprising if target rate announcements were unpredictable. However, given the monetary policy transparency recently favored by central banks of developed economies, this is far from being the case. We show that expectations about policy rate changes can be extracted from prices of fixed income instruments such as overnight index swaps (OIS), allowing investors to forecast upcoming decisions and reap the FX pre-announcement returns. Although the quality of the OIS-implied forecasts that we report here arguably represents the lower bound on the quality of forecasts at investors' actual disposal because of the standard liquidity concerns and contamination of OIS rates with risk premia,[1] it is both high in terms of classification accuracy and consequences for construction of a pre-announcement trading strategy. When constructing the pre-announcement trading strategy with the OIS-based forecasts of rate cuts and hikes substituted for the previously assumed perfect foresight, we report a 3-5% annualized return depending on the strategy specification, with t-statistics of above 2 and the Sharpe ratios of 0.5-0.7.

Quantitatively comparable to other FX investment vehicles such as the carry trade, the pre-announcement trading strategy (either in its perfect-foresight or forecast-based form) thus poses a monetary policy-related asset pricing anomaly on the FX market. We leave a rigorous theoretical explanation of the documented anomaly for further research, but try to rationalize it with a possible exposure to commonplace risk premia and by studying the dynamics of spot returns in the aftermath of "neutral" policy meetings that have been incorrectly predicted to result in a cut or hike. We find that controlling for the equity factors of E. Fama and French (2014), FX carry trade, momentum and value portfolios or hedge fund literature staple trend-following factors of Fung and Hsieh (2015) does not reduce the strategy alpha below economically and statistically significant levels. At the same time, we document a positive relation between monetary policy shocks and currency spot returns after status quo retentions, as currencies appreciate after a positive shock (a cut predicted beforehand) and depreciate after a negative one (a hike predicted beforehand), resulting in losses to an investor who incorrectly predicted the rate change direction.

Summing up, we find that currency return dynamics exhibits a pattern ahead of monetary policy target rate changes and conditional on their direction, which – combined with strong predictability of the latter – allows to construct a profitable trading strategy uncorrelated with common reported risk factors. Our work is inspired by and contributes to the growing body of research on the dynamics of asset returns around monetary policy announcements. In this strand of literature, Lucca and Moench (2015) find strong positive returns of the S&P500 index in the 24-hour window before FOMC releases its policy decisions; Cieslak, Morse, and Vissing-Jorgensen (2019) document cyclicality of US stock returns anchored to FOMC meetings; Neuhierl and Weber (2018) mention a drift of US and international equities over days before and conditional on the direction of FOMC monetary *surprises*. Mueller, Tahbaz-Salehi, and Vedolin (2017) report significant positive returns of being short the US dollar in the hours around FOMC announcements,

[1] Additional sources of information that could improve the quality of forecasts are other derivatives' prices, analysts' surveys and event the regulators' own words: for example, Norges Bank adds monetary policy projections with an in-house view on the future policy rates in its quarterly reports.

finding 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. Distinguishing ourselves from their work, we use a longer pre-meeting window of several days and condition the drift *ex ante* rather than *ex post*; moreover, of all monetary policy announcements we are solely interested in *target rate* decisions, and are silent about unconventional policy tools.[2] Our research is most similar to Karnaukh (2016), who finds depreciation of the US dollar over days before FOMC rate cuts and appreciation before rate hikes, and goes on to show how investors can exploit this pattern in a trading strategy using Federal funds futures-implied forecasts. In contrast to that, we bring the currencies and policy announcements of regulators rather than FOMC into the picture and develop a more flexible methodology for detecting the conditional drift and measuring its investment relevance; we show that the US dollar joins most other currencies in exhibiting the conditional pre-announcement drift, but does not rank high among them in terms of its magnitude.

Our secondary contribution goes to the the literature on forecasting monetary policy path. While evidence on the forecasting power of the Federal funds futures for upcoming target rate changes is abundant – Krueger and Kuttner (1996) and Piazzesi and Swanson (2008) to name a few – our paper is to the best of our knowledge the first example of using overnight index swaps for the same purpose. We show that the expectations of target rate changes priced in into OIS rates have been rather accurate since early 2000s, allowing for high true positive and low false positive rates in the case of all considered economies.

The rest of the paper is organized as follows. Section 2.2 summarizes our dataset; Section 2.3 presents the pre-announcement pattern in currency spot returns through the lens of an event study and panel regression; Section 2.4 details a trading strategy that exploits the above pattern both when investors know upcoming target rate changes with perfect certainty or rely on market-implied forecasts; Section 2.5 features an attempt at a risk-based explanation of the anomaly; Section 2.6 concludes.

[2]The latter have recently enjoying elevated popularity among regulators, and indeed, the FX pattern that we document might to some extent be shaped by announcements of policy easing and tightening, which include but are not limited to target rate changes. For example, on January 22, 2015 the ECB announced its expanded asset purchase program with actual purchases starting in March 2015 simultaneously leaving the key rates unchanged. The euro depreciated against the US dollar by more than 2 percent over ten days preceding the announcement and by more than one percent on the announcement day. Although OIS correctly predicted no target rate change they completely missed the more important unconventional component of the announcement. At the same time, predicting the sentiment of a generic announcement is much more difficult and subject to data manipulation than predicting a rate hike or cut. With that being said, we prefer to concentrate on a subset of monetary policy announcements, arguably the most important and also forecastable in a systemic way.

2.2 Data

We collect data on policy rate announcements for the following economies: United States, United Kingdom, Australia, Canada, New Zealand, Switzerland, Sweden, Norway and the Eurozone. The sample thus covers G10 currencies extensively used in the FX literature, less Denmark and Japan for reasons to be discussed shortly. By November 2000, all of respective central banks adopted the practice of pre-scheduled policy meetings culminating in a target rate announcement. Our sample starts in August 2001, ending in December 2018 over which period all overnight index swap rates to be used to forecast announced rate changes are available.

In our analysis, we consider currency-events satisfying the following conditions: (i) the announcement schedule of the corresponding central bank should be published in advance and known to investors; (ii) an interest rate level or range should be an operational target of the central bank; (iii) explicit exchange rate targeting is not part of the monetary policy framework. Conditions (i, ii) emphasize our focus on announcements with outcomes that can be forecasted in a systematic way several days beforehand. Additionally, condition (i) rules out meetings occurring outside of the schedule: 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), the policy actions undertaken during such unscheduled meetings constitute a small fraction of all target rate changes.[3] Condition (iii) ensures we do not run into endogeneity problems and forces us to leave out Denmark: since 1999, the country's monetary policy objective has been to maintain a fixed exchange rate regime against the euro, such that Danmarks Nationalbank does not even have a schedule for policy decisions, therefore also violating condition (i);[4] for the same reason, we also exclude Switzerland for the period of the minimum exchange rate to the euro (from September 2011 to January 2015). We exclude Japan, since the Bank of Japan has been switching between different monetary policy tools over the past 20 years, and satisfied condition (ii) for a limited number of time intervals over the course

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

[4] When the ECB implements a change in interest rates, Danmarks Nationalbank typically responds with a similar change on the same date. More importantly, in order to maintain the currency peg, the bank resorts to unilateral rate changes in response to persistent changes in the euro-krone exchange rate, explicitly acknowledging it in the announcements. Furthermore, the unilateral rate changes are not scheduled and hence not known to investors in advance. See Spange and Toftdahl (2014) for a detailed overview of Denmark's fixed exchange rate policy. In unreported results we find that the krone's behavior around the ECB announcements is indistinguishable from that of the euro.

of our sample.[5,6] Table 2.1 reports the summary of such announcements (the data are from the websites of the regulators); below, we briefly discuss the policy rates for each economy.

Table 2.1: Monetary policy meetings summary

Country	Schedule	Target rate	Events	Hikes	Cuts	O/N rate	OIS data
Australia	1980s	Cash rate	191	19	21	AONIA	08/2001
Canada	11/2000	Overnight rate	139	23	20	CORRA	08/2001
Eurozone	01/1999	Average of marginal lending and deposit facilty rate	195	11	20	EONIA	01/1999
New Zealand	04/1999	Cash rate	136	23	19	NZIONA	09/2002
Norway	06/1999	Sight deposit rate	133	22	23	-	-
Sweden	10/1999	Repo rate	115	23	23	STIBOR	09/2002
Switzerland	01/2000	3-month CHF LIBOR	56	9	4	TOIS	08/2000
United Kingdom	06/1998	Bank rate	199	12	15	SONIA	12/2000
United States	02/1994	Federal funds rate	140	26	14	EFF	08/2001
Total			1304	168	159		

This table summarizes monetary policy announcements, target rates and OIS underlying rates across several economies. Column "Schedule" contains approximate dates of adoption of the pre-scheduled policy meetings and announcements practice; column "Target rate" mentions the one operational rate whose changes we treat as events; column "Events" shows the number of meetings since August 2001 until December 2018, and columns "Hikes" and "Cuts" – those meetings that resulted in a policy rate hike and cut respectively (here, September 2011 to January 2015 has been omitted for the Swiss franc as the period of exchange rate targeting); column "O/N rate" contains the underlying rate of the corresponding currency OIS; the last column shows the earliest point that OIS data is available.

Australia. The Reserve Bank of Australia (RBA) targets the overnight money market loan rate. From 1998 onwards, announcements fall 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 (BoC) introduced pre-scheduled interest rate announce-

[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 account balances at the Bank of Japan. Subsequently, the Bank resumed targeting the average call rate on March 9th 2006, switched to targeting a band on October 5th 2010, and abandoned interest rate targeting in favor of monetary base targeting on April 4th 2013. Finally, the Bank introduced negative interest rates on the current account balances 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.

[6] Over the course of our sample during the periods when the Bank of Japan had the overnight rate as its operational target, there were only two target rate hikes (in 2006 and 2007), and three rate cuts (one in 2001 and two in 2008). While these data points are unlikely to alter our main results, we nevertheless consider this discretionary choice as a limitation to our dataset.

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

Eurozone. From 1999 to 2001, the European Central Bank (ECB) used to hold monetary policy meetings twice a month, and from 2002 – once a month, before switching to the one meeting per six weeks 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 banking system; (iii) the main refinancing operations (or MRO) rate is the rate at the ECB injects liquidity using repo operations. We take the center of the deposit–lending rate corridor as the proxy for the one target rate. The Bank announces its interest rate decisions on the meeting day, and the changes 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 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 bank adopted the one meeting per six-week cycle. 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 (SNB) 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, and we exclude this period for the Swiss franc.[7]

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

[7]Note that exclusion of the Swiss franc for the period of the exchange rate cap does not have any impact on the the event study results – the only two target rate changes implemented within the cap period were announced at unscheduled meetings. Furthermore, to counter appreciation of the franc the Swiss National Bank conducted foreign exchange interventions outside the 2011–2015 period in 2009, 2010, and 2016. Excluding these years or even the entire currency does not affect our results: as we demonstrate below, the Swiss franc tends to weakly appreciate before rate cuts and depreciate before hikes.

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.

We use Bloomberg daily bid and ask spot exchange rates of the Australian, Canadian and New Zealand dollar, Japanese yen, Norwegian krone, Swedish krona, Swiss franc, British pound and Euro, all against the US dollar. When necessary, we construct cross-rates thereof relying on absence of triangular arbitrage. Tomorrow/next swap points used to roll positions over 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 event study and panel regressions in Section 2.3, we use quotes sampled at different fixing times to ensure that the announcement day is not overlapped for any of the currencies: 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.

To come up with forecasts of target rate changes, we use overnight index swap data, collecting 1-month OIS rates from Bloomberg and extending them with rates supplied by ICAP (provided by Datastream). The availability of the OIS data is described in 2.1 above.

In asset pricing tests we construct long-short tercile currency portfolios: value VAL sorted on purchasing power parity rates taken from the World Bank; momentum MOM, sorted on spot returns over the past 12 months without the very previous month; carry CAR sorted on 1-month forward discounts taken from Bloomberg; and equally-weighted dollar index DOL. E. Fama and French (2014) factors as well as Fung and Hsieh (2015) factors: the trend-following stock STK, interest rate IR, commodity COM, currency FX and bond BD portfolios, – are from the websites of the authors.

2.3 Spot exchange rates before MPAs

2.3.1 Event study

To highlight the pre-announcment drift of exchange rates, we use an event study framework,[8] taking exchange rates to be the test assets and monetary policy announcements of respective central banks to be the events, such that each test asset is associated with multiple events. Specifically, we study the dynamics of spot log-returns over the period of 10 days before an event, from day -10 to and including day -1. Within this window, "normal returns" are assumed to be have mean zero, akin to the constant mean model discussed i.a. in S. Brown and Warner (1980). Seemingly restrictive, the assumption is necessary as the true mean of currency returns cannot be precisely estimated on the sample of 17 years minus event windows that we have; moreover, even on longer datasets spot returns appear close to zero-mean and indistinguishable from such at the standard significance levels, as documented i.a. by Lustig, Roussanov, and Verdelhan (2011).

We define $d_{i,k}$ to be the date of announcement $k \in \{1, \ldots, K_i\}$ relating to currency $i \in \{1, \ldots, N\}$. As discussed above, the event window spans $w = 10$ days. We cut the series of spot returns of currency i into K_i subsamples of length w each and reindex the elements of the subsamples to have incremental ordinal indexes

$$\{t\} = \{-w, \ldots, -1\}.$$

A cumulative s-day abnormal return (CAR) of currency i before event k is defined as the sum of individual daily returns $R_{i,k,t}$:

$$R^{ca}_{i,k,s} = \sum_{t=s}^{-1} R_{i,k,t}. \tag{2.1}$$

Figure 2.1 depicts the average cumulative abnormal returns for each currency (in blue):

$$\overline{R^{ca}_{i,s}} = \frac{1}{K_i} \sum_{k=1}^{K} R^{ca}_{i,k,s}, \tag{2.2}$$

as well as the cross-sectional weighted average thereof (in red), the weights being proportional to the number of policy announcements for each currency. We bootstrap quantiles of cumulative normal returns by block-sampling with replacement from the panel of appreciation rates with all event windows, days of events and the first post-event days of each currency-event dropped. On each resampling, a new instance of the event study is

[8] Event studies in finance have not changed much since E. Fama, Fisher, et al. (1969); recent applications of the methodology to the FX market include Hayward (2018) and Kearns and Manners (2006), among others.

run, the average CAR estimated and recorded. After 1001 iterations we calculate the empirical quantiles of the distribution of estimated "normal" CARs.[9]

The top panel in Figure 2.1 shows that with the exception of the Swiss franc, all currencies depreciate against the US dollar ahead of local interest rate cuts, such that the cross-currency average depreciation is statistically and economically significant at about 8 basis points per day over ten days or half that over five days. With the same exception of the Swiss franc, currencies are found to appreciate before interest rate hikes, as evident from the middle panel: the average cross-sectional pre-announcement spot return is positive at about 5 basis points per day over ten days or 3.5 basis over five days. Results in the bottom panel indicate that ahead of "neutral" announcements no drift is present, as the cumulative returns are tightly gathered around zero. The common carry currencies: the Australian, New Zealand and Canadian dollar – are among the strongest depreciating and appreciating currencies ahead of rate cuts and hikes.

Confirming the findings of Mueller, Tahbaz-Salehi, and Vedolin (2017), we report comparatively large spot returns on the day of the announcement, visible in the Figure at day 0. On average, a 25 bps depreciation and 10 bps appreciation is detectable on the day of announced interest rate cuts and hikes respectively. In relation to our work, these indicate that although currencies anticipate policy rate decisions, the information does not get fully incorporated in the exchange rate in the pre-event period.

Our findings are not driven by the crisis period. Figure A.1 in Appendix A shows that the same conclusions can be drawn from a sample with the June 2007 – June 2009 period excluded. Nor are they driven by the choice of the US dollar as the common counter currency, as we will demonstrate in a counter currency-agnostic panel regression framework in the following section.

2.3.2 Panel regression

In order to complement the results of the event study and quantify the drift of *currencies* rather than currency *pairs*, we set up a panel regression framework, whereby the log-appreciation rate of currency a with respect to currency b is modeled as follows (all

[9] We also calculate simple confidence intervals as follows: first, we assume that the variance of the "normal" spot log-returns can be estimated as the variance between the end of the previous event window and start of the current one; second, we assume that returns are i.i.d. such that variance of a cumulative return is the sum of individual variances, and that of a cross-sectional average scales down as the square root of the number of currencies. This alternative methodology does not lead to meaningfully different conclusions.

variables are contemporaneous, so we omit the time subscripts):

$$\begin{aligned}\Delta s_{ab} = &+ \beta_a^{cut} d_{a,h}^{cut} + \beta_a^{hike} d_{a,h}^{hike} + \beta_a^{zero} d_{a,h}^{zero} \\ &- \beta_b^{cut} d_{b,h}^{cut} - \beta_b^{hike} d_{b,h}^{hike} - \beta_b^{zero} d_{b,h}^{zero} \\ &+ \gamma X_{ab},\end{aligned} \tag{2.3}$$

where $s_{ab} = \log S_{ab}$ is the log of the exchange rate of currency a against currency b; $d_{i,h}^{p}$ is the dummy variable equal to 1 if within h days from date t there is a currency i-specific policy rate decision of type $p = (cut, hike, zero)$ for rate hikes, cuts and status quo retentions respectively; X is a set of control regressors, in which we include the interest rate difference for the two currencies and the spot log-appreciation rate the day before. The counter currency-specific variables are given a negative sign since whichever factor makes b appreciate with respect to a, it trivially makes a depreciate with respect to b, the log-returns securing that Jensen's inequality does not pose a problem.

As a heuristic motivation for equation (2.3), consider a two-economy complete-markets representative agent framework implying that the log-appreciation rate of currency a with respect to currency b is linearly related to the stochastic discount factors (SDF) of agents whose consumption is denominated in each currency:

$$\Delta s_{ab,t} = m_{a,t} - m_{b,t}, \tag{2.4}$$

where m_i is the logarithm of the SDF of investors consuming in currency i. Additionally assuming that the log-SDF can be represented as a linear function of currency i-specific upcoming policy rate decisions and other control variables (e.g. the risk-free rate r^f and the previously realized log-SDF):

$$m_{i,t} = \beta_i^{cut} d_{i,h}^{cut} + \beta_i^{hike} d_{i,h}^{hike} + \beta_i^{zero} d_{i,h}^{zero} + r_i^f + m_{i,t-1}, \tag{2.5}$$

which in combination with 2.4 leads to equation (2.3).

An example of regressors for a =CAD, b=USD, $h = 10$ days is given in Table 2.2: therein, the variables are subdivided into the CAD-specific ones (corresponding to the top row of equation (2.3)), the USD-specific ones (the middle row) and the ones specific to the currency pair (the bottom row). There was a policy rate hike in Canada on September 6, 2017, and a meeting in the US on September 20, 2017 that did not result in a rate change: hence, $d_{CAD,10}^{hike}$ consists of 10 values "+1" before (and not including) September 6, and $-d_{USD,10}^{zero}$ of 10 values "-1" before September 20, the other dummies being zero.

Table 2.2: Example of regressors in equation (2.3) for a=CAD, b=USD, h=5.

		CAD			USD			CADUSD	
		cut	hike	zero	cut	hike	zero	fd_{t-1}	Δs_{t-1}
					…				
August	30, 2017	0	1	0	0	0	0	−0.04	−0.03
	31, 2017	0	1	0	0	0	0	−0.04	−0.87
September	01, 2017	0	1	0	0	0	0	−0.03	1.11
	04, 2017	0	1	0	0	0	0	−0.02	0.69
	05, 2017	0	1	0	0	0	0	−0.03	−0.18
	06, 2017	0	0	0	0	0	-1	−0.02	0.37
	07, 2017	0	0	0	0	0	-1	−0.01	1.20
	08, 2017	0	0	0	0	0	-1	0.00	0.92
	11, 2017	0	0	0	0	0	-1	0.00	−0.34
					…				

Equation (2.3) therefore can be written as:

$$Y_t = X_t B + \varepsilon_t, \qquad (2.6)$$
$$\text{s.t. } RB = 0,$$

where Y is the $N \times N$ block-diagonal matrix constructed of dependent variables (log-appreciation rates); X is the $N \times KN$ block-diagonal matrix constructed of transposed regressors corresponding to each dependent variable (K regressors per equation); B is the KN-vector of coefficients; R is the $M \times KN$ matrix of M restrictions imposed on the variables (such as β_i^p being the same in whichever equation it is present).

We estimate the model on a panel of seven non-redundant currency pairs[10] using seemingly unrelated regressions, which is a natural estimator for the purpose. Logically, coefficient β_i^p for a given (p, i) pair is constrained to stay unchanged across equations (e.g. β_{usd}^{hike} in the equation for AUDUSD is the same as in the equation for CADUSD). A positive estimated coefficient $\beta_{p,i}$ means currency i tends to appreciate before local policy decisions of type p.

In another specification of (2.3), we restrict the coefficients to only differ across policy decision types (i.e. hikes and cuts), but not across currencies: this way the Australian dollar is postulated to depreciate before Australian rate cuts by as much as the US dollar

[10] With eight currencies, all possible cross-exchange rates can be represented using seven currency pairs with the single counter currency.

does before the US rate cuts. The model looks as follows:

$$\begin{aligned}\Delta s_{ab} = &+ \beta^{cut} d^{cut}_{a,h} + \beta^{hike} d^{hike}_{a,h} + \beta^{zero} d^{zero}_{a,h} \\ &- \beta^{cut} d^{cut}_{b,h} - \beta^{hike} d^{hike}_{b,h} - \beta^{zero} d^{zero}_{b,h} \\ &+ \gamma X_{ab},\end{aligned} \tag{2.7}$$

where the only difference is the absence of coefficient subscripts. Table 2.3 presents the results of the estimation.

Consistent with the event study results, the estimated pre-announcement drift β_i^p is negative for most currencies for $p = cut$, implying an expected depreciation of currencies before policy decisions of this type. Unrestricted, the expected log-depreciation reaches 8-13 bps per day, being the most pronounced for the typical carry trade investment currencies: the Australian, Canadian and New Zealand dollar, – and the least pronounced (occasionally reversing to an appreciation) for the Swiss franc and British pound. The joint t-statistic of all pre-cut betas being different from zero is 1.90 or higher depending on the horizon h and hence, given the overall negative sign of the effect, is indicative of a significant pre-cut depreciation. Before policy rate hikes, currencies tend to appreciate with respect to others, the most responsive being the Norwegian krone, British pound and New Zealand dollar (appreciation of 2-4 bps per day). Albeit economically meaningful, the coefficients are jointly not significant, in part due to a lower number of policy decisions of this type. Before announcements of policy status quo retentions, the currency dynamics is the least pronounced, slightly negative and statistically indistinguishable from zero.

When the pre-announcement down- or upward drift is restricted to be equal in magnitude across currencies, we detect average depreciation before policy rate cuts of 4-6 bps and average appreciation before rate hikes of about 1-1.5 bp per day, as indicated by row “restr” in Table 2.3. That said, over a 10-day period ahead of a rate cut announcement by the local central bank, a random currency is expected to lose 40-60 bps with respect to another random currency; similarly, before a decision to raise rates, it is expected to gain 10-15 bps.

Again, we check if these results are robust to exclusion of the recent Financial crisis period from the estimation sample. Table A.1 shows this to be the case: the pre-cut dynamics is measured as negative, and the pre-hike dynamics – as positive for most currencies with familiar exceptions; the hypothesis of all individual pre-cut average returns being jointly equal to zero is strongly rejected; the restricted coefficients β^{cut} and β^{hike} estimated in the vicinity of -4 and $+1$ respectively.

Table 2.3: Currencies before local policy rate decisions.

	Panel A: h=5 cut	Panel A: h=5 hike	Panel A: h=5 zero	Panel B: h=10 cut	Panel B: h=10 hike	Panel B: h=10 zero	Panel C: h=15 cut	Panel C: h=15 hike	Panel C: h=15 zero
aud	-6.85 (-0.40)	1.42 (0.11)	0.42 (0.11)	-7.13 (-0.66)	2.11 (0.25)	-0.20 (-0.07)	-2.77 (-0.30)	0.83 (0.15)	-0.75 (-0.33)
cad	-12.51 (-0.94)	-0.06 (-0.01)	-2.01 (-0.45)	-10.95 (-1.13)	2.38 (0.42)	-0.09 (-0.03)	-9.57 (-1.40)	1.76 (0.38)	0.56 (0.23)
chf	2.82 (0.09)	-4.97 (-0.56)	1.45 (0.29)	0.89 (0.04)	-0.36 (-0.06)	-0.87 (-0.19)	-2.89 (-0.18)	1.22 (0.24)	-0.38 (-0.11)
eur	-1.55 (-0.12)	2.11 (0.22)	-1.61 (-0.34)	-2.77 (-0.27)	2.34 (0.42)	-1.23 (-0.43)	-1.96 (-0.23)	2.10 (0.46)	-1.16 (-0.52)
gbp	-0.70 (-0.03)	2.23 (0.31)	-1.68 (-0.51)	0.33 (0.02)	4.03 (0.79)	-0.87 (-0.37)	-3.89 (-0.40)	3.04 (0.65)	-1.26 (-0.55)
nok	-3.69 (-0.21)	3.72 (0.54)	-2.61 (-0.65)	-5.34 (-0.52)	4.31 (0.76)	-2.56 (-0.82)	-3.29 (-0.41)	2.88 (0.73)	-2.27 (-0.84)
nzd	-11.64 (-0.70)	3.05 (0.30)	-0.80 (-0.14)	-10.12 (-0.97)	1.97 (0.28)	-1.33 (-0.34)	-8.41 (-0.85)	2.57 (0.49)	-0.78 (-0.24)
sek	-4.98 (-0.40)	1.36 (0.17)	-0.26 (-0.05)	-7.80 (-0.93)	-1.67 (-0.29)	-1.94 (-0.53)	-4.76 (-0.68)	-0.77 (-0.15)	-0.94 (-0.31)
usd	0.31 (0.02)	1.01 (0.19)	-3.13 (-1.03)	-3.78 (-0.34)	-2.76 (-0.74)	-1.54 (-0.66)	-3.72 (-0.44)	-2.10 (-0.66)	-1.40 (-0.76)
joint	(1.90)	(2.33)	(0.86)	(3.69)	(1.42)	(2.30)	(3.85)	(1.59)	(1.84)
restr	-4.95 (-0.80)	1.49 (0.47)	-1.36 (-0.64)	-5.82 (-1.45)	1.24 (0.53)	-1.13 (-0.77)	-4.36 (-1.31)	1.12 (0.60)	-1.01 (-0.76)

This table shows the estimated β_i^p in (2.3) (rows "aud"–"usd") for different currencies i (across rows), policy rate decision types p (across columns) and lookback horizons h (across panels). The heatmap is a gradient from deep blue (negative values of the coefficient estimates) to orange (positive values) centered at zero. The coefficients can be interpreted as the average daily log-appreciation rate of currency i before policy rate decision of type p over h days before (and not including) the decision. The corresponding t-statistics, calculated using the Newey and West (1987) HAC covariance matrix estimator with automatic lag selection, are presented in parentheses. Row "joint" shows the joint t-statistic of testing $H_0 : \beta_{aud}^p = \ldots = \beta_{usd}^p = 0$. Row "restr" contains the estimated β^p from eq. (2.7), where the coefficients are assumed to be equal both across currencies. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD, SEK and USD for the period from August 2001 to December 2018.

2.4 Pre-announcement trading

We construct a simple trading strategy (denominated in the US dollar, that is, assuming a US investor's perspective), based on signals about upcoming policy rate changes. Let there be a meeting on date T; on date $T-h-1$ we establish a long (short) position in a currency if the signal is for the meeting to result in a rate hike (cut). The position is held for h days and reversed one day before the announcement at $T-1$. The position is rolled over daily as standard on ForEx. The payoff over h days realized at time $T-1$ is therefore:

$$R_{T-1}(h) = d_{T-h-1}(T)\left(S_{T-1} - S_{T-h-1} + \sum_{t=T-h-1}^{T-1} p_t\right), \tag{2.8}$$

where S is the spot exchange rate of the currency against the US dollar; d is the signal equal to $+1(-1)$ if a hike (cut) is signaled to cap off the upcoming policy rate meeting on date T; p is the tom/next swap points used for rolling the position over. We use bid and ask quotes both for spot rates and swap points as necessary. Whenever more than one position must be opened because pre-event windows overlap, we divide the available capital (calculated as the margin closeout value of the currency portfolio) equally, partially closing a position at the corresponding bid or ask spot quote. Starting with \$1 of capital, we rebalance our currency portfolio as time progresses, reporting on each day the unrealized profits and losses of the strategy.

2.4.1 Perfect foresight

First, we assume that investors are able to forecast policy rate changes with perfect accuracy, such that the signal d in equation (2.8) is the actual realized outcome of an upcoming policy meeting, and the strategy is solely parametrized by the holding horizon h. Just as in Section 2.3.2, we let h take the values of 5, 10 and 15 days. Results are depicted in Figure 2.2, where the adjacent table shows the performance of strategies with different holding horizons h.

Reinforcing the conclusions of the event study in Section 2.3, all strategies earn significant returns, exhibit positive skewness and high Sharpe ratio. For example, the 10-day strategy earns 5.2 percent per year ($t > 3$), has a skewness of 2.6 and a Sharpe ratio of 0.83. A large portion of the strategies' returns since 2001 appear to be earned during the period of the recent Financial crisis; however, as the bottom panel of Figure 2.2 makes clear, it is the abnormally high amount of monetary policy decisions around this time rather than abnormally high returns themselves that explain the spike in performance. The bottom panel of the Figure plots the cumulative return in event time (event-by-event rather than day-by-day), such that the positive return is distributed visibly smoother.

To check if these results are not driven by the US dollar being the counter currency, we re-run the trading strategy with the Japanese yen as the counter currency, which also allows us to include the US monetary policy events. The outcome of this exercise is depicted in Figure 2.3 below.

By following the strategy over the 17 years of our sample, a Japanese investor earns a 50-125% cumulative return on her unleveraged portfolio depending on the pre-announcement holding horizon. Annualized, this translates to values only slightly lower (by 0.5-1 percentage point) than what the US dollar-based strategy exhibits, and still highly statistically significant ($t > 2$) for the 10- and 15-day holding period, the Sharpe ratios exceeding 0.5.

2.4.2 OIS-based forecasts

Next, we argue that policy rate changes are forecastable with sufficiently good precision to make the trading strategy like the one above feasible. For this purpose, we use overnight index swaps (OIS) to come up with rather crude predictions of upcoming policy meeting outcomes and use these predictions as signals in equation (2.8).

OIS are fixed/floating interest rate swaps where the floating leg pays the cumulative underlying overnight rate, e.g. the effective federal funds rate in the US or SONIA in the UK. On settlement day T, the payoff of the floating leg of an OIS with \$1 notional amount quoted at t with start date at $t+1$ is as follows:[11]

$$\pi_{t+1,T} = \prod_{s=t+1}^{T} (1 + r_s) - 1, \tag{2.9}$$

where r_s is the underlying rate in fractions of 1 per day. On the same day T, the swap buyer pays 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 $\pi_{t+1,T} - w_t$. Clearly, as the instrument is forward-looking, its price must convey the expectation about future path of the underlying rate. Furthermore, if the underlying rate closely follows the policy target rate, the price also conveys information about the latter.

In absence of arbitrage opportunities, the price of the swap[12] at date t is equal to the

[11] This is a simplified version, as the actual payoff also depends on the day count and fixing lag, and reads:

$$\pi_{t+1,T} = \prod_{s=t+l_v}^{T} (1 + c_s r_{s+l_f}) - 1,$$

where c_t is the number of days that the rate is effective (e.g. 3 on Fridays), l_v is the value date lag, or the number of periods to "wait" until the first accrual of the underlying rate and l_f is the currency-specific fixing lag, or the lag with which the official rate is determined for the accrual date. In our calculations we honor these complications, but for the sake of brevity omit them in the description of methodology.

[12] 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.

risk-neutral expectation of (2.9):

$$w_t = E_t \pi_{t+1,T} = E_t \prod_{s=t+1}^{T} (1 + r_s) - 1, \tag{2.10}$$

where the expectation is taken under the risk-neutral measure. Let us assume a policy meeting takes place at date t^*, and the rate is possibly changed by a constant value p whose magnitude is known in advance, to $r + p_{t^*}$, effective as of $t^* + 1$. Let r_t be i.i.d with a known mean, such that:

$$\begin{aligned} w_t &= E_t \prod_{s=1}^{t^*} (1 + r_s) \prod_{s=t^*+1}^{T} (1 + r_s + p_{t^*}) - 1 \\ &= (1 + \bar{r})^{t^*} (1 + \bar{r} + p_{t^*})^{T-t^*} - 1, \end{aligned} \tag{2.11}$$

With these assumptions made, we arrive at the time-t expected rate change to happen at the announcement date:

$$\hat{p}_t(t^*) = \left((w_t + 1)(1 + \bar{r})^{-t^*}\right)^{\frac{1}{T-t^*}} - 1 - \bar{r}. \tag{2.12}$$

In empirical application, we proxy $\bar{r}$ as the 10-day rolling median of the underlying rate.

We adopt the following rule to forecast the direction of upcoming policy rate decisions: if the estimate of p is high enough in magnitude, we forecast a hike if it is above zero and a cut if it is below zero; if it is too close to zero, we forecast no change in the policy rate:

$$d_t(\tau) = \begin{cases} 1, & \text{if } \hat{p}_t(t^*) > \tau, \text{ rate hike expected} \\ 0, & \text{if } |\hat{p}_t(t^*)| \leq \tau, \text{ no change expected} \\ -1, & \text{if } \hat{p}_t(t^*) < -\tau, \text{ rate cut expected,} \end{cases}$$

where d is the signal in equation (2.8), and τ is the threshold determining how much "too close to zero" is. Now, the strategy is parametrized by holding period h as before, but also by the threshold that we use to discriminate between hikes, cuts and status quo retentions.

Clearly, OIS-based forecasts are but crude estimates of investors' beliefs about the path of monetary policy. First, OIS underlying rates are not always the same as respective policy target rates, although we equate the two. For instance, as mentioned previously, the Swiss National Bank targets the 3-month LIBOR whereas the OIS are priced based on the expectation of the overnight rate. Second, since OIS underlying rates are unsecured loan rates, liquidity issues might start to play a role whenever the unsecured market gets frozen.[13] Third, the risk-neutral expectation differs from the objective expectation by

[13] For example, De Fiore, Hoerova, and Uhlig (2018) report that the share of the unsecured segment

the amount of the risk premium, although in unreported calculations we find little reliable difference between the *ex post* cumulative overnight rate and OIS rates, a reassuring evidence of negligible risk premium. Finally, the threshold-based rule makes little sense when the threshold is the same for all currencies, as different central banks have executed hikes and cuts of different magnitude, some being more conservative than others. However, flawed as they are, OIS-implied forecasts are a handy tool to predict policy rate decisions in a systemic way. Appendix B quantifies this claim: therein, we report a 0.7-0.9 probability that before a policy rate change the exceedance (in the absolute value) of the OIS-implied rate over the recent overnight rate is larger than before a policy status quo retention, and thus the implied rate can be successfully used as a classifier; we conclude that despite obvious room for improvement, OIS rates are well suited as source of signals for our trading strategy, albeit the profitability of a trading strategy based thereon would constitute a lower bound on what investors could have really achieved in terms of return.

Several days in advance, we use the forecasts to classify the outcome of the approaching policy rate meeting as a hike, cut or a status quo retention, and open long and/or short positions accordingly. Figure 2.4 displays the results for strategies constructed using several different thresholds. At the shortest holding period, the strategies perform poorly, which is not surprising given that there are simply nor enough days to meaningfully exploit the drift even in the perfect-forecast case; however, returns are economically and statistically significant when positions are kept open for 10 and 15 days ahead of policy meetings. On average, the longer holding period strategies produce between 4% and 6% net return per year, about a percentage point lower than what the perfect-foresight strategies do, whereas Sharpe ratios are about 0.2 lower at 0.5-0.6.

A repetition of the familiar exercise, in Figure 2.5 we present several forecast-based strategies (using the same holding period and threshold values as in Figure 2.4) with the Japanese yen as the counter currency. Just as is the case for the US dollar-denominated strategy, performance is robust for longer holding horizons, with the average returns being statistically and economically significant at 4-5% with t-statistics of mostly above 2.0. Interestingly – and rather reassuringly – the strategy returns do not deteriorate after 2012 compared to the strategies in Figure 2.4, keeping a steady upward slope and suggesting a US dollar-specific explanation of the profitability drop therein.

Overall, the forecast-based strategies appear less profitable than the perfect-foresight ones, but produce statistically and economically significant returns, even with the rather crude target rate change forecasting model. Any improvement of the latter is bound to boost the strategies' performance and further strengthen our results.

of the total quarterly interbank money market turnover in the euro area fell from 45% in 2003 to 5% in 2017.

Table 2.4: Explaining strategy performance.

Panel A: Currency strategies							
	α		VAL	MOM	DOL	CARRY	adj. R^2
baseline	5.67		−0.08	−0.00	−0.02	−0.08	0.00
	(2.59)		(−1.09)	(−0.02)	(−0.19)	(−1.05)	
perfect foresight	6.43		−0.08	0.12	−0.09	−0.13	0.08
	(3.43)		(−1.11)	(2.15)	(−1.01)	(−1.89)	
Panel B: E. Fama and French (2014) factors							
		CMA	RMW	HML	SMB	MKT	
baseline	4.96	0.16	0.06	−0.09	0.09	−0.09	0.04
	(2.43)	(1.48)	(0.66)	(−1.14)	(1.78)	(−1.14)	
perfect foresight	6.00	0.14	−0.04	−0.11	0.08	−0.11	0.06
	(3.19)	(1.59)	(−0.59)	(−1.49)	(1.87)	(−1.45)	
Panel C: Fung and Hsieh (2015) factors							
		STK	IR	COM	FX	BD	
baseline	5.47	0.00	0.02	0.02	0.02	−0.02	0.15
	(3.09)	(0.12)	(1.48)	(2.79)	(2.36)	(−1.31)	
perfect foresight	5.71	−0.00	0.02	0.02	0.02	−0.02	0.19
	(3.50)	(−0.06)	(1.42)	(3.04)	(2.37)	(−1.92)	

This table presents the output of regressing returns of the baseline (10-day holding period, 10 basis points classificaiton threshold) and perfect foresight pre-announcement trading strategies (10-day holding period) onto three sets of factors. In panel A, the factors are the dynamically rebalanced FX portfolios: value, momentum, dollar index and carry trade, – constructed on the same set of USD-denominated exchange rates as the response portfolios; in panel B, E. Fama and French (2014) stock factors are used; in panel C, Fung and Hsieh (2015) trend-following factors are used. The first (leftmost) column in each panel keeps the intercept (in percent per year), and the rightmost – the adjusted R^2 values, in fractions of one. The corresponding t-statistics of coefficients, calculated using the Newey and West (1987) HAC covariance matrix estimator with 2 lags are given in parentheses. Returns are at the monthly frequency and non-overlapping. The currencies used to construct FX trading strategies are AUD, CAD, CHF, EUR, GBP, NZD and SEK. Sample size from August 2001 to December 2018.

2.5 Strategy up close

2.5.1 Time series asset prising tests

In attempt to explain the strategy performance through the prism of risk-return tradeoff, we regress its returns onto several established risk factors at the monthly frequency. To be generic, we use both the perfect foresight and the baseline version of the strategy, each with the holding period of 10 pre-announcement days. The baseline strategy is constructed with the threshold of 10 basis points to discriminate between upcoming policy rate cuts, hikes and status quo retentions. We use three sets of regressors: first, staple FX market portfolios; second, the E. Fama and French (2014) stock market factors; third, the Fung and Hsieh (2015) lookback straddle-based trend-following factors. Results are presented in Table 2.4.

First, there is little difference between results for the perfect foresight and the baseline version of the pre-announcement FX trading strategy, so we just refer to either version of it as “the strategy”. Second, the strategy does not load meaningfully on either the stock market factors of E. Fama and French (2014) or currency strategies such as the carry trade, as indicated by the statistically insignificant betas in panels A and B. Third, the strategy resembles a hedge fund return profile, justified by its significant loadings on the trend-following commodity and FX factors in panel C, as well as by the high R^2 values of the corresponding regression. Finally and most importantly, neither set of regressors is able to span the pre-announcement trading strategy returns, as regression alphas are statistically and economically significant across panels, estimated at about 5% per year with t-statistics of 2.40 and higher. We conclude that the strategy, even when based on the rather crude policy rate forecasts, presents an FX market anomaly.

2.5.2 Currency returns and monetary policy shocks

As a quick look at another potential explanation of the pre-announcement trading robust performance, let us consider what happens to currencies on the days of policy meetings that do not result in a rate change. If a subset of market participants need to hold currency positions beyond the policy meeting event day, they face a risk of negative returns whenever mistaken about the policy rate change direction, and hence would require higher average returns on their positions in the pre-announcement period. Figure 2.6 shows the 1-day appreciation rates on such days, conditional on the average value of OIS-implied policy rate change forecasts $\hat{p}$ over the 10-day period beforehand.

The positive monetary shocks that occur when a rate cut of magnitude larger than 10 basis points is forecasted but status quo is eventually retained, are associated with spot appreciation of all considered currencies except the New Zealand dollar, for which no such cases have been recorded; the cross-currency depreciation (in cell “x-avg” to the bottom right) is estimated at 25 basis points, which is statistically significant at the 5% level. In the opposite case a negative monetary policy shock occurs when a rate hike of magnitude larger than 10 basis points is forecasted before a “neutral” meeting, and currencies react by depreciating on the meeting day by about 10 basis points, although the estimate is not significant. Altogether, this means that whenever long (short) positions are “erroneously” opened because $\hat{p}$ is largely positive (negative), currencies move against the investor on the meeting day.

2.6 Conclusion

We describe a persistent pattern in the dynamics of currency returns ahead of policy rate announcements in respective economies: currencies start to moderately appreciate days before declared interest rate hikes and significantly depreciate before rate cuts, showing little recognizable dynamics otherwise. Investors can profit from this knowledge by opening long and short positions in currencies ahead of local target rate hikes and cuts respectively: indeed, we report staggering profitability of a pre-announcement trading strategy for an investor with perfect foresight about the direction of target rate changes. Curiously, when the perfect foresight is substituted with market-implied forecasts extracted from the prices of fixed income derivatives, the strategy does not appear dramatically less profitable, retaining the statistical and economic significance of returns and high Sharpe ratios. Since the market-implied forecasts are arguably inferior in quality to those at the actual disposal of the FX market participants, the pre-announcement trading has been a foresight saga so far. We cannot explain these anomalous returns with exposure to established FX, stock, bond and commodity risk factors; as of completion of our paper, they pose an asset pricing anomaly and a challenge to the existing and new asset pricing theories.

Figure 2.1: Exchange rates around local policy rate decisions.

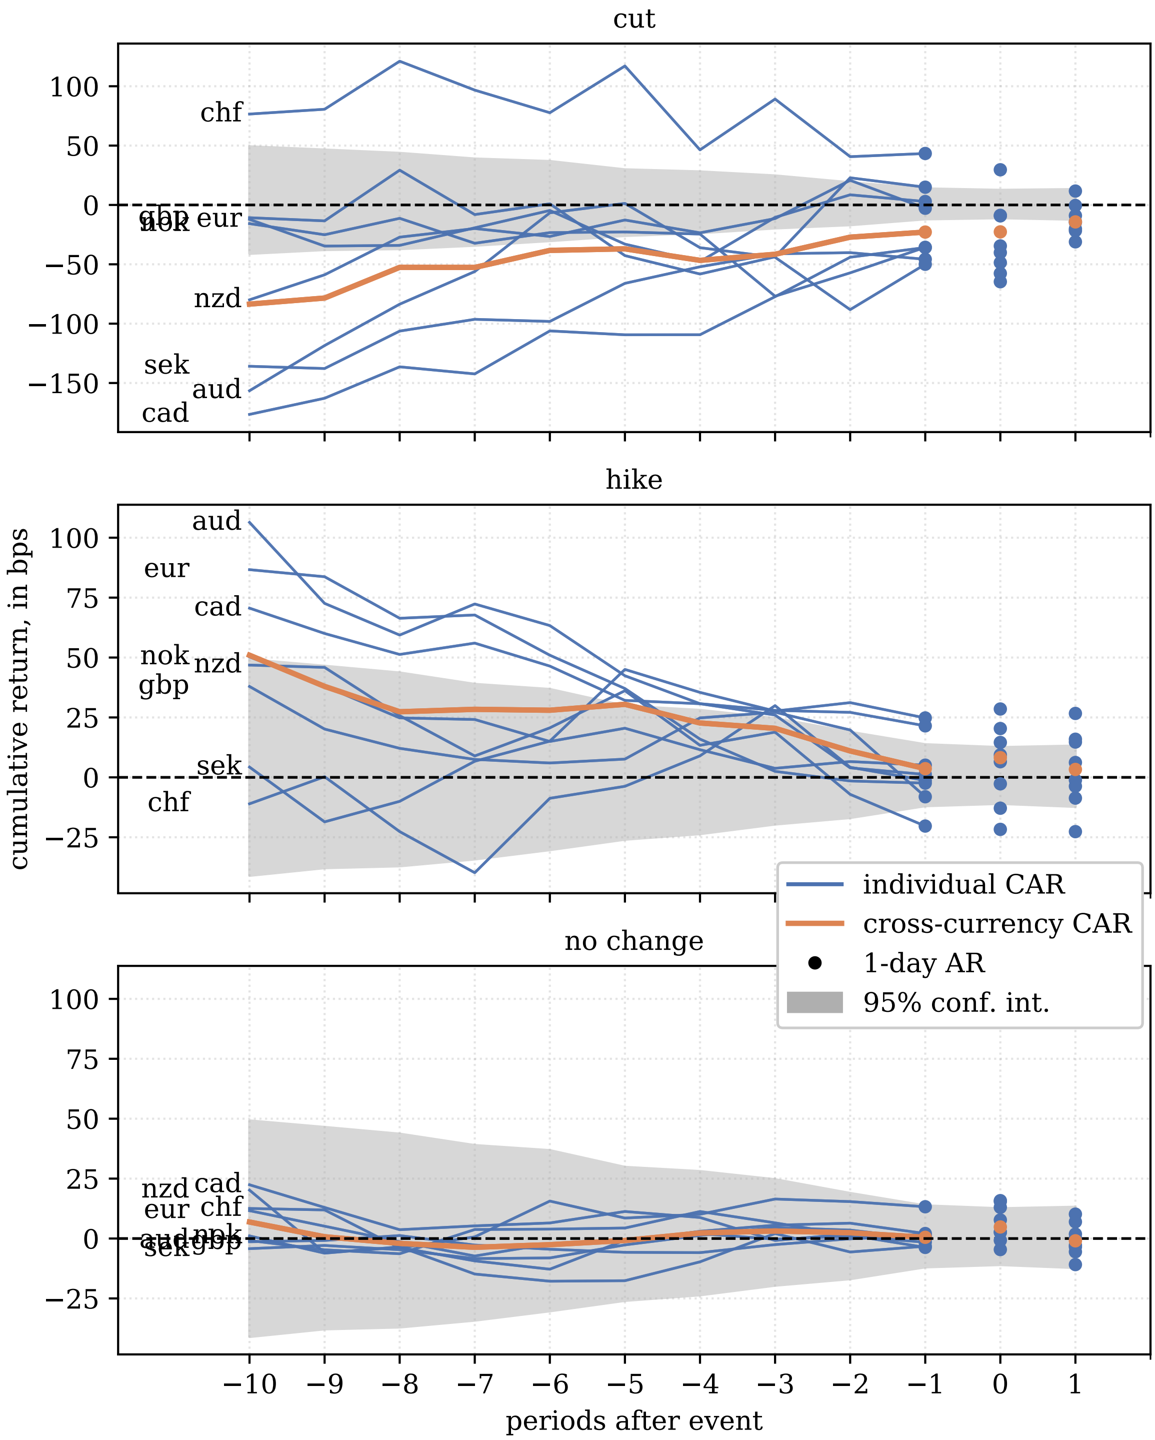

This figure depicts the cumulative appreciation $\sum_{t=x}^{-1} \Delta s_{i,t}$ of currencies against the US Dollar before local policy rate decisions, as well as the average appreciation on the day of and after the latter. At each point x, the Figure shows by how much a currency appreciates over x days preceding policy rate cuts, hikes and no-changes (in panel A, B and C respectively); the thick orange line in each panel shows the weighted average cumulative return across all currencies, with weights proportional to the number of currency-specific events. The day of announcement corresponds to $x = 0$. The shaded area represents the 95% bootstrapped confidence interval for the cross-currency average value (the orange line) around zero. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from August 2001 to December 2018.

Figure 2.2: Pre-announcement trading (perfect foresight).

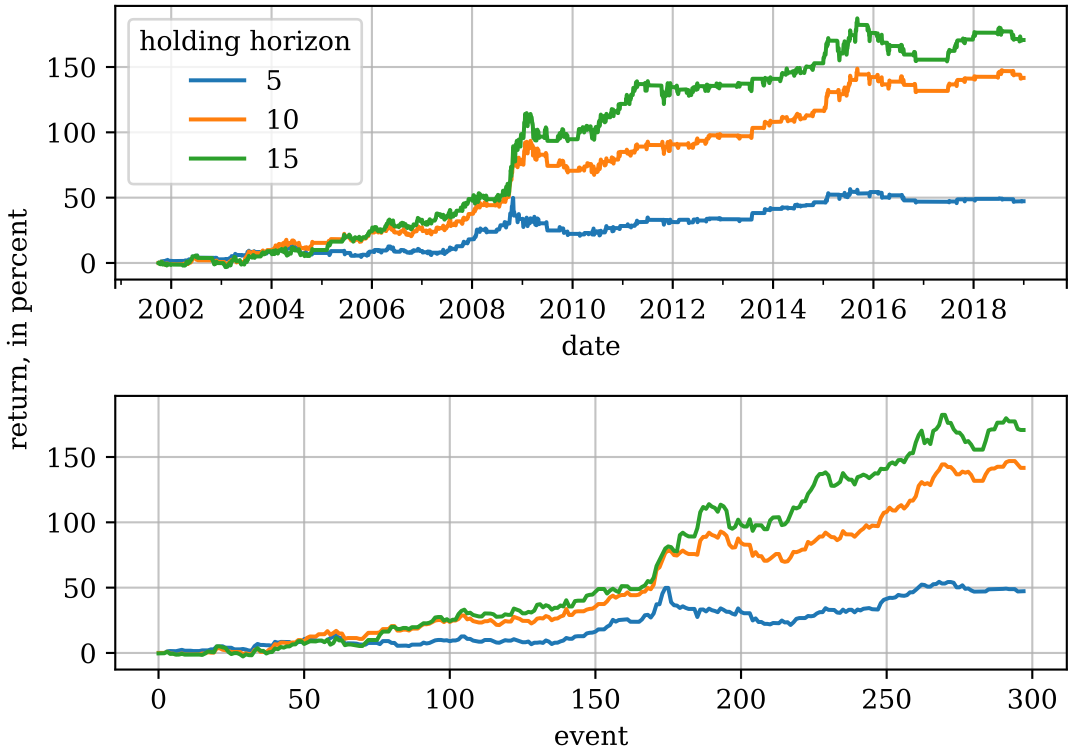

	holding period h		
	5	10	15
mean	2.26	5.19	6.01
t-stat	(2.27)	(3.36)	(3.24)
standard deviation	4.30	6.28	7.07
skewness	0.22	2.59	1.74
Sharpe ratio	0.53	0.83	0.85

This figure depicts the cumulative returns of a trading strategy that goes long (short) currencies against the US dollar $n + 1$ days before local interest rate hikes (cuts) and reverses the position on the pre-announcement day, for holding horizons $n = 5, 10$ and 15 days. The positions are rolled over daily using tom/next swap points. Long (short) trades and rollovers on corresponding positions are initiated at ask (bid) quotes. The upper panel of the figure shows returns plotted in calendar, and the lower panel – those in event time. The table below the figure shows descriptive statistics of annualized returns of the strategies. The standard error of the mean is calculated based on monthly returns and is Newey and West (1987) HAC with 2 lags. The sample includes AUD, CAD, CHF, EUR, GBP, NZD, and SEK; USD is the common counter currency. Sample size is August 2001 to December 2018.

Figure 2.3: Pre-announcement trading (perfect foresight) in JPY.

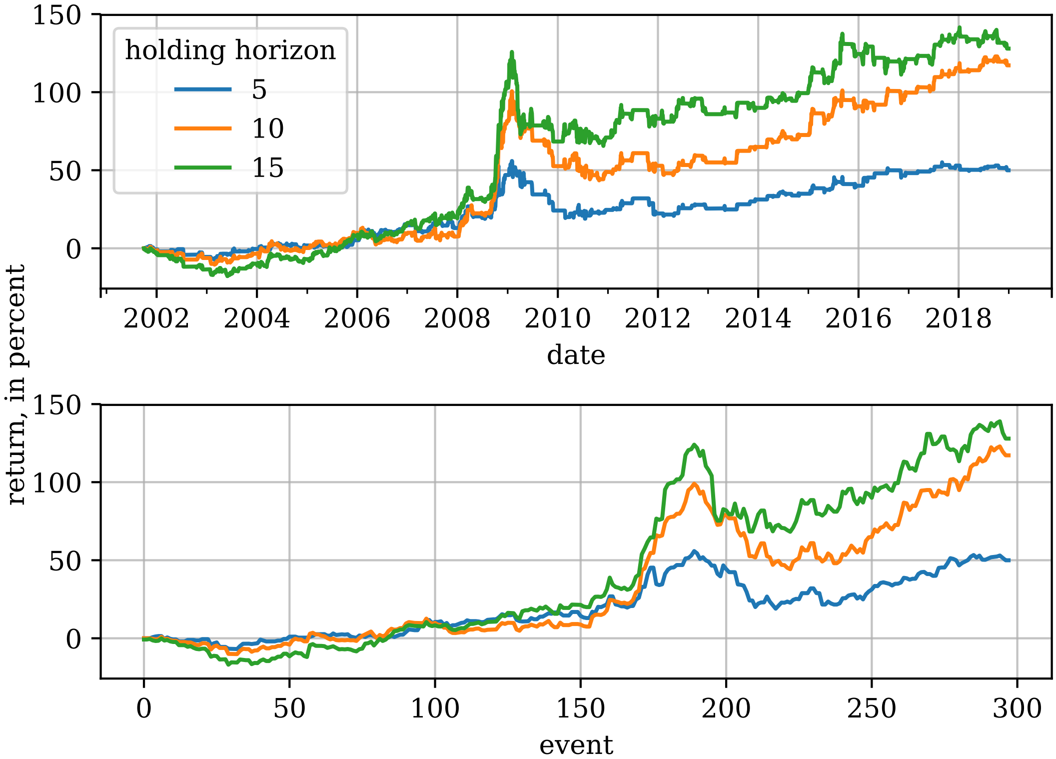

	holding period h		
	5	10	15
mean	2.62	4.47	5.85
t-stat	(1.72)	(2.08)	(2.32)
standard deviation	6.59	8.47	9.34
skewness	0.92	1.59	1.13
Sharpe ratio	0.40	0.53	0.63

This figure depicts the cumulative returns of a trading strategy that goes long (short) currencies against the Japaneses yen $n+1$ days before local interest rate hikes (cuts) and reverses the position on the pre-announcement day, for holding horizons $n = 5, 10$ and 15 days. The positions are rolled over daily with tom/next swaps. Long (short) trades and rollovers on corresponding positions are initiated at ask (bid) quotes. The upper panel of the figure shows returns plotted in calendar and the lower panel those in event time. The table shows descriptive statistics of returns of the strategies, in percent annualized. The standard error of the mean is calculated based on monthly returns and is Newey and West (1987) HAC with 2 lags. The sample includes AUD, CAD, CHF, EUR, GBP, NZD, SEK and USD; JPY is the common counter currency. Sample size is August 2001 to December 2018.

Figure 2.4: Pre-announcement trading (forecast-based).

(a)

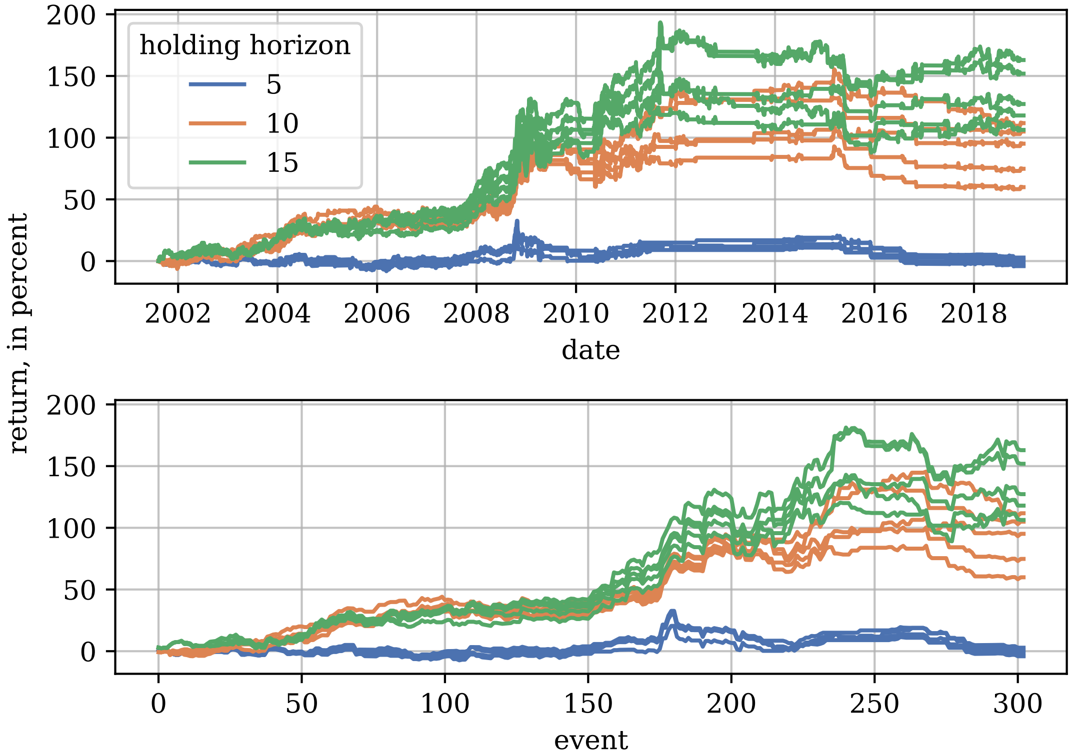

(b)

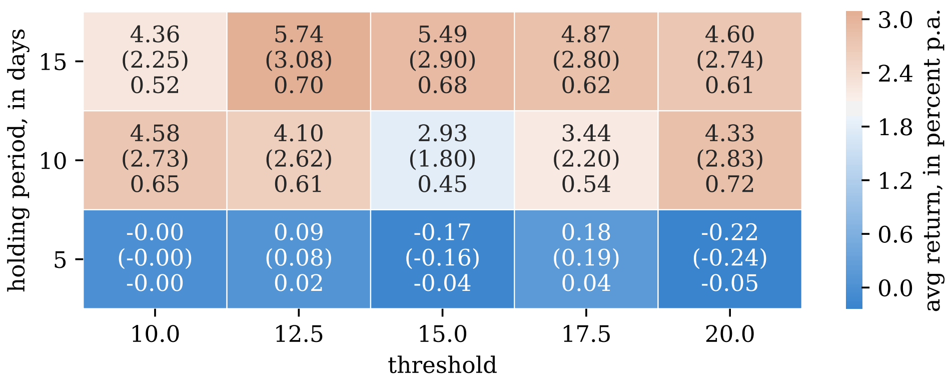

This figure depicts the cumulative returns of trading strategies that go long (short) currencies against the US dollar $n+1$ days before *predicted* local interest rate hikes (cuts) and reverses the position on the pre-announcement day, for holding horizons $n = 5, 10$ and 15 days and several thresholds, in basis points). The positions are rolled over daily with tom/next swaps. Long (short) trades and rollovers on corresponding positions are initiated at ask (bid) quotes. The upper panel of the figure shows returns plotted in calendar, the middle panel – those in event time. The heatmap in the lower panel shows the means (in percent per year), t-statistics thereof and Sharpe ratios (annualized) of the strategies' returns. The standard error of the mean is calculated based on monthly returns and is Newey and West (1987) HAC with 2 lags. The sample includes AUD, CAD, CHF, EUR, GBP, NZD, SEK; USD is the common counter currency. Sample size is August 2001 to December 2018.

Figure 2.5: Pre-announcement trading (forecast-based) in JPY.

(a)

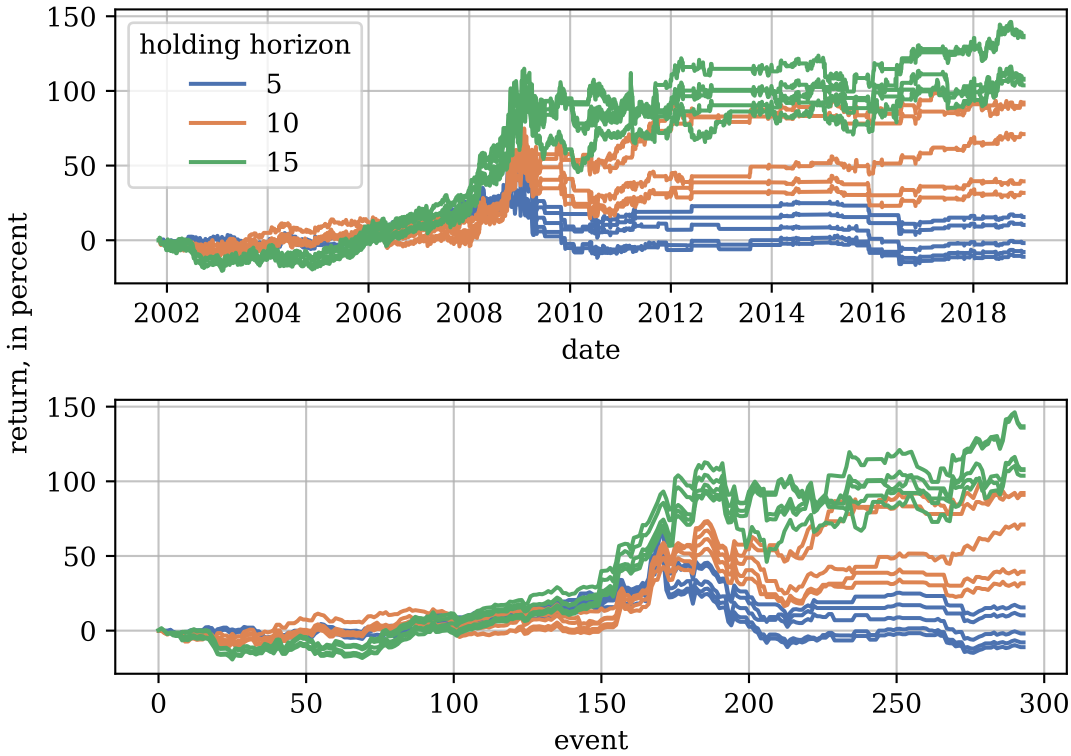

(b)

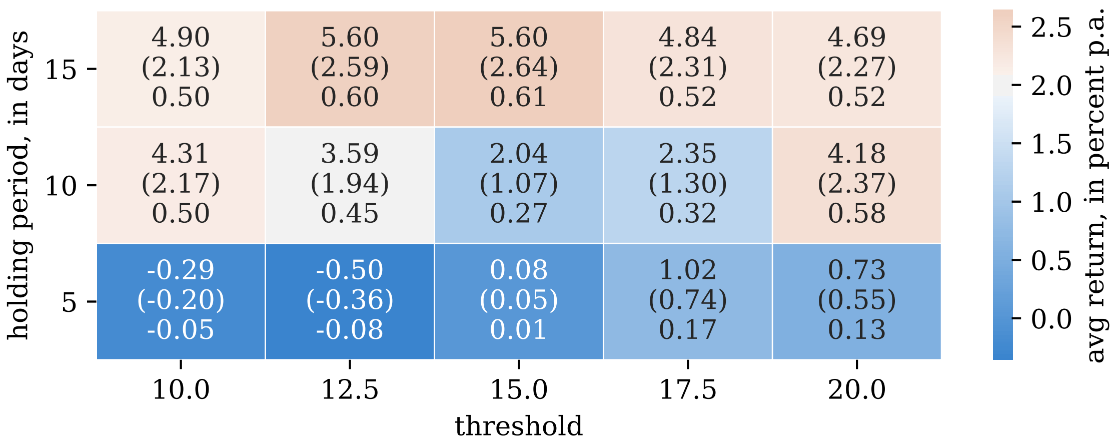

This figure depicts the cumulative returns of trading strategies that go long (short) currencies against the Japanese yen $n+1$ days before *predicted* local interest rate hikes (cuts) and reverses the position on the pre-announcement day, for holding horizons $n = 5, 10$ and 15 days and several thresholds, in basis points). The positions are rolled over daily with tom/next swaps. Long (short) trades and rollovers on corresponding positions are initiated at ask (bid) quotes. The upper panel of the figure shows returns plotted in calendar, the middle panel – those in event time. The heatmap in the lower panel shows the means (in percent per year), t-statistics thereof and Sharpe ratios (annualized) of the strategies' returns. The standard error of the mean is calculated based on monthly returns and is Newey and West (1987) HAC with 2 lags. The sample includes AUD, CAD, CHF, EUR, GBP, NZD, SEK and USD; JPY is the common counter currency. Sample size is August 2001 to December 2018.

Figure 2.6: Currency returns on local policy meeting days.

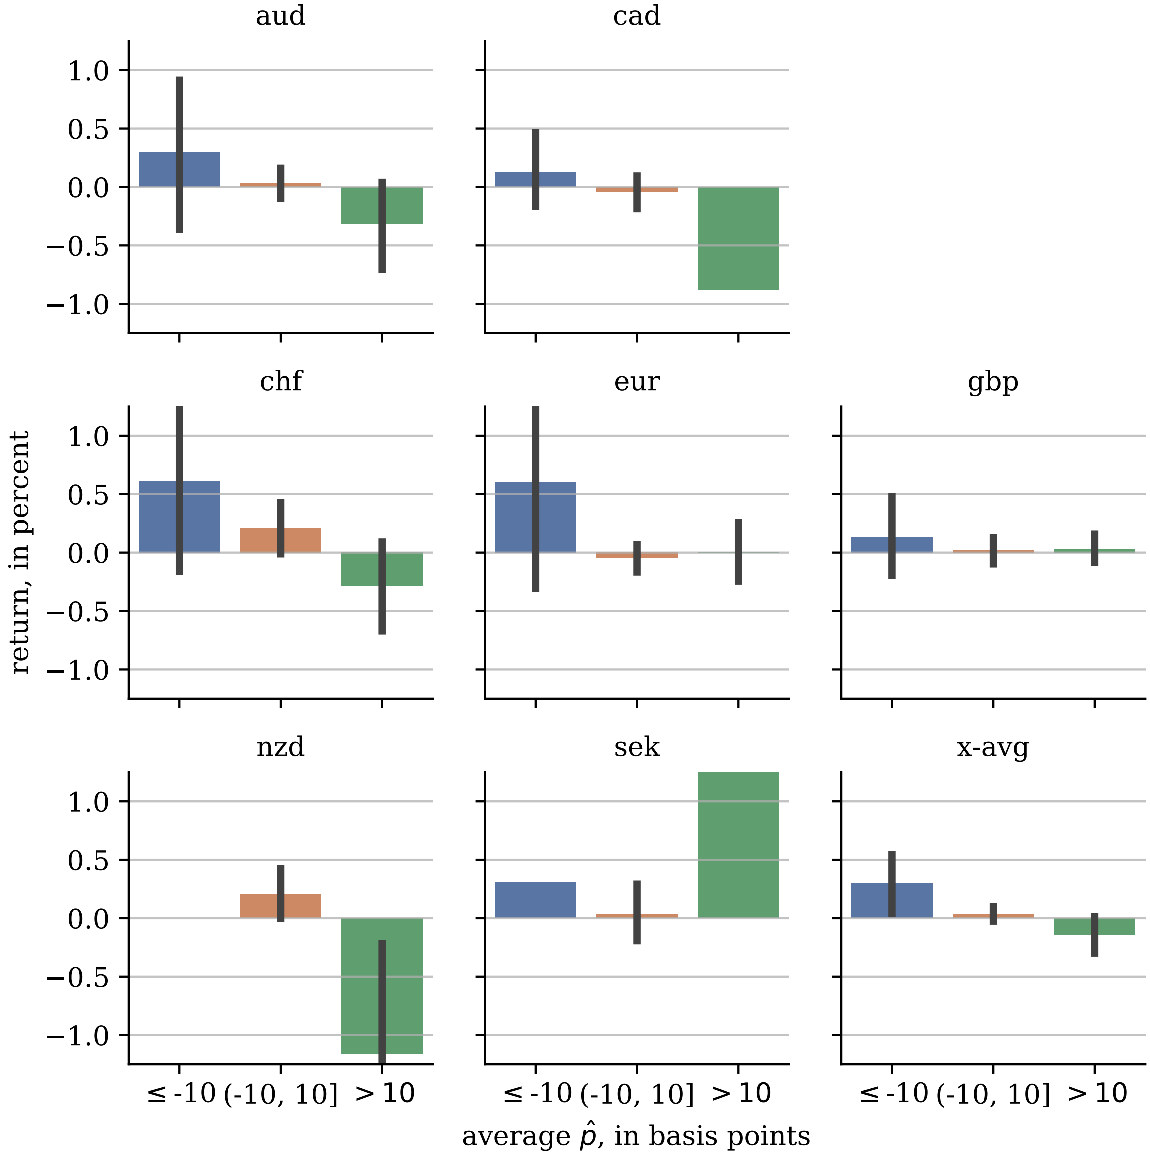

In this figure, the colored thick bars show the average spot appreciation of currencies against the US dollar recorded on those policy meeting days of respective regulators when the target rate remained unchanged. The average is taken conditional on the average magnitude of the OIS-based forecast of the target rate change $\hat{p}$ in equation (2.12) falling into one of three bins (on the x-axis): below -10, between -10 and $+10$ and above $+10$ basis points. The dark narrow bars span the 95% bootstrapped confidence intervals of the mean. The bottom right cell "x-avg" depicts the cross-currency average values. Sample size is August 2001 to December 2018.

Appendix A

Robustness checks

Figure A.1 shows the results of the event study similar to that in Section 2.3.1 with the crisis period (June 2007 to June 2009) excluded.

Figure A.1 shows the results of the panel regression estimation similar to that in Section 2.3.2 with the crisis period (June 2007 to June 2009) excluded.

Figure A.1: Exchange rates around local policy rate decisions.

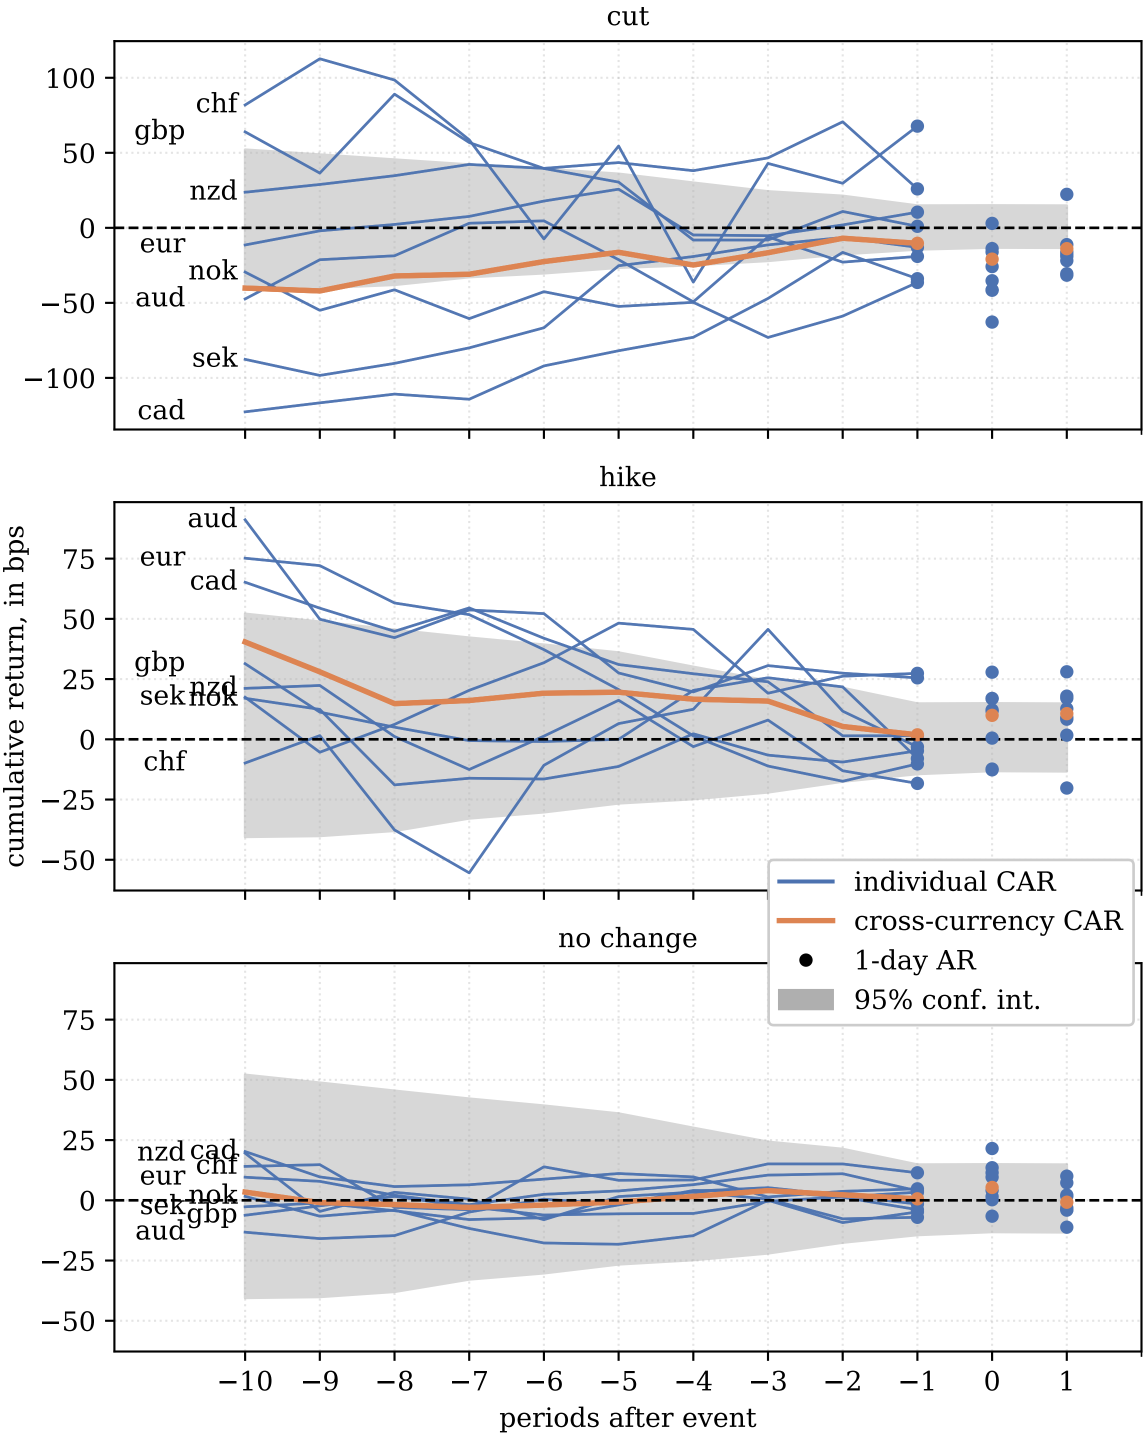

This figure depicts cumulative appreciation of currencies against the US Dollar before local announced policy rate decisions $\sum_{t=x}^{-1} \Delta s_{i,t}$, as well as the average return on the day of and after the latter. At each point x, it shows by how much a currency appreciates over x days preceding policy rate cuts, hikes and no-changes (in panels A, B and C respectively); the thick orange line indicates the weighted average cumulative return across all currencies, with weights proportional to the number of decisions for each. The day of announcements is marked by zero. The shaded area represents the 95% bootstrapped confidence interval for the cross-index average value around zero. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD and SEK for the period from August 2001 to December 2018, with the period between June 2007 and June 2009 excluded.

Table A.1: Currencies before local policy rate decisions.

	Panel A: h=5			*Panel B: h=10*			*Panel C: h=15*		
	cut	hike	zero	cut	hike	zero	cut	hike	zero
aud	-8.54 (-0.54)	1.27 (0.09)	0.06 (0.01)	-4.18 (-0.39)	1.94 (0.21)	-0.33 (-0.10)	-1.63 (-0.20)	0.95 (0.15)	-0.66 (-0.26)
cad	-8.91 (-0.70)	-0.34 (-0.03)	-1.95 (-0.39)	-7.22 (-0.68)	2.23 (0.36)	-0.49 (-0.14)	-5.50 (-0.74)	1.79 (0.36)	0.12 (0.05)
chf	16.36 (0.63)	-3.68 (-0.40)	1.73 (0.36)	2.17 (0.11)	1.19 (0.16)	-1.09 (-0.23)	3.43 (0.24)	2.62 (0.47)	-0.98 (-0.30)
eur	-2.96 (-0.29)	1.52 (0.14)	-1.20 (-0.24)	-2.32 (-0.30)	2.09 (0.30)	-1.12 (-0.35)	-1.72 (-0.28)	1.94 (0.38)	-1.14 (-0.47)
gbp	1.98 (0.18)	3.11 (0.38)	-1.44 (-0.40)	6.24 (0.56)	4.16 (0.71)	-0.95 (-0.37)	1.76 (0.25)	2.83 (0.53)	-1.40 (-0.56)
nok	-7.50 (-0.61)	2.24 (0.31)	-2.18 (-0.49)	-5.70 (-0.65)	3.74 (0.51)	-2.25 (-0.67)	-5.71 (-0.74)	2.18 (0.45)	-2.08 (-0.69)
nzd	-4.90 (-0.24)	0.91 (0.08)	-1.00 (-0.15)	-6.24 (-0.47)	1.14 (0.14)	-1.26 (-0.29)	-2.81 (-0.26)	2.40 (0.40)	-0.69 (-0.19)
sek	-1.08 (-0.08)	3.32 (0.32)	-0.34 (-0.06)	-5.10 (-0.57)	-0.27 (-0.04)	-1.63 (-0.40)	-3.77 (-0.56)	1.24 (0.19)	-0.59 (-0.19)
usd	4.46 (0.47)	0.76 (0.13)	-1.95 (-0.61)	-6.87 (-0.71)	-2.70 (-0.69)	-2.06 (-0.86)	-5.59 (-0.83)	-2.13 (-0.59)	-2.29 (-1.25)
joint	(2.37)	(1.20)	(0.50)	(3.09)	(1.41)	(1.56)	(3.25)	(2.27)	(1.43)
restr	-4.27 (-0.78)	1.23 (0.35)	-1.11 (-0.50)	-4.25 (-1.14)	1.16 (0.44)	-1.15 (-0.76)	-3.11 (-1.09)	1.19 (0.56)	-1.10 (-0.80)

This table shows the estimated β_i^p in (2.3) (rows "aud"–"usd") for different currencies i (across rows), policy rate decision types p (across columns) and lookback horizons h (across panels). The heatmap is a gradient from deep blue (negative values of the coefficient estimates) to orange (positive values) centered at zero. The coefficients can be interpreted as the average daily log-appreciation rate of currency i before policy rate decision of type p over h days before (and not including) the decision. The corresponding t-statistics, calculated using the Newey and West (1987) HAC covariance matrix estimator with automatic lag selection, are presented in parentheses. Row "joint" shows the joint t-statistic of testing $H_0 : \beta_{aud}^p = \ldots = \beta_{usd}^p = 0$. Row "restr" contains the estimated β^p from eq. (2.7), where the coefficients are assumed to be equal both across currencies. The sample includes AUD, CAD, CHF, EUR, GBP, NOK, NZD, SEK and USD for the period from August 2001 to December 2018, with the period between June 2007 and June 2009 excluded.

Appendix B

Forecasting power of OIS

In this section, we quantify the forecasting power of overnight index swaps. Let us assume that there is never any ambiguity about the alternative to a status quo retention, in other words, that there is never an option for three possible outcomes of any policy meeting – arguably a plausible assumption. With that in mind, investors face a two-class classification problem which they price in into overnight index swap rates. Define those events that result in a policy rate change as “positive”, and the others – as “negative” outcomes; the difference between the OIS-implied rate and the current average level of the underlying rate in equaton (2.12) is then used to classify an upcoming event as one of the two classes: positive whenever the difference is higher than a certain threshold, and negative otherwise. Clearly, some thresholds, e.g. 0.0 would produce a high true positive rate, when all rate changes would be forecast perfectly, but at the expense of missing all status quo retentions, or a high false positive rate. Other values, e.g. 0.5 (in annualized terms) would allow to correctly classify no-changes while missing successes altogether. Thresholds would also be different for different currencies. The best classifier in that respect takes on values that are consistently higher ahead of policy rate changes than ahead of no-changes: this concept can be summarized in the area under the receiver operating characteristic (ROC) curve, which equals 0.5 for a random classifier and 1.0 for a perfect one (for details, see e.g. Hastie, Tibshirani, and Friedman (2009)). Figure B.1 depicts the ROC curves and reports respective area values for all currencies in our sample.

While certain events are better forecastable than others, the OIS-implied rates are uniformly better than the random classifier. Interpreting the area under the ROC curve as the probability that a randomly drawn member of class 0 produces a score lower than the score of a randomly drawn member of class 1, we document this to be true in 70-90 percent of cases, the highest fraction manifesting itself for the Canadian, New Zealand and US dollar, and the lowest – for the Australian dollar (predominantly driven by the very start of the sample), Swiss franc and euro. We conclude that OIS contain important information about the markets' expectation of future policy rate changes and can serve the purpose of forecasting the latter.

Figure B.1: Forecasting power of OIS-implied rate.

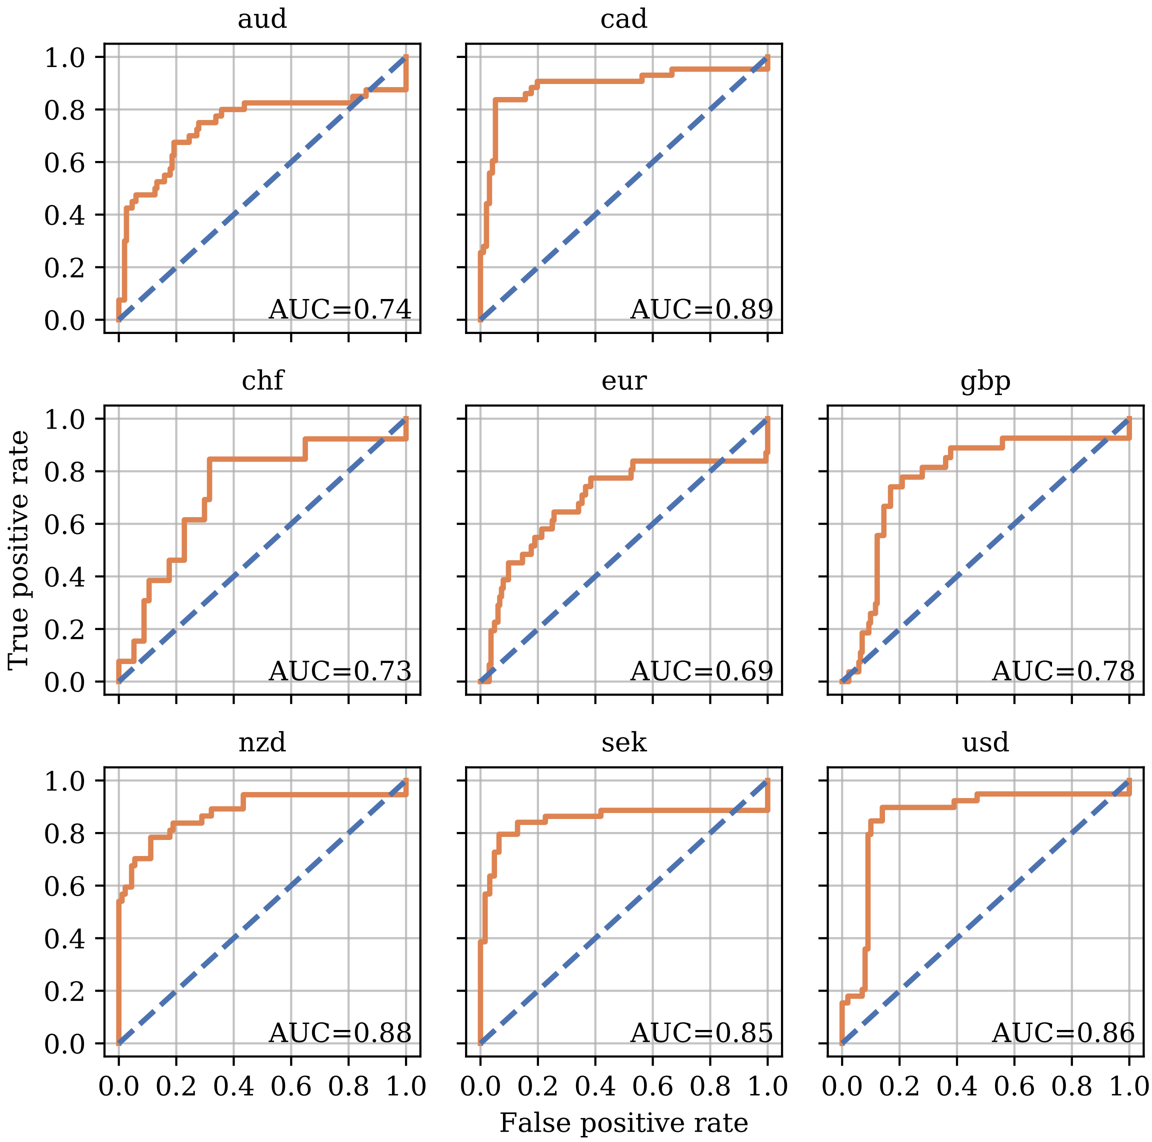

This figure shows receiver operating characteristic curves and areas below them for the OIS-implied classifier used to discriminate between upcoming policy rate changes ("positive" outcomes) and status-quo retentions ("negative" outcomes). It is assumed that before each event, only two outcomes are possible, as investors are aware about the sign of a possible change before actual hikes and cuts, and completely agnostic otherwise. The classifier is the difference between the OIS-implied and the 10-day rolling median (outside the previous event window) of the OIS underlying overnight rate. The sample is from August 2001 to December 2018, with the sample start slightly varying as described in Table 2.1.

Chapter 3

Verbal Interventions and Exchange Rate Policies: The Case of Swiss Franc Cap*

Nikola Mirkov†, Igor Pozdeev‡ and Paul Söderlind§

Abstract

We ask whether verbal interventions by the Swiss National Bank (SNB) affected market beliefs in the desired direction during the period from 2011 to 2015, when the SNB imposed a cap on the Swiss franc at 1.20 against the euro. A verbal intervention was a speech by a member of the SNB Governing Board containing the wording "utmost determination" and/or "unlimited quantities". We show that these verbal interventions lowered option-based measures of uncertainty regarding the future value of euro/Swiss franc exchange rate and steered market beliefs toward franc depreciation, therefore reinforcing the credibility of the Swiss franc cap.

*The views expressed in this paper are those of the authors, and do not necessarily reflect the views of the Swiss National Bank. We would like to thank Katrin Assenmacher, Christoph Meyer, Nicolas Stoffels, Attilio Zanetti, Carlos Lenz, Nina Karnaukh, Thomas Nellen, Hanno Lustig, Adreas Fischer, the participants of the 5th Workshop on financial determinants of exchange rates and the 20th ICMAIF conference for their valuable comments.

†Swiss National Bank, *E-mail*: nikola.mirkov@snb.ch

‡University of St. Gallen, *E-mail*: igor.pozdeev@unisg.ch

§University of St. Gallen, *E-mail*: paul.soderlind@unisg.ch

3.1 Introduction

Verbal interventions are generally shown to be a valuable tool in conducting foreign exchange policies. Beine, G. Janssen, and Lecourt (2009) finds that issuing commentary statements during foreign exchange interventions tends to reduce exchange rate volatility. Fratzscher (2008) shows that communication policies by the central banks of the US, euro area and Japan have influenced the exchange rates in the desired direction. Burkhard and Fischer (2009) find that repeated references of non-sterilized interventions by the Swiss National Bank (SNB) during the period 2002 – 2005 depreciated the domestic currency. On the other hand, comments and verbal interventions by officials can sometimes trigger sudden exchange rate jumps and hence temporarily increase market uncertainty, see Dewachter et al. (2014). Our study adds to this literature by examining the case of the Swiss franc cap against the euro enforced by the SNB during the period from September 2011 to January 2015.

In particular, the SNB set a 1.20 cap on the Swiss franc against the euro to avert the risk of deflation resulting from a "massive overvaluation of the Swiss franc", see SNB (2011). The franc had appreciated significantly against the euro and the US dollar in the months prior to the announcement amid the intensifying euro area crisis. The appreciation posed a threat to the small open Swiss economy. With the policy rate already being at zero, the SNB decided to announce a minimum euro/Swiss Franc exchange rate by promising to buy foreign currency in unlimited quantities if necessary.

In principle, a commitment to keeping the domestic currency low is always credible because a central bank can print unlimited amount of its own currency to buy foreign exchange.[1] However, as the balance sheet of the central bank grows in size and becomes increasingly volatile, the central bank becomes exposed to various financial risks. Excessive volatility could lead to significant balance sheet losses and even to negative equity positions. As noted by Danthine (2012) and Jordan (2011), this is not a problem in the short term but could generate doubts regarding credibility of the policy in the long term.

In this paper, we ask whether SNB's verbal interventions reinforced the credibility of the Swiss franc cap by affecting the market beliefs in the desired direction. We consider a verbal intervention to be a speech made by a member of the SNB Governing Board that contained the wording "utmost determination" and/or "unlimited quantities" when discussing the Swiss franc cap. Market beliefs we analyse are threefold: the perceived uncertainty of the future euro/Swiss franc rate and the skewness of the implied distribution extracted from option prices, and market liquidity measured by bid-ask spreads.

We find that SNB's verbal interventions reinforced the Swiss franc cap. They significantly reduced perceived uncertainty of the future euro/Swiss franc rate and had a positive

[1] Which is different from a situation in which the central bank is buying its own currency and selling foreign exchange in a hope that its reserves will not deplete before the goal is achieved.

effect on skewness of the implied distribution i.e. steering the market beliefs towads franc depreciation. The verbal interventions also had a positive effect on market liqudity by reducing the average bid-ask spread of the Swiss franc against a number of currencies excluding euro. All in all, they affected market beliefs in the desired direction and, more generally, have shown that verbal interventions can be useful in the implementation of foreign exchange policies.

Several papers study the period when the SNB imposed the cap on the Swiss franc cap against the euro. A. Janssen and Studer (2014) employ the Krugman (1991) model to show that market expectations that the SNB will intervene on the foreign exchange market if the euro/Swiss franc surpasses the announced bounds was most of the time sufficient to stabilize the euro/Swiss franc rate during the cap enforcement. Hertrich and Zimmermann (2017) explore the credibility of the Swiss franc cap and find that the cap was never perfectly credible, as the estimated probability of the euro/Swiss franc exchange rate being below 1.20 was high. Hanke, Poulsen, and Weissensteiner (2015) and Jermann (2017) report increasingly lower estimates of the break probability and conclude that the market's confidence in the SNB's commitment increased over time, especially from late-2012 until the end of their samples. We show in the data section that the latter result might be partially driven by a general downward trend in the global uncertainty in the FX market during the same period.

The remainder of the paper is organized as follows: section 3.2 describes the dataset, section 3.3 explains how we measure the effect of verbal interventions and section 3.4 reports our findings.

3.2 Data

The section describes the data set. It includes option data, SNB's verbal interventions, various financial market variables, and liquidity data from the currency spot markets.

3.2.1 Option data

The option data comprise 9 implied volatilities on the following instruments: 35-, 25-, 15- 10-delta risk reversals and butterflies, as well as a delta-neutral straddle. We use composite indicative quotes from Bloomberg, compiled daily at 17:00 New York time. The maturities in the sample are 1, 3, 6, 9 and 12 months and the period for our main sample is 6 September 2011 until 15 January 2015, when the euro/Swiss franc cap was in place.

These instruments, together with the spot exchange rate and riskfree rates (proxied by LIBOR rates from Bloomberg) are used to calculate implied volatilities and strike prices for nine different call options for each maturity (see Malz (2014) and Wystup (2006)). Figure 3.1 illustrates the available strike prices for the 1-month and 12-month horizons.

Figure 3.1: Strike prices of the currency options.

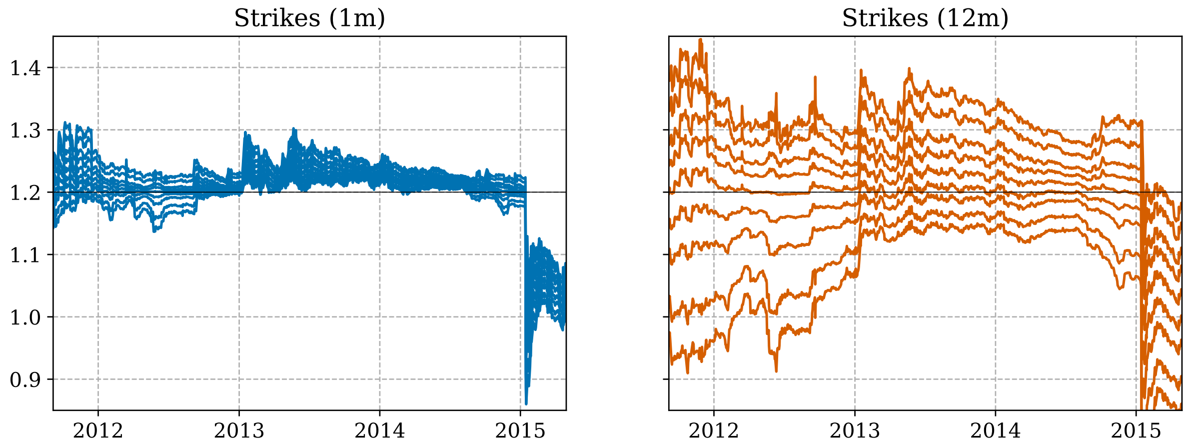

The figure shows for which strike prices our data set has (synthetic plain vanilla call and put) option prices. Only maturities of 1 month (left panel) and 12 months (right panel) are shown. The strike prices are obtained as a on-to-one mapping from the observed deltas of option contracts, following Wystup (2006)

The span of strike prices changes over time, being considerable in mid 2012 but more modest in mid 2014. In fact, the currency option market is organized in such a way that a higher perceived uncertainty about the future euro/Swiss franc leads to a wider span of observed strike prices. Except for the two shortest horizons (1- and 3-month), there are almost always several strike prices below 1.20. Further analysis (not tabulated) suggests that the liquidity of these instruments was good during the sample period (for instance, trading costs were arguably lower in 2014 than in either 2009 and 2016).

The prices of options are most often expressed in terms of the option-implied volatility, obtained by inverting the Black-Scholes formula. As an illustration, Figure 3.2 shows the implied volatilities for the 3-month horizon. Implied volatilities for low strike prices were particularly high at the inception of the Swiss franc cap and in 2012, normalizing in 2013 and staying relatively low until the end of cap enforcement.

Measuring uncertainty

To measure the uncertainty about future levels of the euro/Swiss franc spot exchagne rate, we use two different methods: (a) the implied volatility of the call option with a strike price closest to the spot exchange rate ("at the money"); and (b) the square root of the model-free implied variance (see Jiang and Tian (2005)) which is (loosely speaking) a weighted average of all implied volatilities. Both measures can be interpreted as the market's subjective beliefs about (annualized) volatility over the period until maturity.

The uncertainty measure from the first method is illustrated (for the 3-month horizon

Figure 3.2: Implied volatilities of the currency options.

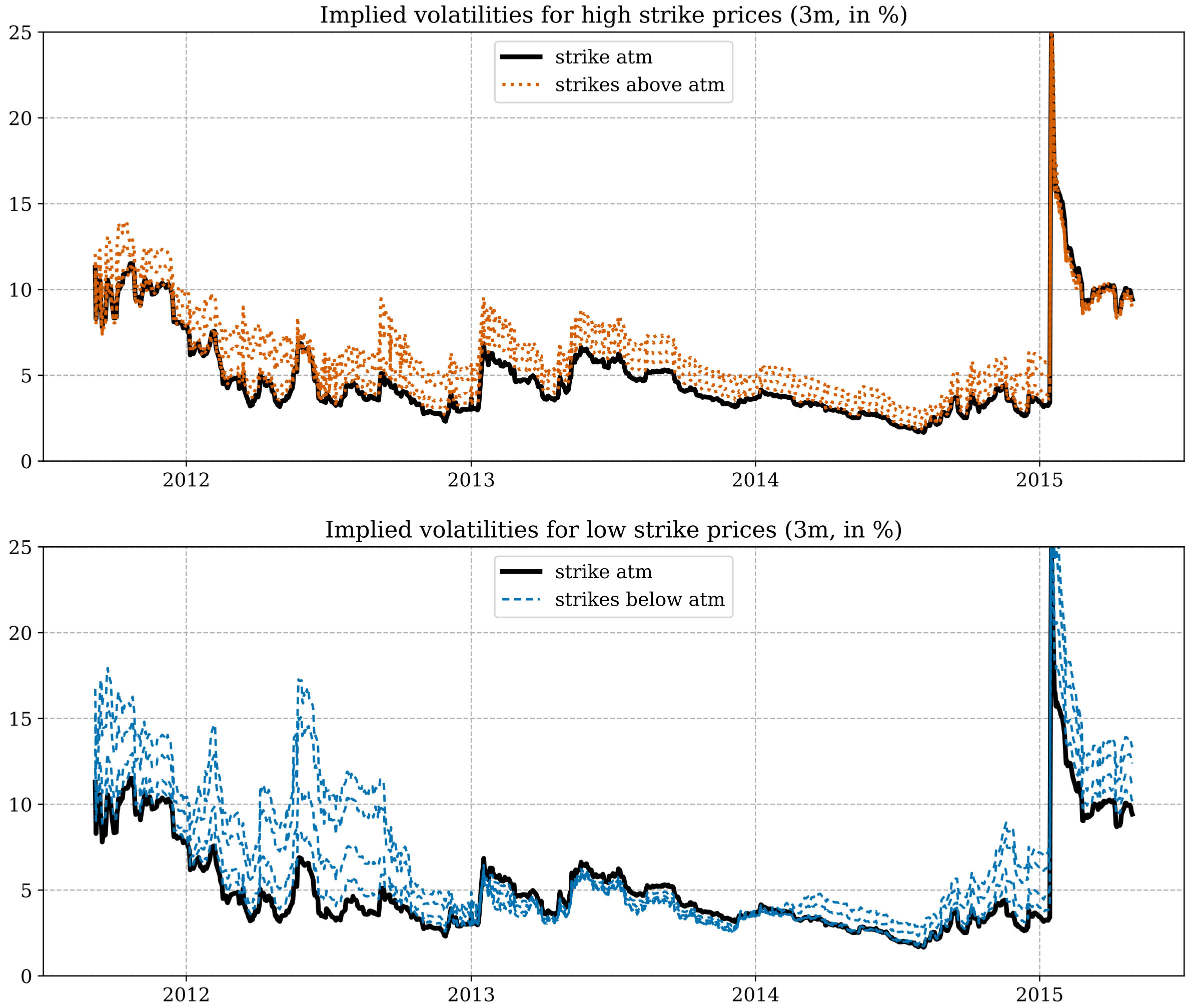

The figure shows all 9 implied volatilities of the synthetic 3-month plain vanilla call and put options on the euro/Swiss franc spot exchange rate. The upper panel covers in-the-money, and the lower panel – out-of-the money options. The volatilities are recovered from option contracts as described in Malz (2014).

only) in Figure 3.3. As a comparison, we also plot an uncertainty measure for the global currency markets excluding the Swiss franc (see below for a precise definition). Interestingly, the euro/Swiss franc uncertainty is strongly linked to the global currency market uncertainty[2]:

$$\Delta U_{eur,t} = \underset{(-0.80)}{-0.0059} + \underset{(16.07)}{1.2086}\overline{\Delta U}_t,$$
$$R^2_{adj.} = 0.54,$$

where U_i is the uncertainty for the exchange rate of currency i against the Swiss franc, and $\overline{\Delta U} = 1/N \sum_{i \neq eur}^{N} \Delta U_i$. Both the euro/Swiss franc uncertainty and the global currency

[2]The t-statistics are shown in parentheses below respective coefficients. The sample period is 6 September 2011 to 14 January 2015.

market uncertainty series experience a relatively steady downward trend during the period of cap enforcement. This stylized feature in the data might be driving partially the results of the previous studies, who estimate that the probability of SNB discontinuing the cap was falling over time, see for example Hertrich and Zimmermann (2017). Figure 3.3 also shows considerable idiosyncratic movements in the euro/Swiss franc uncertainty, for instance, in early 2013.

Figure 3.3: Uncertainty and skewness.

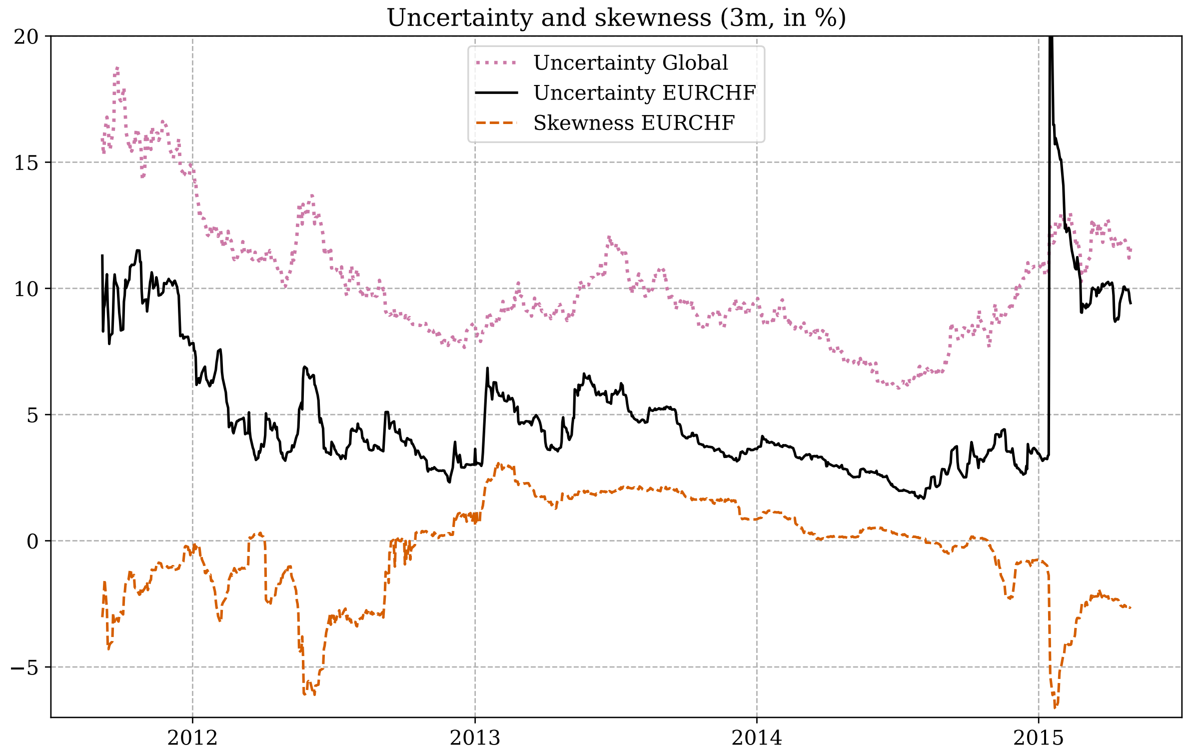

The figure shows the 3-month uncertainty and skewness measures for the euro/Swiss franc spot exchange rate, as well as the measure for the global FX uncertainty. The uncertainty measure is the observed Black-Scholes implied volatility of the at-the-money call option; the skewness measure is calculated as the average risk-reversal implied volatility. The global FX uncertainty is defined as the average of the individual uncertainty measures for bilateral exchange rates of 9 currencies (G10 currencies except the euro) against the Swiss franc.

Measuring skewness

To measure skewness of the implied distribution, we apply two related methods: (a) the implied volatility of options with high strike prices (above the spot rate) minus that of options with low strike prices (this is the average risk reversal); and (b) the model free upside standard deviation minus the downside standard deviation. Both measures can be interpreted as the "tilt" of the market beliefs: towards a higher exchange rate (positive skewness) or a lower one.

The skewness measure from the first method is illustrated in Figure 3.3. It was mainly negative at the beginning of the cap period, corresponding to high implied volatilities of options with low strike prices as seen in Figure 3.2. It started drifting upwards in autumn 2012 and remained high throughout most of 2013, but then started drifting downwards in December 2013 (on 5 December, ECB decided to keep its policy rates unchanged, but reiterated its forward guidance that key ECB rates are expected to remain low "for an extended period of time."). It turned slightly negative in late 2014.

3.2.2 Liquidity data

We study bid-ask spreads of the major currency pairs involving the Swiss franc, but excluding the euro. In general, high bid-ask spreads tend to be related to uncertainty and asymmetric information. In the case of the euro/Swiss franc cap, it can also be conjectured that an expected lifting of the Swiss franc cap would make liquidity providers cautious of an upcoming large-scale rebalancing—thus prompting them to widen the bid-ask spreads.

Our analysis of liquidity uses bid and ask spot quotes of the 9 most traded currency pairs against the Swiss franc: Australian, Canadian, New Zealand and the US dollar, British pound, Japanese yen, Norwegian and Danish krone, and Swedish krona. We exclude the euro/Swiss franc since the SNB is a dominant player on that market and was known to intervene heavily.

Similar to option data, we use the composite indicative bid and ask quotes (from Bloomberg) compiled at 17:00 New York time. We focus on the de-seasonalized weighted average of the bid-ask spreads of the nine currency pairs.[3] The data is de-seasonalized by removing month, week and weekday components.

The liquidity data is illustrated in Figure 3.4. The bid-ask spreads are very volatile, so the figure also shows a moving average: before (after) 15 January 2015 it is a backward (forward) looking 10-day moving average. Overall, the spreads appear to follow the same pattern as the uncertainty measures discussed before.

3.2.3 SNB's communications

Announcing the decision to impose a cap of 1.20 on the Swiss Franc against the euro, the SNB pledged to "enforce this minimum rate with the utmost determination and [...] buy foreign currency in unlimited quantities." The keywords were subsequently used in a variety of speeches by the members of the SNB Governing Board. A verbal intervention

[3]The weight for each currency pair is the share of its base currency in global foreign exchange market turnover in 2013 as taken from the BIS Triennial Central Bank Survey (http://www.bis.org/publ/rpfx13.htm). For instance, the share of all USDXXX pairs is 87%, that of all JPYXXX pairs is 23%, and so forth. We normalize the weights such that they sum up to one.

Figure 3.4: Bid-ask spreads.

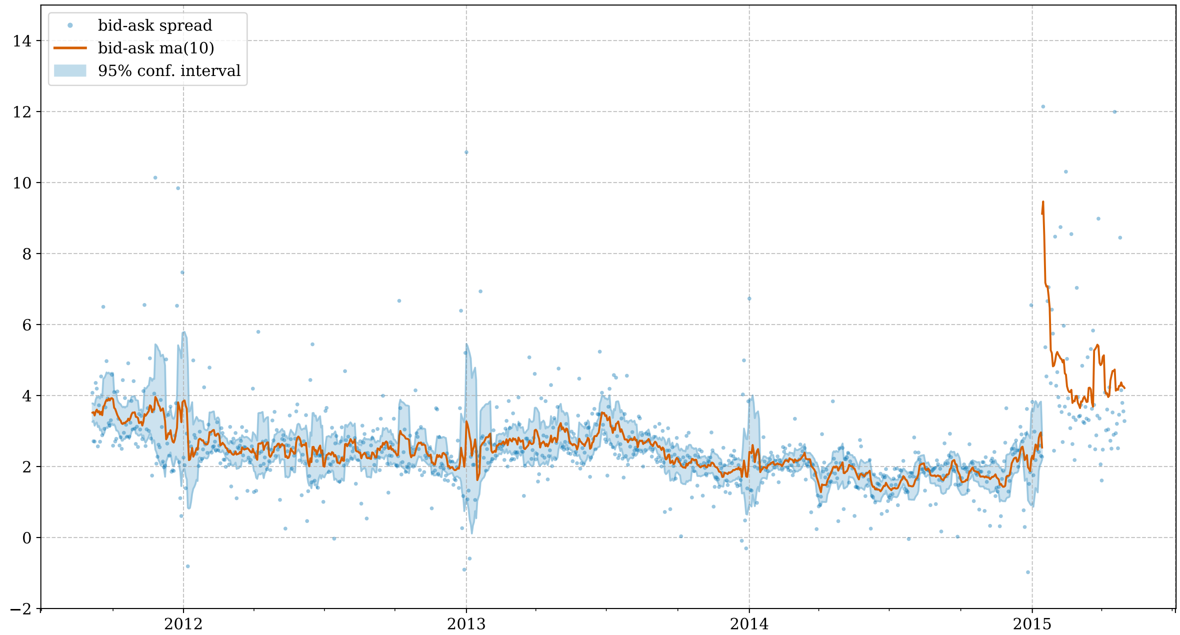

The figure shows the volume-weighted cross-sectional average spread (in pips) of bilateral exchange rates of 9 currencies (G10 currencies except the euro) against the Swiss franc, and the 10-day moving average thereof. The data is de-seasonalized by removing the month, week and weekday components. The blue shaded area – the 95% confidence interval around this value, calculated under assumption that daily average spread values are i.i.d.

is considered to be a speech made by a member of the SNB Governing Board that used this key wording.

There were 26 verbal interventions during the period from 6 September 2011, to 14 January 2015, see the left panel of Table 3.1. Most of these were delivered by Thomas Jordan, four of them by Fritz Zurbrügg, three of them by Jean-Pierre Danthine, and one of them by Philipp Hildebrand. Several speeches on the same day are counted as one event. The empirical analysis uses these 26 dates as dummy variables in regressions of either uncertainty, skewness or liquidity. On 7 ocassions (in June and December), the speech was presenting the regular "Monetary Policy Assessments" (henceforth MPA). In later sections, we split the verbal intervention dummy into one variable for MPA dates and another for the rest.

A second dummy variable indicates the 7 MPAs that were communicated by press releases (March and September), not the speeches. They all contain the "utmost determination" or "unlimited quantities" keywords. See the mid panel of Table 3.1 for details.

Finally, a third dummy variable indicates the 18 policy related speeches that did not contain the "utmost determination" or "unlimited quantities" keywords, see the mid panel of Table 3.1. Only a few non-policy speeches are excluded, for instance, the 100 year celebration of the SNB office building in Berne.

Table 3.1: List of the SNB's verbal interventions and ECB announcements.

SNB, Verbal Interventions		SNB, Other Speeches	ECB	
8-Nov-11		28-Sep-2011	8-Sep-11	regular
15-Dec-11	MPA	3-Nov-2011	6-Oct-11	regular
24-Jan-12		28-Feb-2012	3-Nov-11	regular
7-Feb-12		14-May-2012	8-Dec-11	LTRO
22-Mar-12		5-Sep-2012	12-Jan-12	regular
10-Apr-12		3-Oct-2012	9-Feb-12	regular
27-Apr-12		25-Oct-2012	8-Mar-12	regular
31-May-12		21-Nov-2012	4-Apr-12	regular
14-Jun-12	MPA	21-Mar-2013	3-May-12	regular
3-Sep-12		16-Apr-2013	6-Jun-12	regular
8-Nov-12		23-Sep-2013	5-Jul-12	regular
16-Nov-12		11-Oct-2013	26-Jul-12	"whatever it takes"
28-Nov-12		14-Nov-2013	2-Aug-12	OMT
13-Dec-12	MPA	16-Jan-2014	6-Sep-12	regular
19-Feb-13		8-May-2014	4-Oct-12	regular
26-Apr-13		20-May-2014	8-Nov-12	regular
20-Jun-13	MPA	9-Oct-2014	6-Dec-12	regular
8-Oct-13		23-Nov-2014	10-Jan-13	regular
21-Nov-13			7-Feb-13	regular
12-Dec-13	MPA		7-Mar-13	regular
27-Mar-14		SNB MPA, Press Releases	4-Apr-13	regular
25-Apr-14		15-Sep-2011	2-May-13	regular
19-Jun-14	MPA	15-Mar-2012	6-Jun-13	regular
20-Nov-14		13-Sep-2012	4-Jul-13	regular
11-Dec-14	MPA	14-Mar-2013	2-Aug-13	regular
18-Dec-14		19-Sep-2013	5-Sep-13	regular
		20-Mar-2014	2-Oct-13	regular
		18-Sep-2014	7-Nov-13	regular
			5-Dec-13	regular
			9-Jan-14	regular
			6-Feb-14	regular
			6-Mar-14	regular
			3-Apr-14	regular
			8-May-14	regular
			5-Jun-14	negative depo rate & TLTRO
			3-Jul-14	regular
			7-Aug-14	regular
			4-Sep-14	regular
			2-Oct-14	regular
			6-Nov-14	regular
			4-Dec-14	regular

The first column (SNB verbal interventions) of the table reports the speeches by the members of the SNB's Governing Board used to create the dummy variable d_t in equation (3.1). The speeches presenting the "Monetary Policy Assessment" are marked "MPA". It also reports other SNB speeches, SNB MPA press releases (both in column 2) and the regular monetary policy decisions by the ECB and several key policy announcements (ECB, column 3) used to create the dummies in D_t in equation (3.1).

3.2.4 Other data

The dataset includes a number of financial market variables that we control for when regressing option-implied and liquidity data on SNB's verbal interventions. First, we collect major announcements and regular policy decisions by the ECB that might have considerably affected the option market for the euro/Swiss franc rate, see the right panel of Table 3.1. In addition, our dataset includes the TED spread (from the Federal Reserve Bank of St. Louis), the 5-year European sovereign CDS spread (from Bloomberg), a measure of global currency market uncertainty mentioned before (the cross-sectional average of the square root of model-free implied variance of US dollar exchange rates of G10 currencies excluding the Swiss franc), the 10-year US term premium of Kim and Wright (2005) (from the Federal Reserve's official website), and the spot euro/Swiss franc franc exchange rate (from Bloomberg).

3.3 Empirical strategy

3.3.1 Regressions

We run the following type of regression of the $1+n$ day change of uncertainty (or skewness or liquidity)

$$y_{t+n} - y_{t-1} = \alpha + \beta d_t + \delta' D_t + \gamma' \left(X_{t+n} - X_{t-1} \right) + \varepsilon_{t+n}. \tag{3.1}$$

In these regressions, y_t is the uncertainty (or skewness or the average bid-ask spread) on day t, d_t is a dummy variable that takes the value of one on the day of the verbal SNB intervention (speeches that included the keywords "utmost determination" or "unlimited quantities") and zero otherwise, D_t is a vector of three dummy variables for the MPAs communicated by press releases, other SNB speeches and ECB policy decisions (see Table 3.1 for the dummies), and X_t is a vector of other control variables. The latter include the TED spread, the Eurozone CDS spreads, global currency market volatility, the 10-year US term premium and the spot euro/Swiss franc rate.

3.3.2 Event studies

The regression results are reinforced by an event study approach along the lines of E. F. Fama et al. (1969). To study the evolution of the effect the verbal interventions had on the uncertainty and skewness, we calculate daily changes in the two measures and define the "normal" changes as those prevailing on average over day $t-45$ to day $t-1$, where

t is the day of a speech (similar to S. J. Brown and Warner (1980)).[4] If two speeches are less than 45 days apart, we reduce the estimation window to avoid overlap. That said, the "abnormal" change on day t is the actual change minus the "normal" change. For each day, we compute the cumulative abnormal change in the uncertainty and skewness $n = 1, 2, \ldots, 5$ days forward. We use the standard deviation on days $t - 45$ to $t - 1$ (or in the shorter interval as above) to gauge significance of the the average (over all speeches) cumulative abnormal change.

3.4 Results

This section is structured as follows. We first report the results of estimated effects of verbal interventions on uncertainty. Then, we show the results for skewness of the implied distribution and for market liquidity. Finally, we show and discuss the robustness checks we performed.

3.4.1 Verbal interventions and uncertainty

The SNB's verbal interventions (speeches mentioning "utmost determination" or "unlimited quantities") significantly reduced the uncertainty regarding the future value of the euro/Swiss franc rate during the period of cap enforcement, see Table 3.2. The table shows the regression coefficients and t-statistics for the *two-day change of uncertainty* ($n = 1$ in equation 3.1). The t-statistics are based on a Newey-West approach with a bandwidth of 2 (using a Hansen-Hodrick approach gives very similar results).

The effect is stronger for shorter maturities: it is -0.20 for the 1-month horizon and -0.12 for the 12-month horizon. The same goes for most of the other explanatory variables: the slope coefficients for the 1-month uncertainty are higher than the ones for the 12-month uncertainty. This fact might suggest that the short-term market beliefs are much more responsive to news relative to the long-term ones. Comparing the coefficients with the movements of uncertainty in Figure 3.3, suggests that the economic effect is non-trivial but not very large: uncertainty jumps from 5% to 6% (or similarly) are not unusual in the figure, but the verbal intervention would only account for up to 0.2%.

The MPAs, communicated by press releases and containing the keywords, also have strongly negative effects. For longer maturities, their effect is twice as larger relative to verbal interventions. Interestingly, this pattern does not show up in the regressions for skewness, see the next section.[5] Other SNB speeches that did not contain the keywords "utmost determination" or "unlimited quantities" had no effect on uncertainty.

[4] Changing the estimation window (for instance, to 22 days) has a negligible impact on the results.

[5] See also the section 3.4.4, where we split all the verbal interventions to those made during the MPAs and the rest.

Table 3.2: The SNB's verbal interventions and market uncertainty.

	1m	3m	6m	9m	12m
Verbal interventions	−0.20	−0.14	−0.12	−0.12	−0.12
	(−2.27)	(−2.02)	(−2.51)	(−2.67)	(−2.68)
MPA (press release)	−0.33	−0.24	−0.24	−0.23	−0.25
	(−2.01)	(−2.09)	(−2.53)	(−2.86)	(−3.16)
Other speech	−0.04	−0.06	−0.05	−0.05	−0.04
	(−0.34)	(−0.68)	(−0.76)	(−0.92)	(−0.91)
ECB	0.12	0.16	0.09	0.07	0.05
	(1.37)	(2.63)	(1.97)	(1.80)	(1.59)
TED	−0.87	−0.60	−0.65	−0.41	−0.11
	(−0.66)	(−0.61)	(−0.78)	(−0.56)	(−0.15)
EU CDS	−0.17	−0.12	0.00	0.06	0.07
	(−0.87)	(−0.74)	(0.03)	(0.43)	(0.55)
FX vol	0.53	0.41	0.31	0.27	0.24
	(7.56)	(7.35)	(5.86)	(5.25)	(4.72)
10y US TP	−1.45	−0.76	−0.34	−0.16	−0.07
	(−3.44)	(−2.83)	(−1.58)	(−0.79)	(−0.37)
Spot	0.71	0.44	0.28	0.21	0.18
	(6.50)	(6.29)	(5.21)	(4.76)	(4.40)
R^2	0.36	0.32	0.26	0.24	0.21

The table reports estimated coefficients of equation (3.1) and t-stats (in parentheses) for $n = 1$. The dependent variables are the 2-day changes of uncertainty (implied volatility at the money) from options with 1, 3, 6, 9 and 12 months to expiration (different columns). The right-hand variables are (from top to bottom) a dummy for SNB's verbal interventions, three other dummy variables: SNB "Monetary Policy Assessment" (MPA) press releases, other SNB speeches and ECB key policy decisions and also two-day changes of a number of control variables: the TED spread, the European sovereign CDS spread, a measure of global FX market uncertainty (FX vol), 10-year term premium on US Treasuries (10y US TP), and the the spot euro/Swiss franc rate. The intercept is not reported. The t-stats are estimated using the Newey and West (1987) approach with 2 lags. The sample goes from 6 September 2011 to 14 February 2015.

For the control variables, both the global FX market volatility and the spot euro/Swiss franc exchange rate have positive and strongly significant coefficients. Similarly, the announcements made by the ECB tend to increase euro/Swiss franc uncertainty, but the coefficients are significant only for the 3, 6 and 9 month horizons. In contrast, TED spread and the average CDS spread across EU countries, have insignificant coefficients.

Table 3.3 shows results for one-day changes ($y_t - y_{t-1}$ so $n = 0$) to six-day changes ($y_{t+5} - y_{t-1}$ so $n = 5$). For brevity, only the coefficients for SNB's verbal intervention dummies are shown—with indicators of statistical significance. It seems as if the effect of the speeches builds up during a few trading days, but starts to evaporate after four days. The statistical significance is stronger for the longer maturities.

As a support for the latter two findings, Figure 3.5 illustrates the event study. It shows

Table 3.3: The dynamic effect of the SNB's verbal interventions on uncertainty.

Change over	1m	3m	6m	9m	12m
$t-1$ to t	-0.14	-0.07	-0.08^{**}	-0.08^{**}	-0.09^{***}
$t-1$ to $t+1$	-0.20^{**}	-0.14^{**}	-0.12^{**}	-0.12^{***}	-0.12^{***}
$t-1$ to $t+2$	-0.15	-0.10	-0.10^{*}	-0.11^{**}	-0.12^{***}
$t-1$ to $t+3$	-0.19	-0.13	-0.13^{*}	-0.13^{**}	-0.14^{**}
$t-1$ to $t+4$	-0.18	-0.12	-0.15^{*}	-0.14^{**}	-0.15^{**}
$t-1$ to $t+5$	-0.01	-0.02	-0.09	-0.10	-0.12

The table shows the coefficient estimates on the dummy variable for SNB's verbal inteventions in equation (3.1) for $n = 0$ to $n = 5$. The regressors and the sample are the same as in Table 3.2. For brevity, only the coefficients for the SNB dummy are presented. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively, using t-stats estimated by the Newey and West (1987) approach with $1 + n$ lags.

the average (across speeches) cumulative changes in the uncertainty at different days after a speech, as well as a 95% confidence interval around zero. In line with the regression findings, we see that the first day's change in uncertainty is negative and marginally significant. During the next three days, uncertainty continues to decline significantly. The magnitude of the changes are similar to those of the regression coefficients, which is explained by the fact that the verbal intervention dummy is virtually uncorrelated with the other regressors in equation (3.1), with correlations between -0.07 and 0.02.

3.4.2 Verbal interventions and skewness

We find that SNB's verbal interventions steered the market beliefs toward depreciation of the Swiss franc by having a positive and statistically significant effect on skewness of the implied distribution, see Table 3.5. The table again reports regression results for the *two-day change of skewness*. The regressors are the same as before, except that uncertainty is added to the set of regressors (to control for any direct effects of uncertainty on skewness). Similarly to before, the economic effect is not large, but still strongly significant.

The MPAs communicated by press releases had a strong positive effect only on the shortest maturity. The rest of the maturities were basically unaffected. The same goes for other speeches made by the members of the Governing Board that did not contain the key wording.

Among the control variables, troubles in the euro zone as measured by the euro zone CDS spreads and global currency market uncertainty tilt the beliefs towards a lower euro/Swiss franc rate, whereas a higher current spot rate does the opposite. In contrast, ECB announcements, the TED spread and the euro/Swiss franc uncertainty have no (or only small) significant effect.

Figure 3.5: The dynamic effect of SNB's verbal interventions on uncertainty: an event study.

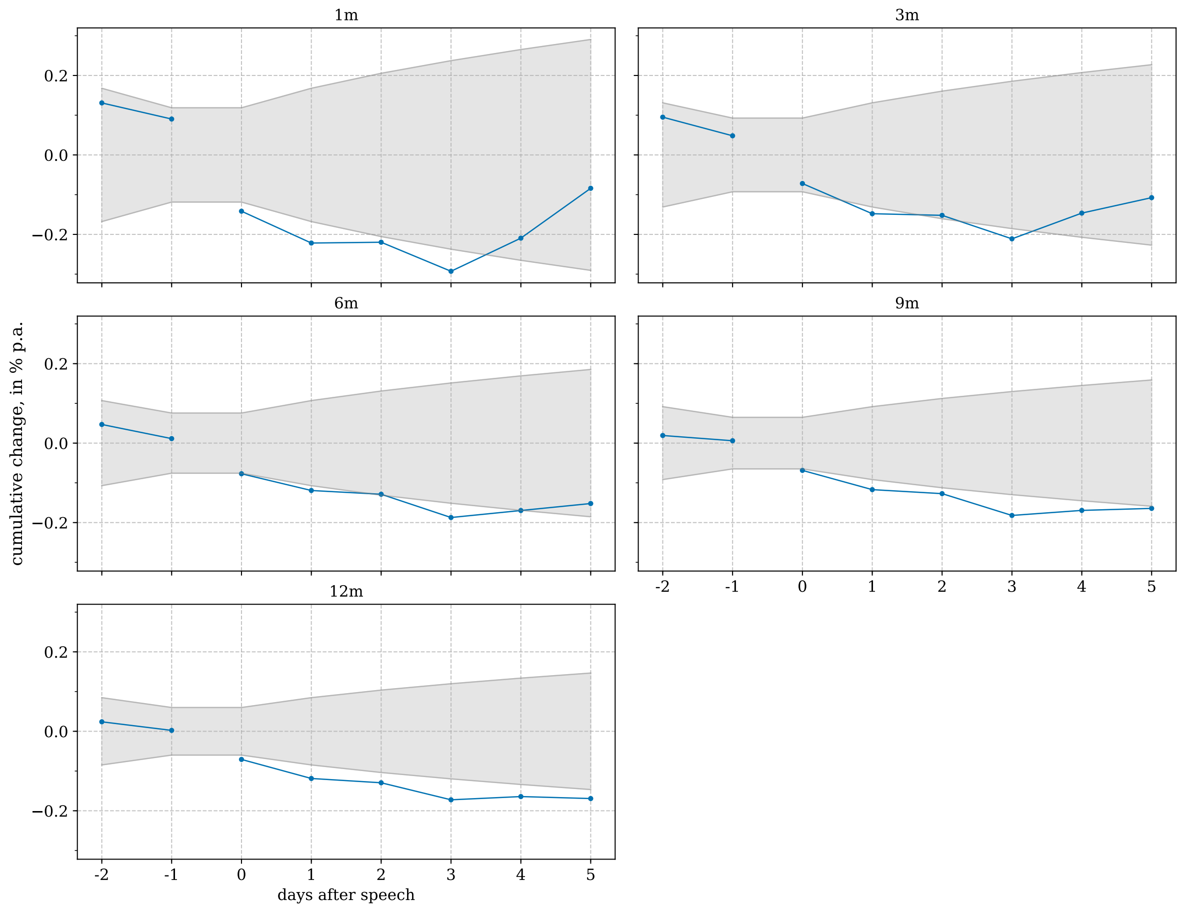

The figure shows the cumulative average difference in uncertainty since the speech day (day 0) and the 95% confidence band around zero, for the 1-, 3-, 6-, 9- and 12-month horizon. The uncertainty at any given day is defined as the Black-Scholes implied volatility of the at-the-money call option observed on that day. The confidence interval is calculated assuming cross-sectional and intertemporal independence of differences in uncertainty.

Table 3.4 shows that the effect of SNB's verbal interventions accumulates for a few days. The event study reported in Figure 3.6 shows the same pattern.

3.4.3 Verbal interventions and market liquidity

Table 3.6 reports regression results for the *two-day change of bid-ask spreads*. The coefficient on SNB's verbal interventions is negative and significant on the 8% level. The fit of this regression is poor (low R^2) as should be expected from such volatile data. It can also be shown (not tabulated) that the results for changes over fewer ($n = 0$) or more ($n \geq 2$) are similar. This lends some support to the earlier results that the verbal interventions contributed to lower market uncertainty.

Table 3.4: The dynamic effect of the SNB's verbal interventions on skewness.

Change over	1m	3m	6m	9m	12m
$t-1$ to t	0.03	0.04*	0.07**	0.05*	0.04
$t-1$ to $t+1$	0.04	0.09*	0.10**	0.10***	0.11***
$t-1$ to $t+2$	0.11	0.10*	0.11**	0.09*	0.05
$t-1$ to $t+3$	0.13*	0.09	0.11*	0.10*	0.09
$t-1$ to $t+4$	0.23**	0.20**	0.20***	0.17**	0.14*
$t-1$ to $t+5$	0.20**	0.22**	0.17**	0.12	0.10

The table shows the coefficient estimates on the dummy variable for SNB's verbal interventions in equation (3.1) for $n = 0$ to $n = 5$. The sample and the right-hand variables are the same as in Table 3.5. For brevity, only the coefficients for the SNB dummy are presented. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively, using t-stats estimated by the Newey and West (1987) approach with $1 + n$ lags.

3.4.4 Extended regressions

News conferences of the SNB on the occasions of MPAs represent major market events globally. The introductory remarks by the members of the Governing Board during these news conferences likely enjoyed a higher media coverage relative to other speeches. Therefore, we split all the verbal interventions to those made during the MPAs and the rest and look into their effects separately, see Table 3.7. As expected, the effects of MPA-interventions are much stronger and more statistically significant relative to non-MPA speeches.

As a robustness check, we also re-run the regressions by using other measurements of the dependent variable and also allow for time-varying effects. The regression results remain broadly unchanged, if we use the model-free uncertainty and skewness, see Table 3.8.

Finally, we investigate the stability over time of the regression results. We re-run the previous regressions but add a dummy for the first verbal intervention, then (in a second regression) dummy for the 2nd verbal intervention instead, and so forth until the last (26th) intervention. The results are illustrated in Figure 3.7, which show the sum of the average effect (coefficient on the previous dummies for SNB's verbal interventions) and the "extra" effect on each intervention days. On most days the effect is negative and often significantly so. There are only three days of significantly positive effects: the first and last intervention days and also the 16 November 2012.

Table 3.5: The SNB's verbal interventions and market skewness.

	1m	3m	6m	9m	12m
Verbal interventions	0.04	0.09	0.10	0.10	0.11
	(1.01)	(1.85)	(2.12)	(2.62)	(2.66)
MPA (press release)	0.18	0.06	0.07	0.01	0.01
	(1.83)	(1.33)	(1.39)	(0.20)	(0.21)
Other speech	0.10	0.08	0.03	0.01	0.00
	(1.08)	(1.15)	(0.50)	(0.11)	(0.01)
ECB	0.10	−0.06	−0.03	−0.03	−0.02
	(1.02)	(−0.69)	(−0.58)	(−0.61)	(−0.56)
TED	1.31	0.21	−0.78	−0.89	−1.61
	(1.36)	(0.25)	(−1.03)	(−1.14)	(−1.75)
EU CDS	−0.22	−0.33	−0.25	−0.26	−0.09
	(−1.28)	(−2.04)	(−1.90)	(−2.05)	(−0.48)
FX vol	−0.09	−0.11	−0.11	−0.11	−0.10
	(−1.70)	(−2.16)	(−2.38)	(−2.55)	(−2.30)
10y US TP	−0.23	0.14	0.13	0.19	0.22
	(−0.57)	(0.43)	(0.42)	(0.74)	(0.87)
Spot	0.19	0.15	0.11	0.11	0.12
	(3.24)	(2.84)	(2.09)	(2.11)	(2.16)
Uncertainty	0.12	0.03	0.04	0.04	0.02
	(1.25)	(0.34)	(0.59)	(0.54)	(0.31)
R^2	0.10	0.09	0.09	0.10	0.07

The table reports estimated coefficients of equation (3.1) and t-stats (in parentheses) for $n = 1$. The dependent variables are the 2-day changes of skewness (implied volatility of options with high strike prices minus the implied volatility of options with low strike prices). The sample and the right-hand variables are the same as in Table 3.2, except that the two-day change of uncertainty is added to the set of control variables. The t-stats are estimated using the Newey and West (1987) approach with 2 lags.

3.5 Conclusion

We examine whether the use of the wording "utmost determination" and/or "unlimited quantities" by the members of the SNB Governing Board significantly affected market beliefs during the period from 2011 to 2015 when the Swiss franc has an upper bound at 1.20 against the euro. We find that such verbal interventions by the SNB reduced uncertainty regarding the future value of euro/Swiss franc spot exchange rate and had a positive effect on skewness of the implied euro/Swiss franc distribution. Therefore, they reinforced the Swiss franc cap and, more generally, provided further evidence of how useful the verbal interventions can be in the implementation of exchange rate policies.

Table 3.6: The SNB's verbal interventions and market liquidity.

	Bid-ask spread
Verbal interventions	−0.25
	(−1.73)
MPA (press release)	−0.55
	(−1.26)
Other speech	0.26
	(0.83)
ECB	0.41
	(1.76)
TED	−9.15
	(−3.39)
EU CDS	−0.44
	(−0.82)
FX vol	0.33
	(2.21)
10y US TP	0.82
	(1.03)
Spot	0.02
	(0.14)
R^2	0.03

The table reports estimated coefficients of equation (3.1) and t-stats (in parentheses) for $n = 1$. The dependent variable is the 2-day changes of de-seasonalized average bid-ask spreads for 9 currencies against the Swiss franc (excluding the euro), in pips. The sample and the right-hand variables are the same as in Table 3.2. The t-stats are estimated using the Newey and West (1987) approach with 2 lags.

Table 3.7: Different types of verbal interventions.

	1m	3m	6m	9m	12m
Uncertainty regression					
Verbal interv. non-MPA	−0.12	−0.13	−0.09	−0.08	−0.09
	(−1.14)	(−1.45)	(−1.56)	(−1.67)	(−1.86)
Verbal interv. MPA	−0.43	−0.15	−0.17	−0.20	−0.21
	(−2.63)	(−2.44)	(−3.20)	(−2.52)	(−2.07)
MPA (press release)	−0.34	−0.24	−0.24	−0.23	−0.25
	(−2.02)	(−2.09)	(−2.53)	(−2.87)	(−3.17)
Skewness regression					
Verbal interv. non-MPA	0.02	0.08	0.08	0.07	0.08
	(0.63)	(1.43)	(1.31)	(1.57)	(1.80)
Verbal interv. MPA	0.01	0.07	0.14	0.17	0.16
	(0.08)	(1.49)	(2.12)	(2.44)	(2.39)
MPA (press release)	0.14	0.05	0.05	−0.00	0.00
	(1.59)	(1.48)	(1.31)	(−0.06)	(0.08)

The table reports estimated coefficients of equation (3.1) and t-stats (in parentheses) for $n = 1$. The sample and structure is the same as Tables 3.2 and 3.5, except that SNB's verbal interventions are divided into two seperate dummies: speeches presenting monetary policy assessments (MPA) and on other days (non-MPA). All speeches contained the keywords "utmost determination" and "unlimited quantities." For brevity only the results for the SNB dummies are presented.

Table 3.8: Robustness check. The SNB's verbal interventions and an alternative measure of market uncertainty and skewness.

	1m	3m	6m	9m	12m
Uncertainty regression	−0.18 (−2.02)	−0.12 (−1.67)	−0.12 (−2.35)	−0.12 (−2.71)	−0.13 (−2.66)
Skewness regression	0.01 (0.34)	0.04 (1.32)	0.05 (1.69)	0.05 (2.20)	0.06 (2.22)

The table reports estimated coefficients of equation (3.1) and t-stats (in parentheses) for $n = 1$. The sample and structure is the same as Tables 3.2 and 3.5, except that the dependent variable is the model free volatility and skewness, respectively. For brevity only the results for the SNB verbal intervention are presented.

Figure 3.6: The dynamic effect of SNB's verbal interventions on skewness: an event study.

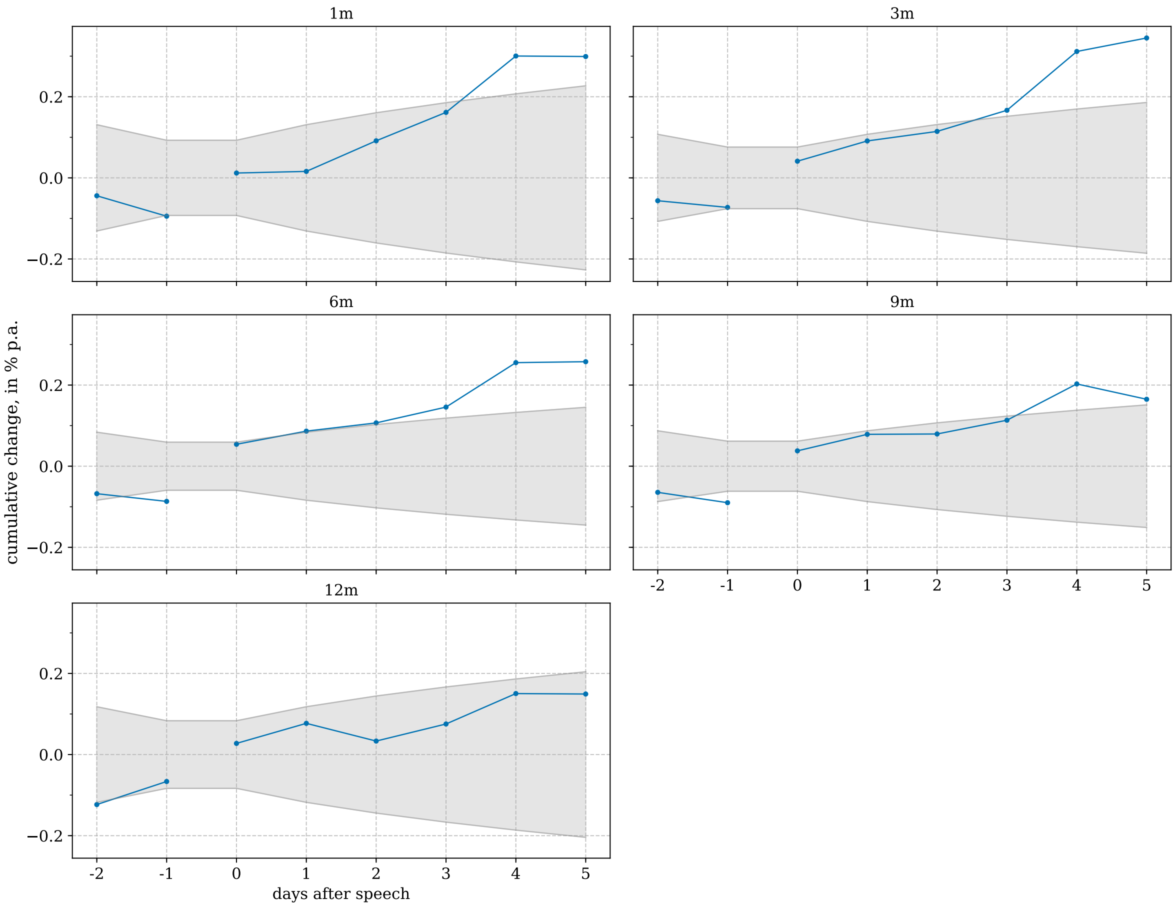

The figure shows the cumulative average difference in skewness since the speech day (day 0) and the 95% confidence band around zero, for the 1-, 3-, 6-, 9- and 12-month horizon. The skewness at any given day is defined as the average risk-reversal implied volatility observed on that day. The confidence interval is calculated assuming cross-sectional and intertemporal independence of differences in skewness.

Figure 3.7: The effect of SNB's verbal interventions on uncertainty: separate effects for each date.

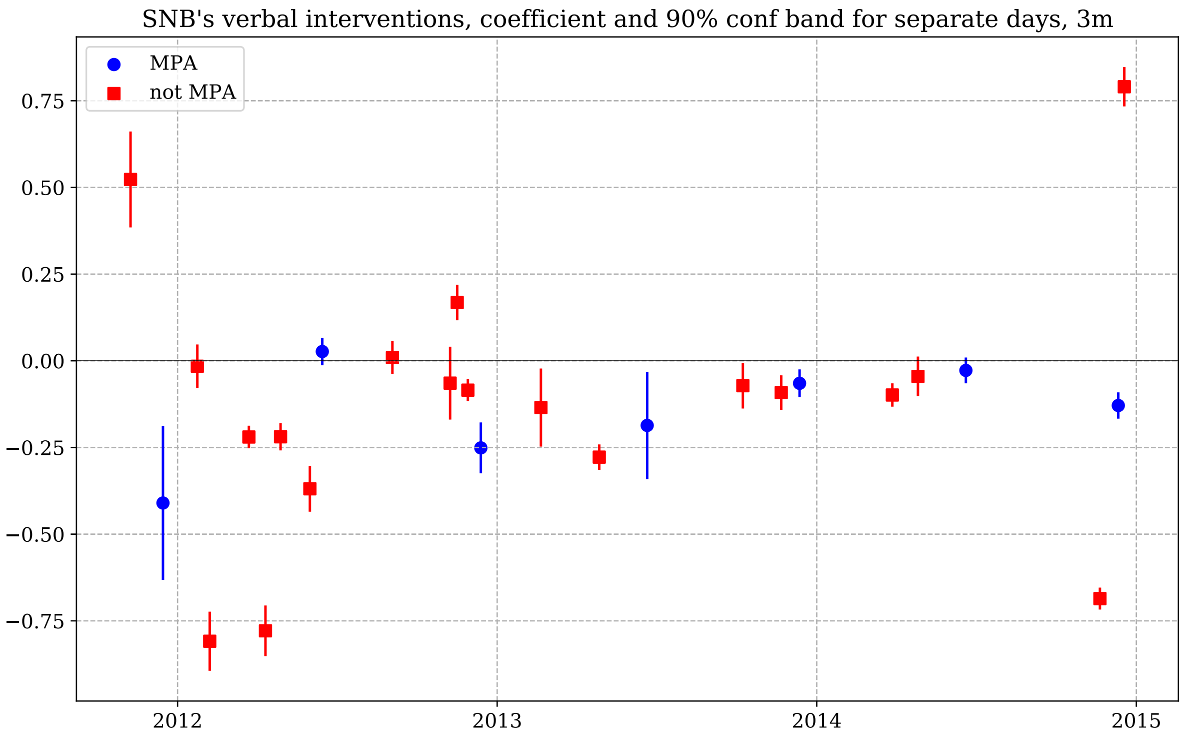

The figure shows the regression coefficient with a 90% confidence band on SNB's verbal interventions. The dependent variable is 2-day change of uncertainty for the 3 month horizon. The regressors and the sample are the same as in Table 3.2, except that we add a dummy variable for intervention i. The figure illustrates the results from 26 regressions, setting $i = 1$, then $i = 2$, etc.

Chapter 4

Overnight Index Swap Rates as Forecasts of Monetary Policy

Dmitry Borisenko* and Igor Pozdeev†

Abstract

Despite growing popularity of overnight index swaps (OIS), which are used by practitioners and regulators as a tool of measuring markets' expectations of future monetary policy, these instruments have gained limited attention from academics. In this paper, we document OIS to be unbiased predictors of future short rates in developed economies, bearing no significant risk premium for maturities up to one year. We show that the OIS underlying overnight rates accurately reflect the target rates set by central banks, making the swaps capable of accurately forecasting the future course of monetary policy. We extract the implied future target rates from the OIS prices to predict the outcome of monetary policy meetings around the world. In the US, a randomly selected triplet of a target rate hike, cut and no-change is correctly classified using the OIS-implied rates in 99.9 and 98 percent of cases five and ten days before a FOMC announcement respectively, which exceeds the prediction accuracy of the federal funds futures-implied rates. We report similarly high prediction accuracy for other developed countries.

*University of St. Gallen, *E-mail*: dmitry.borisenko@student.unisg.ch

†Swiss Institute of Banking and Finance, University of St. Gallen, *E-mail*: igor.pozdeev@unisg.ch

4.1 Introduction

Gauging the expected path of the short-term rate is crucial for financial market participants and regulators. The former use the (expected) risk-free rate as a key element of derivatives pricing, while the latter need to assess the market expectations about the future course of monetary policy. Federal funds futures contracts[1] have become a staple for predicting the future path of the short rate in the US among both practitioners and academics. However, no similar exchange-traded instrument is available in other major economies. Recently, overnight index swaps (OIS)[2] have become an important tool to take a view on the outcome of monetary policy meetings and to value securities using OIS rates as a risk-free benchmark. The question whether OIS rates provide an accurate measure of the expected future short rate has not been thoroughly researched.

Introduced to the markets worldwide in the early 2000s, OIS have enjoyed growing notional values and high degree of liquidity. In 2016 OIS comprised 50, 40 and 30 percent of the total interest rate swap (IRS) turnover in the Australian dollar, US dollar, and the euro (Kreicher, McCauley, and Wooldridge (2017)), with the average daily total IRS turnover in this year being 105, 898, and 445 USD billion respectively, up from 4, 100, and 173 USD billion in 2001 (BIS (2016b)). Similarly to the federal funds futures, the bulk of OIS liquidity is concentrated in maturities of less than one year (Fleming et al. (2012), Ehlers and Eren (2016)). Furthermore, since the 2007-2009 financial crisis reliability of LIBOR as a "risk-free" benchmark rate has been repeatedly questioned by practitioners and regulators. Indicative nature of quotes, absence of active underlying markets and thus proneness to manipulation, led the rate to be phased out in 2021. With the imminent demise of LIBOR, investment community is looking for an alternative benchmark. Hull and White (2013) provide formal arguments favoring OIS rates over LIBOR as risk-free benchmarks for discounting derivatives. In April 2017 the Bank of England recommended SONIA (the underlying rate for the sterling OIS) as the near risk-free reference rate benchmark. International Swaps and Derivatives Association started to promote OIS rates as such benchmarks by advancing standardization of collateral and margin requirements for the swaps, in particular by proposing to use cash as the sole eligible collateral for variation margin, thus bringing the OIS closer to their exchange-traded counterparts in terms of credit risk (ISDA (2013)).[3] The standardization in interest rate OTC derivatives, on the other hand, is furthered by authorities and regulators seeking to reduce systemic risks and thus promoting central counterparty (CCP) clearing of OTC

[1] The federal funds futures contracts are traded at the Chicago Mercantile Exchange. A month M federal funds futures contract settles at the average effective federal funds rate over the month M.

[2] An OIS is an over-the-counter interest rate swap where the floating leg pays the underlying overnight rate compounded over lifetime of the contract, for example in case of the US the underlying is the effective federal funds rate.

[3] Sundaresan, Wang, and Yang (2017) provide an extensive discussion of reasons behind the low credit risk and liquidity premium in OIS contracts. They also provide plethora of examples of OIS rates adopted for discounting among clearing houses and derivative trading venues.

derivatives. Wooldridge (2016) estimates the total amount of interest rate derivatives to be cleared via CCP to have increased from 40 to 85 percent from 2004 to 2016, with around 80 percent of all interest rate swaps cleared centrally in 2016. Despite their increasing likeliness to exchange-traded interest rate instruments and rising popularity among practitioners and regulators, OIS have gained little attention from academics.

In this paper, we show that OIS rates are unbiased predictors of future short rates in major economies. In particular, we demonstrate that for the most liquid maturities from one month two one year, OIS bear neither statistically nor economically significant risk-premium with a few exceptions: for the period from early 2000s to June 2017 excess return on one-year OIS contracts ranges from 3 basis points for Australia (statistically insignificant) to 26 basis points for Switzerland (significant at 1% level), with the USD one-year OIS having statistically insignificant excess return of 20 basis points. We provide empirical evidence that the OIS underlying overnight rates closely follow the interest rate targets set by central banks. Thus, taking into account the absence of risk premium, we argue that OIS are accurate predictors of the future course of monetary policy. We further exploit this predictability to extract the implied future target rates before monetary policy meetings. In a statistical classification framework we show that the implied policy rates extracted from the OIS contracts can accurately predict the outcome of monetary policy meetings around the world. So for instance, in the US a randomly selected triplet of target rate hike, cut, and no-change is correctly inferred from OIS rates in 99.9 and 98 percent of the cases five and ten days before a FOMC announcement respectively. The federal funds futures-implied rates give the prediction accuracy of 99.8 and 94.5 percent for the same horizons, and the LIBOR-implied rates correctly infer the direction of the target rate change only in 43 and 36 percent of the cases. We report similarly high prediction accuracy for the other developed countries.

Our study extends the literature on measuring monetary policy expectations from market prices. Market-based predictions are found to be superior both to those produced by macroeconomic models such as Taylor rules (Evans (1998)) and time-series models (Gürkaynak, Sack, and Swanson (2007)). Gürkaynak, Sack, and Swanson (2007) also demonstrate that federal funds futures dominate other market instruments in terms of predictive power. Federal funds futures have been a popular tool among academics and practitioners to gauge the expected short rate path,[4] however, Piazzesi and Swanson (2008) demonstrate that fed funds futures contain significant time-varying risk premium and thus need a proper adjustment to be unbiased forecasts of the future short rate. We contribute to this strand of literature by taking the monetary policy prediction exercise to the international setting and examining eight major economies whose central banks employ interest rates targeting as monetary policy tool. First, given the above

[4]Krueger and Kuttner (1996), Sack (2004), Kuttner (2001), Bernanke and Kuttner (2005), Piazzesi and Swanson (2008) are a few academic studies. The CME FedWatch Tool uses the futures contracts to compute the probability of target rate changes in the upcoming FOMC meetings.

mentioned negligible risk premium in the OIS rates, we argue that they are good predictors of future *underlying* rates. Second, we link the OIS underlying and *target* rates of respective central banks,[5] by showing that the former closely follow the latter. Hence, we find OIS rates to be an adequate tool for recovery of monetary policy expectations.

Our second contribution is to provide a convenient venue for further research on behavior of asset prices around target rate decisions. There is growing literature documenting abnormal dynamics of asset returns around monetary policy announcements. Lucca and Moench (2015) and Mueller, Tahbaz-Salehi, and Vedolin (2017) document significant excess returns on the US large cap stocks and portfolio of currencies against USD before the FOMC announcements, Cieslak, Morse, and Vissing-Jorgensen (2019) find cyclicality in the US stock returns with the cycle spanned by the FOMC meetings. Karnaukh (2016) further incorporates monetary policy expectations derived from federal funds futures showing that the US dollar appreciates (depreciates) against a portfolio of currencies several days in advance of anticipated US target rate hikes (cuts). Her analysis thus recognizes predicting the outcome of monetary policy meetings as a classification problem with a discrete set of outcomes.[6] Building on our findings and using a statistical classification framework with a threshold-based rule, we demonstrate that the OIS-implied expected rates are capable of accurately predicting the direction of target rate changes in developed countries, with predictive power being on par with that of federal funds futures in the US.

The remainder of the paper is structured as follows. Section 4.2 describes our dataset, the payoff structure of OIS contracts and computation of excess returns, as well as provides empirical evidence of OIS bearing no significant risk premium in the major economies. In this section we also demonstrate that the overnight rates underlying the OIS are tightly linked to the target rates set by central banks, thus providing evidence of OIS being accurate predictors of the future course of monetary policy. Section 4.3 presents the methodology of recovering the expected target rate changes from the prices of OIS, takes a closer look at the required assumptions, and assesses the forecasting power of the implied rates in predicting the outcomes of upcoming target rate decisions. Section 4.4 concludes.

[5]With an exception of Australia and New Zealand where target rates are held constant between monetary policy meetings and overnight rates correspond exactly to the target rates, the target rates are either allowed to float within a band like the federal funds rate in the US, or not explicitly tied to OIS underlying rates at all, with an extreme example being Switzerland whose monetary authority targets the 3-month Swiss franc LIBOR.

[6]Mueller, Tahbaz-Salehi, and Vedolin (2017) also report diverging patterns in currency returns against USD immediately after the FOMC announcements resulting in target rate changes – dollar appreciates after hikes, and depreciates otherwise.

4.2 Excess returns on overnight index swaps

An overnight index swap (OIS) is a contract requiring one party to pay the other the accumulated underlying rate – most often an overnight rate such as the federal funds rate in the US – in exchange for a predetermined fixed rate payment. OIS are priced based on the expectation of the future rate and have zero value at inception, hence the expected return from entering into such a contract can be thought of as a risk premium, which we find to be statistically and economically insignificant for most OIS specifications in our sample. The OIS underlying rates often coincide with the target rates in respective countries; when they do not, we show that the two are tightly linked. When indexed with the target rates instead of the true underlying rates, the prices of OIS are still found to be unbiased predictors of the future compound rate levels.

4.2.1 Overnight index swaps

In an OIS, one party agrees to pay a fixed rate, and the other a floating short rate accumulated over the lifetime of the contract, the payments being exchanged at maturity of the swap.[7] Each day, a new portion of the floating leg return is accumulated based on the underlying rate level which is determined according to the local market conventions – this process is called "fixing" – while the fixed leg portion is a constant negotiated at inception. By the risk-neutral pricing argument, the fixed leg rate of an OIS of maturity m negotiated at time t is determined as follows:

$$W_{m,t} = E_t \prod_{s \in S_{m,t}} (1 + r_s \delta_s), \tag{4.1}$$
$$W_{m,t} := 1 + w_{m,t} D_m,$$

where the expectation is taken under the risk-neutral measure, $w_{m,t}$ is the annualized fixed rate, D_m is the contract-specific day count factor such as Act/360 for USD OIS, $S_{m,t}$ is the set of all fixing days (usually banking days in the respective country) until the swap matures, r_s is the per-day fixing rate on day s, and δ_s is the "length of overnight", or one plus the number of bank close days following day s. For example, if s is Friday and the next Monday is not a bank holiday, $\delta_s = 3$; if the next Monday is a bank holiday, but the Tuesday is not, $\delta_s = 4$ and so on. For two OIS contracts on different currencies traded on the same day, S_m and D_m will be different. The differences come from the fixing and trading conventions of the underlying rates such as the holiday calendar, the effective date and the fixing lag. For example, for a sterling OIS traded on October 15th the first fixing occurs on the same day (effective date equal to quote date), while for a USD

[7] The OIS are subject to netting, that is notional is not exchanged, furthermore at the maturity the difference between the two legs is transferred to the counterparty with the positive payoff.

OIS it occurs in two business days (effective date two days away from quote date). This means that the one-month sterling OIS has the effective lifetime from October 15th to November 15th, while its dollar counterpart starts accruing the floating rate on October 17th and matures on November 17th (provided that October 16th and 17th are banking days in the US). The fixing lag refers to the number of days between the fixing day of the floating leg an the value date of the underlying. For instance, the effective federal funds rate for a given day is published by the Federal Reserve Bank of New York on the following day, on the other hand the underlying rates for the Swiss franc and Swedish krona are tomorrow/next rates, hence the rate from the previous day is used for fixing.

We collect the fixed leg rates and the corresponding underlying rates of OIS contracts written on currencies of the following countries: Australia, Canada, Eurozone, Japan, New Zealand, Sweden, Switzerland, United Kingdom, and United States. For each currency we use one-, three-, six-, nine- and twelve-month tenors. All data comes from Bloomberg. Table 4.1 summarizes specifications of the contracts and data availability. For each currency the table reports the underlying rate of the floating leg;[8] number of business days before the first floating leg indexing, fixing lag, day count convention, and the date since which all tenors are available.[9] Our sample ends in June 2017.

Table 4.1: OIS contracts summary

Currency	Underlying rate	Floating indexing start (days)	Fixing lag (days)	Day count	Data since
AUD	AONIA	1	0	Act/365	Oct-2001
CAD	CORRA	0	1	Act/365	May-2002
CHF	TOIS	2	-1	Act/360	Aug-2000
EUR	EONIA	2	0	Act/360	Jan-2000
GBP	SONIA	0	0	Act/365	Dec-2000
JPY	TONA	2	1	Act/365	Mar-2002
NZD	NZIONA	2	0	Act/365	Sep-2002
SEK	STIBOR T/N	2	-1	Act/360	Aug-2004
USD	Effective federal funds rate	2	1	Act/360	Dec-2001

This table summarizes the structure of OIS contracts. For each currency in the first column, column two contains the underlying rate, column three shows number of business days after trade date until the first floating leg indexing happens. Column four contains the fixing lag, meaning that each day of a contract's life its floating leg is indexed by todays' value of the underlying if the fixing lag is zero, by yesterday's value if it is -1, and by tomorrows' value if it is 1. Column five reports day count convention, and column six reports data availability for all tenors considered in this paper (except for the Canadian dollar where the one-year tenor is available from April 2003).

[8]For Australia and New Zealand the underlying rates are equal to the Interbank Overnight Cash Rate and the Official Cash Rate respectively. Both rates are effectively equivalent to the corresponding central banks' operational targets. AONIA and NZIONA are the names used in the ISDA Definitions.

[9]With a single exception of the Canadian dollar for which the one-year tenor is available from April 2003.

4.2.2 OIS risk premium

In order to be unbiased predictors of the cumulative future underlying rate, the expected return of either OIS counterparty should be zero, that is, OIS should not contain a risk premium. We can test this assumption by calculating the realized return of an OIS fixed rate receiver:

$$rx_{t+m} = W_{m,t} - \prod_{s \in S_{m,t}} (1 + r_s \delta_s), \tag{4.2}$$

and estimating its expected value. Under certain conditions, this expectation can be approximated as the time series average of rx_t. Since neither counterparty pays anything to enter the contract, the realized return is a zero-cost portfolio, so we further refer to this return as the excess return, or risk premium.

Table 4.2 reports the sample estimates of the average excess returns on different OIS, for maturities from one month to one year. For each contract the sample starts as defined in Table 4.1 and ends in June 2017. For each currency (in columns) and for each maturity (in rows) we compute floating leg return according to the contract specification in Table 4.1, and then subtract it from the fixed leg rate. The excess returns are in basis points p.a. The numbers in parentheses show t-statistics, calculated using Newey and West (1987) standard errors with the number of lags determined according to Newey and West (1994).

At the one-month maturity, the return is less than one basis point p.a. on average, with the single exception being the one-month CHF OIS, whose average excess return is 1.53 basis points. The returns are economically insignificant, and the apparent statistical significance of the returns on the one-month AUD swap is misleading: the underlying rate for this OIS is effectively the Cash Rate set by the Reserve Bank of Australia that is held constant for prolonged periods of time, which results in the low variance of the underlying, the low standard error of the mean and consequently the inflated t-statistic. For maturities higher than one month, the excess return is most pronounced for the Swiss franc swaps: from 5.88 basis points for the maturity of 3 months to 26.56 basis points for the twelve months maturity. For the US dollar swaps, the excess returns are several times smaller than those documented by Piazzesi and Swanson (2008) for federal funds futures over 1984-2005. For example, they report "unannualized"[10] excess returns of 2.9, 10.5, and 30.7 basis points for one-, three-, and six-months futures, all statistically significant at 1% level. Standard deviation of OIS excess returns ranges from less than two basis points for Japan to 18.7 basis points for Switzerland at one month maturity and increases with maturity of the swaps for all currencies to as much as 60.4 basis points p.a. for a one-year NZD swap. In general, the distribution of excess returns tends to be rather tight

[10]It is difficult to directly compare excess returns on federal funds futures and OIS beyond the one-month maturity, since payoff of the former for any maturity is determined by annualized effective federal funds rate as reported by the Federal Reserve Bank of New York. A closer counterpart would be one-month OIS with forward start, which lie beyond the scope of this paper.

Table 4.2: Excess returns of an OIS fixed rate receiver.

		AUD	CAD	CHF	EUR	GBP	JPY	NZD	SEK	USD
1M	mean	−0.92	−0.07	1.53	0.72	0.83	0.28	0.27	0.54	0.93
	(t)	(−2.44)	(−0.27)	(0.98)	(0.69)	(0.42)	(1.54)	(0.94)	(0.44)	(0.94)
	std.	5.28	5.08	18.70	7.99	12.09	1.92	4.81	10.11	7.44
	5%	−8.85	−6.06	−11.25	−10.07	−15.41	−1.51	−5.07	−8.49	−4.05
	50%	−0.54	−0.07	−0.41	0.29	−0.14	0.06	0.00	0.32	0.03
	95%	5.75	4.45	18.73	13.12	21.68	2.69	7.90	10.78	5.83
3M	mean	−2.55	0.72	5.88	2.83	3.11	0.66	1.52	2.39	3.21
	(t)	(−1.20)	(0.80)	(1.96)	(1.03)	(0.49)	(1.14)	(0.82)	(1.10)	(0.75)
	std.	13.81	7.83	27.86	12.21	18.58	3.30	12.14	16.32	14.51
	5%	−24.04	−10.11	−11.63	−9.83	−10.48	−2.57	−12.68	−11.83	−5.94
	50%	−1.62	−0.28	1.06	0.57	−0.09	0.17	−0.02	0.11	0.45
	95%	13.42	14.03	34.84	23.96	21.18	6.31	20.27	19.95	20.45
6M	mean	−2.75	4.63	13.19	7.18	7.67	1.22	4.65	6.36	7.56
	(t)	(−1.28)	(3.29)	(6.67)	(0.85)	(0.50)	(0.98)	(0.60)	(1.01)	(0.79)
	std.	28.99	17.90	40.94	23.30	33.32	5.47	27.92	28.55	25.17
	5%	−39.09	−16.96	−13.23	−16.96	−12.92	−4.93	−24.70	−19.58	−13.53
	50%	−3.73	0.26	3.67	2.78	0.49	0.26	0.42	1.26	1.60
	95%	26.91	39.80	73.59	44.20	45.89	12.94	43.81	31.88	56.00
9M	mean	−0.79	9.76	20.52	12.91	13.72	2.14	9.41	12.21	12.96
	(t)	(−0.06)	(3.37)	(4.38)	(0.95)	(0.59)	(1.23)	(0.78)	(0.66)	(0.93)
	std.	42.61	27.51	47.12	33.99	46.17	7.53	44.44	41.90	36.06
	5%	−51.14	−21.64	−16.38	−18.72	−20.92	−6.62	−35.97	−27.63	−24.65
	50%	−5.59	3.04	6.62	5.64	3.25	0.57	0.85	3.92	4.25
	95%	49.95	62.13	112.25	63.31	69.76	20.08	86.35	63.02	107.79
1Y	mean	3.19	18.33	26.56	19.87	20.96	3.51	15.60	19.40	19.86
	(t)	(0.17)	(1.69)	(2.82)	(1.10)	(0.75)	(1.32)	(0.51)	(0.73)	(0.96)
	std.	55.17	36.77	49.67	43.28	56.95	9.18	60.43	53.07	48.22
	5%	−57.44	−19.31	−18.81	−21.62	−31.00	−5.06	−47.14	−28.14	−37.51
	50%	−4.79	7.51	11.51	9.72	9.07	1.34	2.52	7.89	7.10
	95%	83.89	90.95	131.89	85.36	107.42	26.54	130.50	100.94	144.65

The table presents descriptive statistics of returns on receiving the fixed and paying the floating leg of an OIS on currency i (columns) with maturity m (rows), ranging from one month to one year. For each day when the fixed leg rate of an OIS is available, we establish the lifetime of the swap and calculate its future floating leg return, depending on particular contract specifications for each currency (see Table 4.1). We then subtract the realized floating leg return from that of the fixed leg, and report the sample average, standard deviation, median, and 5th and 95th percentiles of the empirical distribution in basis points p.a. The t-statistics of means (in parentheses) are calculated using Newey and West (1987) standard errors with the number of lags determined according to Newey and West (1994). The sample start for each currency is reported in Table 4.1. The sample ends in June 2017.

at low maturities. [11]

Overall, although the excess returns increase on average with maturity they are either economically or statistically insignificant (or both), with an exception of the Swiss franc swaps.

[11] For the Sterling OIS of short maturities the relatively high standard deviation estimate is almost entirely due to the pre-March 14, 2005 period. On this day the Bank of England narrowed the rate corridor on its deposit and lending facilities from +/-100 to +/-25 basis points around the Bank Rate, thus significantly reducing the volatility in the SONIA rate.

4.2.3 OIS underlying rates and central bank policy rates

The underlying rates of OIS contracts on a number of currencies such as the Australian, New Zealand and US dollars coincide with the target rates of the respective central banks. Given the insignificant risk premium documented in the previous section, a natural question arises whether OIS rates can provide an insight into the market expectations of the future course of monetary policy in different countries. Under the "future course of monetary policy" we understand the *target rate compounded overnight* over the lifetime of the swap, akin to the OIS floating leg return, and not the *level of the target rate.*

We start with a simple conjecture: if OIS rates are unbiased predictors of future compound target rates, there should be no difference between indexing the floating leg with the real underlying rate or the central bank target rate. We collect data on target rates for all countries in Table 4.1 except for Japan.[12] The data comes from the central banks' web pages. For each currency, the first three columns in Table 4.3 report the corresponding monetary authority and the target interest rate. In the countries where the central bank targets a range such as the US or Switzerland, we set the corresponding target rate to the band midpoint.

Table 4.3: Central Banks' Policy Meetings Summary

	Monetary Authority	Target Rate	Events	Hikes	Cuts
			(unscheduled)		
AUD	Reserve Bank of Australia	Cash Rate	173 (0)	19 (0)	20 (0)
CAD	Bank of Canada	Target for the Overnight Rate	122 (1)	17 (0)	17 (1)
CHF	Swiss National Bank	3-month CHF LIBOR	79 (11)	10 (1)	15 (10)
EUR	European Central Bank	Main Refinancing Operation Rate	225 (1)	17 (0)	22 (1)
GBP	Bank of England	Bank Rate	197 (0)	10 (0)	20 (0)
NZD	Reserve Bank of New Zealand	Official Cash Rate	117 (0)	19 (0)	18 (0)
SEK	Riksbank	Repo Rate	82 (2)	20 (0)	18 (2)
USD	Federal Reserve System	Federal Funds Rate	130 (5)	21 (0)	13 (2)

This table summarizes target rate announcements across currencies. The first three columns report currencies, corresponding monetary authorities and the target rates. The last three columns contain the total number of monetary policy meetings, and the numbers of hikes and cuts, with values inside parentheses showing meetings that occurred outside meeting calendars. For each currency the sample starts once the OIS data is available (see Table 4.1) and ends in June 2017.

Panel A of Table 4.4 reports the excess returns (in basis points p.a.) of receiving the fixed and paying the floating leg of an OIS on different currencies (in columns) for maturities

[12] We exclude Japan since the Bank of Japan has been using various monetary policy instruments over the past two decades. So for instance, on March 19th 2001 the Bank of Japan changed its operating target from the overnight call rate (MUTAN) to current accounts of Japanese banks at the Bank. The interest rate targeting was restored for the period from March 2006 to April 2013, with average rate being targeted until October 2010, and a band being set thereafter. On April 2013 the Bank of Japan began targeting monetary base. In 2016 the Bank introduced negative interest rates on the current accounts and started to implement the "yield curve control". Also see Kuttner (2014) for a thorough overview of Japan's monetary policy from 1980 to 2012.

ranging from one month to one year (in rows), with the floating leg being indexed by the target rate set by the monetary authority of the corresponding country. The average excess return is computed for each day when the fixed leg rate of an OIS is available, by first, determining the lifetime of the contract according to the contract specification of each currency (see Table 4.1), second, subtracting the realized floating leg return from that of the fixed leg, and third, taking the sample average. Panel B reports the average difference between the excess returns on the OIS with overnight- and the target rate-indexed floating legs, or effectively the difference between the floating leg return indexed by the target and overnight rates. The numbers in parentheses are the t-statistics calculated using the HAC estimator of Newey and West (1987) with optimal number of lags according to Newey and West (1994). For each currency the sample starts with availability of the corresponding OIS data as specified in Table 4.1 and ends in June 2017.

For the Australian, Canadian, and US dollars the risk premium estimates of target rate-indexed floating leg OIS deviate from their overnight rate-indexed counterparts by a fraction of a basis point for any maturity, and for the New Zealand dollar[13] the premium is exactly the same. For the euro, British pound, Swedish krona, and Swiss franc the discrepancies between the compound overnight and target rates indexing are more pronounced, but either economically or statistically insignificant, except for the Swiss franc.

What might be the cause of the discrepancies? For the first three of the four latter currencies the target rates are repo rates, whereas the OIS underlying rates are unsecured money market rates, so discrepancies are expected due to the credit and liquidity risks and/or limited access of market participants to the unsecured credit market, accessible by the banks domiciled in the corresponding currency area only, in addition to other institutional characteristics. For example, the negative difference between the target and overnight rates for Swedish krona in Panel B of Table 4.4 means that, on average, the STIBOR compounded over maturities from one month to one year exceeded Riksbank's repo rate by around eight and a half basis points.[14] Similarly, SONIA was below the Bank of England's bank rate by four and a half basis points.[15] For the euro the positive

[13]We find a single data discrepancy for New Zealand: Bloomberg reports the jump in the underlying rate on April 28th 2004, while the Reserve Bank of New Zealand reports the change in its official cash rate on April 29th. We choose the latter data point over the former.

[14]Although being able to keep the overnight rate stable at the target rate, the Riksbank has no direct way to steer the tomorrow/next rate underlying the SEK OIS. As a result the kronor tomorrow/next market experienced periods of elevated volatility in 2008 and during the phase-out of extraordinary monetary policy measures in 2010 (Riksbank (2014)). Which together with constant spread of 10 basis points over the target rate results in the OIS forecasts being contaminated both with constant and time-varying risk premiums.

[15]The Bank of England's Sterling Monetary Framework explicitly states its objective to keep both unsecured and secured market rates in line with the Bank Rate. BoE (2017) points out that the persistent gap of four to five basis points between SONIA and the Bank Rate over the recent years is likely due to lenders without access to reserves accounts at BoE willing to accept lower rates, and unwillingness of

Table 4.4: OIS underlying rates vs. central bank target rates.

	AUD	CAD	CHF	EUR	GBP	NZD	SEK	USD
	Panel A: OIS excess returns, floating leg indexed by target rates (t-statistics)							
1M	−0.95	−0.01	−11.46	−15.61	−3.77	0.28	8.59	0.60
	(−2.52)	(−0.05)	(−4.60)	(−1.33)	(−5.09)	(0.94)	(5.57)	(0.94)
3M	−2.57	0.74	−8.64	−13.56	−1.41	1.52	10.68	2.85
	(−1.21)	(0.81)	(−1.26)	(−1.03)	(−0.41)	(0.82)	(4.26)	(1.67)
6M	−2.77	4.63	−3.80	−9.32	3.25	4.65	14.86	7.16
	(−1.29)	(3.26)	(−1.93)	(−0.59)	(0.29)	(0.60)	(2.24)	(1.09)
9M	−0.81	9.74	2.17	−3.71	9.40	9.41	20.99	12.53
	(−0.06)	(3.34)	(0.63)	(−0.21)	(0.48)	(0.78)	(1.10)	(1.05)
1Y	3.18	18.48	9.00	3.10	16.75	15.60	28.33	19.39
	(0.17)	(1.72)	(0.78)	(0.15)	(0.68)	(0.51)	(1.05)	(1.09)
	Panel B: Overnihgt vs. target rate floating leg indexing (t-statistics)							
1M	0.03	−0.05	12.86	16.21	4.61	0.00	−8.05	0.33
	(1.26)	(−0.06)	(2.68)	(0.44)	(1.40)		(−0.78)	(0.04)
3M	0.03	−0.02	14.45	16.27	4.52	0.00	−8.29	0.36
	(1.95)	(−0.03)	(3.63)	(0.15)	(1.01)		(−0.13)	(0.02)
6M	0.02	−0.01	17.07	16.39	4.41	0.00	−8.50	0.39
	(4.98)	(−0.01)	(4.40)	(0.22)	(0.66)		(−0.14)	(0.01)
9M	0.02	0.01	18.35	16.49	4.32	0.00	−8.78	0.43
	(5.12)	(0.01)	(5.03)	(0.22)	(0.96)		(−0.45)	(0.01)
1Y	0.01	−0.16	17.56	16.64	4.21	0.00	−8.94	0.47
	(5.31)	(−0.08)	(4.61)	(0.22)	(1.00)		(−0.62)	(0.01)
	Panel C: Regression of overnight on target rates							
α	−0.02	0.47	−4.64	−38.31	−5.18	0.00	1.79	−0.18
(s.e.)	(0.02)	(0.19)	(2.33)	(2.96)	(0.57)		(0.75)	(0.76)
β	1.00	1.00	0.89	1.12	1.01	1.00	1.05	1.00
(s.e.)	(0.00)	(0.00)	(0.02)	(0.01)	(0.00)		(0.01)	(0.00)
$\chi^2(2)$	7.68	7.32	58.60	183.82	98.93		232.24	1.37
(p-value)	(0.02)	(0.03)	(0.00)	(0.00)	(0.00)		(0.00)	(0.50)

Panel A of this table presents the average excess returns, in basis points p.a., of receiving the fixed and paying the floating leg of an OIS on currencies (in columns) for maturities ranging from one month to one year (in rows), with the floating leg being indexed by the target rate set by the monetary authority of the corresponding country. The average excess return is computed for each day when the fixed leg rate of an OIS is available, by first, determining the lifetime of the contract according to the contract specification of each currency (see Table 4.1), second, subtracting the realized floating leg return from the fixed leg return, and third, taking the sample average. Panel B reports the average difference between the excess returns on the OIS with overnight- and target rate-indexed floating legs, or effectively the difference between the floating leg return indexed by the target and overnight rates. Panel C reports the intercept (in basis points p.a.) and slope estimates from the regression of overnight rates on target rates $r_{on,t} = \alpha + \beta r_{tgt,t} + \varepsilon_t$. The last two rows of Panel C report χ^2 statistic and p-value for the joint test of zero intercept and unit slope. The numbers in parentheses are the t-statistics (Panels A and B), standard errors and p-values (Panel C), all calculated using the HAC estimator of Newey and West (1987) with optimal number of lags according to Newey and West (1994). For each currency the sample starts with availability of the corresponding OIS data as specified in Table 4.1 and ends in June 2017.

difference between the average compounded rates is caused by the fact that since 2009 the overnight unsecured rate was stuck at the level of ECB's deposit rate[16] – before that, it closely followed the refinancing rate which we use as the target rate in this study, and was, on average, slightly above the refinancing rate similarly to Sweden.[17,18] The Swiss franc is the notable exception among other currencies. The Swiss National Bank sets a target range for the three month CHF LIBOR without an explicit concern for the overnight money market. Besides the maturity mismatch between the CHF OIS underlying and the central bank's target, SNB explicitly stated its desire to keep the LIBOR rate at the lower bottom of the target range after 2009, thus pushing the unsecured overnight lending rate further downwards, which results in the observed significant positive difference. Nevertheless, the CHF OIS fixed rate predicts the compound future SNB target rate rather accurately slightly underestimating it for lower forecast horizons which is justified by the LIBOR-OIS spread.

To sum up, OIS fixed rates capture the future target rates of the corresponding central banks, being more accurate predictors of the future course of monetary policy when central banks directly target overnight money markets. Yet, as we discuss above, when the target and market rates differ, the institutional characteristics of the monetary policy framework and/or risk premia limit the OIS ability to produce unbiased forecasts of the future compounded target rate. Although the discrepancies between the underlying and target rates are *unconditionally* small in magnitude, do not vary much across swap maturities, and rarely statistically significant, the necessary adjustments should be made in practical applications on country-by-country basis.

Given that the OIS prices themselves bear no significant risk premium, an alternative approach to demonstrate that OIS rates predict future target rates is to show that the underlying rates of the former closely follow the latter. The right hand side of the OIS pricing equation (4.1) implies that if these rates are same, the OIS fixed rates are also unbiased forecasts of the compound target rate, which is the exact case for New Zealand. On the other hand if the underlying rate fluctuates around the target (as in the US), being the same on average and comoving with the target rate in lockstep – that is, if

reserves account holders to arbitrage the difference away due to higher leverage to be reported.

[16] Main refinancing operations refer to (reverse) REPO transactions through which the ECB provides the bulk of liquidity to the banks of the Eurosystem. Although the refinancing operations normally have maturity of one week, the main refinancing rate is effectively capped (floored) by the overnight rates on the standing lending (deposit) facilities.

[17] In October 2008 the ECB introduced the fixed-rate full allotment policy effectively providing unlimited amounts of liquidity against eligible collateral at the main refinancing rate, the subsequent expansion of excess reserves, that are remunerated at the rate on the deposit facility, drove the market rates toward the rate on the facility.

[18] In contrast to the ECB, the Riksbank uses the same repo rate not only to provide liquidity to the banking system via repo but also to withdraw it by issuing Riksbank certificates, typically, with maturity of one week. More importantly, in order to stabilize the overnight market rate around the target the Riksbank offers overnight credit or deposits at the repo rate +/- 10 basis points (also known as the fine-tuning transactions) depending on whether the entire system is in deficit or surplus respectively. See Riksbank (2014) for additional institutional details on the monetary policy implementation in Sweden.

regression of one rate on another produces zero intercept and unit slope coefficients – the result would depend on the variance of the deviations due to the Jensen's term.

Panel C of Table 4.4 reports the intercept (in basis points p.a.) and slope estimates from regression of overnight on target rates $r_{on,t} = \alpha + \beta r_{tgt,t} + \varepsilon_t$. The last two rows of Panel C report χ^2 statistic and the p-value for the joint test of zero intercept and unit slope. Computation of standard errors and test statistics as well as the sample are the same as in the upper panels of the table.

The regression estimates in the Panel C of Table 4.4 further corroborate the results discussed earlier in this section. The slope estimates are one and the intercept estimates are close to zero for the Australian, Canadian, New Zealand and US dollars. For the British pound and Swedish krona the slopes are close to one and intercepts indicate that the overnight rates are slightly higher and lower than the corresponding targets respectively, though in economic terms the differences are negligible. For the euro the estimates reflect the situation where the overnight rate was stuck at the lower bound set by the ECB's deposit rate since 2009 remaining below the main refinancing rate (negative intercept) and not responding to the cuts in this rate in 2013 and 2014, when the deposit rate was set at zero (estimated slope grater than one). Furthermore, each regression in Panel C produces R^2 of over 99%, meaning that the nonlinearity in the OIS payoff is unlikely to have any sizable impact.

Overall, OIS rates are capable of accurately reflecting the future path of monetary policy measured by the central bank target rates compounded over the OIS lifetimes. Moreover, our inference in this section is likely to be conservative in the sense of overrejecting the null of zero risk premium: since our returns are overlapping, Hodrick (1992) standard errors would be more restrictive. However, given the small economic magnitude of the risk premium estimates achieving statistical significance, we deem this exercise to be excessive.

4.3 OIS-implied rates and outcomes of monetary policy meetings

Under certain assumptions, it is possible to extract the underlying rate that the OIS market participants expect to prevail after a certain date within the lifetime of a contract. The assumptions have to do with the Jensen's inequality and the relation between the risk-neutral and natural measures. An additional assumption about the relation between the underlying and the target rate changes allows to extract the expected policy rate several days ahead of monetary policy announcements. These OIS-implied expectations have substantial power to predictively classify upcoming announcements into hikes, cuts and no-change decisions, on par with the federal funds futures in the US.

4.3.1 Recovering expected target rate changes

Consider the risk-neutral pricing formula in equation (4.1) and let the fixing rate of the floating leg be possibly subject to change on a predetermined day $s^* \in S_{m,t}$. We will refer to s^* as the "event" or "announcement".[19] That said, equation (4.1) can be restated as the product of the accrued floating rates before and after s^*:

$$W_{m,t} = E_t \prod_{s<s^*} (1 + r_s \delta_s) \prod_{s \geq s^*} (1 + r_s \delta_s), \tag{4.3}$$

The following assumption will allow to isolate the post-event part of the floating leg rate by swapping (no pun implied) the expectation and product operators.

Assumption 1. *The Jensen's inequality holds as equality with accuracy sufficient for the practical purposes.*

In this case, we can express equation (4.3) as follows:

$$W_{m,t} = \prod_{s<s^*} (1 + \underbrace{E_t r_s}_{\bar{r}_t} \delta_s) \prod_{s \geq s^*} (1 + \underbrace{E_t r_s}_{\bar{r}_t^*} \delta_s), \tag{4.4}$$

where $\bar{r}_t$, $\bar{r}_t^*$ is the short-hand notation for expected pre- and post-announcement rates. The next assumption will allow to retain the post-announcement expected rate as the only unknown variable in (4.4).

Assumption 2. *The expected pre-announcement OIS underlying rate is known.*

A straightforward way to estimate $\bar{r}_t$ is to take the sample average of the OIS underlying rate since the last announcement. If the rate is credibly kept constant by the central bank such as in New Zealand or tends to fluctuate around a constant value from one announcement to another such as in the US, Assumption 2 holds. Now, equation (4.4) can be solved numerically for the post-announcement expected rate $\bar{r}_t^*$. Another assumption is needed to equate this risk-neutral expectation to the expectation under the natural measure.

Assumption 3. *The expectation of the underlying rate taken under the risk-neutral measure coincides with that taken under the natural measure.*

Now, the solution to equation (4.4) is the "real world" expectation of the OIS underlying rate after the upcoming announcement. The final assumption is needed to connect the OIS underlying rates to the central bank target rates.

[19]Most often, the OIS fixing rate changes in response to a target rate change several days after the announcement, because of a delay in the target rate becoming effective, quote and/or fixing lags. For example, a new target rate communicated by the Fed on Wednesday becomes effective on Thursday and is likely to affect the Thursday's effective federal funds rate (EFFR), but because of a 1-day fixing lag, the first time the new rate enters the calculations is Friday. This is known in advance to investors, and, while we honor it in the empirical part, here we neglect it for notational simplicity.

Assumption 4. *The expected post-announcement OIS underlying rate differs from the expected pre-announcement rate by the expected target rate change.*

$$\bar{r}_t^* = \bar{r}_t + E_t \Delta r_{\text{tgt}}, \tag{4.5}$$

where r_{tgt} is the central bank target rate.

In other words, the differences in the expected levels of the pre- and post-announcement OIS underlying rates must equal the expected target rate changes.

Under these assumptions, the difference between the OIS-implied rate and the expected pre-announcement rate equals the market expectation of the target rate change. Several days before each announcement, we recover the implied rate and calculate the bespoke difference:

$$E_t \Delta r_{\text{tgt}} = \bar{r}_t^* - \bar{r}_t. \tag{4.6}$$

4.3.2 Assumptions close-up

Assumption 1 allows us to equate the expectation of the product of (gross) daily interest rates and the product of their expectations. To be able to say anything about its restrictiveness, we run a simulation exercise. In particular, we calculate the following difference, which can be recognized as the Jensen's term:

$$E \prod_{t=1}^{30} (1 + r_t/360) - \prod_{t=1}^{30} (1 + E r_t/360), \tag{4.7}$$

where r_t is conditionally Normally distributed and has the autocorrelation of 0.95, a value often estimated in the data:

$$r_t | \mathcal{F}_{t-1} \sim \mathcal{N}(\mu, \sigma),$$
$$\rho[r_t, r_{t-1}] = 0.95.$$

We let $\mu \in (0, 900)$ basis points p.a. and $\sigma \in (10, 100)$ basis points p.a., which subsumes the actually observed range of rate values in our sample. We proxy the left expectation in equation (4.7) by means of the sample average over 1000 simulations. In the extreme case of a 100 basis points standard deviation in the daily short rate and a mean value of 900 basis points, the Jensen's term is as small as 0.003 basis points p.a. in magnitude and hence can safely be ignored in the applications.

Assumption 3 is essentially equivalent to absence of risk premium in OIS rates. In Section 4.2.2 we demonstrated that the risk premium in the swaps is either economically or statistically insignificant (or both) for all currencies except for the Swiss franc.

Assumption 4 says that the two rates – the OIS underlying and the central bank target – are expected to move one-for-one. In the previous section we reported that the two rates closely correspond to each other in levels and move in lockstep for all currencies except for the euro and Swiss franc. However, we emphasize here that the two rates should jump by an equal amount on the announcement day without systematic errors. We test it by comparing the difference in the average post- and pre-announcement rates to the announced target rate change.

We project the underlying rate changes onto the target rate changes. Denote $\Delta\bar{r}_a = \bar{r}_a^* - \bar{r}_a$, as the change in average OIS underlying rate prevailing after ($\bar{r}_a^*$) and before ($\bar{r}_a$) the target rate announcement a, and $\Delta r_{\text{tgt},a}$ as the announced change in the target rate. If the OIS underlying rates indeed reflect the target rate, then the intercept in the regression below should be zero and the slope coefficient should be one:

$$\Delta\bar{r}_a = \alpha + \beta\Delta r_{\text{tgt},a} + \varepsilon_a, \;\; a \in A, \tag{4.8}$$

where A is the set of all announcements of a central bank. In other words, the underlying and target rates should move in lockstep and there should be no systematic errors. For example, if the OIS-underlying rate consistently changes less than one for with the target rate ($\beta < 1$), then the OIS-implied future rate forecasts will be biased downwards.

By the end of 2000 all major central banks adopted announcements of target interest rates on pre-scheduled dates. We collect the monetary policy announcements data for all countries in Table 4.1 except for Japan. The data comes from the web pages of the central banks. The last three columns of Table 4.3 report the total number of announcements, target rate hikes and cuts respectively. The numbers in parentheses are unscheduled policy meetings which occurred outside the pre-specified calendars. The sample for each country begins with availability of the corresponding OIS data (the last column in Table 4.1) and ends in June 2017.

One additional piece of information required to test the equivalence of changes in the OIS underlying and target rates is the effective date of the target rate change. In the following countries the target rate change becomes effective on the day of the announcement: Australia (before 2008), Canada, New Zealand, Switzerland. For Australia (after 2008), the United Kingdom and the United States, the change in the target rate comes into effect on the day following the meetings. In the Eurozone and Sweden the effective date is announced at the meeting, and is usually in a week from the announcement date.

Table 4.5 reports the estimates of regression 4.8 for target rate announcements of central banks represented by currency codes in columns. The χ^2 statistics is for the joint hypothesis of zero intercept and unit slope. The intercept is in basis points p.a. The numbers in parentheses report standard errors for coefficient estimates and p-value of the joint hypothesis, both derived using HAC estimator of Newey and West (1987) with optimal number of lags according to Newey and West (1994). For each currency sample starts with availability of the corresponding OIS data as specified in Table 4.1 and ens in June 2017.

For each currency the slope coefficients are statistically indistinguishable from one, meaning that, on average, the underlying rates move in lockstep with target rates. Furthermore, the joint hypothesis of zero slope and unit intercept can not be rejected at any conventional significance level for all currencies except for the Australian dollar, which is also the single currency with statistically significant intercept of two hundredth of a basis point. As we mentioned earlier, there is almost no inter-event variability in the AUD OIS underlying over the course of the sample. However, the underlying is the Reserve Bank of Australia's estimate of the overnight unsecured interbank rate in the domestic market and deviated by one to four basis points from the official target rate, with the bulk of deviations concentrated before 2003, which together with virtually no discrepancy between the underlying and the target in the following years leads to the inflated statistical significance.

Table 4.5: **Projection of overnight rate changes on policy rate changes.**

	AUD	CAD	CHF	EUR	GBP	NZD	SEK	USD
α	0.02	0.05	−0.49	−0.28	0.36	0.00	0.16	−0.09
(s.e.)	(0.01)	(0.09)	(1.01)	(0.45)	(0.45)		(0.45)	(0.70)
β	1.00	0.99	0.87	0.93	1.13	1.00	1.10	0.84
(s.e.)	(0.00)	(0.01)	(0.07)	(0.12)	(0.08)		(0.08)	(0.09)
$\chi^2(2)$	11.57	1.34	4.18	0.95	2.70		1.52	3.88
(p-value)	(0.00)	(0.51)	(0.12)	(0.62)	(0.26)		(0.47)	(0.14)

This table presents the intercept and slope estimates for the following regression:

$$\Delta \bar{r}_a = \alpha + \beta \Delta r_{\text{tgt},a} + \varepsilon_a, \;\; a \in A,$$

where $\Delta \bar{r}_a$ is the change in the average OIS underlying rate prevailing after and before announcement a, $\Delta r_{\text{tgt},a}$ is the respective target rate change, and A is the set of monetary policy meetings of a central bank. The regression is estimated by each central bank in the sample (represented by currencies in columns). The χ^2 statistics is for the joint hypothesis of zero intercept and unit slope. The intercept is in basis points p.a. The numbers in parentheses report standard errors for coefficient estimates and p-value of the joint hypothesis, both derived using the HAC estimator of Newey and West (1987) with optimal number of lags according to Newey and West (1994). For each currency the sample starts with availability of the corresponding OIS data as specified in Table 4.1 and ends in June 2017.

To obtain a deeper insight into the relation between the underlying rate changes and the target rate changes, in Figure 4.1 we plot the distribution of the former conditional on the latter for each currency. The target and underlying rate changes (in basis points p.a.) are along horizontal and vertical axes respectively, the average rates are computed over the inter-announcements periods. For each target rate change bracket we draw a box plot, with the box covering middle 50% of the empirical distribution, whiskers covering additional 1.5 interquartile ranges in both directions, and outliers beyond the whiskers drawn as dots. The red diamonds represent the perfect fit, as if intercept and slope were equal to zero and one for every target rate change within a bracket and the underlying moved in lockstep with the target rate. The sample is the same as in Table 4.5.

Figure 4.1 shows that the post-announcement OIS underlying rates changes are indeed concentrated around the target rate changes. For a one-to-one relation, the distribution collapses to the single point of mass coinciding with the red diamonds in the subfigures, as in the case of New Zealand and Australia. The mean underlying rate change, depicted by the solid line within the body of each box, is most close to the announced target rate change in the majority of cases, with several large deviation for a handful cuts in Switzerland.

Figure 4.1 also highlights that the deviations from the equivalence between the OIS underlying and target rate changes (in terms of unit slope point estimates) are driven by a few extreme events. For example, in the US the cut of 87.5 basis points on December 16 2008[20] is arguably the main source of the slope point estimate deviating from one, ignoring this announcement alone results in the slope estimate of 0.95. Similarly, the 50 basis point target rate hike in the Eurozone on June 6, 2000 and 100 basis point cut on November 20, 2008 in Switzerland are responsible for the deviations of the slope estimates for these countries. The same reasoning also applies to the United Kingdom, where the 150 basis point cut was announced on November 6, 2008. In Sweden the deviation is solely due to the underlying's overreaction to 50 basis points target rate cuts. Three out of five 50 basis points happened during the acute stage of the 2007-2009 financial crisis. On each of the six meetings from October 8, 2008 to Jule 2, 2009, the Riksbank announced reduction in its target rate slashing it from 4.75% to 0.25%. More importantly, at its monetary policy announcements, the Bank publishes the projected path of the policy rate, which was substantially revised downwards at each announcement during this period.[21] For the Australian, Canadian and New Zealand dollar OIS, the underlying and target rates are virtually the same.

Overall, we find empirical evidence supporting the assumptions stated at the beginning of this section, so the OIS-implied future rates are likely to reflect the future policy rates. Next we proceed to inferring the expected direction of upcoming target rate changes from the quantity in equation (4.6).

[20] On this day the Fed set the target range for the federal funds rate from zero to 25 basis points, before this date it targeted the level of the rate.

[21] Similarly, the 50 basis point cut announced on July 3, 2014 was accompanied with a significant downside revision of the projected target and followed by subsequent reductions in the target rate. The Riksbank is also quite explicit in communicating its view of the future target rate to the public: for instance in the release accompanying the October 22, 2008 cut, the Bank communicated that the rate "will need to be cut further by 0.5 percentage points" over the coming half a year.

Figure 4.1: **Changes in the OIS underlying rates vs. target rate changes.**

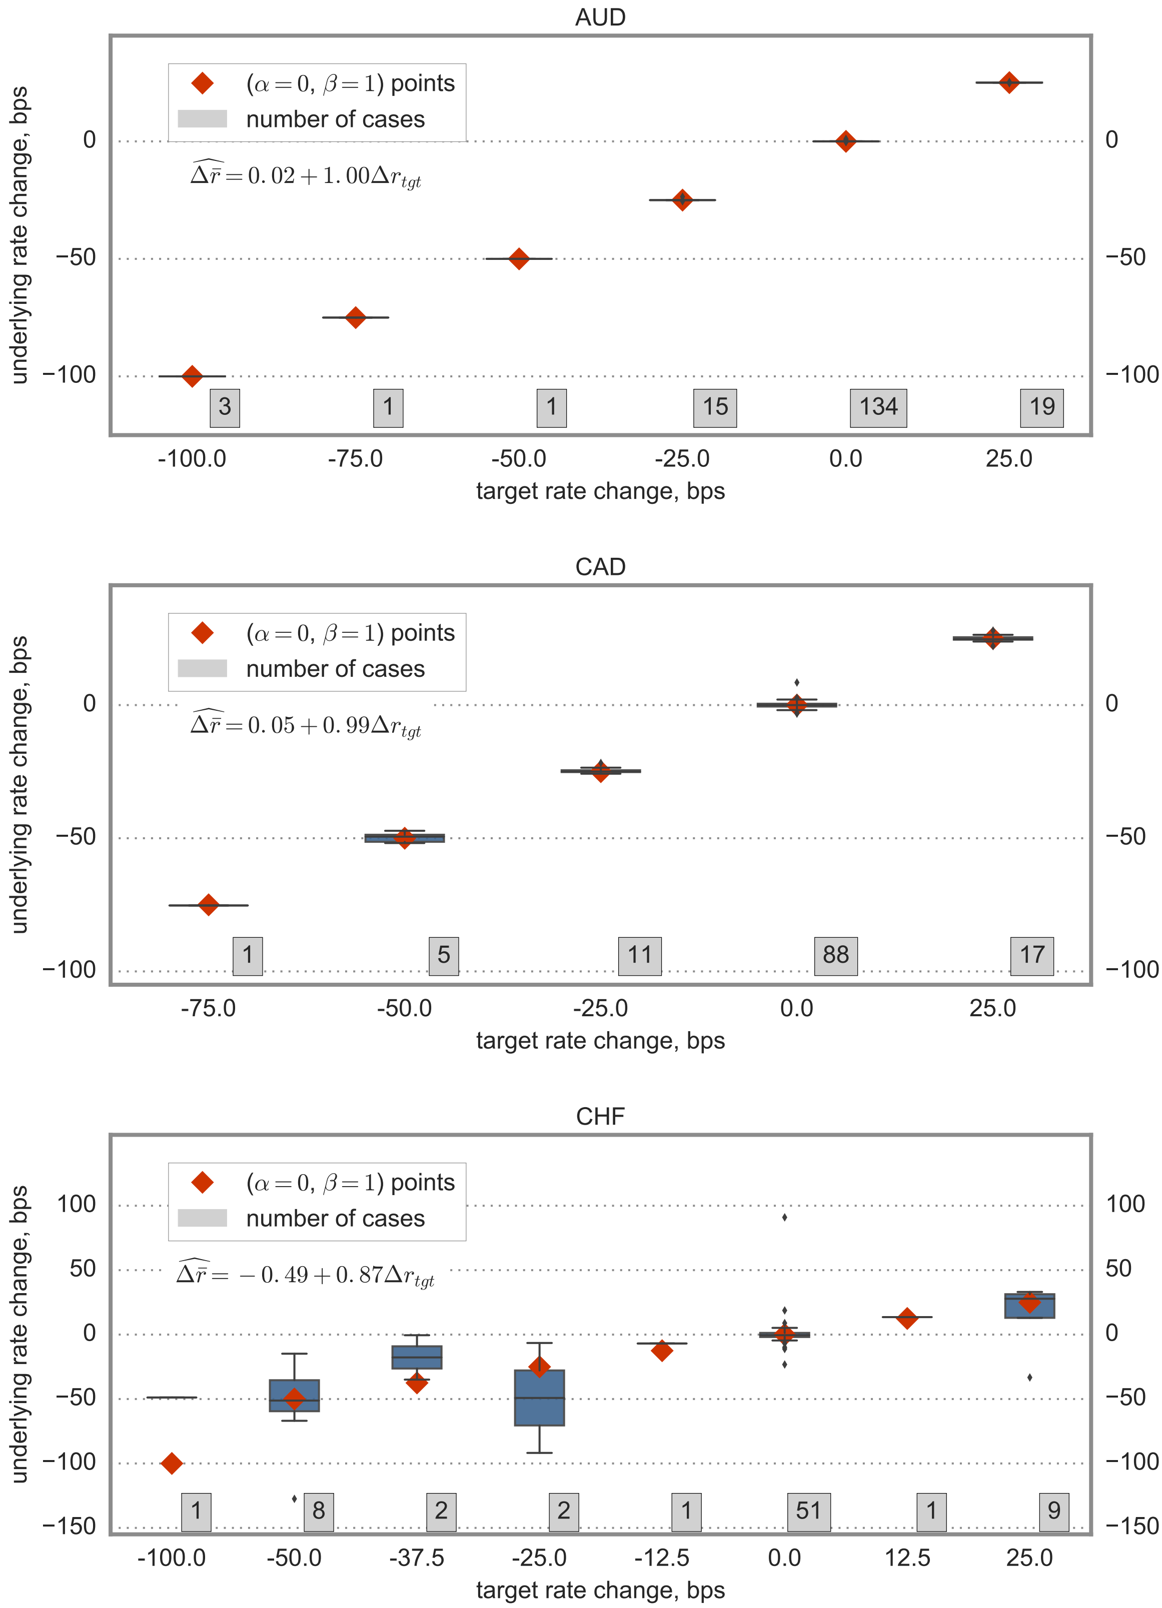

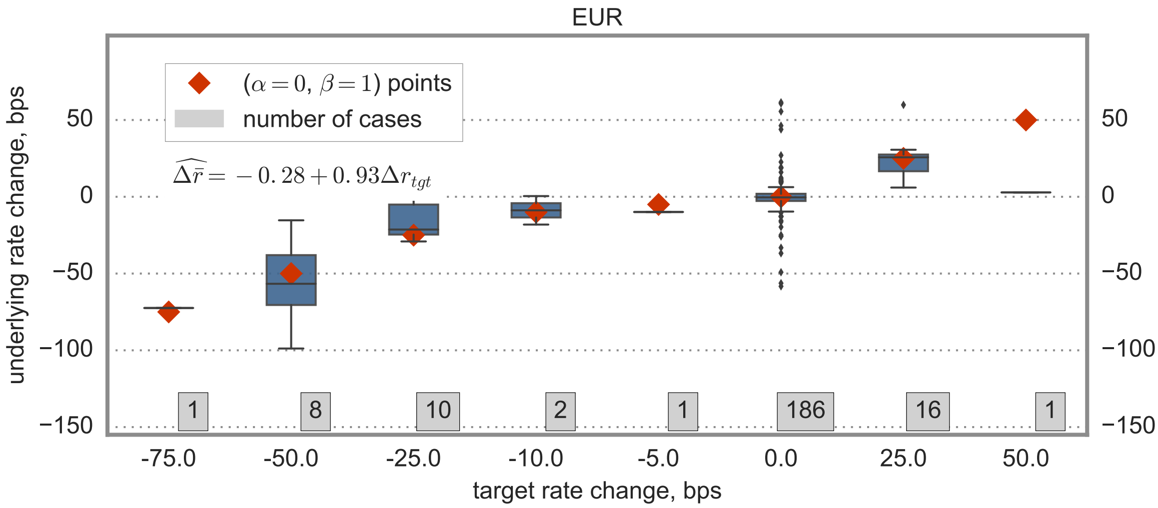
EUR
(α = 0, β = 1) points
number of cases
$\widehat{\Delta\bar{r}} = -0.28 + 0.93\Delta r_{tgt}$
underlying rate change, bps
target rate change, bps
1
8
10
2
1
186
16
1
-75.0
-50.0
-25.0
-10.0
-5.0
0.0
25.0
50.0

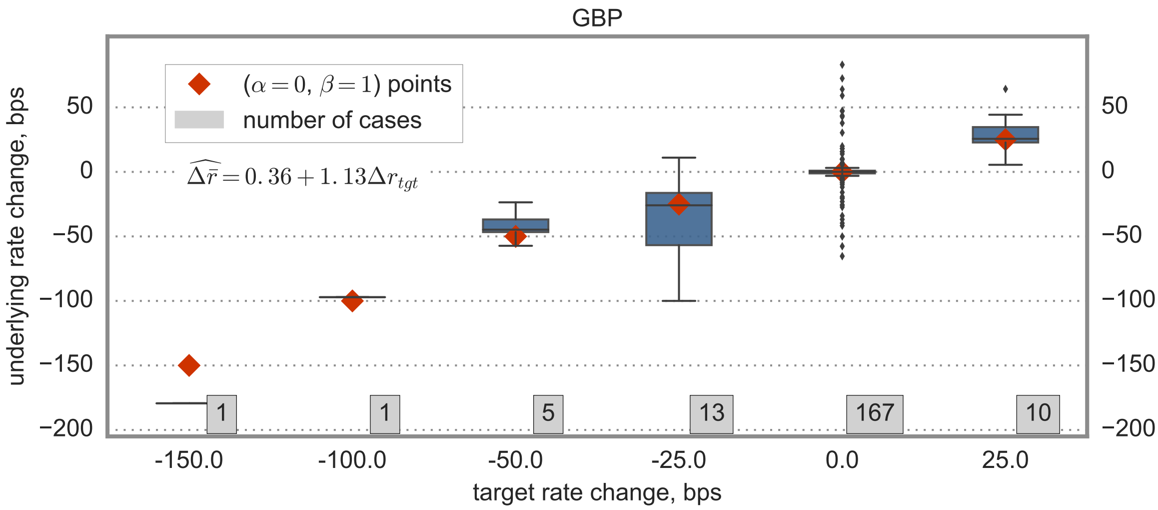
GBP
(α = 0, β = 1) points
number of cases
$\widehat{\Delta\bar{r}} = 0.36 + 1.13\Delta r_{tgt}$
underlying rate change, bps
target rate change, bps
1
1
5
13
167
10
-150.0
-100.0
-50.0
-25.0
0.0
25.0

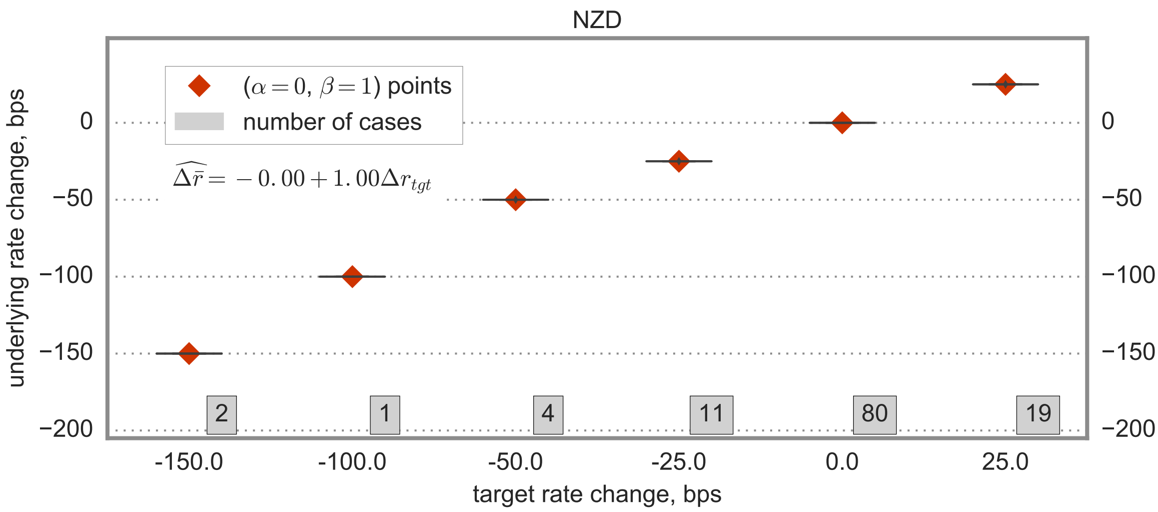
NZD
(α = 0, β = 1) points
number of cases
$\widehat{\Delta\bar{r}} = -0.00 + 1.00\Delta r_{tgt}$
underlying rate change, bps
target rate change, bps
2
1
4
11
80
19
-150.0
-100.0
-50.0
-25.0
0.0
25.0

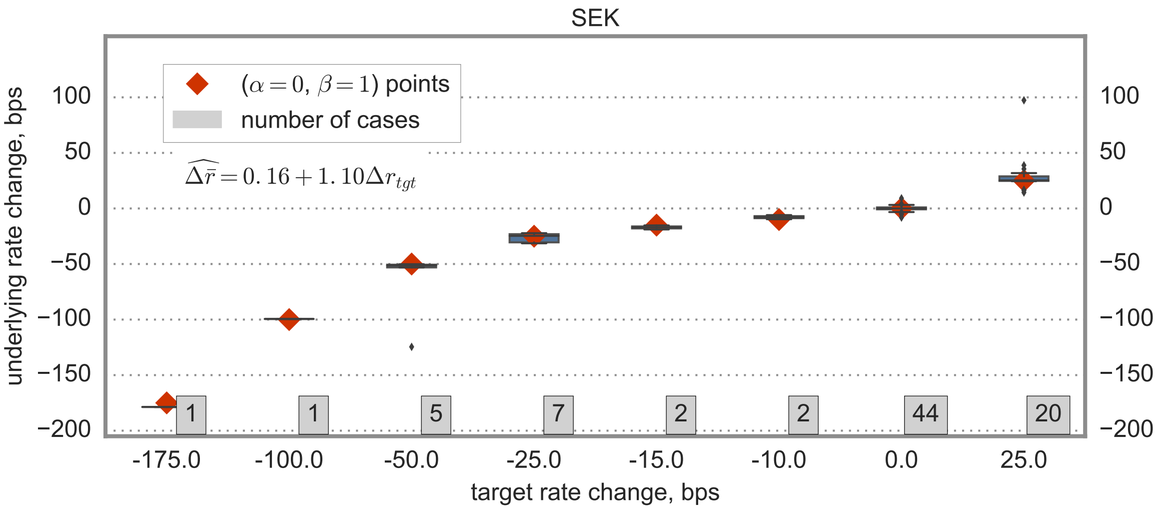

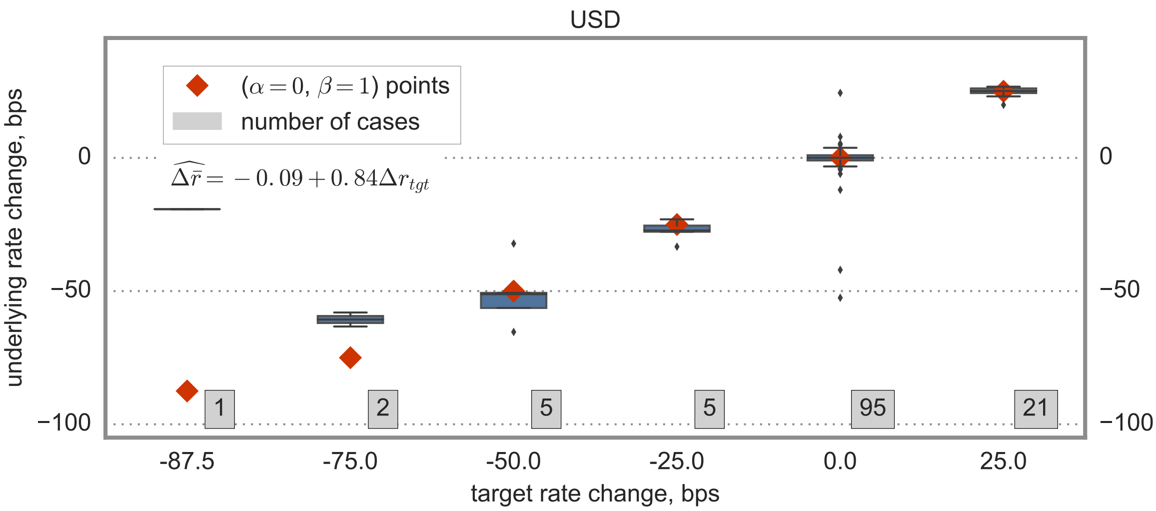

These figures present response of the OIS underlying rates to changes in the target rates of the respective central banks, conditional on the magnitude of the latter. The target rate changes are along the horizontal axis, and the changes between pre- and post-announcement averages of the OIS underlying rate are along the vertical axis. The average rates are computed over the whole period between events. All rate changes are in basis points. Each box covers the middle 50% of the distribution (between the first and the third quartiles of the dataset) while the whiskers cover additional 1.5 interquartile ranges in both directions. The outliers beyond the whiskers are shown as dots. The red diamonds represent the points of perfect fit, as if the underlying and target rates would comove perfectly. The number of observations of a certain target rate change is framed in a dark gray box in the bottom of each plot. For each currency the sample starts with availability of the corresponding OIS data as specified in Table 4.1 and ends in June 2017.

4.3.3 Classifying upcoming target rate decisions

Intuitively, one would expect a rate hike if the implied post-meeting OIS underlying rate exceeds the current level, a rate cut in the opposite case, and no announced change if the rates are equal. In reality, the third case almost never happens, as the implied future rates exceed or undershoot the current underlying rate levels because of noise or risk premium in the market quotes of OIS, estimation errors, failure of Assumptions 1 through 4 and other random and non-random factors. Also, target rate changes have historically been observed over a very limited set of values, most often equaled 25 basis points, and never gone below 10 basis points. That said, instead of deeming any positive (negative) difference to be a hike (cut), we choose a threshold which the difference should exceed (fall short of) to be classified as a target rate change. Following the vast statistics literature on such classification problems, we call the sign of target rate changes *response*, hikes, cuts and no-changes *classes*, and the difference between the implied and the current underlying rate *marker*. The classification rule is stated as follows:

$$\widehat{a_t} = \begin{cases} -1, & \text{if } \Delta_t \leq h^-, \\ 0, & \text{if } h^- < \Delta_t < h^+, \\ +1, & \text{if } \Delta_t \geq h^+, \end{cases}$$

where $\widehat{a_t}$ is the time t estimate of the response, and $\Delta_t = E_t \Delta r_{\text{tgt}}$ as in equation (4.6) is the marker.

The values of h^- and h^+ are arbitrary classification thresholds. For each pair thereof, it is possible to compare the predicted target rate decisions to the actually observed ones, which is best illustrated in form of a confusion matrix such as the one in Figure 4.2.

In the two examples presented therein, we choose two different pairs of thresholds (±25 and ±10 in Panel 4.2a and Panel 4.2b respectively) to show how this choice determines the power to discriminate between different types of announcements five days in advance. In theory, for any series of recovered implied rates, there exists an infinite number of such confusion matrices, one for each possible pair of thresholds, although many of them look alike. A natural question to ask is whether one series, e.g. the rates implied in the prices of federal funds futures, results in a better marker than another, e.g. the OIS-implied rates.

To answer this question, we calculate the volume-under-the-surface (VUS) measure of classifying power, which is a three-class extension of the area-under-the-curve approach to choosing the best threshold-based classifier. As discussed in detail by Mossman (1999) and Dreiseitl, Ohno-Machado, and Binder (2000), this measure is equal to the probability that a randomly drawn triplet of observations, one from each class, is sorted in the correct order (that is, the marker value corresponding to the actually observed cut is less than

Figure 4.2: **Two confusion matrices in forecasting the sign of FOMC announcements.**

(a) $h^- = h^+ = 25$ bps

predicted \ actual	-1	0	1
-1	10	0	0
0	4	92	7
1	0	0	13

(b) $h^- = h^+ = 10$ bps

predicted \ actual	-1	0	1
-1	14	1	0
0	0	90	0
1	0	1	20

Panel 4.2a (Panel 4.2b) in this figure shows the confusion matrix of the target rate change direction classifier five days ahead of announcements, with thresholds set to ± 25 (± 10) basis points. Entry (x, y) (x denotes rows) in the matrices shows the number of announcements of type y that were classified as type x, such as the columns sum up to the total number of announcements of each type.

that corresponding to the no-change, and both are less than that corresponding to the hike). A non-parametric unbiased estimator of VUS is as follows:

$$VUS = \frac{1}{n_h n_c n_z} \sum_{p=1}^{n_h} \sum_{q=1}^{n_c} \sum_{m=1}^{n_z} I(x_{h,p} x_{c,q} x_{z,m}), \tag{4.9}$$

where h, c, z stand for rate hikes, cuts and no-change decisions respectively; n_h, n_c and n_z are the number of rate hikes, cuts and no-changes respectively; $x_{i,j}$ denotes j^{th} marker value from class i; $I(\cdot)$ is a function that takes value 1 if its three arguments are correctly ordered. VUS equals 1 if the classifier is perfect, i.e. if it is possible to select a pair of thresholds (h^+, h^-) such that the measurements are perfectly separated into classes, and equal to 1/6 if the classification is done randomly, e.g. if the marker values are i.i.d. distributed across classes.

Table 4.6 presents the VUS estimates for classifiers constructed as in equation (4.6) for a number of currencies.

Table 4.6: Forecasting power of OIS and federal funds futures.

lag	AUD	CAD	CHF	EUR	GBP	NZD	SEK	USD	USD (FFF)
1	0.9726	0.9816	0.8071	0.8138	0.8764	0.9812	0.8479	0.9932	0.9447
2	0.9686	0.9063	0.7950	0.9291	0.8317	0.9879	0.9032	0.9980	0.9694
3	0.9643	0.9207	0.7763	0.9232	0.7912	0.9864	0.8476	1.0000	0.9890
4	0.9666	0.9195	0.7822	0.9087	0.6474	0.9795	0.8190	0.9990	0.9960
5	0.9442	0.9698	0.7566	0.8842	0.6274	0.9863	0.8090	0.9990	0.9970
6	0.9501	0.9633	0.7515	0.8449	0.6246	0.9809	0.8001	0.9978	0.9924
7	0.9387	0.9383	0.7420	0.8219	0.5778	0.9804	0.8043	1.0000	0.9918
8	0.9419	0.8828	0.6601	0.7787	0.6202	0.9848	0.7440	1.0000	0.9839
9	0.9383	0.8849	0.7127	0.7342	0.6299	0.9804	0.7266	0.9957	0.9783
10	0.9309	0.8232	0.7018	0.7928	0.6909	0.9729	0.7156	0.9753	0.9283

This table presents the volume-under-surface (VUS) values of classifiers based on the OIS-implied rates, by prediction lag (in days). On each day before an announcement, we recover the OIS underlying rate expected to prevail after it, proxy the rate expected to prevail before it as the sample average of the rates since the last announcement and calculate the difference between the two. The time-series of such differences becomes the marker for the upcoming target rate changes, and the VUS estimates thereof are obtained on the full sample of hikes, cuts and no-change decisions using eq. (4.9). The last column reports the VUS values of the classifier based on the federal funds futures-implied rates. Sample size start is different for each contract, as described in Table 4.1, the end is in June 2017.

Not surprisingly, prediction accuracy deteriorates as we move farther from announcements, although the drop is barely evident for the announcements in the countries where the best classification power is achieved: Australia, New Zealand and the US. In the US, investors could have been able to tell upcoming hikes from cuts and no-changes almost every time since 2001. Interestingly, even in Switzerland, where the central bank target a rate most different from the OIS underlying rate, classification power is high at 0.81 immediately before announcements.

Table 4.7: Forecasting power of LIBOR rates.

lag	AUD	CAD	CHF	EUR	GBP	NZD	SEK	USD
1	0.4980	0.3292	0.5124	0.4116	0.3089	0.5356	0.3676	0.5013
2	0.4338	0.3207	0.4526	0.4027	0.3025	0.5636	0.3581	0.4906
3	0.4054	0.3017	0.5030	0.3970	0.1896	0.5540	0.3456	0.4494
4	0.3892	0.2593	0.5124	0.3959	0.1737	0.5125	0.3234	0.4341
5	0.3702	0.2309	0.5124	0.3784	0.1576	0.4510	0.3396	0.4304
6	0.3411	0.2033	0.5308	0.3611	0.1827	0.3800	0.3330	0.3607
7	0.3149	0.1782	0.5291	0.3471	0.1945	0.3256	0.3170	0.3571
8	0.2853	0.1677	0.5047	0.3175	0.2265	0.2532	0.2961	0.3695
9	0.2717	0.1684	0.4748	0.3066	0.2296	0.2072	0.2621	0.3564
10	0.2850	0.1824	0.4171	0.2760	0.2577	0.1890	0.2392	0.3647

This table presents the volume-under-surface (VUS) values of classifiers based on the LIBOR-implied rates, by prediction lag (in days). At each prediction lag before an announcement, we recover the implied overnight rate expected to prevail after it, proxy the rate expected to prevail before it as the 5-day rolling average of the previous rates and calculate the difference between the two. The time-series of such differences becomes the marker for the upcoming target rate changes, and the VUS estimates thereof are obtained on the full sample of hikes, cuts and no-change decisions using eq. (4.9). Sample size start is different for each contract, as described in Table 4.1, the end is in June 2017.

To compare the OIS-based marker to a literature staple, in the last column of the Table we show the VUS estimates of the marker constructed from the federal funds futures-implied rates. We use the standard methodology of recovering the implied rates (see Kuttner (2001) for a reference) and proxy the expected target rate change as in equation (4.6). Interestingly, OIS provide better material for predicting the course of monetary policy in the US: at every lag the OIS-implied rates dominate their federal funds futures-implied counterparts in terms of the VUS value.

Table 4.7 shows the outcome of the similar exercise with the LIBOR-implied rates. We collect the one-month LIBOR quotes from Bloomberg[22] and compute the LIBOR-implied in the same way, we compute the OIS-implied rates.

With the LIBOR rates the forecasting power in terms of VUS decreases dramatically for all currencies and forecast horizons: a randomly selected triplet of hike, cut and no change in the US target rate is correctly predicted in 43 percent of cases five days in advance of the FOMC announcements, comparing to 99.9 and 99.7 percent when using the OIS- and the federal funds futures-implied rates respectively. At the five day forecast horizon the prediction accuracy in the United Kingdom is only 15.8 percent, which is worse than a random guess – in that case one would expect to correctly classify one sixth of all triplets.

[22] After 2013 when LIBOR ceased to exist in a number of countries, we use one-month Bank Bill Swap Rate for Australia, Canadian Dollar Offered Rate for Canada, Bank Bill yields for New Zealand, and Stockholm Interbank Offered Rate for Sweden as substitutes.

4.4 Conclusions

We study the information content of OIS rates and find them to be unbiased predictors of future compound short rates across developed countries. We document that for the overwhelming majority of currencies the prices of OIS do not bear significant risk premium. We further demonstrate that the OIS underlying rates accurately reflect the corresponding target rates set by central banks, thus providing empirical evidence favoring OIS rates as unbiased forecasts of the future target rates compounded over lifetime of the contract.

We critically assess the assumptions needed to extract the expected future policy rates from the OIS prices, finding support for their validity in the data. We find that the OIS-implied future policy rates can accurately predict the outcomes of monetary policy meetings across the major economies.

Our findings have important implications for gauging financial markets' expectations of future policy rates. Specifically, we advocate OIS as a novel tool for evaluating the future path of monetary policy that can be used by academics, practitioners, and regulators. More importantly, we demonstrate robustness of this tool by bringing international dimension to the monetary policy prediction exercise.

Bibliography

Aït-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini (2018). *The Term Structure of Variance Swaps and Risk Premia*. Research paper 18-37. Swiss Finance Institute.

Ammann, Manuel and Ralf Buesser (2013). "Variance risk premiums in foreign exchange markets". In: *Journal of Empirical Finance* 23, pp. 16–32. ISSN: 0927-5398.

Andersen, Torben G. et al. (2003). "Modeling and Forecasting Realized Volatility". In: *Econometrica* 71(2), pp. 579–625.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen (2013). "Value and Momentum Everywhere". In: *The Journal of Finance* 68(3), pp. 929–985.

Baba, Naohiko and Frank Packer (Nov. 2009). "Interpreting deviations from covered interest parity during the financial market turmoil of 2007-08". In: *Journal of Banking & Finance* 33(11), pp. 1953–1962.

Bansal, Ravi and Amir Yaron (2004). "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles". In: *The Journal of Finance* 59(4), pp. 1481–1509.

Beine, Michel, Gust Janssen, and Christelle Lecourt (Sept. 2009). "Should central bankers talk to the foreign exchange markets?" In: *Journal of International Money and Finance* 28(5), pp. 776–803.

Bekaert, Geert and George Panayotov (2018). *Good carry, bad carry*. Research paper 15-53. Columbia Business School.

Bernanke, Ben S and Kenneth N Kuttner (2005). "What explains the stock market's reaction to Federal Reserve policy?" In: *The Journal of finance* 60(3), pp. 1221–1257.

BIS (2016a). *Triennial Central Bank Survey of Foreign Exchange and Derivative Market Activity in 2016*. Tech. rep. Bank for International Settlements Editorial Group.

Triennial Central Bank Survey OTC Interest Rate Derivatives Turnover in April 2016 (2016b). Bank for International Settlements, Basel.

Black, Fischer and Myron Scholes (1973). "The Pricing of Options and Corporate Liabilities". In: *Journal of Political Economy* 81(3), pp. 637–54.

Bloom, Nicholas (2009). "The impact of uncertainty shocks". In: *Econometrica* 77(3), pp. 623–685.

Sterling Monetary Framework Annual Report 2016-2017 (2017). Bank of England.

Bollerslev, Tim, James Marrone, et al. (2014). "Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence". In: *Journal of Financial and Quantitative Analysis* 49(3), pp. 633–661.

Bollerslev, Tim, George Tauchen, and Hao Zhou (2009). "Expected Stock Returns and Variance Risk Premia". In: *The Review of Financial Studies* 22(11), pp. 4463–4492.

Breeden, Douglas T. and Robert H. Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Option Prices". In: *The Journal of Business* 51(4), pp. 621–651. ISSN: 00219398, 15375374.

Britten-Jones, Mark and Anthony Neuberger (2000). "Option Prices, Implied Price Processes, and Stochastic Volatility". In: *The Journal of Finance* 55(2), pp. 839–866. ISSN: 00221082, 15406261.

Brown, Stephen J. and Jerold B. Warner (1980). "Measuring Security Price Performance". In: *Journal of Financial Economics* 8, pp. 205–258.

Brown, Stephen and Jerold B. Warner (1980). "Measuring security price performance". In: *Journal of Financial Economics* 8(3), pp. 205–258.

Buraschi, Andrea, Robert Kosowski, and Fabio Trojani (2014). "When There Is No Place to Hide: Correlation Risk and the Cross-Section of Hedge Fund Returns". In: *The Review of Financial Studies* 27(2), pp. 581–616.

Burkhard, Lukas and Andreas M. Fischer (Sept. 2009). "Communicating policy options at the zero bound". In: *Journal of International Money and Finance* 28(5), pp. 742–754.

Caballero, Ricardo J. and Joseph B. Doyle (Dec. 2012). *Carry Trade and Systemic Risk: Why are FX Options so Cheap?* Working Paper 18644. National Bureau of Economic Research.

Carr, Peter and Dilip Madan (1998). "Towards a Theory of Volatility Trading". In: *Reprinted in Option Pricing, Interest Rates, and Risk Management, Musiella, Jouini, Cvitanic*. University Press, pp. 417–427.

Carr, Peter and Liuren Wu (2009). "Variance Risk Premiums". In: *The Review of Financial Studies* 22(3), pp. 1311–1341.

Choi, Hoyong, Philippe Mueller, and Andrea Vedolin (2017). "Bond Variance Risk Premiums*". In: *Review of Finance* 21(3), pp. 987–1022.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen (2019). "Stock returns over the FOMC cycle". In: *Journal of Finance, Forthcoming*.

Danthine, Jean-Pierre (2012). "Monetary policy is not almighty". In: *Speech at "Journée Solutions Bancaires"*(Geneva, 31 May 2012).

De Fiore, Fiorella, Marie Hoerova, and Harald Uhlig (2018). *Money markets, collateral and monetary policy*. Tech. rep. National Bureau of Economic Research.

Della Corte, Pasquale, Tarun Ramadorai, and Lucio Sarno (2016). "Volatility risk premia and exchange rate predictability". In: *Journal of Financial Economics* 120(1), pp. 21–40. ISSN: 0304-405X.

Dew-Becker, Ian et al. (2017). “The price of variance risk”. In: *Journal of Financial Economics* 123(2), pp. 225–250.

Dewachter, Hans et al. (2014). “The intra-day impact of communication on euro-dollar volatility and jumps”. In: *Journal of International Money and Finance* 43(C), pp. 131–154.

Drechsler, Itamar and Amir Yaron (2011). “What’s Vol Got to Do with It”. In: *The Review of Financial Studies* 24(1), pp. 1–45.

Dreiseitl, Stephan, Lucila Ohno-Machado, and Michael Binder (2000). “Comparing Three-class Diagnostic Tests by Three-way ROC Analysis”. In: *Medical Decision Making* 20(3). PMID: 10929855, pp. 323–331.

Driessen, Joost, Pascal J. Maenhout, and Grigory Vilkov (2009). “The Price of Correlation Risk: Evidence from Equity Options”. In: *The Journal of Finance* 64(3), pp. 1377–1406.

Driscoll, John C. and Aart C. Kraay (1998). “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data”. In: *The Review of Economics and Statistics* 80(4), pp. 549–560.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan (2018). “Deviations from Covered Interest Rate Parity”. In: *The Journal of Finance* 73(3), pp. 915–957.

Ehlers, Torsten and Egemen Eren (2016). “The changing shape of interest rate derivatives markets”. In: *BIS Quarterly Review, December 2016*, pp. 53–65.

Evans, Charles L (1998). “Real-time Taylor rules and the federal funds futures market”. In: *Economic Perspectives-Federal Reserve Bank of Chicago* 22, pp. 44–55.

Fama, Eugene F. et al. (1969). “The Adjustment of Stock Prices of New Information”. In: *International Economic Review* 10, pp. 1–21.

Fama, Eugene, Lawrence Fisher, et al. (1969). “The Adjustment of Stock Prices to New Information”. In: *International Economic Review* 10(1), pp. 1–21. ISSN: 00206598, 14682354.

Fama, Eugene and Kenneth French (2014). *A Five-Factor Asset Pricing Model.* Working Paper.

Fernández-Villaverde, Jesús et al. (Oct. 2011). “Risk Matters: The Real Effects of Volatility Shocks”. In: *American Economic Review* 101(6), pp. 2530–61.

Fleming, Michael J et al. (2012). “An analysis of OTC interest rate derivatives transactions: implications for public reporting”. In: *Federal Reserve Bank of New York Staff Reports* 557.

Fratzscher, Marcel (Dec. 2008). “Communication and exchange rate policy”. In: *Journal of Macroeconomics* 30(4), pp. 1651–1672.

Fung, William and David A. Hsieh (June 2015). “The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers”. In: *The Review of Financial Studies* 14(2), pp. 313–341. ISSN: 0893-9454.

Gürkaynak, Refet, Brian Sack, and Eric Swanson (2007). "Market-based measures of monetary policy expectations". In: *Journal of Business & Economic Statistics* 25(2), pp. 201–212.

Hanke, Michael, Rolf Poulsen, and Alex Weissensteiner (2015). "Where Would the EUR/CHF Exchange Rate be Without the SNB's Minimum Exchange Rate Policy?" In: *Journal of Futures Markets* 35(12), pp. 1103–1116. ISSN: 1096-9934.

Hastie, T., R. Tibshirani, and J.H. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer series in statistics. Springer. ISBN: 9780387848846.

Hayward, Rob (2018). "Foreign Exchange Speculation: An Event Study". In: *International Journal of Financial Studies* 6(1), p. 22.

Hertrich, Markus and Heinz Zimmermann (2017). "On the Credibility of the Euro/Swiss Franc Floor: A Financial Market Perspective". In: *Journal of Money, Credit and Banking* 49(2-3), pp. 567–578.

Hodrick, Robert J (1992). "Dividend yields and expected stock returns: Alternative procedures for inference and measurement". In: *The Review of Financial Studies* 5(3), pp. 357–386.

Hull, John C and Alan White (2013). "LIBOR vs. OIS: The derivatives discounting dilemma". In: *Journal of Investment Management* 11(3), pp. 14–27.

Standard Credit Support Annex (2013). International Swaps and Derivatives Association.

Janssen, Alexandra and Rahel Studer (Aug. 2014). *The Swiss franc's honeymoon*. ECON - Working Papers 170. Department of Economics - University of Zurich.

Jermann, Urban J. (2017). "Financial Markets' Views about the Euro-Swiss Franc Floor". In: *Journal of Money, Credit and Banking* 49(2-3), pp. 553–565.

Jiang, George J. and Yisong S. Tian (2005). "The Model-Free Implied Volatility and Its Information Content". In: 18(4), pp. 1305–1342.

Jordan, Thomas (2011). "Does the Swiss National Bank need equity?" In: *Statistisch-Volkswirtschaftliche Gesellschaft*(Basel, 28 September 2011).

Jurek, Jakub W. (2014). "Crash-neutral currency carry trades". In: *Journal of Financial Economics* 113(3), pp. 325–347. ISSN: 0304-405X.

Karnaukh, Nina (2016). "The Dollar Ahead of FOMC Target Rate Changes". In: *Job Market Paper*.

Kearns, Jonathan and Phil Manners (Dec. 2006). "The Impact of Monetary Policy on the Exchange Rate: A Study Using Intraday Data". In: *International Journal of Central Banking* 2(4).

Kim, Don H. and Jonathan H. Wright (2005). "An Arbitrage-Free Three-Factor Term Structure Model and the Recent behavior of Long-Term Yields and Distant-Horizon Forward Rates". In:

Kreicher, Lawrence, Robert N McCauley, and Philip Wooldridge (2017). "The bond benchmark continues to tip to swaps". In: *BIS Quarterly Review, March 2017*(March), pp. 69–79.

Krueger, Joel T and Kenneth N Kuttner (1996). "The Fed funds futures rate as a predictor of Federal Reserve policy". In: *Journal of Futures Markets* 16(8), pp. 865–879.

Krugman, Paul R. (1991). "Target Zones and Exchange Rate Dynamics". In: *The Quarterly Journal of Economics* 106(3), pp. 669–682.

Kuttner, Kenneth N (2001). "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market". In: *Journal of monetary economics* 47(3), pp. 523–544.

Kuttner, Kenneth N (2014). "Monetary policy during Japan's Great Recession: From self-induced paralysis to Rooseveltian Resolve". In: *PIIE Briefing* 14(4), pp. 66–79.

Lucca, David O and Emanuel Moench (2015). "The Pre-FOMC Announcement Drift". In: *The Journal of Finance* 70(1), pp. 329–371.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (2011). "Common Risk Factors in Currency Markets". In: *Review of Financial Studies* 24(11), pp. 3731–3777.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan (2014). "Countercyclical Currency Risk Premia". In: *Journal of Financial Economics* 111(3), pp. 527–553.

Malz, Allan M (June 2014). *A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions.* Federal Reserve Bank of New York Staff Report 677.

Menkhoff, Lukas et al. (Dec. 2011). "Currency Momentum Strategies". In: (366).

Mirkov, Nikola, Igor Pozdeev, and Paul Söderlind (Jan. 2016). *Toward Removal of the Swiss Franc Cap: Market Expectations and Verbal Interventions.* Research Paper 2016/14.

Moreira, Alan and Tyler Muir (2017). "Volatility-managed portfolios". In: *The Journal of Finance* 72(4), pp. 1611–1644.

Mossman, Douglas (1999). "Three-way ROCs". In: *Medical Decision Making* 19(1). PMID: 9917023, pp. 78–89.

Mueller, Philippe, Andreas Stathopoulos, and Andrea Vedolin (2017). "International correlation risk". In: *Journal of Financial Economics* 126(2), pp. 270–299. ISSN: 0304-405X.

Mueller, Philippe, Alireza Tahbaz-Salehi, and Andrea Vedolin (2017). "Exchange Rates and Monetary Policy Uncertainty". In: *The Journal of Finance* 72(3), pp. 1213–1252. ISSN: 1540-6261.

Munk, Claus (2013). *Financial Asset Pricing Theory.* 1st ed. Oxford University Press.

Neuhierl, Andreas and Michael Weber (June 2018). *Monetary Momentum.* Working Paper 24748. National Bureau of Economic Research.

Newey, Whitney K and Kenneth D West (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". In: *Econometrica* 55(3), pp. 703–708.

Newey, Whitney K and Kenneth D West (1994). "Automatic Lag Selection in Covariance Matrix Estimation". In: *The Review of Economic Studies* 61(4), pp. 631–653. ISSN: 00346527, 1467937X.

Piazzesi, Monika and Eric Swanson (2008). "Futures prices as risk-adjusted forecasts of monetary policy". In: *Journal of Monetary Economics* 55(4), pp. 677–691.

Reiswich, Dimitri and Uwe Wystup (2009). *FX Volatility Smile Construction*. CPQF Working Paper Series 20. Frankfurt School of Finance and Management, Centre for Practical Quantitative Finance (CPQF).

The Riksbank's operational framework for the implementation of monetary policy – A review (2014). Sveriges Riksbank.

Sack, Brian (2004). "Extracting the expected path of monetary policy from futures rates". In: *Journal of Futures Markets* 24(8), pp. 733–754.

SNB (2011). "Swiss National Bank sets minimum exchange rate at CHF 1.20 per euro". In: *Communications*(September 2011).

Spange, Morten and Martin Wagner Toftdahl (2014). "Fixed exchange rate policy in denmark". In: *Danmarks Nationalbank monetary review, 1st Quarter*, pp. 49–60.

Sundaresan, Suresh M, Zhenyu Wang, and Wei Yang (2017). *Dynamics of the Expectation and Risk Premium in the OIS Term Structure*. Working Paper.

Tee, Chyng Wen and Christopher Ting (2017). "Variance Risk Premiums of Commodity ETFs". In: *Journal of Futures Markets* 37(5), pp. 452–472.

Trolle, Anders B and Eduardo S Schwartz (2010). "Variance Risk Premia in Energy Commodities". In: *The Journal of Derivatives* 17(3), pp. 15–32. ISSN: 1074-1240.

Verdelhan, Adrien (2018). "The Share of Systematic Variation in Bilateral Exchange Rates". In: *The Journal of Finance* 73(1), pp. 375–418.

Walter, Christian and Jose Lopez (2000). *Is implied correlation worth calculating? Evidence from foreign exchange options and historical data*. Working Paper Series 2000-02. Federal Reserve Bank of San Francisco.

Wooldridge, Philip (2016). "Central clearing predominates in OTC interest rate derivatives markets". In: *BIS Quarterly Review, December 2016*, pp. 22–24.

Wystup, Uwe (2006). *FX Options and Structured Products*. Wiley.

Wystup, Uwe (2007). *FX Options and Structured Products*. The Wiley Finance Series. Wiley. ISBN: 9780470057926.

Zhou, Hao (2017). *Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty*. Review article for the Annual Review of Financial Economics.

Igor Pozdeev

igor.pozdeev@unisg.ch +41-78-612-0928 www.igorpozdeev.me
igor-pozdeev ipozdeev igor-pozdeev

Ph.D. candidate with hands-on experience in a broad spectrum of quantitative activities, deep interest in financial markets, strong coding skills and intention to pursue career at the cutting edge of investment research.

EXPERIENCE

Kraus Partner Investment Solutions — Quantitative analyst — 02/2016 – present
Zurich, Switzerland

- improve existing and develop new econometric models for financial forecasting in diverse asset classes
- guide evolution of offered products by enhancing them with client-requested features
- execute transition to a Python-based environment from MATLAB-, Excel- and VBA-based engine
- support portfolio management team in risk assessment and performance evaluation projects

Swiss Institute of Banking and Finance — Research assistant — 04/2014 – present
St. Gallen, Switzerland

- prepare financial and economic data obtained from professional vendors and public sources
- tutor courses (Financial econometrics, Theory of finance, Introduction to Julia programming)

Third Year — Quantitative research analyst — 10/2015 – 02/2016
Munich, Germany

- designed and implemented in R a nowcasting engine for low-frequency macroeconomic data

KNOWLEDGE & SKILLS

Financial econometrics	time series analysis, volatility modeling, model calibration, nowcasting
Investment theory	asset pricing models, portfolio theory, fixed income, FX, derivatives
Programming	OOP, testing, scientific programming, optimization, HPC, RegEx, visualization
Project management	version control, collaboration tools (Slack), documentation (Sphinx)
Text	plain documents/beamer presentations with MS Office and LaTeX; markdown

PROGRAMMING EXPERIENCE

Python (+NumPy, pandas, matplotlib)	5 years	SQL (+sqlalchemy)	3 years
R, MATLAB, Julia, VBA, hansl	3-4 years	HTML/CSS, LaTeX	1-4 years

EDUCATION

Visiting scholar (host: Robert Engle)	NYU Stern School of Business	2018
Ph.D. in Economics and Finance	University of St. Gallen	2014 – 2019
M.A. in Banking and Finance	University of St. Gallen	2011 – 2014
B.A. in International Economic Relations	Moscow State Institute of International Relations	2008 – 2011

ACADEMIC AWARDS

26th Finance forum Best paper in fixed income	2018
10th FIW research conference Best paper	2017
SFI Best discussant doctoral award	2017

LANGUAGES

English (fluent), German (fluent), Spanish (conversational), Russian (native)