



Københavns Universitet

Event-Related Exchange Rate Forecasts Combining Information from Betting Quotes and Option Prices

Hanke, Michael; Poulsen, Rolf; Weissensteiner, Alex

Published in:

Journal of Financial and Quantitative Analysis

DOI:

[10.1017/S002210901800042X](https://doi.org/10.1017/S002210901800042X)

Publication date:

2018

Document Version

Early version, also known as pre-print

Citation for published version (APA):

Hanke, M., Poulsen, R., & Weissensteiner, A. (2018). Event-Related Exchange Rate Forecasts Combining Information from Betting Quotes and Option Prices. *Journal of Financial and Quantitative Analysis*, 53(6), 2663-2683 . <https://doi.org/10.1017/S002210901800042X>

Event-related Exchange Rate Forecasts Combining Information from Betting Quotes and Option Prices

Michael Hanke, Rolf Poulsen, and Alex Weissensteiner*

Abstract

Betting quotes provide valuable information on market-implied probabilities for outcomes of events like elections or referendums, which may have an impact on exchange rates. We generate exchange rate forecasts around such events based on a model that combines risk-neutral event probabilities implied from betting quotes with risk-neutral exchange rate densities extracted from currency option prices. Its application to predict exchange rates around the Brexit referendum and the U.S. presidential elections shows that these forecasts – conditional on the respective outcomes – were accurate, and markets were able to separate their views on the likelihood and the impact of these events.

*Hanke (corresponding author), Michael.Hanke@uni.li, University of Liechtenstein, Institute for Finance, 9490 Vaduz, Liechtenstein; Poulsen, rolf@math.ku.dk, University of Copenhagen, Department of Mathematical Sciences, 2100 Copenhagen, Denmark; Weissensteiner, Alex.Weissensteiner@unibz.it, Free University of Bozen-Bolzano, School of Economics and Management, 39100 Bozen, Italy. We thank Jennifer Conrad (the editor) and an anonymous referee for their helpful comments, and we thank Johannes Sivén, who provided part of the data and commented on earlier versions of the paper. Poulsen was partially supported by the research center HIPERFIT funded by contract number 10-092299 from the Danish Strategic Research Council.

I Introduction

In a referendum held on June 23, 2016, 51.9% of the voting population of the United Kingdom voted in favor of leaving the European Union (EU), or *Brexit*¹ for short. The net economic effects on the United Kingdom caused by this result were predominantly viewed as negative, which explains the weakening of the pound against major world currencies observed in the days and weeks after the referendum.² The situation of the Mexican peso around the U.S. presidential elections on Nov. 8, 2016, can be viewed in a similar way. During his election campaign, Donald Trump announced that – if elected – he would impose tariffs of 35% on many goods imported from Mexico to protect U.S. industry,³ and he would have a wall built along the U.S.-Mexican border. Although observers questioned the credibility of this claim,⁴ which was quickly downplayed to “just a fence, not a wall” after the election,⁵ remarks like these led many observers to view a

¹Strictly speaking, Brexit is the act or the process of invoking Article 50 in the Lisbon Treaty. When referring specifically to the outcome of the referendum, we will use the terms Leave and Remain.

²See, e.g., articles in the Economist, Feb. 27, 2016, <http://tinyurl.com/zs4o84v> for a view several months before the referendum, and in The Independent, Aug. 23, 2016, <http://tinyurl.com/jyrbwq7> for an assessment two months afterwards.

³Reuters, Mar. 24, 2016, <http://tinyurl.com/jj4bvoy>.

⁴See, e.g., BBC News on Sept. 1, 2016, <http://tinyurl.com/z52yejx>.

⁵The Independent on Nov. 13, 2016, <http://tinyurl.com/z9zcmnl>.

possible win by Trump as detrimental to both U.S.–Mexican relations and the Mexican economy. A foreseeable consequence of a majority voting for Trump was therefore the weakening of the Mexican peso observed in the days and weeks after the elections.

Both the Brexit referendum and the U.S. presidential elections are examples of important events with likely consequences on exchange rates. In situations like these, conditional exchange rate forecasts depending on the outcome of the event would be interesting for investors, internationally active companies, central banks, and other financial market participants. In this paper, we present a novel approach to derive such forecasts, which combines data from financial markets with data from betting markets. As a second contribution, our derivation shows that the mixture of lognormal densities model, which was developed as a purely descriptive model by Ritchey (1990), has a natural economic interpretation in our setting. In a first step, betting quotes are interpreted as state prices and translated to risk-neutral event probabilities. Using these probabilities as weights, risk-neutral densities implied in foreign exchange (FX) option prices are modeled as linear combinations of conditional densities, one for each possible outcome. Based on these, we derive exchange rate forecasts conditional on the respective outcomes. While this approach is used to analyze effects of political events on exchange rates in this paper,

it can be applied more generally to situations where an upcoming event with a discrete set of outcomes is expected to have an effect on the market price of an asset.

The paper is organized as follows: Section II relates this paper to the literature on modeling and forecasting of exchange rates around major events. Section III presents our model, showing in detail how it combines information from betting quotes and option prices. Section IV describes our data and estimation, and Section V presents and discusses the results. Section VI concludes.

II Relation to the Literature: Event Effects in Exchange Rate Markets

This paper contributes to the literature on event-related forecasting of exchange rates. A number of papers in this area use information from option prices to estimate models for the development of exchange rates around major events. In what follows, we relate the main contributions of our paper to this strand of literature.

A well-known paper in this area is Malz (1996), who models the development of the GBPDEM exchange rate in the years 1990–1992. In 1990, the British pound became part

of the so-called Exchange Rate Mechanism. A target rate was fixed towards the European Currency Unit, together with a $\pm 6\%$ fluctuation band. Political and economic shocks in the early 1990s led to speculation about a possible realignment of this target rate. Malz (1996) models the exchange rate with a jump-diffusion process, which he estimates using option price data. However, contrary to the events used here (the Brexit referendum and the U.S. presidential elections), the timing of a possible realignment was uncertain. Leahy and Thomas (1996) analyze changes in risk-neutral densities extracted from exchange-traded FX options around the Quebec sovereignty referendum in Oct. 1995. Using a mixture of three lognormal densities to model the option-implied exchange rate density, they estimate expected values and standard deviations as well as the weights (mixture parameters) from options data. The shape of the resulting implied densities varies markedly over time, ranging from unimodal densities that are close to lognormal to multimodal densities with up to three peaks. This variation seems to be caused by the large number of estimated parameters. Castrén (2005) analyzes the effects of monetary policy decisions on the risk-neutral densities (RNDs) extracted from options on currencies of new EU member states. Similar to Leahy and Thomas (1996), his work is quite close in spirit to the present paper in that it also analyzes the impact of external

events on risk-neutral densities. His events, however, are more frequent, and have a range of possible outcomes (e.g., interest rate changes of unknown magnitudes). Moreover, whereas we are interested mainly in the expected exchange rate conditional on one of the possible outcomes (which leads us to focus on one of the component densities of the RND, see Section III), Castrén analyzes changes in the moments of the entire RNDs prior to these events.

More recently, in Sept. 2011, the Swiss National Bank (SNB) announced a minimum exchange rate policy regarding the EURCHF exchange rate. The Swiss franc had appreciated considerably in the years before, particularly since the beginning of the financial crisis in 2007. This unusual measure was taken to counteract and/or delay a further strengthening of the Swiss franc, whose apparent overvaluation was considered to be harmful to the Swiss economy. From the very beginning, it was clear that this measure would be temporary, but similar to the situation analyzed by Malz (1996), the date of the revocation of this policy was not known to the market. A number of papers (Hanke, Poulsen, and Weissensteiner (2015), Hanke, Poulsen, and Weissensteiner (2016), Hertrich and Zimmermann (2017), Jermann (2017), Mirkov, Pozdeev, and Söderlind (2016)) analyze various aspects of the EURCHF exchange rate development during this

period. In particular, the question where the exchange rate would have been without the SNB's minimum exchange rate policy is closely related to the topic of the present paper, because it can be framed as forecasting the new level of the exchange rate after a revocation of this policy. Mirkov et al. (2016) analyze whether FX options data indicate that markets anticipated the SNB's revocation of its minimum exchange rate policy. To this end, they estimate risk-neutral densities from FX option prices based on a mixture of lognormal densities model.

All of these papers use information extracted from option prices in their estimations.⁶

Although, none of them use a separate source of information for the event probability, which is the main innovation of the present paper. Using only one source of information leads to the estimation problems described above, making it difficult to separate the influence of different parameters (see, e.g., Malz (1996), pp. 732 ff.). A common way to address this problem is by making simplifying assumptions, which impose more structure

⁶Whereas all other papers directly use information from traded FX options, Hanke et al. (2016) follow Hanke et al. (2015) in viewing part of the observed exchange rate during the SNB's minimum exchange rate policy period as an option on the latent, unobserved exchange rate that would have been observed in the absence of the SNB's policy. Whereas the latter paper combines this modeling approach with the prices of traded options, Hanke et al. (2016) use only the observed exchange rate process itself for the estimation of the latent exchange rate.

on the problem. E.g., Hanke et al. (2015, pp. 1111 ff.) assume a linear dependence of event probabilities on time to stabilize the estimation, and Mirkov et al. (2016, Section 3.1) mention that they need to impose some restrictions to mitigate adverse effects on the parameter estimation. Despite this simplifying assumption by Hanke et al. (2015), they still report noisy estimates due to substitutive effects from estimating four parameters simultaneously with option prices as the only source of information. Hertrich and Zimmermann (2017, p. 571) require a specific assumption on the implied volatilities in order to be able to estimate the barrier level in their model, a parameter that is closely related to our event probability. Similar to Hanke et al. (2015), but based on a binomial tree model, Jermann (2017) also estimates the latent exchange rate and the event probability simultaneously from option prices. His Figure 1 (p. 560) illustrates heavy zig-zagging of probabilities, particularly for the first half of his sample period. Also in his setting, the complex nonlinear interactions of parameters make it difficult to separate the effect of event probabilities when using only option prices in the estimation process, leading to unstable parameter estimates. The approach presented here overcomes this problem by using betting quotes as a separate source of information for risk-neutral event probabilities, which reduces the number of parameters to be estimated from option prices

and stabilizes the estimation.

Our approach, which is also based on a mixture of lognormal densities, is closely related to Leahy and Thomas (1996) and Mirkov et al. (2016). However, there are two important differences: First, whereas these papers estimate the mixture parameter together with the means and standard deviations of the two component densities from option prices, we use risk-neutral event probabilities implied from betting quotes. This reduces the number of parameters to be estimated from option prices and avoids the need for any restrictions in the estimation. Second, they use the mixture of lognormals in a purely descriptive way, i.e., to model nonnormal risk-neutral return densities while still retaining analytical tractability. A second contribution of the present paper is the interpretation of the component densities in our mixture of lognormals model as conditional densities for each of the possible event outcomes. A detailed description of our model is provided in the following section.

III Methodology

In the first part of this section, we show how betting quotes can be translated to risk-neutral event probabilities. Using these probabilities as weights, risk-neutral densities implied in FX option prices are modeled as linear combinations of conditional densities, one for each possible outcome. The component densities used in this model allow for a straightforward calculation of conditionally expected exchange rates. The last part of this section discusses further potential applications of our approach.

A Risk-neutral Event Probabilities from Betting Quotes

Our approach relies on using a separate source of information (different from option prices) to infer event probabilities. From a financial economics point of view, the data most directly related to event probabilities are quotes from liquid betting markets.

Section IV will address the liquidity and other important aspects of the betting markets for the Brexit referendum and the U.S. presidential elections 2016. Here, we show how betting quotes can be transformed to risk-neutral event probabilities by interpreting them as state prices for the possible outcomes of the events. Afterwards, we argue why

betting quotes are superior in this regard compared to alternatives like polls.

For ease of exposition, we explain our model based on the Brexit case. Later in this section, we will discuss the similarities and differences regarding our second application case, the U.S. presidential elections. We will use the U.S. dollar as the numeraire currency in our model. Let time τ be the day after the event, when information about its result is known by the market. There are two possible states, Leave (L) and Remain (R). Betting payoffs at time τ can be related to state prices ψ_L (ψ_R) at time t , i.e., the prices of the Arrow securities paying off 1 in case of state Leave (Remain) and 0 otherwise. The state prices can be computed as the payoffs' discounted expected values under a risk-neutral measure:

$$(1) \quad \begin{aligned} \psi_L &= e^{-r(\tau-t)} E_t^Q [1_L] \\ &= e^{-r(\tau-t)} q_L, \end{aligned}$$

where r is the U.S. risk-free interest rate, 1_L denotes the indicator function taking on a value of 1 U.S. dollar if the referendum outcome is Leave and 0 otherwise, and q_L is the risk-neutral Leave probability (q_R is defined analogously). Equation (1) directly links

betting quotes (state prices) and risk-neutral event probabilities. Computing these probabilities from betting quotes and using them in the process of extracting risk-neutral densities from option prices is a distinctive feature of our approach.

Aside from betting markets, polls are an alternative indicator for the prospective outcomes of events like popular referendums or elections. The question which of the two approaches provides better *forecasting* results has not yet been answered definitively (see, e.g., Erikson and Wlezien (2012)). Our goal of *inferring event probabilities*, however, is different from forecasting. We use betting quotes because they can be directly translated to (risk-neutral) event probabilities as shown in equation (1). In contrast, the relation between poll results and event probabilities does not only depend on the exact phrasing of the question asked (intended voting behavior vs. expected result), but also on voters' preference distributions. For this reason, poll results cannot be directly translated to event probabilities. In addition, betting money on a particular outcome reflects much more of a commitment than simply answering a question, and polls on political events frequently suffer from biases because of people's unwillingness to reveal their opinion or planned voting behavior.

B Risk-neutral Densities Implied in Option Prices

A weakening of the British pound could in principle be measured against any major world currency. We choose the GBPUSD exchange rate because it has by far the highest trading volume among all these currency pairs. Another reason for favoring the GBPUSD rate over the alternative of the EURGBP rate are potential indirect effects of Brexit on the euro aside from its consequences for the pound, which would lead to a combination of these effects being reflected in the EURGBP rate. In principle, such follow-on effects might also influence the GBPUSD rate through the exchange rate triangle via the euro, but with a much smaller magnitude.

Denote the GBPUSD exchange rate (which, by market convention, is the amount of U.S. dollars corresponding to 1 British pound) at time t by S_t , and denote the risk-neutral density of S_τ as seen at time t by $\phi_{t,\tau}$. If a smooth strike-continuum of expiry τ call option prices $C_t(K, \tau - t, \cdot)$ were observable at time t , then this density could be inferred in a model-free way. As shown first by Breeden and Litzenberger (1978) we have

$$(2) \quad \phi_{t,\tau}(y) = e^{r(\tau-t)} \frac{\partial^2 C_t(K, \tau - t, \cdot)}{\partial K^2} \Big|_{K=y}.$$

In reality, only a discrete set of strikes are available, which means that a considerable amount of smoothing and/or curve fitting is needed before equation (2) can be used to produce sensible results. Figlewski (2010, Section 2) provides a taxonomy of methods that can be used for this purpose. We base our model on the mixture of lognormal densities approach, which has been suggested by Ritchey (1990) and used in the present context by Leahy and Thomas (1996), Castrén (2005), and Mirkov et al. (2016). We modify the approach to incorporate risk-neutral event probabilities, which gives the weights and the component densities a straightforward economic interpretation. This distinguishes our model from the sources cited before, which use the mixture of densities as a purely descriptive model.

Following Ritchey (1990), we decompose the risk-neutral *return* density $\phi'_{t,\tau}(\cdot)$ into the sum of two normal densities (this corresponds to a mixture of lognormal densities for exchange rate *levels* and is therefore often referred to as a *mixture of lognormal densities* model):

$$(3) \quad \begin{aligned} \phi'_{t,\tau}(\omega, \mu_1, \sigma_1, \mu_2, \sigma_2, \tau - t) &= \omega f(\mu_1(\tau - t), \sigma_1 \sqrt{\tau - t}) \\ &\quad + (1 - \omega) f(\mu_2(\tau - t), \sigma_2 \sqrt{\tau - t}), \end{aligned}$$

where ω is the mixture parameter, and $f(\cdot)$ denotes the density of a normal distribution. In Ritchey (1990), the mixture of lognormal densities approach is motivated as a purely descriptive model: Starting from the problem of pricing options for observed nonnormal return densities (with nonzero skewness and excess kurtosis), a mixture of normals can approximate such densities while retaining analytic tractability. In our model, however, both the component normals and the weights have a straightforward economic interpretation: We assume that conditional on the outcome of the event, risk-neutral exchange rates are lognormally distributed. The risk-neutral return density before the event is a mixture of two normal densities, weighted by the risk-neutral event probabilities. Based on this interpretation, equation (3) can be rewritten as follows:

$$(4) \quad \phi'_{t,\tau}(\cdot) = q_L \phi'_{t,\tau}(\cdot|L) + q_R \phi'_{t,\tau}(\cdot|R).$$

Put differently, we assume that the Garman and Kohlhagen (1983) model⁷ holds

conditional on each possible outcome of the event, and the resulting option prices are

⁷The Garman and Kohlhagen (1983) model is a variant of the Black and Scholes (1973) model for European currency options. The corresponding pricing formulas are sometimes stated in terms of forward exchange rates, which corresponds to the Black (1976) model. Independent of the notation, this is the standard market model for currency options.

linear combinations of Garman and Kohlhagen (1983) option prices with different parameters:

$$(5) \quad C^{\text{MLN}}(\phi'_{t,\tau}, \cdot) = q_L C_L(f(\mu_L, \sigma_L), \cdot) + q_R C_R(f(\mu_R, \sigma_R), \cdot),$$

where $C^{\text{MLN}}(\phi'_{t,\tau}, \cdot)$ denotes the call price in the mixture of lognormal densities model, and $C_L(\cdot)$ ($C_R(\cdot)$) are the Garman and Kohlhagen (1983) option prices conditional on the respective event outcomes.

In previous applications (see, e.g., Mirkov et al. (2016)), all the parameters in equation (3) (mixture parameter ω , expected returns μ_j and standard deviations σ_j) are estimated from derivatives data. In contrast, we estimate only expected returns and standard deviations from derivatives data. As the mixture parameter, we use the risk-neutral event probability q_L derived from betting quotes, resulting in weights of q_L and $q_R = 1 - q_L$. This reduces the number of parameters to be estimated by one and leads to stable parameter estimates. Our modeling approach implies an additional restriction on the difference between the means of the conditional densities, which saves another parameter. As a result, we do not need to impose restrictions as in Mirkov et al. (2016,

Section 3.1), which they require to mitigate adverse effects on the parameter estimation caused by noisy data. Details of the estimation procedure will be described in Section IV.

Based on the parameters of the conditional return distributions, the calculation of conditional exchange rate expectations is straightforward: The risk-neutral expected exchange rate for an outcome of Leave (analogously for Remain) is

$$(6) \quad E_t^Q(S_\tau | L) = S_t e^{(\mu_L + 0.5\sigma_L^2)(\tau-t)}.$$

With all our calculations carried out under risk neutrality, an important question is whether the risk-neutral conditional expectations $E_t^Q(S_\tau | \cdot)$ from equation (6) are accurate estimators of their corresponding real-world expectations, $E_t^P(S_\tau | \cdot)$, and thus in turn of realized values. For longer forecasting horizons, this depends on the exchange rate risk premium: For a (near) zero risk premium, risk-neutral and conditional expectations would (almost) coincide. Popular approaches to estimating this risk premium include analyzing the performance of carry trades (see, e.g., James, Fullwood, and Billington (2015)), or extracting implied risk premia from options data (Jurek and Xu (2014)). We will discuss the possible effect of a nonzero risk premium on our estimated conditional

exchange rate expectations in Section V.

C Required Adaptations for the Case of the U.S. Presidential Elections

For ease of exposition, the description of our model in the previous subsections used the terminology of the Brexit referendum case. Regarding our second application, the development of the USD/MXN rate around the U.S. presidential elections in Nov. 2016, equations (1) to (6) and the corresponding discussion carry over directly to the MXN/USD rate, which uses the U.S. dollar as the numeraire. Options on the USD/MXN rate use the Mexican peso as the numeraire. The premium currency for these options, however, is the U.S. dollar. This must be taken into account by making a premium adjustment to the deltas when converting USD/MXN implied volatilities to option prices (see Section IV). The analogues of Leave and Remain are Trump and Not Trump. Selecting “Not Trump” rather than “Clinton” reflects the fact that there were more than two candidates. Moreover, only the potential win of one of these candidates (Trump) was viewed as potentially detrimental to the Mexican economy because of the remarks he made during

his election campaign.

D Further Applications

While the applications presented in this paper focus on exchange rate forecasts related to political events, we note that the approach presented here can be viewed more generally.

It can be applied also in other situations that satisfy the following three criteria regarding events, betting markets, and option markets: (i) There is an upcoming event with two possible outcomes.⁸ The event date is fixed and publicly known in advance, and the event is expected to have an effect on the market price of an asset. (ii) There exists a liquid betting market for the possible outcomes of the event. (iii) There exists a liquid market for options on the affected asset with a sufficient variety of strike prices. An example for such a situation would be Federal Open Market Committee (FOMC) meetings and their effects on interest rates. However, whereas there are liquid markets for options on interest rates and interest-sensitive securities, the liquidity of betting markets on FOMC decisions is lower compared to major political events like elections or important referendums.

⁸Our model is formulated for two possible outcomes, but the extension to more than two outcomes is straightforward. However, the limited number of available strikes may lead to difficulties when trying to account for more outcomes, which requires estimating more parameters.

In our applications to the Brexit referendum and the U.S. presidential elections 2016, the direction of the change in the exchange rate after the event (conditional on the possible outcomes) seems clear ex ante, and we use the approach to assess primarily the magnitude of the effect. In other applications, it might be less clear which outcome would be expected to have a positive/negative effect, or whether or not the market expects different event outcomes to have any effect at all. The approach presented in this section can also be used to answer these questions: If the market expects little to no difference regardless of the outcome that occurs, the conditional densities should be essentially the same.

IV Data and Estimation

In the first part of this section, we describe the betting quotes and options data we use in this paper. The second part discusses in detail how we use these data in the estimation process.

A Data

We retrieve daily exchange rate data from Bloomberg (GBPUSD) and from Datastream (USDMXN). The development of the exchange rates around the date of the respective events is shown in Figure 1. For the GBPUSD rate, the level shift after the referendum leads to exchange rates around 1.29–1.33, appr. 10% lower than in the first half of the year and the lowest levels in the past three decades. The USDMXN rate increased from 18.55 on Nov. 8 to values around 20.5 (with considerable variation, appr. plus 10%) in the following days, the highest levels in its history taking currency reforms into account.

Figure 1 about here

The betting quotes are retrieved from <http://betdata.io>. This Web site provides the entire available history of quotes, averaged across a range of platforms and bookmakers.

While cross-checking the data with other sources, similar estimates of the betting quotes were retrieved.⁹ The popularity of betting in various regions of the world (including the UK) is not limited to sports events but extends also to other areas like politics.

Internet-based betting markets attract punters from all over the world. The size and

⁹In particular, we have no indication of systematic differences across geographical regions although theoretically, the risk-neutral event probabilities might differ slightly depending on the numeraire currency.

liquidity of these markets make quotes for bets on major events competitive and readily available. Given the importance of both the Brexit referendum and the U.S. presidential elections, the popularity of betting on these events does not come as a surprise. Betfair, the largest internet betting exchange, reported a total betting volume of GBP 127 million for the Brexit referendum¹⁰ and an all-time-record of just under GBP 200 million for bets on the U.S. presidential elections.¹¹ While these numbers are high enough to indicate reasonable liquidity, they are dwarfed by the trading volume in FX options. In April 2016, the average daily turnover for options on the GBPUSD (USDMXN) currency pair was around USD 21.7 billion (USD 5.2 billion) according to the BIS Triennial Central Bank Survey 2016.¹²

Figure 2 depicts the development of Leave quotes (state prices ψ_L , left) and Trump quotes (ψ_T , right) over the months before the respective events. The betting quotes are denominated in GBP. Before they can be used in equation (1), they are converted to USD using the (theoretical) forward GBPUSD rate based on the interest rate difference.

Figure 2 about here

¹⁰Betfair, <https://tinyurl.com/lgjzo3a>.

¹¹Betfair, <https://tinyurl.com/m4mrkpn>.

¹²Bank for International Settlements, <http://www.bis.org/publ/rpfx16fxt.pdf>.

For most of the time between March and mid-May 2016, risk-neutral Leave probabilities implied in these betting quotes were between 25% and 35%. In the second half of May, they took a dive and hit a low around 17% almost exactly one month before the referendum. Afterwards, they increased strongly, reaching a peak of 40% about 10 days before the referendum, before falling sharply to 23% immediately before the referendum. Risk-neutral Trump probabilities, in contrast, fluctuated more and reached their low at around 11% in April 2016. After a peak at around 34%, they declined to 15% in October and increased again to just below 30% at the beginning of November, before falling again to values between 15% and 20% just before the elections. Interestingly, both Leave and Trump probabilities implied in betting quotes declined in the last few days before the respective events. For both cases, this move is in line with media reports on opinion polls shortly before the events, which indicated a majority for Remain¹³ and for the runner-up in the presidential elections, Hillary Clinton.¹⁴

Implied option volatilities are retrieved from the same sources as the respective exchange rates. Compared to exchange-traded FX options, OTC-traded FX options have the

¹³Reuters, June 23, 2016, <http://tinyurl.com/1f7dtag>.

¹⁴New York Times, Nov. 8, 2016, <http://tinyurl.com/128drbt>.

advantage of constant maturities, i.e., every day “new” options with certain maturities such as one month (1M) or three months (3M) are traded. Prices and implied densities of such options are easier to compare over time than those of options with decreasing maturities. There are some peculiarities regarding quotes in these markets, which affect the way options data are reported (for a general introduction, see e.g. Clark (2010), Reiswich and Wystup (2010)). Currency options are frequently traded as combinations of options, which are known as straddles, strangles, and risk reversals. Implied at-the-money volatilities are quoted for delta-neutral straddles. Exercise prices for in- and out-of-the-money options are defined in terms of certain values of delta: 10D means a delta of 0.1, 25D a delta of 0.25, etc. The details of the delta definitions vary across currency pairs because of market conventions. Following Clark (2010, Chapter 3), we use the so-called pips delta for the GBPUSD options and the percentage delta for USDMXN options, which accounts for the differences between numeraire and premium currency we discussed in Section III. The interest rates we use are the respective 1M and 3M LIBOR rates for USD and GBP for the Brexit case, and the 1M and 3M U.S. T-bill and 28-/91-day Mexican CETES yields for the U.S. election case.

We obtain at-the-money implied volatilities, together with implied volatility spreads for

10D and 25D strangles and risk reversals. From these, we compute implied volatilities of calls with moneyness 10D, 25D, ATM, 75D, and 90D, together with their respective strikes (see, e.g., Clark (2010), Chapter 3). The development of the ATM, 10D and 90D implied volatilities is shown in Figure 3. For the GBPUSD rate (left), there is a slight increase in implied volatilities of both 1M and 3M options in the first quarter.

Afterwards, implied volatilities of 1M options hover around 10%, whereas those of 3M options shoot up exactly on the day when 3M options expiring after the referendum were traded for the first time. The same happens to 1M options two months later. These increases were caused by the high uncertainty around the referendum. After the referendum, volatilities decline again and reach pre-referendum levels within four to five weeks. The USDMXN implied volatilities (right) show similar jumps on the days when options expiring after the elections were first traded. The magnitude of the jumps is higher for 1M options in both cases, which we attribute to the higher relative importance of the event relative to other factors of influence during the options' lifetimes. The difference in magnitudes between 1M and 3M options is markedly more pronounced for USDMXN options, which corresponds well to changes in betting probabilities for these events: Whereas Leave probabilities implied in betting quotes decrease between the dates

when affected 3M and 1M options on the GBPUSD rate were first traded, Trump probabilities increase between the respective dates (see below and Figure 2).

Figure 3 about here

With betting quotes implying a (risk-neutral) Leave probability of around 1/3 between Feb. and mid-May 2016 (see Figure 2), the risk of Brexit was widely known and clearly nonnegligible. This suggests that many market participants with FX exposures involving the pound at least considered hedging their exposures. Fig. 4 illustrates that the put–call ratio for the open interest in GBPUSD options traded on the Chicago Mercantile Exchange (CME) showed a clear upwards trend in the first half of 2016.¹⁵ The higher trading activity in puts indicates an increase in hedging against and/or speculation on a drop in the exchange rate.

Figure 4 about here

¹⁵Note that this put–call ratio is computed from exchange-traded calls and puts, as opposed to the OTC-traded option combinations used in this paper, which are quoted for combinations of equal numbers of calls and puts.

B Estimation

To estimate the parameters of the conditional densities in equation (5), we first convert implied volatilities for calls with moneyness 10D, 25D, ATM, 75D, and 90D to call prices and corresponding strikes using the Garman and Kohlhagen (1983) model. This gives us five call prices C_k^{MKT} , $k = 1, \dots, 5$ per trading day. The decomposition described in equation (5) for call prices must also hold for the exchange rate itself, i.e., the observed exchange rate at time t must correspond to its discounted expected value under the risk-neutral measure:

$$\begin{aligned}
 (7) \quad S_t e^{(r-i)(\tau-t)} &= q_L E_t^Q(S_\tau | \text{L}) + q_R E_t^Q(S_\tau | \text{R}) \\
 &= q_L S_t e^{(\mu_L + 0.5\sigma_L^2)(\tau-t)} + q_R S_t e^{(\mu_R + 0.5\sigma_R^2)(\tau-t)},
 \end{aligned}$$

where i denotes the foreign risk-free interest rate. Solving (arbitrarily) for μ_R , this provides a restriction on the estimation, which saves one additional parameter to be estimated. We estimate the parameters in equation (5) for each trading day by

minimizing the (weighted) sum of squared relative pricing errors:

$$(8) \quad \min \sum_{k=1}^5 w_k \left(\frac{C_k^{\text{MKT}} - C_k^{\text{MLN}}(\mu_L, \sigma_L, \sigma_R, \cdot)}{C_k^{\text{MKT}}} \right)^2,$$

where C_k^{MKT} are the market prices obtained from implied volatility data, C_k^{MLN} are the corresponding model prices computed from equation (5), and w_k are weights. In the literature, minimizing either absolute or relative squared deviations has been suggested (Christoffersen and Jacobs (2004)). We prefer to minimize relative instead of absolute squared deviations to avoid implicitly overweighting highly priced options, which would lead to higher pricing errors for out-of-the-money calls. Following Christoffersen and Jacobs (2004, equation (2)), we give all options equal weight by using the same value for all w_k . As a robustness check, we will also show results for the alternative of using implied vegas as weights, which is another weighting scheme commonly used.

V Results

Throughout the results section, we repeatedly focus on 1M and 3M currency options on the first day when such options were traded that expired after the respective events.

Conditionally expected exchange rates derived from these options can be interpreted as the market's exchange rate forecasts for the first trading day after the Brexit referendum (the U.S. presidential elections), depending on the respective outcomes.

A Model Fit and Parameter Estimates

We estimate the model using equation (8) three months and one month before the respective events. The resulting relative price differences for each of the five options used in the estimations are shown in Table 1. GBPUSD options are in the left part of the table, and USDMXN options in the right part. The mixture of normal densities model in equation (5) fits the data well. The highest differences between market and model prices are less than 1.5%, indicating high accuracy of the mixture of normals model. The first two columns for each currency show results for equal weighting of all options, while the following two columns provide a robustness check using implied vegas as weights. The errors are very similar for both parts, indicating robustness with respect to different weighting schemes.

Insert Table 1 about here

For comparison, we repeat the estimation in equation (8) using the Garman and

Kohlhagen (1983) model instead of the mixture of lognormal densities model, assuming a flat volatility across strike prices. The model fit is markedly worse, leading to relative pricing errors of up to 30% for GBPUSD options and up to 32% for USDMXN options. These systematic deviations between market prices and Garman and Kohlhagen (1983) model prices are taken into account by traders through quoting different implied volatilities across strike prices, leading to the well-known volatility smiles.

B Risk-neutral Densities

The parameter values of the conditional normal densities estimated using equation (8) are shown in Table 2. For each of our two events, both the means and the standard deviations of these densities differ markedly depending on the respective outcomes. The conditional volatility σ_L (σ_T), which denotes the volatility in case of Leave (Trump), is greater than that for Remain (Not Trump) across all option maturities and weighting schemes. The market seems to have expected higher uncertainty after an outcome of Leave compared to Remain, and after a win by Trump compared to Not Trump. The influence of the weighting scheme is very small, indicating again high robustness.

Insert Table 2 about here

Figure 5 shows RNDs from the mixture of lognormal densities model in equation (5) (solid curves), estimated using equation (8). The parameters of the respective component densities are shown in Table 2. The moments of the RNDs depicted in Figure 5 are summarized in Table 3. In Figure 5, the estimated mixed densities are compared to the best-fitting densities from the standard Garman and Kohlhagen (1983) model (dashed curves). Whereas the latter assumes normally distributed returns, the former assumes normality for the conditional return densities, which allows for nonzero skewness and excess kurtosis in the mixed or aggregate densities. The top (bottom) panel in this figure compares RNDs from 3M (1M) options, again estimated on the first day when such options were traded that expired after the respective events (solid curve). GBPUSD densities are shown on the left, and USDMXN densities on the right. Comparing the dashed to the solid curves, there are noticeable differences: For GBPUSD, the mixed densities have higher standard deviations (5.25%/8.60% for 1M/3M options) than the normal densities implied from the Garman and Kohlhagen (1983) model (time-scaled to the respective option maturities of 1M/3M: 4.06%/6.79%). In contrast, the standard deviations of the Garman and Kohlhagen (1983) model for USDMXN, with values of 6.72%/8.46%, are higher compared to the corresponding values for the mixed normal

densities in Table 3.

Table 3 about here

Whereas the USDMXN densities show only moderate deviations from normality with skewness around 0.5,¹⁶ the GBPUSD densities deviate more strongly from normality with skewness below -1 and higher kurtosis. In addition, the 1-month GBPUSD density (lower left in Figure 5) shows a small hump in its left tail, caused by a large difference in means relative to the standard deviations of its components (see column 2 in Table 2). For the USDMXN options, the 1-month density is closer to normal, whereas for GBPUSD options, the 1-month density shows higher skewness, but less kurtosis compared to the 3-month density.

Figure 5 about here

C Conditional Exchange Rate Expectations

Plugging the parameter estimates provided in Table 2 into equation (6), we calculate the expected exchange rates conditional on the outcomes that actually occurred 3 months (1

¹⁶Note that our use of the common notation conventions leads to negative skewness for the GPBUSD rate, but positive skewness for the USDMXN rate.

month) later. For the Brexit referendum, we obtain conditionally expected exchange rates of 1.33 (1.34). Note that these forecasts are made at points in time that are two months apart and for different spot exchange rates of 1.41 (1.46). The forecasts are very close to the exchange rates observed in the days after the referendum, which were in the range of 1.29–1.33. For the U.S. presidential elections, we obtain conditionally expected exchange rates of 19.71 (20.36). These values are also close to the observed exchange rates after the elections, which were 19.97 (20.63) on the first (second) day after the elections.

As discussed briefly in Section III, risk-neutral and real-world conditional exchange rate expectations should be very close to each other if the exchange rate risk premium and/or the time to maturity are small. The typical size of exchange rate risk premia over short forecasting horizons is much smaller than the magnitude of the expected event-induced changes we calculated for our two applications. Bounds on the difference between risk-neutral and real-world exchange rate expectations can be derived using the two approaches to estimating risk premia mentioned in Section III. Regarding the first approach based on analyzing carry trades, James et al. (2015) conduct an exhaustive investigation of G10 countries and conclude that *on average, the carry trade delivers almost exactly the carry* (i.e. the interest rate differential). This corresponds to exchange

rates being almost drift-free (i.e. martingales) under the real-world probability measure, implying that the risk-neutral drift rate is close to the carry. For the time period we consider, the GBPUSD interest rate differential was close to 0. In such a setting, risk-neutral expectations should be very close to expected realized values. For the USDMXN rate, the interest rate differential was higher, on the order of 3% p.a., which would imply a bias between risk-neutral and real-world exchange rate expectations of 0.75% (0.25%) for the three-months (one-month) horizon. Regarding the second approach based on risk premia implied from FX options, Jurek and Xu (2014, Fig. 3) estimate a risk premium implied in GPBUSD options of less than 2% per year. Accounting for a risk premium of 2% would lead to a shift in conditional expected exchange rates estimated from 3M (1M) options by 0.5% (0.167%) between the risk-neutral and real-world probability measures. In the light of predicted changes in the GBPUSD exchange rate on the order of 10%, this is a small magnitude for potential biases. We could not find similar option-implied estimates for the USDMXN risk premium. For other currency pairs relative to the U.S. dollar, Jurek and Xu (2014) estimate risk premia of less than 4% for most currencies, with few exceptions reaching risk premia of up to 6% annually. Using these values, the resulting difference between

risk-neutral and real-world expectations would still be small compared to the magnitude of the changes in rates we observe. Given that risk premia show considerable day-to-day fluctuations, and acknowledging that our estimations are subject to uncertainty from other sources as well, we interpret short-term risk-neutral conditional expectations of the GBPUSD rate as good (i.e., close to unbiased) estimators for the corresponding real-world conditional expectations. Hence, the risk-neutral conditional expectations are used in the rest of this section without any adjustments.

As discussed before, in the absence of extraordinary events such as the Brexit referendum, the GBPUSD exchange rate over short periods of time can be regarded as essentially a random walk without drift. As the interest rate differential was very small during our sample period, this holds approximately for both the risk-neutral and the real world. This allows us to estimate the expected GBPUSD exchange rate – conditional upon a majority voting for “Leave” at the referendum – from a rolling window, which starts on Mar. 23 and extends until the referendum. The values shown in the left panel of Figure 6 are three-months ahead forecasts of this conditionally expected GBPUSD rate. For most of April and May, the range of the estimated values is reasonably close in level to the realized exchange rates after the referendum. Only in June do we see larger

deviations, first below 1.30, and then above 1.40 on the last day before the referendum.

Figure 6 about here

Three months before the referendum, our conditional GBPUSD expectation for the day after the referendum was 1.33, which decreased to values slightly below 1.30 one month before the referendum. The rolling window estimates in the left panel of Figure 6 end at 1.41 on the day of the referendum, which is the forecast for three months after the referendum. Since the rolling estimates in this figure are derived under the assumption that the real-world exchange rate process follows a random walk, 1.41 can be interpreted as the market's best and most recent estimate for the conditional exchange rate in case of Leave. Both the conditionally expected values three months and one month before the referendum estimated at 1.33 (1.34) and most of the rolling estimates in Figure 6 are inside of the interval of the exchange rate values actually observed in the days and weeks after the referendum. Only the rolling estimates in late April/early May and on two of the last few days before the referendum are higher.

The right panel of Figure 6 shows the analogous estimates for the expected USDMXN rate conditional upon a majority voting for Trump.¹⁷ From September onwards, these

¹⁷Due to the higher interest rate differential of this currency pair, the risk-neutral drift of the USDMXN

conditional forecasts are between 20 and 21.3, which is again close to the realized exchange rates after the elections. The conditional forecasts three months and one month prior to the U.S. presidential election are not far apart at 19.71 (20.36), despite marked differences in the corresponding Trump probabilities. Both forecasts, as well as most of the rolling estimates in the right panel of Figure 6, are quite close to the exchange rates actually observed on the days after the election, which increased from 19.97 on Nov. 9 to 20.63 on Nov. 10.

An interesting question is which forecasts our model would have given for the counterfactual event outcomes, i.e., “Remain” for the Brexit referendum and “Not Trump” for the presidential elections. This is provided by Figure 7, which compares conditional expectations for the two possible outcomes Leave/Remain (left) and Trump/Not Trump (right). For the Remain case, our model would have predicted GBPUSD levels about 4% higher than observed in most of 2016. For the counterfactual outcome of Trump not winning the election, our model gives conditional USDMXN forecasts very close to observed exchange rates in the run-up to the elections. These ranges are a result of equation (4), which implies that the observed exchange rate can be

rate will be slightly further away from 0.

viewed as the discounted risk-neutral expected exchange rate, when taking the respective risk-neutral event probabilities into account (see Figure 2).

Figure 7 about here

D Further Robustness Considerations

Our results depend on two inputs to our model: Option prices and betting quotes. A natural question is which of the two inputs may be more important, e.g., are the results primarily driven by the betting probabilities? Also, our combination of betting quotes and option prices may raise some skepticism: Is there really a connection between betting quotes and financial markets?

We start by noting that the sample correlation between the levels of risk-neutral Leave probabilities and the conditional GBPUSD exchange rate expectations shown in Figure 6 is only 33% (0.4% for USDMXN, not shown), and the sample correlation of relative changes in these quantities is 37% (-37%). These values indicate that our results are not primarily driven by the betting probabilities. It may be tempting to interpret the rather weak correlation to indicate that betting quotes and financial markets are only loosely related. However, note that even though Leave probabilities implied in betting quotes

fluctuated considerably between 17% and over 40%, the conditionally expected exchange rate was remarkably stable, showing values between 1.30 and 1.36 for most of the time during the three months before the referendum. Our interpretation of this is that markets were able to correctly separate the probability of Brexit from its potential consequences: Irrespective of changes in current estimates for Leave probabilities, *Brexit meant Brexit* for financial markets in terms of expected consequences, to modify a quote from British prime minister Theresa May. For our analysis, this means that the combination of information implied in option prices and in betting quotes yields insights that could not have been derived from either of the two data sources used in isolation.

As a further check for our approach, we compute the volatilities of our conditional GBPUSD exchange rate forecasts for the post-referendum time windows shown in Figure 6. Other approaches from the literature that generate exchange rate forecasts from option prices suffer from excessive volatility in these forecasts (see, e.g., Hanke et al. (2015), Fig. 3), which require some form of ex-post smoothing to arrive at sensible forecasts. For the volatility of expected exchange rates (annualized) in case of Leave (Remain), we get values of 21.2% (8.1%), which are quite close to normal volatility levels for the exchange rate itself: The GBPUSD realized volatility between Jan. 1 and June 23

was 10.3%, and an increase (decrease) in volatility in case of Leave (Remain) is in line with economic intuition. We conclude that the approach presented here produces forecasts with very low levels of noise.

VI Conclusion

We showed how risk-neutral event probabilities implied in betting quotes can be used in the estimation of risk-neutral densities from FX option prices. Our model delivers not only the risk-neutral exchange rate density but also the conditional densities depending on the respective outcomes. From these conditional densities, conditional exchange rate forecasts can be derived. We applied the model to the recent cases of the U.K. Brexit referendum and the U.S. presidential elections. The forecasts we derived conditional on the outcomes that actually occurred are quite close to the exchange rate values observed after these events. A more detailed analysis reveals that the forecasts are not mainly driven by betting quotes: Fairly stable exchange rate forecasts despite large variations in betting quotes are taken as an indication that markets had a fairly clear view on the size of a potential drop in case of Leave (Trump), and they adjusted option prices and implied

densities accordingly when betting probabilities changed. The approach presented here is applicable to similar situations, given the existence of liquid betting markets for the event and of liquid option markets for the asset affected by the event.

References

Black, F. “The pricing of commodity contracts.” *Journal of Financial Economics*, 3 (1976), 169–179.

Black, F., and M. Scholes. “The pricing of options and corporate liabilities.” *Journal of Political Economy*, 81 (1973), 637–659.

Breeden, D. T., and R. Litzenberger. “Prices of state-contingent claims implicit in option prices.” *Journal of Business*, 51 (1978), 621–651.

Castrén, O. “Estimating and analyzing currency options implied risk-neutral density functions for the largest new EU member states.” European Central Bank working paper series 440 (2005).

Christoffersen, P., and K. Jacobs. “The importance of the loss function in option valuation.” *Journal of Financial Economics*, 72 (2004), 291–318.

Clark, I. J. *Foreign Exchange Option Pricing: A Practitioner’s Guide*. West Sussex, UK: Wiley (2010).

Erikson, R., and C. Wlezien. “Markets vs. polls as election predictors: An historical assessment.” *Electoral Studies*, 31 (2012), 532–539.

Figlewski, S. “Estimating the implied risk neutral density for the U.S. market portfolio.” In *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, T. Bollerslev, J. Russel, and M. Watson, eds. Oxford, UK: Oxford University Press (2010).

Garman, M., and S. Kohlhagen. “Foreign currency option values.” *Journal of International Money and Finance*, 2 (1983), 231–237.

Hanke, M.; R. Poulsen; and A. Weissensteiner. “Where would the EUR/CHF exchange rate be without the SNB’s minimum exchange rate policy?” *Journal of Futures Markets*, 35 (2015), 1103–1116.

Hanke, M.; R. Poulsen; and A. Weissensteiner. “Analyzing the Swiss National Bank’s euro exchange rate policy: A latent likelihood approach.” SSRN working paper #2515210, 2016.

Hertrich, M., and H. Zimmermann. “On the credibility of the euro/Swiss franc floor: A financial market perspective.” *Journal of Money, Credit and Banking*, 49 (2017), 567–578.

James, J.; J. Fullwood; and P. Billington. *FX Option Performance*. West Sussex, UK: Wiley (2015).

Jermann, U. “Financial markets’ views about the euro/Swiss franc floor.” *Journal of Money, Credit and Banking*, 49 (2017), 553–565.

Jurek, J., and Z. Xu. “Option-implied currency risk premia.” Working paper, University of Pennsylvania and NBER (2010).

Leahy, M., and C. Thomas. “The sovereignty option: The Quebec referendum and market views on the Canadian dollar.” Board of Governors of the Federal Reserve System, International Finance Discussion Papers No. 555 (1996).

Malz, A. “Using option prices to estimate realignment probabilities in the European Monetary System: The case of sterling-mark.” *Journal of International Money and Finance*, 15 (1996), 717–748.

Mirkov, N.; I. Pozdeev; and P. Söderlind. “Toward removal of the Swiss franc cap: Market expectations and verbal interventions.” Swiss National Bank Working Papers 2016–10 (2016).

Reiswich, D., and U. Wystup. “A guide to FX options quoting conventions.” *The Journal of Derivatives*, 18 (2010), 58–68.

Ritchey, R. J. “Call option valuation for discrete normal mixtures.” *The Journal of Financial Research*, 13 (1990), 285–196.

Table 1: Relative Differences between Market and Model Prices of Options

Table 1 shows the relative price differences (in percent) between market prices and prices calculated using the mixture of normal return densities model. The results shown are for the optimal choice of parameters in equation (8). Densities are fitted for the first day when one-month (three-months) options were traded that expired after the respective events. Columns 2–3 (6–7) show results when using equal weights, while columns 4–5 (8–9) provide a robustness check when using implied vegas as weights.

	GBPUSD				USDMXN			
	Equal-weighted		Vega-weighted		Equal-weighted		Vega-weighted	
Delta	1M	3M	1M	3M	1M	3M	1M	3M
10D	-0.31	0.09	-0.32	0.08	0.01	0.03	-0.02	0.05
25D	1.04	-0.41	1.00	-0.32	0.02	-0.15	0.03	-0.20
ATM	-1.42	0.81	-1.45	0.71	-0.09	0.34	-0.09	0.38
75D	1.09	-0.39	1.08	-0.51	0.30	0.20	0.30	0.24
90D	-0.56	-1.42	-0.56	-1.47	-0.42	-1.39	-0.42	-1.36

Table 2: Parameters of Conditional Densities

Table 2 shows the parameters (in percent) of the conditional densities in equation (5), estimated from equation (8). L, R, T, and N denote the possible outcomes Leave, Remain, Trump, and Not Trump. For ease of comparison, we provide the time-scaled values here. Densities are fitted for the first day when one-month/three-months options were traded that expired after the respective events. Columns 2–3 (7–8) show results when using equal weights, while columns 4–5 (9–10) provide a robustness check when using implied vegas as weights.

Parameters	GBPUSD				Parameters	USDMXN			
	Equal-weighted		Vega-weighted			Equal-weighted		Vega-weighted	
	1M	3M	1M	3M		1M	3M	1M	3M
$\mu_L(\tau - t)$	-8.66	-6.31	-8.65	-6.29	$\mu_T(\tau - t)$	7.49	6.28	7.49	6.17
$\mu_R(\tau - t)$	1.87	2.94	1.87	2.93	$\mu_N(\tau - t)$	-1.38	-0.87	-1.39	-0.84
$\sigma_L\sqrt{\tau - t}$	4.49	10.98	4.50	10.93	$\sigma_T\sqrt{\tau - t}$	6.35	10.95	6.34	11.03
$\sigma_R\sqrt{\tau - t}$	2.85	4.34	2.85	4.37	$\sigma_N\sqrt{\tau - t}$	5.35	6.31	5.35	6.32

Table 3: Moments of the Mixture of Normal Risk-neutral Return Densities

Table 3 provides the moments for the mixed risk-neutral return densities extracted from option prices. Densities are fitted for the first day when 1M/3M-options were traded that expired after the respective events. Left: Moments of the RNDs extracted from GBPUSD options, right: from USDMXN options. Moments correspond to the respective time to maturity of one month/three months (not annualized). “kurt” denotes kurtosis (3 for a normal distribution).

$(\tau - t)$	GBPUSD					USDMXN				
	date	mean	std	skew	kurt	date	mean	std	skew	kurt
1M	05/26	-0.0014	0.0525	-1.18	4.28	10/10	0.0016	0.0648	0.40	3.40
3M	03/23	-0.0034	0.0860	-1.06	4.64	08/09	0.0067	0.0809	0.59	4.18

Figure 1: Exchange Rate Development around Major Political Events

Figure 1 shows the development of the GBPUSD and USDMXN exchange rates around the respective events. Left: GBPUSD rate from Jan. 1, 2016, to Aug. 17, 2016; daily data from Bloomberg. The vertical solid line indicates the date of the U.K. Brexit referendum (June 23, 2016). Right: USDMXN rate from Jan. 1 to Nov. 23, 2016; daily data from Datastream. The vertical solid line indicates the date of the U.S. presidential elections (Nov. 8, 2016).

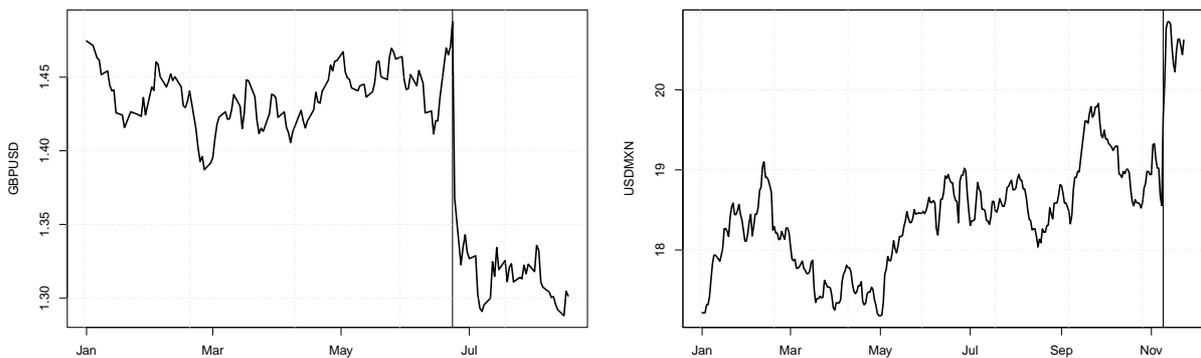


Figure 2: Quotes from Betting Markets

Figure 2 presents quotes from betting markets, expressed as state prices. Left: State prices for Leave, ψ_L ; right: state prices for Trump, ψ_T . The vertical dashed lines indicate the first day when 3-months (1-month) options traded that expired after the respective events.

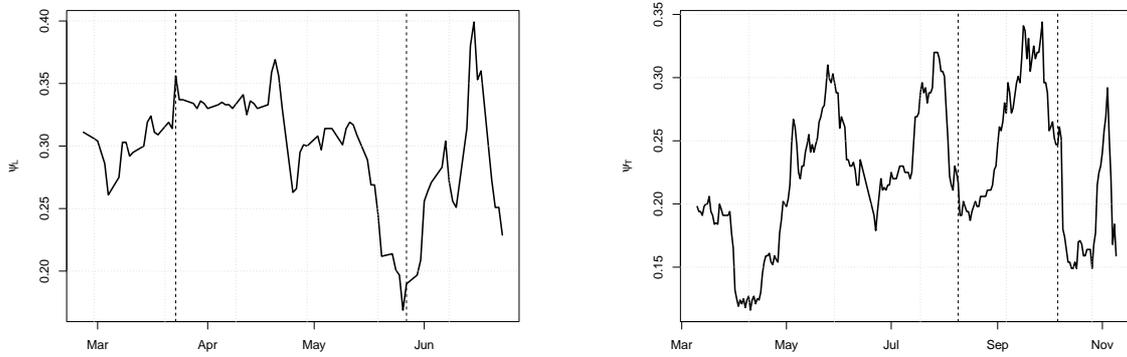


Figure 3: Implied Options Volatilities

Figure 3 presents implied volatilities (in percent) of 3-months options (top) and 1-month options (bottom) on the GBPUSD rate (left) and the USDMXN rate (right). For all panels, the solid line indicates ATM volatilities, while dashed (dotted) lines give the 10D (90D) implied volatilities. Vertical lines: The solid line indicates the event date (left: Brexit referendum, right: U.S. presidential elections), and the dashed line indicates the first day when 3-months (1-month) options traded that expired after the respective events.

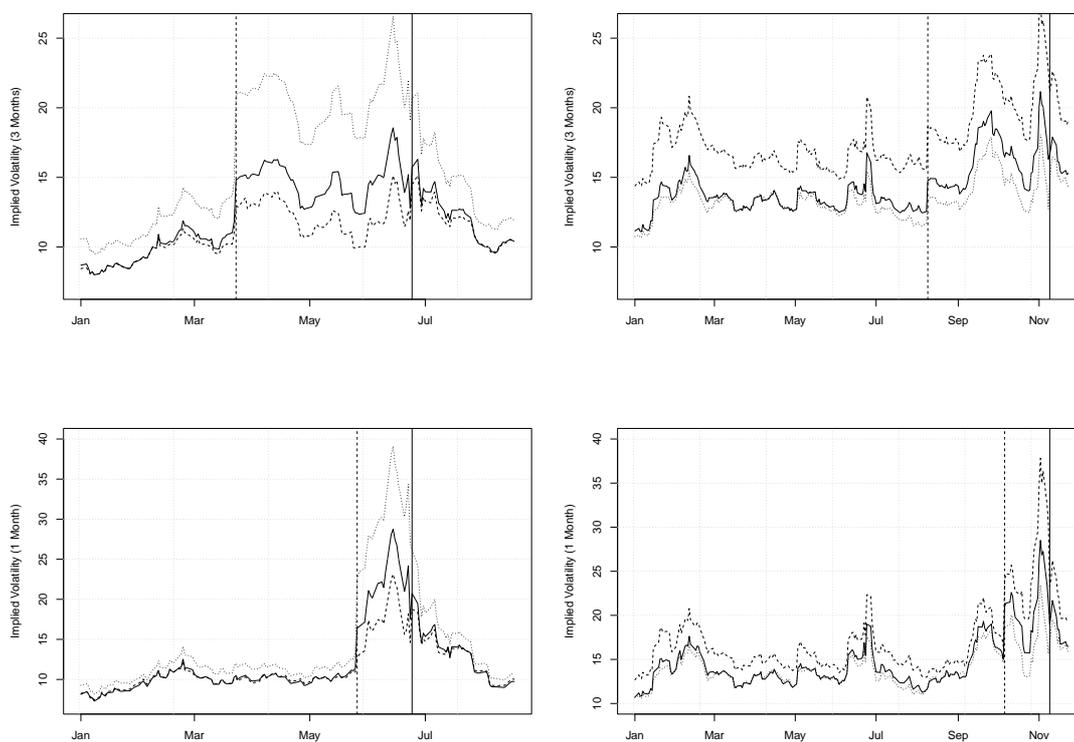


Figure 4: Put–Call Ratio for GBPUSD Options

Figure 4 shows the put–call ratio calculated from the open interest of all GBPUSD options traded on the Chicago Mercantile Exchange (CME), Jan.–Sept. 2016. Data source: CME.

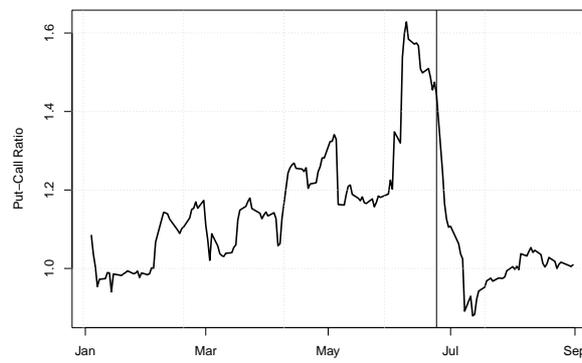


Figure 5: Risk-neutral Densities Estimated from Options Data

Figure 5 presents risk-neutral densities (RNDs) estimated from our options data. Top: RNDs from 3-months options, bottom: RNDs from 1-month options. Left: GBPUSD options, right: USDMXN options. The solid line shows the implied return density estimated using the mixture of normal return densities model. For comparison, the dashed line shows the normal density with the best fit. All curves show the RNDs estimated on the day when 3-months (1-month) options first traded that expired after the event.

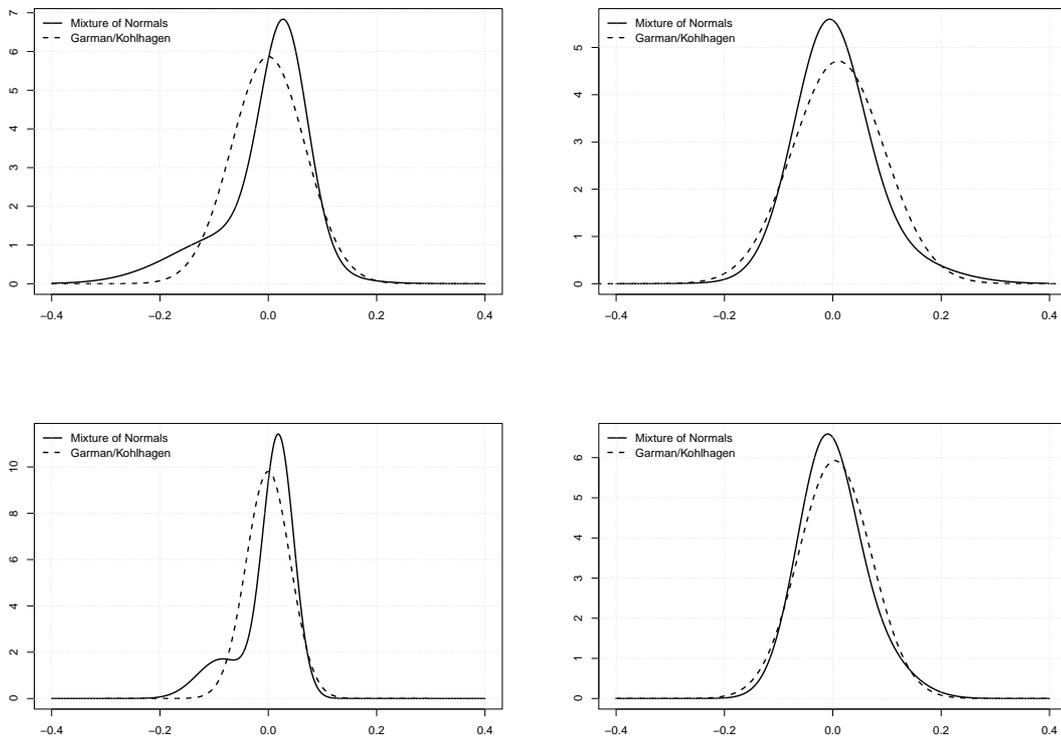


Figure 6: Rolling Forecasts of Expected Exchange Rates

Figure 6 presents rolling forecasts of expected exchange rates, calculated from RNDs of 3-months options, for the time after the respective events. Left: GBPUSD forecasts, right: USDMXN forecasts. The figures show the risk-neutral Leave (Trump) probability (dashed curve, y-axis on the right) and the three-months-ahead conditionally expected exchange rate in case of Leave (Trump), $E_t^Q[S_{t+0.25}|\cdot]$ (solid curve, y-axis on the left). Forecasts are shown for time t , i.e., when the forecast is made, and relate to time $t + 0.25$.

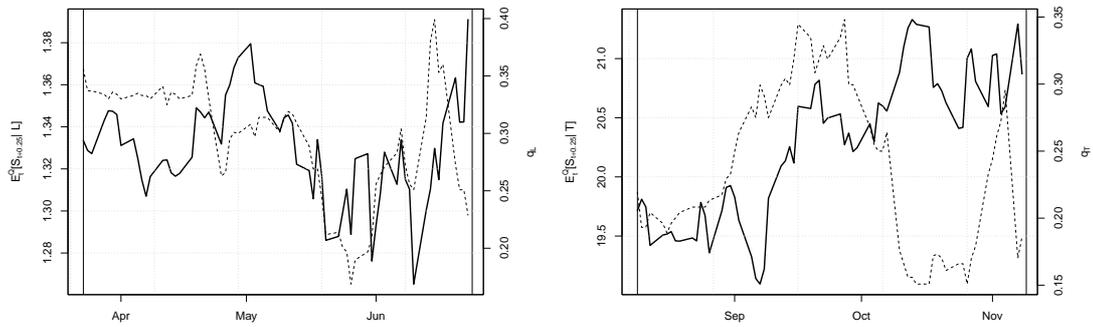


Figure 7: Expected Exchange Rates for the Time after the Brexit Referendum / the U.S. Presidential Elections

Figure 7 presents expected exchange rates from RNDs of three-months options for the time after the respective events. Left: The solid curve is the GBPUSD spot exchange rate. The dashed (dotted) curve shows the three-months-ahead conditionally expected exchange rate $E_t^Q[S_{t+0.25}|\cdot]$ in case of Leave (Remain). Right: The solid curve is the USDMXN spot exchange rate. The dashed (dotted) curve shows the three-months-ahead conditionally expected exchange rates $E_t^Q[S_{t+0.25}|\cdot]$ in case of Trump (Not Trump). All forecasts are shown for time t , i.e., when the forecast is made, and relate to time $t + 0.25$.

