Dynamics of SOFR vs Fed Funds

Karol Gellert* and Erik Schlögl**

*University of Technology Sydney
**University of Technology Sydney, Quantitative Finance Research Centre.
Erik.Schlogl@uts.edu.au

August 30, 2019— Incomplete and preliminary draft, comments welcome

Abstract

In this paper, an interest rate term structure model reflecting the key empirical features of the dynamics of the Secured Overnight Funding Rate (SOFR) is constructed and fitted to market data, both in the sense of cross–sectional calibration and time series estimation. As an overnight rate for (effectively risk–free) lending in US dollars, the dynamics of SOFR are closely linked to the dynamics of the Fed Funds overnight rate, which in turn is the interest rate most directly impacted by US monetary policy decisions. Therefore, these rates feature jumps at known jump times (Federal Open Market Committee meeting dates), and market expectations of these jumps are reflected in prices for futures written on these rates.

1 Introduction

As the Secured Overnight Funding Rate (SOFR) is currently in the process of becoming the key Risk–Free Rate (RFR) benchmark in US dollars, interest rate term structure models need to be updated to reflect this. Historically, interest rate term structure modelling has been based on rates of substantially longer time to maturity than overnight, either directly as in the LIBOR Market Model,[1] or indirectly, in the sense that even models based on the continuously compounded short rate (i.e., with instantaneous maturity)[2] are typically calibrated to term rates of longer maturities, with any regard to a market overnight rate at best an afterthought. Of course, all diffusion–based arbitrage–free interest rate term structure models can be nested in the general framework of Heath, Jarrow and Morton (1992), which describes the stochastic dynamics of the entire term structure of instantaneous forward rates. As such, they *should* be calibrated to interest rates of all liquidly traded maturities.

However, with SOFR this situation is reversed: The overnight rate now is the primary market observable, and term rates (i.e., interest rates for longer maturities) are less readily available and therefore must be inferred (for example from derivatives prices). Furthermore, even when such term rates can be inferred, they will not be equivalent to the LIBOR benchmark which SOFR is mooted to be replacing. Rather, the relationship between SOFR and derivatives–implied term rates is more akin to the relationship between the overnight Fed Funds rate and the associated term rates implied by overnight index swaps (OIS). Thus one would expect a spread between LIBOR and SOFR–based term rates, in the same manner that we observe a LIBOR/OIS spread.[3]

Thus the empirical idiosyncracies of the overnight rate cannot be ignored when constructing interest rate term structure models in a SOFR–based world, and more than longer term rates, these idiosyncracies are driven by monetary and regulatory policy. In this paper, we identify the key effects which a SOFR–based interest rate term structure model must reflect and construct a model which satisfies these requirements. The model is both calibrated to and estimated on market data — in our context, we will speak of "calibration" to mean a cross–sectional fit of the model to available market data (including futures prices) at a given point in time, while "estimation" refers to econometric estimation of the model parameters on time series data.

The rest of the paper is organised as follows. Features of SOFR which will inform our modelling are discussed in Section 2 — in particular, aspects of SOFR dynamics are related to the Fed Funds rate, which is most directly impacted by monetary policy decisions. A model reflecting these features is introduced in Section 3, and calibrated to and estimated on market data in Section 4.

[1]See Miltersen, Sandmann and Sondermann (1997), Brace, Gatarek and Musiela (1997) and Musiela and Rutkowski (1997).

[2]Of these, Hull and White (1990) is the most prominent example.

[3]The LIBOR/OIS spread, and more generally basis spreads between interest rates of different payment frequencies (the "multicurve" phenomenon) can be explained by the presence of "roll–over risk," i.e. risk that an entity will face a higher spread to the market benchmark when they attempt to refinance (roll over) borrowing in the future. This "roll–over risk" has a credit and a funding liquidity component, see e.g. Alfeus, Grasselli and Schlögl (2018).

2 Short rates

2.1 Effective Fed Funds rate (EFFR)

Federal Reserve accounts exist to facilitate the regulatory requirement of depository institutions[4] to hold a proportion of certain deposits in the form of liquid reserves[5]. Transactions between these accounts take place as institutions borrow and lend overnight funds to each other as part of their daily management of reserve balances and operational cash flows. The rate at which this occurs, referred to as the Fed Funds rate, can vary throughout the day and may differ between counterparties. A volume weighted median is calculated from the day's fed fund rates and published the next morning as the effective Fed Funds rate (EFFR)[6]. Because of its role in determining the short term funding costs, the EFFR has an impact on very short term interest rates. Combined with its role as a fixing rate for certain swaps and futures, this transfers to the interest rate term structure particularly at the short end.

The Fed Funds market is also an important lever for the Federal Reserve in the implementation of monetary policy[7], which is formally determined by the Federal Open Market Committee (FOMC) at eight scheduled meetings per year. The policy includes the setting of a target range for the Fed Funds rate. The target range was first introduced in 2008, as part of the Federal Reserve's post financial crisis strategy[8], prior to which the target was expressed as a single rate. The target range is achieved through various forms of interventions by the Federal Reserve, including direct participation in the Fed Funds market, controlling the rate of interest on the reserve accounts, as well as providing repo facilities at a rate designed to support the target range. As a result of the Federal Reserve's intervention, the EFFR remains effectively within the target range and therefore changes to the Federal Reserve's monetary policy as expressed through the range are a key driver of EFFR dynamics.

Prior to the introduction of the target range, Federal Reserve activities in the Fed Funds market were aimed at maintaining the EFFR near a target rate. During this period the EFFR tended to gravitate around the target rate with varying degrees of volatility, see Figure 1. The fact that the EFFR closely tracked the target rate can be attributed to the Federal Reserve's trading desk's activities in the Fed Funds market, formally called open market operations. The factors driving the level of volatility of the EFFR in relation to the target rate are examined by Hilton (2005) and include the daily volumes of reserve account transactions, particularly on high payment days. The ability of the Federal Reserve's trading desk to maintain the EFFR near the target rate significantly deteriorated during the financial crisis. This was acknowledged by the Federal Reserve[9] as one of the factors

[4]financial institutions licensed to hold deposits

[5]Liquid reserves are either vault cash or Federal Reserve account balances.

[6]See Federal Reserve Bank of New York (2015) for the detailed methodology for calculating EFFR.

[7]For a detailed description of monetary policy see Federal Reserve Board (2011).

[8]See Federal Open Market Committee, December (2008).

[9]See Federal Open Market Committee, December (2008) page 9.

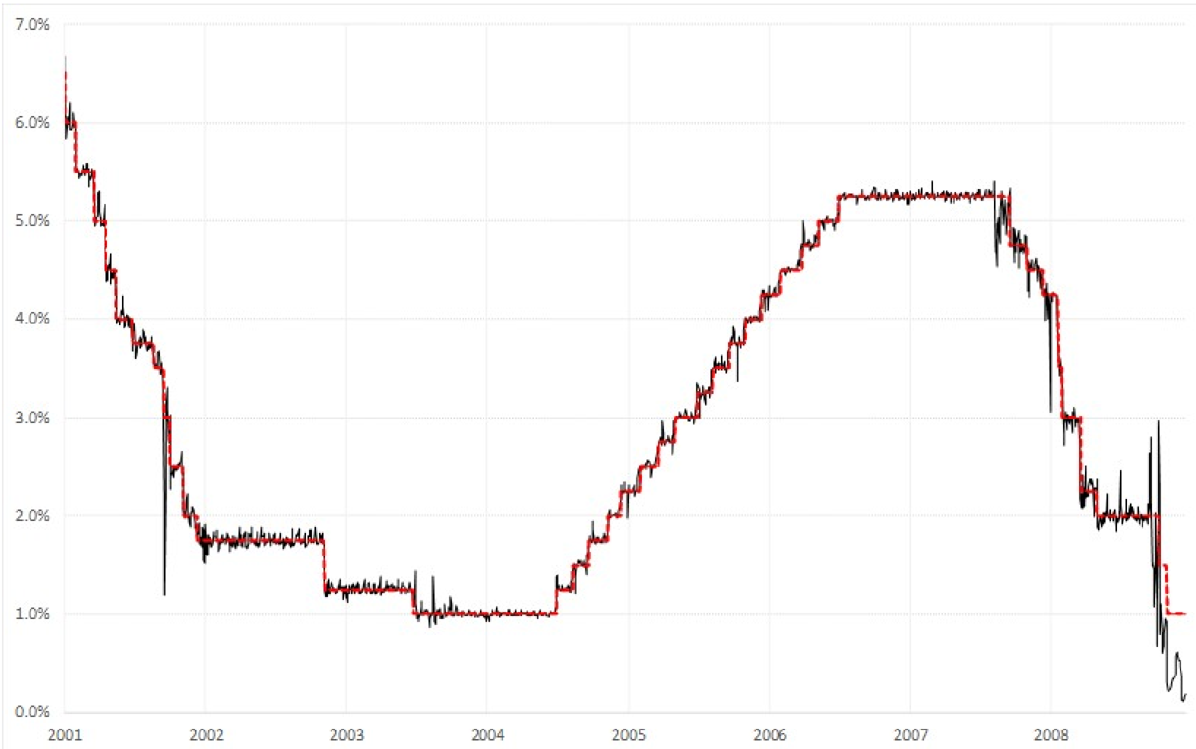

Figure 1: EFFR (black) vs target range (red)

considered when switching to a target range, initially set between 0 and 25 basis points. As can be seen in Figure 2, since the introduction of the range the behaviour of EFFR can be divided into three periods. In the period 2008-2014 EFFR appears to follow a dynamic within the bounds of the target range, which remained at 0 to 25 basis points during this period. Occasional end of month downward spikes[10] do exist in this period but begin to appear regularly in 2014 and continue until 2017, coinciding with the beginning of the normalisation period[11] marked by regularly timed increases in the target range levels. The year 2017 marks the beginning of the third phase where the EFFR is increasingly bound to the interest on excess reserves rate (IOER).

In October 2008, the Federal Reserve began paying IOER[12] to help control the EFFR in an environment of balance sheet expansion resulting in increasing balances in excess reserves. Under normal circumstances, the IOER should act as a lower bound for the EFFR, since no institutions should want to lend below this rate. As such, effective from October 9 the IOER was set 75 basis points below the lowest EFFR over each reserve

[10] For a detailed explanation of the spikes see Hartely (2017).

[11] Normalisation refers to the Federal Reserve's strategy to revert to monetary policy prior to the extreme measures taken during the financial crisis. Discussions of normalisation plans appear regularly in FOMC minutes starting from Federal Open Market Committee, April (2014) page 2.

[12] See Federal Open Market Committee, October (2008) page 7.

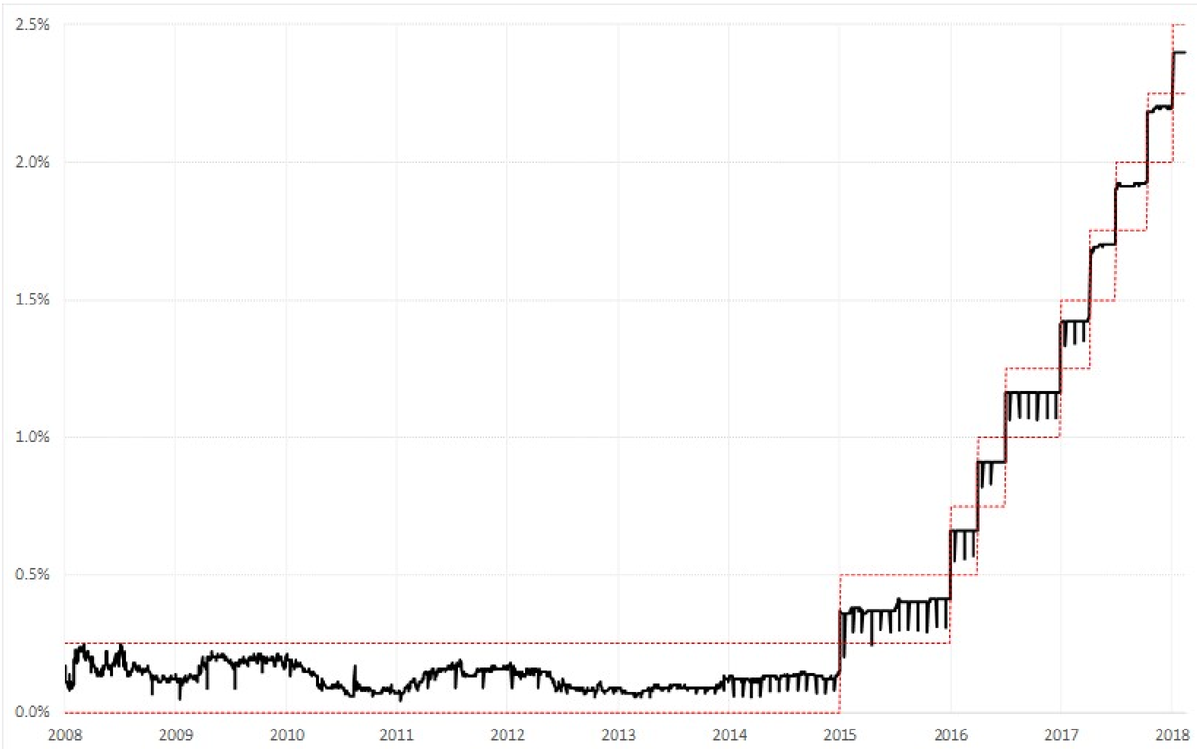

Figure 2: EFFR (black) vs target range (red)

maintenance period,[13] but this was quickly changed on October 22 to 35 basis points below the lowest EFFR over each reserve maintenance period. In the FOMC immediately following the introduction of the IOER, it was noted that institutions not eligible to receive it were willing to sell (lend) funds at rates below the IOER.[14] In December 2008, together with the introduction of the target range, the IOER was set at the target range upper limit of 25 basis points in recognition that due to unique circumstances the IOER was acting as an upper bound for the EFFR.

Normally, one would expect that the IOER would act as a lower bound to the EFFR, with institutions not willing to lend excess funds at a rate below what they could earn as IOER. The initial setting of IOER as outlined above suggests this was also the initial assumption of the Federal Reserve. One of the factors explaining how the IOER has become an effective upper bound for EFFR is that government sponsored institutions (GSIs), which are eligible to hold reserve accounts, are not eligible to earn the IOER[15], and therefore willing to lend excess funds at a lower rate. The second factor is the elevated excess reserves environment resulting from the Federal Reserve's injection of liquidity in response to the financial crisis.

[13]The reserve maintenance period is a two week time frame in which banks and other depository institutions must maintain a specified level of funds.

[14]See Federal Open Market Committee, October (2008) page 2.

[15]See Federal Open Market Committee, October (2008) page 2.

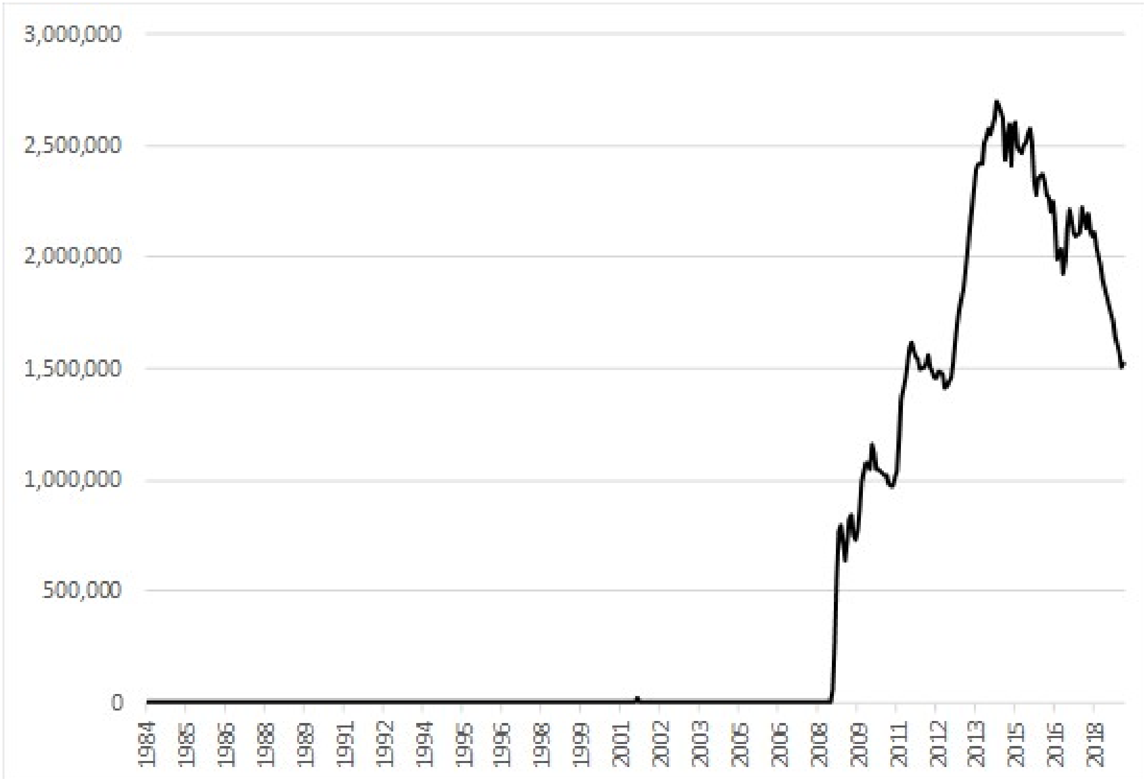

Figure 3: Excess reserves (millions), source: Board of Governors, St. Louis Fed

By historical standards, the rise in excess reserves was extreme and without precedent. As can seen in Figure 3, it increased from under $2 billion in September 2008 to $1 trillion by November 2009, before reaching a high of over $2.5 trillion in October 2015. The large surpluses in excess reserves have eliminated demand for reserve loans. Instead the Fed Funds rate is driven by GSIs lending their excess reserves at below the IOER to institutions who then earn the difference between the Fed Funds rate and the IOER. In effect, by paying the IOER in a market flush with liquidity, the Federal Reserve acts as the borrower, rather than the lender, of last resort.

The Federal Reserve's strategy in response to the financial crisis centred around two key policies: near zero interest rates and quantitative easing. The phases of quantitative easing became known as QE 1/2/3 and involved selling Treasury bonds and purchases of various other credit risky assets[16] in a bid to boost liquidity and credit conditions. Plans for reversal of the post financial crisis expansionary policy were formally laid out at the FOMC September 2014 meeting as the Policy Normalization Principles and Plans.[17] The aim of the reversal strategy is to bring the effective Fed Funds rate back to normal levels and

[16]Such as Agency Debt, Mortgage Backed Securities and Term Auction Facilities, see Binder (2010).
[17]See Federal Open Market Committee, September (2014) page 3.

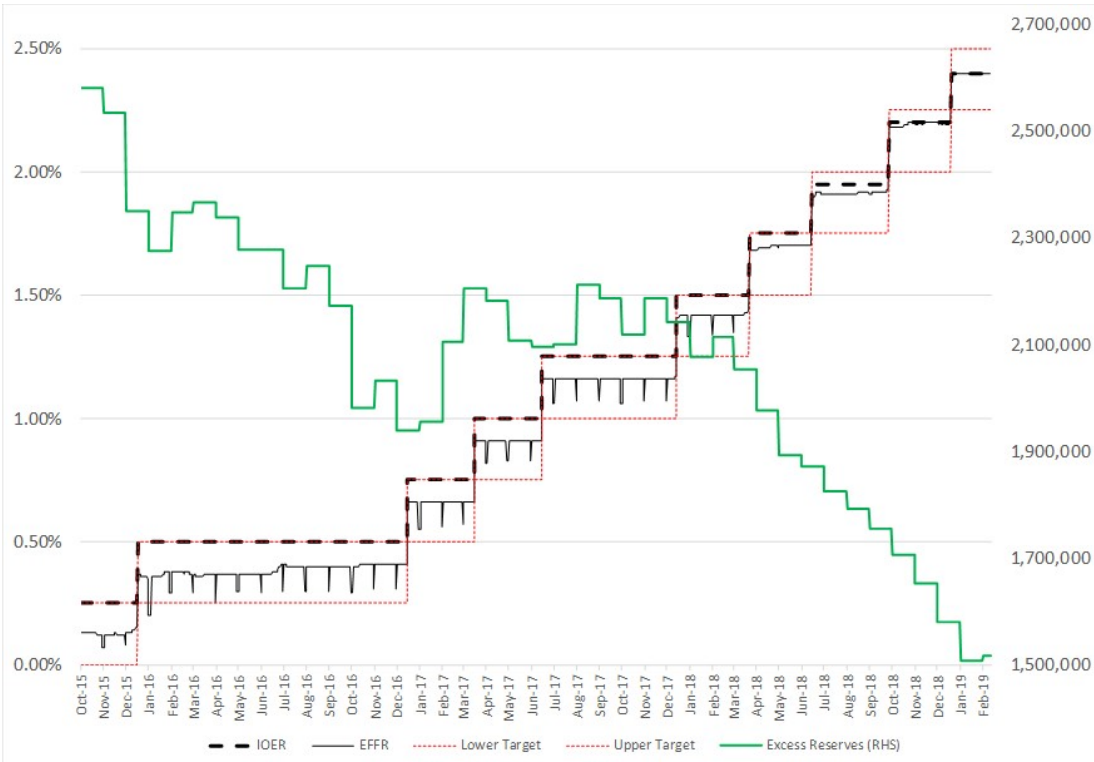

Figure 4: EFFR, IOER, target range and excess reserves during the normalisation phase

reduce the securities held by the Federal Reserve, thereby unwinding the excess reserves held by banks. Prior to the financial crisis, controlling the supply of reserves via open market operations was a key tool in controlling the Fed Funds rate. However the Federal Reserve has adopted the view that with banks using reserves for liquidity more so than prior to the crisis it might be hard to predict demand for reserves and therefore open market operations would not be effective at precisely controlling the EFFR.[18] Instead, the new normal will constitute the Federal Reserve keeping excess reserves just large enough to remain on the flat part of the demand curve, a prerequisite condition for the use of the IOER to control the EFFR.

In terms of outcomes, as can be seen in Figure 4, the normalisation phase has been characterised by regularly increasing EFFR as a result of increases in the target range, increasing IOER and falling reserve balances. Additionally the IOER has been moved below the upper target range level.[19] Another outcome of the normalisation phase is the convergence of EFFR to the IOER, resulting in the IOER acting as an effective target rate for the EFFR. This is consistent with remarks in FOMC minutes, suggesting the target

[18]See Federal Open Market Committee, November (2018) page 3.

[19]Based on FOMC discussions regarding the possibility of EFFR moving above IOER as excess reserves drop, see for example Federal Open Market Committee, December (2017) page 3.

range is used as a communication tool rather than a hard target and the IOER used to control the EFFR.[20] As we will show, the IOER acting as the target for EFFR is also consistent with expectations in the futures market.[21]

From an empirical perspective, the dominating feature of the EFFR dynamics consists of jumps with known jump times, regardless of whether the Federal Reserve is targeting a specific rate or a range. In the period prior to the financial crisis shown in Figure 1, the remaining dynamic could be explained as noise with time dependent variance. Therefore a model consisting of jumps with known jump times and noise captures the most prominent stylized facts contained in the data. In the period after the financial crisis, prior to normalisation, it could be argued that the EFFR follows some other dynamic contained within the target range. However, during normalisation the EFFR has converged to the IOER rate and therefore the IOER could now be considered as the new target rate. The Federal Reserve has adopted a strategy of keeping excess reserves in a state where the IOER could be used to control the EFFR, in place of open market operations. This has resulted in an elimination of noise from the EFFR which is currently effectively equivalent to the IOER. Therefore the new stylized facts are very similar to when the target rate was in place prior to the financial crisis, with the exception that the variance of the noise is close to zero. Even if a specific target rate is unknown, for example in the period prior to normalisation, Fed Funds futures provide enough information to calculate the market expectation of the target rate and expected jumps. In order to capture the outlined stylized facts we propose to model the Fed Fund target rate as a pure jump process with known jump times and the EFFR as the target rate plus noise. This model can be calibrated to Fed Funds futures and is robust to changes in Federal Reserve strategy of controlling the EFFR.

2.2 Secured overnight financing rate (SOFR)

Shortly following the well publicised LIBOR manipulation scandals, the Financial Stability Board and Financial Stability Oversight Council highlighted one of the key problems to be the decline in transactions underpinning LIBOR and the associated structural risks to the financial system.[22] As argued in Schrimpf and Sushko (2019), partly to blame for the decline in interbank term lending are the inflated excess reserves discussed in the previous section.[23] In response, the Federal Reserve convened the Alternative Reference Rates Committee (ARRC)[24] to explore alternative reference rates. In June 2017 the ARRC formally announced the Secured Overnight Financing Rate (SOFR) as the replacement for LIBOR. A key criterion for the choice was the large volume of transactions behind SOFR, translating to it being more representative of bank's funding costs and less susceptible to manipulation. The calculation of SOFR is based on a broad base of overnight repo

[20]See Federal Open Market Committee, July (2014) page 2.

[21]This is also acknowledged in Federal Open Market Committee, June (2018) page 2.

[22]See The Alternative Reference Rates Committee (2018) page 1.

[23]This suggests an interesting causal link between the financial crisis, the Federal Reserve response and the emergence of SOFR by linking the decline in LIBOR transactions to excess reserves.

[24]See https://www.newyorkfed.org/arrc

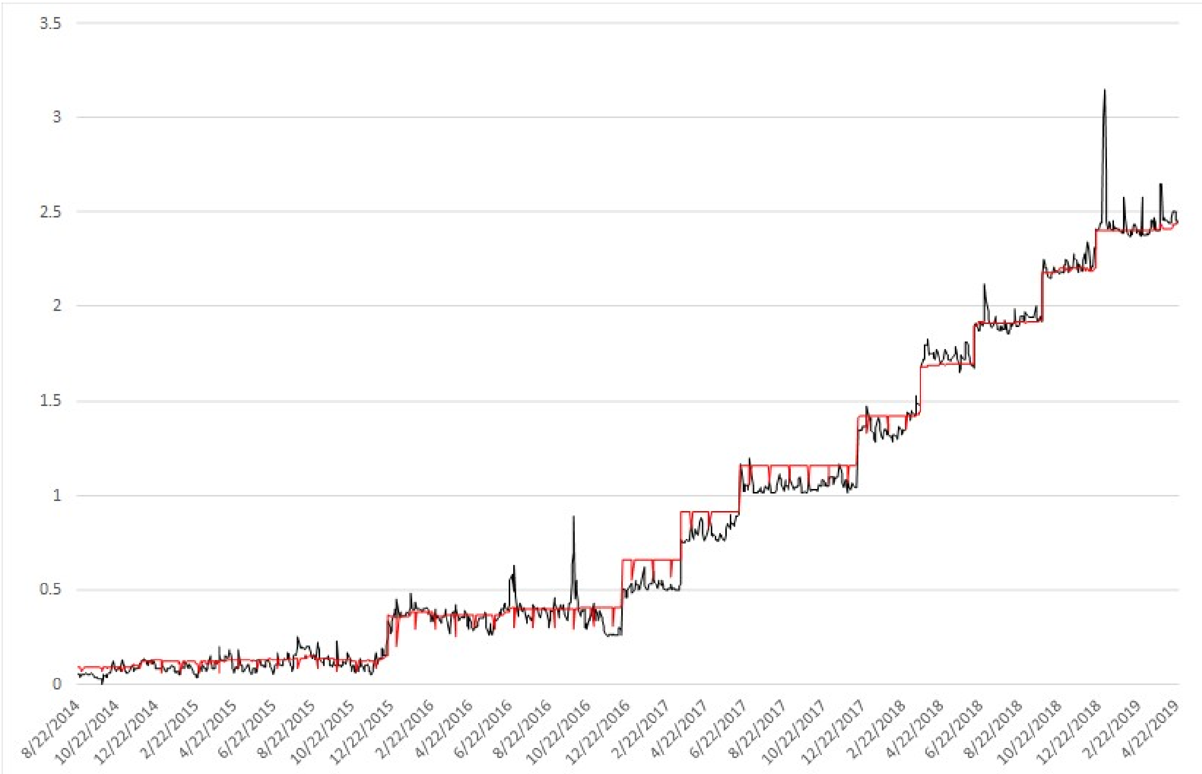

Figure 5: EFFR(red) vs SOFR (black)

transactions, which in 2017 averaged around \$700b daily[25] (compared to less than \$1b for LIBOR transactions).

Official SOFR fixings have been calculated as far back 2014 and can be seen in comparison to EFFR in Figure 5. Three features stand out, firstly EFFR and the SOFR rate both appear to follow the same stepwise function, suggesting that similarly to EFFR, the Fed Funds target rate plays an important role in the SOFR dynamic. Another aspect is that SOFR is significantly more volatile than EFFR. A third feature is the prominence of end of month spikes, which are especially pronounced at end of quarter and end of year.

The end of month spikes are related to the measurement of dealers' balance sheet exposures at month end for regulatory purposes. This single snapshot approach incentivises the management of exposures around reporting dates, which, as explained in Schrimpf and Sushko (2019), has been resulting in increases in the SOFR rate on end of month dates. This is demonstrated in Figure 6, which displays the difference between the end of month SOFR fixing and the average of the SOFR fixings for the corresponding month. A closer look at the behaviour of the SOFR fixing around the end of month spikes reveals that fol-

[25] For details see The Alternative Reference Rates Committee (2018) page 7.

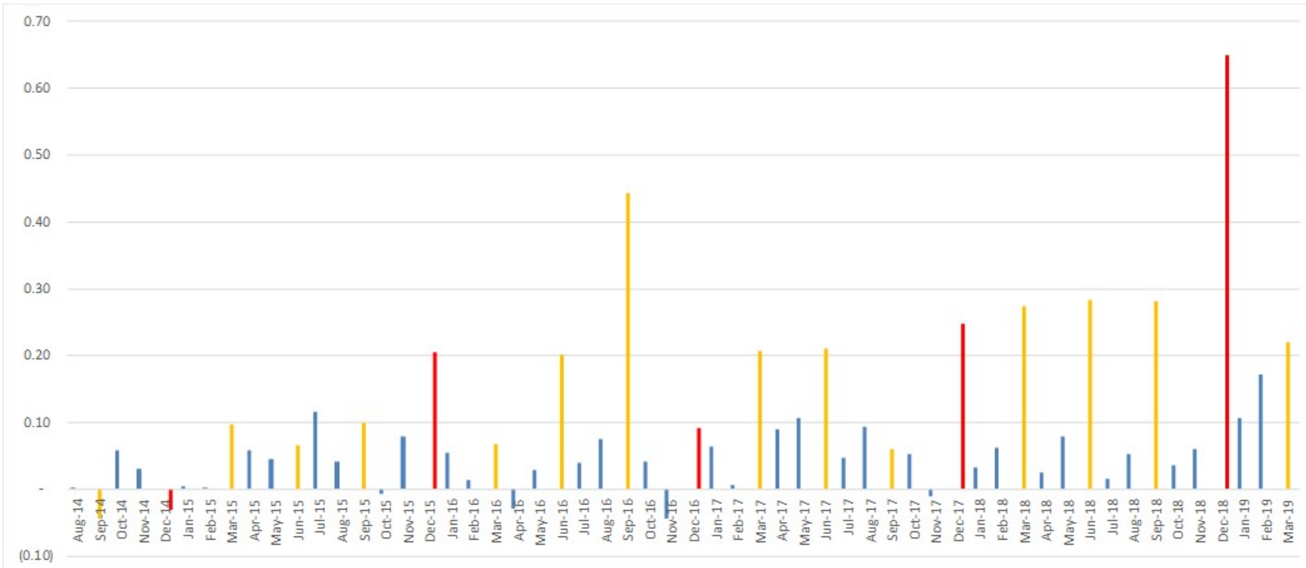

Figure 6: EOM SOFR fixing minus month's average, EOM marked as blue, EOQ marked as yellow, EOY marked as red

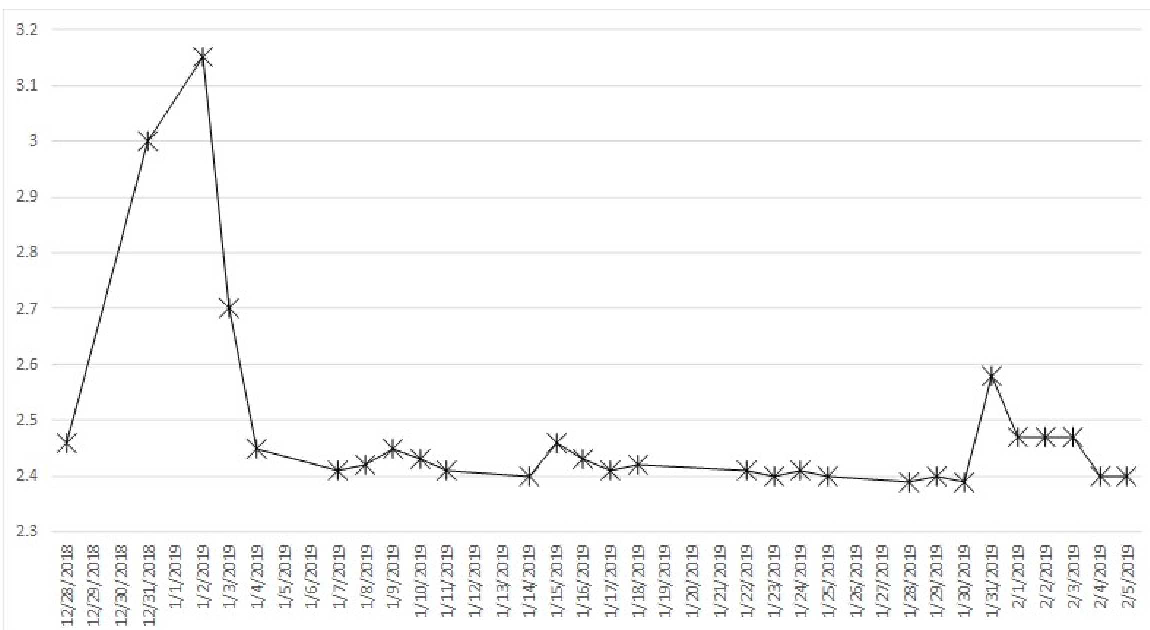

Figure 7: SOFR fixings over an approximately one month period

lowing the spike the SOFR rate tends to return to the level prior to the spike over several days. To reflect the stylized facts of the SOFR discussed so far, we model the SOFR spread to the Fed Funds target rate as an Ornstein-Uhlenbeck process extended with jumps with known jump times. The jumps model the end of month spikes while the mean reversion component of an Ornstein-Uhlenbeck process captures the reversion of SOFR following the spike.

3 Model

3.1 Fed target rate model

As discussed in the previous section, we model the target rate as a pure jump process with known jump times. Denoting r_p as the target rate, T_P as the set of FOMC meeting dates, the expected value of the target rate at time $u > t$ is the current target rate plus the sum of expected target rate jumps:

$$E_t\Big[r_p(u)\Big] = r_p(t) + \sum_{v \in [t,u]} \mathbf{1}_{T_P}(v) E_t[J(v)] \tag{1}$$

where $J(t)$ is the distribution of a jump size at t. The definition of the distribution of the jump size allows some freedom to match option prices, which do exist for Fed Funds futures. However, at the time of writing, potentially due to lack of realised daily volatility in the EFFR, the volume of option contracts traded is extremely low — it appears most option contracts are not at all traded. Therefore throughout the paper we do not define the jump size distribution, it is sufficient that we are able to calculate the expected value of jumps from Fed Funds futures, which is then used for the calibration of the SOFR futures.

3.2 EFFR model

The EFFR evolves in a market in which the Federal Reserve participates for the specific purpose of keeping the rate within the target range. The Fed is generally successful in keeping the rate within the range. However, since we are modelling a target rate rather than a range, we model the EFFR as the target rate plus noise. The noise reflects the variations of the EFFR within the range. Denoting r_f as the EFFR:

$$r_f(t) = r_p(t) + \epsilon_f \ where \ \epsilon_f \sim N(0, \alpha) \tag{2}$$

The expected value at t of an EFFR fixing at u for $u > t$ can be written in terms of the target rate:

$$E_t\Big[r_f(u)\Big] = E_t\Big[r_p(u) + \epsilon_f\Big] = E_t\Big[r_p(u)\Big] \tag{3}$$

3.3 SOFR model

We model the SOFR rate r_s as the Fed Funds target rate plus a spread q. The dynamics of q follow an Ornstein-Uhlenbeck process with the addition of an end of month jump. Define $Z = (z_1, ..., z_n)$ as the set of sequential end of month dates, $J^u(t)$ the spike distribution and

$\Delta H(t) = \mathbf{1}_Z(t)$. The dynamics of the spread q are defined as

$$dq(t) = \theta(\mu - q(t))dt + \sigma(t)dW(t) + J^u(t)\Delta H(t) \tag{4}$$

with

$$r_s(t) = r_p(t) + q(t) \tag{5}$$

The expected value at t of a SOFR fixing at u for $u > t$ can be written as:

$$E_t\Big[r_s(u)\Big] = E_t\Big[r_p(u)\Big] + q(t)e^{-\theta(u-t)} + \mu(1 - e^{-\theta(u-t)}) + \sum_{i=1}^{k} e^{-\theta(u-z_i)} E\Big[J^u(z_i)\Big] \tag{6}$$

where $z_1 > t$ and $z_k < u$

Proof:
Let $f(q(t), t) = q(t)e^{\theta t}$, then:

$$\begin{aligned} df(q(t), t) &= \theta q(t)e^{\theta t}dt + e^{\theta t}dq(t) \\ &= \theta\mu e^{\theta t}dt + \sigma(t)e^{\theta t}dW(t) + e^{\theta u}J^u(t)\Delta H(t) \end{aligned} \tag{7}$$

Integrating from t to u:

$$q(u)e^{\theta u} = q(t)e^{\theta t} + \int_t^u \theta\mu e^{\theta u}du + \int_t^u \sigma(u)e^{\theta u}dW(u) + \sum_{i=0}^{k} e^{-\theta(u-z_i)}J^u(z_i) \tag{8}$$

therefore

$$q(u) = q(t)e^{-\theta(u-t)} + \mu(1 - e^{-\theta(u-t)}) + \int_t^u \sigma(u)e^{-\theta(u-t)}dW(u) + \sum_{i=0}^{k} e^{-\theta(u-z_i)}J^u(z_i) \tag{9}$$

where $z_1 > t$ and $z_k < u$. For the Itô integral we have $E[\int_t^u \sigma(u)e^{-\theta(s-u)}dW(u)] = 0$, therefore equation (6) follows.

4 Calibration

4.1 30 day fed fund futures

The 30 day Fed Funds futures contract[26] pays $4167 \times F_m(\tau_{m,n_m})$, where

- $m :=$ months from current trading month, $m = 0$ indicates current month

[26] Source: https://www.cmegroup.com/trading/interest-rates/stir/30-day-federal-fund_contract_specifications.html

- $\tau_{m,i}$:= date corresponding to day i, month m
- n_m := total days in month m

The futures contract pays $F_m(\tau_{m,n_m}) = 100 - R$ where R is the arithmetic average of the daily EFFR fixing during the contract month, settled on the first business day after the final fixing date. Define $R_m := \frac{1}{n_m}\sum_{i=1}^{n_m} r_f(\tau_{m,i})$, the terminal payoff is:

$$F_m(\tau_{m,n_m}) = 100 - R_m = 100\left(1 - \frac{1}{n_m}\sum_{i=1}^{n_m} r_f(\tau_{m,i})\right)$$

Using the generic futures pricing theorem,[27] the value at t of a futures contract F is

$$F_m(t) = E_t[F_m(\tau_{m,n_m})|\mathcal{F}_t] = 100\left(1 - \frac{1}{n_m}\sum_{i=1}^{n_m} E_t[r_f(\tau_{m,i})|\mathcal{F}_t]\right) \tag{10}$$

where the expectation is taken under the spot risk neutral measure (a.k.a. cash account measure). Define the implied average rate $A^f_m(t) = 1 - \frac{F_m(t)}{100}$

$$A^f_m(t) = \frac{1}{n_m}\sum_{i=1}^{n_m} E_t[r_f(\tau_{m,i})|\mathcal{F}_t]$$

The current futures continues to trade during the observation month, therefore the valuation needs to account for already observed values of r_f:

$$A^f_0(t) = \frac{1}{n_0}\sum_{i=1}^{n_0} \mathbf{1}_{(t>\tau_{0,i})} r_f(\tau_{0,i}) + \frac{1}{n_0}\sum_{i=1}^{n_0} \mathbf{1}_{(t\leq\tau_{0,i})} E_t[r_f(\tau_{0,i})|\mathcal{F}_t] \tag{11}$$

Using equation (3):

$$n_0 A^f_0(t) - \sum_{i=1}^{n_0} \mathbf{1}_{(t>\tau_{0,i})} r_f(\tau_{0,i}) = \sum_{i=1}^{n_0} \mathbf{1}_{(t\leq\tau_{0,i})} E_t[r_p(\tau_{0,i})|\mathcal{F}_t]$$

Using equation (1):

$$n_0 A^f_0(t) - \sum_{i=1}^{n_0} \mathbf{1}_{(t>\tau_{0,i})} r_f(\tau_{0,i}) = \sum_{i=1}^{n_0} \mathbf{1}_{(t\leq\tau_{0,i})}\left(r_p(t) + \sum_{u\in[t,\tau_{0,i}]} \mathbf{1}_{T_P}(u) E_t[J(u)]\right)$$

Define n_t such that $\tau_{0,n_t} = t$

$$\sum_{i=n_t}^{n_0} \sum_{u\in[t,\tau_{0,i}]} \mathbf{1}_{T_P}(u) E_t[J(u)] = n_0 A^f_0(t) - \sum_{i=1}^{n_t-1} r_f(\tau_{0,i}) - \sum_{i=n_t}^{n_0} r_p(t)$$

[27] See Hunt and Kennedy (2004) theorem 12.6.

At most one FOMC is scheduled per month, if a meeting is scheduled to occur in the current observation period define $\widehat{n}_0$ such that $\tau_{0,\widehat{n}_0} \in T_P$, assuming $\widehat{n}_0 \geq n_t$

$$\sum_{i=\widehat{n}_0}^{n_0} E_t[J(\tau_{0,\widehat{n}_0})] = n_0 A_0^f(t) - \sum_{i=1}^{n_t-1} r_f(\tau_{0,i}) - (n_0 - n_t + 1) r_p(t)$$

therefore

$$E_t[J(\tau_{0,\widehat{n}_0})] = \frac{1}{(n_0 - \widehat{n}_0 + 1)} \left(n_0 A_0^f(t) - \sum_{i=1}^{n_t-1} r_f(\tau_{0,i}) - (n_0 - n_t + 1) r_p(t) \right) \tag{12}$$

For subsequent contracts proceed in an iterative fashion using:

$$A_m^f(t) = \frac{1}{n_m} \sum_{i=1}^{n_m} E_t[r_f(\tau_{m,i})|\mathcal{F}_t]$$

Using equation (3):

$$n_m A_m^f(t) = \sum_{i=1}^{n_m} E_t[r_p(\tau_{m,i})|\mathcal{F}_t]$$

Using equation (1):

$$n_m A_m^f(t) = \sum_{i=1}^{n_m} \left(r_p(t) + \sum_{u \in [t, \tau_{m,i}]} \mathbf{1}_{T_P}(u) E_t[J(u)] \right)$$

Isolating jumps for which expectations have been derived from preceding futures:

$$n_m(A_m^f(t) - r_p(t)) = \sum_{i=1}^{n_m} \left(\sum_{u \in [t, \tau_{m-1, n_{m-1}}]} \mathbf{1}_{T_P}(u) E_t[J(u)] + \sum_{u \in [\tau_{m,1}, \tau_{m,i}]} \mathbf{1}_{T_P}(u) E_t[J(u)] \right)$$

Define already calculated jump expectations as $\widetilde{J}_m = \sum_{u \in [t, \tau_{m-1, n_{m-1}}]} \mathbf{1}_{T_P}(u) E_t[J(u)]$

$$n_m(A_m^f(t) - r_p(t) - \widetilde{J}_m) = \sum_{i=1}^{n_m} \sum_{u \in [\tau_{m,1}, \tau_{m,i}]} \mathbf{1}_{T_P}(u) E_t[J(u)]$$

Define $\widehat{n}_m$ such that $\tau_{m,\widehat{n}_m} \in T_P$ and $1 \leq \widehat{n}_m \leq n_m$

$$\sum_{i=\widehat{n}_m}^{n_m} E_t[J(\tau_{m,\widehat{n}_m})] = n_m(A_m^f(t) - r_p(t) - \widetilde{J}_m)$$

therefore

$$E_t[J(\tau_{m,\widehat{n}_m})] = \frac{n_m(A_m^f(t) - r_p(t) - \widetilde{J}_m)}{n_m - \widehat{n}_m + 1} \tag{13}$$

4.2 SOFR futures

4.2.1 1m

Apart from the reference rate the 1m SOFR futures are defined identically to the 30 day Fed Funds futures[28] That is, the contract pays $4167 \times F_m^{s1}(\tau_{m,n_m})$, with $F_m^{s1}(\tau_{m,n_m}) = 100 - R$ where R is the arithmetic average of the daily SOFR fixing during the contract month, settled on the first business day after the final fixing date. The payoff can be written as:

$$R_m^{s1} := \frac{1}{n_m} \sum_{i=1}^{n_m} r_s(\tau_{m,i})$$

The terminal payoff is

$$F_m^{s1}(\tau_{m,n_m}) = 100(1 - R_m^{s1}) = 100\left(1 - \frac{1}{n_m} \sum_{i=1}^{n_m} r_s(\tau_{m,i})\right)$$

Again using the generic futures pricing theorem, the value at t of a futures contract F is

$$F_m^{s1}(t) = E_t[F_m^{s1}(\tau_{m,n_m})|\mathcal{F}_t] = 100\left(1 - \frac{1}{n_m} \sum_{i=1}^{n_m} E_t[r_s(\tau_{m,i})|\mathcal{F}_t]\right) \tag{14}$$

Define the implied average rate $A_m^{s1}(t) = 1 - \frac{F_m^{s1}(t)}{100}$

$$A_m^{s1}(t) = \frac{1}{n_m} \sum_{i=1}^{n_m} E_t[r_s(\tau_{m,i})|\mathcal{F}_t]$$

The current futures continues to trade during the observation month, therefore the valuation needs to account for already observed values of r_s:

$$A_0^{s1}(t) = \frac{1}{n_0} \sum_{i=1}^{n_0} \mathbf{1}_{(t>\tau_{0,i})} r_s(\tau_{0,i}) + \frac{1}{n_0} \sum_{i=1}^{n_0} \mathbf{1}_{(t\leq\tau_{0,i})} E_t[r_s(\tau_{0,i})|\mathcal{F}_t]$$

rewriting:

$$n_0 A_0^{s1}(t) - \sum_{i=1}^{n_0} \mathbf{1}_{(t>\tau_{0,i})} r_s(\tau_{0,i}) = \sum_{i=1}^{n_0} \mathbf{1}_{(t\leq\tau_{0,i})} E_t[r_s(\tau_{0,i})|\mathcal{F}_t]$$

We use (6) and assume that the expectation $E_t\left[r_p(\tau_{0,i})\right]$ is inferred from Fed Funds futures. Also noting that there is one end of month date z per month, therefore only one possible spike in the reference month of each futures contract, we have

$$n_0 A_0^{s1}(t) - \sum_{i=1}^{n_0} \left(\mathbf{1}_{(t>\tau_{0,i})} r_s(\tau_{0,i}) + \mathbf{1}_{(t\leq\tau_{0,i})} E_t\left[r_p(\tau_{0,i})\right]\right)$$

[28]Source:https://www.cmegroup.com/trading/interest-rates/stir/one-month-sofr_contract_specifications.html

$$= \sum_{i=1}^{n_0} \mathbf{1}_{(t \leq \tau_{0,i})} \left(q(t) e^{-\theta(\tau_{0,i}-t)} + \mu(1 - e^{-\theta(\tau_{0,i}-t)}) \right) + E\left[\Delta J^u(z_0) \right]$$

therefore

$$E\left[\Delta J^u(z_0)\right] = n_0 A_0^{s1}(t) - \sum_{i=1}^{n_0} \left(\mathbf{1}_{(t>\tau_{0,i})} r_s(\tau_{0,i}) + \mathbf{1}_{(t \leq \tau_{0,i})} \left(E_t\left[r_p(\tau_{0,i}) \right] + q(t) e^{-\theta(\tau_{0,i}-t)} + \mu(1 - e^{-\theta(\tau_{0,i}-t)}) \right) \right) \tag{15}$$

For subsequent contracts proceed iteratively using:

$$A_m^{s1}(t) = \frac{1}{n_m} \sum_{i=1}^{n_m} E_t[r_s(\tau_{m,i}) | \mathcal{F}_t]$$

We use (6) and assume that the expectation $E_t\left[r_p(\tau_{m,i})\right]$ is inferred from Fed Funds futures, and also separating spikes for which expectations have been calculated in from prior futures, we have

$$n_m A_m^{s1}(t) - \sum_{i=1}^{n_m} E_t\left[r_p(\tau_{m,i}) \right] = \sum_{i=1}^{n_m} \left(q(t) e^{-\theta(\tau_{m,i}-t)} + \mu(1 - e^{-\theta(\tau_{m,i}-t)}) + \widetilde{J}_m^u \right) + E\left[\Delta J^u(z_{n_m}) \right]$$

therefore

$$E\left[\Delta J^u(z_{n_m}) \right] = n_m A_m^{s1}(t) - \sum_{i=1}^{n_m} \left(E_t\left[r_p(\tau_{m,i}) \right] + q(t) e^{-\theta(\tau_{m,i}-t)} + \mu(1 - e^{-\theta(\tau_{m,i}-t)}) + \widetilde{J}_m^u \right)$$

4.3 Empirical results

4.3.1 Jump expectations from leading fed fund futures

We begin our empirical analysis by applying (12) to daily prices of leading Fed Funds futures (next to expire) for the period starting from August 2000 to April 2019. Application of (12) to historical data yields the expected jump values leading up to that month's FOMC meeting, implied from the futures price. For the months since the introduction of the target range, the target rate $r_f(\tau_{0,i})$ is not directly available and has to be inferred from the data. This section employs the simple approach of assuming it is equal to the most recently observed EFFR. However, as we demonstrate in Section 4.3.3, there is evidence that since 2018 the IOER has been acting — and therefore can be used — as the effective target rate. Expected jumps using leading futures are only relevant for months coinciding with an FOMC meeting. Months during which FOMC does not occur are analysed in Section 4.3.2. The results for each month are summarised as an average of the daily implied jump expectations leading up to the FOMC meeting and are compared to actual target rate change in Figure 8.

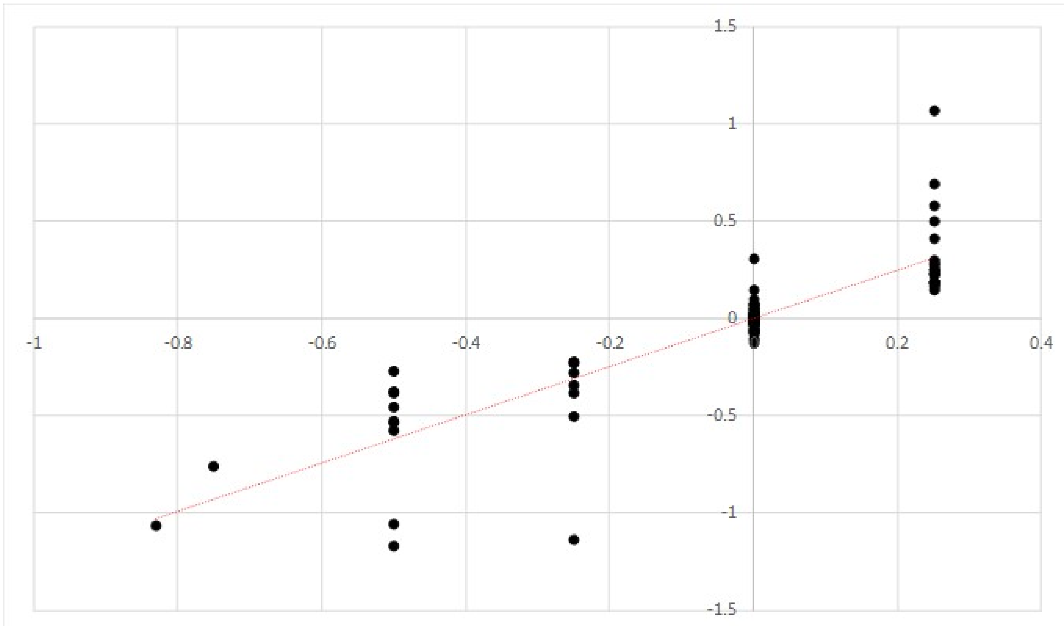

Figure 8: Actual jumps (horizontal axis) versus average expected jumps (vertical axis)

In Figure 8, we can see that all positive changes to the target rate have been of the same 25 basis point magnitude, while negative changes are split between 25 and 50 basis points and include two occasions with negative jumps larger than 50 basis points. The results reveal a high correlation between the expected and actual jumps. This is not surprising given the FOMC strategy to communicate its intention, at least in terms of direction of move, prior to any monetary policy changes. This is apparent in the data, all the results show the market anticipates the correct direction, with only uncertainty regarding the magnitude of the target rate change. For the rest of the section we provide some individual examples highlighting different aspects of futures behaviour. Of particular interest is the behaviour of futures contracts around the time of the financial crisis.

The first example is the the May 2005 contract, which demonstrates the most accurate monetary policy change expectation in our result set. As can be seen in Figure 9, in December 2004 the policy target rate was 2% while the May 2005 future was implying an average EFFR for May 2005 of just under 2.8%, anticipating at least three 25 basis point rate increases. Following the first increase in December, the futures price moved towards anticipation of another three increases. The price did not react to the February and March target rate increases, implying the policy changes were occurring in line with expectation. The final target rate increase in the future's reference month was accurately priced in by the beginning of April, with the futures price exhibiting only minor price changes from that point.

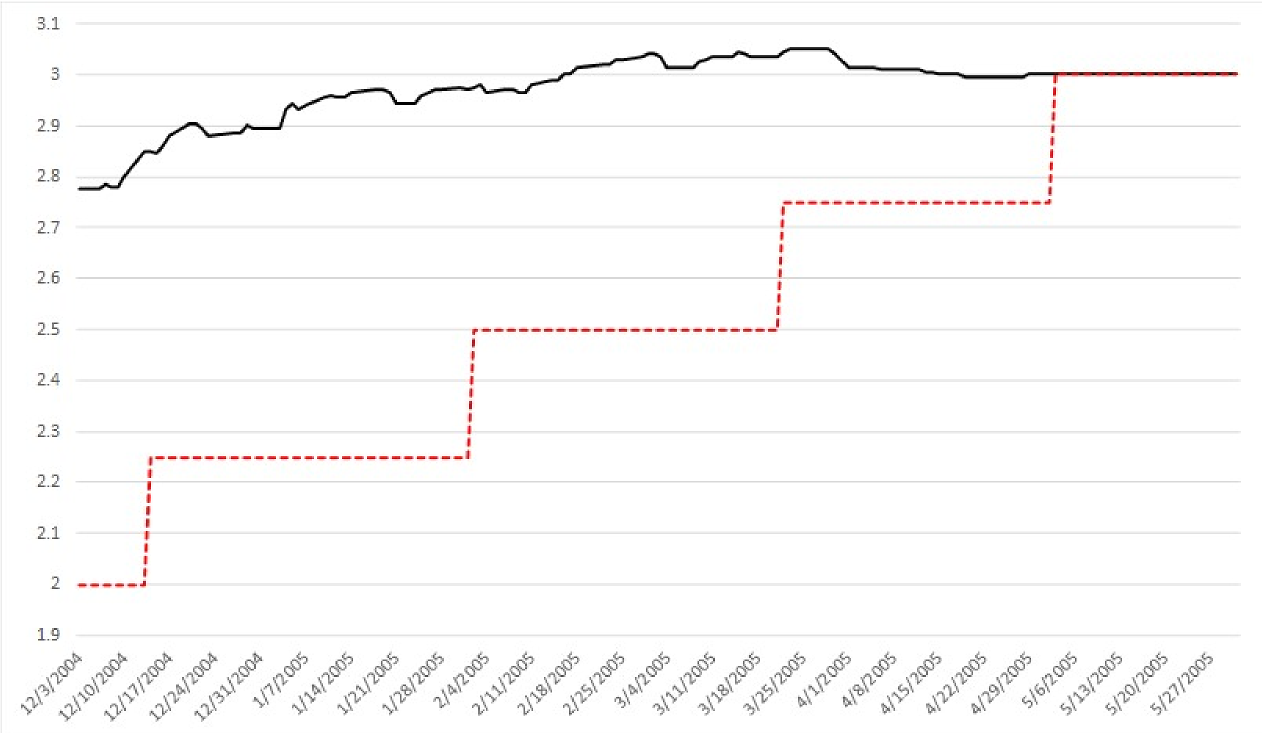

Figure 9: May 2005 Fed Funds future implied average(black) vs Fed Funds target rate (red)

The EFFR, particularly prior to 2008, exhibited significant fluctuation around the target rate. We account for this in the EFFR model with a zero mean noise term. It is however possible that part of the fluctuation is due to systematic factors such as changing activity on specific days due to certain regulatory requirements. This shows up in data prior to 2008 as systematic end of month fluctuations. There is also evidence in the leading month futures price data that these systematic factors are anticipated by the market, the June 2005 contract provides the best example of this. The first noticeable feature of this contract, as can be seen in Figure 10, is that the price and therefore the implied average does not change for the entire reference month, as well as some time prior to that. The implied average is only slightly above the target rate. However, since the FOMC is scheduled for the last day of the month, all of the increase in the month's average is based on one day. The implied jump, based on our model, is over 100 basis points (against an actual change of 25 basis points). It is highly unlikely that the market would be anticipating such a change, especially since rates have not increased by more than 25 basis points at least since the year 2000. Also if the market was anticipating such an increase, the price would react when the expectation did not materialise. Instead, noticing that the realised EFFR fixings were well above the target rate for several final days of the month, it appears the market perfectly anticipated these elevated EFFR levels well in advance. This suggests that on occasion there is some spread between the target rate and the EFFR, which is anticipated by the market but not captured in our model.

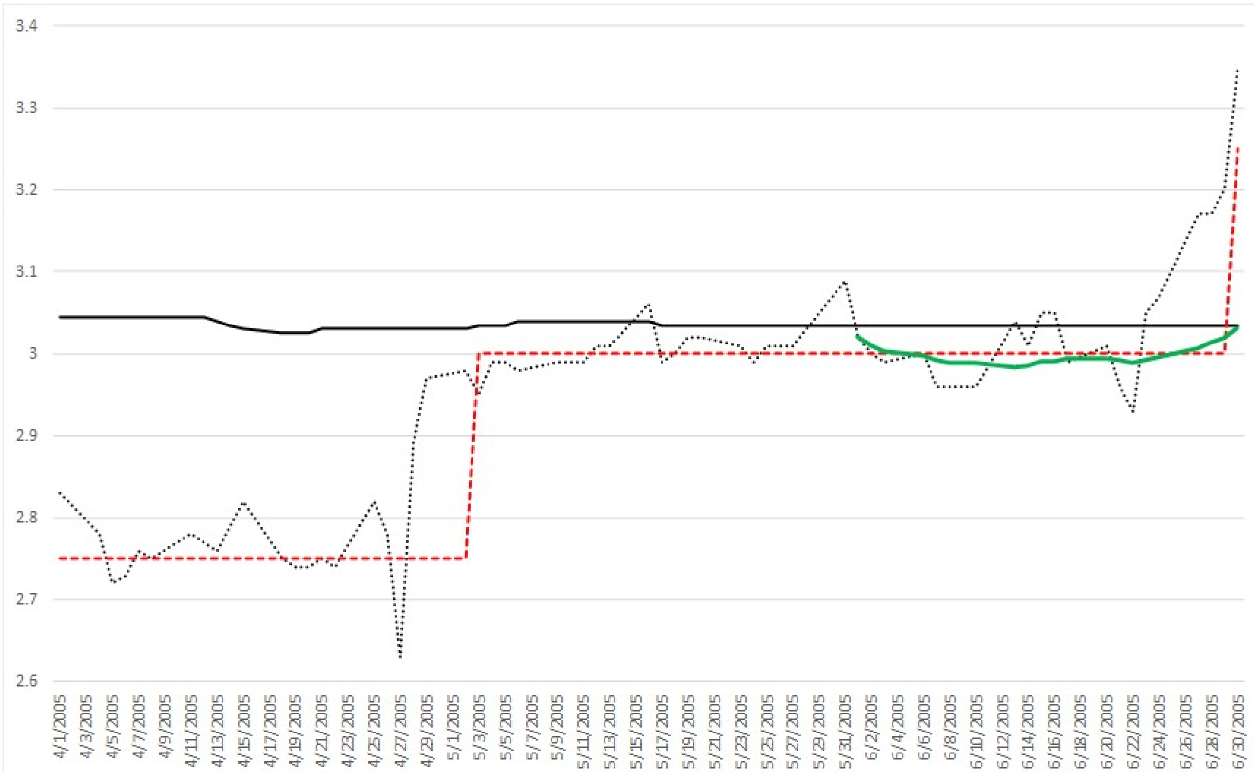

Figure 10: June 2005 Fed Funds future implied average(black), Fed Funds target rate (red), EFFR (black dotted), EFFR fixing running average(green)

The September 2007 futures, shown in Figure 11, is the first reference month impacted by the commencement of aggressive financial easing as a result of the financial crisis. Following a period of steady rate increases, in August of 2006 the FOMC was primarily concerned on controlling high inflation[29] and the September 2007 futures contract was anticipating further target rate increases. The inflation concerns eased slightly by the September 2006 meeting, but apart from a slowdown in home construction and real estate sales the FOMC remained optimistic about the economic outlook.[30] On the other hand, the futures market, as can be seen in Figure 11, was pricing in 25 to 50 basis points of easing over the following year. This expectation had mostly disappeared by December 2006, with the Fed continuing to signal inflation concerns with an economic backdrop of falling activity in real estate balanced by robust consumer spending.[31] Following some poor economic data and the Federal Reserve communicating a softer outlook for inflation, at the beginning of 2007 the futures market again began to anticipate some easing to impact the September 2007 contract month. This anticipation was again extinguished by the March 2007 FOMC, which, while acknowledging declining activity in real estate, remained sanguine about the economic

[29]See Federal Open Market Committee, August (2006).
[30]See Federal Open Market Committee, September (2006).
[31]See Federal Open Market Committee, December (2006).

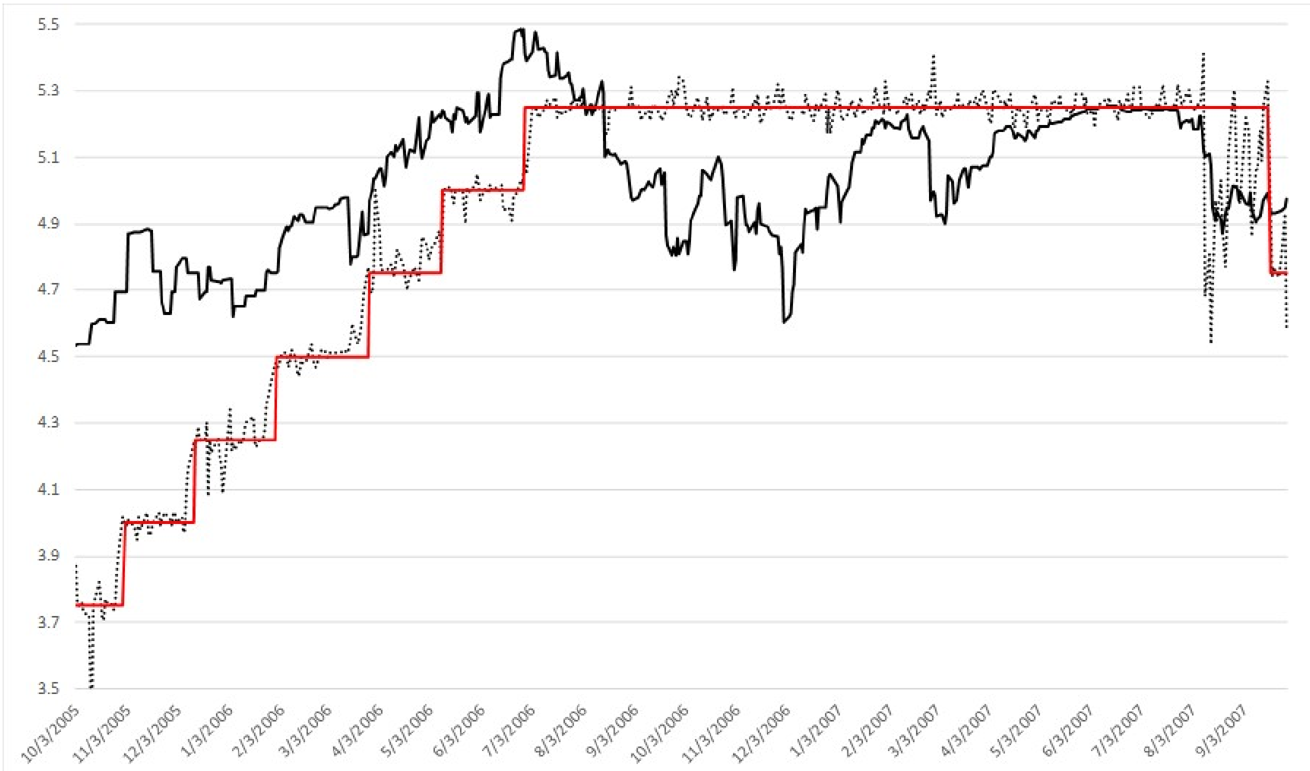

Figure 11: September 2007 Fed Funds future implied average(black), Fed Funds target rate (red), EFFR (black dotted)

outlook and appeared to dispel any expectation about any forthcoming rate decreases.[32] By August 2007 the slowdown in real estate had accelerated into credit market turmoil with several key events[33], eventually spilling over to lack of liquidity in money markets resulting in a jolt in volatility in the EFFR and immediate anticipation of monetary easing at the September 2007 FOMC.

The final example coincides with the largest monthly change in the target rate, which occurred in January of 2008 and consisted of a decrease of 75 basis points after an emergency meeting followed by another 50 basis points announced at the conclusion of the regularly scheduled FOMC.[34] The behaviour of the January 2008 futures with respect to the target rate and EFFR is shown in Figure 12. The green line in the figure represents the expected final target rate derived from the futures prices during the reference month. At the beginning of January and prior to the initial emergency rate drop, the futures market anticipated

[32]See Federal Open Market Committee, March (2007).

[33]These being liquidation and halting of redemption from mortgage back security funds, subprime mortgage lenders' bankruptcy filings and widespread mortgage related bond downgrades, see the financial crisis timeline https://www.stlouisfed.org/financial-crisis/full-timeline

[34]The reasons for the aggressive easing were rapidly decreasing economic activity, money market and credit market stress, see Federal Open Market Committee, January (2008).

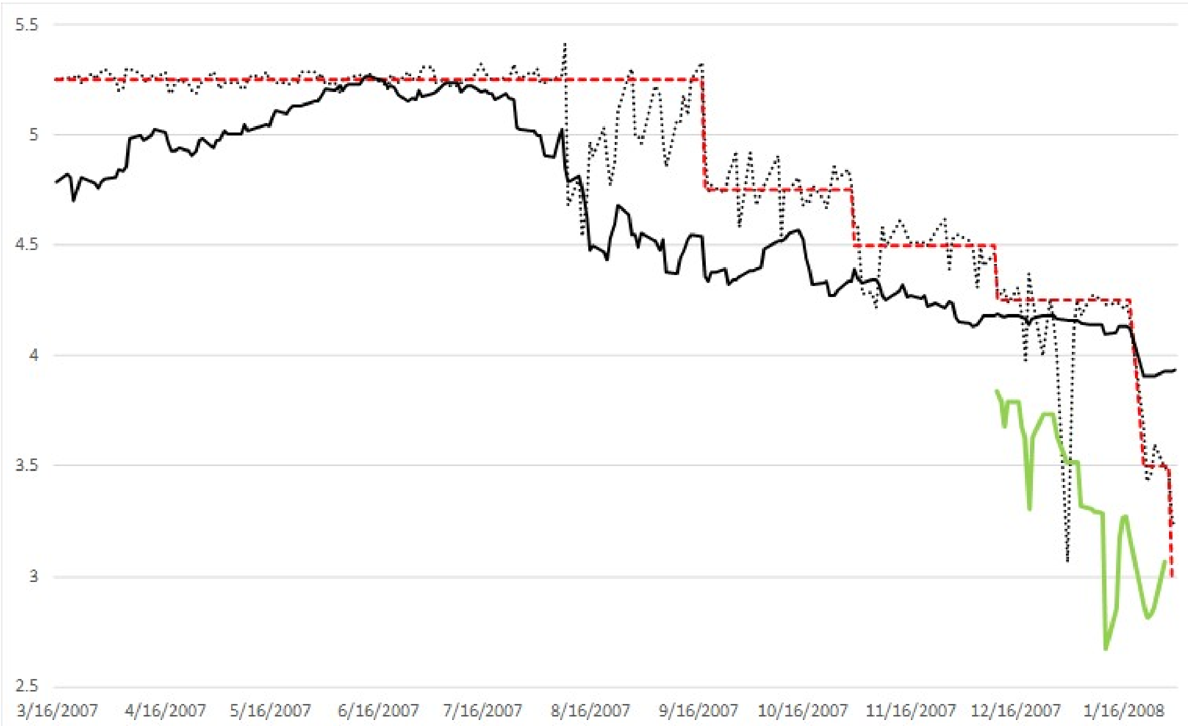

Figure 12: January 2008 Fed Funds future implied average(black), Fed Funds target rate (red), EFFR (black dotted), implied target rate(green)

a 50 to 75 basis point target rate decrease. After the emergency meeting the futures price correctly anticipated another 50 basis point decrease to occur at the regular meeting. Notably, the expected final target rate provides a good example of the increased sensitivity of the expected jump due to changes in the implied average when the FOMC is scheduled on the last day of the month, i.e. small changes in the implied average result in large changes of the expected jump.

4.3.2 Target rate expectations from leading fed fund futures

In this section we examine the correspondence of the implied target $\widehat{r}_p(t)$ and the actual target rate $r_p(t)$ for months without any scheduled FOMC meetings prior to the introduction of the target range. The implied target rate is obtained by setting $E_t[J(\tau_{0,\widehat{n}_0})] = 0$ in (12), which yields:

$$(n_0 - n_t + 1)\widehat{r}_p(t) = n_0 A_0^f(t) - \sum_{i=1}^{n_t-1} r_f(\tau_{0,i})$$

therefore

$$\widehat{r}_p(t) = f(A_0^f(t); n_0, n_t, r_f) = \frac{n_0 A_0^f(t) - \sum\limits_{i=1}^{n_t-1} r_f(\tau_{0,i})}{n_0 - n_t + 1}$$

The accuracy of the implied target rate is limited to the minimum change (tick size) of the futures contract price. During the course of the reference month, as more of the average payoff becomes known and fixed, the sensitivity of the implied target rate to changes in implied average increases, that is as $\left\| \frac{\partial \widehat{r}_p(t)}{\partial A_0^f(t)} \right\| \uparrow$ as $n_t \to n_0$. Since $A_0^f(t)$ is a function of the futures price, the minimum changes in the futures price have an increasing impact on the implied target rate. To account for this, we measure the error as the square of the distance of the actual target rate outside a minimum price change boundary for the implied target rate, that is:

$$error = max\left[max\left(\widehat{r}_p^{+}(t) - r_p(t), 0 \right)^2, min\left(\widehat{r}_p^{-}(t) - r_p(t), 0 \right)^2 \right]$$

where, given a tick size δ, $\widehat{r}_p^{+}(t) = f(A_0^f(t) + \delta/100; n_0, n_t, r_f)$ and $\widehat{r}_p^{-}(t) = f(A_0^f(t) - \delta/100; n_0, n_t, r_f)$. The error is measured on each day of each month for which there are no scheduled FOMC meetings, therefore no expected jumps.

The most prominent feature of the results, shown in Figure 13, is an extreme dislocation as a result of the financial crisis, providing another aspect of stressed market conditions at the time of the crisis. Other deviations are associated with expectations related to deviations in EFFR end of month fixings. Similarly to what is currently happening with SOFR, the EFFR exhibited large deviations related to end of month dates prior to the financial crisis.

4.3.3 Fed Funds convergence to IOER

The use of the IOER as the target rate by the FOMC is one of the interesting outcomes of the financial crisis, particularly since its introduction was not intended for this purpose. The effectiveness of IOER appears to be closely linked to total excess reserves. While excess reserves were extremely high, the IOER served as an upper bound for the EFFR. However since normalisation has effectively reduced the amount of excess reserves, the EFFR has converged to the IOER and has spent the latter part of 2018 and early 2019 effectively glued to the IOER with very little variance. We demonstrate this in Figure 14, using the same approach as in the previous section, by replacing the target rate with the IOER. The key difference to the previous section is that the IOER was not initially used as the target rate, but rather as the upper bound. It was only while deliberating the normalisation strategy, the FOMC decided to use the IOER as a target rate. Therefore, in distinction to the previous section, the results show the transformation of the IOER from the upper bound to the target rate. In the month of April 2019, the last month in the data set, the EFFR has moved above the IOER, suggesting that this could be an equilibrium point where the amount of excess reserves is just right to allow for the use of IOER as the target rate.

4.3.4 Fed Funds futures cross sectional calibration

We now examine the performance of the calibration methodology on a cross sectional term structure of twelve monthly Fed Fund futures contracts prices for each day from 2003 to

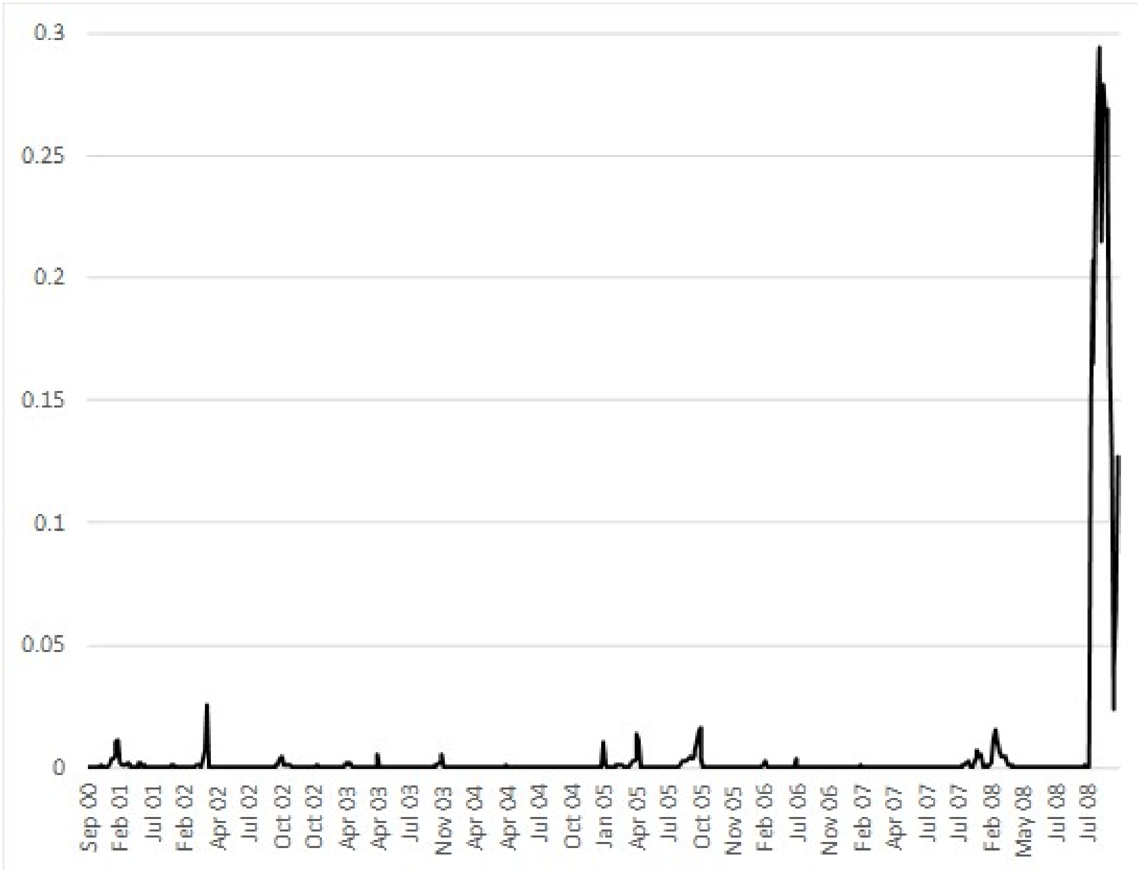

Figure 13: Squared error between actual target rate and implied target rate boundaries for months without an FOMC meeting

2018, effectively all available Fed Fund futures data. The assessment is focused on two desirable features of the calibration: consistency in the direction of expected jumps and consistency between net expected jumps and non-FOMC months.

The Federal Reserve tends to operate in easing and tightening cycles, translating to sustained periods of persistent direction of target rate changes. The monetary policy views are carefully communicated to the market for the purpose of making clear the likely direction and urgency of the next rate change. This should be reflected in the calibration by persistent direction of expected target rate change. At most one change of direction within the one year period should be present marking the time when the market is expecting a change in the cycle. It is highly unlikely that the market would expect the target rate to change direction any more than once within this period. The presence of any more expected rate direction changes suggests a problem with the model and/or calibration. To measure this aspect of the calibration, the number of expected rate change direction turning points are counted for each calibrated day.

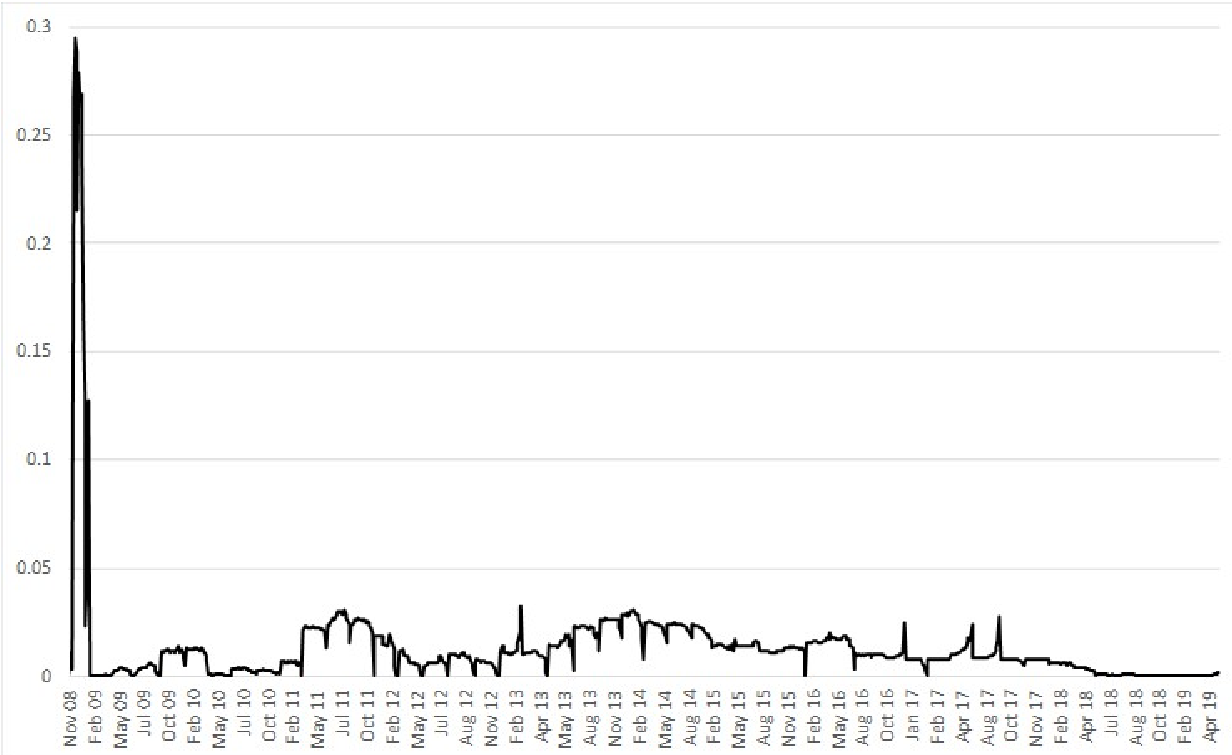

Figure 14: Squared error between IOER and implied target rate boundaries for months without an FOMC meeting

Each calendar year contains four non-FOMC months, i.e. months for which there is no scheduled FOMC meeting. According to our model, the expected target rate for non-FOMC months should be equal to the average rate implied by the contract. Any inconsistency can be quantified for each non-FOMC month m^* as follows:

$$e_{m^*}(t) = A^f_{m^*}(t) - r_p(t) - \widetilde{J}_{m^*} \tag{16}$$

The absolute sum of these errors within each calibrated day is used to measure this aspect of calibration performance. In general, the model performs well over the 15 year test period, however some heuristics are required to deal with certain structural aspects as will be reported in this the section.

The first example of calibration results is shown in Table 1, where months for which an FOMC meeting is scheduled are indicated by the presence of the meeting date in the FOMC column. For the FOMC months, the expected jump $E[J]$ is calculated using the approach outlined in Section 4.1. The table also shows the expected target rate $E_t[r_p(u)]$ for the end date of each month, and for comparison the realised target rate $r_p(u)$. This calibration is performed for each day from 2013 to 2018 and the defined performance measures are recorded for analysis.

Contract	Price	FOMC	$E[J]$	$E_t[r_p(u)]$	$r_p(u)$
Jan 2018	98.5875	31-Jan	-0.1325	1.3675	1.5
Feb 2018	98.58		0.0	1.3675	1.5
Mar 2018	98.51	21-Mar	0.3453	1.7128	1.5
Apr 2018	98.355		0.0	1.7128	1.75
May 2018	98.34	2-May	-0.0545	1.6582	1.75
Jun 2018	98.25	13-Jun	0.1529	1.8112	1.75
Jul 2018	98.175		0.0	1.8112	1.95
Aug 2018	98.14	1-Aug	0.0488	1.86	1.95
Sep 2018	98.14	26-Sep	0.0	1.86	1.95
Oct 2018	98.015		0.0	1.86	2.2
Nov 2018	97.99	8-Nov	0.1957	2.0557	2.2
Dec 2018	97.95	19-Dec	-0.0135	2.0422	2.2

Table 1: calibration results from 17th January 2017

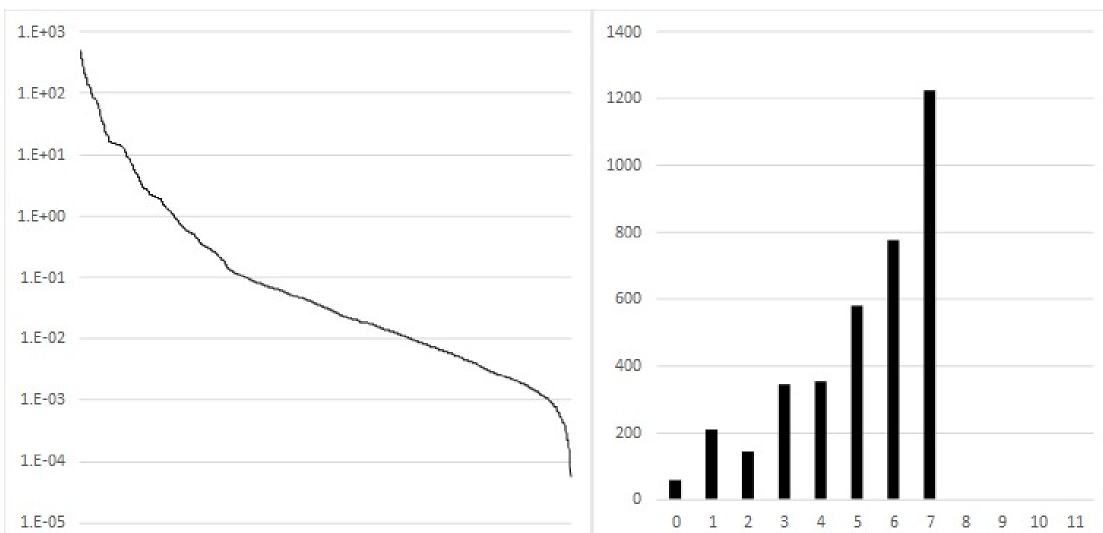

Figure 15: Sorted inconsistency error in log-scale (left), histogram showing number of calibration turning points (right)

The initial results summary, shown in Figure 15, contains large errors for a significant number of calibration days. A histogram of the number of turning points per calibration is also shown, revealing very few of the days fitting the criteria that there should be at most one turning point. For a closer examination, the calibration results shown in Table 2 focus on one of the extreme examples of expected jump direction changes and non-FOMC month discrepancy from the above results. The expected jump oscillates between contracts and diverges rapidly due to a compensatory behaviour of the bootstrap algorithm. That is, any discrepancies in the magnitude of the target rate change are compensated for by the next

Contract	Price	FOMC	$E[J]$	$E_t[r_p(u)]$	$r_p(t)$
Aug-2007	95.005	7-Aug	0.0	5.2500	5.25
Sep-2007	95.13	18-Sep	-0.8769	4.3731	4.75
Oct-2007	95.225	31-Oct	12.4596	16.8327	4.50
Nov-2007	95.425		0.0	16.8327	4.50
Dec-2007	95.505	11-Dec	-18.2128	-1.3801	4.25
Jan-2008	95.565	30-Jan	90.1341	88.7540	3.00
Feb-2008	95.665		0.0	88.7540	3.00
Mar-2008	95.69	18-Mar	-186.9831	-98.2291	2.25
Apr-2008	95.735	30-Apr	3074.8230	2976.5939	2.00
May-2008	95.77		0.0	2976.5939	2.00
Jun-2008	95.77	25-Jun	-14861.8195	-11885.2256	2.00
Jul-2008	95.76		0.0	-11885.2256	2.00

Table 2: calibration results from 21st August 2007

contract and compounded by any new discrepancies. In the case demonstrated in Table 2, this leads to divergent behaviour.

As a temporary fix to remove the divergent behaviour and gauge its impact on the results, consider the model based on (13) in the context of non-FOMC months. Since $E_t[J(\tau_{m,\widehat{n}_m})] = 0$ we have $A_m^f(t) = r_p(t) + \widetilde{J}_m$. However the accumulated jump expectation $\widetilde{J}_m$ is entirely based on the futures price from prior contract months, therefore $E_t[J(\tau_{m,\widehat{n}_m})] = 0$ only occurs if futures prices are perfectly consistent in the context of the model. Any inconsistency is passed on to the next contract allowing it to be accumulated and compounded with further inconsistencies, leading to the divergence observed in Table 2. The simplest way to ensure that $A_m^f(t) = r_p(t) + \widetilde{J}_m$ holds is to feed the error calculated in (16) into the calibration as an expected jump for the non-FOMC month, i.e. letting $A_m^f(t) = r_p(t) + \widetilde{J}_m + e_m(t)$, thus ensuring that for the following month, the accumulated jumps are consistent with the previous months price. In this setup the non-FOMC month acts as an absorber of any inconsistency effectively resetting any occurring divergence.

The results after applying the above fix are compared to the original results in Figure 16, showing that a large section of the largest errors are eliminated. This confirms that the divergent behaviour is the cause of the largest errors in terms of magnitude. The distribution of turning points per calibration has also shifted to the left indicating an improvement in this metric. The improvement is also directly evident in Table 3, showing the same calibration as Table 2 but including the new treatment for non-FOMC months. The October 2007 contract in the table is still showing an unrealistically large expected jump, however the divergent behaviour is eliminated because all of the rate change is absorbed by the following November 2007 contract.

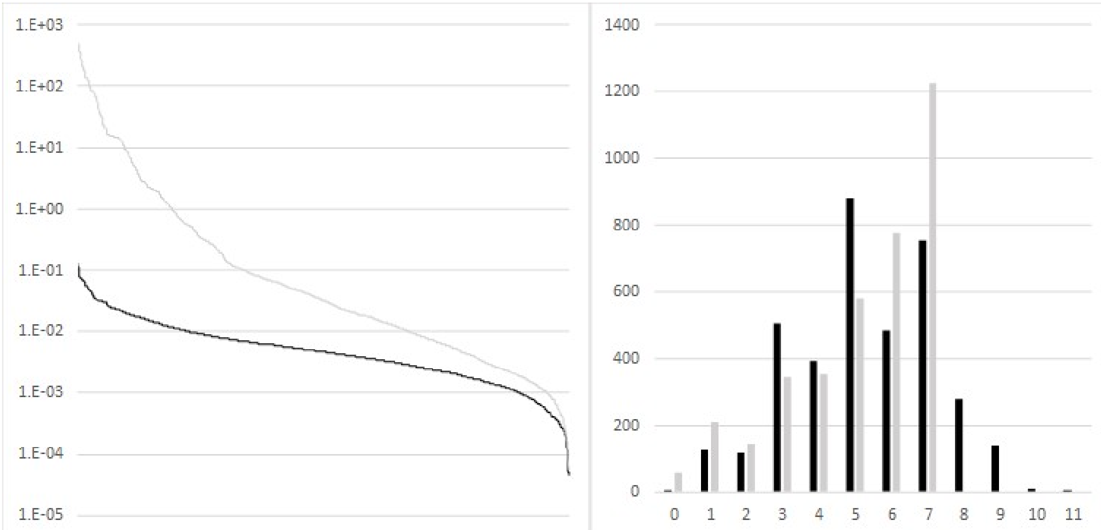

Figure 16: Sorted inconsistency error in log-scale (left), histogram showing number of calibration turning points (right) before(grey) and after applying jump errors to non-FOMC months(black)

Contract	Price	FOMC	$E[J]$	$E_t[r_p(u)]$	$r_p(t)$
Aug-2007	95.005	7-Aug	0.0	5.2500	5.25
Sep-2007	95.13	18-Sep	-0.8769	4.3731	4.75
Oct-2007	95.225	31-Oct	12.4596	16.8327	4.50
Nov-2007	95.425		-12.2577	4.5750	4.50
Dec-2007	95.505	11-Dec	-0.1181	4.4569	4.25
Jan-2008	95.565	30-Jan	-0.3394	4.1175	3.00
Feb-2008	95.665		0.2175	4.3350	3.00
Mar-2008	95.69	18-Mar	-0.0553	4.2797	2.25
Apr-2008	95.735	30-Apr	-0.4400	3.8397	2.00
May-2008	95.77		0.3903	4.2300	2.00
Jun-2008	95.77	25-Jun	0.0	4.2300	2.00
Jul-2008	95.76		0.0100	4.2400	2.00

Table 3: calibration results from 21st August 2007 with jumps in non FOMC months

The October 2007 contract from Table 3 reveals another problematic aspect of this calibration. FOMC meetings occurring on the last day have an amplifying effect on the expected jump from small deviations in the futures price. This is because all else being equal, the average rate for the month implied from the futures prices has to be absorbed by just the one final day of the month. To assess the impact of this amplification effect those meetings are temporarily removed and the contracts are treated as occurring on non-FOMC months. The resulting improvement, shown in Figure 17, is this time more apparent in the distribution of turning points which has again shifted to the left. Only some of the largest

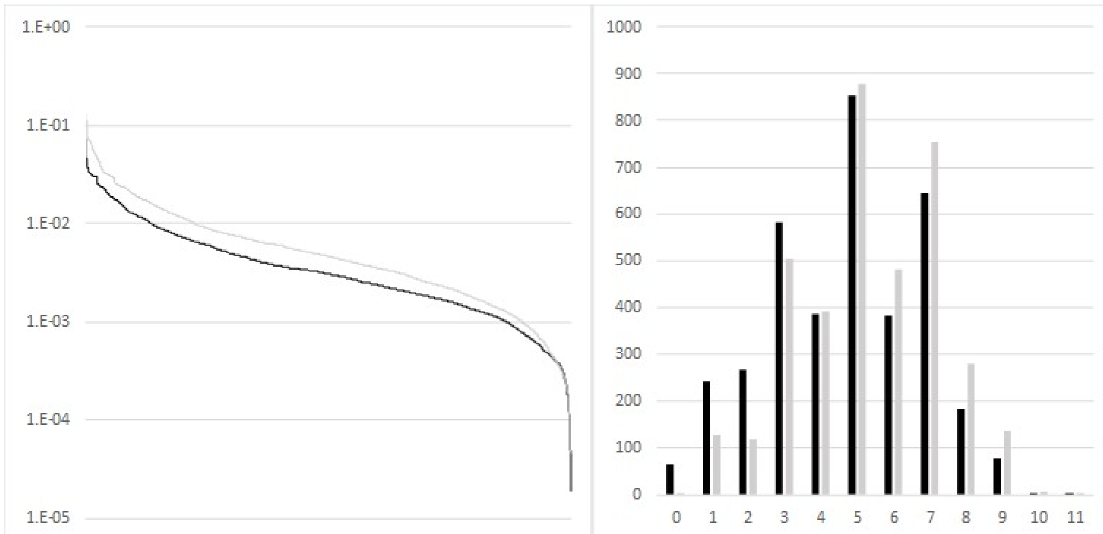

Figure 17: Sorted inconsistency error (left), histogram showing number of calibration turning points (right) before (grey) and after removing jumps occurring on the last day of the month (black)

Contract	Price	FOMC	$E[J]$	$E_t[r_p(u)]$	$r_p(t)$
Aug-2007	95.005	7-Aug	0.0	5.2500	5.25
Sep-2007	95.13	18-Sep	-0.8769	4.3731	4.75
Oct-2007	95.225		0.4019	4.7750	4.50
Nov-2007	95.425		-0.2000	4.5750	4.50
Dec-2007	95.505	11-Dec	-0.1181	4.4569	4.25
Jan-2008	95.565	30-Jan	-0.3394	4.1175	3.00
Feb-2008	95.665		0.2175	4.3350	3.00
Mar-2008	95.69	18-Mar	-0.0553	4.2797	2.25
Apr-2008	95.735		-0.0147	4.2650	2.00
May-2008	95.77		-0.0350	4.2300	2.00
Jun-2008	95.77	25-Jun	0.0	4.2300	2.00
Jul-2008	95.76		0.0100	4.2400	2.00

Table 4: calibration results from 21st August 2007 with jumps in non FOMC months, EOM FOMC dates removed

errors have disappeared and therefore the impact is not as obvious by simple inspection of Figure 17.

The calibration from 21st August shown in Table 4 provides a more direct example of the impact of treating months with FOMC meetings occurring on the last day as non-FOMC months. Removing the largest sources of discrepancy have revealed that the remaining error can be largely attributed to an overestimation of the magnitude of jump size in the September 2007 contract. The reason for the overestimation can be gleaned by inspection of the price of the first contract, which given that the calibration date is past the FOMC date implies a significant spread between the EFFR and the target rate. The calibration date coincides with the onset of the financial crisis where the Federal Reserve began to lose the ability to control the EFFR with open market operations, the size of the implied spread in this case is entirely consistent with the realised EFFR to target rate spread. Applying a spread for the front two contracts of this calibration and removing the temporary fixes described up to this point actually eliminates the chain reaction leading to divergent results.

The spread between EFFR and the target rate reflects the ever changing structure of the Fed Funds market. It varies from very short term fluctuations to extended periods of persistent spread of similar magnitude and direction. The magnitude of the spread is reflected in futures prices, the period of time the spread is expected to persist embeds itself in a spread term structure in the futures prices. It is challenging to precisely disentangle expectation of spread and target rate change, we do find three heuristics for dealing with the evolving aspect of the Fed Funds market which significantly improve the performance of the calibration.

During the extended period of near zero interest rates between 2008 and 2015, the EFFR persisted to set below the IOER target rate. It is reasonable to assume that during this period deviations in futures prices were due to changing expectations of the realised spread rather than any expectation of target rate changes. This period is treated as a special case where the spread is equal to the difference between the target rate and the futures implied average for each of the twelve calibrated contracts, thereby effectively eliminating any expectation of target rate changes.

Spreads expected to last for a shorter time, such as the instability at the onset of the financial crisis, are dealt with a spread applied to only the front futures contracts. Trial and error over the 15 year test period suggests that for best results the spread should be applied to the front three contracts. In this case the remaining contracts are also allowed a spread to take into account minimum price increment related noise.[35] It is assumed that the settlement price reflects either the bid or the offer with a spread of one price increment. Tick size changes in the settlement price don't necessarily reflect changing expectations but rather normal trading fluctuations between the bid and offer price, sometimes referred to as the bid ask bounce. This is dealt with in the calibration by allowing a spread with magnitude bounded by the tick size for the back 9 contracts in the calibration. The spreads applied to all the contracts, as described above, are calibrated in a bootstrap fashion to minimise the error described in (16). The results after removing the temporary fixes and applying the spread are shown in comparison to previous results in

[35] The minimum price increment for non leading month futures is 0.005, see https://www.cmegroup.com/trading/interest-rates/stir/30-day-federal-fund_contract_specifications.html

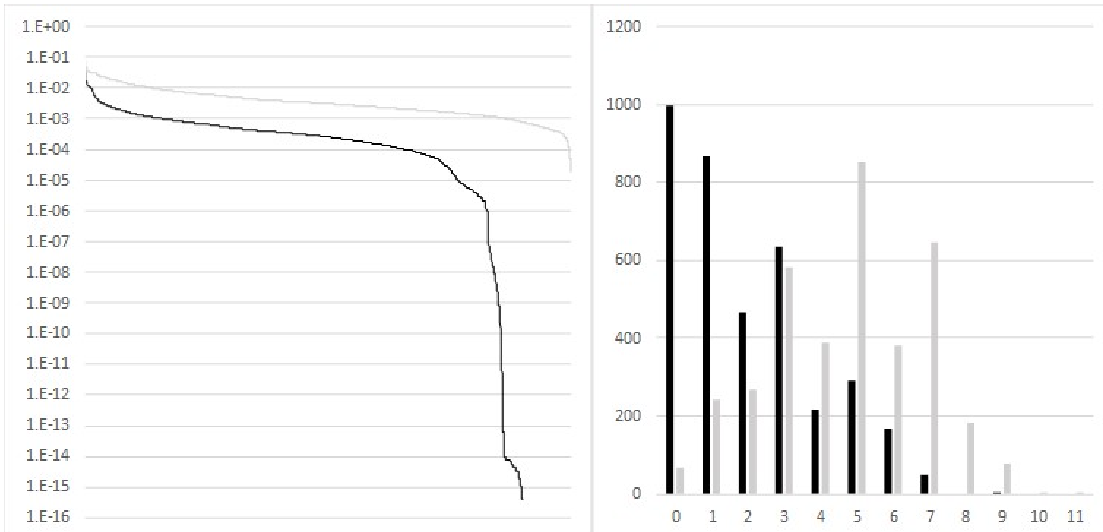

Figure 18: Sorted inconsistency error (left), histogram showing number of calibration turning points (right) before (grey) showing latest results with temporary fixes and after (black) removal of temporary fixes and applying spreads

Figure 18. There is notable improvement in both the errors and the distribution of turning points. This demonstrates that even a fairly crude general treatment of spread is enough to achieve reasonable calibration performance. More elaborate approaches could be developed to further improve the results.

Finally, we examine the ability of the Fed Funds futures market to predict one year ahead changes in the actual Fed Funds target rate. To do this we compare the one year ahead Fed Funds target rate change expectation derived using our calibration approach with the actual change over the same time period. The results, shown in Figure 19, reveal two main features. Firstly, the last 15 years cover one rate cycle of monetary tightening and easing, with the first sequence of rate increases beginning in 2003, followed by easing to an unprecedented near zero target range, with the next cycle of increases not beginning until late 2014. Secondly, over this period the Fed Funds market tends to underestimate rate changes in either direction, this is evident in the initial tightening, the aggressive easing related to the financial crisis and the eventual normalisation period.

4.3.5 SOFR parameter estimation

The calibration for the SOFR futures begins with a maximum likelihood estimation of the Ornstein-Uhlenbeck process parameters from equation (4) based on the methodology described in Franco (2003). The maximum likelihood estimate is made complicated by the inclusion of end of month spikes. Explicitly including the spikes in the maximum likelihood estimate would require establishing a distribution of the spikes, which at the least would have seasonal mean and variance. Instead of attempting to disentangle the

Figure 19: Actual (red) vs expected (black) one year ahead target rate changes

Ornstein-Uhlenbeck and spike distribution variance we exclude the impact from end of month days from the estimate. The estimated parameters are then used to obtain the expected Ornstein-Uhlenbeck contribution to realised changes for end of month days, which we can use to extract the estimate for the realised spikes from the data.

Following from Franco (2003) the conditional density f_i expectation of the Ornstein-Uhlenbeck process is:

$$f_i(q_{t_i};\mu,\theta,\sigma) = (2\pi)^{-\frac{1}{2}}\left(\frac{\sigma^2}{2\theta}(1-e^{-2\theta(t_i-t_{i-1})})\right)^{-\frac{1}{2}}\times$$

$$\exp\left[-\frac{(q_{t_i}-\mu-(q_{t_{i-1}}-\mu)e^{-\theta(t_i-t_{i-1})})^2}{2\frac{\sigma^2}{2\theta}(1-e^{-2\theta(t_i-t_{i-1})})}\right]$$

Given $n+1$ observations of $q=\{q_{t_0},...,q_{t_n}\}$, the log-likelihood function, less the constant terms is:

$$\ell(\mathbf{q};\mu,\theta,\sigma) = -\frac{n}{2}\log\left[\frac{\sigma^2}{2\theta}\right] - \frac{1}{2}\sum_{i=1}^{n}\log\left[1-e^{-2\theta(t_i-t_{i-1})}\right]$$

$$-\frac{\theta}{\sigma^2}\sum_{i=1}^{n}\frac{(q_{t_i}-\mu-(q_{t_{i-1}}-\mu)e^{-\theta(t_i-t_{i-1})})^2}{1-e^{-2\theta(t_i-t_{i-1})}}$$

The maximum log-likelihood estimators $\hat{\mu}$, $\hat{\theta}$, $\hat{\sigma}$ function must satisfy the following first order conditions:

$$\left.\frac{\partial \ell(\mathbf{q};\mu,\theta,\sigma)}{\partial \mu}\right|_{\hat{\mu}} = 0$$

$$\left.\frac{\partial \ell(\mathbf{q};\mu,\theta,\sigma)}{\partial \theta}\right|_{\hat{\theta}} = 0$$

$$\left.\frac{\partial \ell(\mathbf{q};\mu,\theta,\sigma)}{\partial \sigma}\right|_{\hat{\sigma}} = 0$$

The solution is simplified by rewriting the first order conditions as a function of θ.

$$\frac{\partial L(\mathbf{q};\mu,\theta,\sigma)}{\partial \mu} = \frac{2\theta}{\sigma^2}\sum_{i=1}^{n}\frac{q_{t_i}-\mu-(q_{t_{i-1}}-\mu)e^{-\theta(t_i-t_{i-1})}}{1+e^{-\theta(t_i-t_{i-1})}}$$

$$\hat{\mu} = f(\hat{\theta}) = \sum_{i=1}^{n}\frac{q_{t_i}-q_{t_{i-1}}e^{-\hat{\theta}(t_i-t_{i-1})}}{1+e^{-\hat{\theta}(t_i-t_{i-1})}}\left(\sum_{i=1}^{n}\frac{1-e^{-\hat{\theta}(t_i-t_{i-1})}}{1+e^{-\hat{\theta}(t_i-t_{i-1})}}\right)^{-1}$$

$$\frac{\partial L(\mathbf{q};\mu,\theta,\sigma)}{\partial \sigma} = -\frac{n}{\sigma}+\frac{2\theta}{\sigma^3}\sum_{i=1}^{n}\frac{(q_{t_i}-\mu-(q_{t_{i-1}}-\mu)e^{-\theta(t_i-t_{i-1})})^2}{1-e^{-2\theta(t_i-t_{i-1})}}$$

$$\hat{\sigma} = g(\hat{\mu},\hat{\theta}) = \sqrt{\frac{2\hat{\theta}}{n}\sum_{i=1}^{n}\frac{(q_{t_i}-\hat{\mu}-(q_{t_{i-1}}-\hat{\mu})e^{-\hat{\theta}(t_i-t_{i-1})})^2}{1-e^{-2\hat{\theta}(t_i-t_{i-1})}}}$$

$$V(\theta) = -\frac{n}{2}\log\left[\frac{g(f(\theta),\theta)^2}{2\theta}\right]-\frac{1}{2}\sum_{i=1}^{n}\log\left[1-e^{-2\theta(t_i-t_{i-1})}\right]$$
$$-\frac{\theta}{g(f(\theta),\theta)^2}\sum_{i=1}^{n}\frac{(q_{t_i}-f(\theta)-(q_{t_{i-1}}-f(\theta))e^{-\theta(t_i-t_{i-1})})^2}{1-e^{-2\theta(t_i-t_{i-1})}}$$

The end of month movement is excluded by assuming that any instances where t_i is the last observation date of the month $q_{t_i} = q_{t_{i-1}}e^{-\theta(t_i-t_{i-1})}$. This removes the impact from end of month spikes. However, to include the mean reversion following an end of month spike, instances where t_{i-1} is the last observation date of the month are included in the calculation. Maximising $V(\theta)$, yields the maximum likelihood estimate for parameters μ, θ and σ.

The results, summarised in Table 5 are produced for three cases using all SOFR rates since publication, dating back to April 2018. The first two cases compare results including and excluding the end of month data point. The comparison reveals that the impact of including spikes on estimation is consistent with expectation. The value μ to which the mean reversion tends, increases since all the spikes are in the positive direction. There is also a significant increase in the volatility parameter σ. Interestingly the mean reversion speed parameter θ increases suggesting in most cases the spread is below the mean reversion target and therefore some of the spike movement is accounted for by rapid reversion to the mean.

Parameter	Incl. EOM	Excl. EOM	Beg. of Month
μ	0.0002	-0.0002	0.0001
σ	0.0113	0.0078	0.0077
θ	89.64	46.08	153.32

Table 5: OU parameter estimate summary for SOFR from 2-Apr-2018 to 22-Jul-2019

The third case shows the maximum likelihood for data including only the first seven days of each month. The spikes occur at the end of each month and therefore decay during this period. The larger number indicates that this decay is faster than the dynamics not related to the spike for the remainder of the month. Implementation of the calibration to futures prices discussed in the next section suggest that the mean reversion speed parameter needs to be large enough to eliminate the residual expectation after the occurrence of a spike within a few days. Too slow decay following a spike leads to empirically inconsistent near term expectations of the SOFR rate and results in poor calibration results, i.e. the following spike expectations become unrealistic.

Estimating the θ parameter from only the front of the month leads to an overestimate, however the consequence of overestimation on days not related to the spike don't have a material impact on the calibration. The calibration is sensitive to the speed of decay of the spike particularly for large spikes, as such we use the higher mean reversion speed for the calibration described in the next section. This choice is further justified by the characteristics of the likelihood function. As a function of θ, as described above, the likelihood function around the maximum point is quite flat for values of θ between 50 and 250, meaning relatively equal likelihood over this domain.

4.3.6 End of month spike expectations from leading SOFR futures

This section examines the performance of the calibration to 1m SOFR futures defined in Section 4.2.1 for leading SOFR futures. (15) is applied to every day of available data since the inception of SOFR futures trading in May 2018, to yield the expected spike for the corresponding month. The results are initially summarised in Figure 20 as the average of the daily expected spike for each day of the month in comparison to the actual spike for that month.

The results reveal an inconsistent correspondence between predicted and actual spike. The direction of the predicted spike is generally positive which is consistent with every realised spike, however in general the expectation tends to overestimate the eventual size of the spike. The inconsistency is due to both model mis-specification as well as legitimate market uncertainty with respect to the magnitude of the spike. We will explore the evidence for both cases for the remainder of this section.

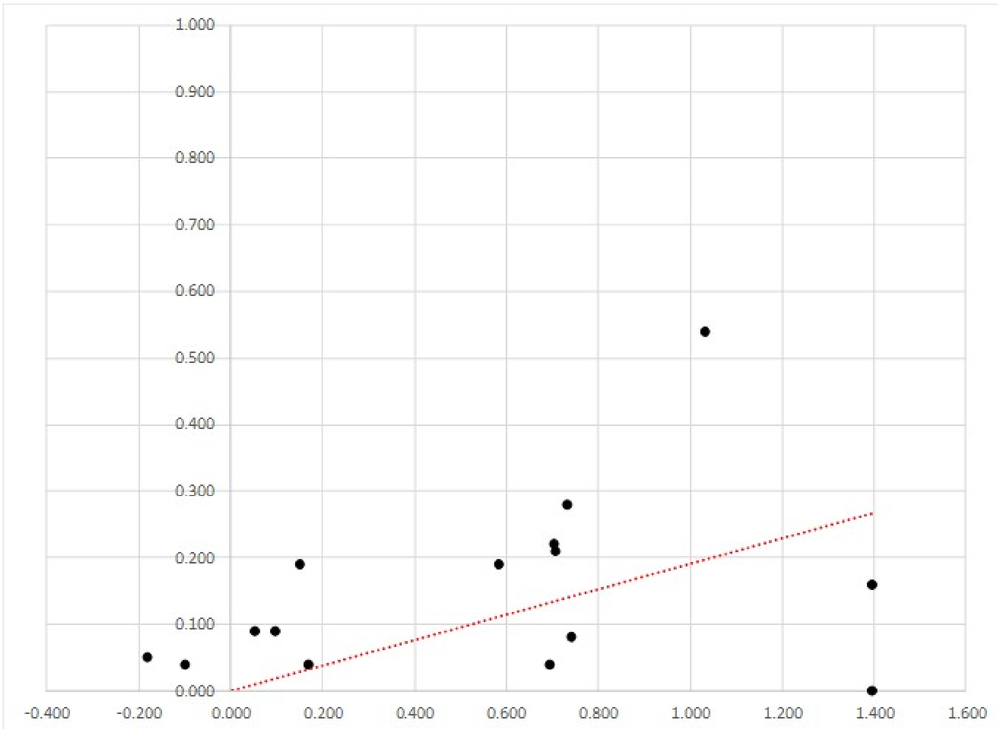

Figure 20: Actual (vertical axis) vs average expected SOFR end of month spike, least squares trend(red)

There is little reason to expect that futures market for SOFR is precisely pricing the magnitude of the spike well ahead of time. Not only are both the SOFR index and its futures market new, the spike is result of a complicated supply and demand imbalance motivated by regulatory measures at the end of the month. The magnitude of the spike may not be possible to predict ahead of time. The predicted spike does tends to converge to the actual spike over the course of the month, as evidenced in Figure 21 which compares the expected spike based on the last day of the month futures closing price. The convergence is consistent with the notion that as the end of the month approaches the conditions that cause the spike become clearer to market participants. On the very last day of trading in the month, since the SOFR rate is based on repo activity during the day, it is entirely possible that the futures market participants have better information regarding the state of the underlying repo market as evidenced by the futures price.

Another factor which explains the convergence behaviour is misspecification of the model parameters, which could lead to misrepresentation of the spike expectations. In the OU model described by (4), the average expected spread is a function of the parameters θ and μ as well as the expected spike and current spread $q(t)$. The results presented so far were produced by calculating the expected spike given constant θ and μ. Empirically the spread is not constant but rather trending up throughout the test period. It is safe to assume

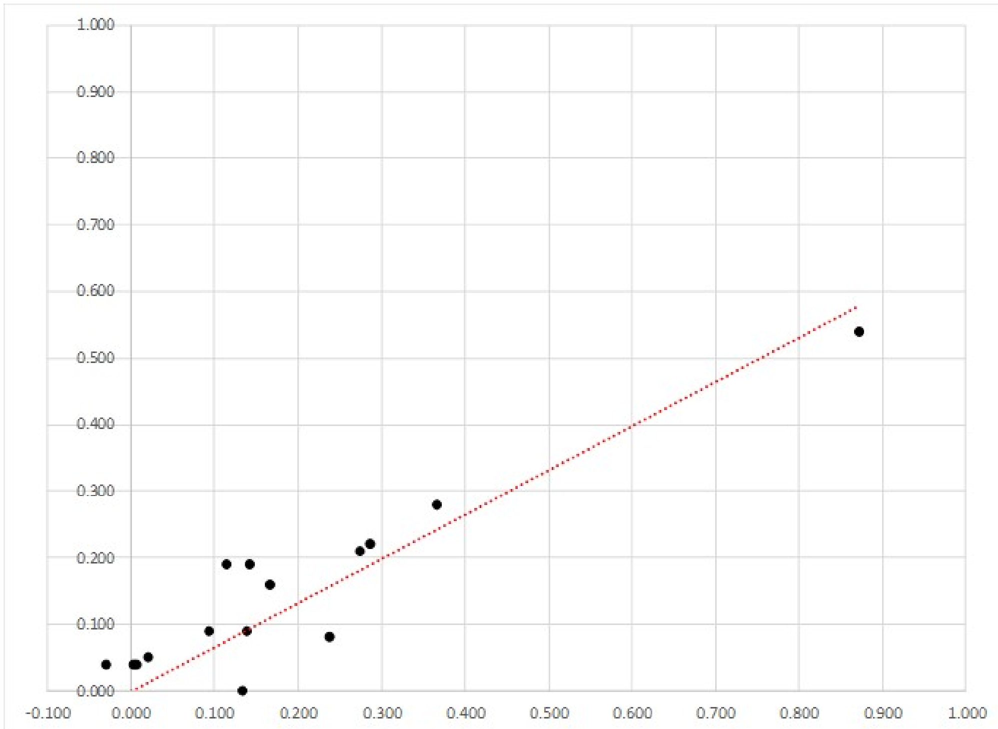

Figure 21: Actual (vertical axis) vs expected SOFR end of month spike implied from the last day futures closing price, least squares trend(red)

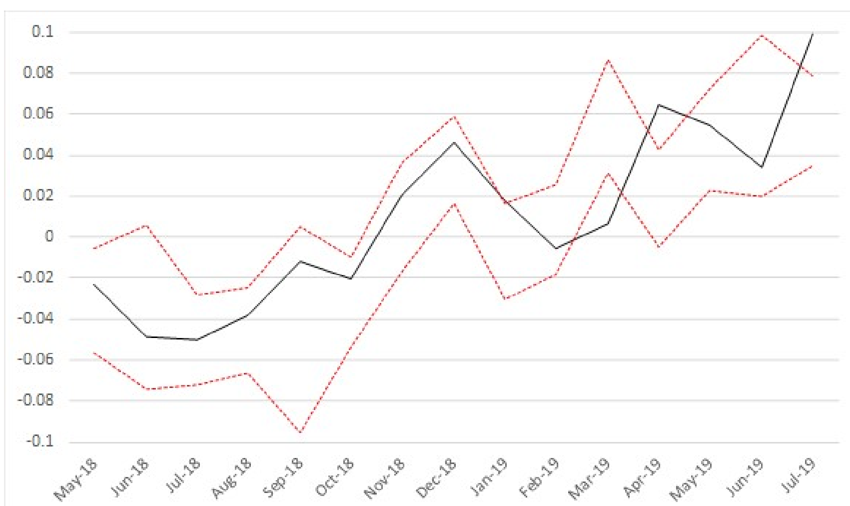

Figure 22: Monthly average realised spread (black) vs implied μ corresponding to a fixed spike of 0 and 0.5%

that the market adapts its expectation of the spread in response to changing realisations. Therefore keeping constant parameters, particularly the mean reversion target μ, results in a miscalculation of the expected spread, which ends up being compensated by expected spike calculation. This suggest that the mean reversion target should also be calibrated, however the 1m futures do not contain enough information to disentangle the expected spike from the expected spread and therefore it may be useful to simultaneously calibrate to 3m futures.

It is still possible to explore the evidence that the market is reflecting both a spread and spike in futures prices without using 3m futures by considering an economically reasonable domain for the spike estimate. Instead of fixing the μ parameter, we obtain two implied μ parameters for each calibration by fixing the expected spike to zero and 50 basis points, which we interpret as economically reasonable bounds based on the empirical behaviour of the SOFR rate. The resulting implied μ values are compared to realised spreads averaged for each month in the test period in Figure 22. Inspection of the results suggests two key conclusions. Firstly the implied μ corresponding to the zero spike is mostly above the realised spread, suggesting the possible existence of a spike premium embedded in the SOFR futures prices. Secondly, the trend of the implied μ parameters is consistent with the realised spread. This suggests that using a μ parameter which closely resembles the realised spread should yield spike expectations within economically reasonable bounds.

The implied μ parameter for a zero spike in comparison to realised spread is further examined for each calibration day in Figure 23. The spread implied on the last day of each month demonstrates market expectations of the spike on the last day of each month. The implied spread closely follows the pattern and magnitude of the realised spikes. An interesting implication is that on the last day there is awareness of the state of the repo market underlying the SOFR fixing in the futures market. It also provides clear evidence for the existence of spike expectation embedded in futures prices. It is reasonable to assume that if spike expectations are present in the futures prices for the last day of the month, they are also present for the lifetime of the futures contract.

4.3.7 SOFR cross sectional calibration

4.3.8 SOFR to Fed Funds futures spread

5 Conclusion

References

Alfeus, M., Grasselli, M. and Schlögl, E.: 2018, A consistent stochastic model of the term structure of interest rates for multiple tenors, *Technical report*, SSRN Working Paper.

Binder, A. S.: 2010, Quantitative easing: Entrance and exit strategies, *Federal Reserve Bank of St. Louis Review* **92**(6), 465–79.

Brace, A., Gatarek, D. and Musiela, M.: 1997, The market model of interest rate dynamics, *Mathematical Finance* **7**(2), 127–155.

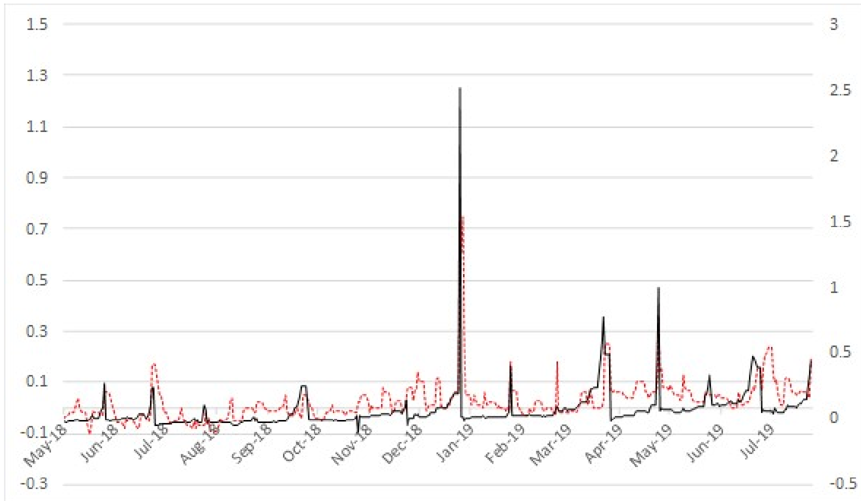

Figure 23: Realised spread (red) vs implied spread assuming zero spike (black)

Federal Open Market Committee, April: 2014, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, August: 2006, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, December: 2006, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, December: 2008, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, December: 2017, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, January: 2008, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, July: 2014, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, June: 2018, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, March: 2007, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, November: 2018, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, October: 2008, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, September: 2006, Minutes of the Federal Open Market Committee.

Federal Open Market Committee, September: 2014, Minutes of the Federal Open Market Committee.

Federal Reserve Bank of New York: 2015, Technical note concerning the methodology for calculating the effective federal funds rate, *Technical report*, Federal Reserve Bank of New York.

Federal Reserve Board: 2011, Conducting monetary policy.

Franco, J. C. G.: 2003, Maximum likelihood estimation of mean reverting processes, *Working paper* .

Hartely, J.: 2017, How european regulators are hindering the fed's ability to raise interest rates, *Forbes Magazine* .

Heath, D., Jarrow, R. and Morton, A.: 1992, Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation, *Econometrica* **60**(1), 77–105.

Hilton, S.: 2005, Trends in Federal Funds Rate Volatility, Federal Reserve Bank of New York, *Current Issues in Economics and Finance* .

Hull, J. and White, A.: 1990, Pricing interest–rate derivative securities, *Review of Financial Studies* **3**(4), 573–592.

Hunt, P. J. and Kennedy, J. E.: 2004, *Financial derivatives in theory and practice, 2nd ed.*, Wiley & Sons.

Miltersen, K. R., Sandmann, K. and Sondermann, D.: 1997, Closed form solutions for term structure derivatives with log-normal interest rates, *The Journal of Finance* **52**(1), 409–430.

Musiela, M. and Rutkowski, M.: 1997, Continuous–time term structure models: A forward measure approach, *Finance and Stochastics* **1**(4).

Schrimpf, A. and Sushko, V.: 2019, Beyond libor: a primer on the new reference rates, *Technical report*, BIS Shrapnel.

The Alternative Reference Rates Committee: 2018, Second report.