Quantification of feedback effects in FX options markets

Benjamin Anderegg, Didier Sornette, Florian Ulmann

SNB Working Papers
3/2019

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK

Legal Issues

DISCLAIMER

The views expressed in this paper are those of the author(s) and do not necessarily represent those of the Swiss National Bank. Working Papers describe research in progress. Their aim is to elicit comments and to further debate.

COPYRIGHT©

The Swiss National Bank (SNB) respects all third-party rights, in particular rights relating to works protected by copyright (information or data, wordings and depictions, to the extent that these are of an individual character).

SNB publications containing a reference to a copyright (© Swiss National Bank/SNB, Zurich/year, or similar) may, under copyright law, only be used (reproduced, used via the internet, etc.) for non-commercial purposes and provided that the source is mentioned. Their use for commercial purposes is only permitted with the prior express consent of the SNB.

General information and data published without reference to a copyright may be used without mentioning the source. To the extent that the information and data clearly derive from outside sources, the users of such information and data are obliged to respect any existing copyrights and to obtain the right of use from the relevant outside source themselves.

LIMITATION OF LIABILITY

The SNB accepts no responsibility for any information it provides. Under no circumstances will it accept any liability for losses or damage which may result from the use of such information. This limitation of liability applies, in particular, to the topicality, accuracy, validity and availability of the information.

ISSN 1660-7716 (printed version)
ISSN 1660-7724 (online version)

© 2019 by Swiss National Bank, Börsenstrasse 15,
P.O. Box, CH-8022 Zurich

Quantification of feedback effects in FX options markets *

Benjamin Anderegg[a], Didier Sornette[b,c], Florian Ulmann[a,b]

[a] *Swiss National Bank, Boersenstrasse 15, 8001 Zuerich*

[b] *ETH Zurich, Department of Management, Technology, and Economics, Scheuchzerstrasse 7, 8092 Zurich, Switzerland*

[c] *Swiss Finance Institute, c/o University of Geneva, 40 blvd. Du Pont d'Arve, 1211 Geneva 4, Switzerland*

Abstract

We model the feedback effect of delta hedging for the spot market volatility of the forex market (dollar-yen and dollar-euro) using an economy of two types of traders, an option market maker (OMM) and an option market taker (OMT), whose exposures reflect the total outstanding positions of all option traders in the market. A different hedge ratio of the OMM and OMT leads to a net delta hedge activity that introduces market friction and feedback effects. This friction is represented by a simple linear permanent impact model for the net delta hedge volumes that are executed in the spot market. This approach allows us to derive the dependence of the spot market volatility on the gamma exposure of the trader that hedges a larger share of her delta exposure and on the market impact of the delta hedge transactions. We reconstruct the aggregated OMM's gamma exposure by using publicly available DTCC trade repository data and find that it is negative, as expected: the OMT usually buys options with either a view on the spot price or with the desire to hedge other positions and, thus, is net long on options. As the OMM provides liquidity as a service to the market, their position is reversed compared with the OMT. Our regressions show a high goodness of fit, a highly significant parameter for the gamma exposure of the OMM and, as expected, that the volatility is increased by the OMM's short gamma exposure. Quantitatively, a negative gamma exposure of the OMM of approximately -1000 billion USD (which is around what we observe from our reconstructed OMM data) leads to an absolute increase in volatility of 0.7% in EURUSD and 0.9% in USDJPY. If we assume that the hedge ratios in the two markets are the same, the difference can be directly explained by the higher market impact of a transaction in the USDJPY spot market compared to the EURUSD spot market, as the liquidity of the EURUSD spot market is higher than that of the USDJPY spot market. Our results are in line with and empirically confirm previous theoretical work on the feedback effect of delta hedging strategies on spot market volatility.

Keywords: , DTCC, FX options, delta hedging, price impact, trade repository data, volatility modeling
JEL: G12, G15, G21, G23, G32, C32, C51, C55, C58, C55, C80

*We are grateful to all of the members of the Foreign Exchange and Gold OU of the SNB for their helpful comments and discussions on this topic. The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the authors. They do not necessarily reflect the views of the Swiss National Bank. The Swiss National Bank takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in this paper.

Email addresses: `Benjamin.Anderegg@snb.ch` (Benjamin Anderegg), `dsornette@ethz.ch` (Didier Sornette), `Florian.Ulmann@snb.ch` (Florian Ulmann)

1. Introduction

This paper focuses on the empirical quantification of the effect of delta hedging strategies on spot market volatility. By accounting for a linear, permanent market impact of spot market trades, we show that, in theory, spot market volatility is influenced by the net delta hedge demand arising from the options market. This theory is subsequently empirically validated with the use of trade repository data.

Option pricing theory has been studied extensively since the seminal works of Black and Scholes (1973) and Merton (1973). Their theory is fundamentally based on an explicit trading strategy for the underlying asset and riskless bonds whose final payoff is equal to the payoff of the derivative security at maturity. As this trading strategy provides an insurance against the risk that comes with buying and selling the derivative, it is called a dynamic hedging strategy (dynamic, since the replication should ideally be continuous) Sircar and Papanicolaou (1998).

A crucial assumption of the Black-Scholes model is a frictionless and elastic market, i.e., a market in which trades have no price impact, neither temporary nor permanent Gueant and Pu (2015). However, a vast body of work concerned with market microstructure clearly shows that one can reject the assumption of a frictionless market Farmer et al. (2006). In the context of dynamic hedging strategies, this implies a feedback effect on market dynamic parameters such as the spot price and volatility structure: the dynamic hedging strategy requires the investor who sells (buys) an option to sell (buy) the asset in the spot market when its price declines and to buy (sell) it when its price goes up. From the nature of such a trading pattern, one can expect an increase (decrease) in spot market volatility Frey and Stremme (1997) Sircar and Papanicolaou (1998). Indeed, spot market volatility has been reported to have changed due to the introduction of futures and options Miller (1997) MacKenzie (2008). Note that this feedback effect can even result in two extreme price dynamics, market crashes and pinning. For example, Almgren (2015) attributes the Oct. 15, 2014, US Treasury price action, in which the US Treasury market underwent the largest intraday move since 2009, to the hedging of short option positions held by the dealer community. Avellaneda and Lipkin (2003), on the other hand, discuss the case where hedge funds, which have a long option position, pin the spot price to the strike at expiration due to their hedging behavior. Furthermore, since changes in market dynamic parameters (e.g., spot market volatility) also change the derivative's price and therefore its dynamic hedging strategy, market dynamic parameters cannot be regarded as exogenous. In this context, Lions and Lasry (2006) cite practitioners who observe that the future values of the parameters they use for the pricing of options might be modified by their own activities (trading/hedging).

Different work streams have emerged from the discussion of market frictions in the context of dynamic hedging strategies of options. Some are more focused on accounting for market frictions in the pricing of options (and in deriving an optimal hedging strategy), while others are more concerned with describing the actual feedback effect that dynamic hedging has on market dynamics. Our paper contributes to the latter stream, so we will mainly focus on the effect of delta hedging on the observed spot market volatility. Nevertheless, we will present a brief literature review for both streams.

Accounting for feedback effects in the Black-Scholes pricing model has been the objective of a large body of academic literature. First, Brennan and Schwartz (1989), Sircar and Papanicolaou (1998) and Frey and Stremme (1997) focus on the derivation of the price dynamics through a clearing condition. Without taking into account the liquidity costs explicitly, they work with single-period models in which the supply and demand curves arise from two types of investors, one that upholds the Black-Scholes hypotheses and one that trades to insure other portfolios, which leads to-modified price dynamics. Sircar and Papanicolaou (1998) even introduce an option pricing model that accounts self-consistently for the effect of dynamic hedging on the volatility of the underlying asset. They find that spot market volatility increases by approximately 10%-18% for traders who write options and hedge their delta in the spot market.

The other approach is to account for the price impact of hedging trades in the spot market directly. In this approach, the price impact is often decomposed into temporary and permanent price impact components Almgren and Chriss (2000). Temporary price impact considerations are mostly relevant in the context of optimal hedging strategies, as presented in Rogers and Singh (2010) and Naujokat and Westray (2011). The basic problem is that continuous hedging is prohibitive due to execution costs stemming from temporary market impact, whereas low-frequency hedging leads to large tracking error of the frictionless Black-Scholes delta hedge Gueant and Pu (2015). Nevertheless, only models that consider a permanent market impact find a change in spot market dynamics (e.g., volatility) Rogers and Singh (2010). Thus, the discussion of feedback effects is typically modeled with the assumption of permanent market impact.

Among others, Abergel and Loeper (2016), Loeper (2013) and Bouchard et al. (2016) consider such a financial market with a linear, permanent price impact, and all three analyses conclude that the market dynamics of the spot market are changed by dynamic hedging of derivatives. The consideration of a permanent market impact is dedicated to the hedging of derivatives in situations where the notional of the option hedged is such that the delta-hedging is nonnegligible compared to the average daily volume traded on the spot market Bouchard et al. (2017). Note that this large position may be the position of a single large trader or may be the aggregate position of a collection of traders, for example, the entire sell-side community, who all hold similar positions and who all hedge while their counterparties do not Li and Almgren (2016). Lions and Lasry (2006) also write about the impact of trading and hedging on the dynamics of an asset, particularly its volatility. They also consider a linear, permanent impact but no temporary impact. In their framework, the trader's position is always perfectly hedged, but there is an observable effect on the realized volatility. More specifically, they consider a situation where a single large investor influences the price dynamics by her trading activity based on the assumption of a linear elastic law (as in e.g., Kyle (1985)). Following a utility maximization approach, they prove that the optimal hedging strategy impacts upon the volatility of the asset through the gamma of the option. Li and Almgren (2016) take both temporary (with quadratic execution costs) and permanent price impacts into account (following Almgren and Chriss (2000)) and find that permanent market impact causes an increase or decrease in the realized volatility of the market price depending on whether the large investor or the entire community of hedging traders is net long or short options.

While much work has been done on the theoretical side, empirical work on feedback effects arising from dynamic hedging has been limited. This gap is probably because an empirical study of such feedback effects requires proprietary data, e.g., data that allow for the identification of the hedging demand. Our paper bypasses the need for proprietary data by reconstructing it (see below), thus providing a procedure to close this gap. Based on these data, we will focus on the theoretical and empirical discussion of the feedback effect stemming from delta hedging strategies in the FX spot market.

First, similar to the approach presented in Frey and Stremme (1997), we work in a model economy with only two agents. Our agents are represented by two types of traders, namely, an aggregated option market maker (OMM), who provides liquidity to the options market, and an aggregated option market taker (OMT), who consumes liquidity. As a consequence of this setup, the OMM and the OMT always have exact opposite positions from one another, and each of their books represent the total outstanding position on the options market.

We then assume that the two traders' hedging trades in the spot market have a linear, permanent market impact and no temporary market impact. This framework results in a simple volatility model that shows that the spot market volatility increases above its fundamental value (which is the volatility in case of no market impact of delta hedging trades) if the hedge ratio of the trader, who is net short options (and therefore has a negative gamma exposure) is larger than the hedge ratio of the trader who is net long options. Note that this model economy is motivated by the idea that typical OMMs, e.g., large banks, dynamically hedge their positions while OMTs, e.g., investors, do not. This would result in an asymmetric net delta hedge demand, which will lead to an increase (if the OMM is short options) or decrease (if the OMM is long options) in the spot market volatility. Thus, our model framework and our theoretical results are similar to those presented by Lions and Lasry (2006) and Li and Almgren (2016). Nevertheless, nowhere in the paper do we a priori assume that the OMM hedges a larger share of her exposure compared to the OMT.

Second, we want to validate this model empirically for the FX market. However, to validate our model, we first need to have the data of the aggregated OMM and OMT. The BIS triennial survey 2016 states that the FX market turnover is composed of 33% spot transactions and 67% derivatives such as FX swaps, outright forwards and options BIS (2016). Since 2016, those 67% are subject to mandatory reporting in the US, for which our main data source in this analysis, the Depository Trust and Clearing Corporation (DTCC) trade repository, is an authorized trade repository covering approximately 20% of the worldwide options market transactions Shtauber and Marone (2014). These transaction data allow for profound insights on the market positioning in terms of new and outstanding positions. DTCC positioning data have been used previously to analyze spillover effects from the FX options market on the FX spot market but mostly with a focus on the spillover effect on the spot price rather than its volatility Weng and Grover (2017) Shtauber and Marone (2014) Winkler (2017). The downside of the DTCC data is that we do not know the side behind each trade, and thus we do not know whether the transaction is buy or sell initiated. However, this information is necessary to reconstruct the aggregated exposure of the OMT and OMM, as we assume that the OMM is always the liquidity provider and the OMT is always the liquidity taker. We solve this problem by using a simple trade classification algorithm similar to the Lee-Ready algorithm to classify the aggressor side behind each option trade stored in the DTCC repository. [2] Thus, we can identify the OMM (liquidity provider) and OMT (liquidity taker) behind each trade and can therefore also calculate their aggregated position at each point in time. This allows us to estimate our volatility model. [3]

The remainder of this paper is structured as follows. The second section introduces our volatility model. The third section briefly explains how we reconstruct the OMM's position behind each transaction of the DTCC trade repository data. In the fourth section, we show our empirical results, supporting the validation of our volatility model. In the fifth section, we discuss our results and draw conclusions from our findings.

2. A model for the feedback effects of delta hedging

2.1. Derivation of the volatility model

In this section, we derive an expression for the spot market volatility that includes the price impact of the delta hedging activity. The observable exchange rate is denoted by S_t and is given by the following general

[2] The equity literature offers a wide range of work on how to classify the aggressor side behind transactions, e.g., Lee and Ready (1991) Chakrabarty et al. (2015) Easley et al. (2016).

[3] Note that we only examine vanilla options, as the details of exotic options in the DTCC trade repository are not sufficient.

stochastic differential equation.

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW \tag{2.1}$$

We now first assume that the nonobservable fundamental value of the exchange rate F_t, which has the interpretation of the market price unaffected by delta hedging, follows a geometric Brownian motion.

$$dF_t = \alpha F_t dt + \nu F_t dW \tag{2.2}$$

Applying Ito's lemma to the log-price changes $df_t = d(ln(F_t))$ results in $df_t = \frac{dF_t}{F_t} - \frac{1}{2}\nu^2 dt$, with which we can express the corresponding stochastic differential equations for the log-price.

$$df_t = \underbrace{(\alpha - \frac{1}{2}\nu^2)}_{\alpha'} dt + \nu dW = \alpha' dt + \nu dW \tag{2.3}$$

Now we assume that there are only two market participants, the OMM and the OMT. Their net delta exposures are $N^{OMM}\Delta^{OMM}$ and $N^{OMT}\Delta^{OMT}$, respectively, where N is the notional (always positive) and Δ is the option delta (the delta is positive for a long call and a short put and negative for a short call and a long put). Both hedge their overall delta exposure to a certain degree in the spot market, given by the hedge ratios h^{OMM} and h^{OMT}. If the price changes by a certain percentage between time t and $t+1$, the net traded delta hedge volume in the spot market at time $t+1$ is given by the negative change of the delta exposure between t and $t+1$, multiplied by the hedge ratio, so $V^{OMM}_{t+1} = h^{OMM}N^{OMM}(-(\Delta^{OMM}_{t+1} - \Delta^{OMM}_t))$ for the OMM and $V^{OMT}_{t+1} = h^{OMT}N^{OMT}(-(\Delta^{OMT}_{t+1} - \Delta^{OMT}_t))$ for the OMT. Note that in a market with two participants, the notional is the same for both participants ($N^{OMM} = N^{OMT} = N$), and the delta of the OMM is always equal to the negative delta of the OMT ($N\Delta^{OMM} = -N\Delta^{OMT}$). Furthermore, the positions of each trader represent the total outstanding position in the market (which is why our two agents are the OMM and the OMT, as by definition the aggregated positions from all OMMs sum up to the total outstanding positions in the market). Defining the asymmetry in the hedge ratio of the OMM and the OMT as $h = h^{OMM} - h^{OMT}$, the total net traded delta hedge volume between t and $t+1$ is given by

$$
\begin{aligned}
V_{t+1} &= V^{OMM}_{t+1} + V^{OMT}_{t+1} && \text{(2.4a)}\\
&= h^{OMM}N^{OMM}(-(\Delta^{OMM}_{t+1} - \Delta^{OMM}_t)) + h^{OMT}N^{OMT}(-(\Delta^{OMT}_{t+1} - \Delta^{OMT}_t)) && \text{(2.4b)}\\
&= (h^{OMM} - h^{OMT})N(-(\Delta^{OMM}_{t+1} - \Delta^{OMM}_t)) && \text{(2.4c)}\\
&= hN(-(\Delta^{OMM}_{t+1} - \Delta^{OMM}_t)) && \text{(2.4d)}
\end{aligned}
$$

In the continuous limit, equation 2.4 has the following form:

$$dV_t = hNd\Delta^{OMM}_t \tag{2.5}$$

Now we assume that the traded delta hedge volume impacts the market price in a linear and permanent fashion, with β defined as the log-price change (market impact) per traded volume. We assume that β is a constant. This assumption basically reflects Kyle's price impact model (the price impact is permanent, linear and depends on the volume of the transaction only), presented in Kyle (1985). Note that we neglect any form of temporary price impact, as we validate our model on large time scales only. Thus, the assumption that the permanent price impact scales linearly with the transaction volume is in line with current market impact theory Bouchaud (2010). [4]

Consequently, the permanent impact depends on the total volume executed, dV_t:

$$ds_t = df_t + \beta dV_t \tag{2.6a}$$

$$= df_t - \beta hNd\Delta_t^{OMM} \tag{2.6b}$$

This equation can be expressed in terms of dS_t (again, by using $ds_t = \frac{dS_t}{S_t} - \frac{1}{2}\sigma_t^2 dt$):

$$dS_t = S_t df_t + \frac{1}{2}\sigma_t^2 S_t dt - \beta hNS_t d\Delta_t^{OMM} \tag{2.7}$$

Applying Ito's lemma to $d\Delta_t^{OMM}$ leads to

$$d\Delta_t^{OMM} = \underbrace{(\underbrace{\frac{\partial\Delta_t^{OMM}}{\partial t}}_{-Charm_t^{OMM}} + \mu_t S_t \underbrace{\frac{\partial\Delta_t^{OMM}}{\partial S}}_{\gamma_t^{OMM}} + \frac{1}{2}\sigma_t^2 S_t^2 \underbrace{\frac{\partial^2\Delta_t^H}{\partial S^2}}_{Speed^{OMM}})}_{\gamma_t'} dt + \sigma_t S_t \underbrace{\frac{\partial\Delta_t^{OMM}}{\partial S}}_{\gamma_t^{OMM}} dW_t \tag{2.8}$$

The option greeks are labeled with brackets. γ_t^{OMM}, which is important in our theory, is the option gamma of OMM at time t. Inserting 2.8 into 2.7 results in

$$dS_t = (\alpha' S_t + \frac{1}{2}\sigma_t^2 S_t - \beta hNS_t\gamma_t')dt + (\nu S_t - \beta hNS_t^2\sigma_t\gamma_t^{OMM})dW_t \tag{2.9}$$

Comparing equation 2.9 with 2.1 results in

$$\sigma_t = \frac{\nu}{(1 + \beta hNS_t\gamma_t^{OMM})} \tag{2.10}$$

2.2. Discussion of the volatility model

In this subsection, we discuss equation 2.10 in more detail. We see that the spot market volatility of a market unaffected by delta hedging, ν, is rescaled by a factor $(1 + \beta hNS_t\gamma_t^{OMM})^{-1}$. Note that it is not the OMM's gamma exposure that determines whether the volatility is scaled up or down but rather the gamma exposure of the trader who hedges a larger part of his delta exposure compared to the other trader, which is given by $hNS_t\gamma_t^{OMM}$. If both traders hedge the same share of their delta exposure, $h = 0$, and the volatility is at its fundamental value. However, if the OMM hedges a larger share of her delta exposure compared to the

[4] We can think of the temporary impact as connected to the liquidity cost faced by the agent but the permanent impact as linked to information transmitted to the market by the agents' trades Li and Almgren (2016).

OMT, $h > 0$, the gamma exposure of the OMM determines if the volatility is scaled up or down. Conversely, if $h < 0$, the gamma exposure of the OMT is relevant (since $-\gamma_t^{OMM} = \gamma_t^{OMT}$). These theoretical results are in line with the work of Lions and Lasry (2006) and Li and Almgren (2016).

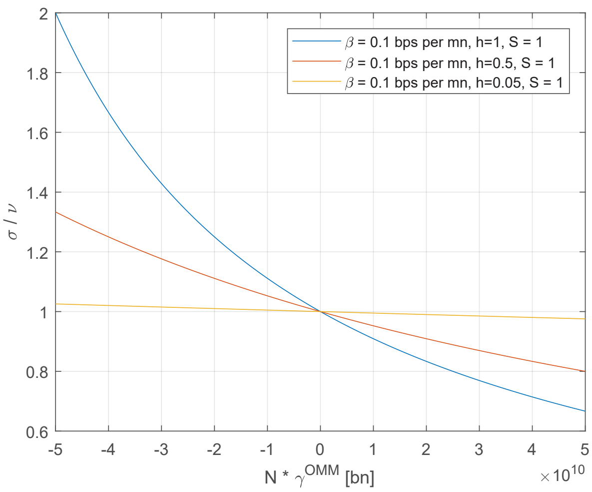

Figure 1: This figure shows the dependence of $\frac{\sigma_t}{\nu}$ as a function of $N\gamma_t^{OMM}$ as predicted by equation 2.10. On the vertical axis, we display the observable spot market volatility normalized by the volatility of a market unaffected by delta hedging, so $\frac{\sigma_t}{\nu}$. On the horizontal axis, we display the factor $\beta hNS\gamma^{OMM}$, where we fix the market impact to $\beta = 0.1[bps/mn\ base\ currency]$ (motivated by an empirical estimation of the history-dependent impact model of Bouchaud (2010)) and show three values of the hedge ratio. We see that the volatility is scaled up (down) by the negative (positive) gamma exposure of the trader who hedges a larger share of their delta exposure.

As we assume a constant market impact per unit traded volume, β, the equation mainly depends on the difference of the hedge ratio of the OMM to the OMT, h, and on the gamma exposure of the OMM, $SN\gamma_t^{OMM}$. If we assume that the OMM hedges a larger portion of her exposure than the OMT (so $h > 0$), we see that the volatility is scaled up if the OMM has a short gamma exposure and scaled down if the OMM has a long gamma exposure. This comes from the fact that a short gamma exposure of the OMM is associated with a short option position, which implies a delta hedging strategy in which the OMM has to buy the asset if the spot price increases and sell the asset if the spot price decreases, which leads to an increase in the spot market volatility given that the hedging trades have a price impact (vice versa for the long gamma exposure). The strength of this feedback effect thus depends on the difference of the hedge ratio of the OMM to the OMT (h), on the market impact of the traded volumes (β) and on the gamma exposure of the OMM. Figure 2 discusses our model for the case when the OMM holds a short call option and $h = 1$.

As the volatility cannot be below zero by definition, the factor $\beta hNS_t\gamma_t^{OMM}$ must be in $[-1, +\infty]$.

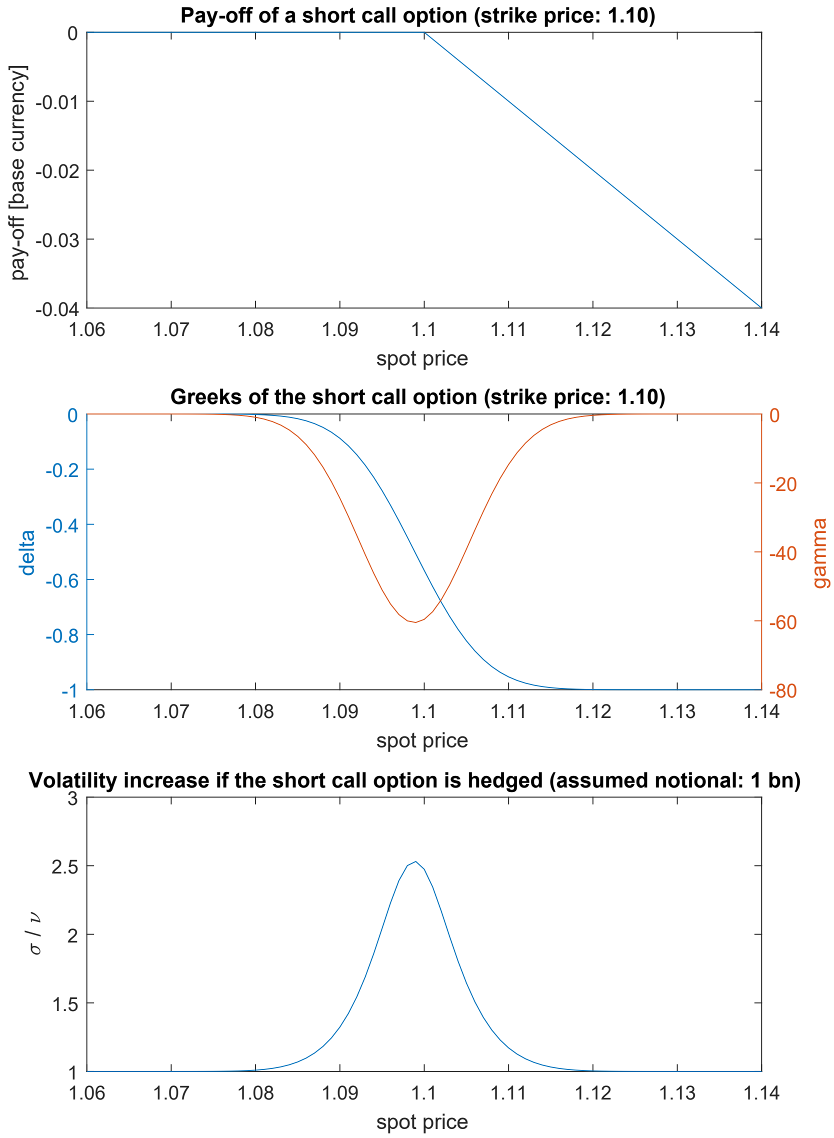

Figure 2: We assume that the OMM has a short call position, which she hedges and that the OMT behind the same position does not hedge her delta exposure. The upper figure shows the pay-off at maturity of this short call position with a strike price at $K = 1.10$. The middle figure displays the corresponding delta and gamma profile of the short call option with a volatility of 6% and a time-to-expiry of three months. We see that delta is a decreasing function in the spot price, and thus the corresponding delta hedge profile (which is the negative of the delta profile) is an increasing function in the spot price. This implies that the OMM needs to buy the underlying if the spot price goes up and sell the underlying if the spot price goes down. The gamma profile is negative. Now we assume that the notional behind the short call option is one billion and that the OMM hedges her total position in the spot market, while the OMT does not hedge her position at all (hedge ratio of $h = 1$). The resulting increase in the volatility ($\frac{\sigma}{\nu}$, based on equation 2.10 with a permanent market impact of $\beta = 0.1[bps/mn]$) is displayed in the lower panel. We see that the volatility locally increases at spot prices where the gamma profile of the OMM is negative.

3. Derivation of the OMM's positioning from DTCC data

Let us now empirically validate equation 2.10. To do so, we need data for the gamma exposure of the OMM ($NS_t\gamma_t^M$ in equation 2.10). We consequently focus on the FX options market.

The Dodd-Frank Act, which was signed into law in July 2010, requires the reporting of all over-the-counter FX option transactions that were traded in the US or involve US persons to a data repository. Note that there are several data repositories that market participants can report to. We use the Depository Trust and Clearing Corporations repository (DTCC), which covers approximately 20% of the worldwide options market transactions Shtauber and Marone (2014).

Option trades stored in the DTCC trade repository are the result of OTC transactions. Large banks mostly act as liquidity providers for their clients, e.g., hedge funds, corporations, institutional investors, asset managers or private clients, which then act as liquidity takers. The banks provide liquidity by constantly offering a bid and ask price (which is quoted in terms of implied volatility) to their clients, and the clients then take the liquidity by either buying on the ask price (this deal is called a paid) or selling on the bid price (this deal is called a given). The trade is then reported to the DTCC trade repository, but the information on whether the deal happened on the bid or the ask price is not included.

As presented before, our model economy assumes that there exist only two traders, the option market maker (OMM) and the option market taker (OMT). The OMM represents the aggregate of the liquidity providers (mostly banks) and the OMT represents the aggregate of the liquidity takers (hedge funds, corporations, institutional investors, asset managers and private clients). We thus attribute all deals to the aggregated position of the OMT, whereas the OMM then just has the exact opposite position. The motivation for this model economy lies in the idea that the OMM hedges a larger share of their delta exposure compared to the OMT. Nevertheless, this is not assumed a priori in this paper.

The main problem is that DTCC does not provide information on whether a deal was a paid or a given. We have to reconstruct this information in order to quantify the OMM and OMT's positioning in the options market. We thus basically follow a trade classification procedure. In the first step, the implied volatility of each option transaction stored in the DTCC trade repository is reconstructed by numerically minimizing the difference between the premium that was paid (which is reported to the DTCC trade repository) and the premium that Black's model suggests for the trade (based on the specifications that are reported to the DTCC trade repository). [5] Next, this reconstructed implied volatility, which was traded at a certain time that is also reported to the DTCC trade repository, is matched onto the simultaneous reference bid-ask price stream that we take from Bloomberg. We label the transaction a paid (buy order on the ask price) if the reconstructed implied volatility is closer to the ask price compared to the bid price or a given (sell order on the bid price) if the reconstructed implied volatility is closer to the bid price compared to the ask price. As the OMM always provides liquidity, it follows that she has a short position if the transaction is labeled a paid and a long position if the transaction is labeled a given. Note that this matching procedure works very well, as the spreads in the FX options market are typically still quite large.

This approach is an improvement on recent work presented by Weng and Grover (2017), Shtauber and Marone (2014) or Winkler (2017), who also worked with DTCC options data but assumed that the OMM is short every transaction reported to the repository. Our reconstructed data set, on the other hand, suggests

[5]For other asset classes, other option pricing models might be appropriate.

that the OMM is short 75% of all traded option contracts in EURUSD and 78% of all traded option contracts in USDJPY. Thus, to assume that the OMM is short every contract is an oversimplification.

As we now know the OMM's position behind each transaction, we can calculate the delta and gamma exposure of the OMM at any given point in time. We simply calculate the delta and gamma of the OMM's open positions at any given point in time, multiply it by the signed notional of the corresponding position (i.e., the negative notional if the OMM is short the option) and take the corresponding sum over all open positions. This results in the aggregated delta and gamma exposure of the OMM's portfolio at any given point in time ($NS_t\gamma_t^H$), which we now can use to validate our volatility model.

A more detailed discussion on the underlying DTCC data quality and on how the OMM's positioning is derived from the DTCC data can be found in the appendix.

4. Results

In this section, we present and discuss our three main contributions to the existing literature. First, we discuss our findings of the OMM's delta and gamma exposure that we derived from DTCC data. Second, we empirically discuss the spillover effect of the OMM's delta hedging activity on the spot market volatility. Third, we validate the volatility model presented in the previous section (equation 2.10) with the reconstructed OMM gamma exposure.

4.1. Discussion of the OMM's exposure reconstructed from the DTCC data

The analysis presented here focuses on the FX market, namely, the FX pairs EURUSD and USDJPY, as for those two FX pairs the outstanding volume in the FX options market is large enough to expect significant spillover effects on the FX spot market. The DTCC trade repository data, based on which we estimate the OMM's delta and gamma exposure, are a reliable and representative source that cover approximately 20% of the total options market turnover (a more detailed discussion of the DTCC data quality can be found in appendix Appendix A.1). As described in section 3, we reconstructed the side of the OMM behind each transaction from the DTCC data and subsequently calculated the delta and gamma exposure of the OMM at all points in time. The corresponding results are displayed in figure 3, where each trading day is split into 12 two-hour intervals. [6]

[6] UTC 00:00-02:00, 02:00-04:00, 04:00-06:00, 06:00-08:00, 08:00-10:00, 10:00-12:00, 12:00-14:00, 14:00-16:00, 16:00-18:00, 18:00-20:00, 20:00-22:00, 22:00-24:00

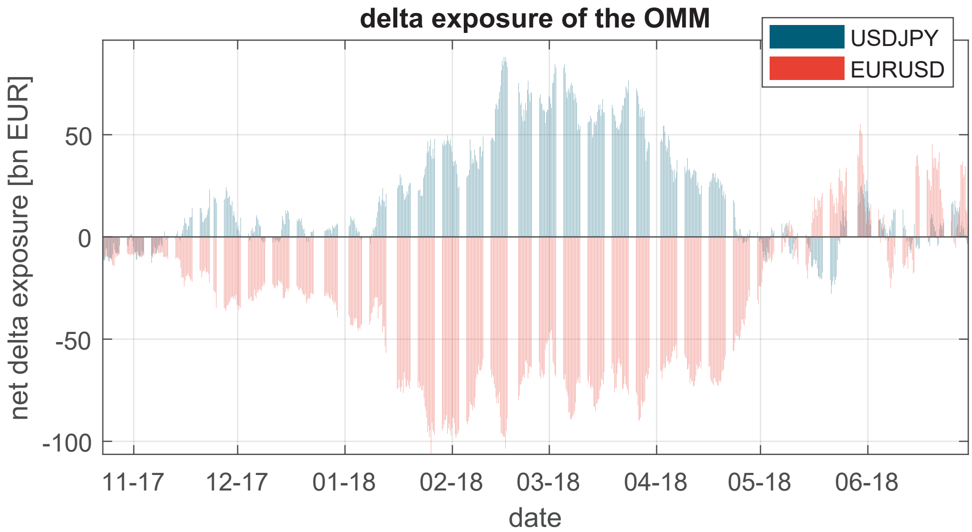

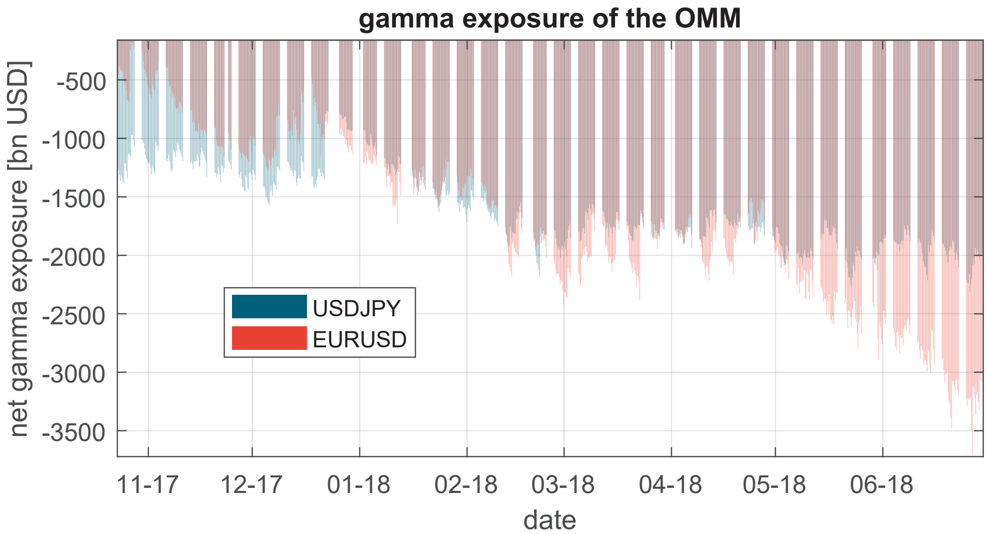

Figure 3: The figure shows the evolution of the accumulated net delta and gamma exposure of the OMM over the analyzed time frame for EURUSD and USDJPY.

As we see in figure 3, the OMM is short gamma at all times. This implies that the OMM is shorting put and call options (on a net basis). This is in line with our expectation. First, most investors are purchasers of options rather than sellers. They use options as a hedge or with consideration of the spot price and, therefore, mostly purchase options. Furthermore, they refrain from going short an option contract (put or call), as the risk involved in doing so is too high. Those traders most certainly trade as market takers in the market, e.g., as clients of banks. On the other side of each transaction of a market taker sits a market maker, e.g., the banks; consequently, the market maker is expected to have a short option position and therefore a short gamma position. Second, a short position in an option implies a short vega position, and as the implied volatility usually trades with a premium to the realized volatility, this is one way for option market makers to make money. For those two reasons, the OMM is expected to have a net short gamma position, which is exactly what we observe. As stated above, we see from our data that the OMM is short 75% of all traded option contracts in EURUSD and 78% of all traded options in USDJPY.

For the delta exposure, we have no a priori expectation of how it should look, since the delta exposure is not directly related to the positioning (if the OMM is long or short options) but rather is a function of where the spot price trades relative to the notional weighted strike price of the aggregated OMM's positions.

If the OMM would hedge her total position in the spot market, her delta hedge strategy would be to trade the negative difference of her delta exposure between two points in time in the spot market. The OMM would then be delta neutral by the end of each time period. Note that, as gamma is the change in delta per unit change in the spot price (derivative of the delta with respect to the spot price), gamma indicates how much delta would have to be hedged if the spot price moved one unit. In that sense, gamma is a proxy for the intensity of the delta hedge activity.

In the next section, we empirically discuss the feedback effect on the spot market volatility independent of our model framework. We subsequently validate our volatility model.

4.2. Empirical discussion of the feedback effect

As outlined in the previous sections, the gamma exposure of the trader who hedges a larger share of her delta exposure, or rather her corresponding delta hedging strategy, is assumed to generate feedback effects: a positive gamma exposure implies that she sells the underlying asset in a rising market to hedge her delta exposure, and as the executed hedging trades in the spot market have a market impact, this activity should have a negative effect on the realized spot market volatility. Following the same line of argument, a negative gamma exposure should have a positive effect on the realized volatility. Our volatility model (equation 2.10) reflects this logic. In this section, we will estimate this effect with the data on the gamma exposure of the OMM that we reconstructed from the DTCC data.

4.2.1. Estimation with daily data

To this end, we first estimate the linear regression model presented in equation 4.1, with α_1 and α_2 as model parameters and ϵ_t as the error term.

$$\sigma_t[\%] = \alpha_1[\%] + \alpha_2[\frac{\%}{\text{bn}}] \cdot NS_t \gamma_t^M[\text{bn}] + \epsilon_t \tag{4.1}$$

σ_t is the realized spot market volatility (annualized) for the time period between t and $t + 1$, and $NS_t\gamma_t^M$ is the OMM's gamma exposure at time t. As the volatility strongly depends on the intraday liquidity cycle, we estimate this linear regression model with daily data. [7]

	EURUSD	USDJPY
α_1 (%)	7.0^{***}	4.4^{***}
	(15.8)	(6.5)
α_2 ($\frac{\%}{\text{bn}}$)	-0.0010^{***}	-0.0014^{***}
	(-4.0)	(-3.2)
R-squared	0.09	0.06
F-test vs. const. model	15.8^{***}	10.5^{***}

Table 1: This table shows the regression results of the linear regression model presented in equation 4.1 for EURUSD and USDJPY. (t-statistics are presented in brackets. The stars *** denote significance at the 1% level.)

Table 4.3 displays the results of the estimation of equation 4.1. We see that the goodness of fit is high for both FX pairs and the parameters are highly significant. The value of α_2 is negative and has almost the same amplitude for the two currency pairs. This suggests that the OMM hedges a larger share of her delta exposure than the OMT and thus increases the spot market volatility by approximately 0.001% per billion of the base currency due to her negative gamma exposure (and the corresponding delta hedge activity). Note that the constant α_1, which can be interpreted as the fundamental volatility, is higher for EURUSD than for USDJPY, reflecting that the volatility level is higher during the European trading session compared to the Asian trading session.

4.2.2. Estimation with intraday data

As we can reconstruct the gamma exposure of the OMM on an intraday frequency, we can also quantify the effect with this intraday data. One difficulty of this intraday estimation is that the fundamental volatility (the constant of the linear regression model) depends on the intraday time, as it is influenced by the intraday liquidity cycle (trades influence prices and therefore also the volatility). We therefore cannot just mix the data from different intraday times and estimate the linear regression model. Thus, we use a mixed-effect model, namely, a random intercept linear mixed-effect model, as presented in equation 4.2. As when we presented the gamma exposure, each trading day is split into 12 two-hour intervals. [8] We define a dummy variable for the corresponding intraday time stamp, I_m (where $m = \{1, ..., 12\}$), that indicates the corresponding intraday interval. I_m then is our grouping variable for the constant $\alpha_1'[\%]$, as the fundamental volatility depends on the intraday time stamp. This construct, the constant conditional on the grouping variable, $(\alpha_1'[\%]|I_m)$, is called the random effect, while the fixed effect is given by $\alpha_1'[\%]$ and $\alpha_2'[\frac{\%}{\text{bn}}]$. As the volatility can have extreme values on an intraday level, due to e.g., data releases, we remove outliers from

[7] For EURUSD, we look at the European trading session (UTC 09:00 to UTC 16:00), and for USDJPY, we look at the Asian trading session (UTC 2:00 to UTC 9:00). The OMM's gamma exposure before each time interval is regressed onto the volatility that is realized during the time interval. In equation 4.1, we regress the gamma position of the OMM in EURUSD at 09:00 UTC onto the realized spot volatility between 09:00 and 16:00.

[8] UTC 00:00-02:00, 02:00-04:00, 04:00-06:00, 06:00-08:00, 08:00-10:00, 10:00-12:00, 12:00-14:00, 14:00-16:00, 16:00-18:00, 18:00-20:00, 20:00-22:00, 22:00-24:00

the intraday data set. [9]

$$\sigma_{t,m}[\%] = \alpha_1^{'}[\%] + (\alpha_1^{'}[\%]|I_m) + \alpha_2^{'}[\frac{\%}{\text{bn}}] \cdot NS_{t,m}\gamma_{t,m}^{M}[\text{bn}] + \epsilon_{t,m} \tag{4.2}$$

Table 4.2.2 displays the results of the estimation of equation 4.2. Looking at the fixed effect, we see that the results confirm those obtained with the daily data. Note that now the fundamental volatility is nearly equal between EURUSD and USDJPY, as we account for the different intraday volatility regimes.

	EURUSD	USDJPY
$\alpha_1^{'}$ (%)	5.7***	6.2***
	(10.8)	(11.1)
$\alpha_2^{'}$ ($\frac{\%}{\text{bn}}$)	-0.0007***	-0.0008***
	(-8.6)	(-3.3)
R-squared	0.34	0.17
F-test vs. const. model	74.5***	10.8***

Table 2: This table shows the regression results of the random intercept linear mixed-effect model presented in equation 4.2 for EURUSD and USDJPY. (t-statistics are presented in brackets. The stars *** denote significance at the 1% level.)

4.2.3. What can we learn from the estimation of the models

First, we want to highlight that both estimations presented above, the one based on daily data and the one based on intraday data, reveal very similar parameter estimates for both EURUSD and USDJPY, and both estimations also show high significance and high explanatory power. This is strong evidence that the quantification of the feedback effect works well with our reconstructed OMM data set.

The quantification of the feedback effect is given by the parameter α_2 of the daily estimation (equation 4.1) or the parameter $\alpha_2^{'}$ of the intraday estimation (equation 4.2). As both estimations yield very similar estimates, we will subsequently discuss only α_2. The significance of α_2 tells us whether the gamma exposure of the OMM actually has relevance for the spot market volatility, and its value tells us what the corresponding linear effect is per billion gamma exposure of the OMM.

However, the interpretation of α_2 is not straightforward: if the OMM is short gamma, her delta hedging strategy implies buying the asset in a rising market and selling the asset in a falling market. Such a strategy increases the volatility if trades have a price impact. The OMM has the exact opposite option position of the OMT, so the delta hedging activity of the former has the exact opposite effect on the volatility. The volatility is therefore only influenced by delta hedging if one trader (OMM or OMT) hedges a larger share of her delta than the other trader (otherwise the net delta hedge activity would be zero). The sign of the gamma exposure of the trader that hedges a larger share of her delta thus determines whether the volatility is increased or decreased. Our empirical estimation shows that a negative gamma exposure of the OMM increases the volatility (as α_2 has a negative sign). From this we can conclude that the OMM hedges a larger share of her delta exposure than the OMT. Assuming the hedging ratios of the OMM and the OMT are stationary, the overall feedback effect stemming from the OMM and OMT's gamma exposure and from the corresponding delta hedge activity are then just given by the value of α_2. Looking at the value of α_2,

[9] We remove the highest and lowest 2.5 percentiles.

our results suggest that a short gamma exposure of the OMM of -1 billion in the base currency results in an increase of the realized spot market volatility of approximately 0.001% for both EURUSD and USDJPY.

This does not look like a large effect, but as we see from the data displayed in figure 3, the gamma exposure of the OMM is of order -10^3 billion in both EURUSD and USDJPY. Consequently, the resulting effect is actually an increase in the spot market volatility of approximately 1% ($-0.001[\frac{\%}{\text{bn}}] \cdot -10^3[bn] = 1[\%]$). [10] Again, this stems from the fact that a short gamma exposure of the OMM (who is found to hedge a larger share of her delta exposure than the OMT) is associated with a short option position, which implies a delta hedging strategy in which the OMM has to buy the base currency if the spot price increases and sell the base currency if the spot price decreases. This leads to an increase in the spot market volatility given that the hedging trades have a price impact (vice versa for the long gamma exposure). Apparently, the feedback effect is stronger (per billion gamma exposure of the OMM) in USDJPY than in EURUSD, which can be explained by the higher liquidity of the EURUSD spot market and thus the lower market share of the delta hedging strategy in EURUSD (and thus the lower impact of the traded volume) BIS (2016). Note that the feedback effect of the gamma exposure on the spot market volatility is (most probably) not linear, as we know from our volatility model. Therefore, the quantification of the feedback effect presented in this section is only a linear approximation. Also note that the true effect per billion is smaller when we adjust for the fact that our OMM positioning is only based on the DTCC data, which represent only approximately 20% of the total options market Weng and Grover (2017). If we therefore simply scale up the OMM positioning, the OMM gamma exposure would be higher by a factor of 5, and accordingly, the parameter α_2 would be lower by a factor of 5.

The constant α_1 is also found to be highly significant. α_1 can be interpreted as the spot market volatility in the absence of any delta hedging activity, the fundamental volatility. Note that there is a rather strong difference in the estimate of α_1 for the daily and intraday USDJPY estimate. The spot market volatility in the absence of any delta hedging activity is found to be lower for the daily estimate (4.4%) compared to the intraday estimate (6.2%). This difference stems from the fact that the volatility data for the daily estimate for USDJPY originate from the Asian trading hours, in which the volatility is usually lower than during the European or American trading hours. As the intraday estimate takes the intraday time explicitly into account (and therefore also the intraday volatility regime), we obtain a more meaningful estimate of α_1 for USDJPY in the intraday estimation.

We also tested other regression methods, such as Bayesian linear regression models, which all confirmed our results. However, for the sake of clarity and interpretability, we chose to present the results of the rather simple linear regression models.

[10] Note that our regression model (equation 4.1) relates the OMM gamma exposure in the beginning of the time period to the realized volatility during the time period. Thus, our regression suggests that based on the publicly available trade repository data we use in this report (and the subsequent estimation of the OMM gamma exposure), a good forecast of the realized spot market volatility is possible. As the realized volatility is strongly correlated with the implied volatility, this could suggest a potential trading strategy. If we see that the OMM increased her short gamma exposure by i.e., $10bn$ EURUSD from time $t-2$ to time $t-1$ (so the net gamma exposure changes by $-10bn$), we know that the realized spot market volatility of EURUSD in the time period $t-1$ to t will go up by $(-10bn) \cdot (-0.001\frac{\%}{bn}) = 0.01\%$. Furthermore, we know that the realized spot market volatility is correlated with the implied volatility. If we therefore buy an option at $t-1$ and sell the option again at t, we sell 0.01% higher (option prices are usually quoted in terms of volatility). The converse is true if we see that the OMM decreases her short gamma exposure by $10bn$. Of course, this trading strategy is based on an isolated view that only takes the observed feedback effect that the OMM gamma exposure has on the spot market volatility into account. We empirically tested this trading strategy with the data behind the results presented in table 4.3 and found that the observed effect has a Sharpe ratio close to zero, which indicates efficient market behavior. Additionally, the bid ask spreads in the options markets are usually wider than 0.01% anyway, so taking transaction costs into account, the strategy would even be unprofitable.

This section served as purely empirical discussion of the feedback effect, independent of our volatility model framework. The next section now bridges the empirical discussion with our volatility model framework.

4.2.4. Relation to the volatility model

Note that the regression model presented in the previous section is just the linear approximation of equation 2.10. Equation 4.3 shows the derivation of the linear regression model, where we first write the series expansion and then a linear approximation of it (the last step in the derivation is justified by $|\beta h N S_t \gamma_t^M| <= 1$).

$$\begin{aligned}
\sigma_t &= \frac{\nu}{(1 + \beta h N S_t \gamma_t^M)} && (4.3a)\\
&= \nu \sum_{k=0}^{\infty} (-\beta h N S_t \gamma_t^M)^k && (4.3b)\\
&= \nu(1 - \beta h N S_t \gamma_t^M + O(2)) && (4.3c)\\
&\approx \underbrace{\nu}_{=\alpha_1} + \underbrace{(-\nu\beta h)}_{=\alpha_2} N S_t \gamma_t^M && (4.3d)
\end{aligned}$$

Therefore, the estimations presented above provide support for our volatility model or, rather, its linear approximation.

$N S_t \gamma_t^M$ is the gamma exposure of the OMM. As $\alpha_2 = -\nu\beta h$, where $h = h^M - h^T$ can be positive or negative (depending on whether the OMM or the OMT hedges a larger share of their exposure), $\nu > 0$ (fundamental volatility) and $\beta > 0$ (market impact), the sign of α_2 tells us whether the OMM hedges a larger share of delta exposure compared to the OMT (as explained before). We estimate α_2 to be negative, so we can conclude that $h > 0$ and therefore that the OMM hedges a higher share of her delta exposure compared to the OMT. The value of α_2 quantifies the percentage increase in spot market volatility per gamma exposure of the OMM.

As we estimate $\alpha_1 = \nu$ separately, we could calculate βh from $\alpha_2 = -\nu\beta h$. Nevertheless, we cannot further dissect βh in its parts, as β and h are not further specified. [11]

4.3. Validation of the volatility model

As we found empirical validity for the linear approximation of our volatility model presented in section 4.2, we now estimate the general model (equation 2.10) in this section by performing a parametric fit with the use of a least-squares approach. The results are presented in figure 4. We use daily data for the estimation, as we do not want to mix different intraday volatility regimes. The black points in figure 4 present the data points. The blue line is the model fit based on those data points. The dashed blue lines are the confidence intervals of the regression.

[11] One approach would be to fix one of the two parameters: if e.g., we want to know h, we could take the approach to fix the value of β by using an estimation of the Kyle model for the corresponding FX pair (and further assume that β and h are both time-invariant). This would allow us to solve for h, with which we could subsequently calculate the net delta hedge amounts that are executed in the spot market, which is just the negative difference of the delta exposure of the OMM between two points in time multiplied by h. Nevertheless, we did not find a reliable empirical estimation of the Kyle model for the FX market, and this approach will not be pursued further here for now.

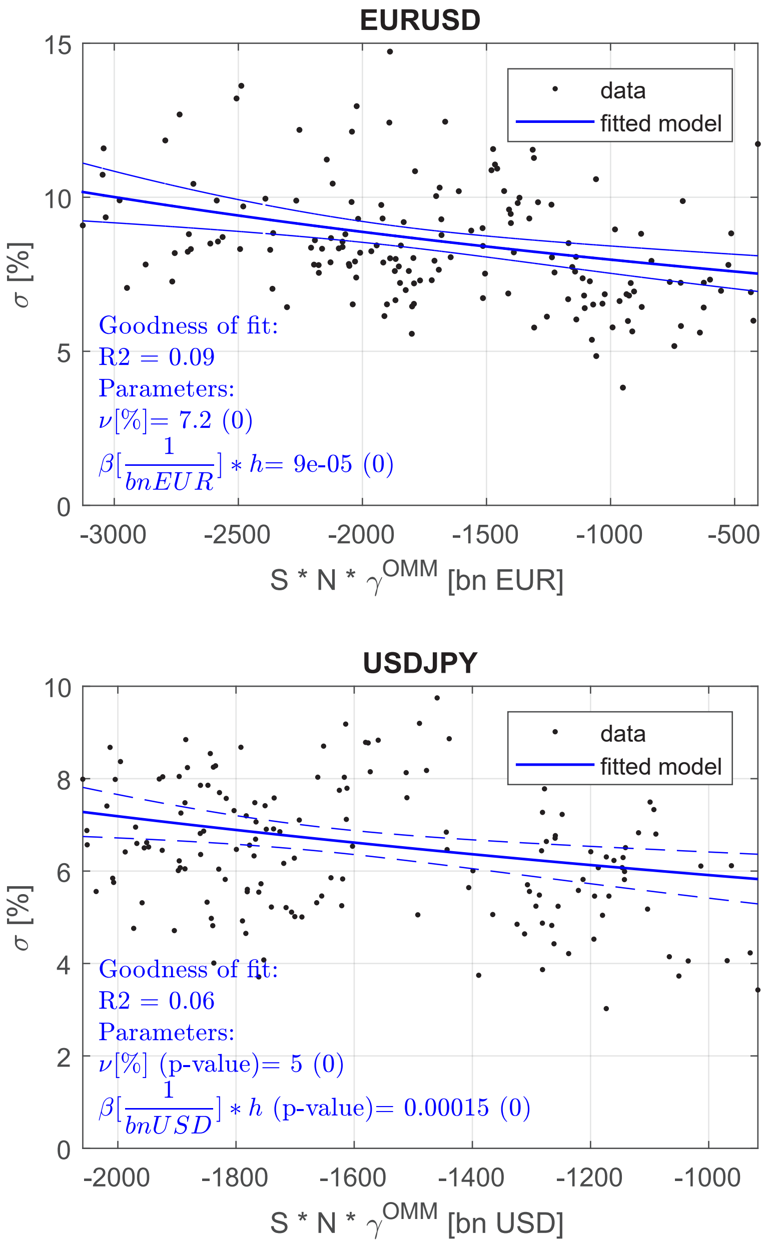

Figure 4: This figure shows the result of the model fit. The black points represent the data points, and the volatility σ is conditional on the gamma exposure. The blue line is the least-square fit of our volatility model (equation 2.10) based on those data points. The result of the parameter fit is displayed in the figure. As the factor βh is very small, the model appears nearly linear.

We see that the goodness of the fit of the nonlinear fit for both EURUSD and USDJPY has the same value as the one for the linear regression model, which is because the nonlinear volatility model appears nearly linear, as the factor βh is very small. Furthermore, the parameter estimates for the factor βh are of similar

	EURUSD	USDJPY
ν (%)	7.2^{***}	5^{***}
	(21.6)	(11.6)
$\beta * h$ ($\frac{\%}{\text{bn}}$)	0.0009^{***}	0.0015^{***}
	(4.8)	(3.8)
R-squared	0.09	0.06
F-test vs. const. model	1430^{***}	1280^{***}

Table 3: This table shows the least-square fit results of the nonlinear model fit for EURUSD and USDJPY. (t-statistics are presented in brackets. The stars *** denote significance at the 1% level.)

size between EURUSD and USDJPY and are in accordance with the results of the linear regression model. Again, we see that the impact of one billion gamma exposure of the OMM is slightly higher in USDJPY than in EURUSD, which, as stated before, can be explained by the fact that the USDJPY spot market is less liquid (and therefore the impact of trades, β, is higher) BIS (2016). Again, the value of βh is not further disentangled. The value of ν is again higher for EURUSD than for USDJPY, as we estimated the model with daily data.

Figure 5 discusses the fitted volatility model for EURUSD and USDJPY with the fitted parameters. In the upper panel, we present the fitted volatility model. The middle panel presents the absolute volatility increase due to the feedback effect, and the lower figure presents the relative volatility increase. We see that a negative gamma exposure of the OMM of approximately -1000 billion in the base currency (which is around what we observe in figure 3) leads to an increase of the EURUSD volatility above the fundamental volatility level of 0.7% and to an increase of the USDJPY volatility above the fundamental volatility level of 0.9%. In the lower panel, we subsequently see that this gamma exposure of approximately -1000 billion in the base currency amounts to a relative increase above the fundamental volatility level of approximately 10% for EURUSD ($\frac{0.7\%}{7.2\%}$) and 18% for USDJPY ($\frac{0.9\%}{5.0\%}$).

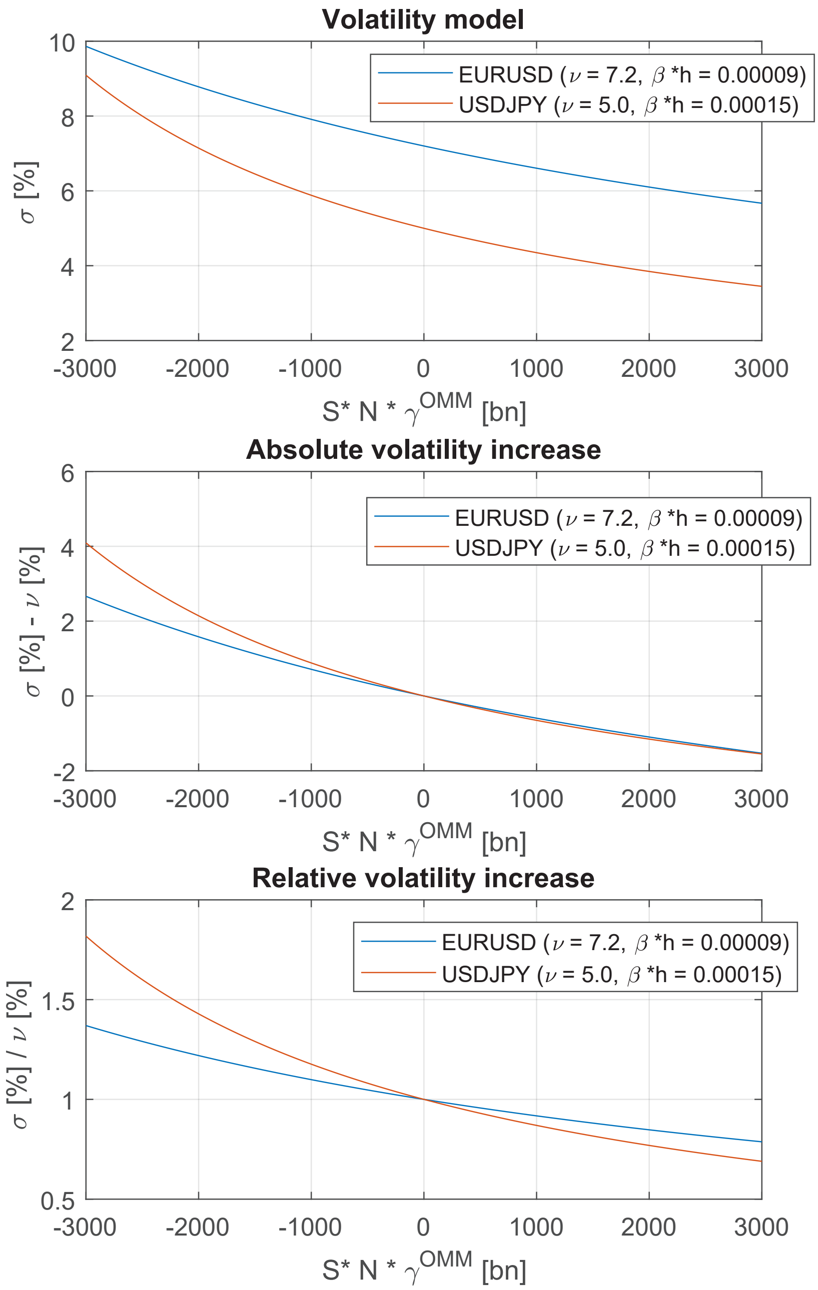

Figure 5: In these panels, we discuss our volatility model with the fitted parameters for EURUSD and USDJPY. We present the fitted volatility model in the upper panel, the absolute volatility increase due to delta hedging activity in the middle panel and the relative volatility increase in the lower panel.

5. Discussion and Conclusion

To model the feedback effect of delta hedging for spot market volatility, we have proposed a model economy of two types of traders, an OMM and an OMT, whose exposures each reflect the total outstanding positions in the market. If the OMM and the OMT have different hedge ratios, this will result in a net delta hedge activity that introduces market friction and feedback effects. We conceptualize this friction by assuming a simple linear, permanent impact model for the net delta hedge volumes that are executed in the spot market. From this model we subsequently derive a spot market volatility model that depends on the gamma exposure of the trader who hedges a larger share of her delta exposure and on the market impact of the delta hedge transactions. For example, if the OMM has a negative gamma exposure and she hedges a higher degree of her delta exposure compared to the OMT (two of our empirical findings), the spot market volatility goes up, as the net hedge amounts would be positive if the spot price moves up and negative if the spot price moves down (which results in an increase in the spot market volatility). The implications of this model are in line with the theoretical results of Lions and Lasry (2006) and Li and Almgren (2016).

Feedback effects of this nature have not been empirically discussed, as the data to do so were not available. In the case of our volatility model, we needed the aggregated OMM's gamma exposure in order to empirically validate the model and therefore quantify the feedback effect. We reconstructed these data by using the publicly available DTCC trade repository data. This reconstructed net gamma exposure of the OMM was found to be negative, as expected: investors usually buy options with either consideration of the spot price or with the desire to hedge other positions. As those traders usually act as OMTs in the markets, we expect the OMT to be net long on options. As the OMM provides liquidity as a service to the market, his position will be the reverse of the OMT. Therefore, a net short position makes sense from our point of view.

With the reconstructed OMM's gamma exposure, we subsequently first empirically discuss the feedback effect. We want to restate the two main points here. First, our regression model shows a high goodness of fit and a highly significant parameter for the gamma exposure of the OMM. Therefore, the gamma exposure of the OMM indeed has high importance for the spot market volatility. Second, the hedge ratio of the OMM is higher than that of the OMT. This, again, is in line with our expectation and in accordance with the line of reasoning presented above: As OMTs are usually investors who buy options with consideration of the spot price or to hedge other positions in their portfolio, they are expected to not hedge their option's delta. As the OMMs provide liquidity to the market as a service but usually without taking a view on the spot, they are assumed to hedge their large portfolio in order to reduce the risk involved in their business model. Therefore, we expected the hedge ratio of the OMM to be higher than that of the OMT, which is exactly what we found. Consequently, we also observe that the volatility is increased by the OMM's short gamma exposure.

Finally, we also validate our volatility model with the reconstructed OMM's gamma exposure. A least-square fit shows a high goodness of fit and very similar parameter estimates compared to the linear regression model. We find that the OMM hedges a larger share of her delta exposure compared to the OMM and that a negative gamma exposure of the OMM of approximately -1000 billion in the base currency (which is around what we observe from our reconstructed OMM data) leads to an absolute increase in volatility of 0.7% in EURUSD and 0.9% in USDJPY. If we assume that the hedge ratios in both markets are the same, this difference can be directly explained by the higher market impact of a transaction in the USDJPY spot market compared to the EURUSD spot market, which makes sense, since the liquidity of the EURUSD spot market is higher than that of the USDJPY spot market. This absolute increase amounts to a relative increase in the volatility above its fundamental level (volatility level in the absence of any delta hedging) of approximately 10% for EURUSD ($\frac{0.7\%}{7.2}$) and 18% for USDJPY ($\frac{0.9\%}{5.0}$). Therefore, our results are in line with previous theoretical work on the feedback effect that delta hedging strategies have on the spot market volatility, as e.g., Sircar and Papanicolaou (1998) suggest a relative increase in the spot market volatility of approximately 18%.

Note that it would be wrong to conclude that the options market introduces a destabilizing or stabilizing effect on the spot market via the delta hedging strategies per se. Whether the effect of the option market on the spot market is stabilizing or not only depends on the net gamma position of the OMM (as the OMM is found to hedge a higher share of her delta exposure compared to the OMT). However, since the OMM is usually short gamma, as other market participants refrain from selling options, the option market is found to lead to an increase in the spot market volatility. Moreover, while these effects appear to exist, implied volatility is still trading at a premium to the realized volatility! One would assume that the OMM would thus only stop selling options if her hedging strategy would drive the realized volatility above the implied volatility.

In future work, we intend to discuss how the persistence of the gamma exposure of the OMM could explain part of the autocorrelation that is observed in the spot market volatility; we also intend to discuss the local volatility phenomena and to examine the connection of an extreme gamma position in the context of market crashes.

Abergel, F. and Loeper, G. (2016). *Option Pricing and Hedging with Liquidity Costs and Market Impact, Proceedings of the International Workshop on Econophysics and Sociophysics*. Springer.

Almgren, R. (2015). Treasury price swings on october 15, 2014. Technical report, Quantitative Brokers. Technical report.

Almgren, R. and Chriss, N. (2000). Optimal execution of portfolio transactions. *Mathematics and Financial Economics*, 3(1):5–39.

Avellaneda, M. and Lipkin, M. D. (2003). A market-induced mechanism for stock pinning. *Quant. Finance*, 3(6):417425.

BIS (2016). Triennial central bank survey of foreign exchange and otc derivatives markets in 2016. Technical report, Bank for International Settlements. Triennial Central Bank Survey.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *J. Polit. Econ.*, 81(3):637–654.

Bouchard, B., Loeper, G., and Zou, Y. (2016). Almost-sure hedging with permanent price impact. *Finance and Stochastics*, 20(3):741–771.

Bouchard, B., Loeper, G., and Zous, Y. (2017). Hedging of covered options with linear market impact and gamma constraint. *SIAM J. Control Optim.*, 55(5):33193348.

Bouchaud, J. P. (2010). Price impact. In *Encyclopedia of Quantitative Finance*. Wiley.

Brennan, M. and Schwartz, E. (1989). Portfolio insurance and financial market equilibrium. *Journal of Business*, 62(4):455–476.

Chakrabarty, B., Pascual, R., and Shkilkoc, A. (2015). Evaluating trade classification algorithms: Bulk volume classification versus the tick rule and the lee-ready algorithm. *Journal of Financial Markets*, 25(C):52–79.

DTCC (2011). Explanation of trade information warehouse data. Technical report, The Depository Trust & Clearing Coorporation. Technical Document.

Easley, D., de Prado, M., and O'Hara, M. (2016). Discerning information from trade data. *Journal of Financial Economics*, 120(2):269–285.

Farmer, J., Gerig, A., Lillo, F., and Mike, S. (2006). Market efficiency and the long memory of supply and demand: Is price impact variable and permanent or fixed and temporary? *Quantitative Finance*, 6(2):107–112.

Frey, R. and Stremme, A. (1997). Market volatility and feedback effects from dynamic hedging. *Mathematical Finance*, 7(4):351–374.

Gueant, O. and Pu, J. (2015). Option pricing and hedging with execution costs and market impact. *Mathematical Finance*, 27(3):803–831.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):13151335.

Lee, C. M. C. and Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance*, 46(2):733–746.

Li, T. M. and Almgren, R. (2016). Option hedging with smooth market impact. *Market Microstructure and Liquidity*, 2(1):1650002.

Lions, P.-L. and Lasry, J.-M. (2006). Towards a self-consistent theory of volatility. *J. Math. Pures Appl.*, 86(6):541551.

Loeper, G. (2013). Option pricing with market impact and non-linear Black and Scholes pdes. *SSRN Electronic Journal*.

MacKenzie, D. (2008). *An Engine, Not a Camera: How Financial Models Shape Markets*. MIT Press, 1 edition.

Merton, R. C. (1973). The theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183.

Miller, M. (1997). *Financial Innovations and Market Volatility*. Basil Blackwell Ldt.

Naujokat, F. and Westray, N. (2011). Curve following in illiquid markets. *Mathematics and Financial Economics*, 4(4):299–335.

Rogers, L. C. G. and Singh, S. (2010). The cost of illiquidity and its effects on hedging. *Mathematical Finance*, 20(4):597615.

Shtauber, A. and Marone, G. (2014). Can DTCC positioning data predict EMFX. Technical report, Deutsche Bank. Deutsche Bank Markets Research.

Sircar, K. R. and Papanicolaou, G. (1998). General Black - Scholes models accounting for increased market volatility from hedging strategies. *Applied Mathematical Finance*, 5(1):45–82.

Weng, N. and Grover, R. (2017). Options flows matter: introducing the deutsche bank skew volume indicators. Technical report, Deutsche Bank. Deutsche Bank Markets Research.

Winkler, R. (2017). Foreign exchange coffee report. Technical report, Deutsche Bank. Deutsche Bank Markets Research.

Appendix A. Data Specification

Appendix A.1. DTCC data

The Dodd-Frank Act, which was signed into law in July 2010, requires the reporting of all over-the-counter FX option transactions that were traded in the US or involve US persons to a data repository. There are several data repositories to which market participants can report. We use the Depository Trust and Clearing Corporations repository (DTCC). DTCC (2011) Weng and Grover (2017)

Most relevant information of the options trade is included in the repository, such as the type of option traded, the notional amount (volume), the strike price, the expiry date, the premium paid, the currency for quoting the notional and the premium and the option trade timestamp. The missing information is the investor type (hedge fund, corporation, asset manager, private client, etc.) and the side of the trade (if it was a paid or given); in addition, trades done outside the US are not represented. Our report fills one of the gaps, namely, the missing information about the side behind a transaction, by reconstructing the implied volatility of each trade and subsequently matching it with the reference bid/ask price at the time of the trade to check if the trade happened on the bid or the ask side (for more detail, see appendix Appendix B.1 and Appendix B.2). The key in this reconstruction and matching procedure therefore is to have accurate time stamps from the DTCC repository in a sense that they need to be reported immediately to the repository. The intraday turnover distribution of the DTCC data time stamps thereby indicates that the DTCC data inherit this accuracy, as the distribution of the time stamps clearly follows the usual intraday turnover pattern observed in all financial markets.

Furthermore, the DTCC data are a representative source of the total options market. The share is approximately 20% of the total turnover for all main FX pairs (so the share of the DTCC transactions among the total turnover in the options market, as published by the BIS, is approximately 20%). Additionally, the volumes traded in EURUSD and USDJPY are sufficiently high for an analysis that mainly focuses on the spillover effects of the options market onto the spot market. Weng and Grover (2017)

All of the above factors make the DTCC repository an ideal source for our analysis purposes, and due to the large volumes traded in the EURUSD and USDJPY options market, we focus on those two instruments. Weng and Grover (2017)

Appendix A.2. General remarks on the data used in this report

- The options price stream onto which the reconstructed implied volatilities are matched comes from Bloomberg. Therefore, we adjust the measures by tenor and delta where needed (by spline interpolation).
- The forward points, interest rates and option price streams that we use in the estimation of the implied volatility and the calculation of the delta and gamma of the OMM at each point in time are also taken from Bloomberg.
- Only vanilla option contracts with a time to expiry below 600 days are taken into account.
- The realized annualized spot market volatility is calculated based on mid-price data from the corresponding primary FX market of the currency pairs.
- Analysis date range: 20.9.2017 to 30.6.2018.

- Intraday Bloomberg price streams (which we need for the reconstruction of the hedger's position) are only available six months into the past. We thus cannot estimate the hedger's position (from the DTCC data) for the period before the analysis period. Therefore, the estimated gamma exposure of the OMM is inaccurate in the beginning of the analysis period, as only the "new" transactions contribute to our estimate of the OMM exposure. However, the farther we go into the analysis period, the more accurately we can know the OMM exposure. As the median time to expiry of the analyzed option contracts turned out to be approximately 30 days, we leave out the first 30 days in our analysis to obtain a representative data set of the OMM positioning. All figures are also displayed accordingly.

Appendix B. Derivation of the hedger's positioning behind each transaction

The DTCC repository provides useful information on each transaction but lacks one key point for our analysis: whether the deal was a paid (buyer initiated) or given (seller initiated). We solve this question by using a trade classification algorithm. First, we reconstruct the implied volatility of each trade by using the data from the DTCC trade repository. Second, using a trade classification algorithm, we match the reconstructed implied volatility onto a reference bid/ask price stream from Bloomberg and estimate whether it was traded on the bid or the ask price. By doing so, we estimate the side behind each transaction, that is, if it was a paid (active buy transaction) or a given (active sell transaction). Third, we derive the OMM's position by assuming that she is always liquidity provider (i.e., if the transaction was identified as paid, we assume the OMM is short the contract).

Appendix B.1. Recipe on how to construct the implied volatility from the transaction data

Here, we outline Blacks model. These formulas are used in the determination of the implied volatility.

$$C(F,\tau) = D \cdot (N(d_+)F - N(d_-)K) \tag{B.1}$$

$$P(F,\tau) = D \cdot (N(-d_-)K - N(-d_+)F) \tag{B.2}$$

$$d_\pm = \frac{1}{\sigma\sqrt{\tau}} \cdot (ln(\frac{F}{K}) \pm \frac{1}{2} \cdot \sigma^2 \tau) \tag{B.3}$$

$$d_\pm = d_\mp \pm \sigma\sqrt{\tau} \tag{B.4}$$

$$\tau = T - t \tag{B.5}$$

$$D = e^{-r_{price\ ccy} \cdot \tau} \tag{B.6}$$

$$F = S \cdot e^{(r_{price\ ccy} - r_{base\ ccy}) \cdot \tau} \tag{B.7}$$

Now we briefly explain the variables and comment on where the information for each variable comes from.

- N: Cumulative distribution function of the standard normal distribution
- T: Expiry date; comes from DTCC data
- t: Trade date; comes from DTCC data
- τ: Time-to-expiry (also called tenor); can be calculated with equation (B.5)
- K: Strike price; comes from DTCC data
- $r_{price\ ccy}$: Interest rate of the price currency for the tenor τ; comes from Bloomberg; obtained by either looking at the bond yield of the corresponding country with tenor τ or, if τ is in between two standard bond contracts, calculating an interpolated yield from the two corresponding bond contracts (e.g., for a tenor of $\tau = 1.3$ years, the bond yield of the $\tau = 1$ year and $\tau = 2$ years is interpolated)
- $r_{base\ ccy}$: Interest rate of the base currency for the tenor τ; comes from Bloomberg; obtained by either looking at the bond yield of the corresponding country with tenor τ or, if τ is in between two standard bond contracts, calculating an interpolated yield from the two corresponding bond contracts.
- $F(\tau)$: Forward price (amount of price currency per unit of base currency in time τ); comes from Bloomberg; obtained by either looking at the forward points of the corresponding currency with tenor τ or, if τ is in between two standard contracts, calculating an interpolated forward point from the two corresponding contracts (Note: We can get the data for the forward point from Bloomberg directly or calculate the forward with the interest rates $r_{price\ ccy}$ and $r_{base\ ccy}$. We prefer the first method, as it probably is a more realistic measure because it includes the cross-currency basis.).
- S: Spot price (amount of price currency per unit of base currency); comes from Bloomberg
- C, P: Premium of the call (C) and put (P) option; comes from DTCC data

Using the above variables and formulas, we see that the only unknown quantity is the implied volatility σ. This quantity is our objective. We therefore use the Black-Scholes formula to obtain an expression of the premium, which only depends on one unknown, namely, the implied volatility of the options contract. The corresponding equation is given by (B.1) for a call option or (B.2) for a put option. In the last step, we numerically minimize the following equation and estimate the implied volatility of each transaction.

$$\min_{\sigma} \|C_{Black\,formula}(\sigma) - C_{observed\ in\ DTCC}\| \rightarrow \sigma_{constructed} \tag{B.8}$$

Appendix B.2. Recipe on how to determine the OMM side behind each transaction

- Once we estimate the implied volatility behind a transaction, we can match the transaction time stamp (which is given by the DTCC data) to a corresponding reference bid/ask price stream for the implied volatility (e.g., from Bloomberg). Note that it is important that the reference price stream is adjusted to match the tenor and delta of the reconstructed implied volatility of the transaction. Thus, the reference price stream onto which we match the reconstructed implied volatility of the transaction results from spline interpolation of the volatility surface (e.g., from Bloomberg) at the time when the transaction occurred. If the difference between the estimated implied volatility and the ask price of Bloomberg is *smaller* than the corresponding difference between the estimated implied volatility and the bid price from Bloomberg, we label the transaction as a paid, as the transaction probably occurred on the ask price. Conversely, if the difference between the estimated implied volatility and the ask price from Bloomberg is *larger* than the corresponding difference between the estimated implied volatility and the bid price from Bloomberg, we label the transaction as a given.

$$|\sigma_{estimated} - \sigma^{ask}_{reference}| - |\sigma_{estimated} - \sigma^{bid}_{reference}| \rightarrow \begin{cases} < 0, & \text{paid} \\ > 0, & \text{given} \end{cases} \tag{B.9}$$

- Now we assume that the market maker is always a liquidity provider, so she always quotes on both sides. This implies that she is short the position if the transaction is a paid and long if the transaction is a given.

Appendix C. Recipe on how to determine the OMM delta and gamma exposure at each time step

Now that we have derived the OMM's position behind each transaction, we can calculate her potential delta and gamma exposure at each point in time using equations (C.1) and (C.2) for the delta and equation (C.3) for the gamma. K is the strike price, S is the current spot price, F is the interpolated forward price for the remaining time, σ is the interpolated implied volatility for the remaining time, and τ is the remaining time.

$$\Delta_{call} = \frac{\partial C(F, K, S, \sigma, \tau)}{\partial S} = N(d_+) \tag{C.1}$$

$$\Delta_{put} = \frac{\partial P(F, K, S, \sigma, \tau)}{\partial S} = -N(-d_+) \tag{C.2}$$

$$\Gamma = \frac{N^{'}(d_1)}{S\sigma\sqrt{\tau}} \tag{C.3}$$

Note that $N^{'}$ is the standard normal probability density function and N is the cumulative distribution function of the standard normal distribution. Both measures, gamma and delta, are subsequently multiplied by the signed notional of each position (e.g., -1 million base currency if we are short an option with a notional of 1 million and +1 million base currency if we are long an option with a notional of 1 million), resulting in the gamma and delta exposure of the OMM. Equation (C.5) describes the delta exposure and equation (C.5) the gamma exposure of one option contract, where A is the notional of the option contract and D is a dummy for the side that the OMM has (as described by equation (C.6)). If we sum up the gamma exposures of all active (not yet expired) positions for a certain point in time, we end up with the net gamma exposure of the OMM for this time.

$$\delta_{exposure} = \Delta * A * D \tag{C.4}$$

$$\gamma_{exposure} = \Gamma * A * D \tag{C.5}$$

$$D = \begin{cases} 1 & \text{if market maker is long} \\ -1 & \text{if market maker is short} \end{cases} \tag{C.6}$$

Recent SNB Working Papers

2019-3 Benjamin Anderegg, Didier Sornette, Florian Ulmann:
Quantification of feedback effects in FX options markets

2019-2 Katrin Assenmacher, Franz Seitz, Jörn Tenhofen:
The demand for Swiss banknotes: some new evidence

2019-1 Darlena Tartari, Albi Tola:
Does the IMF Program Implementation Matter for Sovereign Spreads? The Case of Selected European Emerging Markets

2018-19 Samuel Reynard:
Negative Interest Rate, OE and Exit

2018-18 Hans-Ueli Hunziker, Christian Raggi, Rina Rosenblatt-Wisch, Attilio Zanetti:
The impact of guidance, short-term dynamics and individual characteristics on firms' long-term inflation expectations

2018-17 David R. Haab, Thomas Nitschka:
Carry trade and forward premium puzzle from the perspective of a safe-haven currency

2018-16 Gregor Bäurle, Elizabeth Steiner and Gabriel Züllig:
Forecasting the production side of GDP

2018-15 Christian Grisse, Gisle J. Natvik:
Sovereign debt crises and cross-country assistance

2018-14 Thomas Lustenberger:
Has the American Output Growth Path Experienced a Permanent Change?

2018-13 Stephan Imhof, Cyril Monnet and Shengxing Zhang:
The Risk-Taking Channel of Liquidity Regulations and Monetary Policy

2018-12 Andreas M. Fischer, Henrike Groeger, Philip Sauré and Pinar Yeşin:
Current account adjustment and retained earnings

2018-11 Maxime Phillot, Rina Rosenblatt-Wisch:
Inflation Expectations: The Effect of Question Ordering on Forecast Inconsistencies

2018-10 Alex Cukierman, Thomas Lustenberger:
International Evidence on Professional Interest Rate Forecasts:
The Impact of Forecasting Ability

2018-9 Thomas Nitschka:
Did China's anti-corruption campaign affect the risk premium on stocks of global luxury goods firms?

2018-8 Andreas M. Fischer, Lucca Zachmann:
Do the rich pay their taxes early?

2018-7 Basil Guggenheim, Mario Meichle and Thomas Nellen:
Confederation debt management since 1970

2018-6 Lucas Marc Fuhrer, Basil Guggenheim and Matthias Jüttner:
What do Swiss franc Libor futures really tell us?

2018-5 Barthélémy Bonadio, Andreas M. Fischer and Philip Sauré:
The speed of exchange rate pass-through

2018-4 Simone Auer, Christian Friedrich, Maja Ganarin, Teodora Paligorova, and Pascal Towbin:
International Monetary Policy Transmission trough Banks in Small Open Economies

2018-3 Romain Baeriswyl, Camille Cornand and Bruno Ziliotto:
Observing and shaping the market: the dilemma of central banks

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK