Constrained Liquidity Provision in Currency Markets*

Wenqian Huang[†], Angelo Ranaldo[‡], Andreas Schrimpf[§], and Fabricius Somogyi[¶]

This version: November 2, 2022

Working paper and subject to change.

Abstract
We study dealers' liquidity provision in the currency market. We show that at times when dealers' intermediation capacity is constrained their cost of liquidity provision increases disproportionately relative to dealer-provided volume. As a result, the elasticity of dealers' liquidity provision weakens by at least 80% relative to periods when they are unconstrained. We identify constrained periods based on leverage ratios, Value-at-Risk measures, credit default spreads, and debt funding costs. We interpret our novel empirical findings within a parsimonious model that sheds light on the key mechanisms of how liquidity provision by dealers tends to weaken when intermediary constraints are tightening.

J.E.L. classification: F31, G12, G15

Keywords: Currency markets, dealer constraints, market liquidity, foreign exchange, liquidity provision.

*We welcome comments, including references to related papers we have inadvertently overlooked. We are grateful to Briana Chang (discussant), Stijn Claessens, Semyon Malamud, Paul Söderlind, Dagfinn Rime (discussant), and Jonathan Wallen for helpful comments. We also thank seminar and conference participants at the Bank for International Settlements, 11th Workshop on Exchange Rates (SNB), King's College London, 11th Annual 2022 Stern/Salomon Center Microstructure Conference, 2022 SFI Research Days in Gerzensee, CLS Group, Credit Suisse, University of St.Gallen, VfS Annual Conference 2022, and the Boston Fed. Angelo Ranaldo acknowledges financial support from the Swiss National Science Foundation (SNSF grant 182303). All errors are our own.

[†]Bank for International Settlements, Switzerland. E-mail: `wenqian.huang@bis.org`.

[‡]University of St.Gallen and Swiss Finance Institute, Switzerland. E-mail: `angelo.ranaldo@unisg.ch`.

[§]Bank for International Settlements and CEPR, Switzerland. E-mail: `andreas.schrimpf@bis.org`.

[¶]Northeastern University, United States. E-mail: `f.somogyi@northeastern.edu`.

1. Introduction

Financial intermediaries play a crucial role in supporting the functioning of financial markets. This is especially true for over-the-counter (OTC) markets, which are not organised based on centralised exchanges but rely on the intermediation by dealers providing immediacy to their clients.[1] However, dealers' willingness to provide liquidity depends on their ability to make balance sheet space available. Thus, constraints on dealers' intermediation capacity can reduce dealers' incentives to intermediate trades, increasing the cost of liquidity provision, as well as generating widespread violations of no-arbitrage conditions.[2]

Against this backdrop, the key question that we address in this paper is how dealer constraints affect liquidity provision: do constrained dealers increase the cost of market liquidity, or do they even down-scale their market-making activity? To analyse this question, we study prices and quantities in the foreign exchange (FX) market based on a globally representative trade data set from CLS Group. The FX market is the largest financial market in the world by turnover and often regarded as one of the most liquid ones. Given that many dealer banks that are active in FX also provide intermediation services across a host of financial markets, we expect that our findings about the key mechanisms also apply to other OTC markets.

How intermediation constraints impact market liquidity is an important issue for both policy makers and academics. In fact, it has been much debated whether the policies put in place since the Global Financial Crisis in 2008/09 have had any (unintended) consequence on the functioning of OTC markets by disincentivising liquidity provision (e.g., Dudley, 2018). On top of the regulatory costs, the academic literature seeks to understand the factors influencing dealers' willingness to use their balance sheets for intermediation. Moreover, the monitoring of how liquidity costs fluctuate and are affected by dealer behavior is also beneficial for practitioners as well as central banks operating in such markets.

The contribution of this paper is threefold. First, we provide a simple, yet effective, analytical framework to tease out the costs of providing FX spot market liquidity from a no-arbitrage condition. The framework builds on the well-known triangular relation among FX rates (e.g., Chaboud, Chiquoine, Hjalmarsson, and Vega, 2014; Foucault, Kozhan, and Tham, 2016) and outlines two measures of liquidity costs that both depend on FX dealers' intermediation capacity. VLOOP, which stands for violation of the law of one price, is the deviation of the midquotes from the triangular relation. It captures the shadow cost of intermediary constraints and arises when dealers' marginal valuation of a currency pair deviates from its fundamental value. TCOST is the round-trip transaction cost of performing a triangular arbitrage trade. It represents the dealer's realised compensation to endure inventory risks due to clients' order flow imbalances. Our second contribution is to study VLOOP and TCOST through the

[1] To be clear, we focus on the role of FX dealer banks as liquidity providers rather than cross-market arbitrageurs. This is consistent with the role that these institutions have played after the clampdown on proprietary trading in the aftermath of the Global Financial Crisis.

[2] See "Holistic Review of the March Market Turmoil," Financial Stability Board, November 2020.

lenses of constrained financial intermediaries. In particular, we document a novel nonlinear relation between the cost of liquidity provision and dealer-intermediated volume. We show that dealers' liquidity provision deteriorates along two dimensions when they face intermediation constraints: i) their provided market liquidity becomes increasingly more costly, and ii) dealers' liquidity supply becomes less elastic relative to states when they are unconstrained. Lastly, we provide a tractable model that formalises the economic intuition for these novel empirical findings.

Our paper consists of three parts. We start by proposing a novel analytical approach for measuring the cost of liquidity provision. The conceptual underpinning is a no-arbitrage relation that ties together triplets of FX spot rates. We show that deviations from this no-arbitrage condition represent an amalgam of two liquidity cost components: VLOOP and TCOST. In line with the literature on intermediary asset pricing (e.g., Adrian, Etula, and Muir, 2014; Duffie, 2018; Fleckenstein and Longstaff, 2018; Du, Hébert, and Huber, forthcoming 2022), a natural interpretation of VLOOP is the shadow cost of intermediary constraints.[3] By contrast, TCOST captures dealers' demanded compensation for enduring inventory imbalances. Using FX spot transaction data from CLS Group, we document three new empirical results. First, both liquidity cost measures clearly matter from an economic standpoint and, as expected, VLOOP is an order of magnitude smaller than TCOST. Second, the two liquidity cost components move in tandem over time, albeit their correlation is less than 30% on average. This is consistent with the idea that both are affected by dealers' intermediation capacity, but in different ways. Third, both components, but especially TCOST, tend to rise as intermediated volumes increase. This correlation pattern is in line with the idea that dealers require a higher compensation when providing more immediacy to clients.

In the second part of the paper, we link the cost of liquidity provision to dealers' constraints. The key mechanism is that constrained dealers are less inclined to make balance sheet space available to clients who wish to trade currency positions that create persistent inventory imbalances. We empirically investigate two different types of dealer behaviour: First, dealers can pass their intermediation costs onto clients by charging higher prices for immediacy. As such, this will be reflected in our two measures of liquidity costs, which both increase as intermediated volumes rise. Second, dealers are less inclined to maintain large and unbalanced inventories for market making purposes when their intermediation capacity is lower. The higher trading costs charged by dealers in turn also disincentivise potential liquidity traders from trading and thereby mitigating the imbalance. What this boils down to is that dealer constraints can have a bearing on both the *cost* and *quantity* of FX liquidity provision. In particular, tightening dealer constraints increase the cost of providing liquidity, which in turn renders dealers' liquidity provision less elastic.

In principle, the mechanism of interest could be studied based on any no-arbitrage con-

[3] In line with this strand of literature, one may also refer to these shadow costs as "balance sheet costs" associated with FX spot liquidity provision.

dition (e.g., put-call parity, treasury cash bonds vs futures, etc.). However, the triangular FX no-arbitrage condition offers at least two main advantages: First, it provides us with a clean laboratory to measure the cost of liquidity provision in the FX *spot* market. This is because unlike other no-arbitrage conditions that embed various frictions[4] (e.g., counterparty risk, funding roll-over risk, etc.) the triangular no-arbitrage condition captures the shadow cost of constraints on intermediation activities that are nearly risk-free (i.e., FX spot) and which can be financed by deploying existing balance sheet cash or short-term funding (Andersen, Duffie, and Song, 2019). Second, it allows us to dissect two liquidity cost components (i.e., VLOOP and TCOST) and to attribute them an economic meaningful interpretation. Using our unique FX spot volume data set, we are able to thoroughly quantify these two cost components. Moreover, FX no-arbitrage conditions that involve forward contracts (e.g., covered interest parity) are less accurate measures of the shadow cost of intermediary constraints in FX spot trading compared to VLOOP. This is because these other arbitrages are also driven by constraints that arise from liquidity provision in FX swap and money markets (Rime, Schrimpf, and Syrstad, 2021).

For our empirical analysis, we construct a composite measure capturing both regulatory and market risk constraints faced by FX dealer banks. Specifically, we consider the following variables: i) FX Value-at-Risk, ii) He, Kelly, and Manela (2017) leverage ratio, iii) credit default swap (CDS) premium, and iv) funding costs that are particularly relevant for debt-financed positions and market-making functions (see, e.g., Andersen et al., 2019; Berndt, Duffie, and Zhu, 2020). Our time-varying "dealer constraint measure" DCM is defined as the first principal component of these four variables for the largest 10 FX dealer banks. To rule out reverse causality, DCM enters all our regressions with a lag of 1 day (and our results are robust to using more lags). From an institutional perceptive, the relation between the cost of liquidity provision (i.e., VLOOP and TCOST) and dealer-intermediated volume is unlikely to have any effect on DCM. This is because FX spot trading itself does not affect dealer leverage or funding costs, not least given substantial internalisation of trades and netting.[5]

Equipped with a comprehensive measure of dealers' constraints, we analyse the relation between our two measures of liquidity costs and dealer-intermediated volumes. Two main results emerge concerning liquidity costs and volumes. First, as expected, the *cost* of liquidity provision increases significantly with dealers' constraints. Second, we find that the *quantity* of supplied liquidity crucially depends on dealers' intermediation constraints. More specifically, our analysis highlights how the "elasticity of liquidity provision" (i.e., the correlation between liquidity costs and volumes) decreases as dealers become more constrained.

To establish our main results, we rely on logistic smooth transition regression (LSTAR)

[4]For instance, Siriwardane, Sunderam, and Wallen (2021) show that the low correlation in the violation of seven no-arbitrage conditions (excluding the FX triangular no-arbitrage condition) is attributable to various frictions driving both the segmentation of balance sheets and also the funding abilities of financial firms.

[5]Moore, Schrimpf, and Sushko (2016) document that some of the major FX dealer banks have internalisation ratios of up to 90%. Moreover, CLS claims that multilateral netting shrinks funding needs by over 96%.

methods (e.g., Christiansen, Ranaldo, and Söderlind, 2011; Jeanneret, 2019) that are particularly well-suited to capture correlations across different regimes. The main goal of the regime-dependent analysis is to investigate how the relation between dealer-provided volume and liquidity costs changes across two regimes: i) normal times when dealers are unconstrained and ii) stressed periods when they face constraints to their intermediation capacity. While in the normal regime dealer-intermediated volume and our liquidity cost measures co-move with an average correlation of 12–31%, this correlation falls off sharply in the regime when dealers are constrained.[6] In particular, the correlation between the two liquidity costs (i.e., VLOOP and TCOST) and intermediated volume decreases by 14–15 percentage points in constrained periods. We interpret this result as a drop in dealers' elasticity to provide liquidity. All our findings remain qualitatively unchanged when considering each determinant of the dealer constraint measure DCM as a single regime variable.

In the third part of the paper, we build on inventory management models (see, e.g., Grossman and Miller, 1988; Hendershott and Menkveld, 2014) to rationalise our empirical findings. Our model features two periods, three currency pairs, and two types of agents: i) a risk-averse and debt-financed dealer supplying liquidity and ii) liquidity traders with exogenous trading demands. In particular, heterogeneity in private values among liquidity traders results in demand imbalances across the three pairs (Gabaix and Maggiori, 2015). Such imbalances in turn weigh on the dealer's marginal valuation of the three pairs due to the implied "balance sheet costs." These costs incentivise the dealer to set mid-quotes away from their fundamental value and result in violations of the law of one price (i.e., VLOOP). Thus, VLOOP reflects the shadow cost of the dealer's constraints in FX spot trading.[7] Within this context, TCOST is the demanded compensation of the dealer to accommodate customer flows that engender inventory risk. In normal times, when the dealer is unconstrained, the shadow cost of intermediary constraints (i.e., VLOOP) and also transaction costs (i.e., TCOST) are small. Both liquidity cost measures increase as dealer-intermediated volume rises since a larger trading demand implies that the dealer needs to endure larger inventory risk. But, at times when the dealer is constrained liquidity costs increase disproportionately more compared to equilibrium volume. This is because it is optimal for the constrained dealer to i) set mid-quotes further away from the fundamental value (i.e., higher VLOOP) as well as ii) charge a higher bid-ask spread (i.e., higher TCOST) to clients. Taken together, the increase in both liquidity cost measures leads to a slower increase in equilibrium trading volume and hence, renders liquidity provision to be less elastic.

[6]The key focus of our analysis lies on the supply side of liquidity provision that we capture via our dealer constraint measure DCM. In the robustness section of the paper, we also use a structural vector autoregression with sign restrictions to disentangle liquidity demand and supply dynamics. The setup here closely follows Goldberg (2020) and Goldberg and Nozawa (2020), respectively.

[7]For most market participants it will, however, not be feasible to benefit from such deviations due to high transaction costs (i.e., TCOST) as well as FX quantity conventions on major trading platforms. For instance, on inter-bank trading platforms such as EBS or Reuters, only trade sizes exceeding one million of the base currency will be accepted. This can lead to potential slack when performing triangular arbitrage trades.

2. Related literature

Our paper contributes to three strands of literature. First, our work is related to the broader literature that emphasises the role of intermediary frictions in affecting asset prices and, in particular, risk premia (e.g., Gârleanu and Pedersen, 2011; He and Krishnamurthy, 2011, 2013; Adrian and Boyarchenko, 2012; Adrian et al., 2014; He et al., 2017). Our main contribution is to show in depth how constrained dealers curtail their liquidity provision and charge higher liquidity costs in times of markets stress. This finding is also consistent with the evidence in Nagel (2012) showing that market makers' liquidity supply is increasing in their intermediation capacity and decreasing in the level of risk. Moreover, our paper corroborates the idea that market-wide liquidity conditions depend on intermediaries' balance sheet capacity (e.g., Adrian and Shin, 2010; Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes, 2010) and that intermediary leverage and banks' risk management practices (e.g., following Value-at-Risk methodologies) tend to be pro-cyclical (Adrian and Shin, 2013). Lastly, our findings suggesting that the cost of dealers' balance sheet space affects both the cost and quantity of liquidity provision are consistent with slow-moving intermediated capital being a key factor contributing to distortions in asset pricing relations (Duffie, 2010).

Second, we add to the literature on limits to arbitrage. Our work differs from previous research along two important dimensions. First, while prior research has mostly focused on constrained *arbitrageurs* (e.g., Shleifer and Vishny, 1997; Gromb and Vayanos, 2002; Hombert and Thesmar, 2014) and more recently Du et al. (forthcoming 2022) and Siriwardane et al. (2021), we study the role of constrained *dealers* and how their ability to provide immediacy contributes to market liquidity. Second, we propose a well-known no-arbitrage identity to derive two liquidity cost components with an economically meaningful interpretation. Thus, our key contribution is to analyse arbitrage conditions to shed light on the relation between liquidity costs and trading volume and to show how this relation critically depends on the intermediation capacity of dealers. In addition, a large body of prior research has studied limits to arbitrage in equity markets (see Gromb and Vayanos, 2010). However, many of the frictions considered in that literature, such as short sale constraints which are considered a major explanation in equities (e.g., Chu, Hirshleifer, and Ma, 2020), do not apply to FX. Related to the stock market literature, recent studies document widespread mispricings in stressful conditions (Pasquariello, 2014), commonality in arbitrage deviations (Rösch, Subrahmanyam, and van Dijk, 2016), and limits to arbitrage impacting market liquidity (Rösch, 2021). We add to this branch of the literature by identifying constrained intermediaries as the main driving force behind such commonalities and by showing how their liquidity provision impacts aggregate trading volume.

Lastly, we contribute to the literature on FX microstructure on understanding the role of trading volume. Our key angle here is to study the relation of both quantities (i.e., volumes from CLS Group) and prices. In contrast to the order flow literature (e.g., Evans, 2002;

Evans and Lyons, 2002, 2005), the literature on trading volume is relatively scarce due to the lack of comprehensive data. Earlier research has focused largely on the inter-dealer segment, which is dominated by two platforms: Reuters (e.g., Evans, 2002; Payne, 2003; Foucault et al., 2016) and EBS (e.g., Chaboud, Chernenko, and Wright, 2008; Mancini, Ranaldo, and Wrampelmeyer, 2013; Chaboud et al., 2014). Other sources of FX spot volume are proprietary data sets from specific dealer banks (e.g., Bjønnes and Rime, 2005; Menkhoff, Sarno, Schmeling, and Schrimpf, 2016). The recent public access to CLS data has enabled researchers to study customer-dealer volume at a global scale (Hasbrouck and Levich, 2018, 2021; Ranaldo and Santucci de Magistris, forthcoming 2022; Cespa, Gargano, Riddiough, and Sarno, 2021; Ranaldo and Somogyi, 2021). We contribute to this strand of literature by investigating the impact of dealer constraints on both the cost and quantity dimension of global FX liquidity.

3. Measuring the cost of liquidity provision in FX markets

3.1. Data sources

The empirical analysis employs high-frequency trade and quote data from two main sources. The FX spot volume data come directly from CLS Group (CLS) and are sampled hourly. Note that the dataset excludes any trades between two market makers or two price takers and hence only includes trading activity that is *intermediated* by FX dealer banks. This data set is publicly available from CLS or via Quandl.com, a financial and economic data provider.[8] CLS data have been used in prior research, among others, by Hasbrouck and Levich (2018, 2021), Ranaldo and Santucci de Magistris (forthcoming 2022), Cespa et al. (2021), and Ranaldo and Somogyi (2021). These authors have comprehensively described the CLS volume data.

The full sample period spans from November 2011 to September 2020 and includes data for 18 major currencies and 33 currency pairs. Our goal is to construct measures capturing the cost of dealers' liquidity provision. We derive these measures from the classical triangular no-arbitrage relation involving one non-dollar currency pair (e.g., AUDJPY) and two dollar legs (i.e., USDAUD and USDJPY). Hence, we exclude the USDHKD, USDILS, USDKRW, USDMXP, USDSGD, and USDZAR from our sample because there are no further non-dollar currency pairs involving the respective quote currencies (i.e., HKD, ILS, KRW, MXP, SGD, and ZAR). Furthermore, to maintain a balanced panel, we also remove all currency pairs involving the Hungarian forint (HUF), which enters the data set later, on 7 November 2015.[9] The remaining 25 currency pairs cover at least 75% of global FX spot trading volume according to

[8] CLS operates the world's largest payment-vs-payment settlement system handling up to 50% of global FX volumes. At settlement, CLS mitigates principal and operational risk by settling both sides of the trade at once.

[9] This filtering leaves us 15 non-dollar currency pairs (i.e., AUDJPY, AUDNZD, CADJPY, EURAUD, EURCAD, EURCHF, EURDKK, EURGBP, EURJPY, EURNOK, EURSEK, GBPAUD, GBPCAD, GBPCHF, and GBPJPY) and 10 dollar pairs (i.e., USDAUD, USDCAD, USDCHF, USDDKK, USDEUR, USDGBP, USDJPY, USDNOK, USDNZD, and USDSEK) that are used to synthetically replicate each of the non-dollar pairs.

the Bank for International Settlements (see "Triennial central bank survey — global foreign exchange market turnover in 2019," September 2019).

Next, we pair the hourly FX volume data with intraday spot bid and ask quotes from Olsen, a well-known provider of high-frequency data. Olsen compiles historical tick-by-tick data from various electronic trading platforms, both from the inter-dealer and dealer-customer segments. The indicative bid and ask quotes are directly available for all 25 currency pairs but do not correspond to actually executable transaction prices. This is not an issue for our purposes for two reasons: First, we are interested in measuring the cost of liquidity provision rather than identifying actual triangular arbitrage opportunities. Second, on average, the correlation of Olsen indicative quotes with tradeable EBS best bid and offer prices is around 99% and the mean absolute error is roughly 3.3%.[10]

3.2. Key variables

We measure the cost of liquidity provision in the FX spot market along two dimensions: i) violations of the law of one price (VLOOP), and ii) round-trip transaction costs (TCOST). In a first step, we explain how we derive the two components and, in a second step, elaborate how they relate to the costs that dealers face when providing FX spot liquidity.

Conceptually, VLOOP captures the price dislocations for two assets or trading positions with the same intrinsic value, while TCOST refers to the round-trip trading costs to take advantage of the dislocations. We derive both measures from the well-known triangular arbitrage trade that takes advantage of three FX rates (e.g., Chaboud et al., 2014; Foucault et al., 2016). The VLOOP component of the triangular arbitrage trade can be computed with midquote prices reflecting the intrinsic values of the direct and indirect positions. TCOST in turn is computed from the bid and ask quotes (depending on the base and quote currency).

Deriving VLOOP and TCOST from the triangular no-arbitrage relation. To derive VLOOP, consider a trader exchanging one euro (EUR) to some amount of US dollar (USD), exchanging the amount of US dollar to some amount of Canadian dollar (CAD) and exchanging back the amount of Canadian dollar to euro instantaneously at time t. The final amount of such a round-trip transaction measured in euro is given as:

$$\Delta_t \equiv \prod_{i=1}^{3} P_{i,t}, \tag{1}$$

where $P_{1,t} = \frac{1}{USDEUR_t^{mid}}$, $P_{2,t} = USDCAD_t^{mid}$, and $P_{3,t} = \frac{1}{EURCAD_t^{mid}}$ denote midquote exchange rates expressed as the amount of quote currency per unit of base currency, for instance, 1.255 Canadian dollar per US dollar (i.e., indirect quotation).

[10]To be precise, we estimate correlations and mean absolute errors individually for 25 currency pairs over the full-year of 2016. For brevity, we relegate these results to the Online Appendix.

The trader has identified a violation of the law of one price if Δ_t is different from unity. Note that Δ_t may be positive or negative depending on the direction of the trade but will be identical in absolute terms (if measured in logs) irrespective of the initial endowment of the trader (i.e., CAD, EUR or USD). Clearly, an arbitrageur would take this into account by choosing the direction of the triangular no-arbitrage trade, provided that corresponding arbitrage profits can be reaped. Panel A in Figure 1 provides a schematic overview of how to measure such deviations from triangular no-arbitrage conditions based on midquote prices.

To derive TCOST, we consider the same trader as before but now incorporate transaction costs by accounting for bid-ask spreads. Specifically, for every transaction that a trader makes, they pay the midquote price plus the half-spread. To reflect this, we replace the midquote prices in Eq. (1) by bid and ask prices, that is, $P_{1,t} = \frac{1}{USDEUR_t^{ask}}$, $P_{2,t} = USDCAD_t^{bid}$, and $P_{3,t} = \frac{1}{EURCAD_t^{ask}}$, respectively. The superscripts 'bid' and 'ask' refer to the price at which someone sells and buys one currency for another currency. Panel B in Figure 1 provides an overview of such a triangular arbitrage trade including transactions costs. Note that the bid and ask prices in this example are illustrative and do not correspond to actual data.

Figure 1: Triangular arbitrage trade

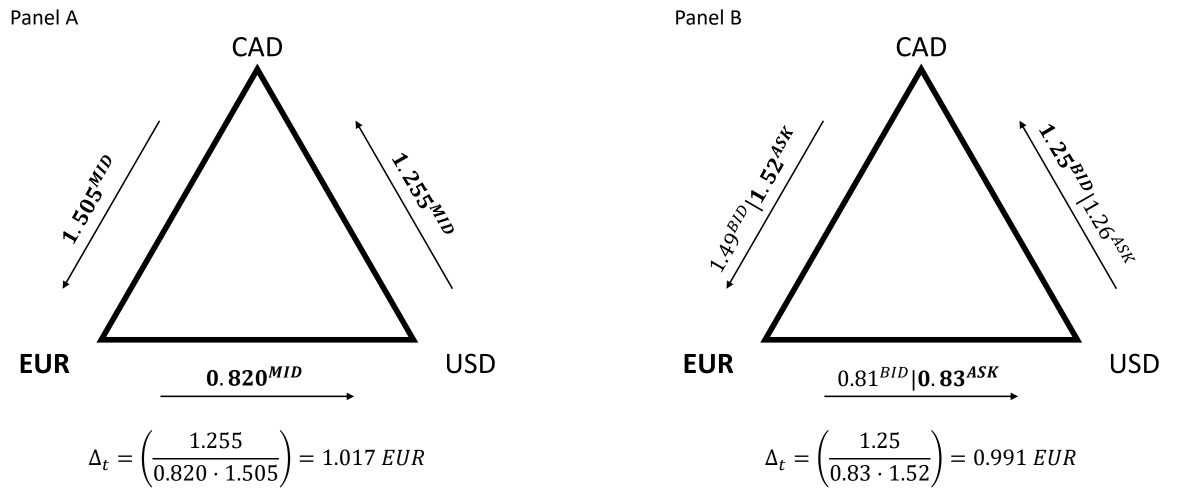

Note: This figure provides a schematic overview of a triangular arbitrage trade prior (Panel A) and after transaction costs (Panel B). The arrows denote the direction of the trade. Panel A shows the prior transaction cost return of a trader starting with one euro, first exchanging it to $\frac{1}{0.820} = 1.220$ US dollars, then exchanging 1.220 US dollars to Canadian dollars at the midquote price of 1.255 Canadian dollars per US dollar. This yields 1.531 Canadian dollars that are exchanged back to euros at the CADEUR midquote that is equivalent to $\frac{1}{EURCAD^{MID}} = \frac{1}{1.505}$. Such a round trip yields 1.017 euros or equivalent a positive return of 1.7% in this example. Panel B shows the return of first exchanging one euro to $\frac{1}{0.83} = 1.21$ US dollars at the ask price, then exchanging 1.21 US dollars to Canadian dollars at the bid price of 1.25 Canadian dollars per US dollar. This yields 1.51 Canadian dollars that are exchanged back to euros at the CADEUR bid price that is equivalent to $\frac{1}{EURCAD^{ASK}} = \frac{1}{1.52}$. Such a round trip yields 0.991 euros or equivalently a negative return of -0.9%.

The last step in the derivation of the two liquidity cost metrics consists of taking the log on both sides of Eq. (1), and leveraging the fact that bid and ask prices are the midquote minus and plus half the bid-ask spread. This yields the following expression:

$$log(\Delta_t) \equiv \underbrace{log\left(\frac{USDCAD_t^{mid}}{USDEUR_t^{mid} \cdot EURCAD_t^{mid}}\right)}_{VLOOP_t} - \underbrace{log\left(\frac{\left(1+\frac{USDEUR_t^{bas}}{2}\right)\cdot\left(1+\frac{EURCAD_t^{bas}}{2}\right)}{1-\frac{USDCAD_t^{bas}}{2}}\right)}_{TCOST_t}, \quad (2)$$

where the superscripts 'mid' and 'bas' denote the midquote price (i.e., the average of the bid and ask price) and relative bid-ask spread (i.e., the difference between ask and bid price relative to the midquote). Note that in this expression $TCOST_t$ is by definition *positive*.

The first part of Eq. (2) (i.e., $VLOOP_t$) captures the violations from the law of one price. Following the literature on intermediary asset pricing (e.g., Adrian et al., 2014; Duffie, 2018), we interpret these no-arbitrage violations as a measure of the lower bound of the shadow cost of intermediary constraints.[11] The second part (i.e., $TCOST_t$) reflects the cumulative round-trip transaction cost of performing such a triangular arbitrage trade. These transaction costs represent the dealer's compensation to endure an inventory imbalance due to the customers' demand for immediacy.

Following our methodology, we compute $VLOOP_t$ and $TCOST_t$ for $k = 1, 2, ..., 15$ triplets of currency pairs. A triplet of currency pairs is defined as one non-dollar currency pair (e.g., EURCAD) plus the two USD legs (e.g., USDEUR and USDCAD).[12] In particular, at every point in time we take the perspective of the arbitrageur above by first, identifying the seemingly profitable direction of the trade by conditioning on $VLOOP_t$ being positive and second, by extracting the associated transaction cost $TCOST_t$. Moreover, we prune the hourly time-series of $VLOOP_t$ and $TCOST_t$, respectively, for heavy outliers, which we define as observations in the top and bottom 1.5 percentiles of the data. Eventually, we can also compute daily measures of $VLOOP_t$ and $TCOST_t$ by summing up hourly observations for each day.

Empirical illustrations. Figure 2 shows the time-series and cross-sectional variation of hourly no-arbitrage violations $VLOOP_t$ (left y-axis) and round-trip transaction costs $TCOST_t$ (right y-axis), respectively. Economically, a higher reading of $VLOOP_t$ coincides with a larger shadow cost of intermediary constraints, whereas $TCOST_t$ captures the realised compensations for providing immediacy. Both measures of dealers' liquidity costs exhibit intuitive properties in the sense that they surge during periods of market stress and mean-revert during calm

[11] Dávila, Graves, and Parlatore (2022) show that these no-arbitrage violations correspond to the marginal social value of executing an arbitrage trade and that the welfare gain of closing arbitrage spreads is higher in more liquid markets such as FX spot.

[12] As a robustness check, we have also constructed triplets of euro-based currency pairs that do not involve any dollar currency pairs (e.g., AUD-EUR-JPY). This leaves us with 6 currency pair triplets: AUD-EUR-JPY, CAD-EUR-JPY, GBP-EUR-AUD, GBP-EUR-CAD, GBP-EUR-CHF, and GBP-EUR-JPY, respectively. All our key empirical results remain qualitatively unchanged when estimated based on this alternative cross-section of currency pair triplets. See the Online Appendix for these additional findings.

periods. The large spike during the Covid-19 market turmoil in March and April 2020 is particularly well pronounced across all 15 triplets of currency pairs and is indicative of the global nature of the stress. The correlation of $VLOOP_t$ and $TCOST_t$ is positive for the entire cross-section and ranges from 12–39%. We interpret this as evidence of commonality in no-arbitrage violations and market liquidity in the broader sense (Rösch, 2021).

Figure 2: No-arbitrage violations and round-trip transaction costs

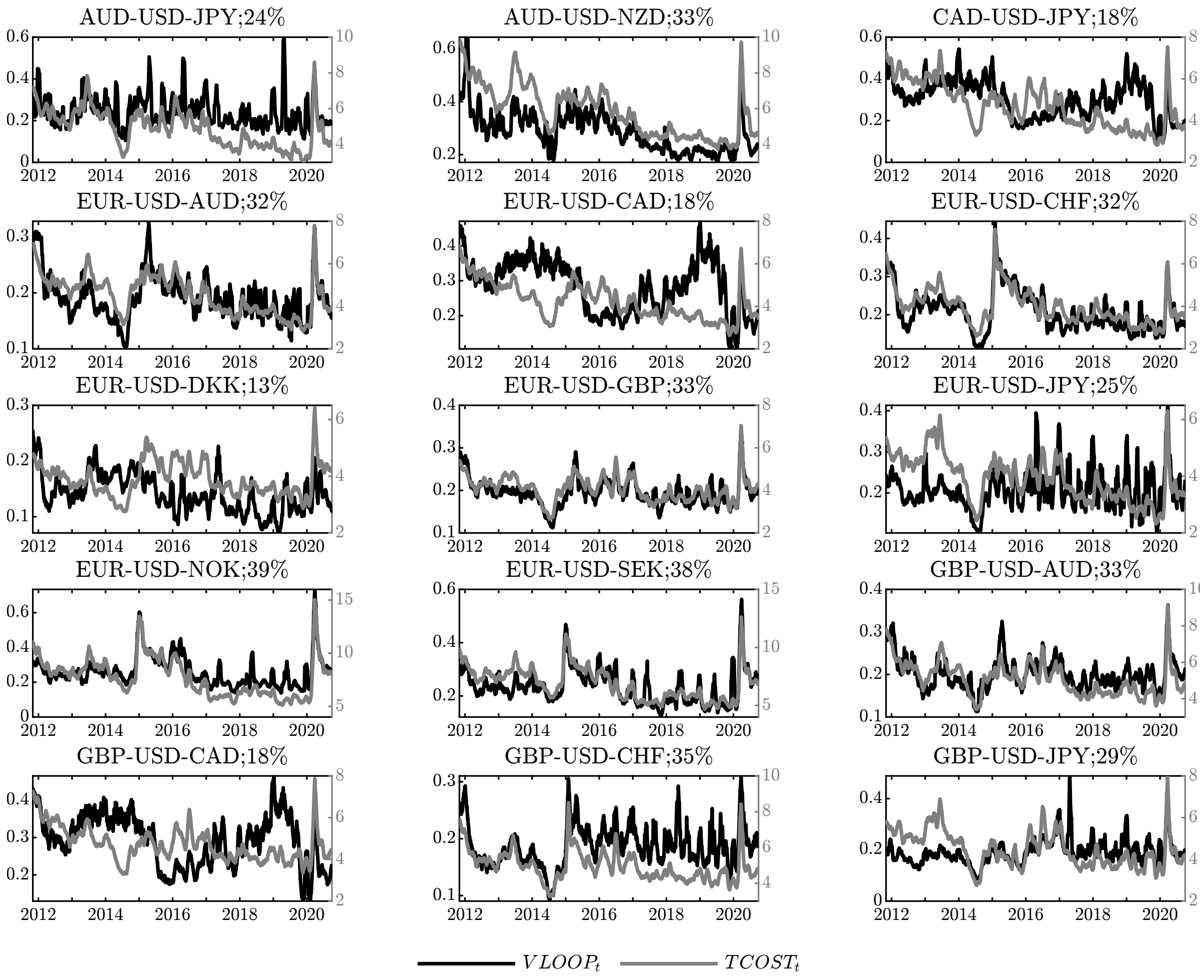

Note: This figure plots the 22-day moving averages of hourly triangular no-arbitrage deviations $VLOOP_t$ (left y-axis) and round-trip trading costs $TCOST_t$ (right y-axis), respectively, for 15 triplets of currency pairs. Both variables are measured in basis points. The numbers in the titles refer to the correlation coefficient of $VLOOP_t$ and $TCOST_t$. The sample covers the period from 1 November 2011 to 30 September 2020.

Table 1 reports the time-series average of hourly no-arbitrage deviations VLOOP and round-trip trading costs TCOST. In addition, it also tabulates hourly averages of direct trading volume in non-dollar currency pairs (e.g., AUDJPY) and synthetic trading volume in dollar pairs. By "synthetic" we refer to the sum of trading volume in two dollar pairs (e.g., USDAUD and USDJPY) within a triplet of currency pairs. Each row corresponds to one currency pair triplet, which we abbreviate as, for instance, AUD-USD-JPY.

This simple summary statistics table conveys three main insights: First, deviations from fundamentals VLOOP are an order of magnitude smaller than round-trip transaction costs TCOST. We interpret this result as suggestive evidence that dealers recharge their intermediation costs on the bid and ask prices offered to their customers. Another implication is that seemingly profitable violations of triangular no-arbitrage are most of the time not exploitable by the average trader as transactions costs are prohibitively high (i.e., there is no free lunch). Second, trading volume in non-dollar currency pairs is considerably smaller relative to the synthetic volume in dollar pairs. This is essentially the case for all 15 currency pair triplets but the effect is less pronounced for those involving the NOK and SEK, where the euro crosses play a bigger role. Lastly, the synthetic relative bid-ask spread is somewhat larger than the direct spread in non-dollar currency pairs. This finding is fully consistent with Somogyi (2021) showing that strategic complementarity in price impact rather than traditional trading costs (e.g., bid-ask spreads) is the key determinant of the cross-sectional differences in trading volume between dollar and non-dollar currency pairs.

Table 1: Summary statistics

	Liquidity cost in bps		Volume in $bn		Bid-ask spread in bps		Volatility in bps
	VLOOP	TCOST	Direct	Synthetic	Direct	Synthetic	Direct
AUD-USD-JPY	0.24	4.88	0.18	5.11	4.15	5.87	14.38
AUD-USD-NZD	0.29	5.85	0.09	2.01	4.44	7.43	9.32
CAD-USD-JPY	0.30	4.67	0.03	5.32	4.29	5.21	12.66
EUR-USD-AUD	0.19	4.52	0.14	7.72	3.54	5.64	11.54
EUR-USD-CAD	0.28	4.25	0.08	7.94	3.55	4.99	10.15
EUR-USD-CHF	0.21	3.98	0.37	6.76	2.62	5.41	6.38
EUR-USD-DKK	0.14	3.89	0.09	6.17	2.54	5.30	1.82
EUR-USD-GBP	0.19	4.07	0.61	8.16	3.19	4.95	9.52
EUR-USD-JPY	0.21	3.90	0.65	9.67	3.14	4.83	11.43
EUR-USD-NOK	0.26	7.69	0.24	6.25	6.25	9.40	11.01
EUR-USD-SEK	0.23	6.86	0.27	6.27	5.41	8.42	9.18
GBP-USD-AUD	0.20	5.08	0.04	3.60	4.22	5.99	12.53
GBP-USD-CAD	0.29	4.69	0.03	3.81	4.00	5.34	10.85
GBP-USD-CHF	0.19	4.94	0.03	2.64	4.09	5.76	10.69
GBP-USD-JPY	0.19	4.47	0.20	5.55	3.85	5.18	12.78

Note: This table reports the time-series average of hourly triangular no-arbitrage deviations $VLOOP$ in basis points (bps), round-trip trading costs $TCOST$ in bps, direct trading volume in non-dollar pairs (e.g., AUDJPY) in $bn, synthetic volume in dollar pairs in $bn, direct and synthetic relative bid-ask spreads, and realised volatility in non-dollar pairs in bps. By "synthetic" we refer to the sum of trading volumes and relative bid-ask spreads in two dollar pairs (e.g., USDAUD and USDJPY) within a currency pair triplet. Each row corresponds to a triplet of currency pairs, for example, AUDJPY, USDAUD, and USDJPY that we abbreviate as AUD-USD-JPY. The sample covers the period from 1 November 2011 to 30 September 2020.

4. Elasticity of liquidity provision and dealer constraints

This section presents evidence consistent with our hypothesis that dealers' liquidity provision becomes increasingly inelastic at times of heightened intermediation constraints. The analysis is split into two parts. We first start with some simple motivating evidence showing that the relationship between dealer-intermediated volumes and the cost of liquidity provision depends on dealers' intermediation constraints. We then draw on smooth transition regressions to study the state-dependent nature of this relation more formally.

4.1. Motivating evidence

For motivational purposes, we first present some descriptive evidence of how trading volume and the costs of liquidity provision (i.e., VLOOP and TCOST) co-move over time. In a second step, we then compute conditional correlations to show how the relation weakens as intermediaries become more constrained.

To measure dealer banks' intermediation constraints we derive a composite dealer constraint measure (DCM) that we compute in two steps. As a first step, we create four time-series based on cross-sectional averages of the top 10 FX dealer banks' (see Euromoney FX surveys) i) FX Value-at-Risk (VaR) of their overall trading book (quarterly), ii) He et al. (2017) leverage ratio (quarterly), iii) credit default swap (CDS) premia (daily) and iv) debt funding costs (daily). Figure 3 shows the time-series variation of these four measures that exhibit correlations ranging from 58–91% percent. See the Online Appendix for details on how we retrieve and compute each of these variables.

All of these factors can have a bearing on financial intermediaries' capacity to absorb customer order flow imbalances on their balance sheet. For example, self-imposed or regulatory-driven VaR-limits force dealers to scale back their market making or proprietary trading. Similarly, dealers' willingness to engage in market making activity and liquidity provision is linked to their risk profile, as reflected, for instance, in higher leverage and CDS premia on the bank's outstanding debt. In addition, elevated risk exposure can lead to an immediate increase in funding costs and valuation adjustments (XVA), which may manifest themselves in debt and funding value adjustments by dealers (Andersen et al., 2019). These factors in turn will affect dealers' assessment of the financing costs for its franchise, and as such will have a bearing on their intermediation activities.

Since large FX dealer banks (e.g., Citi Bank or UBS) typically operate on a global scale and provide liquidity in many currencies and other asset classes at once, they are even more exposed to these issues. As such, intermediaries might be forced to reduce liquidity provision in *all* currency pairs when they endure trading losses in particular positions and/or experience funding constraints affecting the whole dealer franchise.

As a second step, we extract the first principal component as our composite measure of dealer constraints. The first principal component explains around 83% of the total variance of the individual dealer constraint time-series. The key advantage of our dealer constraint measure DCM is that it encompasses a range of different factors that can all impact dealers' risk appetite and ability to warehouse risk.[13] By doing so, we are able to extract common

[13]One might wonder how much our results are driven by market-wide state factors that are not dealer specific.

Figure 3: State variables: Dealer constraint measure (DCM) and its components

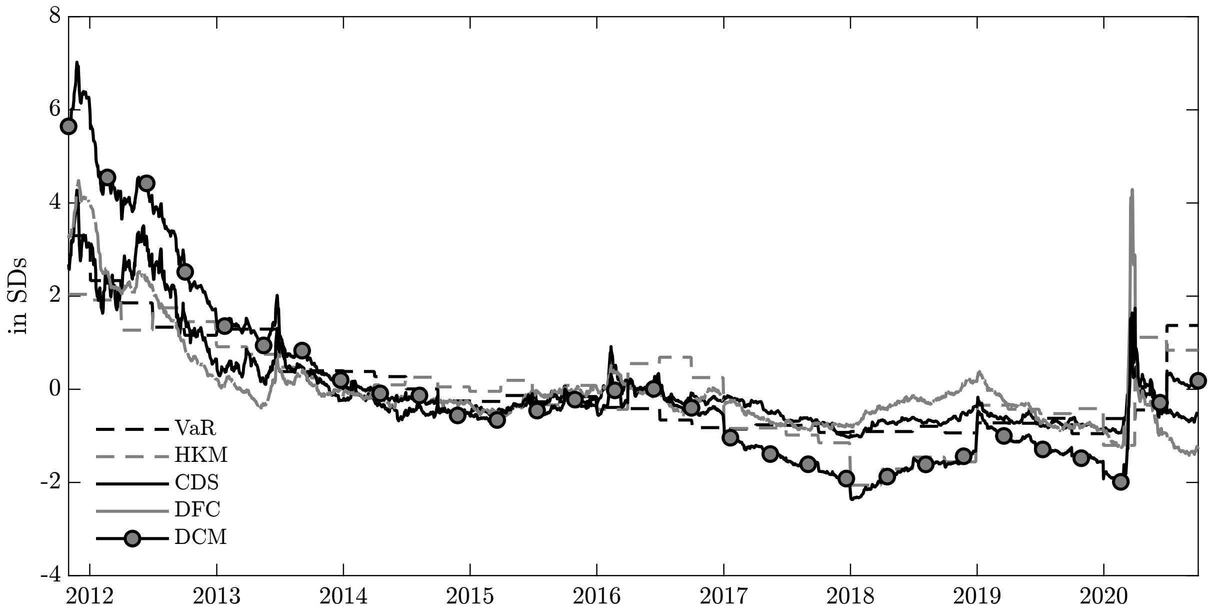

Note: This figure plots different state variables that we observe at the daily and quarterly frequency. Observations have been standardised by subtracting the sample mean and dividing by the sample standard deviation of every variable. The four state variables are primary FX dealer banks' i) quarterly Value-at-Risk measure (VaR, dashed black line), ii) quarterly He et al. (2017) leverage ratio (HKM, dashed grey line), iii) daily credit default spread (CDS, solid black line), and iv) daily funding cost yield (DFC, solid grey line). We define our dealer constraint measure (DCM, black solid line with grey markers) as the first principal component across these four variables. The sample covers the period from 1 November 2011 to 30 September 2020.

information of all these factors and obtain a *daily* measure of economic constraints on global dealers' intermediation activity.[14]

A valid question is whether our dealer constraint measure may also be affected by the amount of intermediated volume itself. There are at least two reasons why this is unlikely to be the case: First, we use the 1-day lagged value of DCM in all our state-dependent regression analyses to rule out any contemporaneous relation or reverse causality issues. Second, FX spot intermediation activity only minimally affects dealer leverage since it is a direct exchange of two currencies (i.e., an accounting exchange on the asset side). And similarly, FX spot trading volume is unlikely to affect the dealers' credit spread and overall funding costs.

Table 2 takes the descriptive analysis one step further by computing conditional correlation coefficients of (log) changes in each of our two liquidity cost measures (i.e., VLOOP and TCOST) and total dealer-intermediated trading volume (i.e., VLM) across the percentiles of

To prove robustness, we use the VIX index, TED spread, gold price, and the LIBOR-OIS spread. We find that these state variables do not appropriately capture the state-dependent relation between liquidity costs and trading volume. See the Online Appendix for output tables.

[14]Ideally, we would be able to measure dealer constraints intraday. Unfortunately, this is not feasible due to the low frequency availability of VaRs, leverage ratios, CDS spreads, and debt funding costs at the bank level.

our dealer constraint measure DCM.[15] Consistent with our intuition, we find that the correlation with each of the two liquidity cost measures weakens substantially as our dealer constraint measure DCM increases. For instance, the conditional correlation based on the highest DCM decile (i.e., when dealers are most constrained) is a mere 17% for TCOST, and hence economically and statistically significantly lower than the full-sample correlation of 31%. Moreover, we also observe a monotonic increase in both liquidity cost measures and also a stronger commonality of VLOOP and TCOST (i.e., $cor(VLOOP, TCOST)$) across the DCM percentiles.

These initial results suggest that market liquidity tends to deteriorate when dealers are more constrained. Our preliminary explanation (which we formalise in Section 5) for these empirical findings is that there are two main forces at play when dealers are constrained: On the one hand, violations of no-arbitrage conditions (i.e., VLOOP) increase as dealers charge a higher mark-up (or mark-down) across currency pairs reflecting the increase in the shadow cost of intermediary constraints. On the other hand, due to the surge in the cost of intermediation capacity dealers also post more conservative bid and ask quotes and thus increase transaction costs (i.e., TCOST). Consequently, a higher spread discourages clients' trading activity and suppresses trading demands. As a result, equilibrium volume increases at a slower pace relative to the rise in liquidity costs, which results into a weakening correlation between the cost of liquidity provision and dealer-provided volume during constrained periods.

Table 2: Liquidity provision cost characteristics across DCM percentiles

DCM percentile		$cor(VLOOP, VLM)$	$cor(TCOST, VLM)$	$cor(VLOOP, TCOST)$	$\overline{VLOOP}$ in %	$\overline{TCOST}$ in %	$\overline{VLM}$ in \$bn
full sample	0.0	0.12	0.31	0.22	0.05	1.12	148.66
least constrained	0.1	0.12	0.31	0.22	0.05	1.14	150.21
	0.2	0.11	**0.30	0.22	0.05	1.18	153.27
	0.3	0.13	*0.30	0.22	0.05	1.21	157.93
	0.4	0.13	**0.30	0.22	0.05	1.23	159.24
	0.5	0.12	***0.28	0.21	0.05	1.23	158.99
	0.6	0.12	***0.22	0.21	0.06	1.28	161.65
	0.7	0.10	***0.16	*0.20	0.06	1.32	169.36
	0.8	***0.06	***0.18	0.22	0.06	1.35	176.74
most constrained	0.9	***−0.03	***0.17	***0.28	0.06	1.44	180.52

Note: This table shows the conditional correlation of our two liquidity cost measures (i.e., $VLOOP$ and $TCOST$) and total trading volume VLM across the percentiles of the dealer constraint measure (DCM, columns 1 and 2). Column 3 reports the conditional correlation between $VLOOP$ and $TCOST$. The underlying data are based on a panel of 15 currency pair triplets. The asterisks *, **, and *** indicate that the correlation is significantly different from the full sample (in the first row) estimate at the 90%, 95%, and 99% levels. The corresponding test statistic for the conditional correlation cor^{τ} being equal to the full sample correlation $cor^{\tau=1.00}$, where $\tau \in 0.1, 0.2, ..., 0.9$ refers to DCM_t deciles, are based on the Fisher z-transformation. Columns 4 to 6 report the within-decile average $VLOOP_{k,t}$ ($\overline{VLOOP}$) in %, $TCOST_{k,t}$ ($\overline{TCOST}$) in %, and the average $VLM_{k,t}$ ($\overline{VLM}$) in \$bn across 15 currency pair triplets. The full sample covers the period from 1 November 2011 to 30 September 2020.

[15] Note that the CLS volume data include the FX trading activity of all top dealer banks listed in the Euromoney FX surveys. In particular, the banks that show up in the Euromoney FX surveys are also the most dominant players on the CLS settlement system. Moreover, the decline in DCM since 2012 is also consistent with the rise of electronic and algorithmic trading activity in the global FX market.

4.2. *Smooth transition regressions*

We next briefly describe our main econometric approach based on smooth transition regression (LSTAR) models (e.g., van Dijk, Teräsvirta, and Franses, 2002; Christiansen et al., 2011). The LSTAR model is particularly well-suited for our analysis as constrained and unconstrained regimes are determined endogenously and may vary smoothly over time.[16]

For the LSTAR model, let $G(z_{t-1})$ be a logistic function depending on the 1-day lagged regime variable z_{t-1}

$$G(z_{t-1}) = (1 + \exp(-\gamma'(z_{t-1} - c)))^{-1}, \tag{3}$$

where the parameter c is the central location and the vector γ determines the steepness of $G(z_{t-1})$. Hence, the LSTAR model is of the form

$$y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}, \tag{4}$$

where the dependent variable is one of our two liquidity cost measures (i.e., VLOOP or TCOST) and $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors. We include both cross-sectional α_k and time-series λ_t fixed effects to control for any unobservable heterogeneity that is constant across triplets of currency pairs k and days t, respectively. For estimation, we use the generalised method of moments (GMM) and determine the optimal parameters γ and c by nonlinear least squares minimising the concentrated sum of squared errors.[17] Note that the slope coefficients in Eq. (4) vary smoothly with the regime variable z_{t-1} from β_1 at low values of $\gamma' z_{t-1}$ to β_2 at high values of $\gamma' z_{t-1}$. There are two interesting boundary cases: First, if $\beta_1 = \beta_2$ we effectively have a linear regression. Second, the limit case where $\gamma \to \infty$ is equivalent to a linear regression with a dummy variable.

The state-dependent explanatory variable $f_{k,t}$ is the total dealer-provided trading volume (i.e., $VLM_{k,t}$) that is defined as the sum of trading volume in one non-dollar as well as two dollar currency pairs within a particular currency pair triplet k. The state-independent variable is the realised variance (i.e., $RV_{k,t}$) in the direct non-dollar currency pair (e.g., AUDJPY) that we estimate following Barndorff-Nielsen and Shephard (2002) as the sum of squared intraday midquote returns. Note that across all regression specifications both LHS and RHS variables are taken in logs and first differences. The obvious advantage of this is twofold: First, regression coefficients can be interpreted as elasticities. Second, FX volume in levels is non-stationary and persistent (see Ranaldo and Santucci de Magistris, forthcoming 2022), hence taking first-differences is an effective remedy to render the series stationary.

Table 3 shows the passage from a linear model with a dummy to a nonlinear smooth

[16] As a robustness check, we have also experimented with Markov chain switching models using DCM as an exogenous state variable and found consistent results. These additional findings are available upon request.

[17] Our inference is based on Driscoll and Kraay's (1998) covariance matrix that allows for random clustering and serial correlation up to 8 lags. We choose the optimal bandwidth using the plug-in procedure for automatic lag selection by Andrews and Monahan (1992) and Newey and West (1994), respectively.

transition regression (LSTAR). To be specific, the first two columns in this table report the results from estimating a linear model (OLS) of the form

$$y_{k,t} = \lambda_t + \alpha_k + \beta_1' f_{k,t} + \delta' f_{k,t} \cdot D_{t-1} + \beta_3' w_{k,t} + \epsilon_{k,t}, \tag{5}$$

where $f_{k,t}$ and $w_{k,t}$ collect all regressors and D_{t-1} is a 1-day lagged interaction variable capturing periods of heightened dealer constraints. Note that the estimate of δ corresponds to the difference between the constrained and unconstrained regime coefficient (i.e., $\beta_2 - \beta_1$) of the LSTAR model in Eq. (4). In column 'Dummy', D_t is equal to one if DCM_t is above its 75% percentile in period t. Note that the specification in 'Dummy' is a simple, yet intuitive, approximation of nonlinear regression model. In column 'Logistic', D_t is a logistic transformation of DCM_t based on $1/[1 + exp(-\gamma DCM_t)]$, where γ determines the steepness of the function. For simplicity, we set $\gamma = 1$ but the results in column 'Logistic' are robust for values of γ ranging from 1 to 12. Lastly, in column 'LSTAR' the table shows results from a smooth transition regression as in Eq. (4), which constitutes our main econometric approach. Again, the logistic function $G(z_{t-1})$ depends on 1-day lagged values of our dealer constraint measure DCM_t.[18] Note that across each of the three specifications we control for the realised variance in the direct non-dollar currency pair (e.g., AUDJPY).[19]

There is a consistent picture that arises across all three specifications in Table 3: the difference between the slope coefficient on total trading volume in constrained and unconstrained periods (i.e., $\beta_2 - \beta_1$) is negative and highly statistically significant for both VLOOP and TCOST. Moreover, the estimated slope coefficients are almost identical for the linear model with dummy (column 'Dummy') and the LSTAR model (column 'LSTAR'). In both cases, the slope coefficient is at least 80% (e.g., $-0.08/0.09 = 89\%$) lower when dealer banks are constrained and hence less willing or able to cater their customers' trading demands. In sum, two findings stand out from this analysis: First, dealer-provided volume covaries significantly less with VLOOP and TCOST during times when dealers are more constrained. Second, it is above all the relation between VLOOP and trading volume that strongly diverges and even exhibits negative (albeit insignificant) coefficients in the constrained regime.

Taken together, these results highlight the state-dependent nature of the relation between our two liquidity cost measures (i.e., VLOOP and TCOST) and dealer-intermediated trading volume.[20] Moreover, the estimated coefficient linking VLOOP and TCOST to DCM can be interpreted as an elasticity due the log and first-difference transformation of all variables. From this perspective, our results suggest that when dealers face constraints their liquidity

[18] In the Online Appendix we show that our findings are robust to using up to 22 lags (see Figures B.4 and B.5).

[19] Moreover, in the Online Appendix we show that our findings are robust to including the cross-currency (CIP) basis (e.g., Akram, Rime, and Sarno, 2008; Du, Tepper, and Verdelhan, 2018) and the Amihud (2002) price impact as a control for FX funding (Andersen et al., 2019; Rime et al., 2021) and market liquidity (e.g., Ranaldo and Santucci de Magistris, forthcoming 2022), respectively.

[20] Note that our estimates for the difference in the slope coefficients across constrained and unconstrained regimes are very similar when using either VLOOP or TCOST as a regressor and VLM as the dependent variable.

supply becomes less elastic. As a result, FX liquidity conditions (i.e., VLOOP and TCOST) deteriorate (see Table 2), since liquidity supply does not keep up with the rise in demand.

Table 3: From linear model with dummies to smooth transition regressions

	VLOOP			TCOST		
	Dummy	Logistic	LSTAR	Dummy	Logistic	LSTAR
γ		1.00	***12.05		1.00	***5.94
c			***−0.14			***0.39
Unconstrained volume	***0.08	***0.16	***0.11	***0.09	***0.12	***0.09
	[2.94]	[3.47]	[3.50]	[11.07]	[8.55]	[10.85]
Constrained volume	−0.07	*−0.09	−0.05	0.01	0.02	0.01
	[1.37]	[1.72]	[1.40]	[0.85]	[1.14]	[0.96]
Realised variance	**0.02	**0.02	**0.02	***0.03	***0.03	***0.03
	[2.05]	[2.01]	[2.02]	[7.98]	[7.95]	[7.95]
Constrained-Unconstrained	***−0.14	***−0.25	***−0.16	***−0.08	***−0.10	***−0.08
	[2.66]	[2.90]	[3.25]	[4.88]	[4.19]	[4.78]
R^2 in %	0.13	0.14	0.15	3.78	3.73	3.78
Avg. #Time periods	2,182	2,182	2,182	2,185	2,185	2,185
#Currency triplets	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes

Note: In columns labelled 'Dummy' and 'Logistic' this table reports results from estimating a linear model (OLS) of the form $y_{k,t} = \lambda_t + \alpha_k + \beta_1' f_{k,t} + \delta' f_{k,t} \cdot D_{t-1} + \beta_3' w_{k,t} + \epsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., *VLOOP* or *TCOST*), $f_{k,t}$ and $w_{k,t}$ collect all regressors and D_{t-1} is a 1-day lagged interaction variable capturing distressed market periods. Note that the estimate of δ corresponds to the difference between the constrained and unconstrained regime coefficient (i.e., $\beta_2 - \beta_1$) in column 'LSTAR'. In column 'Dummy', D_t is equal to one if DCM_t is above its 75% percentile in period t. In column 'Logistic', D_t is a logistic transformation of DCM_t based on $1/[1 + exp(-\gamma DCM_t)]$, where γ determines the steepness of the function. In column 'LSTAR' the table shows results from a smooth transition regression (LSTAR) of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

To hone some further intuition for the LSTAR model, we plot the resulting time path of the fitted regime function $G(DCM_t)$ for VLOOP and TCOST in Panels A and B of Figure 4. Except for the first two years of the sample, the fitted $G(DCM_t)$ is mostly close to 0 (it is less than 0.25 on 57% and 72% of the days in the sample for VLOOP and TCOST, respectively) and occasionally increases above 0.75 (36% and 22% of the days for VLOOP and TCOST, respectively). These upward spikes are particularly pronounced during the European sovereign debt crisis through 2014/15, uncertainty around Brexit and US elections in 2016, and also the Covid-19 period in 2020. Hence, the unconstrained regime (when β_1 is the effective slope coefficient) corresponds to a normal market situation, while the constrained regime (when the effective slope coefficient is close to β_2) represents stressed periods when dealers face constraints on their intermediation capacity.

Besides the LSTAR approach, we also rely on an alternative methodology in the robustness section of the paper. Specifically, we estimate the correlation between each of our two liquidity cost measures (i.e., VLOOP and TCOST) and dealer-intermediated volume in a rolling window fashion and regress it on DCM. Consistent with our baseline results, we find that the rolling correlation between VLOOP or TCOST and volume decreases as DCM increases. Moreover, we extract exogenous liquidity demand and supply shocks from a structural vector autoregression inspired by Goldberg (2020) and Goldberg and Nozawa (2020). We use the demand shocks as an instrument for DCM and the supply shocks as an alternative measure of tightening dealer constraints (see Section 6 for further details).

Figure 4: Time-series of fitted regime function $G(\mathrm{DCM})$

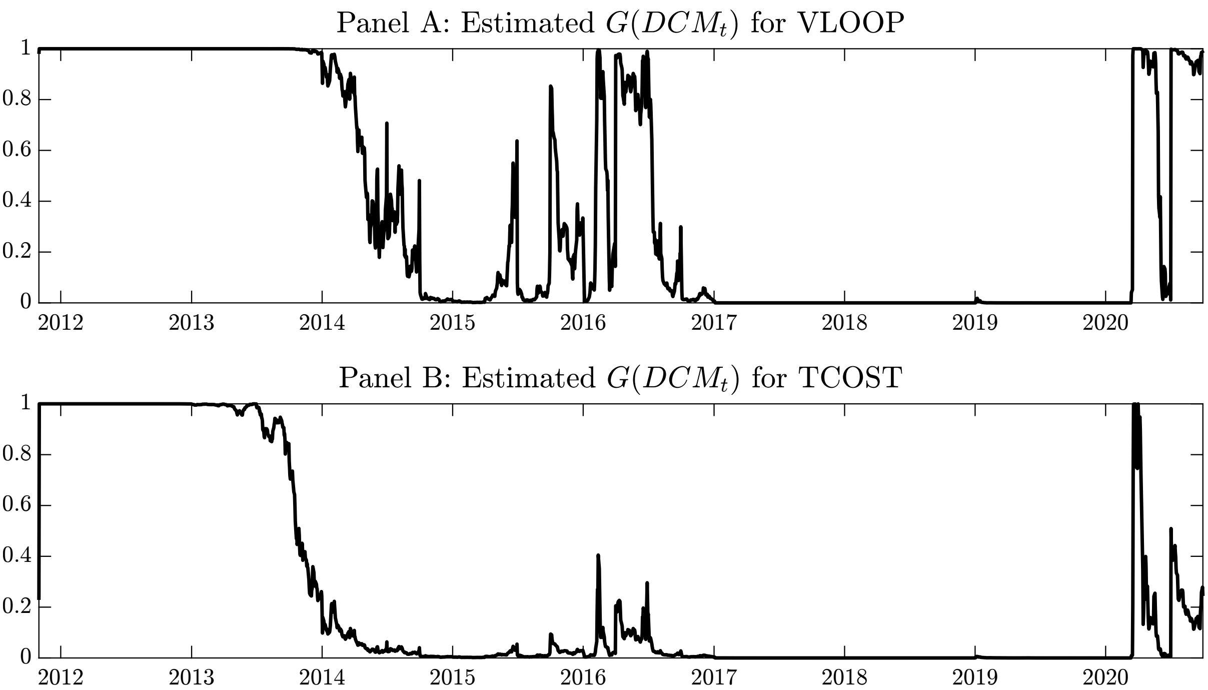

Note: Panel A of this figure shows the fitted regime function $G(DCM_t)$ for VLOOP using the point estimates in column 3 of Table 3, whereas Panel B shows the fitted regime function $G(DCM_t)$ for TCOST using the point estimates in column 6 of Table 3. The sample covers the period from 1 November 2011 to 30 September 2020.

5. A simple model of constrained liquidity supply

This section presents a static partial equilibrium model that rationalises the two main empirical findings from the previous section:

1. The cost of liquidity provision (i.e., VLOOP and also TCOST) is higher when FX dealers are more constrained (see Table 2).

2. Dealer-intermediated volume and liquidity provision costs co-move when dealers are unconstrained but the positive correlation decreases in dealer constraints (see Table 3).

The model features two periods ($t = 0, 1$), three currency pairs (forming a triplet of currencies) and two types of agents: liquidity traders and one representative dealer. At $t = 0$, liquidity traders arrive and trade with the dealer. At $t = 1$, the uncertainty is resolved.

5.1. *Trading environment*

FX spot contracts. Let x, y, and z denote the three currency pairs, for instance, EURCAD, USDEUR, and USDCAD, respectively. Let $\mathbf{p}$ denote the exchange rates of the three currency pairs $[p^x, p^y, p^z]^T$ at $t = 0$. For instance, for currency pair x, 1 EUR = p^x CAD. More specifically, let a^j and b^j denote the ask and bid price of currency pair j, and m^j and s^j denote the mid-quote and the bid-ask spread of currency pair j where $j \in \{x, y, z\}$.

The agents trading these three currency pairs at $t = 1$ receive their fundamental value. Note that in terms of fundamental value there is no difference between the direct FX rate (i.e., EURCAD) and the synthetic rate computed with two indirect rates (i.e., first trading USDEUR and then USDCAD). We denote the fundamental value of the three currency pairs as $\tilde{\mathbf{e}} = [\tilde{e}^x, \tilde{e}^y, \tilde{e}^z]^T$. The mean and variance of the fundamental value are $\mathbf{e} = [e^x, e^y, e^z]^T$ and $\boldsymbol{\sigma} = [\sigma, \sigma, \sigma]^T$. Note that the three fundamental values are intimately linked via $e^x = e^y e^z$.

Liquidity traders. We model liquidity demand in reduced form, following the classic market microstructure literature (see, e.g., Grossman and Miller, 1988; Hendershott and Menkveld, 2014). At $t = 0$, the demand for a currency pair is given by $L = \lambda\sigma(1 - s)$, where λ is a tuning parameter that captures the slope of the demand function. Hence, L increases in volatility σ but decreases in the bid-ask spread s quoted by the dealer.

Trading demand is imbalanced across the three currency pairs due to diverging private values among market participants (i.e., disagreement) following the spirit of Gabaix and Maggiori (2015).[21] For simplicity, we assume that a $\pi > 1/2$ fraction of the liquidity traders in currency pair x are buyers and the rest are sellers. Conversely, for currency pair y, a $1 - \pi$ fraction of liquidity traders are sellers and the rest are buyers. For currency pair z, half of the traders are buyers, whereas the other half are sellers. Thus, the liquidity traders impose net buying pressure $(2\pi - 1)$ in currency pair x and net selling pressure $(1 - 2\pi)$ in currency pair y. As a result, the liquidity traders' demand imbalance (i.e., the net buying/selling pressure)[22]

[21] To account for the effect of disagreement, we have explored regression specifications including common high-frequency measures of disagreement. For instance, we consider the dispersion of order flows of corporates, funds, non-bank financials, and banks (Cespa et al., 2021) as well as the volume-volatility ratio (Liu and Tsyvinski, 2020) as control variables and have found that our key empirical results remain qualitatively unchanged. These additional findings are available upon request.

[22] As the empirical results suggest in the previous section, the violation of the law of one price does not directly imply that there are profitable triangular arbitrage opportunities if transaction costs in the form of bid-ask spreads are sufficiently large. Hence, for simplicity, in the model we abstract away from any cross-market arbitrageurs.

in each of the three currency pairs is given as

$$\mathbf{d} = \lambda\sigma(1-s) \times [2\pi - 1, 1 - 2\pi, 0]^T. \tag{6}$$

Dealer. There is a representative and risk-averse dealer (see Sec. 3.5 Foucault, Pagano, and Roell, 2013) who makes the market at $t = 0$. Being competitive and starting with zero inventory, the dealer decides on their positions $\mathbf{q}$ in the three currency pairs by taking the market-clearing prices $\mathbf{p}$ as given. Note that $q^j > 0$ means the dealer sells the quote and buys the base currency in currency pair j. The dealer is assumed to use debt to finance their market-making activities (e.g., Scott, 1976; van Binsbergen, Graham, and Yang, 2010). The cost of market making increases in the per-unit debt funding cost η and has two components: an inventory cost and a Value-at-Risk component. The inventory cost is proportional to the net positions in the three currency pairs: $\mathbb{1}^T|\mathbf{q}|$, where $\mathbb{1} = [1,1,1]^T$. The Value-at-Risk of the currency positions is $\mathbf{q^T\Sigma q}$ (Duffie and Pan, 1997; Adrian and Shin, 2013), where $\mathbf{\Sigma}$ is the covariance matrix across the triplet of currency pairs.[23] The utility of the dealer is given as

$$U^D = E\Big(\mathbf{p^T q} - \mathbf{e^T q}\Big) - \eta\Big(\underbrace{\mathbb{1}^T|\mathbf{q}|}_{\text{Inventory}} + \underbrace{\mathbf{q^T\Sigma q}}_{\text{Value-at-Risk}}\Big) \tag{7}$$

where the second term in the utility function captures the empirical constraints examined in Section 4.[24] In terms of the link with our empirical analysis, η reflects constraints related to leverage and funding costs (as reflected in their bond financing costs and CDS premia on the issued debt), while σ (embedded in $\mathbf{q^T\Sigma q}$) captures the Value-at-Risk constraint.

Market clearing. The market clearing condition is the following: At $t = 0$, liquidity traders' demand $\mathbf{d}$ for immediacy must be equal to the dealer's liquidity supply $\mathbf{q}$ (i.e., $\mathbf{d} = \mathbf{q}$).

5.2. Equilibrium outcomes

At $t = 0$, the supply function of the dealer is pinned down by their first order condition:

$$\frac{\partial U^D}{\partial q^j} = \begin{cases} \underbrace{a^j - \eta - e^j}_{\text{marginal value of selling}} - \underbrace{2\eta\sigma^2 q^j}_{\text{price impact}} & \text{if } q^j > 0, \\ \underbrace{b^j + \eta - e^j}_{\text{marginal value of buying}} - \underbrace{2\eta\sigma^2 q^j}_{\text{price impact}} & \text{if } q^j < 0. \end{cases} \tag{8}$$

[23] For tractability, we assume that the three currency pairs are i.i.d. and hence the correlations among them are zero. In principle, however, the three currency pairs are closely tied together. Somogyi (2021) discusses a set of realistic assumptions that allows for a non-singular covariance matrix for a triplet of currency pairs.

[24] "Constraints" in our context refer to impaired intermediation capacity due to higher leverage, funding cost, or Value-at-Risk and hence do not necessarily imply (regulatory) binding restrictions.

The first order condition suggests that there are two components in the dealer's supply function. The first one is related to the marginal valuation of buying and selling and reflects the shadow cost of intermediary constraints (Adrian et al., 2014; Duffie, 2018). The second component is the price impact that depends on the size and the direction of the order flow, which stems from the Value-at-Risk constraint. Thus, the supply function is given by

$$q^j = \begin{cases} \frac{a^j - \eta - e^j}{2\eta\sigma^2} & \text{if } q^j > 0, \\ \frac{b^j + \eta - e^j}{2\eta\sigma^2} & \text{if } q^j < 0. \end{cases} \tag{9}$$

Facing the liquidity traders' demand, the bid and ask prices for the three currency pairs are pinned down by the following six market clearing conditions (see Eqs (10) to (12)). There are two equations for each currency pair: one for the case when the dealer is buying (i.e., bid price) and another reflecting the situation when the dealer is selling (i.e., ask price). Taking the first condition as an example, the left-hand side is the amount sold by the liquidity traders and the right-hand side is the amount bought from the dealer in currency pair x. Eventually, in equilibrium, the bid price of currency pair x is determined by market clearing:

$$-\lambda\sigma(1-s^x)(1-\pi) = \frac{b^x + \eta - e^x}{2\eta\sigma^2}, \qquad \lambda\sigma(1-s^x)\pi = \frac{a^x - \eta - e^x}{2\eta\sigma^2}; \tag{10}$$

$$-\lambda\sigma(1-s^y)\pi = \frac{b^y + \eta - e^y}{2\eta\sigma^2}, \qquad \lambda\sigma(1-s^y)(1-\pi) = \frac{a^y - \eta - e^y}{2\eta\sigma^2}; \tag{11}$$

$$-\frac{1}{2}\lambda\sigma(1-s^z) = \frac{b^z + \eta - e^z}{2\eta\sigma^2}, \qquad \frac{1}{2}\lambda\sigma(1-s^z) = \frac{a^z - \eta - e^z}{2\eta\sigma^2}. \tag{12}$$

Solving the system of equations (i.e., subtracting both sides of the equations on the left from those on the right in each line of Eqs (10) to (12)), the bid-ask spreads for the three currency pairs turn out to be the same. The intuition for this hinges on the simplifying assumptions that the dealers' debt financing cost as well as the volatility of fundamental values are the same across the three currency pairs.[25] Therefore, in this setup, the bid-ask spread s is given by the following expression:

$$\mathrm{s} = \frac{2\eta(1+\lambda\sigma^3)}{1+2\eta\lambda\sigma^3}. \tag{13}$$

Thus, we can express TCOST, which is equal to three times the half bid-ask spread (since financing cost, leverage, and volatility are symmetric across currency pairs) as follows:

$$TCOST = \frac{3s}{2} = \frac{3\eta(1+\lambda\sigma^3)}{1+2\eta\lambda\sigma^3}. \tag{14}$$

[25] Relaxing the assumption of homogeneous volatility across currency pairs (i.e., having currency pair specific volatility) does not change the results qualitatively because the market clearing conditions Eqs (10) to (12) are also currency pair specific.

One can see that TCOST captures the dealer's compensation to endure order flow imbalances due to the clients' demand for immediacy. Substituting s into the market clearing conditions yields the following expression for the mid-quote prices in currency pairs x, y, and z:

$$m^x = e^x + \frac{\eta\lambda\sigma^3\,(2\pi-1)\,(1-2\eta)}{1+2\eta\lambda\sigma^3}, \quad m^y = e^y + \frac{\eta\lambda\sigma^3\,(1-2\pi)\,(1-2\eta)}{1+2\eta\lambda\sigma^3}, \quad m^z = e^z. \tag{15}$$

The mid-quotes of currency pair x and y in Eq. (15) deviate from their fundamental values e^x and e^y, respectively, if $\pi \neq \frac{1}{2}$ or $\eta \neq \frac{1}{2}$. On the contrary, the mid-quote for currency pair z is equal to its fundamental value (i.e., $m^z = e^z$).[26] The effect of volatility and the dealer constraints on the mid-quotes are directional and depend on the order imbalance across the three currency pairs. Take currency pair x as an example. The buying pressure from the liquidity traders dominates their selling pressure, which creates an inventory imbalance for the dealer. Maintaining such an inventory imbalance is costly and hence the dealer adjusts their marginal valuation for currency pair x by pushing up the mid-quote price above its fundamental value. In other words, the dealer charges a mark-up reflecting the shadow cost of balance sheet constraints. Thus, the deviations of the mid-quotes (set by the dealer) from the fundamental values represent violations of the law of one price:

$$\text{VLOOP} = m^x - m^y m^z = \lambda\sigma^3(2\pi-1)(1+e^z)\frac{\eta(1-2\eta)}{1+2\eta\lambda\sigma^3}. \tag{16}$$

Note that such deviations from the law of one price are not necessarily profitable arbitrage opportunities due to non-zero transaction costs (i.e., TCOST), which define the so-called arbitrage bounds (Shleifer and Vishny, 1997).[27]

5.3. *Comparative statics*

Next, we use the empirical results in Section 4 to identify the parameter space of interest of the model. In turn, we investigate the different channels in the model and then use comparative statics analysis to support the empirical results.

Taking the first order derivatives of TCOST with respect to η and σ, from Eq. (13) we have

$$\frac{\partial \text{TCOST}}{\partial \eta} = \frac{3(1+\lambda\sigma^3+2\eta\lambda\sigma^3)}{(1+2\eta\lambda\sigma^3)^2} > 0, \tag{17}$$

[26] Note that this choice is for simplicity in the sense that with a slightly different setup on demand imbalances we could also have that $m^z \neq e^z$. In general, as long as the demand imbalance across the three currency pairs is not exactly the same, the size of the deviations of the mid-quotes from the fundamental values is different across the three pairs, generating violations of the law of one price. For brevity, we focus on one particular example.

[27] As indicated in Footnote 22, the model focuses on the cases when there are no actually profitable arbitrage opportunities and is hence consistent with the empirical evidence in Section 4.

$$\frac{\partial \text{TCOST}}{\partial \sigma} = \frac{9\eta\lambda\sigma^2 \overbrace{(1-2\eta)}^{>0 \text{ as } \eta<1/2}}{(1+2\eta\lambda\sigma^3)^2} > 0. \tag{18}$$

Hence, TCOST increases when the dealer is more constrained (i.e., higher η). Moreover, as dealer-provided volume is proportional to $(1-s)$, the model implicitly assumes that $s < 1$, since volume cannot be negative. Thus, Eq. (13) implies that $\eta < 1/2$ (see the discussion above). In this case, TCOST increases in volatility σ, which is the key determinant for the Value-at-Risk measure. Hence, when the dealer's VaR constraints become more binding, the bid-ask spread is also higher. The dashed line in Figure 5 visualises these effects.

Figure 5: The effects of volatility and the dealer's constraint on VLOOP and TCOST

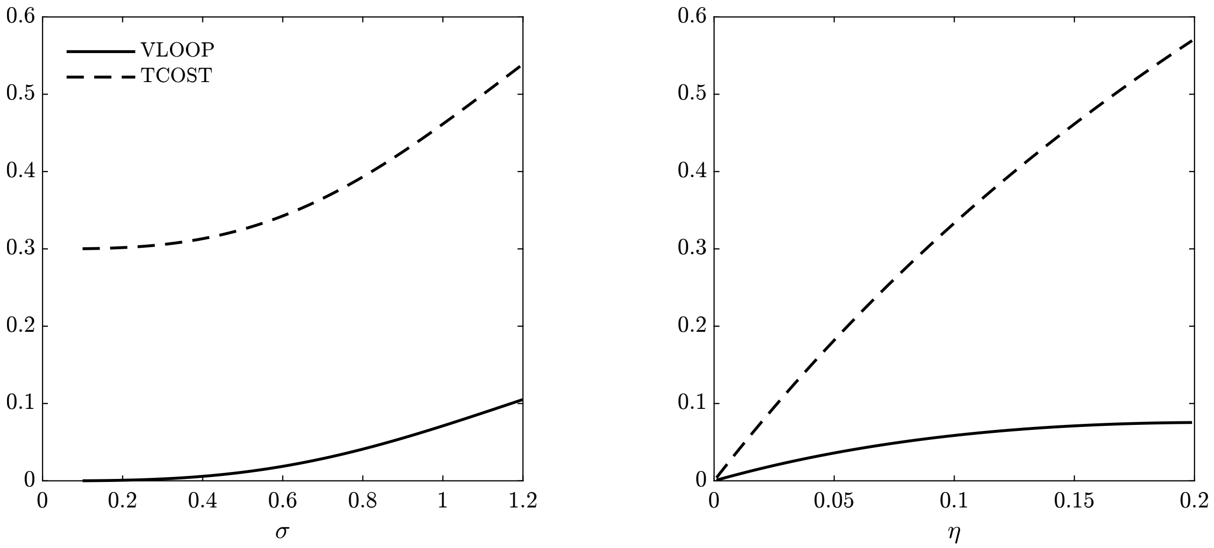

Note: This figure shows how dealer constraints η and volatility σ affect VLOOP and TCOST, respectively. The baseline parameters are $\pi = 0.7$, $\lambda = 1$, $\sigma = 1$, $\eta = 0.15$, $e^x = 1.32$, $e^y = 1.1$, $e^z = 1.2$, where π denotes the fraction of liquidity traders that are buyers (sellers) in currency pair x (y), λ is the trading demand parameter, σ is the fundamental volatility of exchange rates, η is the per-unit debt funding cost, whereas e^x, e^y, and e^z denote the fundamental values of currency pairs x, y, and z, respectively.

Taking the first order derivative of VLOOP in Eq. (16) with respect to η and σ, we get

$$\frac{\partial \text{VLOOP}}{\partial \eta} = \frac{\lambda\sigma^3(2\pi-1)(1+e^z)}{(1+2\eta\lambda\sigma^3)}\left(1-s-2\eta\right), \tag{19}$$

$$\frac{\partial \text{VLOOP}}{\partial \sigma} = (2\pi-1)(1+e^z)\frac{3\eta\lambda\sigma^2(1-2\eta)}{(1+2\eta\lambda\sigma^3)^2} > 0. \tag{20}$$

The impact of the dealer's constraint on VLOOP is more complex due to two offsetting effects. On the one hand, a constrained dealer (i.e., facing higher η or σ) charges a higher mark-up (or mark-down) for currency pair x (or y), which increases VLOOP. On the other hand, a higher η or σ also leads to a higher TCOST as shown in Eq. (17) and Eq. (18). The higher trading cost suppresses additional trading demands and renders order flows less imbalanced, which subsequently dampens VLOOP. In the case of σ, the first channel always dominates the second and VLOOP increases monotonically in σ (see eq. (20)). Contrarily, VLOOP increases in η only if $\eta < (1-s)/2$ as shown in Eq. (19). In our sample period (see Table 2) VLOOP is larger when FX dealers are more constrained, suggesting that empirically the first channel dominates the second one. Moreover, VLOOP increases relatively less compared to TCOST as dealer constraints tighten (see Table 2), which is consistent with these two offsetting forces. Hence, the analysis below focuses on the case where $\eta < (1-s)/2$. The solid line in Figure 5 visualises these effects. Proposition 1 summarises the key predictions.

Proposition 1: *Both VLOOP and TCOST are higher when*

i) the volatility of the currency pairs is higher and the (representative) dealer faces a more stringent Value-at-Risk constraint (i.e., higher σ);

ii) the dealer is more constrained (i.e., higher η) due to higher leverage and funding constraints.

5.4. *Elasticity of liquidity provision*

One of the key empirical finding in Section 4 is that the positive correlation between the cost of liquidity provision and dealer-intermediated volume decreases when dealers are more constrained. Through the lenses of our model, we refer to this correlation as the "elasticity of liquidity provision." This is because trading volume, VLOOP, and TCOST are equilibrium outcomes and the corresponding concept of "elasticity" describes how equilibrium volume co-moves with equilibrium VLOOP and TCOST, respectively, when key structural parameters such as fundamental volatility or dealer constraints are changing.

In particular, dealer-intermediated volume is given by[28]

$$VLM = 3L = 3\lambda\sigma(1-s) = \frac{3\lambda\sigma(1-2\eta)}{1+2\eta\lambda\sigma^3}. \tag{21}$$

Taking the first order derivative with respect to volatility σ, we have that

$$\frac{\partial VLM}{\partial \sigma} = 3\lambda(1-2\eta)\frac{(1-4\eta\lambda\sigma^3)}{(1+2\eta\lambda\sigma^3)^2}. \tag{22}$$

Thus, dealer-intermediated volume increases in volatility only if $\lambda < 1/(4\eta\sigma^3)$. Intuitively, as discussed above, an increase in volatility affects volume via two channels: First, trading demand increases due to a larger dispersion in fundamentals. Second, trading demand is suppressed due to the concurrent rise in the bid-ask spread. When the trading demand parameter λ is small, the former dominates the latter. As shown in Table 2, dealer-intermediated volume increases in volatility, indicating that the parameter space of interest is indeed $\lambda < 1/(4\eta\sigma^3)$. Thus, for the rest of this section, we only focus on the latter case.

[28]Note that the scalar 3 comes from the fact that each currency pair triplet comprises three currency pairs.

From Eq. (22), it is clear that when η is larger, the first order derivative is smaller (yet still positive). Put differently, dealer-intermediated volume (i.e., VLM) increases in volatility (i.e., σ) at a slower pace as the dealer becomes more constrained. Lemma 1 summarises this result.

Lemma 1: *When the dealer is more constrained, volume increases in volatility at a slower pace.*

Both TCOST and VLOOP increase in volatility σ for a given level of the dealer's constraint η (see Eqs (18) and (20)). Specifically, an increase in volatility is associated with a rise in equilibrium volume (i.e., VLM) as well as an increase in both TCOST and VLOOP, respectively. To translate this co-movement into an elasticity of liquidity supply we take the ratio of these two partial derivatives with respect to volatility σ:

$$\frac{\partial VLM}{\partial \sigma} / \frac{\partial TCOST}{\partial \sigma} = \frac{(1-4\eta\lambda\sigma^3)}{3\eta\sigma^2}, \qquad \frac{\partial VLM}{\partial \sigma} / \frac{\partial VLOOP}{\partial \sigma} = \frac{(1-4\eta\lambda\sigma^3)}{\eta\sigma^2(2\pi-1)(1+e^z)}, \tag{23}$$

where the partial derivatives are capturing changes in volume (VLM) together with TCOST (left-hand side of Eq. (23)) as well as changes in VLM and VLOOP (right-hand side of Eq. (23)), respectively. In essence, Eq. (23) captures how the co-movement between dealer-provided volume and the two liquidity cost measures changes conditional on the level of the dealer's constraint η. In our empirical analysis, we capture such changes in the co-movement of volume and liquidity costs via state-dependent regression analyses (see Section 4).

Figure 6 visualises the elasticity of liquidity provision across constrained and unconstrained regimes. The solid lines plot equilibrium volume against the two liquidity cost measures conditional on changing volatility σ. The grey lines indicate unconstrained states with low η, while the black lines refer to constrained states with high η. The grey dotted lines correspond to the unconstrained (and low volatility) periods in the empirical analysis, whereas the black dotted lines depict the constrained (and high volatility) periods.

When the dealer is unconstrained (i.e., small η) the two liquidity cost measures are small as well (see grey lines). Hence, in unconstrained periods liquidity provision is elastic since both variables are positively correlated conditional on different realisations of volatility. However, when the dealer is more constrained, the line shifts towards the upper right (see black line). In this new equilibrium, trading volume is only marginally higher, whereas both liquidity cost measures increase substantially. The grey and black dashed tangent lines capture the elasticity of liquidity provision across the two market states. The slope of the black dashed line is much steeper compared to the grey line, suggesting that liquidity provision is less elastic when the dealer is more constrained. This numerical result is fully consistent with the fact that both derivatives in Eq. (23) are decreasing in η. The underlying mechanism for this finding evolves along two dimensions: First, the constrained dealer charges a higher mark-up (or mark-down) relative to the midquote resulting into a higher VLOOP. Second, the constrained dealer also charges a higher bid-ask spread which directly translates to a higher TCOST. Taken together, the increase in both VLOOP and TCOST leads to a slower increase in equilibrium trading volume. Proposition 2 summarises these results.

Proposition 2: *Dealer-intermediated volume co-moves with both VLOOP and TCOST. However, this co-movement, which captures the elasticity of liquidity provision, weakens when the dealer is more constrained due to higher leverage and/or funding costs (i.e., higher η).*

Figure 6: The elasticity of liquidity provision

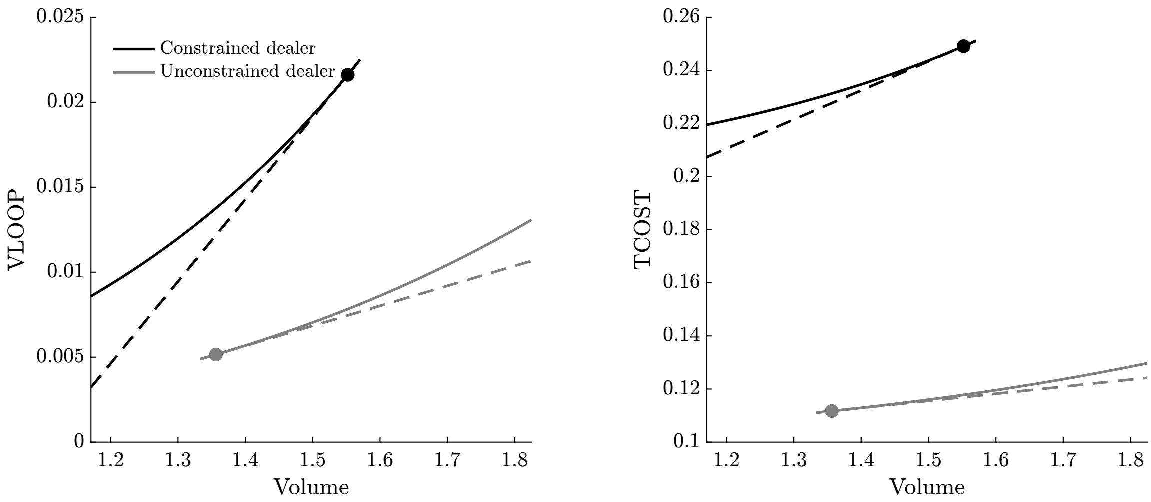

Note: This figure plots dealer-intermediated volume against VLOOP and TCOST, respectively. The baseline parameters are $\pi = 0.7$, $\lambda = 1$, $e^x = 1.32$, $e^y = 1.1$, $e^z = 1.2$, where π denotes the fraction of liquidity traders that are buyers (sellers) in currency pair x (y), λ is the trading demand parameter, whereas e^x, e^y, and e^z denote the fundamental values of currency pairs x, y, and z, respectively. The unconstrained dealer faces $\eta = 0.05$, whereas the constrained one is exposed to $\eta = 0.1$. The solid lines indicate the equilibrium outcomes when varying the volatility of the exchange rates σ from 0.5 to 0.7. The grey dashed line indicates the derivative of volume with respect to VLOOP and TCOST when the dealer is unconstrained and volatility is low. The black dashed line is the derivative of volume with respect to VLOOP and TCOST when the dealer is constrained and volatility is high.

Proposition 2 finds compelling support in our empirical analyses. Specifically, Table 2 shows that dealer-intermediated volume increases even when dealers are more constrained. However, the state-dependent relation between intermediated volume and the cost of liquidity provision weakens when dealer constraints intensify. This result holds for both liquidity cost measures (VLOOP and TCOST) no matter which econometric model is applied (Table 3) and after controlling for volatility and other confounding factors (see Online Appendix).

6. Additional analyses and robustness tests

The key goal of our empirical analysis is to show how the elasticity of dealer banks' liquidity provision (i.e., the correlation between the cost of liquidity provision and dealer-intermediated volume) weakens as dealer constraints tighten. To demonstrate this we have relied on logistic smooth transition regressions (LSTAR) that are particularly well-suited to capture nonlinear relations. However, one might wonder whether our results are robust to using alternative methods and measures of dealer constraints. To address this issue, we employ rolling window correlations and structural vector autoregressions to infer exogenous liquidity demand and supply shocks directly from price and quantity data.

6.1. *Disentangling liquidity demand and supply*

The empirical analysis based on LSTAR has employed our dealer constraint measure (i.e., DCM) as an exogenous proxy for liquidity supply shocks. Here, we take the analysis one step further by explicitly disentangling liquidity demand and supply shocks using a structural vector autoregression with sign restrictions. Specifically, we build on the approach by Uhlig (2005) and others (e.g., Canova and Nicoló, 2002; Rubio-Ramírez, Waggoner, and Zha, 2010), which has become widely used in economics and finance to estimate models with sign restrictions. Eventually, we are using the supply shocks as an alternative measure for tightening dealer constraints. Moreover, we employ demand shocks as an instrument for our dealer constraint measure. The economic intuition following our theoretical framework in Section 5 is that a positive liquidity demand shock causes more imbalanced customer order flows to which dealers respond by increasing the bid-ask spread to dampen volumes. In addition, concurrently, no-arbitrage deviations also increase due to the customer order flows being more imbalanced. Hence, demand shocks can have an impact via dealer constraints on the relation between trading volume and our two liquidity cost measures (i.e., VLOOP and TCOST). Our empirical analysis proceeds in two steps.

In a first step, we estimate a structural (bivariate) vector autoregression (SVAR) model of liquidity cost measures (i.e., VLOOP or TCOST) and dealer-provided volume VLM. To identify liquidity supply and demand shifts, we estimate the SVAR imposing sign restrictions in the spirit of Cohen, Diether, and Malloy (2007), Goldberg (2020), and Goldberg and Nozawa (2020), respectively, using Bayesian methods (see the Online Appendix Section C for a detailed explanation of the setup and the estimation procedure). In particular, the sign restrictions assume that supply shifts lead to changes in liquidity costs and trading volume that have opposite signs (i.e, a fall in dealer-intermediated trading volume corresponds to an increase in liquidity costs). Demand shifts, by contrasts, are assumed to lead to changes in liquidity costs and volume that have the same sign (i.e., higher intermediated trading volume goes in hand with a deterioration of liquidity costs).

In a second step, we estimate the correlation between the cost of liquidity provision (i.e., VLOOP or TCOST) and dealer-intermediated trading volume (i.e., VLM) in a 30-day rolling window[29] fashion (cf. Figure C.3) and estimate the following panel regression model:

$$\rho_{k,t} = \alpha_k + \eta_1 DCM_t + \eta_2 RV_{k,t} + \eta_3 Amihud_{k,t} + \epsilon_{k,t}, \tag{24}$$

where the dependent variable is the 30-day rolling window correlation of a liquidity cost measure (i.e., VLOOP or TCOST) and trading volume, α_k denotes currency triplet fixed ef-

[29] All our results are qualitatively unchanged when using longer or shorter estimation windows.

fects, $RV_{k,t}$ ($Amihud_{k,t}$) the realised variance (Amihud (2002) price impact) in the non-dollar currency pair within each triplet k, and DCM_t is our dealer constraint measure. Throughout this paper, we estimate Amihud as the ratio between daily realised volatility and aggregate daily trading volume following Ranaldo and Santucci de Magistris (forthcoming 2022).

Table 4 documents the results of estimating Eq. (24) by OLS and 2SLS, respectively. In particular, Panel A shows the OLS estimates of Eq. (24), whereas Panel B uses liquidity demand shocks $\delta^d_{k,t}$ from the SVAR as an instrument for DCM.[30] The (unreported) first stage coefficients are highly significant (with an F-statistic well above 10) suggesting that we have identified an economically relevant instrument for our dealer constraint measure DCM. Panel C reports the results of using liquidity supply shocks $\delta^s_{k,t}$ as an alternative measure of tightening dealer constraints.[31] The key takeaway from Table 4 is fully consistent with the LSTAR analysis (see Table 3) and corroborates the idea that soaring dealer constraints are associated with a significantly lower elasticity of liquidity provision (i.e., smaller $\rho_{k,t}$).

6.2. Robustness tests

To investigate the robustness of our main findings we run five additional empirical tests: i) decompose the dealer constraint measure into its constituents, ii) split volume into inter-bank and customer-bank trades, iii) perform a subsample analysis, iv) estimate the LSTAR currency pair triplet by triplet, and v) account for potential bias in the bid-ask spread.

Different components of dealer constraints. We consider the same LSTAR specification as in Eq. (4) but instead of our dealer constraint measure DCM we use its four constituents. In particular, we use the 1-day lagged value of primary FX dealer banks' quarterly Value-at-Risk measure (VaR), quarterly He et al. (2017) leverage ratio (HKM), daily credit default spread (CDS), and daily funding cost yield (DFC) as regime variables. Table 5 reports the estimates of using each of the four aforementioned measures as a state variable. The difference between the constrained and unconstrained coefficients is negative and significant across all four specifications for both VLOOP and TCOST. These estimates are fully in line with our baseline specification based on DCM in terms of economic magnitudes. The robustness of our results is not surprising given the strong co-movement across these four different regime variables.

Inter-bank vs customer-bank volumes. We decompose trading volume into to inter-bank and customer-bank volume to better understand which market segments suffer the most from

[30] We estimate demand and supply shocks individually for every currency pair triplet and then stack them together. Our findings are robust to extracting the shocks from a panel SVAR with currency triplet fixed effects.

[31] In the Online Appendix we also exploit the unexpected removal of the Swiss franc cap on 15 January 2015 by the Swiss National Bank as a quasi-natural experiment. We find that the elasticity of dealer banks' liquidity provision drops significantly in currency pair triplets involving the Swiss franc (i.e., EUR-USD-CHF and GBP-USD-CHF) but not in other triplets. These findings support the idea that dealers face currency (pair) specific risk limits. See the Online Appendix Section D for additional information.

Table 4: Elasticity of liquidity provision and dealer constraints

	cor(VLOOP,VLM)			cor(TCOST,VLM)		
Panel A	(1)	(2)	(3)	(4)	(5)	(6)
DCM	***−0.03	***−0.03	***−0.03	***−0.04	***−0.04	***−0.03
	[3.50]	[3.51]	[3.57]	[2.92]	[2.92]	[2.77]
Realised variance		*0.00	0.00		−0.01	−0.01
		[1.70]	[1.17]		[1.47]	[0.93]
Amihud (2002)			***0.03			***−0.08
			[3.31]			[5.14]
Adj. R^2 in %	1.65	1.70	1.87	1.76	1.93	2.80
Avg. #Time periods	2,159	2,159	2,159	2,159	2,159	2,159
Panel B						
Instrumented DCM	***−0.07	***−0.08	***−0.07	*−0.06	*−0.06	**−0.07
	[3.62]	[3.75]	[3.60]	[1.76]	[1.77]	[2.14]
Realised variance		0.01	0.00		−0.01	−0.01
		[1.57]	[1.26]		[1.47]	[0.89]
Amihud (2002)			**0.02			***−0.09
			[2.39]			[5.54]
Avg. #Time periods	2,159	2,159	2,159	2,159	2,159	2,159
Panel C						
δ^s	***−0.01	***−0.01	***−0.01	***−0.02	***−0.02	***−0.02
	[3.64]	[3.61]	[3.76]	[4.28]	[4.75]	[4.10]
Realised variance		*0.00	*0.00		−0.01	−0.01
		[1.89]	[1.89]		[1.51]	[1.45]
Amihud (2002)			0.00			−0.01
			[0.02]			[1.46]
Adj. R^2 in %	0.10	0.13	0.13	0.38	0.62	0.77
Avg. #Time periods	2,256	2,256	2,256	2,256	2,256	2,256
#Currency triplets	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes
Time series FE	no	no	no	no	no	no

Note: This table reports results from daily fixed effects panel regressions of the form $\rho_{k,t} = \alpha_k + \eta_1 DCM_t + \eta_2 RV_{k,t} + \eta_3 Amihud_{k,t} + \epsilon_{k,t}$, where the dependent variable is the 30-day rolling window correlation of a liquidity cost measure (i.e., $VLOOP$, or $TCOST$) and trading volume (i.e., VLM), α_k denotes cross-sectional fixed effects, $RV_{k,t}$ ($Amihud_{k,t}$) the realised variance (Amihud (2002) price impact) in the non-dollar currency pair within each triplet k, and DCM_t is our dealer constraint measure. Panel A shows the OLS estimates of Eq. (24), whereas Panel B uses liquidity demand shocks δ^d as an instrument for DCM. Panel C reports the results of using liquidity supply shocks δ^s as an alternative measure of tightening dealer constraints. All regressors have been normalised to have unit standard deviation. Hence, the regression coefficients measure the increase in ρ associated with a one standard deviation increase in DCM and δ^s, respectively. The sample covers the period from 1 September 2012 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

reduced liquidity provision when dealer constraints tighten. Specifically, the CLS customer-bank order flow data comprise three groups of customers, that is, corporates, funds, and non-bank financials.[32] Note that bilateral trades between two such customer groups are quasi non-existent given the two-tier structure of the FX market (Rime and Schrimpf, 2013) and

[32] See Cespa et al. (2021) and Ranaldo and Somogyi (2021) for a detailed description of the CLS flow data set.

Table 5: Smooth transition regression with different state variables

	VLOOP				TCOST			
	VaR	HKM	CDS	DFC	VaR	HKM	CDS	DFC
γ	***11.15	***4.97	***9.20	***12.10	***12.08	***4.93	***12.03	***3.92
c	***0.25	***0.59	***−0.10	***−0.35	***0.39	***0.60	***−0.16	***−0.10
Unconstr. volume	***0.09	**0.06	***0.08	***0.16	***0.09	***0.09	***0.10	***0.10
	[3.16]	[2.23]	[2.89]	[3.59]	[11.35]	[10.61]	[11.26]	[7.83]
Constr. volume	**−0.08	−0.03	*−0.06	−0.01	*0.02	**0.03	**0.03	***0.04
	[2.12]	[0.73]	[1.66]	[0.42]	[1.65]	[2.45]	[2.31]	[2.99]
Realised variance	**0.02	**0.02	**0.02	**0.02	***0.03	***0.03	***0.03	***0.03
	[2.07]	[2.13]	[2.11]	[2.09]	[8.35]	[8.39]	[8.37]	[8.00]
Constr.-Unconstr.	***−0.17	*−0.10	***−0.15	***−0.18	***−0.07	***−0.05	***−0.07	***−0.06
	[3.50]	[1.68]	[3.00]	[3.04]	[4.28]	[3.28]	[4.83]	[2.82]
R^2 in %	0.13	0.08	0.11	0.15	3.72	3.65	3.77	3.66
BIC	91.61	91.62	91.61	91.06	49.24	49.25	49.23	48.55
Avg. #Time periods	2,280	2,280	2,280	2,182	2,284	2,284	2,284	2,185
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the regime variable z_{t-1}. The regime variables are the 1-day lagged value of primary FX dealer banks': quarterly Value-at-Risk measure (VaR, columns 1 and 6), quarterly He et al. (2017) leverage ratio (HKM, columns 2 and 7), daily credit default spread (CDS, columns 3 and 8), and daily funding cost yield (DFC, columns 4 and 9). Note that we assign an equal weight to each top 10 FX dealer bank (based on the Euromoney FX survey) when computing a cross-sectional average. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

hence also do not form part of the data that CLS provides. As a result, the customer-bank data only contains trades that pass through an FX dealer bank (e.g., Citi Bank or UBS). Moreover, the inter-bank data include trades between two banks that are members of the CLS system. Some of these banks are GSIBs, whereas others include lower-tier banks outside of the main dealer community (e.g., Danske Bank or Commerzbank).

Table 6 reports the results of estimating the LSTAR model in Eq. (4) based on inter-bank and customer-bank volume rather than total volume. To be precise, we define total volume in each client group as the sum of buy and sell volume in a given currency pair. There is an interesting picture that arises: On the one hand, the coefficients related to unconstrained volume of the inter-bank segment are higher than those of the customer-bank segment suggesting a more elastic liquidity provision in the former. On the other hand, the elasticity of liquidity provision weakens significantly with dealer constraints for both inter- and customer-bank trading activity. However, the economic magnitudes of the constrained minus unconstrained coefficients suggest that large dealer banks mainly curtail their liquidity provision in trades with other banks facing similar constraints. Of course, this does not rule out the possibility that dealers charge higher spreads to their customers when they are more constrained.

Table 6: Smooth transition regression with different counterparty groups

	VLOOP				TCOST			
	Non-bank	Non-bank	Bank	Bank	Non-bank	Non-bank	Bank	Bank
γ	20.01	20.02	***20.03	***20.06	20.04	20.08	***14.96	**20.03
c	***1.10	***0.37	***-0.15	***-0.15	0.03	0.03	***0.57	***0.58
Unconstr. volume	**0.03	*0.03	***0.12	***0.11	***0.03	***0.02	***0.12	***0.09
	[2.11]	[1.87]	[3.68]	[3.14]	[6.77]	[5.40]	[13.70]	[9.53]
Constr. volume	-0.03	-0.02	-0.05	-0.06	0.01	0.00	0.02	-0.01
	[1.16]	[0.81]	[1.07]	[1.38]	[1.51]	[0.33]	[1.07]	[0.23]
Realised variance		**0.02		*0.02		***0.04		***0.03
		[2.27]		[1.68]		[9.02]		[7.14]
Constr.-Unconstr.	**-0.06	-0.05	***-0.17	***-0.17	***-0.02	***-0.02	***-0.09	***-0.09
	[2.06]	[1.62]	[2.98]	[2.96]	[2.62]	[2.86]	[3.84]	[3.74]
R^2 in %	0.05	0.10	0.14	0.16	0.54	2.99	2.49	3.90
Avg. #Time periods	1,979	1,979	1,979	1,979	1,983	1,982	1,983	1,982
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the regime variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 September 2012 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Non-bank liquidity providers. To shed some light on the importance of non-bank liquidity providers (e.g., XTX, HC Tech or Jump Trading) we split our sample period into two halves. The first half concerns the time period from November 2011 until May 2016, whereas the second half runs from June 2016 to September 2020. Our sample split is motivated by the fact that XTX enters the top 10 of the Euromoney FX surveys for the first time in 2016. Table 7 documents the same regression specifications as in our baseline in Eq. (4) except for the time periods being different. The key takeaway from comparing the constrained minus unconstrained coefficients across the first and second half of the sample is that the economic magnitudes of the coefficients are almost twice as large for the first half than for the second half. We interpret this as suggestive evidence in favour of the idea that non-bank liquidity providers are much less affected by our dealer constraint measure and are hence able to provide additional liquidity when dealer banks are more constrained.

Table 7: Sample split: Smooth transition regression with DCM as state variable

	11/2011 – 05/2016				06/2016 – 09/2020			
	VLOOP		TCOST		VLOOP		TCOST	
γ	***4.98	***4.96	***6.51	***7.48	12.03	12.06	12.02	12.10
c	***−0.40	***−0.41	***−0.41	***−0.40	*0.47	*0.46	***0.89	***0.90
Unconstr. volume	0.06	0.05	***0.13	***0.11	***0.13	**0.10	***0.13	***0.10
	[1.42]	[1.07]	[10.72]	[8.62]	[3.10]	[2.38]	[10.89]	[7.36]
Constr. volume	*−0.09	**−0.10	*0.03	0.01	0.10	0.08	***0.08	**0.05
	[1.82]	[2.00]	[1.92]	[0.51]	[1.44]	[1.08]	[3.36]	[2.03]
Realised variance		0.01		***0.02		**0.03		***0.03
		[1.20]		[4.44]		[2.27]		[6.39]
Constr.-Unconstr.	**−0.15	**−0.15	***−0.10	***−0.10	−0.03	−0.03	**−0.05	*−0.05
	[2.26]	[2.24]	[5.35]	[5.18]	[0.36]	[0.32]	[2.04]	[1.85]
R^2 in %	0.06	0.08	2.77	3.77	0.17	0.24	2.41	3.82
Avg. #Time periods	1121	1121	1122	1122	1061	1060	1063	1062
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

LSTAR estimates currency pair triplet by triplet. Thus far, we have mainly focused on the time-series dimension of the relation between trading volume and the cost of liquidity provision but have not delved deeper into the cross-section of currency pair triplets. To explore the cross-sectional heterogeneity, we estimate the LSTAR model individually for 15 triplets of currency pairs. We further contrast the result with a simple linear model that does not distinguish between constrained and unconstrained regimes. We report these analyses in the Online Appendix (see Tables B.1 and B.2).

The currency pair triplet by triplet estimates strongly support the idea that intermediary constraints nonlinearly impact the relation between dealer-provided volume and the cost of liquidity provision. In particular, the difference between the parameter estimates of constrained and unconstrained regimes (i.e., $\beta_2 - \beta_1$) is significantly negative for 10 and 9 out of 15 triplets of currency pairs for VLOOP and TCOST, respectively. In line with this finding, the R^2s of these regressions are rather close to the linear model. This is entirely expected, given that the coefficient with respect to trading volume in constrained regimes is close to zero. In sum, both results are consistent with the idea that in calm periods dealers' liquidity provision is elastic supporting FX market liquidity, however it becomes more inelastic when dealer constraints are tightening.

Bias in the bid-ask spread. Hagströmer (2021) shows that the effective bid-ask spread measured relative to the spread midpoint overstates the true bid-ask spread in markets with discrete prices and elastic liquidity demand (e.g., the currency market). To address this issue, we compute both no-arbitrage violations VLOOP and round-trip transaction costs TCOST using the "weighted midpoint" (i.e., m^{wp}) as an alternative measure of the midquote price:

$$m^{wp} = \frac{b \times q^{buys} + a \times q^{sells}}{q^{buys} + q^{sells}}, \tag{25}$$

where b and a are bid and ask prices, respectively, whereas q^{buys} and q^{sells} are the buy and sell volume in a given currency pair. Table 8 shows the results of estimating the same regression specifications as in our baseline in Eq. (4), but using m^{wp} instead of the spread midpoint m to compute VLOOP and TCOST. The difference between the constrained and unconstrained coefficient on intermediated-trading volume is negative and economically significant for both VLOOP and TCOST across most specifications. Thus, we conclude that our findings are not materially affected by any potential bias in the quoted bid-ask spread in the Olsen data.

To summarise, these additional robustness tests corroborate our previous results and support the main mechanisms of our model. Dealers promote FX market liquidity in normal times through elastic liquidity provision. As such, dealer intermediation contributes to better market liquidity, that is, narrower spreads and more informative prices (i.e., lower transactions costs and tight no-arbitrage conditions). However, during periods of market stress FX dealer banks are more constrained and as a result their intermediation activities cannot keep up with the deterioration of market liquidity.

7. Conclusion

In this paper, we have studied whether dealer constraints have adverse implications on market liquidity. Using a unique data set of prices and volumes in the FX market, we provide a novel analytical method to identify and measure the cost of liquidity provision, and its main components: the shadow cost of intermediary constraints and dealers' realised compensation for enduring inventory imbalances. Equipped with these two measures, we show that at times when dealers' intermediation capacity is constrained (e.g., due to higher leverage ratios, Value-at-Risk measures or debt funding costs) their cost of liquidity provision increases disproportionately. As a result, the elasticity of dealer banks' liquidity provision weakens by at least 80% relative to periods when they are unconstrained. In other words, during such constrained regimes dealers' supply of liquidity is insufficient to curb the deterioration in liquidity conditions. We rationalise our findings with a theoretical model outlining how liquidity costs may deteriorate when markets are more volatile and when financial intermediaries are more constrained.

We obtain our results for the FX spot market, which is commonly regarded as one of

Table 8: Weighted midpoint: Smooth transition regression with DCM as state variable

	VLOOP				TCOST			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	**4.97	***4.99	***4.91	12.09	***4.98	***5.00	***5.49	***4.96
c	***1.10	***−0.41	***−0.41	***0.26	***0.53	***0.54	***0.58	***0.36
Unconstr. volume	*0.03	**0.06	**0.06	0.03	***0.12	***0.12	***0.11	***0.12
	[1.69]	[2.19]	[2.28]	[1.52]	[11.65]	[11.32]	[9.24]	[9.47]
Constr. volume	−0.06	−0.01	−0.01	−0.04	***0.07	**0.06	**0.05	*0.05
	[0.85]	[0.50]	[0.32]	[0.81]	[2.70]	[2.53]	[1.97]	[1.77]
Amihud (2002)		0.00				−0.01		
		[0.05]				[1.33]		
Realised variance			−0.01				***0.02	
			[0.80]				[3.36]	
1M CIP basis				0.00				0.00
				[0.72]				[0.21]
Constr.-Unconstr.	−0.09	*−0.07	*−0.07	−0.07	**−0.06	**−0.06	**−0.06	**−0.07
	[1.27]	[1.67]	[1.69]	[1.35]	[2.07]	[2.07]	[2.02]	[2.05]
R^2 in %	0.02	0.03	0.04	0.03	1.36	1.38	1.56	1.21
Avg. #Time periods	1,981	1,981	1,981	1,786	1,983	1,982	1,982	1,788
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$) computed based on the weighted midquote price (Hagströmer, 2021), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 September 2012 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

the most liquid financial markets in the world. However, we believe that our findings also have implications for other over-the-counter (OTC) markets. For instance, broadly similar mechanisms could be at play when pricing distortions emerge between similar government bonds (Hu, Pan, and Wang, 2013) with pronounced deviations from a smooth yield curve (as observed during the Covid-19 crisis). We leave the study of the role of dealer constraints on the liquidity provision in other important OTC markets (e.g., government and corporate bonds, OTC derivatives) to future research.

References

Adrian, T. and Boyarchenko, N., 2012. Intermediary leverage cycles and financial stability. *SSRN Electronic Journal*.

Adrian, T. and Shin, H. S., 2010. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3): 418–437.

Adrian, T. and Shin, H. S., 2013. Procyclical leverage and value-at-risk. *Review of Financial Studies*, 27 (2):373–403.

Adrian, T., Etula, E., and Muir, T., 2014. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, 69(6):2557–2596.

Akram, Q. F., Rime, D., and Sarno, L., 2008. Arbitrage in the foreign exchange market: Turning on the microscope. *Journal of International Economics*, 76(2):237–253.

Amihud, Y., 2002. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

Andersen, L., Duffie, D., and Song, Y., 2019. Funding value adjustments. *The Journal of Finance*, 74(1): 145–192.

Andrews, D. W. K. and Monahan, J. C., 1992. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica*, 60(4):953–966.

Barndorff-Nielsen, O. E. and Shephard, N., 2002. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(2):253–280.

Berndt, A., Duffie, J. D., and Zhu, Y., 2020. Across-the-curve credit spread indices. *SSRN Electronic Journal*.

Bjønnes, G. H. and Rime, D., 2005. Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics*, 75(3):571–605.

Canova, F. and Nicoló, G. D., 2002. Monetary disturbances matter for business fluctuations in the g-7. *Journal of Monetary Economics*, 49(6):1131–1159.

Cespa, G., Gargano, A., Riddiough, S. J., and Sarno, L., 2021. Foreign exchange volume. *The Review of Financial Studies*, 35(5):2386–2427.

Chaboud, A. P., Chernenko, S. V., and Wright, J. H., 2008. Trading activity and macroeconomic announcements in high-frequency exchange rate data. *Journal of the European Economic Association*, 6 (2-3):589–596.

Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., and Vega, C., 2014. Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5):2045–2084.

Christiansen, C., Ranaldo, A., and Söderlind, P., 2011. The time-varying systematic risk of carry trade strategies. *Journal of Financial and Quantitative Analysis*, 46(04):1107–1125.

Chu, Y., Hirshleifer, D., and Ma, L., 2020. The causal effect of limits to arbitrage on asset pricing anomalies. *The Journal of Finance*, 75(5):2631–2672.

Cohen, L., Diether, K. B., and Malloy, C. J., 2007. Supply and demand shifts in the shorting market. *The Journal of Finance*, 62(5):2061–2096.

Comerton-Forde, C., Hendershott, T., Jones, C. M., Moulton, P. C., and Seasholes, M. S., 2010. Time variation in liquidity: The role of market-maker inventories and revenues. *The Journal of Finance*, 65 (1):295–331.

Dávila, E., Graves, D., and Parlatore, C., 2022. The value of arbitrage. *Unpublished working paper. NBER*.

Driscoll, J. C. and Kraay, A. C., 1998. Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

Du, W., Tepper, A., and Verdelhan, A., 2018. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957.

Du, W., Hébert, B., and Huber, A. W., forthcoming 2022. Are intermediary constraints priced? *The Review of Financial Studies*.

Dudley, W. C., 2018. Regulation and liquidity provision. *Remarks of the President and Chief of the Federal Reserve Bank of New York at the SIFMA Liquidity Forum*.

Duffie, D., 2010. Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance*, 65(4):1237–1267.

Duffie, D., 2018. Post-crisis bank regulations and financial market liquidity. *Lezioni Paolo Baffi di Moneta ae Finanza*.

Duffie, D. and Pan, J., 1997. An overview of value at risk. *The Journal of Derivatives*, 4(3):7–49.

Evans, M. D., 2002. FX trading and exchange rate dynamics. *The Journal of Finance*, 57(6):2405–2447.

Evans, M. D. and Lyons, R. K., 2002. Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1):247–290.

Evans, M. D. and Lyons, R. K., 2005. Do currency markets absorb news quickly? *Journal of International Money and Finance*, 24(2):197–217.

Fleckenstein, M. and Longstaff, F., 2018. Shadow funding costs: Measuring the cost of balance sheet constraints. *Unpublished working paper. NBER*.

Foucault, T., Pagano, M., and Roell, A., 2013. *Market Liquidity*. Oxford University Press.

Foucault, T., Kozhan, R., and Tham, W. W., 2016. Toxic arbitrage. *The Review of Financial Studies*, 30(4): 1053–1094.

Gabaix, X. and Maggiori, M., 2015. International liquidity and exchange rate dynamics. *The Quarterly Journal of Economics*, 130(3):1369–1420.

Gârleanu, N. and Pedersen, L. H., 2011. Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022.

Goldberg, J., 2020. Liquidity supply by broker-dealers and real activity. *Journal of Financial Economics*, 136(3):806–827.

Goldberg, J. and Nozawa, Y., 2020. Liquidity supply in the corporate bond market. *The Journal of Finance*, 76(2):755–796.

Gromb, D. and Vayanos, D., 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2-3):361–407.

Gromb, D. and Vayanos, D., 2010. Limits of arbitrage. *Annual Review of Financial Economics*, 2(1): 251–275.

Grossman, S. J. and Miller, M. H., 1988. Liquidity and market structure. *The Journal of Finance*, 43(3): 617–633.

Hagströmer, B., 2021. Bias in the effective bid-ask spread. *Journal of Financial Economics*, 142(1):314–337.

Hasbrouck, J. and Levich, R. M., 2018. FX market metrics: New findings based on CLS bank settlement data. *SSRN Electronic Journal*.

Hasbrouck, J. and Levich, R. M., 2021. Network structure and pricing in the FX market. *Journal of Financial Economics*, 141(2):705–729.

He, Z. and Krishnamurthy, A., 2011. A model of capital and crises. *The Review of Economic Studies*, 79 (2):735–777.

He, Z. and Krishnamurthy, A., 2013. Intermediary asset pricing. *American Economic Review*, 103(2): 732–770.

He, Z., Kelly, B., and Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.

Hendershott, T. and Menkveld, A. J., 2014. Price pressures. *Journal of Financial Economics*, 114(3): 405–423.

Hombert, J. and Thesmar, D., 2014. Overcoming limits of arbitrage: Theory and evidence. *Journal of Financial Economics*, 111(1):26–44.

Hu, G. X., Pan, J., and Wang, J., 2013. Noise as information for illiquidity. *The Journal of Finance*, 68(6): 2341–2382.

Jeanneret, A., 2019. When do commodity prices matter for the carry trade? The role of FX liquidity. *SSRN Electronic Journal*.

Liu, Y. and Tsyvinski, A., 2020. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34 (6):2689–2727.

Mancini, L., Ranaldo, A., and Wrampelmeyer, J., 2013. Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *The Journal of Finance*, 68(5):1805–1841.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A., 2016. Information flows in foreign exchange markets: Dissecting customer currency trades. *The Journal of Finance*, 71(2):601–634.

Moore, M., Schrimpf, A., and Sushko, V., 2016. Downsized FX markets: Causes and implications. *BIS Quarterly Review*.

Nagel, S., 2012. Evaporating liquidity. *Review of Financial Studies*, 25(7):2005–2039.

Newey, W. K. and West, K. D., 1994. Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, 61(4):631–653.

Pasquariello, P., 2014. Financial market dislocations. *Review of Financial Studies*, 27(6):1868–1914.

Payne, R., 2003. Informed trade in spot foreign exchange markets: an empirical investigation. *Journal of International Economics*, 61(2):307–329.

Ranaldo, A. and Santucci de Magistris, P., forthcoming 2022. Liquidity in the global foreign exchange market. *Journal of Financial Economics*.

Ranaldo, A. and Somogyi, F., 2021. Asymmetric information risk in FX markets. *Journal of Financial Economics*, 140(2):391–411.

Rime, D. and Schrimpf, A., 2013. The anatomy of the global FX market through the lens of the 2013 Triennial Survey. *BIS Quarterly Review*.

Rime, D., Schrimpf, A., and Syrstad, O., 2021. Covered interest parity arbitrage. *Review of Financial Studies*.

Rösch, D., 2021. The impact of arbitrage on market liquidity. *Journal of Financial Economics*, 142(1): 195–213.

Rösch, D. M., Subrahmanyam, A., and van Dijk, M. A., 2016. The dynamics of market efficiency. *The Review of Financial Studies*, 30(4):1151–1187.

Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T., 2010. Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*, 77(2):665–696.

Schrimpf, A. and Sushko, V., 2019. FX trade execution: Complex and highly fragmented. *BIS Quarterly Review*.

Scott, J. H., 1976. A theory of optimal capital structure. *The Bell Journal of Economics*, 7(1):33.

Shleifer, A. and Vishny, R. W., 1997. The limits of arbitrage. *The Journal of Finance*, 52(1):35–55.

Siriwardane, E., Sunderam, A., and Wallen, J., 2021. Segmented arbitrage. *SSRN Electronic Journal*.

Somogyi, F., 2021. Dollar dominance in FX trading. *SSRN Electronic Journal*.

Uhlig, H., 2005. What are the effects of monetary policy on output? results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2):381–419.

van Binsbergen, J. H., Graham, J. R., and Yang, J., 2010. The cost of debt. *The Journal of Finance*, 65(6): 2089–2136.

van Dijk, D., Teräsvirta, T., and Franses, P. H., 2002. Smooth transition autoregressive models — a survey of recent developments. *Econometric Reviews*, 21(1):1–47.

Appendix A. Data sources

CLS data. The CLS system is owned by its 72 settlement members, which include all the dealer banks listed in the Euromoney FX surveys. To protect member anonymity, CLS does not disclose any transaction-level information about settlement activity. Therefore, the CLS data set only contains hourly aggregates of the trading activity in each currency pair and provides no information about traders' identities or executed transaction prices.

The CLS FX spot volume and order flow data sets are interrelated. Volume data include the sum of all inter-dealer and dealer-to-customer trades. Order flow data contain separate entries for buying and selling activity but only for dealer-to-customer transactions. Moreover, the buy and sell volume in a given hour and currency pair refers to how much of the base currency was bought and sold by customers from dealer banks (see Somogyi, 2021).

Customers can be categorised into four groups: corporates, funds, non-bank financial firms, and non-dealer banks. "Funds" may also include principal trading firms (PTFs) such as high-frequency trading firms and electronic non-bank liquidity providers (e.g., XTX or Jump Trading). The majority of these PTFs relies on prime brokers to gain access to the FX market (Schrimpf and Sushko, 2019). Hence, if PTFs trade via a prime broker who is a CLS member, then this trade would appear as a bank-to-bank trade. Inter-bank trades are excluded from the flow (but not from the volume) data set unless one of the counterparties is classified as a non-dealer bank. See Ranaldo and Somogyi (2021) for further details on how CLS categorises market participants into customers and dealer/non-dealer banks, respectively.

Euromoney FX survey. Major FX dealer banks are at the heart of our composite dealer constraint measure. For each year from 2011 to 2020, we retrieve the ranking of the top 10 FX dealer banks from the Euromoney FX surveys, which are publicly available. See Table A.1 for an overview of the top 10 FX dealer banks over the sample period from 2011 to 2020. Note that this implies that we do not include any non-bank financial liquidity providers (i.e., XTX or Jump Trading), which are privately held companies. What follows lists the data source for each of the four subcomponents of our composite dealer constraint measure (DCM).

- **Value-at-Risk (VaR)** is retrieved directly from the financial statements for each of the top 10 dealer banks and is based on the FX risk in banks' overall trading book. Hence the VaR measure captures, among others, risks related to fixed income, equities, commodities, derivatives, and foreign exchange trading positions. The frequency is quarterly.
- **Leverage ratio (HKM)** is computed following the work by He et al. (2017) as book debt (i.e., short plus long term debt) relative to the sum of market equity (i.e., shares outstanding times share price) and book debt that are retrieved from Bloomberg for each dealer bank. The frequency is quarterly.

- **Credit default spread (CDS)** with 5 year maturity is retrieved from Bloomberg for each dealer bank. The CDS premia are denominated in dollars for US banks and in euros for all Europeans banks, including the UK domiciled ones. The frequency is daily.
- **Debt funding cost (DFC)** is retrieved from iBoxx for each dealer bank and corresponds to the average bond issuance cost across different maturities and major currencies (i.e., USD, EUR, and GBP). Note that conceptually our measure of debt funding cost is similar to the across-the-curve credit spread index (AXI) proposed by Berndt et al. (2020). The main difference is that our key measure of bond issuance cost is the annual yield, whereas Berndt et al. (2020) utilise credit spreads. The frequency is daily.

Figure 1: Time-series of top 10 FX dealer share

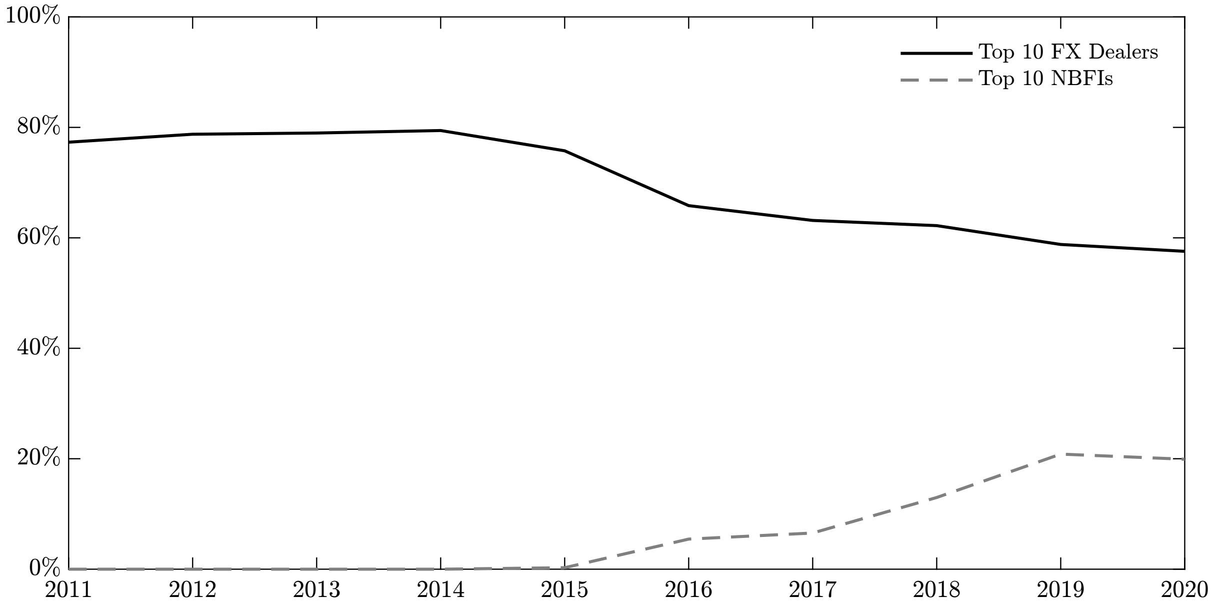

Note: This figure reports the market share of the top 10 FX dealer banks (e.g., Citi Bank or UBS) as well as non-bank financial liquidity providers (i.e., XTX, HC Tech or Jump Trading) for the years 2011 to 2020 from the Euromoney FX surveys.

Table A.1: Top 10 FX dealer banks

Rank	2011	2012	2013	2014	2015
1	Deutsche Bank	Deutsche Bank	Deutsche Bank	Citi Bank	Citi Bank
2	Barclays	Citi Bank	Citi Bank	Deutsche Bank	Deutsche Bank
3	UBS	Barclays	Barclays	Barclays	Barclays
4	Citi Bank	UBS	UBS	UBS	JP Morgan Chase
5	JP Morgan Chase	HSBC	HSBC	HSBC	UBS
6	HSBC	JP Morgan Chase	JP Morgan Chase	JP Morgan Chase	Bank of America
7	Royal Bank of Scotland	Royal Bank of Scotland	Royal Bank of Scotland	Bank of America	HSBC
8	Credit Suisse	Credit Suisse	Credit Suisse	Royal Bank of Scotland	BNP Paribas
9	Goldman Sachs	Morgan Stanley	Morgan Stanley	BNP Paribas	Goldman Sachs
10	Morgan Stanley	Goldman Sachs	Bank of America	Goldman Sachs	Royal Bank of Scotland
Rank	2016	2017	2018	2019	2020
1	Citi Bank	Citi Bank	JP Morgan Chase	JP Morgan Chase	JP Morgan Chase
2	JP Morgan Chase	JP Morgan Chase	UBS	Deutsche Bank	UBS
3	UBS	UBS	Bank of America	Citi Bank	Deutsche Bank
4	Deutsche Bank	Bank of America	Citi Bank	UBS	Citi Bank
5	Bank of America	Deutsche Bank	HSBC	State Street	HSBC
6	Barclays	HSBC	Goldman Sachs	HSBC	Goldman Sachs
7	Goldman Sachs	Barclays	Deutsche Bank	Bank of America	State Street
8	HSBC	Goldman Sachs	Standard Chartered	Goldman Sachs	Bank of America
9	Morgan Stanley	Standard Chartered	State Street	Barclays	BNP Paribas
10	BNP Paribas	BNP Paribas	Barclays	BNP Paribas	Barclays

Note: This table reports the ranking of the top 10 FX dealer banks for the years 2011 to 2020 from the Euromoney FX surveys. Note that this ranking only includes banks and excludes any non-bank financial liquidity providers (i.e., XTX, HC Tech or Jump Trading), which are privately held companies.

Appendix B. Estimating a panel LSTAR model

Figure B.1: Time-series of fitted G(DCM) and VLOOP

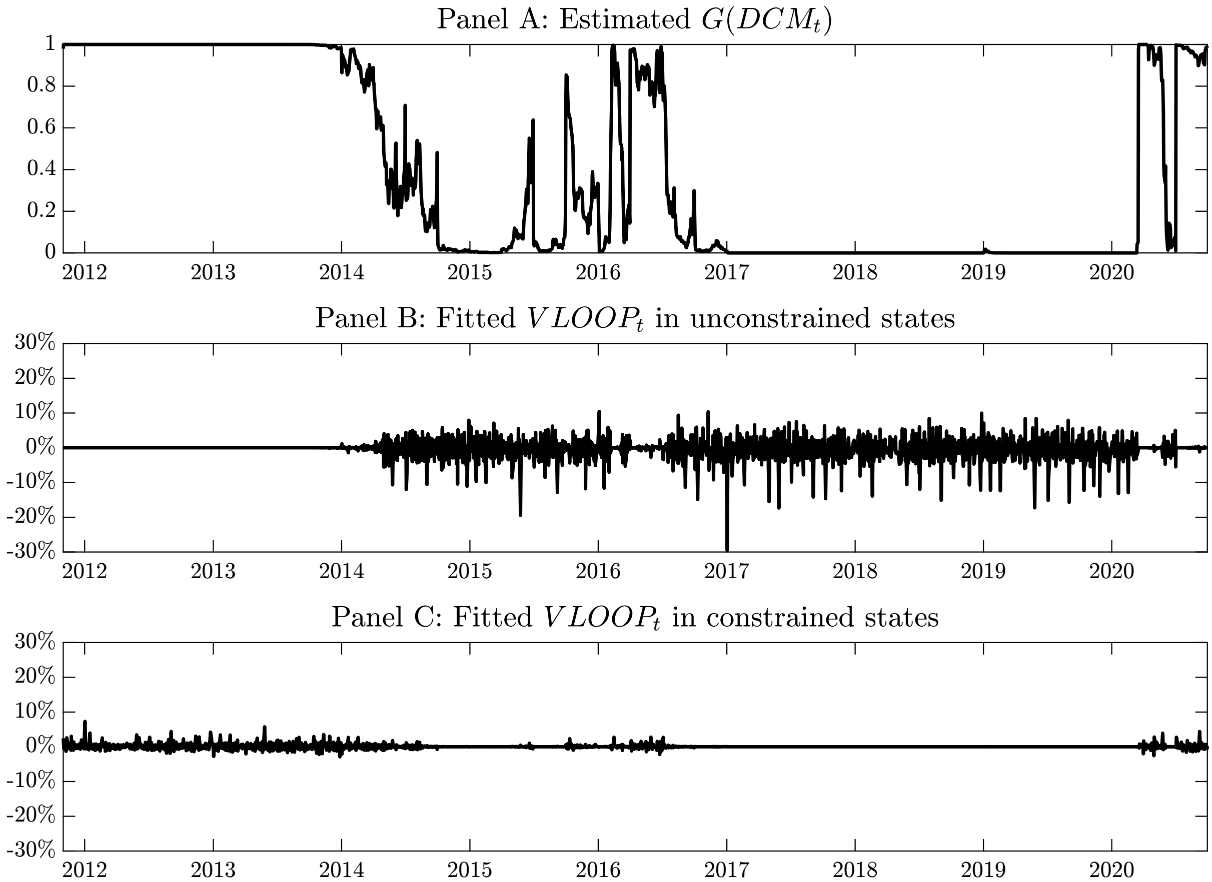

Note: Panel A of this figure shows the fitted regime function $G(DCM_t)$, using the point estimates in column 3 of Table 3. Panel B shows the cross-sectional average of the part of the fitted log changes in $VLOOP_t$ that is driven by unconstrained state coefficients ($[1 - G(z_{t-1})]\beta_1' f_t$). Panel C shows the cross-sectional average of the part driven by the constrained state coefficients ($G(z_{t-1})\beta_2' f_t$). By construction, the fitted values for log changes in $TCOST_t$ are the sum of Panels B and C. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure B.2: Time-series of fitted G(DCM) and TCOST

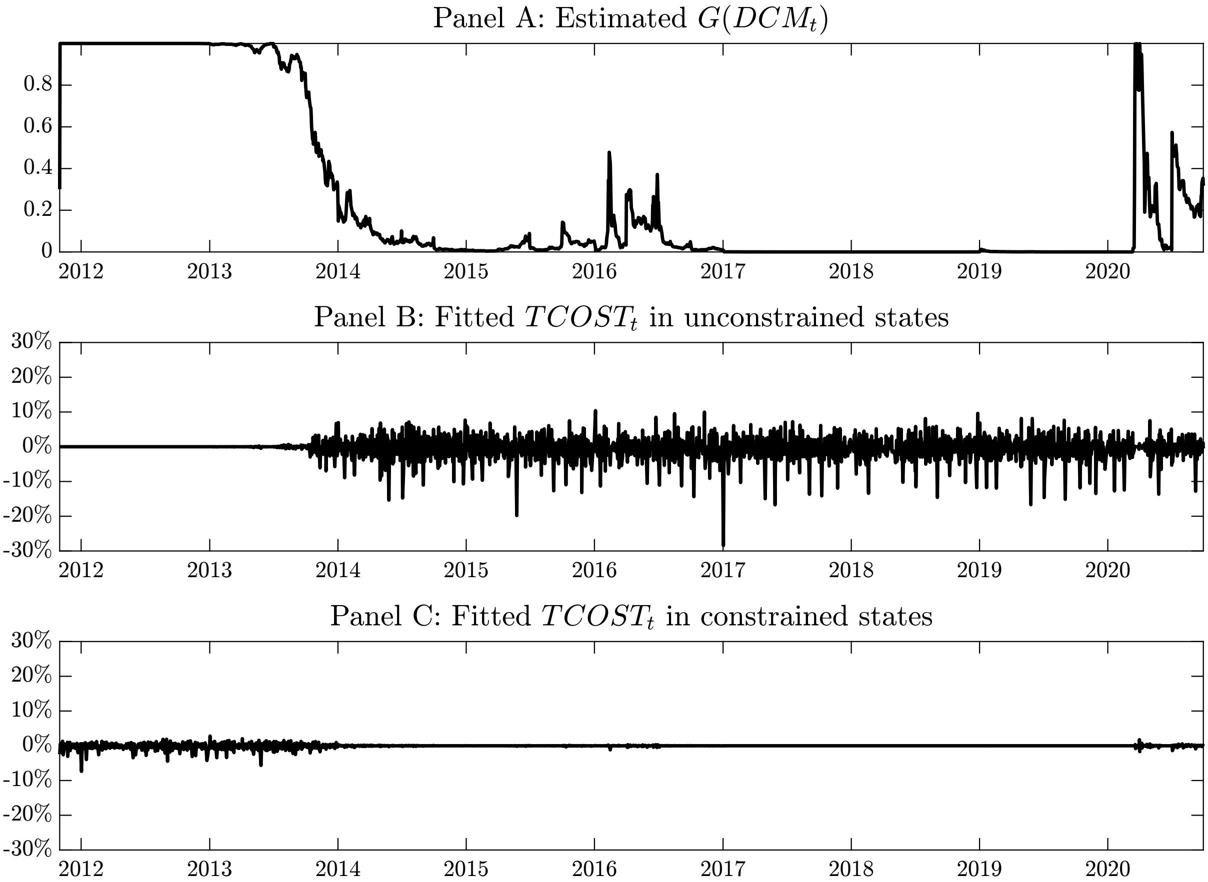

Note: Panel A of this figure shows the fitted regime function $G(DCM_t)$, using the point estimates in column 6 of Table 3. Panel B shows the cross-sectional average of the part of the fitted log changes in $TCOST_t$ that is driven by unconstrained state coefficients ($[1 - G(z_{t-1})]\beta_1' f_t$). Panel C shows the cross-sectional average of the part driven by the constrained state coefficients ($G(z_{t-1})\beta_2' f_t$). By construction, the fitted values for log changes in $TCOST_t$ are the sum of Panels B and C. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure B.3: Principal component analysis

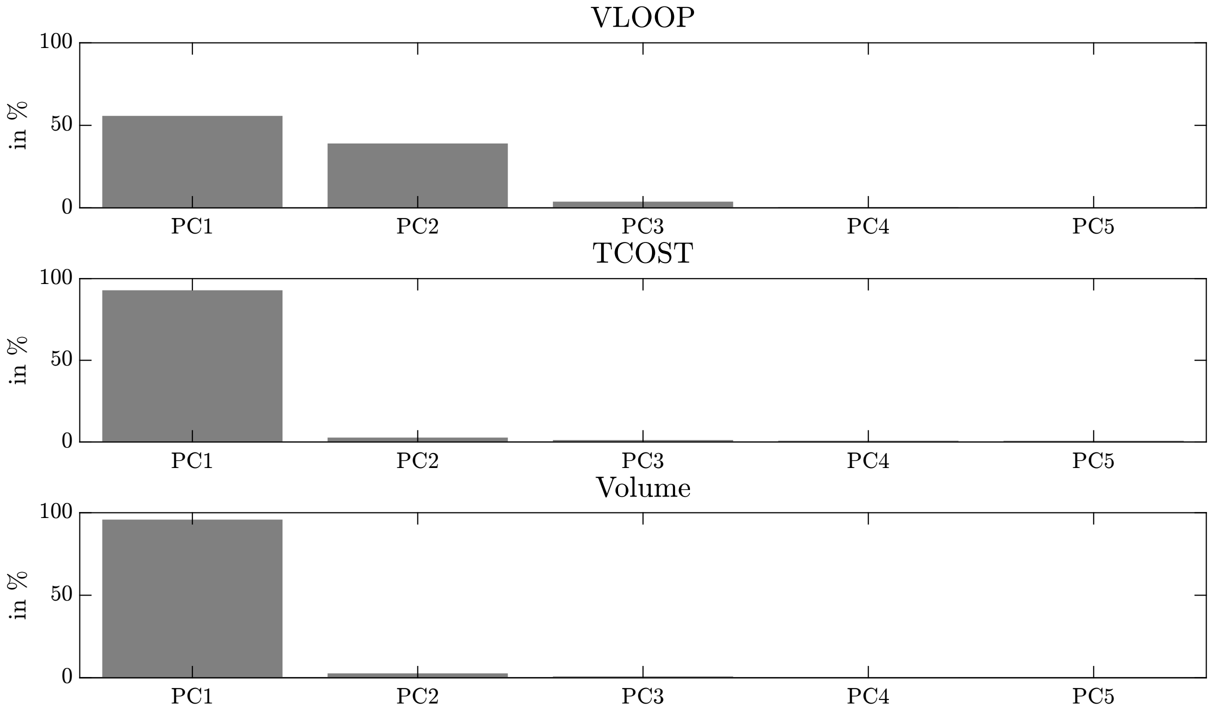

Note: This figure plots the share of variation (in %) across currency pair triplets explained by the first 5 principal components (PCs). The top two figures are based on our two liquidity cost measures (i.e., VLOOP or TCOST), whereas the bottom figure is based on total trading volume. The sample covers the period from 1 November 2011 to 30 September 2020.

Table B.1: Linear model and smooth transition regression for VLOOP with DCM as state variable

Panel A: OLS	AUDJPY	AUDNZD	CADJPY	EURAUD	EURCAD	EURCHF	EURDKK	EURGBP	EURJPY	EURNOK	EURSEK	GBPAUD	GBPCAD	GBPCHF	GBPJPY
Intercept (α)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[0.13]	[0.01]	[0.45]	[0.12]	[0.02]	[0.14]	[0.17]	[0.06]	[0.78]	[0.09]	[0.26]	[0.02]	[0.02]	[0.07]	[0.54]
Volume	***0.22	***0.08	0.02	***0.14	0.00	0.00	0.02	***0.07	***0.18	**0.05	0.04	***0.12	***−0.11	***0.08	***0.17
	[4.10]	[2.78]	[0.43]	[4.31]	[0.03]	[0.07]	[0.65]	[2.66]	[2.97]	[2.06]	[1.38]	[3.24]	[3.12]	[2.63]	[3.84]
Realised variance	***0.08	***0.09	**0.07	*0.06	0.02	**0.05	***0.09	***0.06	**0.07	***0.09	***0.11	*0.05	***0.06	***0.10	***0.10
	[2.84]	[5.90]	[2.42]	[1.96]	[0.76]	[2.50]	[3.97]	[3.32]	[2.12]	[4.99]	[6.15]	[1.65]	[2.75]	[5.03]	[4.01]
Panel B: LSTAR	AUDJPY	AUDNZD	CADJPY	EURAUD	EURCAD	EURCHF	EURDKK	EURGBP	EURJPY	EURNOK	EURSEK	GBPAUD	GBPCAD	GBPCHF	GBPJPY
γ	12.00	12.00	12.00	12.00	12.00	12.00	12.00	3.46	1.00	12.00	1.00	12.00	12.00	12.00	1.29
c	1.56	1.47	1.89	−1.38	−0.64	−1.56	0.78	−0.44	2.00	1.44	−0.23	1.69	−1.70	−0.97	−0.94
Unconstrained volume	***0.32	***0.13	*0.09	***0.26	−0.16	0.06	*0.05	***0.20	**0.31	***0.10	0.20	***0.19	0.19	**0.14	**0.60
	[5.70]	[3.95]	[1.69]	[3.91]	[0.18]	[0.91]	[1.78]	[3.86]	[2.13]	[3.41]	[0.25]	[4.76]	[1.18]	[2.52]	[2.20]
Constrained volume	0.17	−0.02	−0.07	**0.10	0.08	0.00	−0.08	0.00	−0.10	**−0.13	−0.11	0.04	***−0.18	***0.10	0.06
	[1.04]	[0.21]	[1.04]	[2.46]	[0.09]	[0.13]	[1.49]	[0.06]	[0.38]	[2.06]	[0.24]	[0.58]	[3.73]	[2.75]	[0.93]
Intercept (α)	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	*0.01	0.00	0.00	0.00	0.00	0.00	0.00
	[1.02]	[1.09]	[1.01]	[0.21]	[0.00]	[0.12]	[0.00]	[0.93]	[1.72]	[0.74]	[0.11]	[1.15]	[0.86]	[1.00]	[0.78]
Realised variance	*0.05	***0.08	0.04	*0.05	0.02	**0.05	***0.09	***0.05	*0.06	***0.09	***0.10	0.03	***0.07	***0.09	***0.08
	[1.82]	[4.93]	[1.47]	[1.85]	[0.30]	[2.34]	[4.06]	[2.98]	[1.93]	[5.25]	[5.72]	[1.03]	[2.61]	[4.79]	[2.95]
Constrained-Unconstrained	−0.15	−0.15	*−0.17	**−0.16	*0.24	−0.06	**−0.12	**−0.20	***−0.41	***−0.23	−0.31	*−0.15	**−0.36	−0.04	*−0.53
	[0.93]	[1.38]	[1.89]	[1.99]	[1.72]	[0.86]	[2.13]	[2.22]	[2.66]	[3.27]	[0.25]	[1.78]	[2.24]	[0.69]	[1.77]
Adj. R^2 in % - OLS	4.26	2.79	0.33	3.40	−0.04	0.77	3.27	2.82	3.79	2.02	2.81	2.89	0.48	4.15	6.29
Adj. R^2 in % - LSTAR	4.84	3.24	0.43	3.12	0.30	0.68	3.67	3.85	5.03	2.77	3.26	3.65	0.78	4.36	8.49
#Obs	2,273	2,285	2,271	2,285	2,285	2,283	2,283	2,285	2,265	2,285	2,283	2,285	2,285	2,285	2,266

Note: In Panel A this table reports results from estimating a linear model (OLS) of the form $VLOOP_{k,t} = \alpha_k + \beta_1' v_{k,t} + beta_3' w_{k,t} + \epsilon_{k,t}$, where $v_{k,t}$ collects all regressors. In Panel B the table shows results from a smooth transition regression (LSTAR) of the form $VLOOP_{k,t} = \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on robust standard errors (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Table B.2: Linear model and smooth transition regression for TCOST with DCM as state variable

Panel A: OLS	AUDJPY	AUDNZD	CADJPY	EURAUD	EURCAD	EURCHF	EURDKK	EURGBP	EURJPY	EURNOK	EURSEK	GBPAUD	GBPCAD	GBPCHF	GBPJPY
Intercept (α)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[0.19]	[0.22]	[0.10]	[0.23]	[0.28]	[0.29]	[0.07]	[0.09]	[0.15]	[0.13]	[0.14]	[0.18]	[0.23]	[0.21]	[0.09]
Volume	***0.10	***0.09	***0.07	***0.08	***0.07	***0.07	***0.09	***0.06	***0.11	*0.02	***0.03	***0.06	**0.03	**0.03	***0.08
	[3.99]	[7.77]	[4.24]	[5.78]	[6.23]	[6.44]	[8.35]	[4.30]	[5.39]	[1.74]	[2.63]	[2.93]	[2.23]	[2.10]	[2.98]
Realised variance	***0.07	***0.05	***0.06	***0.07	***0.04	**0.04	***0.03	***0.05	***0.06	***0.08	***0.06	***0.07	*0.03	***0.06	**0.05
	[4.06]	[6.19]	[4.64]	[6.08]	[3.33]	[2.50]	[5.87]	[3.35]	[3.57]	[6.71]	[8.32]	[3.89]	[1.82]	[2.85]	[2.01]
Panel B: LSTAR	AUDJPY	AUDNZD	CADJPY	EURAUD	EURCAD	EURCHF	EURDKK	EURGBP	EURJPY	EURNOK	EURSEK	GBPAUD	GBPCAD	GBPCHF	GBPJPY
γ	1.20	12.00	12.00	12.00	1.68	12.00	12.00	12.00	1.00	3.72	12.00	12.00	12.00	12.00	12.00
c	−0.56	1.35	0.91	−1.22	0.35	0.90	−0.10	−1.39	−0.66	0.70	1.41	−1.28	1.57	−2.00	−0.05
Unconstrained volume	***0.24	***0.13	***0.10	***0.15	0.10	***0.08	***0.10	**0.10	0.26	**0.04	***0.04	***0.13	***0.05	−0.07	***0.14
	[3.09]	[9.03]	[5.79]	[6.25]	[1.04]	[5.93]	[8.78]	[2.14]	[0.93]	[2.13]	[3.14]	[4.65]	[2.63]	[0.79]	[3.39]
Constrained volume	0.02	0.02	−0.01	***0.06	0.01	**0.04	*0.05	***0.06	0.01	−0.03	−0.03	**0.05	−0.02	**0.04	**0.05
	[0.52]	[0.64]	[0.63]	[4.14]	[0.42]	[1.99]	[1.88]	[3.38]	[0.10]	[1.07]	[1.09]	[2.34]	[0.51]	[2.36]	[2.30]
Intercept (α)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[0.50]	[1.52]	[0.29]	[0.33]	[0.27]	[0.11]	[0.79]	[0.34]	[0.38]	[0.48]	[0.10]	[0.61]	[0.11]	[0.18]	[0.88]
Realised variance	***0.06	***0.05	***0.06	***0.07	***0.04	**0.04	***0.03	***0.05	***0.06	***0.08	***0.06	***0.07	0.03	***0.06	0.04
	[3.76]	[5.49]	[3.98]	[5.91]	[3.28]	[2.39]	[5.99]	[3.17]	[3.35]	[6.41]	[7.91]	[3.60]	[1.38]	[2.69]	[1.64]
Constrained-Unconstrained	**−0.22	***−0.12	***−0.12	***−0.08	−0.09	−0.04	−0.05	−0.05	−0.25	**−0.08	**−0.07	***−0.08	*−0.06	0.11	**−0.09
	[2.13]	[4.19]	[4.78]	[3.15]	[1.03]	[1.59]	[1.43]	[1.05]	[0.72]	[2.49]	[2.17]	[2.62]	[1.82]	[1.22]	[2.27]
Adj. R^2 in % - OLS	19.18	13.44	14.04	16.31	9.48	10.45	10.18	10.97	19.84	8.75	7.20	11.97	3.09	8.84	11.89
Adj. R^2 in % - LSTAR	20.75	15.07	14.27	16.19	8.69	9.92	8.45	10.30	19.62	9.15	7.30	12.23	2.89	8.30	12.58
#Obs	2,285	2,285	2,278	2,285	2,285	2,285	2,285	2,285	2,285	2,285	2,285	2,285	2,285	2,285	2,277

Note: In Panel A this table reports results from estimating a linear model (OLS) of the form $TCOST_{k,t} = \alpha_k + \beta_1' v_{k,t} + beta_3' w_{k,t} + \epsilon_{k,t}$, where $v_{k,t}$ collects all regressors. In Panel B the table shows results from a smooth transition regression (LSTAR) of the form $TCOST_{k,t} = \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on robust standard errors (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Table B.3: Smooth transition regression with DCM as state variable

	VLOOP				TCOST			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	***12.02	***12.02	***12.05	***12.08	***5.38	***5.38	***5.93	***4.96
c	***−0.14	***−0.14	***−0.14	***−0.20	***0.34	***0.34	***0.39	***0.34
Unconstr. volume	***0.13	***0.12	***0.11	***0.13	***0.12	***0.13	***0.09	***0.13
	[4.09]	[3.96]	[3.50]	[3.33]	[15.01]	[15.04]	[10.85]	[12.17]
Constr. volume	−0.04	−0.04	−0.05	−0.05	***0.04	***0.04	0.01	**0.03
	[1.04]	[1.15]	[1.40]	[1.43]	[2.77]	[2.91]	[0.96]	[2.13]
Amihud (2002)		−0.01				**0.01		
		[1.08]				[2.24]		
Realised variance			**0.02				***0.03	
			[2.02]				[7.95]	
1M CIP basis				0.01				0.00
				[1.07]				[0.65]
Constr.-Unconstr.	***−0.16	***−0.16	***−0.16	***−0.18	***−0.09	***−0.09	***−0.08	***−0.09
	[3.30]	[3.30]	[3.25]	[3.23]	[5.22]	[5.19]	[4.78]	[4.63]
R^2 in %	0.12	0.13	0.15	0.12	2.47	2.51	3.78	2.31
Avg. #Time periods	2,182	2,182	2,182	1,978	2,186	2,185	2,185	1,981
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Table B.4: Smooth transition regression with non dealer specific state variables

	VLOOP				TCOST			
	VIX	XAU	TED	LOIS	VIX	XAU	TED	LOIS
γ	***4.16	*12.08	***12.06	12.01	*7.65	12.04	12.02	12.07
c	***−0.50	***−0.17	***−0.39	**−0.48	***0.50	*−0.11	0.50	−0.31
Unconstr. volume	0.10	**0.06	−0.02	0.00	***0.07	***0.08	***0.08	***0.07
	[1.64]	[2.11]	[0.71]	[0.07]	[9.04]	[9.80]	[9.27]	[5.26]
Constr. volume	0.02	0.00	***0.09	0.05	***0.09	***0.06	***0.07	***0.07
	[0.46]	[0.07]	[2.66]	[1.02]	[5.36]	[4.78]	[4.84]	[5.46]
Realised variance	**0.02	**0.02	**0.02	**0.02	***0.03	***0.03	***0.03	***0.03
	[2.10]	[2.15]	[2.30]	[2.14]	[8.39]	[8.43]	[8.45]	[6.54]
Constr.-Unconstr.	−0.09	−0.06	**0.12	0.05	0.02	*−0.03	0.00	0.00
	[1.03]	[1.32]	[2.31]	[0.83]	[0.98]	[1.81]	[0.14]	[0.12]
R^2 in %	0.09	0.08	0.11	0.09	3.64	3.65	3.65	3.93
BIC	91.33	91.61	91.44	87.90	48.78	49.21	48.99	43.20
Avg. #Time periods	2,221	2,279	2,247	1,853	2,225	2,283	2,251	1,855
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the regime variable z_{t-1}. The regime variables are the 1-day lagged value of the VIX index, which is the Chicago Board Options Exchange's volatility index measuring the stock market's expectation of volatility based on S&P 500 index options; the gold price (i.e., XAU); the TED spread, which is the difference between the interest rates for three-month U.S. Treasuries contracts and the three-month Eurodollars contract; and the LIBOR-OIS spread (i.e., $LOIS$), which is considered to be measuring the health of the banking system. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Table B.5: Smooth transition regression with different state variables using market shares

	VLOOP				TCOST			
	VaR	HKM	CDS	DFC	VaR	HKM	CDS	DFC
γ	***12.05	***12.05	***7.85	***12.08	***12.10	***12.04	***6.65	***5.00
c	***0.25	***0.71	***−0.10	***−0.36	***0.39	***0.72	***−0.15	***−0.09
Unconstr. volume	***0.09	***0.07	***0.09	***0.17	***0.09	***0.09	***0.10	***0.10
	[3.15]	[2.66]	[2.88]	[3.77]	[11.35]	[11.56]	[10.83]	[8.74]
Constr. volume	**−0.08	*−0.09	−0.06	−0.02	*0.02	0.01	**0.03	**0.03
	[2.11]	[1.90]	[1.55]	[0.73]	[1.65]	[0.73]	[2.19]	[2.55]
Realised variance	**0.02	**0.02	**0.02	**0.02	***0.03	***0.03	***0.03	***0.03
	[2.07]	[2.10]	[2.11]	[2.07]	[8.35]	[8.34]	[8.37]	[7.98]
Constr.-Unconstr.	***−0.17	***−0.16	***−0.14	***−0.19	***−0.07	***−0.08	***−0.08	***−0.07
	[3.50]	[2.88]	[2.90]	[3.27]	[4.28]	[4.88]	[4.81]	[3.43]
R^2 in %	0.13	0.11	0.11	0.16	3.72	3.74	3.76	3.69
BIC	91.61	91.61	91.61	91.05	49.24	49.23	49.23	48.54
Avg. #Time periods	2,280	2,280	2,280	2,182	2,284	2,284	2,284	2,185
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors and $G(z_{t-1})$ is a logistic function depending on the regime variable z_{t-1}. The regime variables are the 1-day lagged value of primary FX dealer banks': quarterly Value-at-Risk measure (VaR, columns 1 and 6), quarterly He et al. (2017) leverage ratio (HKM, columns 2 and 7), daily credit default spread (CDS, columns 3 and 8), and daily funding cost yield (DFC, columns 4 and 9). Note that we weight each top 10 FX dealer bank (based on the Euromoney FX survey) by its relative market share when computing a cross-sectional average. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Figure B.4: VLOOP: Constrained–Unconstrained coefficient and t-stat

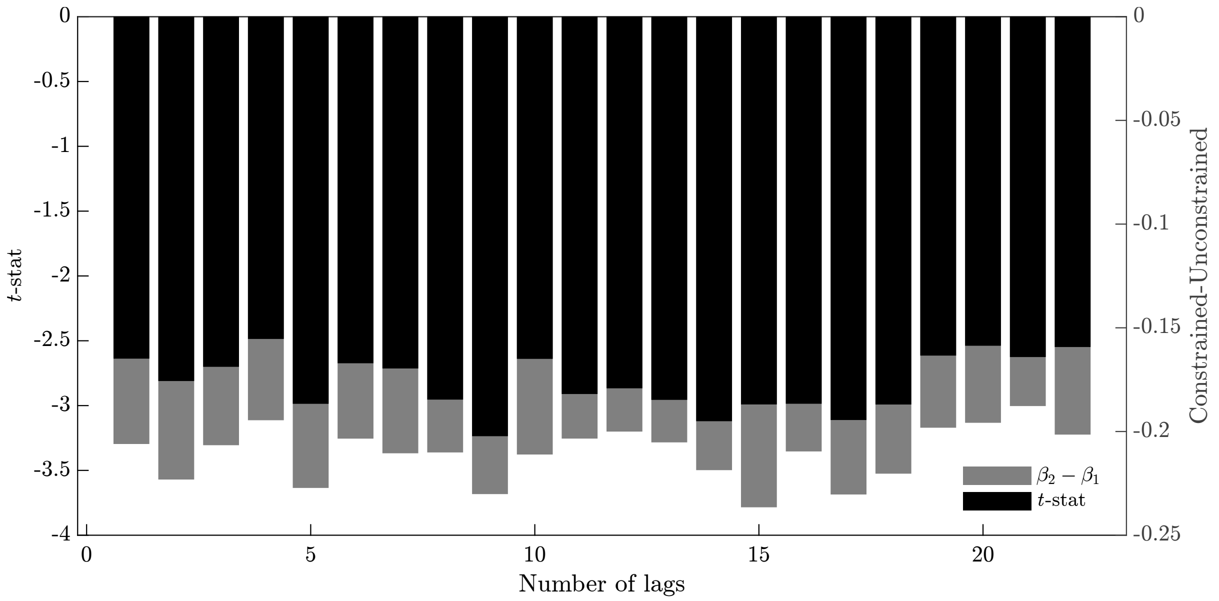

Note: This figure plots the difference between the constrained and unconstrained regime coefficient (i.e., $\beta_2 - \beta_1$) of the LSTAR model in Eq. (4) with VLOOP being the dependent variable and conditional on varying the number of lags in the regime variable DCM_{t-n} for $n = 1, 2, ..., 22$. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure B.5: Constrained–Unconstrained coefficient and t-stat

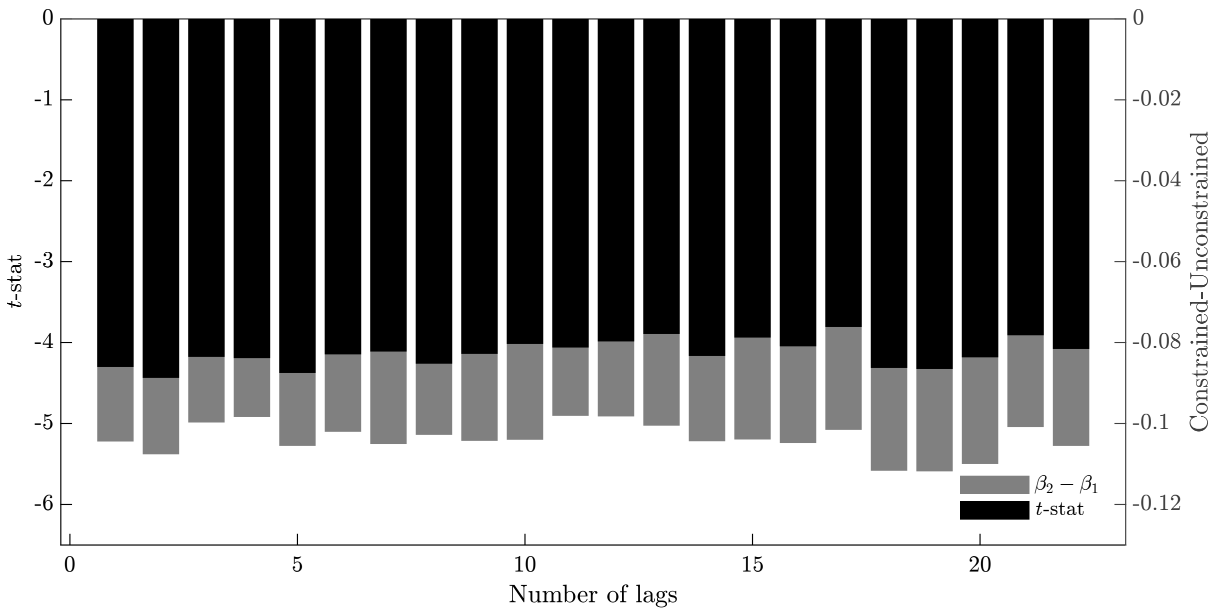

Note: This figure plots the difference between the constrained and unconstrained regime coefficient (i.e., $\beta_2 - \beta_1$) of the LSTAR model in Eq. (4) with TCOST being the dependent variable and conditional on varying the number of lags in the regime variable DCM_{t-n} for $n = 1, 2, ..., 22$. The sample covers the period from 1 November 2011 to 30 September 2020.

Table B.6: London hours: Smooth transition regression with DCM as state variable

	VLOOP				TCOST			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	*12.03	***4.97	***6.04	*12.04	***12.04	***12.06	***12.07	***12.04
c	***0.42	***0.72	***0.75	***0.35	***0.24	***0.24	***0.27	***0.26
Unconstr. volume	***0.11	***0.11	**0.06	***0.10	***0.13	***0.13	***0.10	***0.14
	[3.99]	[4.13]	[2.20]	[3.17]	[11.90]	[11.47]	[7.89]	[10.10]
Constr. volume	−0.07	−0.09	**−0.13	−0.08	*0.03	*0.03	0.00	0.03
	[1.41]	[1.59]	[2.31]	[1.56]	[1.72]	[1.76]	[0.18]	[1.43]
Amihud (2002)		0.01				0.00		
		[1.05]				[0.51]		
Realised variance			***0.05				***0.03	
			[4.13]				[3.80]	
1M CIP basis				0.00				0.00
				[0.97]				[0.46]
Constr.-Unconstr.	***−0.18	***−0.20	***−0.19	***−0.18	***−0.10	***−0.10	***−0.09	***−0.11
	[3.08]	[3.22]	[3.11]	[2.93]	[4.91]	[4.89]	[4.47]	[4.38]
R^2 in %	0.11	0.12	0.24	0.10	1.82	1.82	2.75	1.86
Avg. #Time periods	2,173	2,173	2,173	1,970	2,186	2,185	2,185	1,981
#Currency triplets	15	15	15	15	15	15	15	15
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the regime variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. When aggregating hourly to daily data we omit any observations outside of the main London stock market trading hours (i.e., from 8 am to 5 pm GMT). The sample covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Table B.7: Smooth transition regression with DCM as state variable (euro triplets)

	VLOOP				TCOST			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
γ	12.09	12.01	12.06	12.03	***4.94	***12.08	**12.05	***12.08
c	−0.06	−0.06	−0.06	−0.07	***−0.33	***−0.21	***0.53	***−0.23
Unconstr. volume	***0.20	***0.21	***0.19	***0.18	***0.10	***0.10	***0.06	***0.10
	[4.56]	[4.54]	[4.13]	[4.20]	[8.86]	[9.54]	[6.84]	[9.00]
Constr. volume	**0.12	**0.12	**0.11	**0.11	***0.05	***0.05	0.01	***0.05
	[2.28]	[2.41]	[2.04]	[2.01]	[4.26]	[4.37]	[1.00]	[4.22]
Amihud (2002)		0.01				*0.00		
		[0.88]				[1.69]		
Realised variance			0.01				***0.04	
			[0.95]				[9.07]	
1M CIP basis				0.01				0.00
				[1.24]				[0.28]
Constr.-Unconstr.	−0.09	−0.08	−0.08	−0.07	***−0.06	***−0.05	***−0.04	***−0.05
	[1.21]	[1.16]	[1.13]	[0.99]	[3.28]	[3.22]	[2.93]	[3.04]
R^2 in %	0.71	0.70	0.71	0.61	3.33	3.37	5.91	3.55
Avg. #Time periods	2,183	2,182	2,182	2,092	2,186	2,184	2,184	2,095
#Currency triplets	6	6	6	6	6	6	6	6
Currency triplet FE	yes	yes	yes	yes	yes	yes	yes	yes
Time-series FE	yes	yes	yes	yes	yes	yes	yes	yes

Note: This table reports results from daily fixed effects LSTAR panel regressions of the form $y_{k,t} = \lambda_t + \alpha_k + [1 - G(z_{t-1})]\beta_1' f_{k,t} + G(z_{t-1})\beta_2' f_{k,t} + \beta_3' w_{k,t} + \varepsilon_{k,t}$, where the dependent variable $y_{k,t}$ is a liquidity cost measure (i.e., $VLOOP$ or $TCOST$), $f_{k,t}$ ($w_{k,t}$) are state-dependent (*state-independent*) regressors, and $G(z_{t-1})$ is a logistic function depending on the state variable z_{t-1}. The regime variable is the 1-day lagged value of the dealer constraint measure DCM_t. The optimal parameters γ and c are determined by nonlinear least squares minimising the concentrated sum of squared errors. Both dependent and independent variables are taken in logs and changes. The sample consists of 6 euro-based currency pair triplets that do not involve any dollar currency pairs (i.e., AUD-EUR-JPY, CAD-EUR-JPY, GBP-EUR-AUD, GBP-EUR-CAD, GBP-EUR-CHF, and GBP-EUR-JPY) and covers the period from 1 November 2011 to 30 September 2020. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Andrews and Monahan, 1992; Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Appendix C. Estimating an SVAR with sign restrictions

We estimate a structural vector autoregression (SVAR) model of liquidity cost measures (i.e., VLOOP or TCOST) and dealer-provided volume VLM. Let $Y_{k,t} = [X_{k,t} \; VLM_{k,t}]^T$ be a 2×1 vector containing $X \in \{VLOOP, TCOST\}$ and VLM in currency pair triplet k and day t. The bivariate panel SVAR for $Y_{k,t}$ is:

$$Y_{k,t} = \alpha_k + \sum_{i=1}^{l} B_{k,i} Y_{k,t-i} + \xi_{k,t}, \tag{C.1}$$

where $B_{k,i}$ is a 2×2 matrix of coefficients, l the lag length, $\xi_{k,t} = [\xi_{X;k,t} \; \xi_{VLM;k,t}]^T$ the reduced form error, and α_k is a 2×1 vector of currency triplet fixed effects. The residual $\xi_{k,t}$ is:

$$\begin{bmatrix} \xi_{X;k,t} \\ \xi_{VLM;k,t} \end{bmatrix} = A_k \begin{bmatrix} \delta^s_{k,t} \\ \delta^d_{k,t} \end{bmatrix}, \tag{C.2}$$

where A_k is a 2×2 matrix and $\delta_{k,t} = [\delta^s_{k,t} \; \delta^d_{k,t}]^T$ is a 2×1 vector. Based on Eqs (C.1) and (C.2), the first column of A_k corresponds to changes in liquidity provision costs (i.e., VLOOP or TCOST) and dealer-intermediated volume associated with an increase in $\delta^s_{k,t}$, whereas the second column corresponds to changes in liquidity costs and VLM associated with an increase in $\delta^d_{k,t}$. Following Goldberg (2020), if A_k satisfies the following sign restrictions:

$$sign(A_k) = \begin{pmatrix} + & + \\ - & + \end{pmatrix}, \tag{C.3}$$

then $\delta^s_{k,t}$ can be interpreted as an inward shift in liquidity supply reflecting a tightening of dealer constraints, whereas $\delta^d_{k,t}$ corresponds to an outward shift in customers' liquidity demand. In particular, the sign restrictions in Eq. (C.3) assume that supply shifts lead to opposite-sign changes in liquidity costs and trading volumes, whereas demand shifts are assumed to lead to same-sign changes in liquidity costs and volume.

To identify supply and demand shifts, we estimate Eqs (C.1) and (C.2) imposing the sign restrictions in Eq. (C.3) using Bayesian methods. Specifically, we follow the approach of Uhlig (2005) and others, which has become widely used to estimate models with sign restrictions. Both the liquidity cost measure (i.e., VLOOP or TCOST) and dealer-intermediated volume enter in log levels. Consider the reduced-form SVAR in Eq. (C.1) with parameters $B_k = [B_{k,1}, ..., B_{k,l}]$ and covariance matrix Σ_k for currency pair triplet k. We use a weak Normal-Wishart prior over these parameters. The lag length l is determined according to the Akaike Information Criterion and is equal to 2 in our baseline estimation. The parameters of the panel SVAR are B_k, Σ_k, and A_k, where A_k is the mapping from the liquidity supply and demand shifts $\delta_{k,t}$ to the reduced-form residual $\xi_{k,t}$ given by $\xi_{k,t} = A_k \delta_{k,t}$. The ultimate aim is to draw from the posterior distribution of $\delta_{k,t}$. Hence, we first draw from the posterior distribution over B_k and Σ_k. By definition, A_k has to satisfy $A_k A_k^T = \Sigma_k$. Specifically, we draw A_k by using Cholesky factorisation: $\Sigma_k = chol(\Sigma_k)chol(\Sigma_k)^T$. Next, we draw orthonormal matrices Q_k uniformly from the unit circle and compute $A_k = chol(\Sigma_k)Q_k$. If the resulting A_k satisfies the sign restrictions in Eq. (C.3) over 2 periods then we keep the draw and discard it otherwise. When implementing this estimation procedure we make 500 draws over B_k and Σ_k and, for each B_k and Σ_k, 500 draws over Q_k. Eventually, the liquidity supply and demand shift proxies are normalised to have mean zero and standard deviation equal to one.

For illustrative purposes, Figure C.1 (Figure C.2) shows estimates of the dynamic responses of $VLOOP$ ($TCOST$) and VLM to supply and demand shifts for the EUR-USD-JPY currency pair triplet.[33] By construction, concurrently with a liquidity supply shift, VLOOP (TCOST) rises and VLM positions decline. As shown in Figure C.1, contemporaneous with a liquidity supply shift, VLOOP (TCOST) rises 31% (6%) and VLM positions decline 15% (18%), according to the posterior mean. Contrarily, a liquidity demand shock is associated with an increase in VLOOP (TCOST) and VLM by 25% and 12% (13% and 21%), respectively.

Figure C.1: Dynamic impulse response function for EUR-USD-JPY; VLOOP

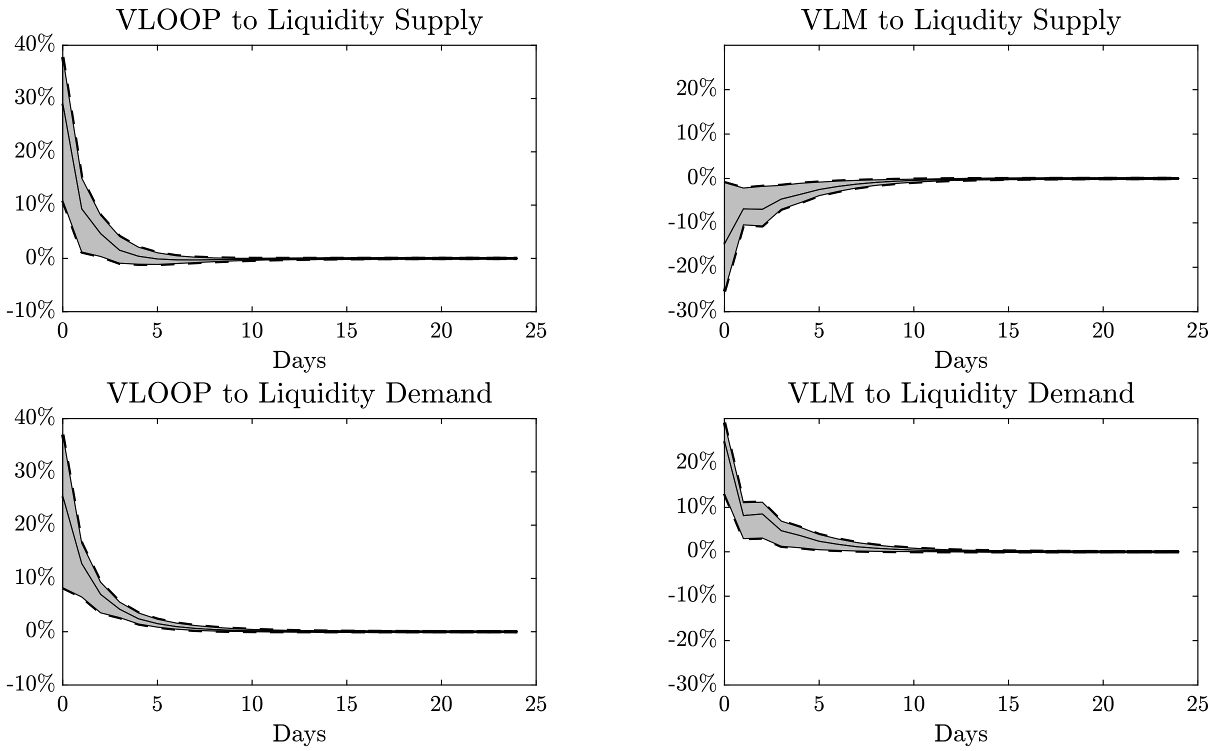

Note: This figure plots the estimated dynamic response of the shadow cost of intermediary constraints (VLOOP) and dealer-intermediated volume (VLM) associated with liquidity supply and demand shifts. The median response is shown by the black solid line. The grey shaded area marks a pointwise 95% confidence interval around the median. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure C.3 presents a scatter plot of the average 30-day rolling window correlation be-

[33] The impulse response functions for the other 14 currency pair triplets exhibit qualitatively similar patters.

Figure C.2: Dynamic impulse response function for EUR-USD-JPY; TCOST

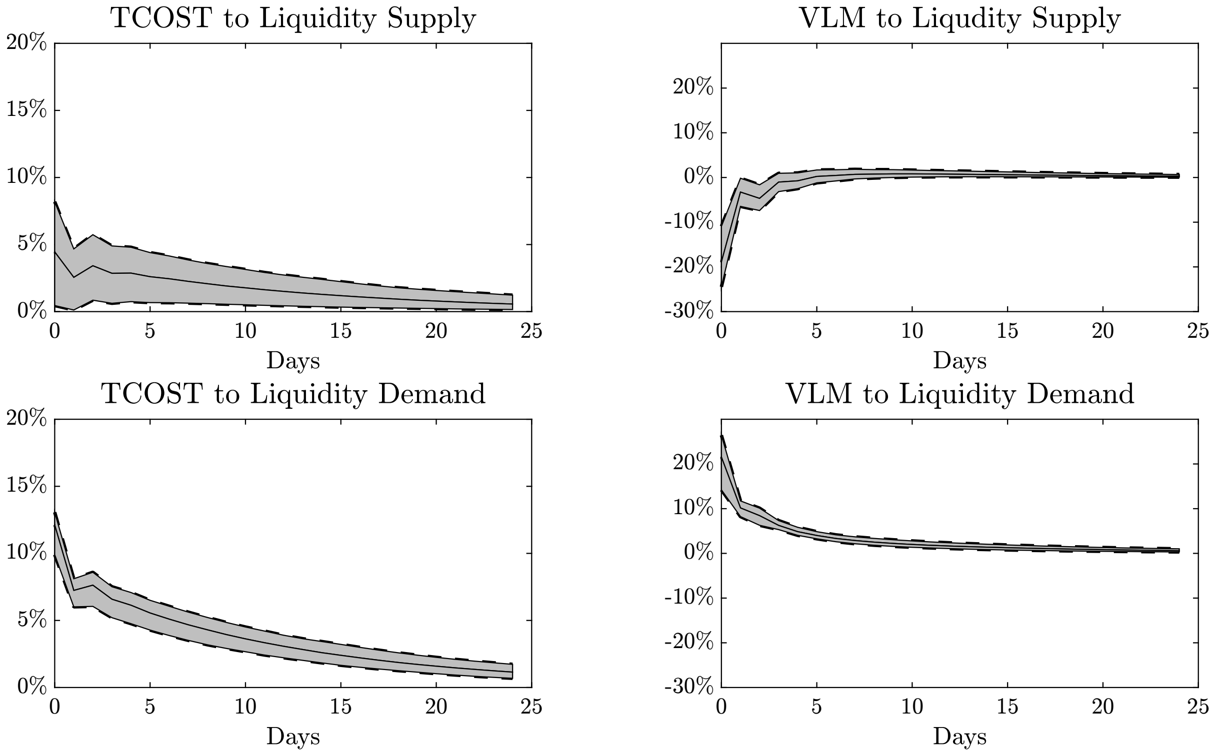

Note: This figure plots the estimated dynamic response of dealers' compensation for enduring inventory imbalances (TCOST) and dealer-intermediated volume (VLM) associated with liquidity supply and demand shifts. The median response is shown by the black solid line. The grey shaded area marks a pointwise 95% confidence interval around the median. The sample covers the period from 1 November 2011 to 30 September 2020.

tween each of our two liquidity cost measures (i.e., VLOOP and TCOST) and total dealer-intermediated trading volume against our dealer constraint measure DCM. For ease of illustration, we show the cross-sectional average of these rolling window correlations across 15 triplets of currency pairs. There are two key takeaways from this figure: First, both dimensions of liquidity costs (i.e., VLOOP and TCOST) covary positively on average with dealer-intermediated trading volume. Second, the correlation between the cost of liquidity provision (i.e., VLOOP and TCOST) and trading volume weakens substantially as DCM increases.

Figures C.4 and C.5 plot the rolling correlation between each of our two liquidity cost measures (i.e., VLOOP or TCOST) and dealer-provided volumes. It is easy to see that the strong positive association between liquidity costs and trading volume breaks down during the Covid-19 market turmoil in March and April 2020 across all 15 currency pair triplets.

Figure C.3: Rolling correlations of liquidity costs and volumes vs dealer constraints

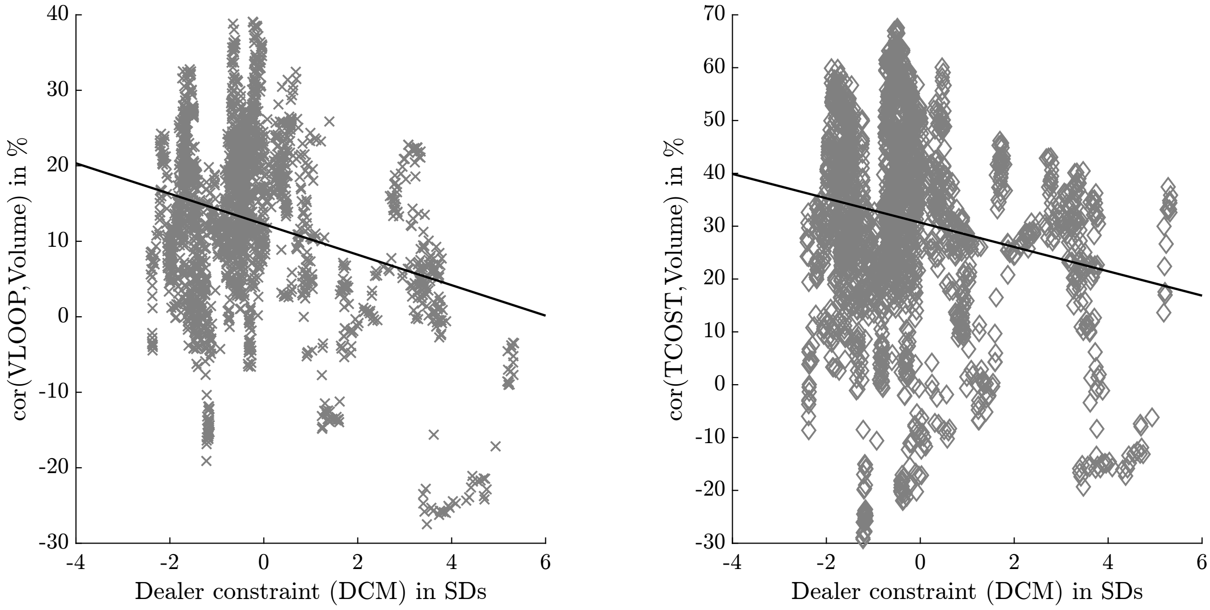

Note: This figure plots the cross-sectional average of the 30-day rolling window correlation between the shadow cost of intermediary constraints and total dealer-intermediated trading volume (i.e., $cor(VLOOP, Volume)$, left figure) as well as between dealers' compensation to endure inventory imbalances and total dealer-intermediated trading volume (i.e., $cor(TCOST, Volume)$, right figure) in percent (%). Our dealer constraint measure (DCM) is in units of standard deviations. We define DCM as the first principal component of the top 10 FX dealers' (based on the Euromoney FX survey) quarterly Value-at-Risk measure (VaR), quarterly He et al. (2017) leverage ratio (HKM), daily credit default spread (CDS), and daily debt funding cost (DFC). The bold black lines are linear regression lines. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure C.4: Rolling window correlation VLOOP and VLM

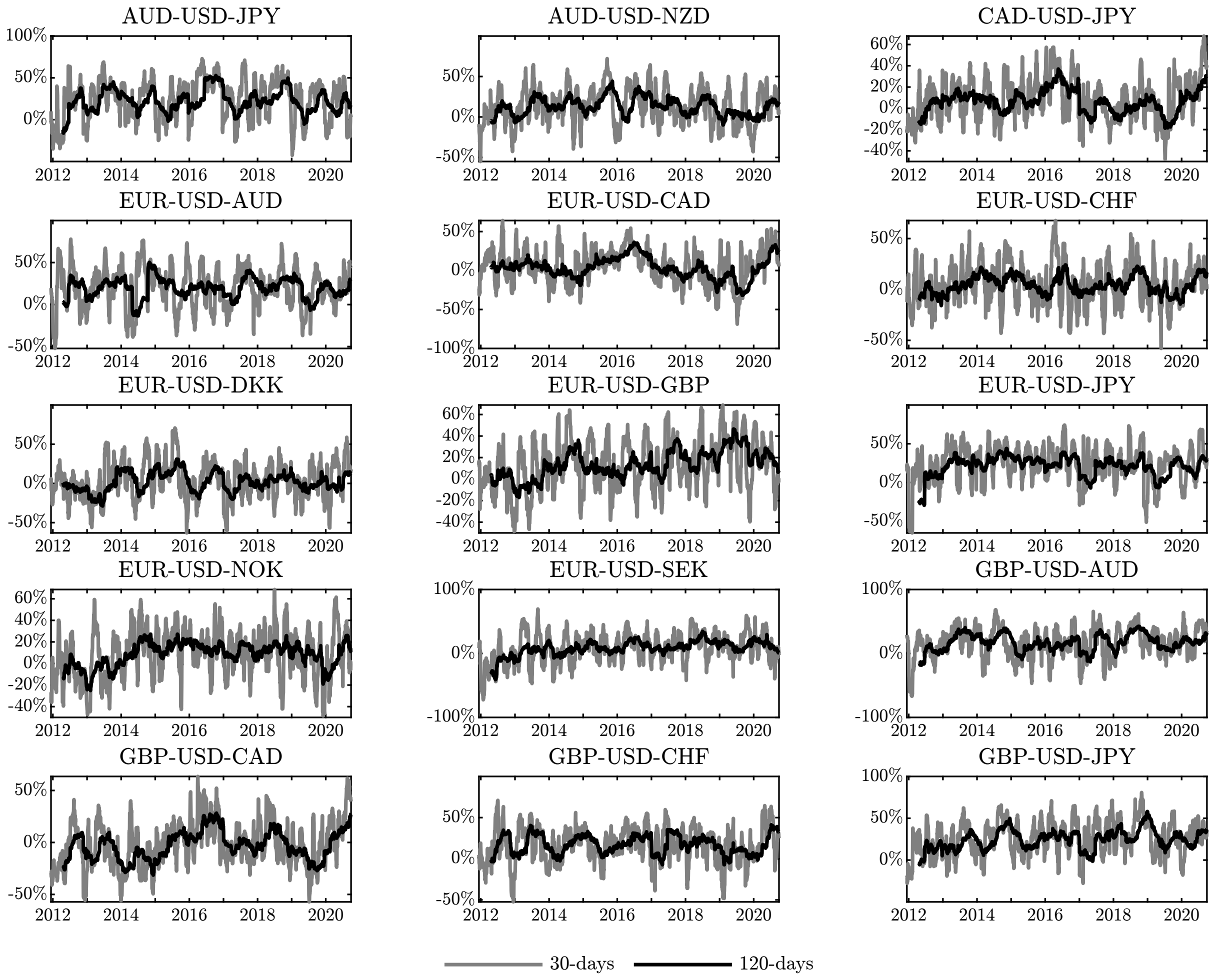

Note: This figure plots the 30- and 252-day rolling window correlation of daily cumulative no-arbitrage deviations VLOOP (i.e., shadow cost of intermediary constraints) and dealer-intermediated trading volume VLM. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure C.5: Rolling window correlation TCOST and VLM

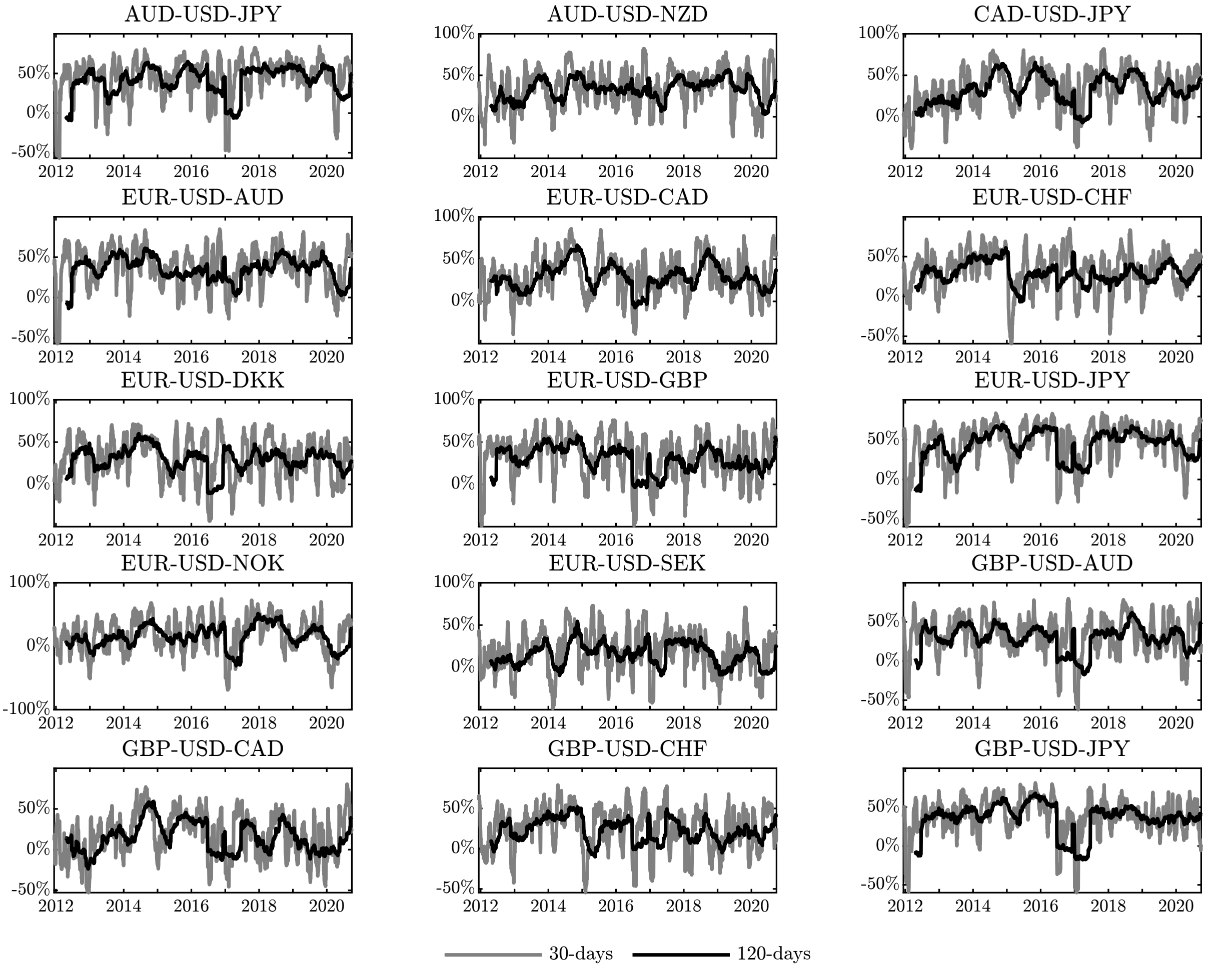

Note: This figure plots the 30- and 252-day rolling window correlation of daily cumulative round-trip transaction cost TCOST (i.e., dealers' compensation for enduring inventory imbalances) and dealer-intermediated trading volume VLM. The sample covers the period from 1 November 2011 to 30 September 2020.

Appendix D. Quasi-natural experiment: Swiss franc decap

Table D.1 reports the results from daily panel regressions of the form:

$$\rho_{k,t} = \alpha + \eta_1 D_{k,t} + \eta_2 Post_t + \eta_3 (D_{k,t} \times Post_t) + \kappa' w_{k,t} + \epsilon_{k,t}, \tag{D.1}$$

where the dependent variable is the 30-day rolling window correlation of liquidity provision costs (i.e., VLOOP or TCOST) and dealer-provided trading volume (i.e., VLM), α denotes the intercept, $D_{k,t}$ is equal to one if currency pair triplet k contains the Swiss franc, $Post_t$ is one for the time period after the removal of the Swiss franc cap on 15 January 2015, and η_3 is the difference-in-differences (DnD) coefficient. $w_{k,t}$ collects additional control variables such as the realised variance or Amihud (2002) price impact measure. Except for the case where $\rho = cor(VLOOP, VLM)$ we find the DnD regression coefficient η_3 to be negative and statistically significant. For instance, after the removal of the Swiss franc cap the correlation between TCOST and VLM is 54 percentage points lower for currency pair triplets involving the Swiss franc. Figures D.1 and D.2 illustrate the drop in the rolling window correlation coefficient based on each of our two liquidity cost measures (i.e., VLOOP or TCOST) after the removal of the Swiss franc cap.

Figure D.1: Event study: $cor(VLOOP, VLM)$ around the Swiss franc decap

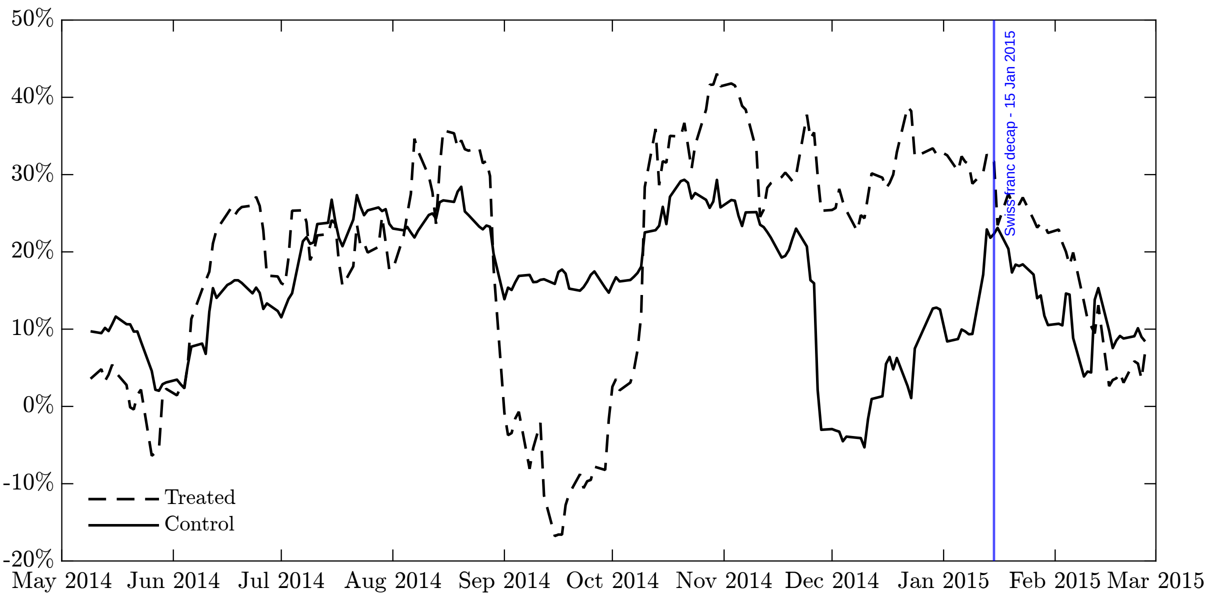

Note: This figure plots the cross-sectional average of 30-day rolling window correlations of daily VLOOP (i.e., shadow cost of intermediary constraints) and dealer-intermediated trading volume VLM. The "Treated" group comprises currency pair triplets that involve the Swiss franc (i.e., EUR-USD-CHF and GBP-USD-CHF) and the "Control" group contains the remaining 13 triplets.

Figure D.2: Event study: $cor(TCOST, VLM)$ around the Swiss franc decap

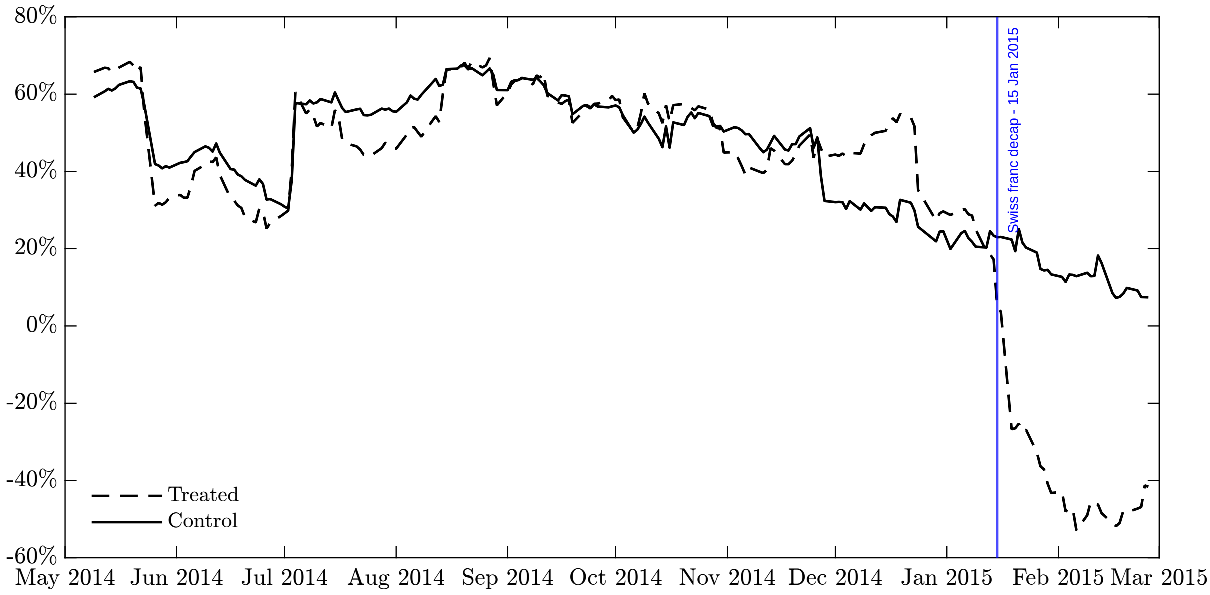

Note: This figure plots the cross-sectional average of 30-day rolling window correlations of daily TCOST (i.e., dealers' compensation for enduring inventory imbalances) and dealer-intermediated trading volume VLM. The "Treated" group comprises currency pair triplets that involve the Swiss franc (i.e., EUR-USD-CHF and GBP-USD-CHF) and the "Control" group contains the remaining 13 triplets.

Table D.1: Event study panel regression: Removal of the Swiss franc cap

	cor(VLOOP,Volume)			cor(TCOST,Volume)		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	***0.16	***0.16	***0.19	***0.49	***0.49	***0.49
	[11.42]	[11.40]	[11.55]	[23.15]	[23.21]	[26.74]
$D_{k,t}$	0.04	0.04	0.04	0.00	0.00	0.00
	[1.49]	[1.46]	[1.63]	[0.13]	[0.05]	[0.14]
$Post_t$	*−0.04	*−0.04	−0.03	***−0.35	***−0.35	***−0.34
	[1.93]	[1.94]	[1.23]	[12.36]	[12.33]	[11.48]
$D_{k,t} \times Post_t$	0.00	0.00	0.05	***−0.54	***−0.54	***−0.53
	[0.03]	[0.02]	[1.17]	[18.99]	[19.33]	[19.00]
Realised variance		***0.00	***0.02		***−0.01	**−0.01
		[6.16]	[5.48]		[6.40]	[2.51]
Amihud (2002)			***−0.04			−0.02
			[4.55]			[1.62]
R^2 in %	29.35	29.37	31.19	86.11	86.19	86.27
Adj. R^2 in %	29.28	29.28	31.08	86.09	86.18	86.25
Avg. #Time periods	207	207	207	207	207	207
#Exchange rates	15	15	15	15	15	15

Note: This table reports results from daily panel regressions of the form $\rho_{k,t} = \alpha + \eta_1 D_{k,t} + \eta_2 Post_t + \eta_3(D_{k,t} \times Post_t) + \kappa' w_{k,t} + \epsilon_{k,t}$, where the dependent variable is the 30-day rolling window correlation of our liquidity cost measure (i.e., VLOOP or TCOST) and trading volume (i.e., VLM), α denotes the intercept, $D_{k,t}$ is equal to one if currency pair triplet k contains the Swiss franc, $Post_t$ is one for the time period after the removal of the Swiss franc cap on 15 January 2015, and η_3 is the difference-in-differences coefficient. $w_{k,t}$ collects additional control variables such as the realised variance or Amihud (2002) price impact measure. The sample covers the period from 9 May 2014 to 26 February 2015. The test statistics based on Driscoll and Kraay (1998) robust standard errors allowing for random clustering and serial correlation (using the plug-in procedure for automatic lag selection by Newey and West, 1994) are reported in brackets. Asterisks *, **, and *** denote significance at the 90%, 95%, and 99% levels.

Appendix E. Additional empirical results

Figure 2: Identifying a triangular arbitrage opportunity

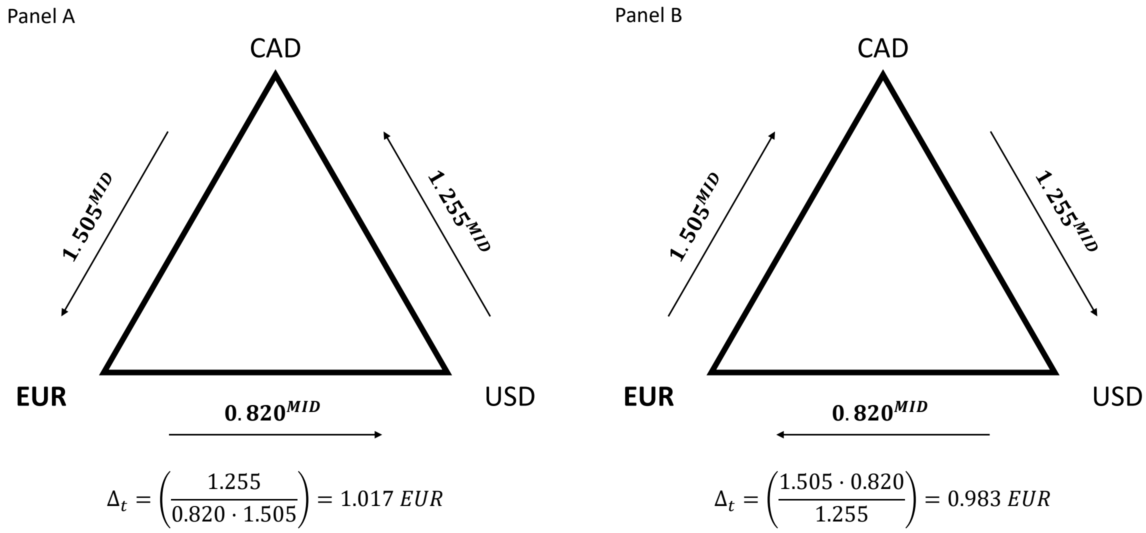

Note: This figure provides a schematic overview of two triangular arbitrage strategies, where the arrows denote the direction. Panel A shows the hypothetical profit of a trader starting with one euro, first exchanging it to $\frac{1}{0.820} = 1.220$ US dollars, then exchanging 1.220 US dollars to Canadian dollars at the midquote price of 1.255 Canadian dollars per US dollar. This yields 1.531 Canadian dollars that are exchanged back to euros at the CADEUR midquote that is equivalent to $\frac{1}{EURCAD^{MID}} = \frac{1}{1.505}$. Such a round trip yields 1.017 euros or equivalent a positive return of 1.7% in this example. Panel B embraces the same logic but going the opposite direction, that is, first from euro to Canadian dollar, to US dollar and then back to euro yielding a negative return of -1.7%.

Figure 3: Triangular arbitrage trade with transaction costs

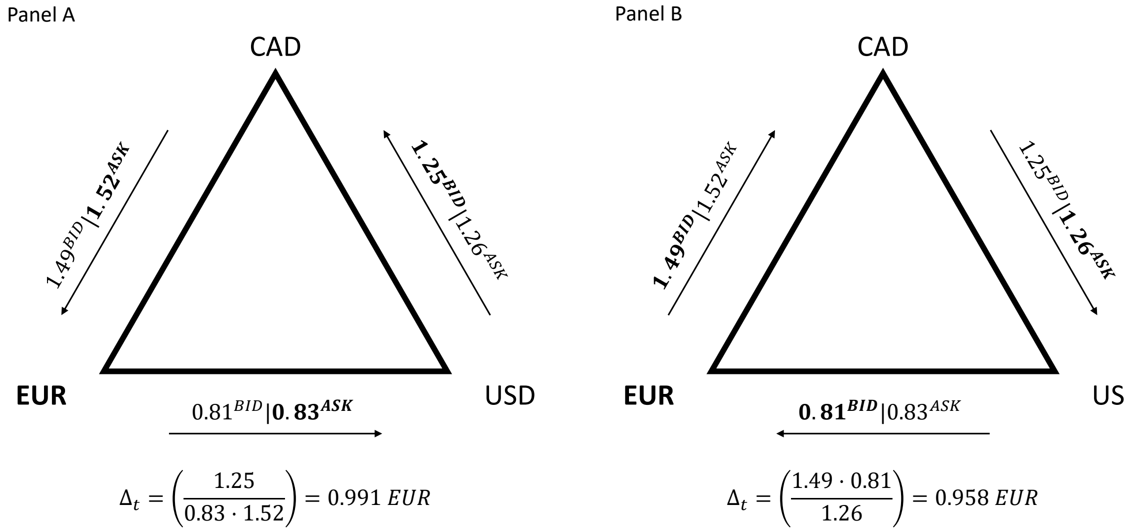

Note: This figure provides a schematic overview of two triangular arbitrage strategies, where the arrows denote the direction. Panel A shows the profit of first exchanging one euro to $\frac{1}{0.83} = 1.21$ US dollars at the ask price, then exchanging 1.21 US dollars to Canadian dollars at the bid price of 1.25 Canadian dollars per US dollar. This yields 1.51 Canadian dollars that are exchanged back to euros at the CADEUR bid price that is equivalent to $\frac{1}{EURCAD^{ASK}} = \frac{1}{1.52}$. Such a round trip yields 0.991 euros or equivalent a negative return of -0.9%. Panel B embraces the same logic but going the opposite direction, that is, first from euro to Canadian dollar, to US dollar and then back to euro.

Figure 4: Law of one price violations, intermediated volume, dealer constraints

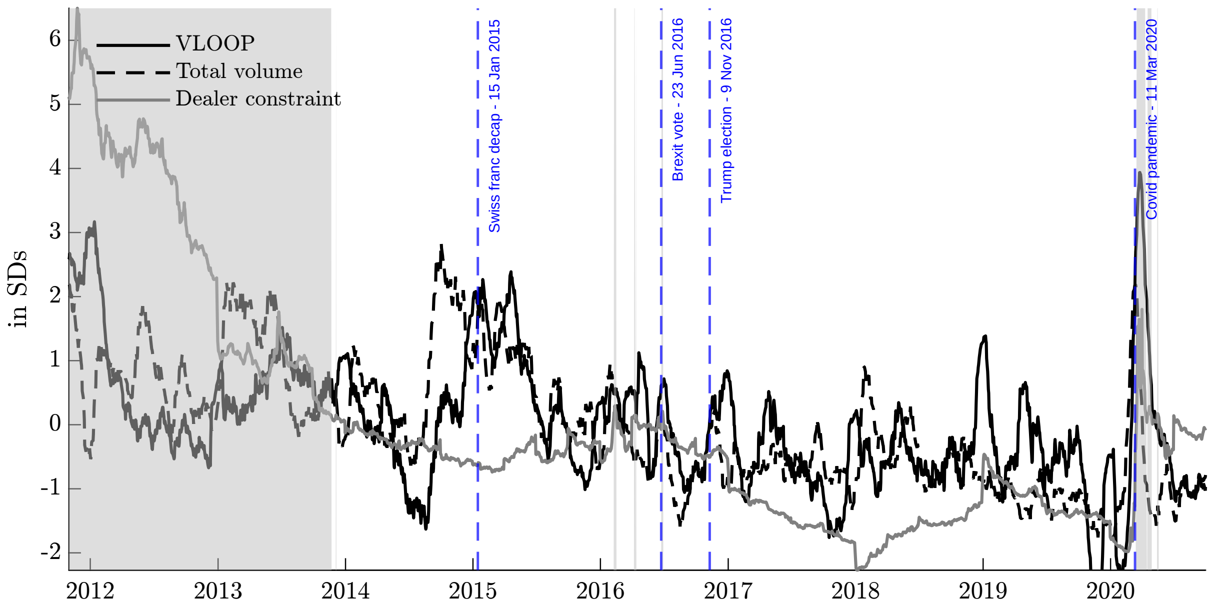

Note: This figure plots the cross-sectional average of the shadow cost of intermediary constraints (VLOOP, black solid line) against total trading volume (Total volume, black dashed line) in units of standard deviations. Both time-series correspond to 22-day moving averages. The grey line plots our dealer constraint measure (DCM). The grey shaded areas correspond to times when DCM exceeds its 75% percentile. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure 5: Round-trip trading costs, intermediated volume, dealer constraints

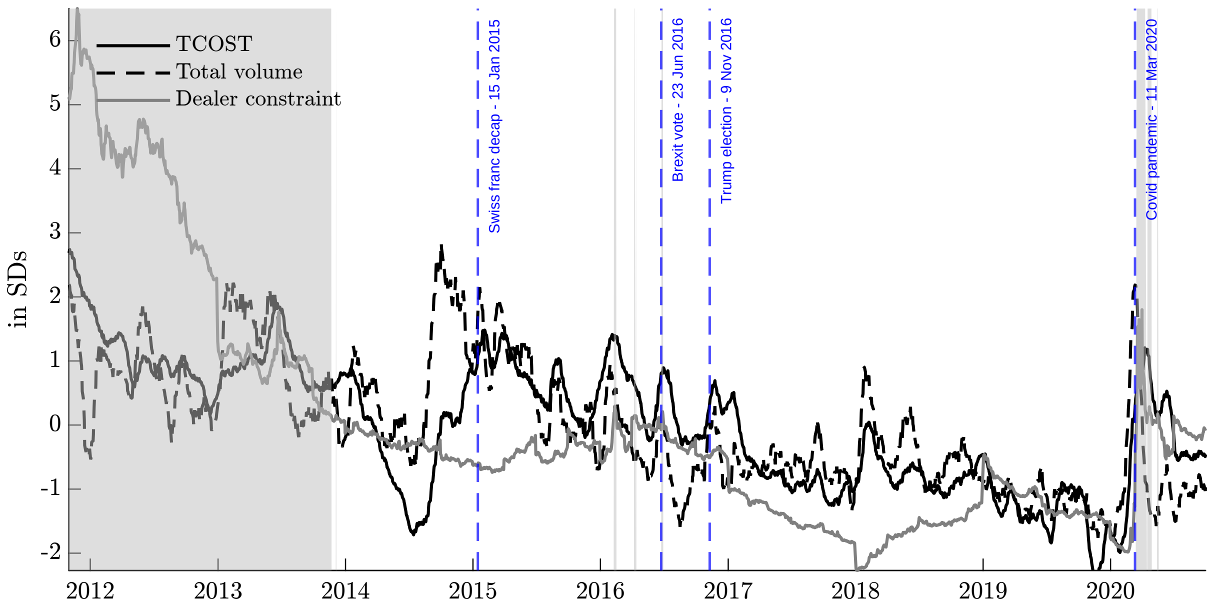

Note: This figure plots the cross-sectional average of dealers' compensation for enduring inventory imbalances (TCOST, black solid line) against total trading volume (Total volume, black dashed line) in units of standard deviations. Both time-series correspond to 22-day moving averages. The grey line plots our dealer constraint measure (DCM). The grey shaded areas correspond to times when DCM exceeds its 75% percentile. The sample covers the period from 1 November 2011 to 30 September 2020.

Table E.1: Summary statistics

	Liquidity cost in bps		Volume in \$bn		Bid-ask spread in bps		Volatility in bps	VLOOP>TCOST in %
	VLOOP	TCOST	Direct	Synthetic	Direct	Synthetic	Direct	
AUD-USD-JPY	0.24	4.88	0.18	5.11	4.15	5.87	14.38	0.18
AUD-USD-NZD	0.29	5.85	0.09	2.01	4.44	7.43	9.32	0.02
CAD-USD-JPY	0.30	4.67	0.03	5.32	4.29	5.21	12.66	0.43
EUR-USD-AUD	0.19	4.52	0.14	7.72	3.54	5.64	11.54	0.04
EUR-USD-CAD	0.28	4.25	0.08	7.94	3.55	4.99	10.15	0.07
EUR-USD-CHF	0.21	3.98	0.37	6.76	2.62	5.41	6.38	0.10
EUR-USD-DKK	0.14	3.89	0.09	6.17	2.54	5.30	1.82	0.05
EUR-USD-GBP	0.19	4.07	0.61	8.16	3.19	4.95	9.52	0.03
EUR-USD-JPY	0.21	3.90	0.65	9.67	3.14	4.83	11.43	0.66
EUR-USD-NOK	0.26	7.69	0.24	6.25	6.25	9.40	11.01	0.05
EUR-USD-SEK	0.23	6.86	0.27	6.27	5.41	8.42	9.18	0.05
GBP-USD-AUD	0.20	5.08	0.04	3.60	4.22	5.99	12.53	0.02
GBP-USD-CAD	0.29	4.69	0.03	3.81	4.00	5.34	10.85	0.05
GBP-USD-CHF	0.19	4.94	0.03	2.64	4.09	5.76	10.69	0.03
GBP-USD-JPY	0.19	4.47	0.20	5.55	3.85	5.18	12.78	0.62

Note: This table reports the time-series average of hourly triangular no-arbitrage deviations $VLOOP$ in basis points (bps), round-trip trading costs $TCOST$ in bps, direct trading volume in non-dollar pairs (e.g., AUDJPY) in \$bn, synthetic volume in dollar pairs in \$bn, direct and synthetic relative bid-ask spreads, and realised volatility in non-dollar pairs in bps. By "synthetic" we refer to the sum of trading volumes and relative bid-ask spreads in two dollar pairs (e.g., USDAUD and USDJPY) within a currency pair triplet. The last column shows the relative share of VLOOP>TCOST in %. Each row corresponds to a triplet of currency pairs, for example, AUDJPY, USDAUD, and USDJPY that we abbreviate as AUD-USD-JPY. The sample covers the period from 1 November 2011 to 30 September 2020.

Table E.2: Comparison EBS vs Olsen bid and ask quotes

	RMSE		MAE		CORR	
	BID	ASK	BID	ASK	BID	ASK
AUDJPY	0.286	0.232	0.131	0.126	0.996	0.997
AUDNZD	0.001	0.001	0.001	0.001	0.999	0.999
AUDUSD	0.001	0.002	0.001	0.001	0.998	0.997
CADJPY	0.335	0.315	0.148	0.147	0.995	0.996
EURAUD	0.003	0.002	0.002	0.002	0.998	0.998
EURCAD	0.002	0.002	0.001	0.001	0.998	0.998
EURCHF	0.001	0.001	0.001	0.001	0.993	0.992
EURDKK	0.001	0.001	0.001	0.001	0.991	0.988
EURGBP	0.001	0.001	0.001	0.001	1.000	1.000
EURJPY	0.211	0.201	0.124	0.124	0.999	0.999
EURNOK	0.016	0.012	0.008	0.008	0.996	0.998
EURSEK	0.009	0.009	0.006	0.006	0.999	0.999
EURUSD	0.002	0.002	0.001	0.001	0.997	0.998
GBPAUD	0.004	0.005	0.003	0.003	1.000	0.999
GBPCAD	0.004	0.004	0.003	0.003	0.999	0.999
GBPCHF	0.003	0.003	0.002	0.002	0.999	0.999
GBPJPY	0.523	0.590	0.251	0.257	0.999	0.999
GBPUSD	0.002	0.003	0.001	0.001	1.000	0.999
NZDUSD	0.001	0.001	0.001	0.001	0.999	0.999
USDCAD	0.002	0.002	0.001	0.001	0.999	0.999
USDCHF	0.001	0.001	0.001	0.001	0.998	0.998
USDDKK	0.007	0.007	0.005	0.005	0.999	0.999
USDJPY	0.178	0.194	0.112	0.111	1.000	0.999
USDNOK	0.014	0.016	0.010	0.010	0.998	0.998
USDSEK	0.018	0.013	0.009	0.008	0.999	0.999

Note: This table reports the root mean squared error ($RMSE$, columns 1 and 2), the mean absolute error (MAE, columns 3 and 4), and the pairwise correlation coefficient ($CORR$, columns 5 and 6) for bid and ask quotes based on EBS and Olsen data, respectively. The sample covers the period from 4 January 2016 to 30 December 2016.

Table E.3: Correlations in percent

	VLOOP	TCOST	VOD	VOS	BAD	BAS	RVD
TCOST	***28.10						
VOD	***−0.48	***6.19					
VOS	***4.93	***16.53	***61.87				
BAD	***23.24	***74.34	***1.85	***9.59			
BAS	***20.68	***75.85	***17.01	***34.91	***83.36		
RVD	***15.52	***37.78	***27.43	***37.57	***54.60	***49.05	
RVS	***13.22	***44.43	***31.47	***55.92	***45.59	***69.84	***74.85

Note: This table reports the pairwise correlation coefficient of hourly triangular no-arbitrage deviations $VLOOP$, trading costs $TCOST$, direct trading volume VOD in non-dollar pairs (e.g., XXXYYY), synthetic trading volume VOS in dollar pairs (e.g., the average across USDXXX and USDYYY), relative bid-ask spread BAD and realised volatility RVD in non-dollar pairs, as well as relative bid-ask spreads BAS and realised volatility RVS in dollar currency pairs in percent (%). Significant correlations at the 90%, 95%, and 99% levels are represented by asterisks *, **, and ***, respectively. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure E.1: Comparison of no-arbitrage violations using EBS vs Olsen data

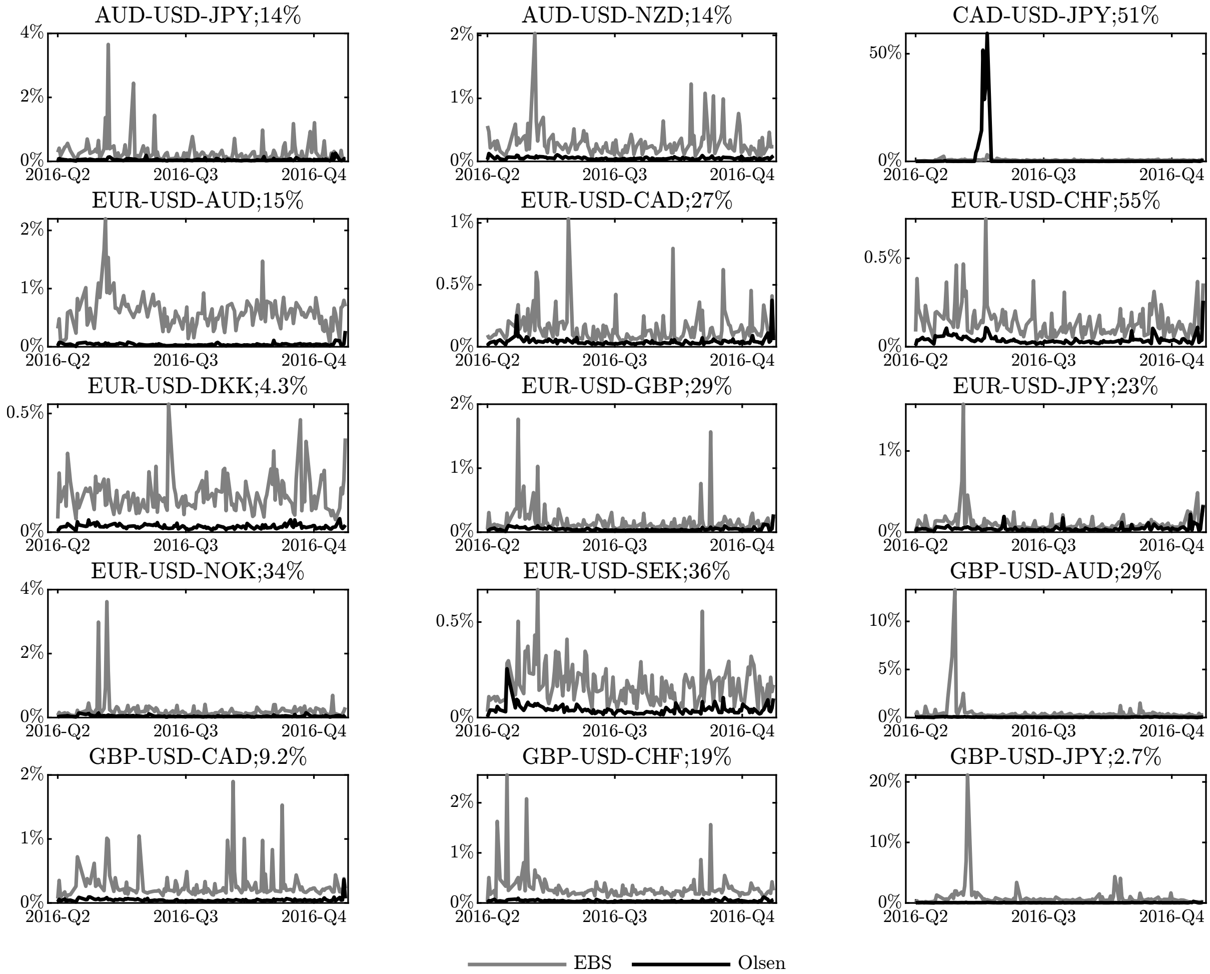

Note: This figure plots the daily cumulative sum of hourly no-arbitrage deviations (VLOOP) computed based on EBS and Olsen data, respectively. The percentages in the titles report the Pearson correlation coefficient between the two time-series. The sample covers the period from 8 June 2016 to 30 December 2016.

Figure E.2: No-arbitrage violations and trading volumes

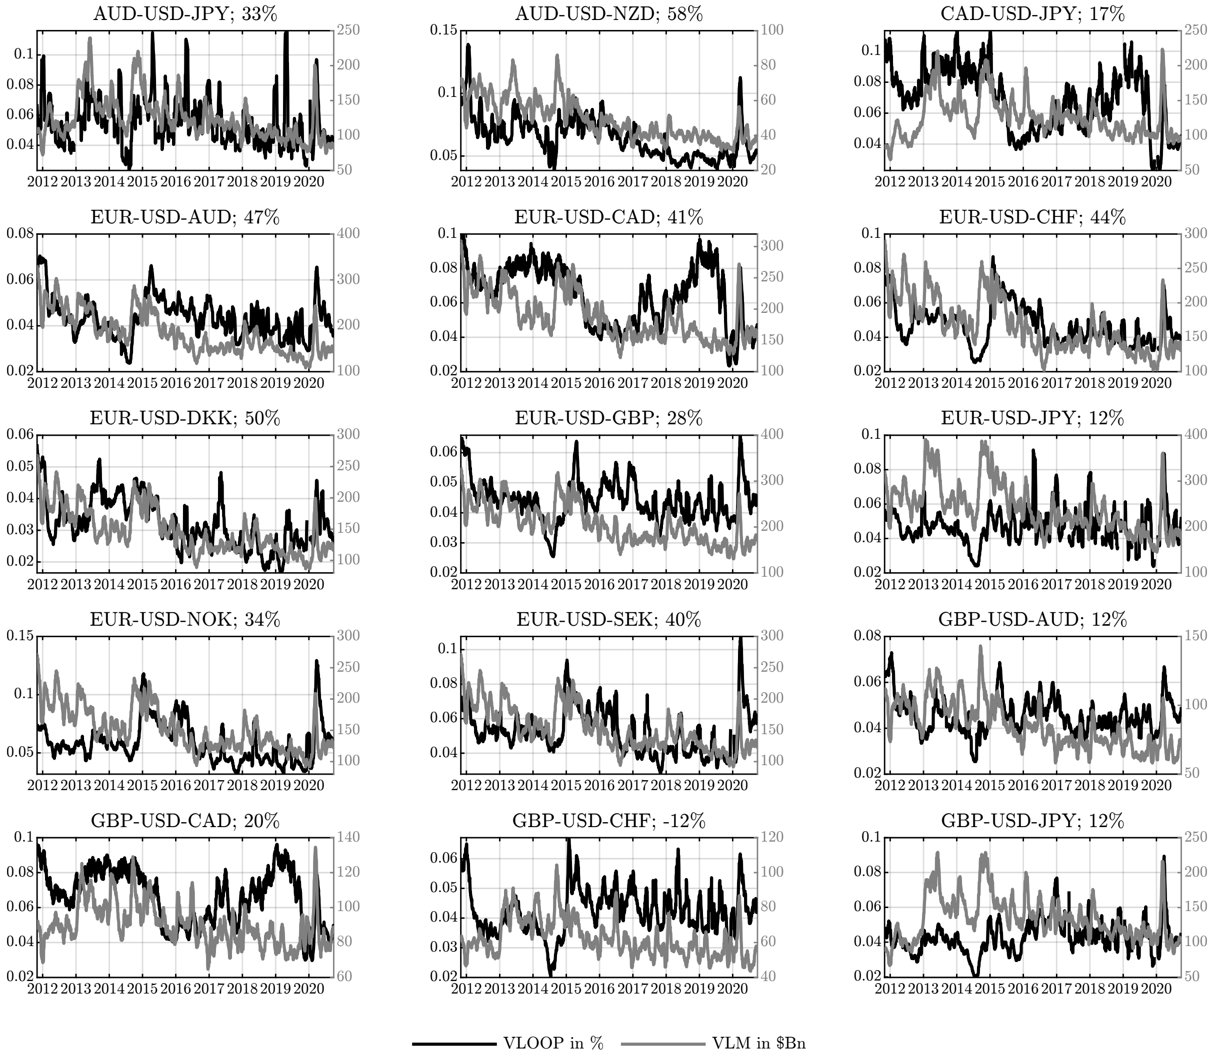

Note: This figure plots total trading volume VLM against no-arbitrage deviations VLOOP (i.e., shadow cost of intermediary constraints) for 15 triplets of currency pairs. Currency pair triplets are denoted as XXX-USD-YYY, consisting of two dollar currency pairs (i.e., USDXXX, and USDYYY) as well as one non-dollar currency pair (i.e., XXXYYY). The percentages in the titles report the Pearson correlation coefficient between VLM and VLOOP. Both time-series correspond to 22-day moving averages. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure E.3: Round-trip transaction costs and trading volumes

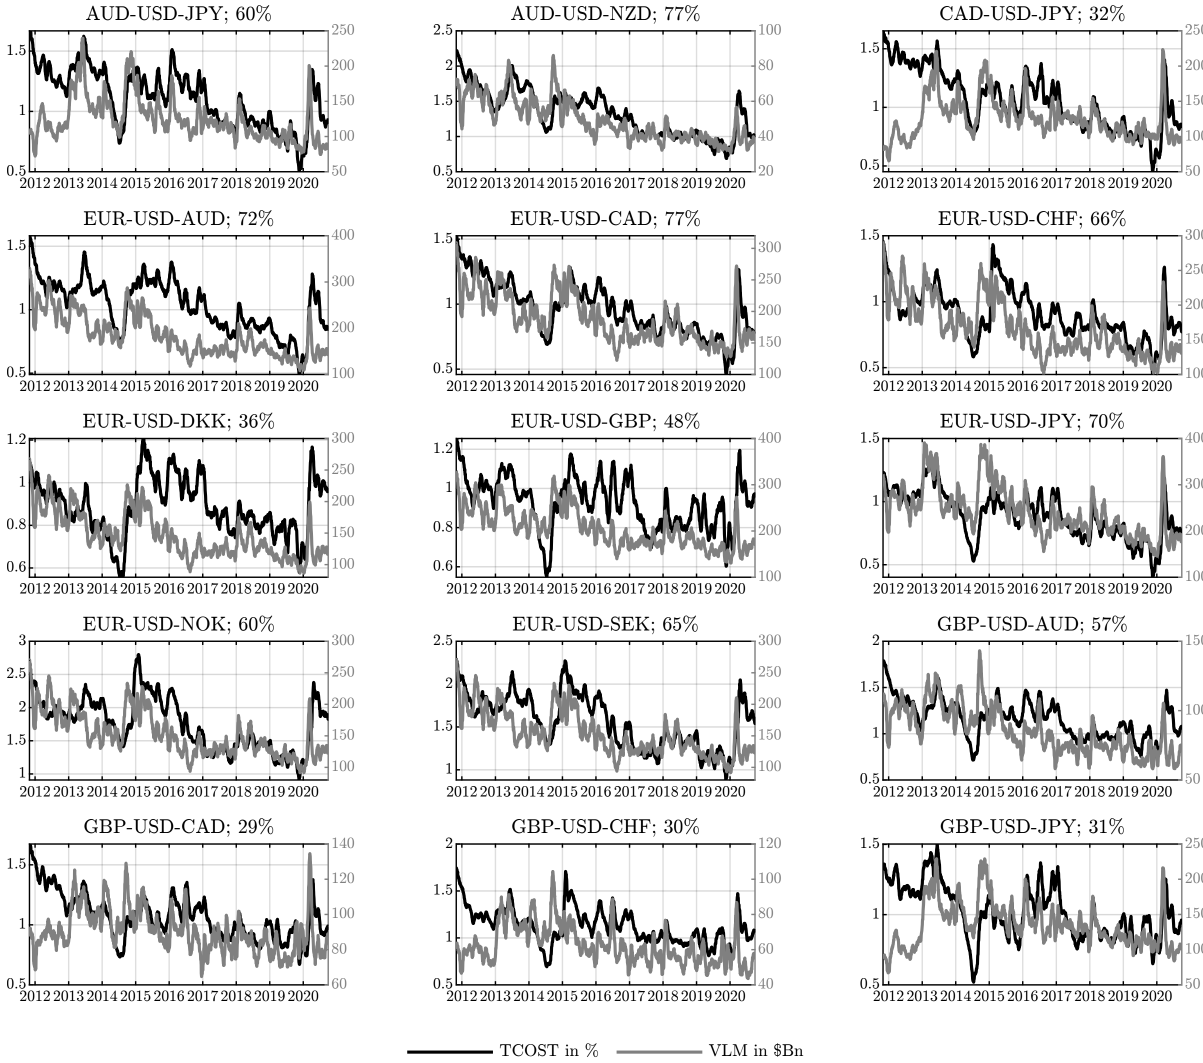

Note: This figure plots total trading volume VLM against round-trip transaction cost TCOST (i.e., dealers' compensation for enduring inventory imbalances) for 15 triplets of currency pairs. Currency pair triplets are denoted as XXX-USD-YYY, consisting of two dollar currency pairs (i.e., USDXXX, and USDYYY) as well as one non-dollar currency pair (i.e., XXXYYY). The percentages in the titles report the Pearson correlation coefficient between VLM and TCOST. Both time-series correspond to 22-day moving averages. The sample covers the period from 1 November 2011 to 30 September 2020.

Figure E.4: Autocorrelated trading volume

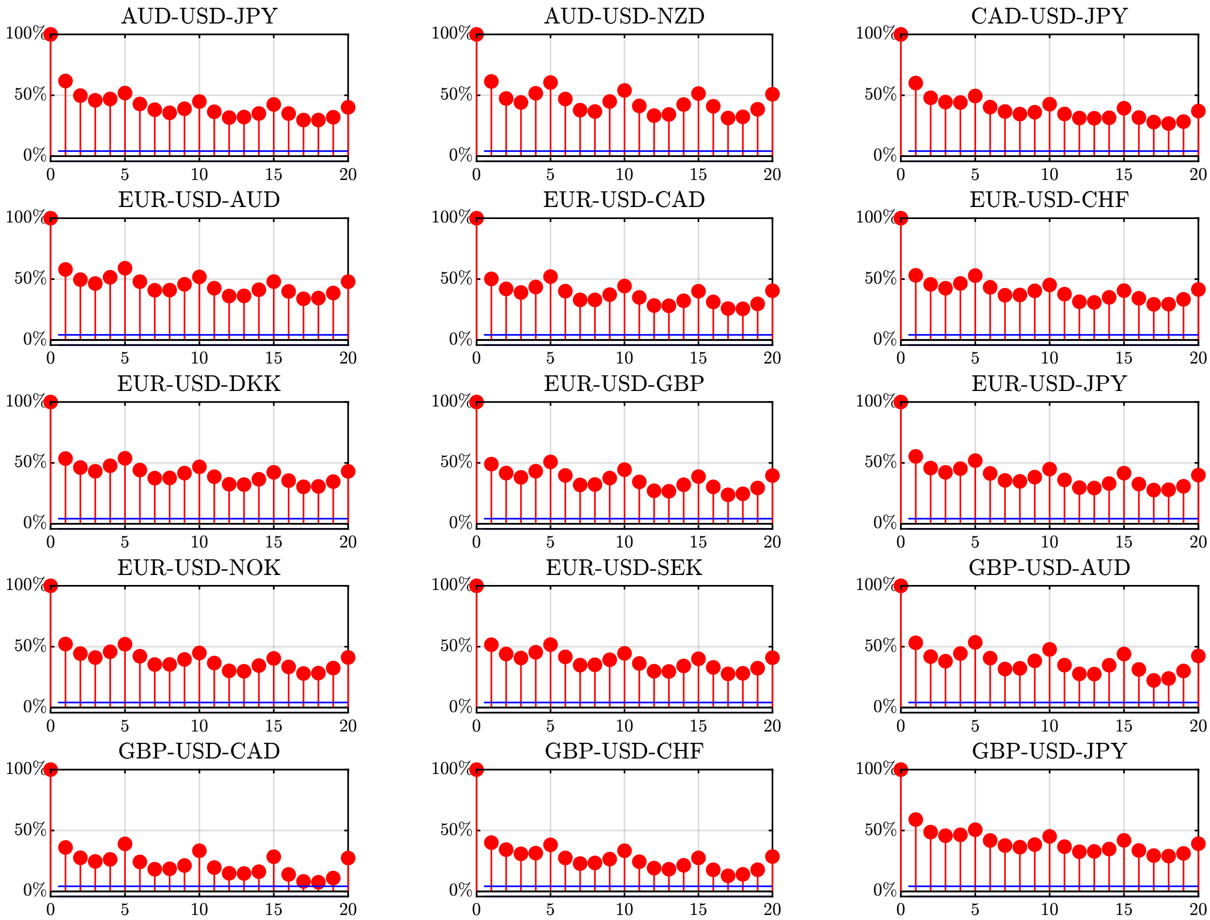

Note: This figure plots the autocorrelation coefficient of total dealer-provided trading volume VLM for 15 triplets of currency pairs. Currency pair triplets are denoted as XXX-USD-YYY, consisting of two dollar currency pairs (i.e., USDXXX, and USDYYY) as well as one non-dollar currency pair (i.e., XXXYYY). The solid lines are approximate 95% confidence bounds. Both time-series correspond to 22-day moving averages. The sample covers the period from 1 November 2011 to 30 September 2020.

References: Online Appendix

Andrews, D. W. K. and Monahan, J. C., 1992. An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica*, 60(4):953–966.

Driscoll, J. C. and Kraay, A. C., 1998. Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

He, Z., Kelly, B., and Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.

Newey, W. K. and West, K. D., 1994. Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, 61(4):631–653.