

# FX Futures Invariance

Torben G. Andersen, Oleg Bondarenko, Eleni Gousgounis, and Esen Onur\*

November 7, 2025

## Abstract

We develop and empirically examine a novel intraday invariance relation among trading activity variables across a set of currencies in the foreign exchange (FX) futures market exploiting tick level data. Our Trading Invariance (TI) hypothesis reflects an equilibrium involving trade-offs between the risk per trade versus market liquidity indicators as captured by the bid-ask spread and market depth. We observe exogenous breaks in the trading costs for many of the currencies, as the exchange lowered the tick size for individual contracts at different points in time. Using these incidents as quasi-natural experiments, we confirm that the remaining components of the TI hypothesis adjust to maintain the equilibrium by shifts in the trading intensity and market depth. In contrast, alternative invariance relations fail dramatically when confronted with this type of exogenous shock. We further test the value of a subset of the coefficients in the TI relation, finding a good correspondence to the theoretically predicted values. Our results point towards a fundamental market equilibrium mechanism that should be accommodated within existing market microstructure paradigms. Moreover, they have direct implications for market design and surveillance.

*JEL classification:* G10, G23

*Keywords:* Testing for Market Invariance, Market Microstructure, Tick Size

---

\*Andersen: Kellogg School, Northwestern University; NBER; CoRE, Aarhus; Canadian Derivatives Institute; and Consultant at the U.S. Commodity Futures Trading Commission; Bondarenko: University of Illinois Chicago and Consultant at the U.S. Commodity Futures Trading Commission; Gousgounis: U.S. Commodity Futures Trading Commission, Washington, D.C. 20581; Onur: U.S. Commodity Futures Trading Commission, Washington, D.C. Andersen acknowledges research support from the Kenney Fund at the Kellogg School. The analyses and conclusions expressed in this paper are those of the authors in their official capacity and do not reflect the views of other Commission staff, the Office of the Chief Economist, or the Commission. All errors and omissions, if any, are the authors' responsibility.

# 1. Introduction

During regular trading hours, the main financial markets brim with activity, facilitating trades among agents with diverse speculative, hedging, portfolio allocation, and liquidity motives. Throughout, investors, dealers, and market makers monitor the news flow and market dynamics – the flow of trades, transaction sizes, price changes, bid-ask spread, order-book depth, and return volatility – for signals regarding innovations to the underlying “efficient” asset price. These market activity indicators are known to be interdependent, yet there is no simple bivariate relation between any pair of these variables that remains stable over time or across asset classes.

We rely on fundamental principles regarding stable economic relations, coupled with empirical scrutiny, to develop a novel trading invariance relation. The resulting equilibrium involves trade-offs between risk and liquidity, involving six separate market activity variables. Although we expect frictions to generate deviations from this theoretical relation, the null hypothesis stipulates that, absent market imperfections, it applies universally across time and assets.

To facilitate empirical analysis, we develop high-frequency measurement and inference procedures to explore invariance-style relations in a multiple asset setting, focusing on the main foreign exchange (FX) futures contracts traded at the CME Group, namely the euro (EUR), British pound (GBP), Japanese yen (JPY), Australian dollar (AUD) and Canadian dollar (CAD). Our approach enables joint testing of existing hypotheses and, more importantly, systematic extensions of prior models, guided by conjectures regarding the inherent features of market invariant quantities. We find strong corroborative evidence for a specification we denote the “Trading Invariance” (TI) hypothesis, involving equilibrium interactions amongst separate dimensionless components related to return volatility per transaction, marginal trading costs, and top-of-book liquidity. This relation generalizes the ubiquitous strong volatility-trading activity covariance which has motivated extensions into the so-called Mixture-of-Distributions Hypothesis (MDH) models. It also embeds alternative high-frequency variants of the market Microstructure Invariance (MI) theory, labeled Intraday Trading Invariance (ITI) by [Andersen et al. \(2018\)](#).

Our focus on a set of CME Group FX futures is motivated by their limited and comparable degree of frictions. One, these contracts trade exclusively on the GLOBEX platform. Hence,

we avoid synchronization issues with processing delays and split orders executed on different exchanges. Two, we have access to messages identifying each marketable order that is executed against one or more resting limit orders. This ensures coherent recording of trade times and sizes by the (active) party initiating the transaction. Three, the trading environment is fairly stable, with changes in regulation and trading costs being readily observed, and the underlying variables not displaying any pronounced low-frequency trends over the sample period. Four, compatibility across the FX contracts is enhanced by the exchange maintaining a rough equivalence in contract size, tick size, order types, and regulation. Five, the markets are liquid and trade almost round-the-clock five days per week. Finally, there are distinct shocks to the average trading costs for individual contracts induced by changes to the minimum tick size. They serve as “pseudo-natural experiments,” allowing us to gauge the robustness of alternative invariance hypotheses.

Hence, the FX futures market lets us explore the cross-asset features of the invariance relations in a transparent setting with liquid and distinct, yet fairly homogenous, assets. This helps us assess the explanatory power of the relations, while sidestepping issues of how they must be modified in the face of rapid evolution in market structure and trading technology. Moreover, the dispersed timing of exogenous changes to the tick size for individual currencies enables us to study the dynamic response of the invariance relations to an abrupt shift in market design.

We hypothesize that the TI relation reflects interacting forces that sustain market equilibrium during the ongoing price discovery process. Specifically, it characterizes longer-term relations among the market activity variables that help us anticipate the implications of changes in market design, the impact of exogenous shocks to the market environment, and the emergence of market tensions by monitoring the deviations from the established equilibrium, thus potentially foreshadowing turbulent or disfunctional market conditions. Finally, we uncover proportional discrepancies in the TI relation across the FX futures contracts that line up with measures of liquidity, suggesting that they proxy for the strength of microstructure frictions.

We initially review time series and intraday features of the FX futures activity variables that are critical for our empirically-guided invariance hypothesis development.

### 1.1. Illustration of the Market Dynamics


In this section, for brevity, we focus solely on the British pound (GBP) and euro (EUR). They provide contrasting pictures concerning the evolution of some key market attributes, yet display mutually consistent patterns for the remaining activity variables.

Figure 1 depicts the daily averages of the futures price for British pound, expressed in U.S. dollar ( $P$ ), along with the return volatility ( $\sqrt{S} = \sigma$ ), number of transactions ( $N$ ), trade size ( $Q$ ), bid-ask spread ( $B$ ), and depth at the top of the order book ( $D$ ). The red curve, superimposed on some panels, represents a 21-day moving average of the underlying series. The measures are based on 5-minute observations across the trading day, as detailed in Section 4.

The top left panel indicates that the pound depreciated versus the U.S. dollar over the sample period, but mostly in an orderly fashion, with the most dramatic moves being the drops around the Brexit vote on June 23, 2016, and the onset of the pandemic in early 2020. The primary activity variables - the volatility, transaction intensity, trade size, and market depth - display strong persistence, yet also pronounced idiosyncratic fluctuations. Nonetheless, they all appear compatible with a stationary dynamic. The most striking features are the strong correlation among the trade size and order book depth and the drop in depth in response to the Brexit and pandemic events. Finally, the bottom left panel shows that the tick size is “large.” The daily average bid-ask spread is only a fraction larger than the tick size, implying that the spread mostly equals a single tick. Thus, the tick size is likely binding throughout most trading days.

Figure 2 provides the identical series for the euro. It also depreciates against the U.S. dollar, but the rate has been quite stable since 2015. The volatility and trading intensity series are correlated and display the same characteristics as in Figure 1. In contrast, the bottom row conveys a very different evolution. The euro tick size was lowered on January 11, 2016. This is accompanied by a large drop in the spread and a plunge in market depth. The marginal transaction cost (bid-ask spread) declines, but there is an offsetting drop in top-of-book depth. This is not surprising – limit orders may now be picked off by swift moving traders following smaller adverse price moves than before, so expected losses for liquidity providers increase, while compensation from the spread shrinks. Likewise, the typical transaction is smaller after the tick change, although the effect is less pronounced than for the depth. Overall, this suggests that a

Figure 1: GBP Futures — Daily Time Series




**Daily Time Series Averages for the GBP Futures Contract over the Full Sample.**  $P$  is the FX Futures Price in USD per one GBP;  $\sigma$  is the annualized RV;  $N$  is the number of transactins in 5 minutes;  $Q$  is the number of contracts per transaction; Bid-Ask Spread (blue) and Tick Size (green) are in USD per one GBP;  $D$  is the number of contracts available at best bid and ask. When shown, the red curve represents a 21-day moving average of the underlying series.

lower tick size leads to deteriorating depth, but without any notable impact on return volatility or trading intensity. In contrast, the endogenous increase in the spread and drop in book depth, experienced at the onset of the pandemic, is accompanied by heightened volatility and trading activity. Evidently, the association between tick size, spread, and liquidity responds differently to endogenous shifts in market conditions compared to exogenous regulatory shocks.

Additional perspectives follow from the intraday activity patterns. Figure 3 depicts the mean value of each variable across the euro sample for all 5-minute intervals during the continuous

Figure 2: EUR Futures — Daily Time Series




**Daily Time Series Averages for the EUR Futures Contract over the Sample.**  $P$  is the FX futures Price in USD per one EUR;  $\sigma$  is the annualized RV;  $N$  is the Number of Trades in 5 minutes;  $Q$  is the number of contracts per transaction; Bid-Ask Spread (blue) and Tick Size (green) are in USD per one EUR;  $D$  is the number of contracts available at best bid and ask. When shown, the red curve represents a 21-day moving average of the underlying series.

GLOBEX trading session from  $-5$  (i.e., 7:00 pm Chicago Time (CT) on the preceding day) to 15 (3:00 pm CT). The dashed vertical lines separate the main trading zones. Early on, activity is concentrated in Asia, then moves to Europe and finally North America. The two curves in each panel reflect two non-overlapping subsamples - from the start of the sample until January 11, 2016, and from that date through 2021 – before and after the reduction in tick size.

Trading activity is particularly subdued early in the day, but we still observe the strong

Figure 3: EUR Futures — Intraday Activity Series



**Intraday Averages for EUR Futures Before and After the Tick Size Change.** Volume  $V$  is the number of contracts traded in 5 minutes. The other variables are as indicated in Figures 1 and 2. The averages are computed for each 5-minute interval across the trading day. The tick change occurred on Jan 11, 2016. The averages before (after) that date is indicated by the blue (green) line.

intraday correlation between the return volatility, trading intensity and volume. This confirms the impression from Figure 2 that the overall levels for  $V$ ,  $\sigma$ , and  $N$  did not change much across the two tick size regimes. On the other hand, the bottom left panel shows that the bid-ask spread is very close to the minimal one tick throughout each subsample, so the reduced tick size has a large permanent effect of the marginal trading cost. Similarly, the market depth drops dramatically, when the spread is lower. Finally, the trade size is lower in the latter subsample, but from Figure 2 it is not evident that this variable undergoes any permanent significant shift

in response to the tick change. The corresponding display for the pound is provided in Figure A1 of Appendix A. There was no tick change for GBP, and we do not see much qualitative difference across the two subsamples, corroborating the hypothesis that the striking discrepancies for some of the euro activity variables, indeed, were due to the lowering of the tick size.

## 1.2. Related Literature

The high-frequency variant of MI, explored in Andersen et al. (2018), constitutes a benchmark for our transaction level invariance hypothesis. The number of empirical MI studies remains limited. Kyle and Obizhaeva (2016b) study portfolio transitions with multiple-day durations. Kyle and Obizhaeva (2016a) rationalize market crashes through large (hidden) selling pressure that disturb the MI equilibrium. Evidence regarding observable proxies for the speed of business time within the MI setting is explored by Kyle et al. (2017) through the number of news releases for a firm per month, by Kyle et al. (2020) via the number of monthly prints for U.S. stocks, and by Bae et al. (2020) using the number of buy-sell switching points in retail trading accounts for South Korean stocks. All these tests concern equity market data over daily or monthly horizons. Moreover, Bucci et al. (2020) explore MI using a large set of metaorders executed within one trading day. Unconditionally, MI works well, but the associated innovations display strong correlation with the trading cost of the metaorder. Their results suggest a critical role for the bid-ask spread and the price impact associated with sequential directional trades to fulfill large orders. In contrast, we test an invariance hypothesis jointly for a set of FX futures markets by exploring the implications of the transaction-level relationship for daily time series.

There are very few prior empirical studies dedicated to this type of transaction-based invariance theory. A high-frequency variant of the MI hypothesis is tested for the e-mini S&P 500 futures by Andersen et al. (2018), who find largely confirmatory evidence. A number of critical features are distinctive in the current work. First, our test involves multiple (five) liquid FX futures contracts. Second, we test the transaction-based invariance relation over a nine year period, much longer than in prior work. Third, at distinct point during the sample period, some contracts experience a halving of the tick size. Since the tick size is almost universally binding, this induces large reductions in trade costs – and lower compensation to liquidity providers –



through a lower bid-ask spread. These changes trigger an adjustment period until a new equilibrium is obtained. Fourth, we test the invariance relation through a set of moment conditions applying jointly across the FX contracts, thus exploiting additional information to investigate the generality of the relation. Fifth, our Trading Invariance (TI) Hypothesis builds on interactions among unitless quantities. This imposes discipline on the equilibrium relation, as any economic quantity denominated in, e.g., seconds or dollars, is unlikely to remain invariant when technology, market structure and organization evolve and alter trading costs, transaction speed, and the degree of cross-market integration. Finally, our approach pins down most of the coefficients governing the equilibrium interaction, but one is left unspecified and subject to empirical determination. In contrast, MI determines all coefficients through theoretical arguments.

[Benzaquen et al. \(2016\)](#) adopt key aspects of the [Andersen et al. \(2018\)](#) approach in studying a number of futures contract and equity markets. Their results are inspirational, suggesting a role for tick size or bid-ask spreads in the high-frequency invariance relation. However, as much of the prior MI-related literature, they rely on regression techniques that lead to biased inference due to endogeneity. We develop formal inference procedures and, in the process, document the shortcomings of the regression-based tests. While we also document a critical role for transaction cost variables, our TI relation is distinct in involving both the spread and order book depth. Moreover, we operate strictly within a homogeneous and stable institutional setting, which ensures high data quality and precise identification of transaction count and trade size, thus effectively controlling for confounding effects. Finally, recently, an alternative MI specification focusing on the determinants of the bid-ask spread is analyzed by [Hou et al. \(2024\)](#). They find corroboratory evidence for this invariance representation, but their regression test is severely biased due to endogeneity, as we document in Sections 2.3.1 and 5.2.2.

## 2. Trading Invariance

This section outlines the main features of our TI hypothesis. We focus on a single asset and given time interval. We address issues of temporal aggregation, multiple variable systems, and estimation of latent variables in Section 4.

## 2.1. Candidate Market Activity Variables

The notion of market invariance originated from the long-recognized correlation between price variation and transaction activity or trading volume. This feature is embodied in the most basic variants of the “Mixture-of-Distributions Hypothesis” (MDH) and has led scholars to introduce the concept of “business time,” reflecting the idea that economic events or financial activity evolve at a rate separate from the calendar clock. Given our objective of exploring stable interactions among market activity variables, a natural component to appear within a trading equilibrium is the  $S$ -to- $N$  ratio,

$$S/N,$$

where the return variation  $S = \sigma^2$  is a number (percent squared) per unit time and  $N$  likewise is a number, or count (transactions) per unit time. Thus, the ratio is dimensionless which, intuitively, should render the relation more robust to variation in external factors over time.<sup>1</sup>

As emphasized in much of the MDH literature, despite the unitless denomination, the  $S/N$ -ratio is unlikely to be (i) invariant for a given asset over long horizons or (ii) identical across all assets. For example, if one currency experiences secular growth in trading, while its return volatility remains stable - a scenario consistent with standard assumptions, as volatility typically is assumed to be stationary - then the  $S/N$ -ratio declines over time. The point is that, in this setting, there is no active economic forces to counterbalance factors that induce an exogenous shift in the  $S/N$  relation. Ergo, a robust equilibrium relation must include components that naturally absorb this type of systematic or secular changes in trading activity.<sup>2</sup>

Market microstructure theory and well-established empirical regularities provide guidelines for extensions to this basic relation.  $1/N$  reflects the average duration before the next transaction. Therefore, a higher  $S/N$  value is indicative of enhanced short term (business-time) risk. Active liquidity providers respond to higher risk by widening the bid-ask spread and reducing exposure by canceling limit orders. Hence, higher  $S/N$  should imply lower liquidity provision, i.e., increased trading costs and lower market depth. Thus, we introduce two additional dimen-

---

<sup>1</sup>In contrast, if the denominator is trading volume  $V = N \cdot Q$  or dollar volume  $P \cdot V = N \cdot P \cdot Q$ , then  $S/V$  and  $S/(P \cdot V)$  are denominated in inverse trade size and inverse nominal dollar units, respectively.

<sup>2</sup>A formal elaboration of this point is provided by [Tauchen and Pitts \(1983\)](#), who explicitly considers an expanding set of market participants along with other factors that may distort and disguise the  $S/N$  relation. A common procedure is to stochastically detrend the trading series and treat the MDH as a relation between volatility and “abnormal” trading activity, see, e.g., [Andersen \(1996\)](#).

sionless ratios that serve as liquidity measures related to cost and depth. First,

$$B/P,$$

captures the marginal trading costs of a market order, as manifest in the percentage bid-ask spread, where  $B$  denotes the dollar gap between the best ask and bid quotes and  $P$  is the dollar price, defined as the midpoint of the bid-ask spread.

Second, we introduce the (top-level) market depth measure,

$$D/Q,$$

where  $D$  denotes the number of contracts available at the best bid and ask quotes and  $Q$  is the average number of contracts exchanged per transaction. Thus, this measure captures book capacity, i.e., how many regular-sized transactions can be accommodated instantly.

## 2.2. Trading Invariance – TI

In specifying our invariance hypothesis, we draw on theory along with existing evidence regarding relevant activity variables. First, our focus on unitless quantities renders ratios of similarly denominated variables natural candidates. Second, we expect strong interdependencies between risk, marginal trading costs, and market depth, motivating our choice of the ratios introduced in Section 2.1. Third, our specification nests several invariance-style representations in the literature that each have documented a degree of empirical merit.

Consequently, for the intraday interval  $\tau$ , a natural invariance specification takes the form,

$$\mathcal{I}_\tau = \left( \frac{S_\tau}{N_\tau} \right) \left( \frac{B_\tau}{P_\tau} \right)^{\gamma_B} \left( \frac{D_\tau}{Q_\tau} \right)^{\gamma_D}, \quad (1)$$

where the stipulated invariant  $\mathcal{I}_\tau$  is a stationary and strictly positive random variable with constant (unconditional) mean  $C$  and finite higher order moments. The hypothesis is that endogenous interaction between the activity variables ensures the invariant – the  $\mathcal{I}_\tau$  realizations – will fluctuate around  $C$  over time. The unitary power coefficient for  $S/N$  reflects a normalization to ensure identification. Further, note that a change to the measurement unit for any of the activity variables, e.g., fractions versus percentage terms or cents versus dollars, merely scales each side of the invariance relation (1) by a multiplicative constant.

Alternative invariance-style hypotheses arise as special cases of equation (1), employing different auxiliary assumptions regarding the system. For example, if  $\gamma_{\mathcal{B}} = \gamma_{\mathcal{D}} = 0$ , the invariance relation only involves  $S/N$ , which we label the SN hypothesis. A related specification stipulates local proportionality between volatility and volume, i.e., the invariant involves  $S/V$ , which we refer to as the SV relation. Such bivariate volatility-trading activity specifications have been subject to a range of distinct theoretical representations and empirical implementations.<sup>3</sup> They typically emphasize that these systems are *ceteris paribus* relations, so the associations apply only locally, as they evolve with changes in the market and economic environment.<sup>4</sup> Moreover, the error term from these specifications are generally allowed to be heteroskedastic and display serial correlation, reflecting temporary shocks and underlying market frictions. In contrast, the theoretical underpinnings for MI stipulate that the invariant is an i.i.d. random variable, although institutional frictions may be considered in empirical work.<sup>5</sup>

Letting  $\mathbb{E}(\log(\mathcal{I}_\tau)) = -c$ , the equivalent log-linear invariance hypothesis is given as,

$$TI_\tau = (s_\tau - n_\tau) + \gamma_{\mathcal{B}} (b_\tau - p_\tau) + \gamma_{\mathcal{D}} (d_\tau - q_\tau) + c, \quad (2)$$

where  $TI_\tau$  is a mean zero stationary process with finite moments. We label this representation “Trading Invariance,” or TI. As discussed above, TI deviations may arise from persistent adjustments stemming from shocks to the institutional, regulatory, exchange, and technological environment along with other market- and asset-specific features. Hence, we allow for serial correlation and heteroskedasticity in  $TI_\tau$ . The aspiration is that the direct incorporation of market liquidity variables will serve to stabilize the invariance features of the system, thus providing a robust extension to the bivariate volatility-trading activity relations. This can be assessed in two distinct ways. First, at the individual asset level, we can check for improvements in the intertemporal fit of the specification. Second, we may explore the stability of the coefficient estimates across assets. If the level of trading intensity is associated with systematic variation in market liquidity, then the estimated system should speak to the endogenous adjustment occurring as more active trading generates lower transaction costs and enhanced market depth.

---

<sup>3</sup>See, e.g., Tauchen and Pitts (1983), Jones et al. (1994), Richardson and Smith (1994), Andersen (1996), Liesenfeld (2001), and Darolles et al. (2015, 2017) for such generalizations.

<sup>4</sup>An early example is Tauchen and Pitts (1983), who explicitly model the relation between return volatility and volume, when there is secular growth in the number of market participants.

<sup>5</sup>See, for example, the assessment of microstructure invariance in the U.S. stock market by Kyle et al. (2020).

The key question is whether, for our FX futures contracts operating within similar market environments, TI provides robust explanatory power, thus serving as a useful approximation to the dominant interactive forces amongst the activity variables. If this can be verified, it speaks to missing ingredients in existing microstructure invariance hypotheses, providing guidance for market design. It should also provide signals regarding emerging market malfunction, if the TI equilibrium condition, characterizing an ordinary market equilibrium, is unraveling.

A few additional comments on the TI relation (2) are warranted. First, most of the activity variables are readily observed within any given intraday interval, as long as it contains at least one transaction. The exception is the latent return variation,  $s$ . For the latter, we apply a standard realized volatility (RV) estimator, as detailed in section 4. Second, a consistent scaling of any of the activity variable across all currencies will induce the identical change to the intercept  $c$  in equation (2).<sup>6</sup> Thus, henceforth, we let  $c$  denote a generic constant whose value may differ across equations. This further implies that our invariance representation remains valid, even if the estimators for the activity variables display a proportional bias. Third, we emphasize that the TI relation (2) stipulates stationarity only for the residual term,  $TI_\tau$ . In general, the individual activity variables could display, e.g., stochastic trends. In that case, TI implies that there is a cointegrating relation among the activity variables.

### 2.3. A General Framework for Alternative Invariance Hypotheses

Our invariance representation (2) is related to a number of prior hypotheses concerning the interaction among market activity variables. To streamline the discussion, we generalize equation (2) by removing all coefficient restrictions, retaining only the normalization  $\gamma_s = 1$ ,

$$\mathbb{E}[(s_\tau + \gamma_n n_\tau) + (\gamma_b b_\tau + \gamma_p p_\tau) + (\gamma_d d_\tau + \gamma_q q_\tau)] = c. \quad (3)$$

This representation is not denomination-free, unless suitable restrictions are enforced, so it does not constitute a natural invariance relation. Nonetheless, it is useful, as it nests a number of existing hypotheses: TI applies if  $\gamma_n = -1$ ,  $(\gamma_b, \gamma_p) = (\gamma_b, -\gamma_b)$ , and  $(\gamma_d, \gamma_q) = (\gamma_d, -\gamma_d)$ . Likewise, the volatility-trading intensity relations SN and SV are obtained by imposing,  $\gamma_b = \gamma_p = \gamma_d = 0$ , along with, respectively,  $(\gamma_n, \gamma_q) = (-1, 0)$  or  $(\gamma_n, \gamma_q) = (-1, -1)$ .

---

<sup>6</sup>For a constant  $\eta > 0$ ,  $\log(\eta A_\tau) = \log(\eta) + a_\tau$ , so the scaling is absorbed into the intercept.

### 2.3.1. Microstructure Invariance

Kyle and Obizhaeva (2016b) summarizes their invariance principle through a couple of alternative specifications. The first states that business-time risk is proportional to the expected dollar bet size, also known as bet activity. The second implies that the relative bid–ask spread is proportional to the product of bet volatility and bet activity to the power of  $-1/3$ . We refer to the first principle as MI1 and the second as MI2.

The first variant stipulates,  $\frac{S_\tau \cdot Q_\tau^2 \cdot P_\tau^2}{N_\tau} = \mathcal{I}_\tau$ , where the invariant  $\mathcal{I}$  assumes the same role as in the TI relation (2). Following Andersen et al. (2018), we translate this MI1 specification into a high-frequency environment in logarithms and allow for time-series dependence of general form to obtain the representation (3), but with specific MI1-restricted parameter values,

$$MI1_\tau = (s_\tau - n_\tau) + 2 p_\tau + 2 q_\tau + c, \quad (4)$$

with the sequence of MI1 deviations,  $MI1_\tau$ , being mean zero and strictly stationary.

Compared to TI, MI1 likewise includes the  $(s - n)$  term while, for the other two pairs, the dimensionless feature is violated through the restrictions  $(\gamma_b, \gamma_p) = (0, 2)$  and  $(\gamma_d, \gamma_q) = (0, 2)$ .<sup>7</sup>

The MI2 relation of Kyle and Obizhaeva (2016b) is explored in Hou et al. (2024), who present a regression analysis based on the following specification,

$$b_\tau - p_\tau - s_\tau/2 = c - (1/3) \cdot w_\tau + u_\tau, \quad (5)$$

where  $w_\tau$  denotes bet activity, defined as  $w_\tau = \log(W_\tau) = \log(P_\tau \cdot N_\tau \cdot Q_\tau \cdot \sigma_\tau)$ , and  $u_\tau$  is a regression residual.

In the literature, expressions of this type have been explored via OLS. We advise against this practice. Inserting  $w_\tau = (p_\tau + n_\tau + q_\tau + s_\tau/2)/3$  into equation (5), we see that  $s_\tau$  appears on both sides of the equation, generating an endogenous regressor problem with associated biases in the parameter estimates and inflated  $R^2$  values.

An alternative specification is obtained by ensuring that  $(s_\tau - n_\tau)$  appears as an isolated component of the relation. Simple computations reveal that equation (5) implies that the MI2

---

<sup>7</sup>The dimensionless property is accommodated in Kyle and Obizhaeva (2016b) by equating the expected value of the invariant  $\mathcal{I}$  in equation (4) to the expected cost of generating an informed bet. The latter is unobserved and assumed constant, generating the MI1 representation featured here.

deviations take the form,

$$MI2_\tau = (s_\tau - n_\tau) - 2(b_\tau - p_\tau) - (b_\tau + q_\tau) + c. \quad (6)$$

### 2.3.2. Leverage Neutrality

MI is unique in prescribing the value of all power coefficients of the invariance system a priori. Interestingly, there is a purely arbitrage-based rationale for the stipulated coefficient of  $S$  relative to  $P$ . [Kyle and Obizhaeva \(2017\)](#) note that, if MI applies for a given asset, it must also be valid for a leveraged variant of the same asset. Specifically, let asset 1 have price  $P$  and variance  $S$ , while asset 2 denotes a portfolio consisting of asset 1 plus  $(L - 1) \cdot P$  in cash, where  $L > 0$ . Absent arbitrage, asset 2 then has price  $L \cdot P$  and variance  $S/L^2$ . For MI to apply for both assets, any representation including  $S \cdot P^\gamma$  must satisfy  $S \cdot P^\gamma = (S/L^2) \cdot (P \cdot L)^\gamma$ . This implies  $\gamma = 2$ , consistent with the manner in which the terms enter the MI relation.

More generally, this reasoning suggests that  $S \cdot P^2$  should feature in any genuine market invariant quantity, not just MI. For our TI representation (2), it implies  $\gamma_B = -2$ .

## 3. Data and Institutional Details

Our data set comprises all regular outright transactions and quotes in seven FX futures contracts spanning January 2013 through December 2021, namely the British pound (GBP), Japanese yen (JPY), Australian dollar (AUD), Canadian dollar (CAD), Euro (EUR), Swiss franc (CHF), and New Zealand dollar (NZD) contracts.<sup>8</sup> However, to ensure a reasonable degree of homogeneity across contracts, our main focus is on the first five currencies, which constitute the most liquid FX futures traded on Globex. Table 1 reports daily averages for a couple of activity metrics for all the contracts. The euro is, by some distance, the most actively traded currency, followed by the yen and pound and, after another notable drop, the CAD and AUD. Finally, by trading volume, CHF and NZD are clearly illiquid relative to the other contracts.

Our transaction data set is constructed from the Transaction Capture Report of the U.S. Commodity Futures Trading Commission (CFTC). This database has detailed information on

---

<sup>8</sup>Table A3 in Appendix B provides details on the specification of the individual FX futures contracts.

Table 1: **Liquidity Metrics of FX Futures Contracts**

| Currency                 | Open Interest | Volume  |
|--------------------------|---------------|---------|
| Euro (EUR)               | 467,095       | 224,951 |
| Japanese yen (JPY)       | 198,226       | 141,425 |
| British pound (GBP)      | 202,718       | 111,767 |
| Canadian dollar (CAD)    | 141,483       | 99,366  |
| Australian dollar (AUD)  | 144,407       | 74,150  |
| Swiss franc (CHF)        | 54,174        | 28,458  |
| New Zealand dollar (NZD) | 41,507        | 23,839  |

The table reports the averages of daily open interest and number of contracts traded over the period January 2013 and December 2021.

every transaction executed through the limit order books on Globex – the world’s largest regulated FX marketplace.<sup>9</sup> Specific information includes transaction time, price, and quantity of every trade, a dummy indicating whether the trade is a leg of a spread, such as, e.g., a calendar spread and a dummy indicating whether a trade is active, i.e., it executes against an order resting on the order book. We focus on outright trades (e.g., excluding calendar spread trades) in the front month futures contracts.<sup>10</sup> Every transaction is also paired with an Order ID, an identifier, linking it to the originating order. This is critical because it is unlikely for even a moderately sized marketable order to be filled against a single resting limit order of equal (or larger) size. The small size of limit orders on the book likely reflects the liquidity providers’ concerns about sniping, see, e.g., [Budish et al. \(2015\)](#). Thus, we focus on marketable orders submitted by the active party, defined as single orders executed (almost) instantaneously in one or more transactions against resting limit orders (the passive side of the trade).

We also use the market-by-price (MBP) data set provided to the CFTC by the CME Group. This database provides a new snapshot of the order book, whenever it changes, indicating how many contracts are available at the ten best bid and ask levels. This allows us to compute the bid-ask spread and number of contracts available at top-of-the-book (market depth). Moreover, we use the mid-quote to proxy for the price level and to estimate intraday volatility.<sup>11</sup> Given the large tick size, active trades execute almost exclusively against orders resting at the top of the book, so short-term liquidity is effectively captured by the spread and depth at this layer.

<sup>9</sup>See <https://www.cmegroup.com/trading/why-futures/welcome-to-cme-fx-futures.html>.

<sup>10</sup>We adhere to the CME Group’s rule to determine the roll-over date for the front month contract.

<sup>11</sup>The estimation procedure for volatility is detailed in Section 4.



The trading week starts Sunday at 17:00 Chicago Time (CT) and runs till 16:00 CT on Friday. We exclude weekends, holidays, and abbreviated trading sessions from our analysis along with a small set of trading days with missing quote data (when data gaps exceed 3 hours). Ultimately, the number of trading days varies from 2,198 (JPY) to 2,224 (CAD). Liquidity is low during 15:00-19:00 CT, which we label the overnight (ON) period, and we exclude it from our main analysis. This leaves 240 five-minute intervals from 19:00 to 15:00 the following day. In all, our sample comprises about 530,000 five-minute observations for each currency.

We previously commented on the substantial variation in activity across contracts and time. In Table A1 of Appendix B, we report a range of summary statistics for the five key currencies in our analysis.<sup>12</sup> Volatility is comparable, except for AUD which displays notably higher volatility than the other currencies. The euro transaction (dollar) volume is about double that of the Japanese yen, nearly fourfold the levels for British pound and Australian dollar, and almost six-fold the Canadian dollar volume.<sup>13</sup> Other measures of trading volume, like contracts exchanged or transaction count yield essentially the same ordering. The characteristics of the remaining variables are fairly similar across currencies. On average, active trades involve between 2.15 and 3.65 contracts, the bid-ask spread is between 1.01 and 1.29 times the minimum tick, the percentage spread ranges from 0.65 to 1.83 basis points, and each trade consumes between 9 and 30% of the depth at the top of the book. In summary, the bid-ask spread is close to the tick size, individual trades are small, and average depth is limited. Hence, a few trades in the same direction may well trigger a change in the prevailing top level quotes.

Overall, our FX futures contracts are well suited for an exploration of TI. First, they trade exclusively on the CME platform, minimizing synchronization errors. Second, they are subject to identical regulatory and technological requirements. Contract-specific details like tick sizes do differ and change at different times, but this exogenous variation serves as a pseudo-natural experiment to explore the robustness of the TI relation. Third, they trade almost around the clock, providing a stable pool of liquidity. These common traits allow us to pool our analysis and meaningfully test our predictions jointly for all major FX futures markets.

---

<sup>12</sup>The corresponding statistics for CHF and NZD are available in Table A2 of Appendix B.

<sup>13</sup>For benchmarking, the euro futures liquidity is roughly comparable to that of gold futures, while those of AUD and CAD are comparable to silver futures.

## 4. High-Frequency Measurement and Inference

To develop feasible estimation and inference procedures for the invariance relations introduced in Section 2.2, we must generate empirical measures for the activity variables over discrete time intervals using potentially noisy high-frequency data. The primary variable subject to measurement error is  $S$ , which is latent and subject to microstructure distortions. Consequently, it can only be assessed with reasonable precision over a non-trivial interval. As is common, we diversify the associated errors through aggregation of  $S$  to the daily level. In contrast, variables like  $N$  and  $Q$  are effectively observed, except for possible misclassification of the number of active trades. The latter should be a minimal source of error given the availability of order ID information. Finally,  $B$ ,  $P$ , and  $D$  are also measured with high accuracy, as we observe high-frequency snapshots of their level throughout the trading day.

The primary obstacle to accurate measurement of the input for the empirical analysis is the nonlinearity of the logarithmic transformation in equation (3), as we observe or estimate the level of the variables over a short, albeit nontrivial, interval. Because the series are subject to pronounced intraday variation, we record the realization of each variable locally. In doing so, we confront key tradeoffs. On the one hand, the shorter the time interval  $\tau$ , the more reasonable is the assumption of a stable value for the underlying variables. On the other hand, the interval must be sufficiently long to allow for multiple transactions to ensure meaningful measurement of  $n$  and  $q$  and, most importantly,  $s$ . We check whether our inference is robust to variation in the measurement of  $s$  in Section 5.5.3.

### 4.1. Basic Notation and Setting

Our sample covers a set of trading days,  $d = 1, \dots, D$ , with each day containing a fixed number of intervals,  $t = 1, \dots, T$ , of equal length  $1/T$ . Our measure for the generic activity variable  $A$  in interval  $t$  on day  $d$  is a strictly positive random variable  $\hat{A}_{d,t}$ , or simply  $\hat{A}_\tau$ , where  $\tau = 1, \dots, T \cdot D$ , so that interval  $(d, t)$  corresponds to  $\tau = T \cdot (d - 1) + t$ . We will use the  $A_{d,t}$  and  $A_\tau$  notation interchangeably throughout. If the realization of  $A_\tau$  is observed without error, then  $\hat{A}_\tau = A_\tau$ , but if  $A_\tau$  is latent or observed with error, then  $\hat{A}_\tau$  constitutes an estimate of  $A_\tau$ . We do not

require  $A_\tau$  to be stationary, as noted in Section 2.2, but only the  $TI_\tau$  series itself, representing the actual deviations from the TI relation, to be strictly stationary with finite fourth moments.

To obtain local measures for the activity variables, we further divide each short interval  $\tau$  into  $K$  subintervals of length  $1/K$ . Within each subinterval, for the “level or stock” variables,  $P, B$ , and  $D$ , we simply record the last observed value,  $\hat{A}_{\tau;k}$ , in line with the so-called previous tick method. We then estimate  $A_\tau$  by averaging, i.e.,  $\hat{A}_\tau = (\hat{A}_{\tau;1} + \dots + \hat{A}_{\tau;K})/K$ . Analogously,  $S_\tau$  is estimated by cumulating the squared log returns implied by the subinterval  $P_{\tau;k}$  observations. Finally, for the “flow” variables  $N$  and  $V$ , we simply cumulate the active trade count and contracts traded within  $\tau$  to obtain  $A_\tau$ . Finally,  $Q_\tau$  is  $V_\tau$  divided by  $N_\tau$ .

The above estimators may feature errors, ranging from non-existent (for, e.g.,  $N$  and  $V$  if active trades are recorded accurately) to large (for, e.g., the RV estimator). Our approach is designed to accommodate the presence of such errors. Specifically, to rationalize our measurement procedure, note that for any strictly positive estimator, we may write, without loss of generality,

$$\hat{A}_\tau = A_\tau \cdot U_\tau^{(A)}, \quad U_\tau^{(A)} > 0, \quad (7)$$

where a realization,  $U_\tau^{(A)} \neq 1$ , represents a measurement error.

If the multiplicative error  $U_{d,t}$  is stationary with limited intraday dependence and, possibly, pronounced intraday heteroskedasticity, the setting is reminiscent of the multiplicative error model (MEM), coined by Engle (2002), for which the innovations are scaled multiplicatively by the level of the activity variable.<sup>14</sup> Importantly, the proportional scaling feature of the noise is known to provide a good fit for key volatility and trading variables. Contrary to MEM, we exploit nonparametric measures for  $A_\tau$  based on high-frequency data. Thus, we inherit the flexible MEM representation, while avoiding concerns regarding parametric misspecification.

Finally, as in Section 2.2, we represent the system through the dynamics of the logarithmic activity variables, shifting the focus from proportional changes to changes in compound percentage terms. In obvious notation, with  $\mathbb{E}[\log(U_{d,t}^{(A)})] = \mathbb{E}[u_{d,t}^{(a)}] = c_a$ ,

$$\hat{a}_{d,t} = a_{d,t} + u_{d,t}^{(a)} = c_a + a_{d,t} + \tilde{u}_{d,t}^{(a)}, \quad (8)$$

---

<sup>14</sup>The MEM class includes successful specifications used for high-frequency return dynamics through the interaction of daily stochastic volatility factors and intraday diurnal features (Andersen and Bollerslev (1997, 1998)), for dynamic modeling of trade durations (Engle (2002)), as well as for high-frequency trading volume, e.g., (Brownlees et al. (2011, 2012)).

where the additive logarithmic residual  $\tilde{u}_{d,t}^{(a)} = u_{d,t}^{(a)} - c_a$  is strictly stationary with mean zero and variance  $\sigma_a^2$ . Notice that if  $A_{d,t}$  is observed with minimal error, then  $\hat{a}_{d,t} = a_{d,t}$  and, furthermore, weak dependence of the proportional noise process in equation (7) translates into serially correlated measurement errors in the corresponding additive log representation (8), which we may mitigate through intertemporal aggregation.

## 4.2. Aggregation to Daily Activity Measures

The intraday activity measures in equation (8) serve as the basic input to our empirical work. We diversify the observation errors via aggregation to a daily level. [Jacod et al. \(2017\)](#) and [Li and Linton \(2022\)](#) strongly reject the hypothesis of i.i.d. noise, but confirm that the intraday noise serial dependence is limited, implying that the noise mitigation will be effective.

To summarize, our individual daily activity measures, for  $d = 1, \dots, D$ , take the form,

$$\hat{a}_d = \hat{a}_d(T; K) = \sum_{t=1}^T \hat{a}_{d,t} = \sum_{t=1}^T a_{d,t} + \sum_{t=1}^T u_{d,t}^{(a)} = a_d + u_d^{(a)}, \quad (9)$$

where  $u_d^{(a)}$  is uncorrelated with  $u_{d'}^{(a)}$  for  $d \neq d'$ . Moreover, notice  $\mathbb{E}[\hat{a}_d] = c_d^{(a)} + a_d$ , where  $c_d^{(a)} = T \cdot c_a$ . Thus, apart from an invariant constant, that will be absorbed into the intercept of the TI relation, our estimate for the realization of the log activity variable is centered.

The inclusion of the index  $K$  in the second term of equation (9) serves as a reminder that the daily realized activity estimates depend on two underlying sampling frequencies. For notional brevity, we omit explicit references to these frequencies. Nonetheless, they may be important due to the nonlinearity of the mapping from the measured level of the activity variables to the log-linear TI relation. We provide additional discussion of these issues in Section 5.

## 4.3. TI, Frictions, and Intercepts

Exploiting the log-linear representation for each activity variable, plugging into equation (2), and letting superscript  $(j)$  denote asset  $j = 1, \dots, J$ , we obtain, for  $d = 1, \dots, D$ ,

$$TI_d^{(j)} = s_d^{(j)} - n_d^{(j)} + \gamma_B (b_d^{(j)} - p_d^{(j)}) + \gamma_D (d_d^{(j)} - q_d^{(j)}) + c_d^{(j)}. \quad (10)$$

This is the implied day-level TI specification where the intercept  $c_d^{(j)}$  absorbs all constants from the individual activity variables. The TI asserts that equation (10) holds for all  $d$  and  $j$  with

identical coefficients  $\gamma_{\mathcal{B}}$  and  $\gamma_{\mathcal{D}}$ , while the realized  $TI_d^{(j)}$  series is mean-zero and stationary, albeit likely serially correlated and heteroskedastic.

In this TI relation, we allow for different intercepts across assets. Requiring  $c_d^{(j)}$  to be identical for all  $j$  is very restrictive condition. Recall that the intercept of the log-linear representation for the individual activity variables absorbs local measurement errors and biases due to the concave logarithmic transformation. For assets with distinct characteristics, trading in segmented markets, and under different institutional settings subject to idiosyncratic frictions, there is no reason to expect the intercepts to be identical. At best, one may strive to identify patterns in the intercept related to the size or structure of underlying frictions across markets - a theme we pursue in Section 5.5.1. In contrast, for securities with similar characteristics and traded in the same market environment, one may expect some homogeneity also in the intercept term.

Corresponding homogeneity issues arise in the time dimension. We expect changes in the market organization, composition of market participants and trading technology to generate trend-like behavior in the scale for some of the system innovations. In this case, we may not expect the intercept to remain invariant across the full sample, but only over subsamples for which the relevant frictions and errors for the activity variables remain homogeneous.<sup>15</sup> Likewise, a sudden exogenous shift in the trading environment may generate a temporary disequilibrium in the trading patterns, generating serial correlation in the deviation from the TI relation. In this context, one may also tentatively interpret discrepancies in the intercept across assets as a measure of the relative degree of frictions in the respective markets.

This discussion highlights advantages of our focus on liquid FX contracts traded in a homogeneous market structure and regulatory environment during a period characterized by a relatively stationary evolution in the activity variables. If TI fails in this setting, it is unlikely to be a viable empirical hypothesis for more complex market structures and across more heterogeneous assets. On the contrary, if TI provides a good approximation to the market dynamics across our long sample, several assets, and multiple exogenous institutional changes, it solidifies the claim that TI captures important equilibrium interactions among key market activity variables.

---

<sup>15</sup>See [Kyle et al. \(2020\)](#) for a study that explicitly explores the impact of shifts in the trading environment on the empirical performance of the MII.

## 5. Empirical results

This section provides empirical evidence regarding TI. We initiate the analysis by exploring the volatility-trading activity relation. Observing systematic deviations from their association, we sequentially control for alternative activity variables that may help enhance the stability and robustness of the relation between volatility and trading, ending up with a final TI style representation that suitably accommodates the interaction amongst the activity variables in the FX futures contract market over our sample period.

### 5.1. Illustrative Evidence on the Volatility-Trading Activity Relation


It is well-known that return volatility and trading volume tend to be strongly correlated for a given asset over short horizons, inspiring MDH-style specifications. At the same time, it is recognized that MDH is a *ceteris paribus* relation. It is not expected to remain stable over longer periods and does not manifest itself in identical coefficients across different assets. To confirm that our FX futures display features consistent with the usual findings, we start our empirical analysis with an informal exploration of the volatility-trading volume association for the British pound futures (GBP), which is the only FX futures contract in our main sample that does not experience a structural break in the average trade size due to a tick size change. Hence, it offers us the opportunity to gauge the stability of the association within a stable institutional environment.

The left panel of Figure 4 displays the fitted OLS regression line for the SV representation  $s = c + \gamma_v (n+q) + \epsilon$ , i.e., daily realized log return volatility is regressed on the contemporaneous log trading volume,  $v = n+q$ . The fitted slope coefficient is near indistinguishable from unity,  $\hat{\gamma}_v \approx 1$ , lending support to a basic SV relation.<sup>16</sup> Of course, as noted in Section 2.2, this specification is incompatible with a dimensionless representation. Turning to the unitless quantity  $s - n$ , the above regression suggests the relation,  $s - n = c + \gamma_q q + \epsilon$  with  $\hat{\gamma}_q \approx 1$ . The OLS regression fit for this specification is presented in the right panel of Figure 4. If anything, we now observe an

---

<sup>16</sup>This regression is qualitatively consistent with prior results by, e.g., Epps and Epps (1976), Tauchen and Pitts (1983), Jones et al. (1994), Richardson and Smith (1994), Andersen (1996), and Liesenfeld (2001). In most of these papers, this basic MDH relation is generalized through more sophisticated theoretical and empirical work, as discussed previously for Tauchen and Pitts (1983), but see also, e.g., Darolles et al. (2015, 2017).

Figure 4: Illustrative OLS-Regressions for GBP



Illustrative OLS regressions for GBP of log realized volatility on log contract volume (left panel) and log realized volatility per transaction on average trade size (right panel). The solid lines are the fitted OLS regression lines and the dashed lines have the slopes constrained to unity ( $\gamma_v = 1$ ) and minus 2 ( $\gamma_v = -2$ ), respectively. Red dots indicate the peak of the COVID crisis, i.e., dates from March 15, 2020 to April 15, 2020.

even tighter linear association but the slope estimate is – economically and statistically – very far from unity, as  $\hat{\gamma}_q \approx -2$ . In other words, while the regression in the left panel is aligned with the traditional evidence, once we control separately for the trading intensity, we find significant confounding effects, grossly violating the SV ( $\gamma_q = 1$ ) and SN ( $\gamma_q = 0$ ) representations. This motivates our expanded framework encompassing alternative theories about the interplay of the volatility-trading activity with auxiliary activity variables.

In Appendix A, we depict the OLS fit corresponding to the left panel in Figure 4 for the other four FX futures contracts, documenting that the strong positive log-linear slope is robust across the currencies. In that sense, our data is well aligned with standard findings in the literature.

## 5.2. MI-Inspired Relations

### 5.2.1. MI1

This section provides formal estimates across our five currencies, indexed by superscript ( $j$ ), for a system embedding variants of the SN and MI1 inspired relations. To preserve the dimensionless feature of the volatility-related variable, we impose the common restriction that  $\gamma_n = 1$ . Since

the coefficient estimate for  $q$  in Figure 4 is consistent with the MI1 hypothesis (4), we test this representation for our set of FX futures contracts. The system takes the form,

$$MI1_d^{(j)} = c^{(j)} + (s_d^{(j)} - n_d^{(j)}) + \gamma_q^{(j)} q_d^{(j)} + \gamma_p^{(j)} p_d^{(j)}, \quad (11)$$

where the error  $MI1_d^{(j)}$  is strictly stationary and has finite fourth moments.

We note that the basic specifications explored in Section 5.1 arise as special cases through imposition of specific coefficient restrictions, e.g., MI1 implies  $\gamma_q^{(j)} = 2$ , and  $\gamma_p^{(j)} = 2$ . As such, the evidence speaks to the applicability of various representations explored in prior work.

Because we view the interaction amongst the activity variables as a trading equilibrium phenomenon, we account for endogeneity and report standard GMM estimation results. Beyond equation (11), the requisite moment conditions reflect orthogonality of the innovations to the  $MI1_d^{(j)}$  series with respect to lagged instruments, here consisting of the average value of  $(s - n)$ ,  $q$  and  $p$  computed over the preceding ten trading days.

Table 2 provides coefficient estimates and associated standard errors for each of the currencies, while the last row reports results for a pooled regression which restricts all coefficients across the currencies to be identical. If the invariance relation is compatible across contracts, this improves the efficiency of the inference and thus succinctly summarizes the overall evidence.<sup>17</sup> Finally, as a simple gauge for the quality of fit, the  $\bar{R}^2$  statistic in the last column of the table indicates the fraction of sample variation in  $(s - n)$  explained by the GMM point estimates and sample realizations for  $q$  and  $p$ .

The estimation results in Table 2 are wildly inconsistent across currencies. The explanatory power, in terms of the estimated coefficients accounting for the sample variation in  $(s - n)$ , is very weak, albeit better for GBP, and partially AUD, than for the other FX contracts. One may speculate that the particularly poor results for EUR, CAD and JPY stem from a regime shift in  $q$  induced by the tick size change. GBP and AUD, respectively, experience no tick change or one very late in the sample period. The extreme variation in the coefficient values imply that all invariance hypothesis encompassed within the given specification are decidedly rejected.

---

<sup>17</sup>When estimating different invariance specifications, we first de-mean all activity variables by subtracting their pooled average (across all day and currencies). This simplifies interpretation of the log-intercept  $c$ , which should equal zero, if the (additive log) scale-adjusted invariance relation is identical across assets.



Table 2: Estimation Results based on the MI1 Equation (11)

|        | Nobs  | $c$    | $\gamma_q$ | $\gamma_p$ | $\text{se}(c)$ | $\text{se}(\gamma_q)$ | $\text{se}(\gamma_p)$ | $\bar{R}^2$ |
|--------|-------|--------|------------|------------|----------------|-----------------------|-----------------------|-------------|
| AUD    | 2127  | -0.663 | 1.786      | 1.113      | 0.106          | 0.242                 | 0.339                 | 0.261       |
| CAD    | 2030  | -0.760 | 1.142      | -3.022     | 0.211          | 0.835                 | 1.258                 | 0.001       |
| EUR    | 2146  | 0.799  | -0.366     | 0.330      | 0.312          | 0.733                 | 1.434                 | -0.006      |
| GBP    | 2150  | -0.368 | 2.240      | 0.739      | 0.108          | 0.277                 | 0.310                 | 0.451       |
| JPY    | 2147  | 0.223  | -0.152     | -2.124     | 0.115          | 0.445                 | 1.326                 | 0.042       |
| Pooled | 10600 | -0.076 | -0.597     | 0.963      | 0.047          | 0.233                 | 0.182                 | 0.122       |

GMM results for MI1 Equation (11) using a constant plus ten-day averages of lagged log activity variables as instruments. The columns with “se” indicate Newey-West standard errors for the coefficients using 63 lags. The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the point estimates and sample realizations for  $q$  and  $p$ .

### 5.2.2. MI2

The above findings suggest that, at least for these “large tick” markets, the reduction in tick size lowers the bid-ask spread and the marginal trading costs, effectively inducing a regime shift, incompatible with standard invariance representations. We now turn to the MI2 specification, which explicitly incorporates the bid-ask spread within the MI setting. Hou et al. (2024) provides supportive evidence for this relation through OLS regressions based on equation (5).<sup>18</sup> However, as noted in Section 2.3.1, that regression is subject to a serious endogeneity bias. Instead, we estimate the equivalent MI2 relation (6) by GMM,

$$MI2_d^{(j)} = c^{(j)} + (s_d^{(j)} - n_d^{(j)}) + \gamma_B(b_d^{(j)} - p_d^{(j)}) + \gamma_{qb}(q_d^{(j)} + b_d^{(j)}), \quad (12)$$

where the error  $MI2_d^{(j)}$  is strictly stationary and has finite fourth moments. We note that, from Section 2.3.1, MI2 implies  $\gamma_B = -2$ ,  $\gamma_{qb} = -1$ .

Table 3 shows that the unconstrained MI2 system improves the fit substantially, capturing about 83% of the sample variation in  $(s - n)$  across the currencies versus 12% in Table 2. Moreover, the currencies with lowest  $\bar{R}^2$  values, GBP and AUD, are those with a limited degree of sample variation in the spread, consistent with  $(b - p)$  providing important explanatory power. Furthermore, the  $\gamma_B$  estimates are close to the theoretical value of  $-2$ , albeit somewhat downward biased. The most troubling aspect is that the estimates for  $\gamma_{qb}$  are significantly positive for four

<sup>18</sup>We have confirmed, in unreported results, that OLS regression results for equation (5) are qualitatively consistent with the findings of Hou et al. (2024). Thus, our very different conclusions reported below based on equation (6) are driven by the correction of the strong endogeneity bias in equation (5).

Table 3: Estimation Results based on the MI2 Equation (12)

|        | Nobs  | $c$    | $\gamma_B$ | $\gamma_{qb}$ | $se(c)$ | $se(\gamma_B)$ | $se(\gamma_{qb})$ | $\bar{R}^2$ |
|--------|-------|--------|------------|---------------|---------|----------------|-------------------|-------------|
| AUD    | 2127  | -0.144 | -2.051     | 1.112         | 0.073   | 0.172          | 0.180             | 0.593       |
| CAD    | 2030  | -0.028 | -2.462     | 0.537         | 0.045   | 0.186          | 0.101             | 0.858       |
| EUR    | 2146  | -0.095 | -2.984     | 1.084         | 0.062   | 0.212          | 0.112             | 0.803       |
| GBP    | 2150  | -0.109 | -2.455     | -0.059        | 0.089   | 0.147          | 0.245             | 0.653       |
| JPY    | 2147  | 0.179  | -3.571     | 1.410         | 0.044   | 0.371          | 0.247             | 0.728       |
| Pooled | 10600 | -0.010 | -2.202     | 0.357         | 0.019   | 0.075          | 0.069             | 0.828       |

GMM results for MI2 equation (12) using a constant plus the ten day average of lagged log activity variables as instruments. The columns with “se” indicate Newey-West standard errors for the coefficients using 63 lags. The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the point estimates and sample realizations for  $(b - p)$  and  $(q + b)$ .

of the currencies as well as the pooled system and all very far from  $-1$ , as prescribed by MI2.

It is evident that the unconstrained MI2 representation, involving the bid-ask spread, provides a performance boost. The second component in the MI2 relation is proportional to the transaction cost of the average trade size. Since  $b$  tends to increase with higher  $(s - n)$ , the negative sign of the coefficient again suggests that  $q$  declines in response to elevated business-time risk. This is also in line with our findings in Table 2 for GBP and AUD, which are largely insulated from tick size variation across the sample. Thus, while the theoretical MI2 specification clearly is rejected, we conjecture that both the spread and  $q$  belong in a suitable invariance representation. Hence, we now turn to the TI specification, including both  $q$  and  $b$ .

### 5.3. TI: Interacting Risk, Trading Costs and Market Depth

TI consists of three denomination-less components, reflecting business-time return variation, or risk, and a pair of variables capturing separate dimensions of liquidity, the bid-ask spread and order book capacity. Our empirical specification of equation (2) takes the form,

$$TI_d^{(j)} = c^{(j)} + (s_d^{(j)} - n_d^{(j)}) + \gamma_B^{(j)} (b_d^{(j)} - p_d^{(j)}) + \gamma_D^{(j)} (d_d^{(j)} - q_d^{(j)}), \quad (13)$$

where the error  $TI_d^{(j)}$  is strictly stationary with finite fourth moments.

As for the MI systems, we provide GMM estimates based on orthogonality of the TI innovations with respect to an average of lagged activity variables. The implementation is detailed in Appendix C. Estimation results for the key parameters are provided in Table 4.

Our TI estimates indicate a dramatic improvement over the prior invariance representation.

Table 4: Estimation Results based on the TI Representation (13)

|        | Nobs  | $c$    | $\gamma_{\mathcal{B}}$ | $\gamma_{\mathcal{D}}$ | $\text{se}(c)$ | $\text{se}(\gamma_{\mathcal{B}})$ | $\text{se}(\gamma_{\mathcal{D}})$ | $\bar{R}^2$ |
|--------|-------|--------|------------------------|------------------------|----------------|-----------------------------------|-----------------------------------|-------------|
| AUD    | 2127  | -0.216 | -1.683                 | 0.645                  | 0.104          | 0.311                             | 0.171                             | 0.737       |
| CAD    | 2030  | -0.094 | -2.030                 | 0.481                  | 0.026          | 0.099                             | 0.106                             | 0.890       |
| EUR    | 2146  | 0.112  | -1.929                 | 0.581                  | 0.118          | 0.309                             | 0.341                             | 0.869       |
| GBP    | 2150  | -0.091 | -2.093                 | 0.190                  | 0.090          | 0.414                             | 0.334                             | 0.727       |
| JPY    | 2147  | 0.053  | -1.904                 | 0.553                  | 0.031          | 0.121                             | 0.099                             | 0.860       |
| Pooled | 10600 | -0.005 | -2.098                 | 0.582                  | 0.048          | 0.140                             | 0.184                             | 0.924       |

GMM results for the TI system (13) using orthogonality of the TI innovation to variables in the lagged information set, as detailed in Appendix C. The columns with “se” indicate White (heteroskedasticity-robust) standard errors with autocorrelation captured through an ARMA(2,1) representation. The  $\bar{R}^2$  column refers to the fraction of the sample variation in  $(s - n)$  explained by the point estimates and sample realizations for  $(b - p)$  and  $(d - q)$ .

We now find the intercepts for all FX contracts to be close to the mean value of zero, while the  $\gamma_{\mathcal{B}}$  estimates are statistically compatible with the prescribed TI value of  $-2$ , and the  $\gamma_{\mathcal{D}}$  point estimates are consistently positive and compatible with the pooled estimate of 0.58. Moreover, the explanatory power for variation in  $(s - n)$  has uniformly improved from the MI2 system, while the two contracts with limited variation in the bid-ask spread continue to have the lowest degree of explanatory power. In terms of goodness-of-fit, the  $J$ -statistic with one degree-of-freedom for our pooled TI system is low, indicating no statistical rejection of the underlying moment conditions.<sup>19</sup> A Wald test for  $\gamma_{\mathcal{B}} = -2$  equals 0.49, again much below the level for rejection. In addition, the hypothesis that both  $\gamma_{\mathcal{B}}$  and  $\gamma_{\mathcal{D}}$  are identical across assets, as stipulated by the genuine invariance principle, also cannot be rejected.<sup>20</sup>

In summary, Table 4 documents striking performance enhancement, consistent with TI delivering an acceptable invariant representation of a trading equilibrium jointly across five FX contracts over a long sample period. In comparison, the invariance specifications explored above generate starkly different coefficient values for our FX contracts and time periods, and they are clearly subject to regime-like shifts in response to an exogenous shock to the bid-ask spread.

Another issue of interest is whether the average deviation from the TI hypothesis, captured by the scale-adjusted intercept  $c$ , displays a systematic pattern across the currencies. For example, is the TI equilibrium characterization similar for more and less liquid FX futures contracts? To

<sup>19</sup>As detailed in Appendix C, all GMM estimation results are obtained with one moment condition in excess of the number of estimated parameters, generating standard  $\chi^2(1)$ -distributed test values.

<sup>20</sup>For example, conjecturing that  $(\gamma_{\mathcal{B}}, \gamma_{\mathcal{D}}) = (-2, 0.5)$  across all currencies leads to a Wald statistic of 0.61, again suggesting a satisfactory statistical fit.

address this question, we estimate the pooled TI system with the slope coefficients restricted to be identical, but allowing for distinct fixed currency effects through separate  $c^{(j)}$  values. This corresponds to estimating the pooled system reported in the last row of Table 4 with identical slope coefficients, but allowing for (five) different currency intercepts. The results for the individual currency intercept and common slope coefficients are reported in Table 5.

Table 5: Estimation Results for the TI Representation (13) with Fixed Effects

|        | AUD    | CAD    | GBP   | EUR   | JPY   | $\gamma_{\mathcal{B}}$ | $\gamma_{\mathcal{D}}$ | $\bar{R}^2$ |
|--------|--------|--------|-------|-------|-------|------------------------|------------------------|-------------|
| Coeff. | -0.089 | -0.100 | 0.118 | 0.001 | 0.053 | -1.912                 | 0.539                  | 0.929       |
| se     | 0.049  | 0.033  | 0.055 | 0.054 | 0.035 | 0.127                  | 0.172                  |             |

GMM estimation results for the TI system (13) with fixed currency effects, but  $\gamma_{\mathcal{B}}$  and  $\gamma_{\mathcal{D}}$  constrained to be equal. The row with “se” indicates White (heteroskedasticity-robust) standard errors with autocorrelation accommodated through an ARMA(2,1) error specification. The  $\bar{R}^2$  column refers to the fraction of the sample variation in  $(s - n)$  explained by the GMM point estimates and the sample realizations for  $(b - p)$  and  $(d - q)$ .


Table 5 reveals several points worthy of reflection. First, the intercepts are quite similar, but with the distinct pattern that the less liquid currencies, AUD and CAD, have negative intercepts, while the more heavily traded currencies have positive intercepts. Second, the explanatory power of the system is only mildly improved by leaving the intercept unconstrained relative to the pooled estimation result in Table 4. Third, the slope coefficients remain statistically indistinguishable from  $-2$  and  $1/2$  (or any other candidate value close to  $1/2$ ).

Overall, our results corroborate the basic predictions of the TI: our trading equilibrium representation based on high-frequency measures of activity variables applies near universally across five separate currency contracts, which display substantive differences in terms of trading volume, bid-ask spreads, and order book depth. Moreover, the theoretical constraint  $\gamma_{\mathcal{B}} = -2$  cannot be rejected, and  $\gamma_{\mathcal{D}}$  is similar across assets with a value around or slightly above  $1/2$ . Of course, our results are obtained for a select group of assets trading exclusively within one exchange, so the market organization and regulatory structure is aligned across the contracts. We expect more substantive deviations from the TI, as we consider more disparate markets and assets. The tendency for the intercept values to roughly line up with the trading intensity is a hint that such factors may be at play, even for our fairly homogenous FX contracts. In fact, a Wald test for whether all intercepts are equal to some common value  $c_0$  along with  $\gamma_{\mathcal{B}} = -2$  and a common freely estimated value for  $\gamma_{\mathcal{D}}$  across the currencies is strongly rejected.

## 5.4. Performance Measures for the Invariance Relations

In this section, we provide a few illustrations designed to convey the relative performance or in-sample accuracy of the alternative invariance relations analyzed above.

Figure 5: EUR Futures — 6 Invariants Over Time



The panels display 10-day moving averages for the alternative invariance relations: SV, SN, MI1, MI2, TI, and TI\*. The first four relations are evaluated for the theoretically prescribed parameters. TI uses the slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ , whereas TI\* uses the estimated slopes  $(\hat{\gamma}_B, \hat{\gamma}_D)$ .

Figure 5 displays estimated TI realizations for the euro. The six panels reflect alternative specifications, two volatility-trading relations SV and SN, two associated with MI, and two implied by different TI coefficients. For the first four panels, we impose the theoretically prescribed parameters. The TI panel imposes the theoretically predicted value of  $\gamma_B = -2$  and sets  $\gamma_D = 1/2$ , motivated by the empirical findings. The TI\* panel relies on the estimated coefficients. The two panels for TI hypothesis are nearly indistinguishable.<sup>21</sup>

While the relatively large and serially correlated errors for the SV and SN representations may be expected, the more striking aspect is the equally large deviations for the MI specifications. For MI1, we already noted the poor fit in Table 2 for the euro, but the MI2 fares no better, because we are imposing the theoretical coefficient in the relation. In other words, the relatively large

<sup>21</sup>The euro results are representative for the other currencies, see Figures A3-A6 in Appendix A.

degree of explained variation captured by MI2 in Table 3 stems from the inclusion of relevant activity variables, but once the theoretical coefficients are imposed in lieu of the estimated ones, it is clear that MI2 is starkly at odds with the evidence. Moreover, the halving of the tick size generates a large and persistent shift in the sign of the MI2 error - reminiscent of a structural break. It is apparent that MI2 fails to capture the interaction of the bid-ask spread with the other activity variables. In contrast, TI evolves quite smoothly around zero, although the serial correlation in the error structure is notable. Visibly, it stands to reason that TI cannot formally be rejected once we account for the strong interdependencies in the system.


We further note that the tick size change is accommodated smoothly by TI, in contrast to the MI specifications. This pattern is observed for all currencies subject to a tick change, although TI does incur a small dip at the tick change for AUD; see Figures A3-A6. Likewise, the onset of COVID does not lead to TI violations. This speaks to the importance of including liquidity in the form of market depth and bid-ask spread in the invariance relation. Finally, for TI, we observe an upward trend towards the end of the sample, corresponding to the initial phase of the inflationary surge associated with supply chain disruptions following the COVID pandemic. The higher TI values likely reflect more friction-plagued or illiquid market conditions, as discussed in more detail in Section 5.6.

To convey the in-sample fit for our main currencies, Figure 6 depicts the root-mean-squared-error (RMSE) for the different invariance specifications. Since there was no tick size change for GBP and only late in the sample for AUD, the lower MI2-RMSE values for those currencies corroborate our finding that MI2 has difficulty capturing the interaction of the bid-ask spread with the remaining activity variables. Finally, we see that TI generates consistently low RMSE measures and uniformly dominates the alternative representations in this metric.

## 5.5. Robustness of the TI

We now assess the robustness of TI to changes in our measurement and estimation procedures. First, we explore the impact of the trading intensity. Second, we check whether TI applies with daily sampling. Third, we evaluate the sensitivity to the frequency used for measuring RV.

Figure 6: RMSE for Alternative Invariance Relations



The figure displays in-sample Root-Mean-Squared-Errors for alternative invariance representation. Theoretical coefficients are used to compute the errors for the first four columns. For TI, the “theoretical” slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$  are employed, with the latter coefficient determined by the general empirical evidence.

### 5.5.1. TI and Trading Intensity

Our focus on a limited set of FX futures contracts is motivated by the relatively frictionless trading environment on the GLOBEX platform. The critical question is whether TI provides a good benchmark for a trading equilibrium in liquid market settings. Moreover, if frictions do matter, what asset or market characteristics invalidate the relation and along what dimensions?

Table 6: TI System (13) with Liquid Contracts and Fixed Effects

|        | EUR   | GBP    | JPY   | $\gamma_B$ | $\gamma_D$ | $\bar{R}^2$ |
|--------|-------|--------|-------|------------|------------|-------------|
| Coeff. | 0.070 | -0.072 | 0.001 | -1.885     | 0.484      | 0.900       |
| se     | 0.057 | 0.077  | 0.040 | 0.149      | 0.273      |             |

GMM results for the TI system (13) with fixed currency effects, but  $\gamma_B$  and  $\gamma_D$  constrained to be equal. Only the three most liquid currency contracts are considered. The row with “se” indicates White (heteroskedasticity-robust) standard errors. The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the GMM point estimates and the sample realizations for  $(b - p)$  and  $(d - q)$ .

For FX futures, one may conjecture that average trading intensity is (negatively) correlated with the degree of market frictions. We check whether our results differ as we vary the underlying homogeneity in trading intensity across the contracts. Specifically, we present results featuring

only the three most actively traded contracts, EUR, JPY, and GBP, and results where we instead expand the set of five currencies to include also CHF and NZD. For brevity, we focus on the pooled estimates with fixed currency effects.

Table 7: TI with Thinly Traded Contracts and Fixed Effects

|        | AUD   | CAD   | EUR   | GBP   | JPY   | CHF    | NZD    | $\gamma_B$ | $\gamma_D$ | $\bar{R}^2$ |
|--------|-------|-------|-------|-------|-------|--------|--------|------------|------------|-------------|
| Coeff. | 0.018 | 0.000 | 0.324 | 0.087 | 0.224 | -0.236 | -0.445 | -1.719     | 0.361      | 0.954       |
| se     | 0.053 | 0.028 | 0.093 | 0.043 | 0.063 | 0.045  | 0.095  | 0.153      | 0.140      |             |

GMM results for the TI system (13) with fixed currency effects, but  $\gamma_B$  and  $\gamma_D$  constrained to be equal. The currency set is expanded to include CHF and NZD. The row with “se” indicates White (heteroskedasticity-robust) standard errors. The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the GMM point estimates and sample realizations for  $(b - p)$  and  $(d - q)$ .


Tables 6 and 7 provide GMM results for the pooled TI system with fixed currency effect. They convey a systematic pattern, with smaller estimated intercepts for less actively traded contracts. Specifically, in Table 6 for a set of homogenous contracts, the intercepts are indistinguishable from zero and the slope coefficients close to the hypothesized values of -2 and 1/2. In contrast, Table 7 reports a notable drop in the (absolute) slope coefficient values relative to Tables 5 and 6, and the intercepts are now statistically distinct. In short, we encounter systematic heterogeneity in the estimates, as we move towards less liquid securities, suggesting the TI intercepts absorb market frictions or, alternatively, it speaks to a missing ingredient in our TI specification, that may accommodate auxiliary liquidity effects. Nonetheless, the slope estimates have not changed significantly in a statistical sense, so TI retains a high degree of explanatory power.

To provide a visual perspective, the left panel of Figure 7 depicts the association between the average TI residual (equal to  $-c$ ) and the liquidity proxy  $v$ . The monotone association suggests that increased market frictions tend to be absorbed into the intercept of the TI relation.<sup>22</sup> If this interpretation is correct, one should expect this negative relation to become even stronger, if we include less liquid assets in the system. In the right panel of Figure 7, we corroborate this hypothesis – the incorporation of the less active currencies, CHF and NZD, generates an even starker negative association with the newly introduced currencies representing outliers positioned in the upper left corner of the display. That is, the results suggest that the intercepts in the

<sup>22</sup>Alternative liquidity proxies like transaction count or bid-ask spread generates qualitatively similar results.



Figure 7: TI Invariant versus  $v$



For each currency, the figure plots the average TI residual against the average log trading volume, where the TI relation is evaluated for the “theoretical” slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ . The left panel includes our five main currencies, while the right panel adds less active NZD and CHF.

pooled-fixed effect TI system serve as proxies for the relative level of frictions in the markets, consistent also with the conclusions of [Benzaquen et al. \(2016\)](#).

### 5.5.2. TI at the Daily Level

Most prior studies of market invariance relations rely on daily observations, i.e.,  $T = 1$ , while we exploit measurements over short intraday intervals,  $1/T = 5$  minutes, with subsequent aggregation of the logarithmic values to the daily level. Section 4.1 outlines the basic tradeoff of shorter versus longer intervals prior to the log transformation. One further potential advantage of a measurement interval as long as one trading day (prior to taking logs) is that it effectively removes the diurnal pattern in the activity variables, which may bolster robustness.<sup>23</sup>

Table 8 displays TI estimates based on the log of the daily activity measures, i.e.,  $T = 1$  trading day. The individual currency and pooled results bear some resemblance to those for the same system with high-frequency aggregation ( $1/T = 5$  minutes) in Table 4. The main distinction is a drop in the  $\gamma_B$  estimate and the degree of explained variation. We now obtain a  $p$ -value around 2% for a Wald test of  $\gamma_B = -2$ , while the joint test for  $\gamma_B = -2$  and identical intercepts generates a  $p$ -value below 5%, suggesting a mild deterioration in TI fit compared

<sup>23</sup>Note that all activity variables are still observed at the underlying frequency of  $1/K = 30$  sec, but for  $T = 1$  they are aggregated to the daily level before taking logs. E.g., the daily observation of  $s$  equals  $\log(RV)$ , with  $RV$  estimated as the cumulative sum of 30-second squared returns.

Table 8: TI System (13) with Daily Aggregation and Fixed Effects

|        | AUD    | CAD    | GBP   | EUR    | JPY   | $\gamma_B$ | $\gamma_D$ | $\bar{R}^2$ |
|--------|--------|--------|-------|--------|-------|------------|------------|-------------|
| Coeff. | -0.159 | -0.083 | 0.154 | -0.004 | 0.051 | -1.733     | 0.483      | 0.884       |
| se     | 0.110  | 0.032  | 0.070 | 0.040  | 0.040 | 0.114      | 0.114      |             |

GMM results for the TI system (13) with fixed currency effects, but  $\gamma_B$  and  $\gamma_D$  constrained to be equal. The activity variables are aggregated to the daily level before taking logarithms. The row with “se” indicates White (heteroskedasticity-robust) standard errors. The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the GMM point estimates and the sample realizations for  $(b - p)$  and  $(d - q)$ .

to high-frequency measurement approach. Nonetheless, given the relatively mild rejections, we find that the TI, to first order, provides a reasonable characterization of the trading equilibrium at the daily frequency. Even so, we stress that these results are obtained with high-frequency monitoring of the bid-ask spread and market depth throughout the trading day along with the cumulative 30-second squared returns. As such, an intraday measurement procedure remains at the core of the approach also for this daily TI relationship.

### 5.5.3. Alternative RV Measurement

While we observe the realizations of most of the activity variables directly at high frequencies, mitigating concerns regarding distortions from measurement error or noise, this is not the case for  $s$ . The return volatility is a latent variable, estimated with sampling error and subject to microstructure noise at very high frequencies.

Table 9 reports estimation results for the TI system (13) with fixed currency effects using three separate high-frequency sampling frequencies for the squared returns within each five-minute window. The 30-second frequency equals the one employed in Table 5, while the alternatives of 10-second and 60-second sampling are sensible alternatives to explore the sensitivity to noise or bias in the RV measure. The most striking feature is the relative stability of the results across the three alternatives. Moreover, if anything, the results indicate improvements from pushing towards the higher sampling frequency of 10-seconds. The estimated coefficients almost uniformly move closer to the theoretical value associated with TI, the estimation precision improves, and the system achieves a higher degree of explained in-sample variation for  $(s - n)$ . Moreover, use of the 10-second RV measure does not eliminate the systematic pattern

in the intercept, implying that the monotonic relation displayed in Figure 7 should remain.

Table 9: TI System (13) with Fixed Effects and Alternative RV Measures

|               | AUD    | CAD    | GBP   | EUR    | JPY   | $\gamma_B$ | $\gamma_D$ | $\bar{R}^2$ |
|---------------|--------|--------|-------|--------|-------|------------|------------|-------------|
| <b>10 sec</b> |        |        |       |        |       |            |            |             |
| Coeff.        | -0.086 | -0.084 | 0.104 | 0.013  | 0.030 | -1.988     | 0.500      | 0.933       |
| se            | 0.046  | 0.027  | 0.037 | 0.039  | 0.027 | 0.077      | 0.105      |             |
| <b>30 sec</b> |        |        |       |        |       |            |            |             |
| Coeff.        | -0.089 | -0.100 | 0.118 | 0.001  | 0.053 | -1.912     | 0.539      | 0.929       |
| se            | 0.049  | 0.033  | 0.055 | 0.054  | 0.035 | 0.127      | 0.172      |             |
| <b>60 sec</b> |        |        |       |        |       |            |            |             |
| Coeff.        | -0.106 | -0.093 | 0.121 | -0.003 | 0.055 | -1.889     | 0.561      | 0.912       |
| se            | 0.053  | 0.029  | 0.043 | 0.042  | 0.028 | 0.092      | 0.111      |             |


GMM estimation results for the TI system (13) with fixed currency effects, but  $\gamma_B$  and  $\gamma_D$  constrained to be equal. The 5-minute RV measures are obtained from cumulative squared returns sampled at different underlying frequencies, 10, 30, and 60 seconds, respectively. The rows with “se” indicates White (heteroskedasticity-robust) standard errors (Newey-West lag = 0). The  $\bar{R}^2$  column refers to the fraction of sample variation in  $(s - n)$  explained by the GMM point estimates and the sample realizations for  $(b - p)$  and  $(d - q)$ .

## 5.6. Longer-Term Equilibrium Adjustments among TI Components

The pseudo-natural experiments afforded by the tick size change for four currency contracts at different times across the sample enables us to gauge the temporal robustness of the TI relation in more depth. If TI applies and the  $(b - p)$  values are relatively stable within each tick size regime, variation in  $(s - n)$  should largely be offset by a concurrent variation in  $(d - q)$ . By assessing the strength of this relation, before and after the tick change, we gain insight into the long-term adaptation to the abrupt shift in marginal trading costs.

Figure 8 depicts the average value for each dimensionless quantity across the trading day before and after the tick change where, for GBP, we define the before-versus-after split to coincide with that for EUR. Hence, the GBP series serves as a placebo-style control. First, the consistent patterns observed across the trading day corroborates the robustness of our empirical measurement technique. In each case, we observe a distinct parallel shift in the average value, but the qualitative intraday pattern, reflecting the different activity levels, regional market openings and closures and timing of macroeconomic announcements across the global trading cycle, is preserved. The larger deviations observed for the AUD reflects a very short sample period after the tick change, rendering estimates for the “after tick change” pattern imprecise.

Figure 8: Dimensionless Variables Before and After the Tick Change (de-meanned)



The figure displays the average value of  $(b-p)$ ,  $(s-n)$ , and  $(d-q)$  across all 5-minute intervals within the continuous segment of the trading day for our five main FX contracts. For those undergoing a tick size change, the averages are taken before (blue) and after (green) the tick change. The tick size did not change for GBP, and the before and after periods are instead aligned with the euro contract.

The long-term response to the drop in the tick size is strikingly similar across currencies. The transaction activity, as signified by  $n$ , increases relative to  $s$ , generating a substantial decline in the return variation measured in business time,  $(s - n)$ . Likewise, we observe a large decrease in our book capacity indicator,  $(d - q)$ . Thus, consistent with the TI predictions, the drop in the spread leads to more active trading, albeit at a smaller trade size, and a substantial drop in the market depth at the top level of the order book.

As a back-of-the-envelope assessment, we see an approximate 0.45 drop in the spread in the euro market after the tick change, indicating that the tick size is slightly less binding at this lower level. According to the TI, this should induce offsetting adjustments in the business time return variation and order book depth. We observe something akin to a 0.65 drop in  $(s - n)$  and a decrease in  $d - q$  of about 0.5. This fits well with the TI prediction:  $-2\Delta(b - p) + \Delta(s - n) + 0.5\Delta(d - q) \approx -2(-0.45) + (-0.65) + 0.5(-0.5) = 0$ . Similar rough relations may be deduced from the other currencies undergoing a tick change.<sup>24</sup>

The GBP offers an additional illustration. We note that there was an average increase in the spread of about 0.1 during the European and American trading segment from the period before to after the EUR tick size change - albeit this was surely not a causal effect. According to TI, this should be compensated through shifts in the other activity variables. Indeed, the volatility per business time rises on average by approximately 0.1, while the depth measure shows a more solid increase of 0.2. Again the proportions are in line with the TI relation. Finally, even the unchanged spread for GBP during the Asian regime is consistent with the proportionally offsetting drop in  $(s - n)$  by around 0.1 and an increase of 0.2 in  $d - q$ . That is, not only do we see the TI capturing the longer-term effects in response to exogenous market design shocks, but we also observe the relationship reflected for less dramatic, but meaningful secular shifts in the individual activity measures.

We conclude that, for this “large tick” market, the average bid-ask spread is largely stable, even if may be slightly lower in the more active North-American (intraday) segment, and it

---

<sup>24</sup>The CME Group pays close attention to the state of liquidity in their FX futures contracts. For example, following the lowering of the AUD tick size in 2020, they conducted a before-versus-after analysis of the spread and order book depth, using their FX Market Profile tool to pinpoint signs of improvements. Relative to TI, their analysis ignores the role of the transaction intensity  $n$  and average trade size  $q$  in the adjustment process to the tick change. See <https://www.cmegroup.com/education/articles-and-reports/reviewing-the-impact-of-aud-usd-futures-tick-size-change.html> for the review of the impact of the tick size change.

clearly serves a critical role in the market’s short-term adjustment to large shocks. Hence, during regular tranquil market conditions, the equilibrium involves substantial offsetting variation in  $(s - n)$  and  $(d - q)$ . Qualitatively, this is intuitive. If liquidity providers react to the evolving market dynamics at the prevailing transaction speed, then rising business-time volatility implies an increased sniping probability, leading to active pairing of risk exposure through withdrawal of liquidity at the top of the book. The quantification of this relationship embodied in the coefficient estimate of around  $1/2$  is intriguing but, as of now, not theoretically founded.

## 6. Conclusion

We provide consistent inference and testing procedures for the multi-asset TI system within a stable institutional setting that avoid the endogeneity bias associated with prior empirical work on related hypotheses. Moreover, we study the impact of external market shocks through quasi-natural experiments, where marginal trading costs drop due to a tick size change. Our stipulation, that market invariants should be dimensionless, guides our extension of prior high-frequency invariance specifications to complement business-time asset risk (volatility) with components reflecting the book capacity and (marginal) trading costs. Our extended TI formulation performs well empirically for the FX futures contracts, while we convincingly reject a set of prior invariance specifications explored in the literature.

The robust empirical findings corroborate the hypothesis that TI captures important aspects of the dynamic trading equilibrium, at least within these large tick markets. It is a first-order task to seek extensions that can corroborate or disprove the hypothesis for alternative asset markets. Unfortunately, this task is complicated by the inevitable frictions that arise when exploring the high-frequency hypothesis for assets traded in fragmented markets subject to rapid shifts in trading technology, regulation, and general market structure. Moreover, it is often difficult to access data that allow for effective identification of the active party in each individual transaction which is critical for identification of the trade size.

Our results have implications for the assessment of market design and the trade-offs involved in changes to the contract specification and trading environment. For example, one may use the

results to gauge the cost and benefits of the changes to the minimum tick sizes implemented for a subset of the FX futures contracts during our sample period, conditional of assigning relative cost and benefit to the lower marginal trading costs versus the deterioration in market depth. Likewise, TI is an equilibrium relation. During periods of stress, it is natural to study deviations from the relation as indications of market malfunction, aiding in the design of market surveillance and the potential establishment of objective rules governing trading halts.

## References

- Andersen, Torben G.** 1996. “Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility.” *Journal of Finance*, 51(1): 169–204.
- Andersen, Torben G., and Tim Bollerslev.** 1997. “Intraday Periodicity and Volatility Persistence in Financial Markets.” *Journal of Empirical Finance*, 4: 115–158.
- Andersen, Torben G., and Tim Bollerslev.** 1998. “DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies.” *Journal of Finance*, 53: 219–265.
- Andersen, Torben G., Oleg Bondarenko, Albert S Kyle, and Anna A Obizhaeva.** 2018. “Intraday trading invariance in the E-mini S&P 500 futures market.” *Available at SSRN*.
- Bae, Kyoungun, Albert S. Kyle, Eun Jung Lee, and Anna A. Obizhaeva.** 2020. “Invariance of Buy-Sell Switching Points.” New Economic School; Working Paper, No. 273.
- Benzaquen, Michael, Jonathan Donier, and Jean-Philippe Bouchaud.** 2016. “Unravelling the Trading Invariance Hypothesis.” *Market Microstructure and Liquidity*, 2(3&4): 1650009.
- Brownlees, Christian T., Fabrizio Cipollini, and Giampiero M. Gallo.** 2011. “Intradaily Volume Modeling and Prediction for Algorithmic Trading.” *Journal of Financial Econometrics*, 9: 489–518.
- Brownlees, Christian T., Fabrizio Cipollini, and Giampiero M. Gallo.** 2012. “Multiplicative Error Models.” In *Handbook of Financial Engineering and Econometrics: Volatility Models and Their Applications*, ed. Luc Bauwens, Christian Hafner and Sebastian Laurent, 225–248. Wiley.
- Bucci, Frédéric, Fabrizio Lillo, Jean-Philippe Bouchaud, and Michael Benzaquen.** 2020. “Unravelling the Trading Invariance Hypothesis.” *Quantitative Finance*, 20(7): 1059–1068.
- Budish, Eric, Peter Crampton, and John Shim.** 2015. “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response.” *Quarterly Journal of Economics*, 130(4): 1547–1621.

- Darolles, Serge, Gaëlle Le Fol, and Gulten Mero.** 2015. “Measuring the Liquidity Part of Volume.” *Journal of Banking and Finance*, 50: 92–105.
- Darolles, Serge, Gaëlle Le Fol, and Gulten Mero.** 2017. “Mixture of Distribution Hypothesis: Analyzing Daily Liquidity Frictions and Information Flows.” *Journal of Econometrics*, 201(Winter): 367–383.
- Engle, Robert F.** 2002. “New Frontiers for ARCH Models.” *Journal of Applied Econometrics*, 17: 425–446.
- Epps, Thomas W., and Mary Lee Epps.** 1976. “The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distributions Hypothesis.” *Econometrica*, 44(2): 305–321.
- Hou, Ai Jun, Lars L. Nordèn, and Caihong Xu.** 2024. “Futures trading costs and market microstructure invariance: Identifying bet activity.” *Journal of Futures Markets*, 1–22.
- Jacod, Jean, Yingying Li, and Xinghua Zheng.** 2017. “Statistical Properties of Microstructure Noise.” *Econometrica*, 85(4): 1133–1174.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson.** 1994. “Transactions, Volume and Volatility.” *Review of Financial Studies*, 7(4): 631–651.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016*a*. “Large Bets and Stock Market Crashes.” Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2023776>.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016*b*. “Market Microstructure Invariance: Empirical Hypotheses.” *Econometrica*, 84(4): 1345–1404.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2017. “Dimensional Analysis, Leverage Neutrality, and Market Microstructure Invariance.” Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2785559>.
- Kyle, Albert S., Anna A. Obizhaeva, and Tugkan Tuzun.** 2020. “Microstructure invariance in U.S. stock market trades.” *Journal of Financial Markets*, 49: 1–32.
- Kyle, Albert S., Anna A. Obizhaeva, Nitish R. Sinha, and Tugkan Tuzun.** 2017. “News Articles and the Equity Trading.” Working Paper, available at <http://dx.doi.org/10.2139/ssrn.1786124>.
- Liesenfeld, Roman.** 2001. “A Generalized Bivariate Mixture Model for Stock Price Volatility and Trading Volume.” *Journal of Econometrics*, 104: 141–178.
- Li Z. Merrick, and Oliver Linton.** 2022. “A ReMeDI for Microstructure Noise.” *Econometrica*, 90(1): 367–389.
- Richardson, Matthew, and Tom Smith.** 1994. “A Direct Test of the Mixture of Distributions Hypothesis: Measuring the Daily Flow of Information.” *Journal of Financial and Quantitative Analysis*, 29(1): 101–116.
- Tauchen, George E., and Mark Pitts.** 1983. “The Price Variability-Volume Relationship on Speculative Markets.” *Econometrica*, 51: 485–505.




# Appendix

## A. Additional Figures

Figure A1 documents the intraday activity patterns for GBP. Similar to Figure 3, it shows the mean value of each variable across the sample for all 5-minute intervals during the continuous GLOBEX trading session, from  $-5$  (i.e., 7:00 p.m. Chicago Time (CT) on the preceding day) to 15 (3:00 p.m. CT). The dashed vertical lines delineate the main trading zones. The two curves in each panel reflect the split into two non-overlapping subsamples - from the start of the sample until January 11, 2016, and from that date through 2021. These subsamples coincide with those used for the euro in Figure 3. There was no tick change for GBP, and the overall levels for all six activity variables remained broadly stable across the two subsamples.


Figure A1: GBP Futures — Intraday Activity Series



**Intraday Averages for GBP Futures Before and After the Tick Size Change.** Volume  $V$  is the number of contracts traded in 5 minutes. The other variables are as indicated in Figures 1 and 2. The averages are computed for each 5-minute interval across the trading day. The tick change occurred on Jan 11, 2016. The blue (green) line shows averages before (after) January 11, 2016.

Figure A2 depicts the OLS fit corresponding to the left panel in Figure 4 for the other four FX futures contracts.


Figure A2: OLS Regression  $s$  vs  $v$



OLS regression of log realized volatility  $s$  on log contract volume  $v$  for four currencies: AUD, CAD, EUR, and JPY. The solid line is the fitted OLS regression line and the dashed line has the slope constrained to unity ( $\gamma_v = 1$ ). Red dots indicate the peak of the COVID crisis, i.e., dates from March 15, 2020 to April 15, 2020.


Figures A3-A6 plot alternative invariance relations for AUD, CAD, GBP, and JPY. The six panels of each figure display 10-day moving averages for SV, SN, MI1, MI2, TI, and TI\*. The first four invariance relations are evaluated for the theoretically prescribed parameters. The TI panel imposes the theoretically predicted value of  $\gamma_B = -2$  and sets  $\gamma_D = 1/2$ , motivated by the empirical findings, but without any theoretical basis. The TI\* panel relies on the estimated coefficients ( $\hat{\gamma}_B, \hat{\gamma}_D$ ) for a given currency.

Figure A3: AUD Futures — 6 Invariants Over Time




The panels display 10-day moving averages for the alternative invariance relations: SV, SN, MI1, MI2, TI, and TI\*. The first four relations are evaluated for the theoretically prescribed parameters. TI uses the slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ , whereas TI\* uses the estimated slopes  $(\hat{\gamma}_B, \hat{\gamma}_D)$ .

Figure A4: CAD Futures — 6 Invariants Over Time




The panels display 10-day moving averages for the alternative invariance relations: SV, SN, MI1, MI2, TI, and TI\*. The first four relations are evaluated for the theoretically prescribed parameters. TI uses the slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ , whereas TI\* uses the estimated slopes  $(\hat{\gamma}_B, \hat{\gamma}_D)$ .

Figure A5: GBP Futures — 6 Invariants Over Time



The panels display 10-day moving averages for the alternative invariance relations: SV, SN, MI1, MI2, TI, and TI\*. The first four relations are evaluated for the theoretically prescribed parameters. TI uses the slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ , whereas TI\* uses the estimated slopes  $(\hat{\gamma}_B, \hat{\gamma}_D)$ .

Figure A6: JPY Futures — 6 Invariants Over Time



The panels display 10-day moving averages for the alternative invariance relations: SV, SN, MI1, MI2, TI, and TI\*. The first four relations are evaluated for the theoretically prescribed parameters. TI uses the slopes  $(\gamma_B, \gamma_D) = (-2, 0.5)$ , whereas TI\* uses the estimated slopes  $(\hat{\gamma}_B, \hat{\gamma}_D)$ .

## B. FX Futures Details

Table A1: Descriptive Statistics

|                  | Mean   | Min   | Max    | Ratio | Night | Asia  | Europe | America |
|------------------|--------|-------|--------|-------|-------|-------|--------|---------|
| <b>AUD</b>       |        |       |        |       |       |       |        |         |
| Volatility (%)   | 9.12   | 4.75  | 20.62  | 4.34  | 6.78  | 8.99  | 10.04  | 9.41    |
| Volume           | 66.04  | 14.35 | 247.77 | 17.27 | 26.84 | 48.70 | 78.16  | 90.67   |
| # Trades         | 17.59  | 3.63  | 62.92  | 17.34 | 7.58  | 13.85 | 20.65  | 23.16   |
| Trade Size       | 3.63   | 3.13  | 4.27   | 1.36  | 3.41  | 3.43  | 3.70   | 3.88    |
| Relative BAS, bp | 1.38   | 1.31  | 1.76   | 1.34  | 1.51  | 1.40  | 1.36   | 1.33    |
| BAS, tick        | 1.14   | 1.08  | 1.46   | 1.36  | 1.25  | 1.16  | 1.13   | 1.10    |
| Book Depth       | 39.47  | 16.84 | 54.88  | 3.26  | 25.08 | 33.49 | 43.35  | 48.68   |
| <b>CAD</b>       |        |       |        |       |       |       |        |         |
| Volatility (%)   | 6.53   | 3.26  | 22.13  | 6.79  | 4.27  | 4.53  | 7.76   | 8.50    |
| Volume           | 48.58  | 7.97  | 301.81 | 37.89 | 15.88 | 15.89 | 57.24  | 90.22   |
| # Trades         | 14.87  | 2.80  | 86.07  | 30.71 | 4.92  | 5.31  | 17.78  | 26.84   |
| Trade Size       | 3.07   | 2.60  | 3.81   | 1.47  | 2.99  | 2.82  | 3.09   | 3.37    |
| Relative BAS, bp | 1.03   | 0.95  | 1.32   | 1.40  | 1.13  | 1.05  | 1.01   | 0.98    |
| BAS, tick        | 1.24   | 1.14  | 1.57   | 1.38  | 1.35  | 1.27  | 1.22   | 1.19    |
| Book Depth       | 21.88  | 10.89 | 29.26  | 2.69  | 15.30 | 20.10 | 23.18  | 25.51   |
| <b>EUR</b>       |        |       |        |       |       |       |        |         |
| Volatility (%)   | 6.44   | 3.23  | 18.27  | 5.66  | 3.80  | 4.70  | 8.39   | 7.58    |
| Volume           | 142.65 | 19.04 | 666.95 | 35.03 | 30.75 | 52.42 | 220.72 | 213.41  |
| # Trades         | 42.35  | 6.71  | 193.35 | 28.83 | 10.15 | 17.78 | 64.76  | 61.25   |
| Trade Size       | 3.19   | 2.68  | 3.73   | 1.39  | 2.95  | 2.87  | 3.37   | 3.46    |
| Relative BAS, bp | 0.65   | 0.62  | 0.79   | 1.28  | 0.71  | 0.65  | 0.64   | 0.63    |
| BAS, tick        | 1.19   | 1.13  | 1.45   | 1.28  | 1.29  | 1.19  | 1.17   | 1.16    |
| Book Depth       | 25.23  | 12.40 | 34.40  | 2.77  | 16.56 | 21.76 | 28.17  | 30.01   |
| <b>GBP</b>       |        |       |        |       |       |       |        |         |
| Volatility (%)   | 7.13   | 3.42  | 16.21  | 4.75  | 4.18  | 4.99  | 9.82   | 8.12    |
| Volume           | 72.10  | 10.13 | 324.27 | 32.01 | 16.68 | 24.05 | 113.99 | 107.53  |
| # Trades         | 20.78  | 3.41  | 84.05  | 24.63 | 5.06  | 7.53  | 33.10  | 29.98   |
| Trade Size       | 3.23   | 2.68  | 4.02   | 1.50  | 3.06  | 2.94  | 3.32   | 3.51    |
| Relative BAS, bp | 0.92   | 0.82  | 1.28   | 1.56  | 1.05  | 0.97  | 0.89   | 0.85    |
| BAS, tick        | 1.29   | 1.15  | 1.80   | 1.57  | 1.47  | 1.35  | 1.24   | 1.19    |
| Book Depth       | 18.63  | 10.20 | 27.68  | 2.71  | 13.24 | 16.28 | 19.33  | 22.96   |
| <b>JPY</b>       |        |       |        |       |       |       |        |         |
| Volatility (%)   | 7.08   | 4.03  | 17.89  | 4.44  | 5.47  | 6.68  | 7.54   | 7.79    |
| Volume           | 94.14  | 18.32 | 412.58 | 22.52 | 37.72 | 68.57 | 107.40 | 134.47  |
| # Trades         | 28.62  | 5.56  | 118.67 | 21.34 | 12.26 | 22.50 | 32.17  | 39.21   |
| Trade Size       | 3.22   | 2.76  | 3.74   | 1.35  | 3.07  | 3.06  | 3.27   | 3.42    |
| Relative BAS, bp | 0.81   | 0.76  | 1.13   | 1.49  | 0.93  | 0.81  | 0.79   | 0.78    |
| BAS, tick        | 1.22   | 1.13  | 1.63   | 1.44  | 1.37  | 1.21  | 1.19   | 1.17    |
| Book Depth       | 27.04  | 11.11 | 36.47  | 3.28  | 17.81 | 24.43 | 29.26  | 31.90   |

Average statistics from January 2013 to December 2021. **Volatility** is measured as realized volatility computed from 30-second squared returns, annualized and stated in percents. **Volume** and **# Trades** are the number of contracts transacted and the number active trades, reported *per minute*. The **Average Trade Size** is given in terms of contracts. The average bid-ask spread (BAS = Ask - Bid) is reported as **Relative BAS** (BAS/*P* in basis points) and **BAS, tick** (in ticks). **Book Depth** is the average of Bid and Ask depth in contracts at the top of the book. The first four columns refer to averages of the mean, minimum and maximum daily values and lastly the ratio of the latter two. The final four columns report average values over different regions, starting with the overnight period, and then the Asian, European and American regimes.

Table A2: Descriptive Statistics: Additional Currencies

|                  | Mean  | Min  | Max   | Ratio | Night | Asia  | Europe | America |
|------------------|-------|------|-------|-------|-------|-------|--------|---------|
| <b>CHF:</b>      |       |      |       |       |       |       |        |         |
| Volatility (%)   | 6.73  | 3.16 | 17.76 | 5.62  | 3.83  | 4.70  | 8.82   | 8.15    |
| Volume           | 18.54 | 3.21 | 83.61 | 26.03 | 4.98  | 6.45  | 26.73  | 29.63   |
| # Trades         | 7.87  | 1.38 | 33.27 | 24.12 | 2.05  | 2.83  | 11.46  | 12.38   |
| Trade Size       | 2.21  | 2.01 | 2.74  | 1.36  | 2.22  | 2.11  | 2.23   | 2.31    |
| Relative BAS, bp | 1.33  | 1.17 | 1.89  | 1.61  | 1.55  | 1.40  | 1.27   | 1.23    |
| BAS, tick        | 1.40  | 1.23 | 1.98  | 1.62  | 1.63  | 1.46  | 1.34   | 1.29    |
| Book Depth       | 15.32 | 7.45 | 21.77 | 2.92  | 10.52 | 13.50 | 16.34  | 18.50   |
| <b>NZD:</b>      |       |      |       |       |       |       |        |         |
| Volatility (%)   | 9.45  | 5.63 | 21.82 | 3.88  | 7.48  | 8.91  | 10.47  | 9.92    |
| Volume           | 15.85 | 5.59 | 73.24 | 13.11 | 8.12  | 10.60 | 18.36  | 22.56   |
| # Trades         | 5.56  | 1.56 | 22.86 | 14.63 | 2.70  | 3.84  | 6.54   | 7.74    |
| Trade Size       | 2.65  | 2.44 | 3.29  | 1.35  | 2.73  | 2.54  | 2.63   | 2.76    |
| Relative BAS, bp | 2.05  | 1.71 | 3.11  | 1.82  | 2.48  | 2.23  | 1.92   | 1.80    |
| BAS, tick        | 1.52  | 1.25 | 2.27  | 1.81  | 1.84  | 1.66  | 1.42   | 1.32    |
| Book Depth       | 19.32 | 9.54 | 24.22 | 2.54  | 13.70 | 17.73 | 20.63  | 22.33   |

Average statistics from January 2013 to December 2021. **Volatility** is measured as realized volatility computed from 30-second squared returns, annualized and stated in percents. **Volume** and **# Trades** are the number of contracts transacted and the number active trades, reported *per minute*. The **Average Trade Size** is given in terms of contracts. The average bid-ask spread (BAS = Ask - Bid) is reported as **Relative BAS** (BAS/ $P$  in basis points) and **BAS, tick** (in ticks). **Book Depth** is the average of Bid and Ask depth in contracts at the top of the book. The first four columns refer to averages of the mean, minimum and maximum daily values and lastly the ratio of the latter two. The final four columns report average values over different regions, starting with the overnight period, and then the Asian, European and American regimes.

Tables A1 and A2 provide a range of summary statistics for the FX contracts used in our analysis. Table A1 documents the statistics for the five currencies used in our main analysis, while Table A2 documents those statistics for the less liquid currencies CHF and NZD. In these tables, the first three columns refer to averages of the mean, minimum, and maximum daily values. The fourth column is the ratio of the latter two, conveying a sense of variability over time. The final four columns report average values over different regions, starting with the overnight period, and then the Asian, European and American regimes.

Table A3: Contract Specification

|                        | AUD   | CAD   | GBP  | EUR   | JPY   | CHF   | NZD   |
|------------------------|-------|-------|------|-------|-------|-------|-------|
| Multiplier, 000s       | 100.0 | 100.0 | 62.5 | 125.0 | 125.0 | 125.0 | 100.0 |
| Tick Size, \$          | 5.00  | 5.00  | 6.25 | 6.25  | 6.25  | 6.25  | 5.00  |
| Price                  | 0.78  | 0.80  | 1.41 | 1.18  | 0.92  | 1.05  | 0.72  |
| Notional Value, \$000s | 78.0  | 79.7  | 87.9 | 147.5 | 115.2 | 131.3 | 71.6  |
| Relative Tick, bp      | 1.22  | 0.84  | 0.72 | 0.55  | 0.68  | 0.95  | 1.38  |

The statistics are computed for the seven FX futures contracts at the end of our sample in December 2021 and include contract multiplier, minimum tick size, price, notional value, and relative tick size (Tick/ $P$ ).

Table A3 provides details on the specification of the individual FX futures contracts. As noted in the main text, the minimum tick size changes during some point of our sample period for all FX futures contracts but GBP.

## C. GMM Estimation

The basic TI representation for a given asset on trading day  $d$  follows from equation (10)<sup>25</sup>

$$TI_d = (s_d - n_d) + \gamma_B (b_d - p_d) + \gamma_D (d_d - q_d) + c, \quad (14)$$

where the deviation  $TI_d$  is a mean-zero, possibly serially correlated error process with finite fourth moments, and the  $3 \times 1$  coefficient vector  $\nu = (c, \gamma_B, \gamma_D)'$  represents our TI parameters of primary interest. If we stipulate that  $TI_d$  is uncorrelated with past information, then we may use lagged realizations of the activity variables as instruments. In that case, we define the  $4 \times 1$  vector  $z_d = (1, (s_d - n_d), (b_d - p_d), (q_d - d_d))'$ , and we denote the average value of the prior  $M$  observations of  $z_d$  by  $\bar{z}_d^{(M)} = \sum_{m=1}^M z_{d-m}/M$ . We then have,

$$\mathbb{E}[TI_d \cdot \bar{z}_d^{(M)}(\nu)] = 0, \quad 0 < M < d, \quad (15)$$

which allows for straightforward application of standard GMM inference techniques.<sup>26</sup>

If the assumption that the  $TI_d$  residual is uncorrelated with the past realized activity variables is untenable - and we show later on that, indeed, this is an issue for our TI system - then we do need additional identification conditions. Our approach is to accommodate the residual serial correlation via a general ARMA( $p, q$ ) representation, letting the coefficients differ across currencies to allow for distinct degrees of friction-induced persistence. That is,  $TI_d$  in equation (14) now follows a mean-zero ARMA( $p, q$ ) process,

$$TI_d - \psi_1 TI_{d-1} - \dots - \psi_p TI_{d-p} = \epsilon_d + \eta_1 \epsilon_{d-1} + \dots + \eta_q \epsilon_{d-q}, \quad (16)$$

and  $\epsilon_d$  is uncorrelated with prior information, has mean zero and finite fourth moments.<sup>27</sup> Of course, if the TI deviations are uncorrelated, then the ARMA( $p, q$ ) process has  $p = q = 0$ , and  $TI_d = \epsilon_d$ .

For each currency, equation (16) introduces the auxiliary ARMA  $(p+q) \times 1$  coefficient vector  $\zeta = (\psi_1, \dots, \psi_p, \eta_1, \dots, \eta_q)'$ . Defining  $u_d = (TI_d, \dots, TI_{d-p+1}, \epsilon_d, \dots, \epsilon_{d-q+1})'$  naturally leads to the following moment conditions,<sup>28</sup>

$$\mathbb{E}[\epsilon_d(\nu, \zeta) \cdot \bar{z}_d^{(M)}(\nu)] = 0, \quad 0 < M < d, \quad (17)$$

$$\mathbb{E}[\epsilon_d(\nu, \zeta) \cdot u_{d-j}(\nu, \zeta)] = 0, \quad k \geq 1. \quad (18)$$

For the ARMA coefficients, we estimate separate coefficients across the currencies to allow for distinct friction-induced dependency structures in the TI residual process.

Hence, for the general case, for each currency, we have four moment conditions in equation (17) reflecting the orthogonality of the ARMA innovation to the lagged activity variables (plus a constant), and then  $k \cdot (p+q)$  orthogonality conditions in equation (18) implied by the ARMA specification. This generates a regular GMM system with  $4 + k \cdot (p+q)$  equations, leading to  $1 + (k-1) \cdot (p+q)$  over-identifying restrictions. Standard GMM estimation and inference

<sup>25</sup>For notational clarity, we drop the asset-specific superscript ( $j$ ) here.

<sup>26</sup>In this scenario, we typically find strong serial correlation in the moment conditions, and we employ the usual Newey-West HAC error correction for inference.

<sup>27</sup>As a consequence, our TI-ARMA representations should eliminate any serial correlation in our moment conditions, and we do not employ any “Newey-West lags” in the inference process for these specifications reported below. As a robustness check, we also found that the results are qualitatively identical with inclusion of such lags.

<sup>28</sup>These conditions correspond to moments imposed during regular likelihood estimation of ARMA processes.

procedures may be applied straightforwardly, with the overall fit assessed through the scaled value of the objective function (the  $J$ -statistic) and coefficient restrictions by regular Wald tests. Moreover, whenever serial dependence in the TI residuals seems problematic due to a correlation with the lagged instruments, we assess the ARMA residuals through standard Ljung-Box statistics.

Finally, we also provide currency-pooled estimates for the invariance systems. In this case, we average both the instruments and the relevant innovation component across currencies to obtain the same set of moments and parameters as for the individual currencies, but now the information is aggregated, enhancing efficiency under the null hypothesis.