Hidden Liquidity in CME Futures

Hidden liquidity is resting volume available in the order book, that is not visible in market data but that can be traded against by a suitable marketable order. There are two types of hidden liquidity:

- *Direct* hidden liquidity results from orders that display only a fraction of their size, commonly called "iceberg" orders. Direct hidden liquidity is typically released into the order book only after a trade has consumed the displayed quantity, unless the displayed quantity is not enough to satisfy the aggressive order.
- *Implied* liquidity results from combinations of calendar and intercommodity spreads. *Implied hidden* liquidity arises from combinations that are computed but not displayed according to the rules of the CME Globex matching engine. It is available for immediate execution by a suitable aggressive order.

We show a way to estimate both direct and hidden liquidity, using special features of CME Globex market data. We estimate the volume of trading that executes against hidden liquidity for all major CME futures products. Finally, we show that traders tend to use "round" numbers (such as 10 or 20 lots, rather than 13) in the `MaxShow` parameter.

1 Hidden liquidity: definition and examples

Figure 1 shows our basic method for estimating hidden liquidity. If all liquidity were visible, then when a trade executes, we would expect the post-trade quote size to be less than the pre-trade quote size by exactly the size of the trade. Any excess of the post-trade quote size over that expected value represents hidden liquidity. This estimation method uses only publicly available market data; we do not use any internal QB execution data.

Matching trade and quote data is a notoriously difficult topic, since trades and quotes usually come through different data pipelines (as they do with CME market data) having different latencies. But we use the sequence numbers provided in CME market data, which are strictly sequential for each date and each symbol although not sequential across different symbols, to precisely and unambiguously determine the quotes immediately before and immediately after each trade.

Furthermore, CME market data allows us to estimate both direct and implied liquidity separately and independently. For quotes, the market data provides both a direct and an implied quote size. For trades, since the introduction of the MDP 3.0 market data format in September 2015, CME provides a breakdown that allows us, for each trade, to determine the quantity that executed against direct quotes in the limit order book for that contract, and the quantity that executed against implied quotes arising from other contracts such as calendar spreads and other outrights.

Thus, direct hidden liquidity is determined by the discrepancy computed from the pre-trade and post-trade direct quote volume, and the direct trade size; implied hidden liquidity is determined using implied quote volumes and the implied trade size.

This market data is "aggregated:" for each direct quote CME disseminates only the total number of lots available and the number of separate orders making up that quote, but not the individual size of each order. For implied quotes they give only the total number of lots, since the number of separate orders is not well defined. Other exchanges do provide, and CME will provide later in 2016, an "unaggregated" market data stream with full details, and then it may be possible to repeat this study in more detail.

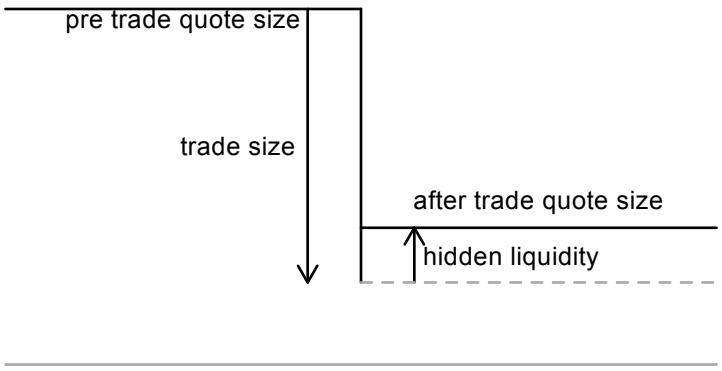

Figure 1: Hidden liquidity schematic. Hidden liquidity is estimated as the discrepancy between the expected and actual post-trade quote size.

ZBU6		Trades			Quotes	
rseq	time (UTC)	prc	dsiz	aggr	dbid	dbsiz
3282828	19:48:13.340				172-08	34
3282831	19:48:13.340	172-08	15	S		
3282833	19:48:13:340				172-08	119

Table 1: Direct hidden liquidity example for the US Treasury Bond futures contract ZBU6 on July 19, 2016. On the trade, "dsiz" is the volume that CME tells us printed against direct quote volume, and "aggress" is the aggressor direction. On the quote, "dbid" is the direct bid price, and "dbsiz" the total number of displayed lots of direct liquidity. The implied bid price was lower than the direct bid at this time and is not relevant.

It is possible for the post-trade quote size to be *smaller* than the pre-trade quote size minus the trade size. The most significant way in which this happens arises from self-trade prevention with the "cancel resting" instruction. We have estimated that in practice, this quantity is negligible compared to hidden liquidity, and therefore we neglect this effect.

1.1 Direct hidden liquidity

Let us give an example of direct hidden liquidity. Table 1 shows trades and quotes on July 19, 2016, for the Sep 2016 US Treasury Bond futures contract ZBU6. These transactions are on the bid side of the order book, at price 172-08.[1]

Before the trade, the direct bid was 6 separate orders, totalling 34 lots. An aggressive sell trade for 15 lots prints at the bid. We would expect the subsequent quote size to be $34 - 15 = 19$. Instead, the subsequent quote size is 119 lots, indicating that 100 lots had been hidden and were revealed following the trade and before CME updated the order book.

The sequence numbers "rseq" are strictly increasing, though not sequential. Gaps are present because market data messages are also used to update other levels in the order book, and because bid and ask updates require two separate messages. Although the exchange time stamps on these trades and quotes appear to be identical at millisecond precision, in fact they were separated by nonzero fractions of milliseconds. However, the trade in this

[1] By Treasury price conventions, "172-08" means $172 + (8/32) = 172.25$.

ZBU6		Trades				Quotes			
rseq	time (UTC)	prc	dsiz	isiz	aggr	dbid	dbsiz	ibid	ibsiz
3144752	17:48:04.346					172-12	21	172-11	70
3144753	17:48:05.575	172-12	21	52	S				
3144755	17:48:05.577					172-11	261	172-11	70

Table 2: Implied hidden liquidity example for the US Treasury Bond futures contract ZBU6 on July 19, 2016. The total trade size was 73 lots, of which 21 printed against visible direct liquidity at 172-12, and 52 printed against undisplayed implied liquidity at 172-12 (implied bid quantity was visible only at the lower value of 172-11).

UBU6		Trades				Quotes			
rseq	time (UTC)	prc	dsiz	isiz	aggr	dbid	dbsiz	ibid	ibsiz
2090918	17:48:04.111					187-08	68	187-08	4
2090920	17:48:05.575	187-08	0	39					
2090922	17:48:05.577					187-08	29	187-08	4

BOB 04-03 U6		Trades				Quotes			
rseq	time (UTC)	prc	dsiz	isiz	aggr	dbid	dbsiz	ibid	ibsiz
1396427	17:48:04.111					−0-01	13	−0-03	5
1396429	17:48:05.575	−0-01	0	13					
1396431	17:48:05.577					−0-02	92	−0-03	5

Table 3: Source of hidden liquidity for ZBU6 in Table 2. At exactly the same time that the 73 lots traded in ZBU6, we see implied trades in the Ultra futures contract UBU6 and the ZB–UB intercommodity spread BOB. See text for discussion.

example was received by our system 7 msec before the quote against which it executed, illustrating the need for an unambiguous way to order market data events.

Direct hidden liquidity cannot be precisely computed in real time, only after the fact when trades are observed. It may be estimated using various statistical techniques; for example, one may suspect that if liquidity is consistently and repeatedly released following trades, that there remains a large reservoir of nondisplayed liquidity.

1.2 Implied hidden liquidity

Implied hidden liquidity comes from implied pricing relationships within the entire product universe. For example, a calendar spread may combine with an outright contract to give a quote in another outright contract. The CME exchange calculates a broad set of these relationships, but displays only a subset of what is calculated.

Table 2 shows an example of implied hidden liquidity, again on the bid side of the order book for the US Treasury Bond futures contract ZBU6 on July 19, 2016. In this trade, 52 lots printed against invisible implied liquidity at 172-12. The direct print of 21 lots exhausted the visible bid size at 172-12, and the best visible bid dropped to 172-11. We do not know whether invisible implied bid liquidity remained at 172-12 (except that shortly after this event, visible *offer* liquidity filled in at 172-12).

Table 3 shows the source of this invisible implied liquidity. It arises from a combination of the Ultra Treasury futures contract UBU6, along with the intercommodity spread BOB 04–

03 U6. Each long unit of this spread represents 4 long units of ZBU6 and 3 short units of UBU6. In symbols, we write

$$[\text{BOB 04-03 U6}] = 4\,[\text{ZBZ6}] - 3\,[\text{UBU6}]$$

The trades shown in Tables 2 and 3 are 13 lots [BOB 04-03 U6], 39 lots [UBU6], and 52 lots [ZBZ6], which are exactly in these proportions. The trades in the other two legs [UBU6] and [BOB 04-03 U6] have no aggressiveness tag, since they were generated through an implied relationship. The coefficients 4 and 3 in the definition of the BOB intercommodity spread are chosen by CME each quarterly cycle, and it is often the case that more than one intercommodity spread may exist between two outright contracts, with different ratios.

The above relationship may be rearranged to give the implied quote for ZBZ6:

$$4\,[\text{ZBZ6}] = 3\,[\text{UBU6}] + [\text{BOB 04-03 U6}]$$

That is, selling 4 lots of [UBU6] at its bid, and selling 1 lot of [BOB 04-03 U6] at its bid, is equivalent to selling 3 lots of [ZBZ6] at an implied bid computed by adding the prices. The price calculation is slightly complicated, since the spread prices are quoted relative to the previous day's settlement prices for the legs. But when the arithmetic is properly done, the price computed is exactly the trade price for ZBZ6 shown in Table 2.

Implied hidden liquidity can easily be calculated by understanding and implementing the relationships that CME uses but does not display. We do do this at QB, to improve price performance in Treasury trading. (Similar techniques also yield substantial improvement in trading of Eurodollars and energy futures.)

2 Hidden liquidity comparison

The above examples do not demonstrate that hidden liquidity is significant in actual trading. To answer this question, we have measured the total quantity of hidden liquidity uncovered on each trade (whether or not that trade executed against hidden liquidity) divided by the total trade volume. This gives us an estimate of the fraction of hidden liquidity that is used "by accident." That is, this is hidden liquidity resulting from orders that likely were not explicitly intending to trade against it. If orders were sent using some estimate of the available hidden liquidity, then it is likely that this fraction would be much higher than it is.

Figure 2 shows the results for all major futures products on CME Group exchanges, from January through July 2016. Energy futures are separated by liquid and not liquid contracts, because traded volume in energy futures is usually divided among a handful of maturities and also with seasonal effects such as concentration in December contract. Particularly, hidden liquidity of Natural Gas (Henry Hub) Futures in illiquid maturity are as high as 26% of daily trading volume. We may make several broad observations:

- Hidden liquidity represents consistent fractions of daily volume across each product category. Agricultural futures and Energy futures in illiquid maturity have 15–20% hidden liquidity. Energy futures in liquid maturity and Metal futures have around 10% hidden liquidity. Rates futures have 5–10% hidden liquidity. Equity index futures and FX futures have 3–5% hidden liquidity. The E-mini S&P futures contract (ES) has around 8% hidden liquidity.
- Direct hidden liquidity is usually larger than implied hidden liquidity, especially there are no implied hidden liquidity in FX futures and Equity Index futures. However in energy futures product, implied hidden liquidity is more than direct hidden liquidity. The potential reason is traders heavily trade across different maturities, so spread is heavily traded which can generate more useful implied relationship.

- Hidden liquidity portion is larger in illiquid products.
  1. Between classes, Agricultural and Energy futures have more hidden liquidity.
  2. Between different products in the same class,
     - RBOB Gasoline (RB), NY Harbor ULSD (HO), and Natural Gas (NG) have more hidden liquidity than Crude Oil (CL);
     - Platinum (PL) and Palladium (PA) have more hidden liquidity than Gold (GC), Silver (SI), and Copper (HG);
     - The two year note (ZT) has less hidden liquidity than other rates products.
  3. Within same instrument and among different maturities, as we mentioned above, illiquid maturities have a higher fraction of hidden liquidity.

3 Displayed quantity distribution

Finally, we estimate the size distribution of the hidden direct liquidity. On CME, hidden liquidity is specified using the `MaxShow` order parameter, that specifies how much of the order is released into the order book after the visible part is consumed. The `MaxShow` parameter must be strictly greater than zero; completely hidden orders are not allowed. Thus, the size of the direct hidden liquidity is roughly the size of the `MaxShow` value for the resting orders that have hidden parts.

Figure 3 shows the distribution of the refreshed quantity for six US Treasury futures contracts, from January . The striking feature is that the sizes have strong clustering at “round” numbers such as 20 or 50 lots. In fact, we suspect that the few less-round numbers such as 120 lots come from two orders with hidden sizes 100 and 20 lots that are simultaneously refreshed.

We propose that this preference for round numbers arises because orders having hidden parts are largely submitted by human traders rather than algorithms, and humans have preferences for round numbers.

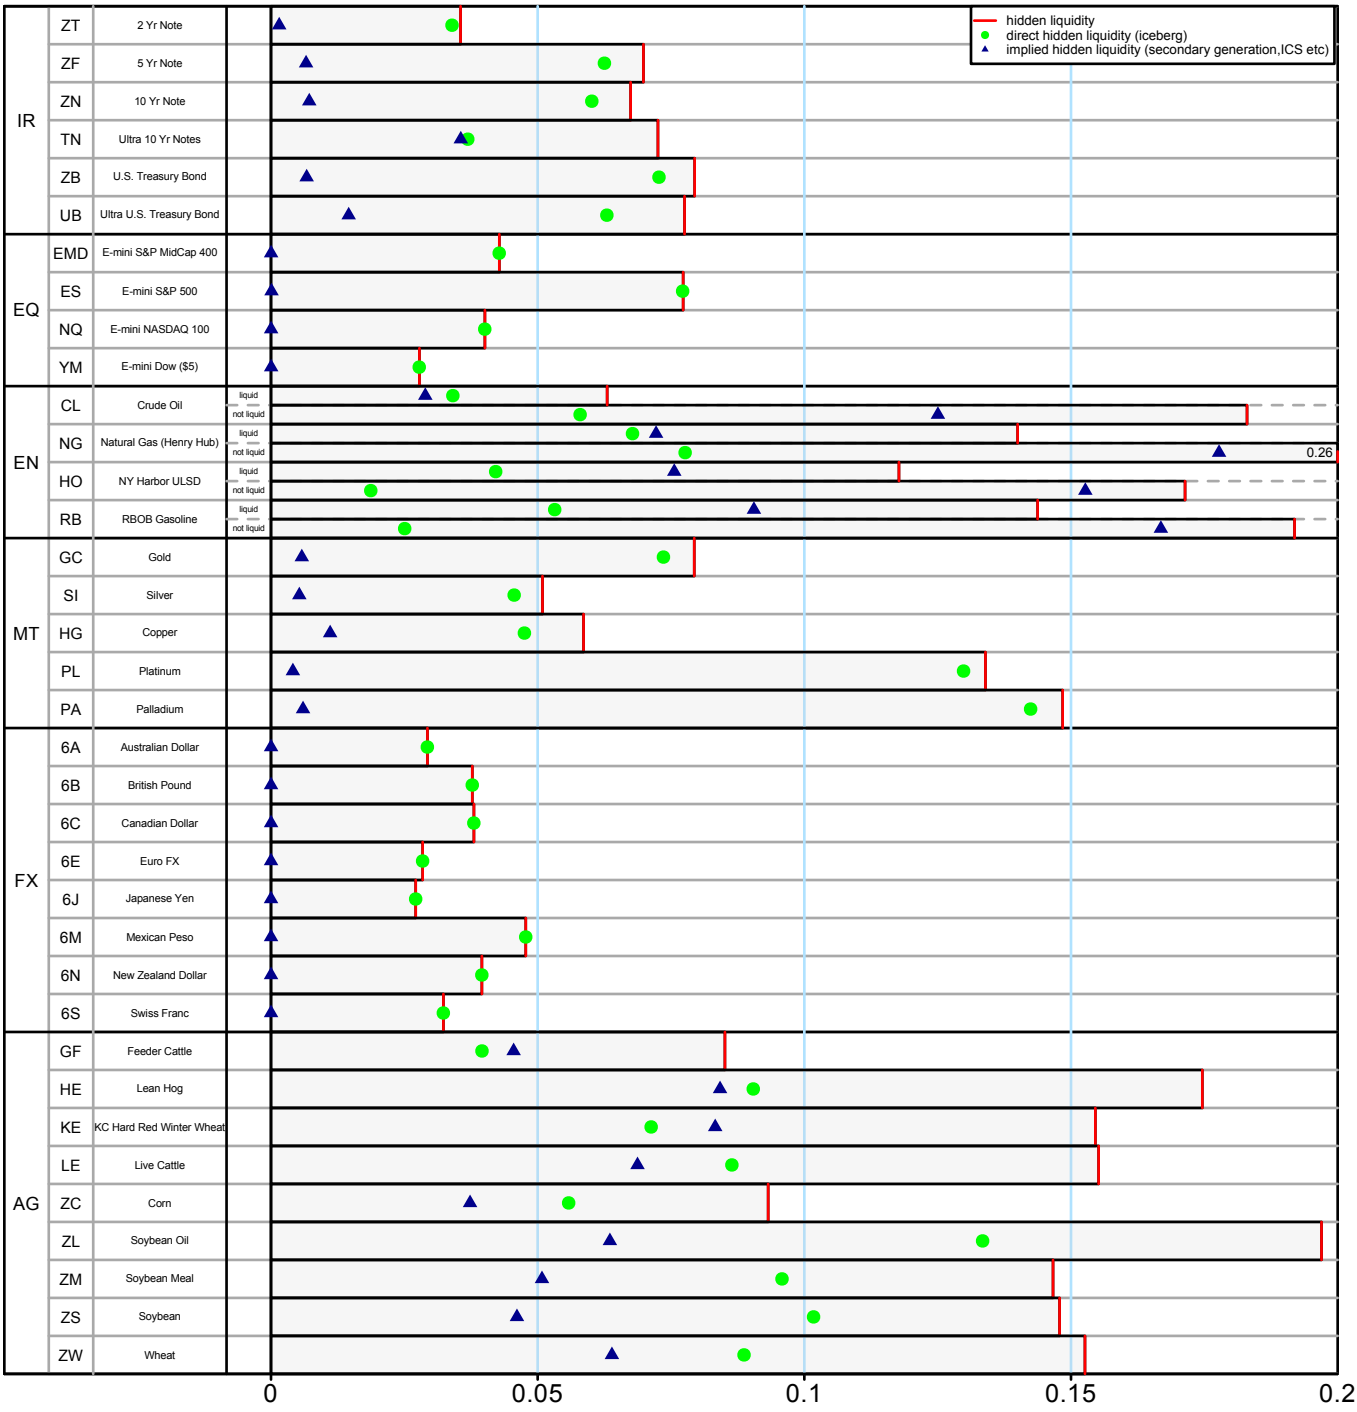

Figure 2: Hidden liquidity across products in CME, The horizontal axis is the hidden liquidity refilled volume to the daily trading volume. e.g. 0.2 means 20% of daily trading volume.

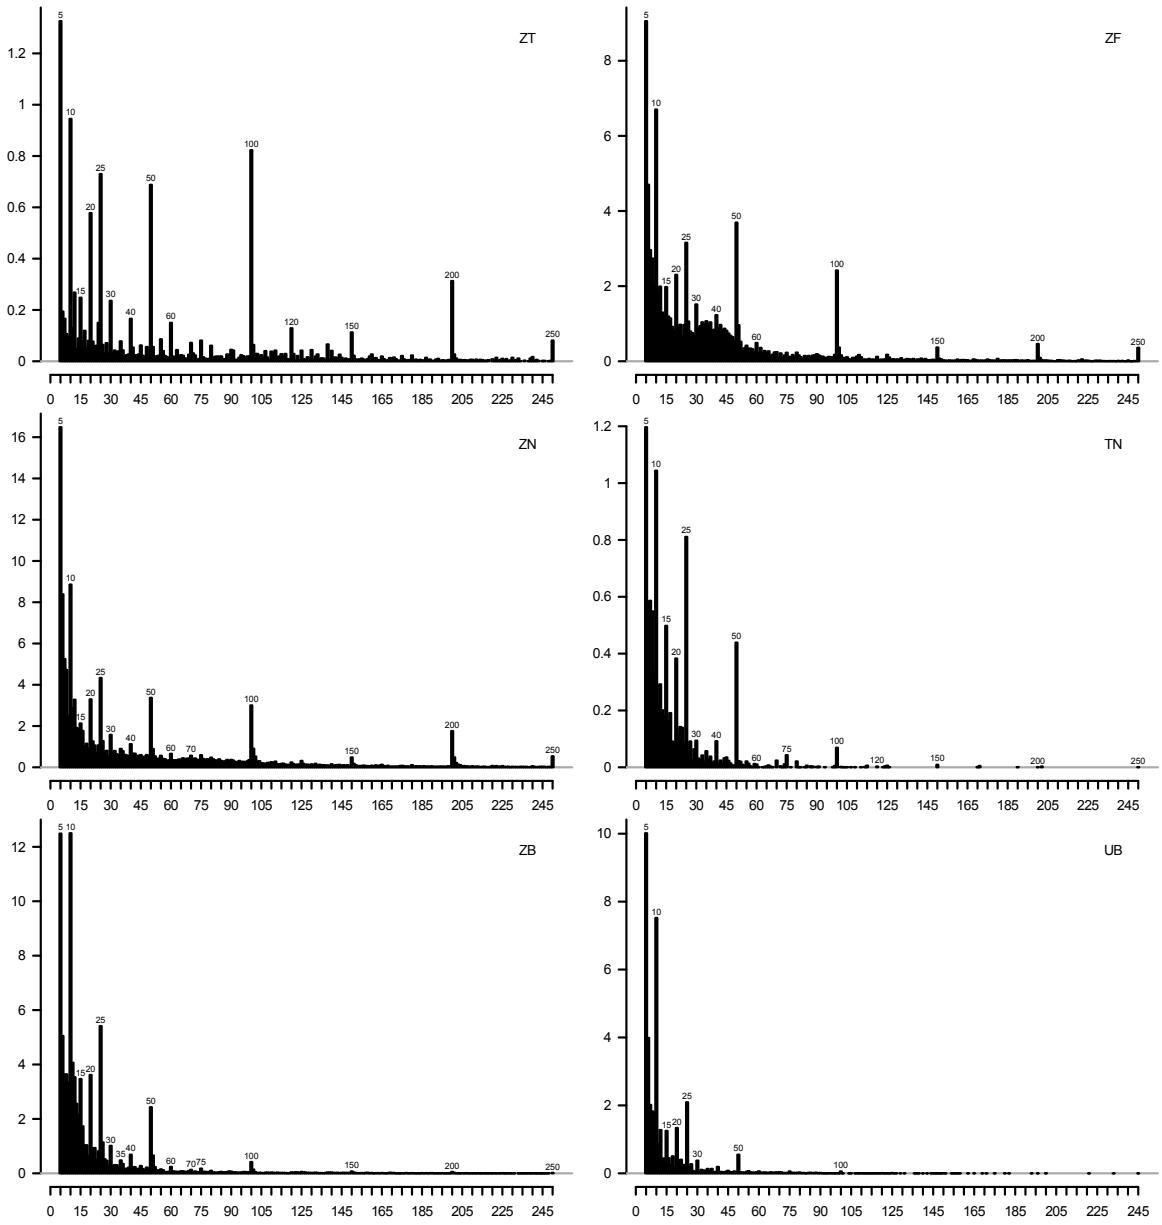

Figure 3: Size distribution of direct hidden liquidity in US Treasury futures from January through July 2016. Horizontal axis is size of refreshed quantity following a trade. Vertical axis is total quantity in thousands of lots. There is a clear preference for round numbers.