Execution Strategies in Fixed Income Markets

Robert Almgren*

April 2, 2013

Reducing trading costs and slippage is a universal concern of asset managers. Although the decision of what assets to hold is stil the most important aspect of investing, poor execution of trade decisions can subtract many basis points from overall return. Conversely, having an effective strategy to execute trades and to measure transaction costs can enhance returns: "A penny saved in slippage is a penny earned in alpha."

Execution techniques in equities have advanced far ahead of those in other markets such as futures, options, foreign exchange, and fixed income. One reason for this is the overall size of the equity markets, and the widespread use of active investment strategies. A second reason is the simplicity of the products themselves: for trading a single name of stock you need very little information beyond its price. Relationships between different stocks are at best weak. As a consequence, quant researchers in equity markets have focused intensively on the details of the execution process.

By contrast, fixed income products are inherently complex, and quantitatively minded researchers in the area have focused on such aspects as yield curve modeling, day counts, etc. Asset managers have not traditionally focused on measuring or managing execution costs, and have few effective tools to do so. However, the Securities Industry and Financial Markets Association (SIFMA) in 2008 (SIFMA Asset Management Group, 2008) noted that "It is clear that the duty to seek best execution imposed on an asset manager is the same regardless of whether the manager is undertaking equity or fixed-income transactions."

This chapter discusses some details of the fixed income markets that present special challenges for best execution in general and automated trading in particular. The focus will be on interest rate mar-

*Quantitative Brokers LLC, and New York University Courant Institute of Mathematical Sciences

kets and in particular on interest rates futures markets, since those are the most highly developed and the most amenable to quantitative analysis.

Following a brief overview of the markets and the products that we consider, the specific features on which we concentrate are the following:

Information events Interest rates markets are strongly affected by events such as economic information releases and government auctions. In contrast to earnings releases in the equities markets, these events generally happen in the middle of the trading day and one must have a strategy for trading through them.

Cointegration Interest rates products generally differ only in their position on the yield curve. Thus they move together to a much greater degree than any collection of equities. To achieve efficient execution in a single product, one must monitor some subset of the entire universe of products.

Pro rata matching Because futures products are commonly traded on a single exchange (in contrast to the fragmentation in the equities markets), the microstructural rules of trading can be much more complex. One example is pro rata matching, in which an incoming market order is matched against all resting limit orders in proportion to their size. This is in contrast to the more common time-priority matching algorithm. This change in match algorithm has dramatic effects on the dynamics of the order book and optimal submission strategies.

1 Fixed Income Products

The fixed income universe is large and varied, from corporate bonds, to municipal debt, mortgage-backed products, and soverign debt instruments, and including various combinations such as swaps, etc. Some of these products are traded only by dealers, and for some of them there is not even a central record of transactions.

We shall focus on the subset of fixed income products that are denoted "interest rate products," that is, products for which default risk is negligible and market risk only comes from changes in the underlying interest rate. Such products are usually, though not always, soverign debt of countries which have the ability to print the currency in which their debts are denominated. However, the STIR products that we discuss below also fall into this category since they are defined in terms of specific rates rather than issuers.

In particular we shall focus on interest rate *futures* since these are centrally cleared, and traded on organized exchanges for which detailed market information is easily available. Participants in all aspects of fixed income trading use interest rate futures to hedge their rates exposure, letting them concentrate on the more idiosyncratic aspects of their preferred products. Interest rate futures are therefore a natural point of departure. As this field develops we hope to be able to extend execution analysis to a broader range of fixed income products.

Short Term Interest Rates STIR products are instruments of very short duration. The largest product in this category is the Eurodollar future. Introduced by the Chicago Mercantile Exchange (CME) in 1981, the Eurodollar was the first cash-settled futures product and is now one of the most heavily traded futures contracts in the world. Each contract represents a forward bet on the LIBOR rate as of the date of expiration; the contract price is defined as $100 -$ LIBOR. The deliverable amount is 3 months interest on a notional amount of $1 million; thus each basis point change in LIBOR represents a mark-to-market cash payment of $25 per contract. The minimum price increment for a CME Eurodollar (except for certain short-dated maturities) is one-half of a basis point, representing a cash value of $12.50 (compared with clearing and execution costs of a dollar or less), and the bid-ask spread is almost always equal to this minimum value. Not crossing the spread becomes one of the most important aspects of trading them.

These products are thus “large-tick” in the sense of Dayri and Rosenbaum (2012), meaning among other aspects that one-tick price moves are often followed by reversals and special techniques are necessary to estimate high-frequency volatility (Large, 2011).

Eurodollar futures are traded with quarterly maturities out to ten years (plus some thinly traded “monthly” contracts that we neglect), of which at least 10–15 are active. This is in contrast to almost all other futures products, for which only the contract closest to expiration, the “front month” is active except during the “roll.” Eurodollar futures are thus inherently multidimensional.

Eurodollars trade using pro rata matching, which we discuss in Section 4. The CME interest rates electronic markets are open 23 hours per day, with consequences that we discuss in Section 2.

Euribor futures and Short Sterling, both primarily traded on the London International Financial Futures Exchange (LIFFE), are similar products in which the underlying rate is respectively a European interbank rate and a UK rate.

Treasury futures The other large category of interest products are more traditional futures contracts, in which the deliverable is a government debt security. For example, the CME Treasury futures complex covers products with underlying maturities from 2 years to 25 years and more. Their prices track very closely those of the underlying products, and can thus to some extent be used as proxies. These contracts also are "large-tick" because of the exchange-specified minimum price increments. They exhibit strong coupling between each other, although for each, only one maturity is active at one time, execept around the roll. On the CME these products, like Eurodollars, trade electronically 23 hours per day.

In Europe, the analogous products are the Euro-Bund and related contracts (Bobl, Schatz, and Buxl), traded on Eurex, which represent European government bonds of varying duration. The UK analog is the Long Gilt contract (Short and Medium Gilt are very thinly traded).

2 Information events

Interest rate markets are strongly affected by information releases and economic events that happen during the day, such as, for example, US Treasury auctions, announcements by the US Federal Open Market Committee (FOMC), and releases such as the Change in Nonfarm Payrolls number from the US Bureau of Labor Statistics (BLS) on the first Friday of every month. These events are the analogs of earnings announcements for equities, but whereas earnings announcements are usually scheduled outside of trading hours, these events generally happen during the trading day.

Trading through information events has been specific to rates markets (and energy and foreign exchange to a lesser extent) but may be being introduced in other markets. When the ICE exchange extended its hours for electronic trading of grain futures in July 2012, CME was obliged to do the same, and traders complained that "Trading now will be open during the release of most of the USDA's supply and demand reports, which will increase volatility and decrease the ability of traders to make informed decisions" (Dreibus and Wilson, 2012).

Figure 1 shows an example of market reaction to an information event: the 10-year front-month (March 2013) Treasury futures contract trading through the Change in Nonfarm Payrolls information event on December 7, 2012. (This event is one component of the Employment Situation cluster of simultaneous releases by the BLS, but it is by far the most important and most people know the cluster by this name.) This event is the most significant of all information releases, since it is the best single indicator of the health of the economy and the likely future direction of interest rates.

It is clear that this event cannot be neglected in the design of an effective trading algorithm. For any "significant" information event, the algorithm must take defensive actions before the event happens, such as removing limit orders from the book and aligning itself with the forecast schedule if the desired order type is based on a schedule. Also, all such events must be included into forecast curves for intraday volume and volatility.

Figure 2 shows another example: the 10-year Treasury futures contract trading through a 30-year bond auction on December 13. The scheduled time of the auction is 1 PM New York time, but that denotes the time at which the Treasury stops accepting bids. The market response comes approximately one-and-one half minutes later, when the auction results are released.

The auction is not the only event on that day: Table 1 shows 8 events that occur on Thursday, December 13, 2012, including the 30-year bond auction. This table shows only US, UK, EC, and German events; if all of Western Europe and Canada are included then there are 32 events on this day.

In order to trade through and around these events, we need to obtain quantitative answers to several questions:

1. Which events are "significant"? That is, for which events should we take potentially costly action such as withdrawing orders from the book. The event database on Bloomberg shows 482 distinct events within 2012, including US, Canada, and Western Europe, and including government auctions and information releases. Only a small fraction of these are significant.

2. When do events happen? Is the time as given by Bloomberg an accurate indication of the actual time as reflected in price action? For example, for an auction, not only is the market response a minute or two after the scheduled time, but also the uncertainty in this time is several seconds. It would be embarrassing to pull limit orders from the book several minutes before or after the actual price move.

3. Do US events affect European products and vice versa?

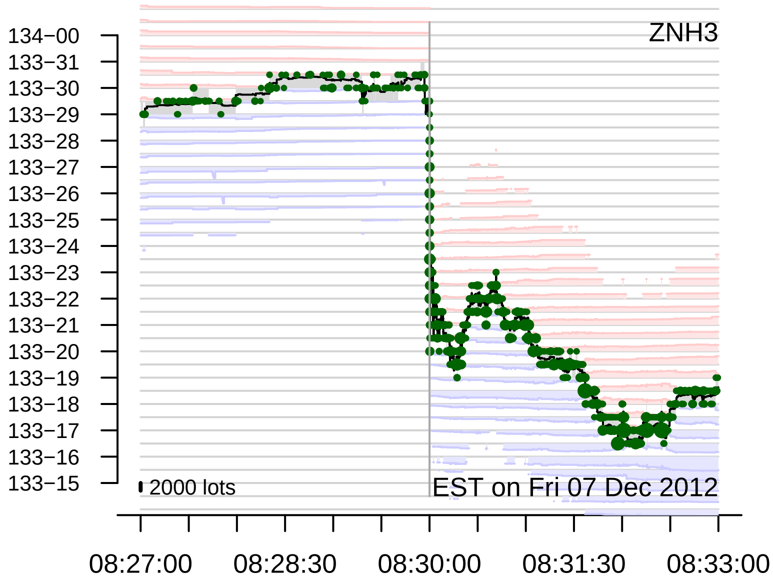

Figure 1: The 10-year US Treasury futures contract trading through the Change in Nonfarm Payrolls information release at 8:30 AM New York time on Friday, December 7, 2012. Vertical axis is price in dollars and 32nds; the minimum price increment for this contract is 1/2 of a 32nd. Gray region is bid-ask spread; green dots are trades, and blue and red shaded region show book depth. Before the event, liquidity thins out and price moves little. When the number is released, the price jumps and activity resumes.

NY Time	Region	Event
06:00	UK	CBI Trends Total Orders
06:00	UK	CBI Trends Selling Prices
08:30	US	Jobless Claims
08:30	US	PPI
08:30	US	Retail Sales ‘Control Group’
09:45	US	Bloomberg Consumer Comfort
10:00	US	Business Inventories
13:00	US	30-Year Bond Auction

Table 1: Events on Thursday, December 13, 2012, from Bloomberg database. Not all of these are significant.

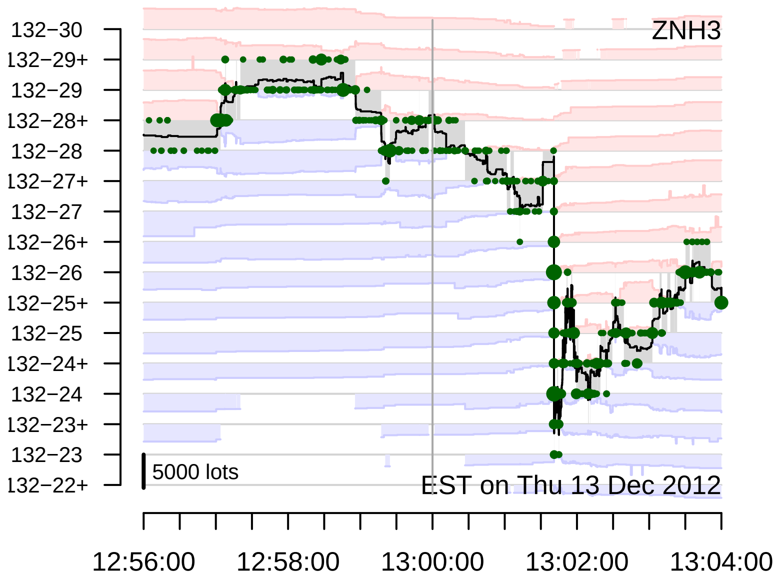

Figure 2: The 10-year US Treasury futures contract trading through a 30-year bond auction, scheduled at 1:00 PM New York time on Thursday, December 13, 2012. The price jump happens about 95 seconds later than the auction time.

2.1 Event microscope

At Quantitative Brokers, we have designed an "event microscope" to look in detail at price jumps around events. A more detailed description of how this work is described by Almgren (2012) but here is a brief summary: Compute exponential moving averages of the midpoint price time series, both from the left (backwards-looking) and from the right (forward-looking), with a variety of different time constants. The difference between the left-moving average and the right-moving average has a peak at the time of the event: the location and magnitude of this peak let us precisely locate the timing and significance of the event response.

Table 2 shows results for significant events for the US 10-year Treasury futures (ZN) and its European equivalent, the Euro-Bund (FGBL), for calendar year 2012. To generate these results, we do the analysis shown in Figure 3 for each instance of each different event type as given by Bloomberg (column n_e is the number of instances). We identify significant instances for which we detect a jump of at least two times the minimum price increment (two "ticks"), since

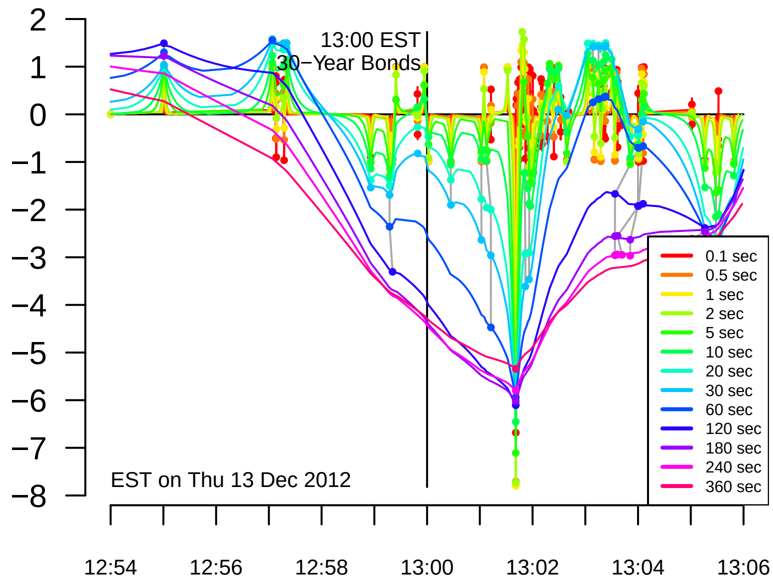

Figure 3: "Event microscope" applied to the Treasury auction shown in Figure 2. The difference of moving averages identifies a price move of -8 "ticks" (minimum price increment) at a time which is 95 seconds after the event time.

an ordinary price change will appear as a step of one time the minimum price increment. We then calculate the median absolute price change across all instances, and define "significant" events to be those for which the median price change is at least four ticks.

Among significant instances, we calculate the median time offset in seconds relative to the scheduled time from Bloomberg (column Δt). We also measure the uncertainty in this time by reporting the interquartile range of jump times (rng Δt).

We reach the following conclusions about event effects for interest rates futures:

- Change in Nonfarm Payrolls is in a class by itself in terms of significance. It consistently causes price jumps of more than ten ticks. The time is always quite close to the scheduled time, within a fraction of a second.

- FOMC Rate Decision is highly significant, at least for US rates products. The time offset is a minute or more, and the uncertainty in time is several minutes.

Event name and region	n_e	n	$\|\Delta p\|$	Δt	rng Δt
ZN (US 10-year)					
Change in Nonfarm Payrolls US	12	12	15.5	0.2	0.4
FOMC Rate Decision US	8	8	6.5	50.8	98.1
Construction & Manufacturing US	12	8	5.2	0.5	0.4
10-Year Notes US	12	10	5.0	98.1	3.5
30-Year Bonds US	12	9	4.6	99.0	3.0
ADP Employment Change US	12	8	4.6	0.0	0.1
5-Year Notes US	12	6	4.1	98.0	7.3
FGBL (Euro-Bund)					
Change in Nonfarm Payrolls US	11	11	18.2	0.3	0.6
ADP Employment Change US	12	11	5.1	-0.0	0.1
PMI Manufacturing GE	24	14	5.0	-119.5	1.1
Consumer Confidence US	12	11	4.3	1.7	60.9
Consumer Confidence Indicator DE	12	7	4.1	-118.4	181.7

Table 2: Event effects on US 10-year Treasury futures (ZN), and its European equivalent the Euro-Bund (FGBL), for Jan–Dec 2012. n_e is the number of times the event occured, for example, 12 for a monthly event. n is the number of instances that were significant for that product. $|\Delta p|$ is the median absolute price change, measured in units of the minimum price increment; this value must be at least 4 to be significant. Δt is the median time of the event, as an offset in seconds relative to the scheduled time. rng Δt is the Q1–Q3 interquartile range of the time offset.

- Treasury auctions are significant for US products. The time offset is generally 95–100 sec, with an uncertainty of several seconds.

- A miscellaneous collection of other information events are significant for US and European products. These events generally happen near the scheduled time, though occasionally events are released with specific offsets. For example, the Purchasing Managers Index (PMI) from Markit, for which the German Manufacturing information release is shown in Table 2, is released on Reuters two minutes earlier than the "standard" release, and it is this pre-release that moves the market. Similarly, the Chicago Purchasing Managers report (not shown) from the Institute for Supply Management is released to subscribers three minutes before the public release, and this offset is clearly visible in the data.

- Almost never are non-US events significant for US rates prod-

Event name and region	n_e	n	$\|\Delta p\|$	Δt	rng Δt
ZB (US Bond)					
Change in Nonfarm Payrolls US	12	12	17.3	0.3	0.5
FOMC Rate Decision US	8	7	7.9	74.9	144.3
30-Year Bonds US	12	11	5.9	99.0	3.2
Construction & Manufacturing US	12	8	5.9	0.5	0.9
ADP Employment Change US	12	8	5.7	0.0	0.1
10-Year Notes US	12	10	4.4	99.0	5.3
UB (US long-term)					
Change in Nonfarm Payrolls US	12	12	25.6	0.5	0.6
FOMC Rate Decision US	8	8	10.3	51.4	132.4
30-Year Bonds US	12	12	9.9	99.0	5.0
10-Year Notes US	12	11	7.0	99.4	7.0
ADP Employment Change US	12	11	4.6	0.1	7.4
Goods US	12	12	4.4	0.9	5.0
Construction & Manufacturing US	12	11	4.3	0.5	0.7
Retail Sales 'Control Group' US	12	12	4.2	0.7	2.5
Consumer Confidence US	12	9	4.1	0.0	0.1
FGBX (Euro-Buxl)					
Change in Nonfarm Payrolls US	11	11	13.3	0.7	1.0
FOMC Rate Decision US	8	7	6.8	89.8	183.2
30-Year Bonds US	12	11	5.3	96.6	70.0
ADP Employment Change US	12	8	4.3	0.1	9.7
Construction & Manufacturing US	11	10	4.2	2.1	112.0
Leading Indicators US	12	9	4.2	0.5	133.4
Minutes of FOMC Meeting US	8	7	4.1	8.3	62.6

Table 3: Event effects on long-term rates futures: US 30-year Treasury futures (ZB), the CME Ultra contract (UB), and the Euro-Buxl (FGBX), for Jan–Dec 2012. Columns as in Table 2.

ucts. For European products, US events are the most important, and only a few European events rise to significance (the exact combination depends on the time period and contract). This is consistent with results found by Andersson et al. (2009) and Cailloux (2007).

3 Cointegration

Cointegration is a widely studied and sought-after property of financial time series; see Alexander (2001) for a broad and detailed discussion. Although strong relationships between different price series are rare in, for example, prices of different stocks, they are absolutely ubiquitous in interest rate products, since the only differences between different products concern the duration of the product and possibly the national origin. An understanding of these interrelationships is essential to obtaining effective execution. An order for even a single asset must be understood within a highly multidimensional market context.

Figure 4 shows an example of intraday price motion. The contracts shown are the 4 primary Treasury futures contracts traded on CME (the 2-year contract is not shown). Since these contracts represent US interest rates at different durations ranging from 5 years to 30 years, they move very closely together. A cointegration model can help us identify short term mispricings. That is, these contracts establish a relationship with each other. When that relationship is disturbed it is likely to reestablish itself, and that provides some amount of short term price predictivity.

The most flexible approach to understanding cointegration is based on the principal components construction of Shintani (2001) and Chigira (2008). In constrast to the traditional approach of Johansen (1991), it does not require estimation of a discrete-time VAR model; it is extremely flexible and robust for real-time continuous market data. The construction is illustrated in Figures 5 through 7.

In Figure 5 we have extracted two of the price series shown in Figure 4, in order to display them against each other on the page. In reality, we would do the analysis on the full n-dimensional price series. The relationship seen in Figure 5 is here reflected by the alignment of the price dots along a diagonal axis. The axes show the actual price values in dollars and 32nds per contract. The axes have been scaled independently to accomodate each contract, but it can be seen that the 10-year contract ZNH3 moves approximately 10/32 through the course of the day (133-10 to 133-20), while the 30-year contract ZBH3 moves approximately 28/32 (148-28 to 149-24). Of course the longer duration contract has higher volatility since its price is much more sensitively affected by changes in underlying yield (this sensitivity is the definition of "duration" for an interest rate contract), and this difference in volatility must be properly handled in the analysis.

The round dot in the middle of Figure 5 denotes the simple mean of all the data points. The two lines denote the first and second principal components of the correlation matrix of the points, taken about this mean point. The principal components are computed using the components normalized by their variance, and are orthogonal in that scaled coordinate system. (In fact, with only two variables, each one scaled, the principal vectors are $(1, 1)$ and $(1, -1)$.) Since the plot axes are also approximately scaled by standard devi-

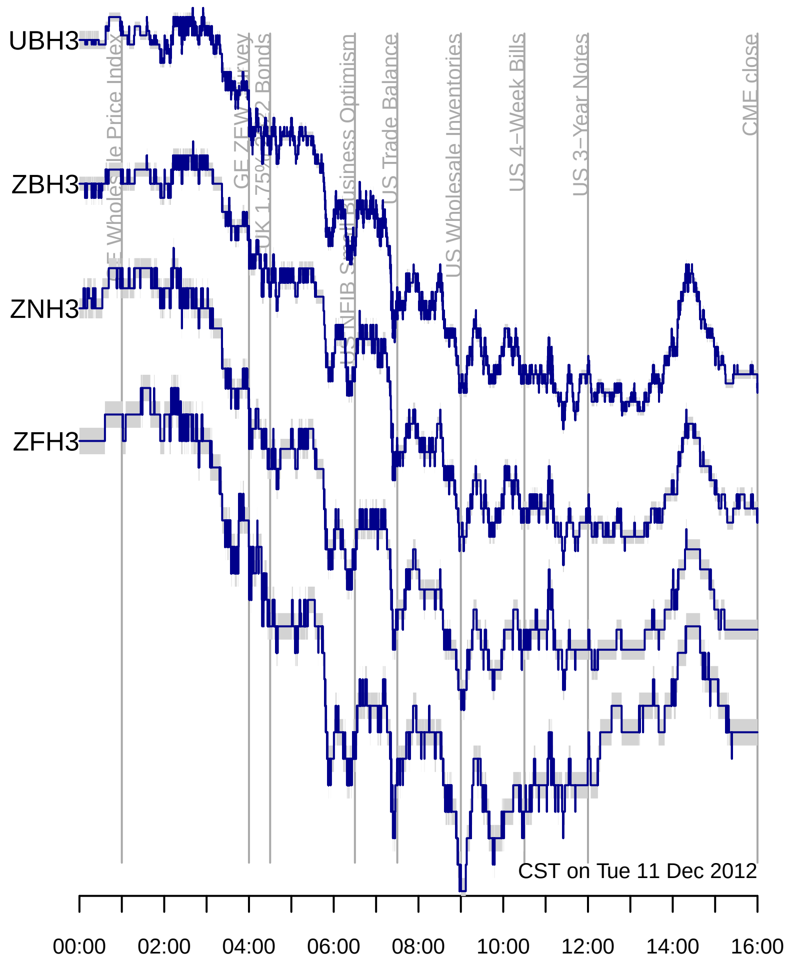

Figure 4: From bottom to top, the CME 5-year Treasury futures (ZF), 10-year (ZN), 30-year (ZB), and Ultra (UB), expiring in March 2013 (H3), from midnight to market close at 4 PM Chicago time, on Dec. 11, 2012. Price scale is arbitrary, to show relationships. Gray band is bid-ask spread; line is bid-ask midpoint. These contracts move very closely together; trading in ignorance of the relationships between them would give poor performance.

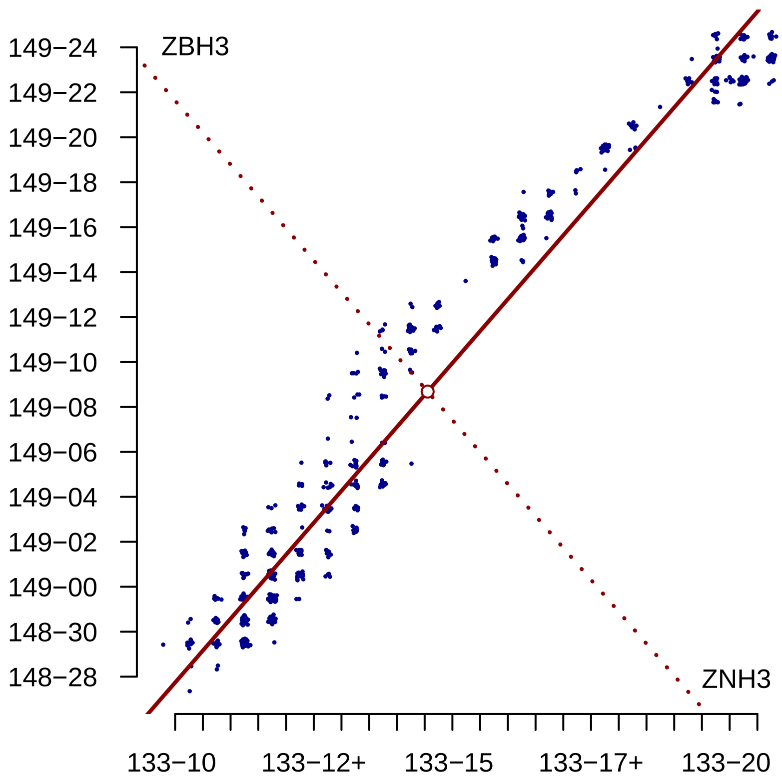

Figure 5: Two of the four contracts in Figure 4 plotted against each other: the 10-year (ZNH3, horizontal axis) and 30-year Treasury futures (ZBH3, vertical axis). Each dot represents a 1-minute sample point. Since the prices move on a grid, a small amount of random noise has been added, so that larger point clusters show larger numbers of data samples. The lines show the axes of a singular value decomposition applied to the correlation matrix of the entire day's data (this forward-looking construction is used only for an example; in practice a rolling volume-weighted average would be used).

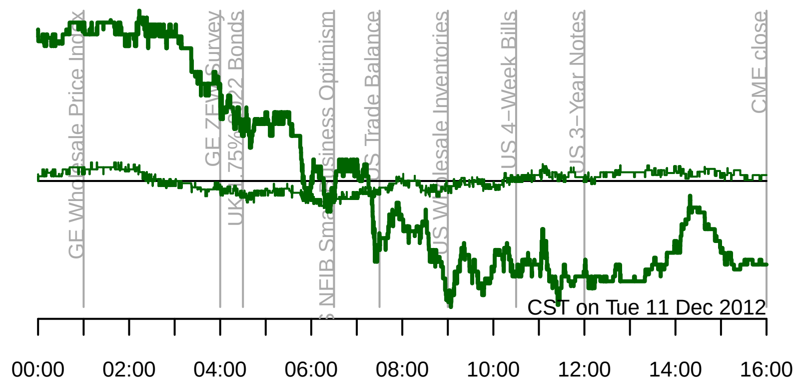

Figure 6: Projection of the price series in Figure 5 onto the principal axes of the correlation matrix. Vertical scale is arbitrary but is identical for the two components. The thick line is the projection onto the primary axis (solid line in Figure 5), reflecting the overall market motion (compare Figure 4) and is largely unpredictable. The thin line is the projection onto the secondary axis (dotted line in Figure 5) which shows mean reversion and is useful for prediction.

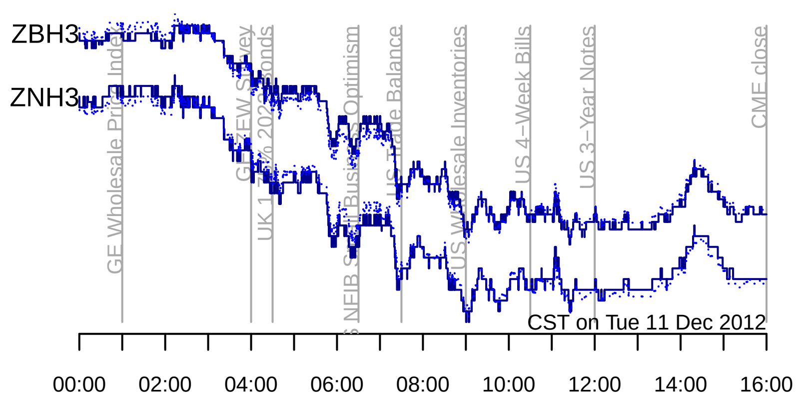

Figure 7: Short term price predictor, using the projections shown in Figures 5 and 6. Lines are the raw price series as in Figure 4. Dots are the forecast assuming that the secondary component in Figure 6 is set to zero, which is equivalent to projecting onto the solid line in Figure 5. This says that when the two contracts are relatively over- or under-valued relative to each other and to their historical relationship, that the contracts will return toward equilibrium.

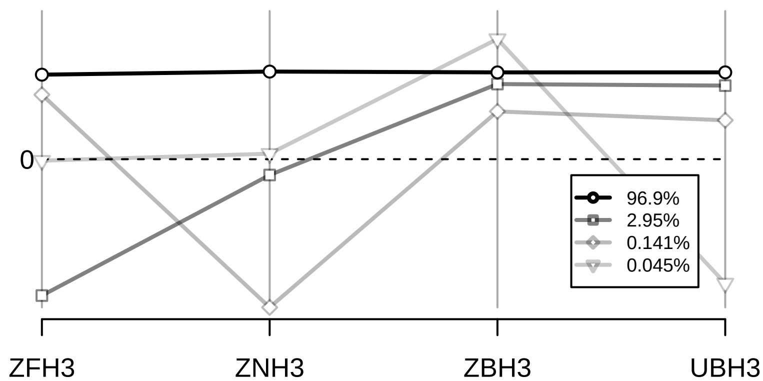

Figure 8: Principal components for the 4 price series (normalized by volatilities) in Figure 4. Inset shows contributions to variance. The first component, carrying 97% of total variance, is a constant component corresponding to overall price shifts. The second component, carrying nearly all of the remaining 3%, is roughly a tilt in the yield curve. The third and fourth components are negligible. Note that not until the fourth component can we distinguish the 30-year bond (ZB) from the Ultra contract (UB).

ation, the axes are nearly orthogonal in the plot.

In this simple example, we perform the analysis using the entire day's data. In practice this would be impossible to compute since it looks forward in time. In a real production application, both the mean and the correlation matrix would be computed as a rolling exponential average. In addition, the averages would be not assign each minute equal weight, but would use some form of weighting by trade volume.

Figure 6 shows the price data from Figure 5, projected along the two principal components. The difference in the components is clear. The projection along the primary component (light line) is essentially a reflection of the overall market movements seen in Figure 4. In this two-dimensional example, it is nothing other than the average of the two prices, appropriately scaled. In contrast, the projection along the secondary component—in this two-dimensional example, the difference of scaled prices–appears to fluctuate around zero. That is, deviations of this component away from zero predict a move back toward zero, and this information is extremely useful for short-term trading on time scales of minutes and hours. Of course, this simple graph does not constitute a rigorous test for cointegration behavior, especially given the forward-looking construction, but it indicates the nature of the price dynamics.

Figure 7 shows the two price series in real time (lines), along with the price predictor derived from the cointegration model (dots). This price predictor is obtained by setting the secondary component to zero, in effect, projecting onto the solid line in Figure 5. That is, it identifies the historical relationship between the two products on an intraday time scale, and supposes that when they deviate from this relationship, future prices will evolve so as to restore the relationship. A systematic test of the accuracy of the cointegration prediction shows that it is far less than perfectly accurate, but still effective enough to add value to real-time trading.

Figure 8 shows the principal components for the full set of 4 price series shown in Figure 4. This corresponds to what would be obtained by a traditional analysis of yield curve dynamics, but here on an intraday time scale. The first component represents the overall market motion, while the other components represent shifts relative to equal changes. The cointegration forecast would project all components after the first to zero.

A similar equilibrium model could have been derived by considering the underlying nature of the products traded, and analysing their sensitivity to changes in interest rates. The advantage of this formulation is that it is extremely straightforward, requiring no fundamental understanding of the products.

One limitation of the cointegration formulation is that it is completely symmetric between products, and has no intrinsic way to capture whether certain components drive others by moving earlier. For example, it is a common belief that futures prices move more quickly than the underlying cash instruments. There is also some empirical evidence that interest rates futures of longer durations lead futures of shorter duration. This distinction can be extracted by careful analysis of the predictive power of the signal, which will be higher for the lagging products than for the leaders.

In actual practice, the models used for price prediction would be more sophisticated than this. The Eurodollar complex is even more tightly coupled than the Treasuries illustrated here, and modeling their interrelationships is essential. It is debatable whether it is better to build a single model encompassing both Eurodollars and Treasuries, or whether to model each asset class separately; the decision must be based on a systematic analysis of the predictive power of the regression signal. Similarly, other markets such as Euribor, Short Sterling, and the European Bund-Bobl-Schatz complex may be modeled independently or together. The choice of how to group the wide range of different products is guided by a mixture of market insight and quantitative analysis.

The overall situation is not quite as simple as we have made it appear here, but cointegration is definitely a feature of interest rate markets that cannot be ignored.

4 Pro rata matching

Match algorithm refers to the process by which an exchange matches limit orders resting in the order book against incoming market orders. Typically, the total quantity of limit bid orders at the best bid price, say, is larger than the size of an incoming market sell order, and therefore some allocation must be made of the market order among the resting orders. The market order will be completely filled, but not all the limit orders will be. Some prioritization must be imposed among the limit orders.

The most obvious matching algorithm, and the one that is used by the overwhelming majority of markets, is *time priority*. Resting orders are maintained in a list in the order in which they were entered. The market order is matched against the earliest order; when that is filled the remaining quantity is matched against the next earliest order, and so on. This algorithm is simple and efficient. In such a market, the trader's main concern is to keep track of his position in the queue, in order to have an estimate of when his order will be filled.

Interest rates futures markets, largely alone among all markets, often use some variant of *pro rata* matching. In pro rata matching, the incoming market order is allocated among the resting limit orders in proportion to the *size* of the limit order, ignoring (in a first approximation) the time sequence in which the orders were entered. That is, a large limit order will receive a large allocation, even if smaller orders were entered much earlier. Field and Large (2008) have surveyed the use of pro rata matching in futures markets, identifying its predominance in rates markets, and provide a simple model for the oversizing that we discuss below.

Pro rata matching is typically used for short duration products, most notably the CME Eurodollar complex. The 2-year Treasury futures contract uses a mixed time/pro rata match algorithm. The Treasury calendar spread contracts use pro rata matching. On LIFFE, the short term Euribor and Short Sterling contracts use a "time pro rata," in which the allocation is weighted by preceding volume as well as by size of the individual order.

All these products have low volatility compared to the exchange-imposed minimum price increment, and as a consequence the bid-

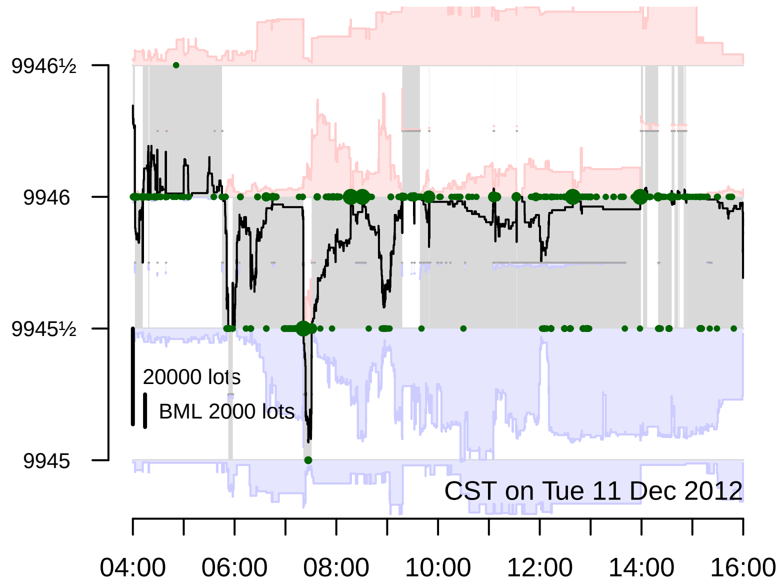

Figure 9: The March 2015 Eurodollar on December 11, 2012. Gray shaded region is bid-ask spread. Green dots are trades. Blue and red shaded regions are quote size on the bid and on the offer respectively (as well as some "implied" bid volume at the midpoint). The black jagged line is the "microprice," a midpoint weighted by bid and ask sizes. The bid and offer prices move only rarely, although substantial trade activity occurs.

ask spread is nearly always equal to its minimum allowable value. That is, as noted above, they are "large-tick" in the sense of Dayri and Rosenbaum (2012). This means that being executed on a limit order is much more valuable than crossing the spread with a market order. Also, the bid and offer prices commonly stay at constant values for substantial lengths of time.

Figure 9 shows an example. The March 2015 contract is the third most heavily traded Eurodollar contract on December 11, 2012 and substantial trade activity is visible (the two heaviest are March and June 2013, three and six months from expiration). This contract represents the LIBOR rate more than two years in the future. One would expect market beliefs about this quantity to change throughout the day, at least by several multiples of the minimum price increment, one-half of one basis point. Nonetheless, the bid and ask prices do not move at all for extended periods, for example for a period of 4 hours between 10 AM and 2 PM Chicago time (though two very short flickers are visible within this interval).

When the quote prices do not move, a time priority match algorithm would excessively weight early arrivals. If a market maker were able to capture the head of the queue with one large order on the bid, and another on the ask, then every subsequent market participant would be obliged to trade with him. Pro rata matching gives later entrants the possibility to execute.

In these markets, the dynamics of the order book is extremely volatile. Since there is no penalty for losing queue position, there is no disincentive to cancel and resubmit limit orders. This can be seen in Figure 9, in the substantial changes in quote volume on the bid and on the ask. The LIFFE time pro rata algorithm is an attempt to partially dampen these wild swings.

An additional consequence of pro rata matching is the "arms race" to oversize orders (the term was introduced by Field and Large (2008)). Since allocation is determined by order size, and since incoming market order volume is typically too small to satisfy the traders who are hoping for passive execution, each limit order participant has an incentive to post a much larger quantity than he actually desires to execute. Typically, in these markets, the average volume on the bid and the ask is several hundred times a typical trade size. In Figure 9, the quantity on the inside quotes is in the range of 10–20,000 lots, whereas the average market order size for this contract is around 20 lots.

The only limitation on oversizing is the risk that a large market order will fill for much more quantity than was desired, but such large orders are rare (see Arora (2011) for an analysis of their statistics and effect on the market). Balancing the risk of overfilling if one does oversize, against the certainty of underfilling if one does not oversize, is the central concern for a trader in a pro rata market.

5 Summary

Traders in fixed income and interest rates market have just as much need for effective execution and transaction cost management as do their counterparts in the equities markets, although the latter have received vastly more quantitative attention. Several features of interest rates futures markets in particular are substantially different from these other markets, and must be taken proper account of in order to achieve good execution results.

References

Alexander, C. (2001). *Market Models: A Guide to Financial Data Analysis*. New York: John Wiley & Sons.

Almgren, R. (2012, October). High-frequency event analysis in European interest rate futures. Technical report, Quantitative Brokers.

Andersson, M., L. Jul Overby, and S. Sebestyén (2009). Which news moves the Euro area bond market? *German Economic Review 10*(1), 1–31.

Arora, R. (2011). An analysis of large trades in Eurodollar futures. Technical report, Quantitative Brokers.

Cailloux, J. (2007). What moves the European bond market? Market responses to surprises in main economic indicators. Technical report, The Royal Bank of Scotland.

Chigira, H. (2008). A test of cointegration rank based on principal component analysis. *Appl. Econ. Letter 15*, 693–696.

Dayri, K. and M. Rosenbaum (2012). Large tick assets: Implicit spread and optimal tick size. arXiv:1207.6325 [q-fin.TR].

Dreibus, T. C. and J. Wilson (2012, July). Grain-pit traders squeezed out as CME expands to match ICE hours. *Bloomberg*.

Field, J. and J. Large (2008). Pro-rata matching and one-tick futures markets. Preprint.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica 59*(6), 1551–1580.

Large, J. (2011). Estimating quadratic variation when quoted prices change by a constant increment. *J. Econometrics 160*, 2–11.

Shintani, M. (2001). A simple cointegrating rank test without vector autoregression. *J. Econometrics 105*, 337–363.

SIFMA Asset Management Group (2008). *Best Execution Guidelines for Fixed-Income Securities*. SIFMA Asset Management Group. White paper.