

Automated Trading in Futures Markets*

Richard Haynes[†] John S. Roberts[‡]

March 13, 2015

Introduction

Over the last decade, U.S. markets have seen an increased use of technology in areas such as trading, order generation, and order routing. Electronic order books, standardized message protocols, and widely distributed data feeds have reduced the cost of building and implementing automated systems. These systems, which often use algorithms to generate trading or routing decisions, are commonly classified under the general market term “automated trading.” Recent studies on automated trading in domestic markets have found that often over half of the trades on securities and futures exchanges make some use of algorithms, a subset of automated trading, to match trades, oversee certain order types (e.g., stop orders) and monitor general market risk. Academic research which looks at the rise of automated trading has considered its effects on market liquidity, efficiency, volatility and the costs or benefits to retail and institutional investors, with varied results.¹ This study will focus on the activity of automated traders in markets regulated by the Commodity Futures Trading Commission (CFTC), summarizing automation across asset classes over the last few years, and detailing how this differs across futures products and product classes.

One commonly discussed subset of automated trading is high-frequency trading (often abbreviated as HFT). It has proven difficult to effectively define a dividing line between those traders who should be included in the “high-frequency” set and those who should not, especially when trying to answer the question across a wide set of products and time periods with differing liquidity characteristics. Often, metrics like the number of trades in a day, the ratio of orders to trades, or the ratio of held positions to participant volume have been used as proxies

*The research presented in this paper was co-authored by Richard Haynes and John Roberts, who are both CFTC employees, in their official capacities with the CFTC. The Office of the Chief Economist and CFTC economists produce original research on a broad range of topics relevant to the CFTC’s mandate to regulate commodity futures markets, commodity options markets, and the expanded mandate to regulate the swaps markets pursuant to the Dodd-Frank Wall Street Reform and Consumer Protection Act. These papers are often presented at conferences and many of these papers are later published by peer-review and other scholarly outlets. The analyses and conclusions expressed in this paper are those of the authors and do not reflect the views of other members of the Office of Chief Economist, other Commission staff, or the Commission itself.

[†]Senior Research Analyst and Corresponding Author, rhaynes@cftc.gov

[‡]Research Analyst, jroberts@cftc.gov.

¹For a general overview of automated trading literature in the context of equity markets, see the SEC’s [Equity Market Structure Literature Review: Part II](#).

for HFT. Because of the difficulty in defining appropriate thresholds between “HFT” and “non-HFT” participants using these or other metrics, the analysis below will focus instead on the full set of automated traders as identified in the data. In the later sections, we will provide a small set of breakdowns within the automated group, separating out the subset of more active participants, as defined by metrics explained below.

In this overview of automated trading in futures markets, we will address aspects of the following general questions:

- How prevalent is automated trading across different product classes, liquidity profiles and time periods?
- What role does automated trading play in futures liquidity provision?
- What is the relationship between automated trading and the speed of trading decisions?
- What is the relationship between automated trading and the speed of position turnover?

We find a large presence of automated trading across a wide set of futures products, often well over half of all trading, and usually proportional to the level of activity in the contract. We also find that, in general, markets have gotten faster over the last few years, with automated traders often trading, or quoting, at a faster level than those who trade manually. However, these findings are not universal. Many contracts, especially physical commodities, have significant manual participation, especially during futures roll periods. We also find that manual trading can be quite fast, and act as a large portion of liquidity provision in key contracts. In brief, futures markets appear to provide an environment conducive to both manual and automated participants, where each often plays both common and distinctive roles.

To answer the questions outlined above, we use trade level data across all of the Chicago Mercantile Exchange (CME) futures products over a time period of approximately two years. The data set includes 1.5 billion trades, across 805 futures products and close to 362 thousand individual accounts.² The data allows for breakdowns by individual trading account or firm, by product or product class and includes millisecond level timestamps for determining event speed within these categories.

Data

This study is based on high frequency regulatory data collected by the CFTC from the CME Group exchanges.³ The CME Group is the largest U.S. futures exchange operator, and offers trading across commodity, foreign exchange (FX), equity and interest rate asset classes. The data set represents the complete record of all trades executed on the exchange, including all those entered into the GLOBEX electronic matching engine. For each trade, reported information includes numerous variables about the product and the participant, and trade specific details, such as

²Although there are 804 different products in the sample period, activity is concentrated in a small subset of products. For example, the top 50 products by number of trades represent 99.6 percent of total sample trade count. The top 50 products by volume represent 98.9 percent of total sample volume. Likewise for unique accounts, the top 2,000 accounts by trading activity represent 75 percent of total sample volume.

³For further academic work utilizing this information, see the CFTC’s [Economic Analysis Research Page](#).

the trade aggressor, the trade quantity and price, the underlying futures product, and order entry and execution timestamps at a millisecond precision. Domestic futures markets trade on a highly consolidated basis, so the statistics below are an essentially comprehensive view of activity within the identified markets. This contrasts with equity markets, where recent statistics indicate that over 30 percent of equity volume is done off-exchange, and on-exchange trading is dispersed across more than ten platforms.

Table 1 provides a breakdown by execution venue type; the table shows that orders entered through the electronic trading platform make up a vast majority of futures trading, almost 95 percent of trading across all products. The remaining 5 percent is comprised of orders executed in the pit or those with specialized characteristics (larger trades which are executed as blocks, or trades that exchange a given product for another). The analysis in this paper focuses strictly on trades executed on the electronic venue.

Table 1: Activity by Venue and Transaction Type

Notes: This table contains percentages of total volume based on venue and transaction type. This table makes clear that manual order entry information is only available for electronic activity, which makes up the vast majority of the total sample volume at the CME. Other trading on Pit and X-Pit venues is small and we will group these into one non-electronic trading category. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Manual Order Entry Information	Venue Type	Transaction Type	Percent of Total Sample Volume
Yes	Electronic	Regular Trade	94.7
No	Pit	Regular Trade	1.1
No	X-Pit	Block Trade	2.3
No	X-Pit	EFP (Exchange for Physical)	1.1
No	X-Pit	EFR (Exchange for Risk)	0.6
No	X-Pit	EFS (Exchange for Swap)	0.1
No	X-Pit	PNT (Privately Negotiated Trade)	0.0

The variable of particular interest for this study is the manual order indicator. CME Group exchanges require that each order entered into GLOBEX includes a binary value that identifies whether the order was entered manually (MAN) or automatically (ATS).⁴ Orders are identified as automated if they are “generated and/or routed without human intervention.” This set of orders is not restricted to those that are directly generated by algorithms, or those associated with HFT firms. Instead, it includes a much broader category, such as those which are generated manually but make use of automated spreading functionality, or even those where manual traders use the order submission management of third-party trading systems. The manual order indicator was first observed in the data on November 12, 2012. Therefore, this study will incorporate all electronic futures data from CME group exchanges from November 12, 2012 through October 31, 2014.

⁴This information is contained in FIX Tag 1028. Further details for the manual/automated trading indicator can be found on the [CME's website](#). The indicator was codified in each Exchange's rulebook on October 8, 2012.

Prevalence of Automated Trading

Trading by, or with the help of, automated systems can reduce the delay between receiving and responding to information relevant to a given market. This increase in speed can often be beneficial in markets where prices adjust quickly or where the frequency of trading is relatively high. Because of this, there is usually a positive correlation between the level of activity in a financial product and the percentage of that activity that makes use of automated systems. In Table 2, we provide a breakdown of futures trading volume across all CME Group products by the manual order indicator; this breakdown summarizes the prevalence of automated trading at the very highest aggregation level. Later tables concentrate on a small set of highly liquid products, chosen from each main product group.

Table 2: Manual Vs. Automated Trading By Product Group

Notes: This table summarizes volume by type as a percentage of total volume within the product group and across the full sample (in parentheses). Each individual transaction is linked to its associated product, which is then aggregated within the associated product group. The Low Volume Group represents a set of products with the least activity — most of which is non-electronic — that are not linked to a product group name. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Product Group Name	Products	Non-Electronic (%)		ATS (%)		Manual (%)	
Low Volume Group	376	98.1	(0.0)	0.1	(0.0)	0.9	(0.0)
Agriculture	36	6.4	(0.5)	38.1	(3.0)	55.6	(4.4)
Energy	285	16.8	(2.2)	46.9	(6.0)	36.4	(4.7)
Equities	25	0.9	(0.2)	66.6	(14.3)	32.5	(7.0)
FX	51	3.4	(0.2)	79.9	(5.8)	16.8	(1.2)
Interest Rate	15	4.2	(2.0)	62.3	(29.7)	33.4	(15.9)
Metals	17	5.6	(0.2)	46.5	(1.4)	48.0	(1.4)

In general, though not universally, the higher the volume in a product group, the greater the presence of automated participants. As seen in Table 2, the three product groups with highest automated participation are, in decreasing order, FX (80 percent), equities (67 percent) and interest rates (62 percent). Though futures volume for FX is lower than other product categories, tools for automated trading have been in use in this asset class since at least the early 1980’s when Reuters began offering an electronic trading platform.⁵ With such a lengthy history, this product class remains one of the most mature for automated execution, widely used in both the currency futures and spot (immediate delivery) markets.

Equity futures, the asset class with the next largest automated presence, are dominated by the E-mini S&P contract, which settles against the price of the S&P 500 index. Like equity index ETFs, the E-mini provides a means of gaining general equity exposure through a single instrument; partly because of this diversification convenience, it is one of the most actively traded futures contracts (second by contract volume after the Eurodollar rates contract). Rounding out the top three automated product categories, the interest rate group includes a number of very active products — in this category are both Treasury futures, tracking US debt, and Eurodollar futures, which reflect

⁵For a more detailed history of early FX electronic trading, see Chaboud, A. and Weinberg, S. “Foreign exchange markets in the 1990s: intraday market volatility and the growth of electronic trading.”

interbank lending spreads. At the other end of the spectrum, the products with the lowest level of automated participation are the physical commodity classes, encompassing energy, metals and agricultural products like oil, copper and wheat. Physical commodity futures trading often has a higher level of both buy- and sell-side end users who consume and produce these goods, and often trade on a manual basis.

A more detailed breakdown of automated trading in futures products is provided in Table 3. This division between manual and non-manual trading further emphasizes the wide spectrum of automation use across futures products, and similarly emphasizes the depth of use in the most highly active categories like domestic equity indices, G10 currencies and US Treasury futures. The chart also shows the diversity and popularity of execution types within given product categories. As an example, all coal and electricity trades within the sample set do not have a manual/automated classifier, given that they are not executed on the electronic platform; in contrast, the proportion of equity index and short term interest rate futures trades executed off the electronic platform are each below 2.5 percent of total trading volume. The chart also includes the number of products within a given product subclass, rising from just one in the lumber and pulp category to 126 in refined products.

Table 3: Manual Vs. Automated Trading By Product Group and Subgroup

Notes: This table summarizes volume as a percentage of total within a product/sub-product pair and across the full sample (in parentheses). Each individual transaction is linked to its associated product, which is then aggregated within the associated product/sub-product category. Here we exclude the low volume group of Table 2. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Product Group and Subgroup Name	Products	Non-Electronic (%)		ATS (%)		Manual (%)	
Agriculture - Commodity Index	6	26.0	(0.0)	22.6	(0.0)	51.5	(0.0)
Agriculture - Dairy	6	4.1	(0.0)	6.8	(0.0)	89.1	(0.0)
Agriculture - Grain and Oilseed	16	5.6	(0.4)	39.0	(2.7)	55.4	(3.8)
Agriculture - Livestock	3	11.3	(0.1)	32.4	(0.3)	56.3	(0.5)
Agriculture - Lumber and Pulp	1	6.9	(0.0)	10.5	(0.0)	82.6	(0.0)
Energy - Biofuels	5	79.4	(0.0)	2.5	(0.0)	18.1	(0.0)
Energy - Coal	8	100.0	(0.1)	-	-	-	-
Energy - Crude Oil	22	5.5	(0.3)	54.3	(3.2)	40.2	(2.4)
Energy - Electricity	62	100.0	(0.5)	-	-	-	-
Energy - Natural Gas	40	21.2	(0.8)	44.2	(1.6)	34.6	(1.2)
Energy - Refined Products	126	15.0	(0.4)	46.0	(1.2)	39.0	(1.0)
Equities - International Index	4	0.4	(0.0)	69.3	(0.4)	30.3	(0.2)
Equities - Select Sector Index	9	18.8	(0.0)	59.2	(0.0)	22.0	(0.0)
Equities - US Index	12	0.9	(0.2)	66.5	(13.8)	32.6	(6.8)
FX - E Micros	8	0.0	(0.0)	66.3	(0.1)	33.7	(0.0)
FX - Emerging Markets	18	12.5	(0.1)	70.3	(0.3)	17.1	(0.1)
FX - G10	25	2.8	(0.2)	80.7	(5.4)	16.4	(1.1)
Interest Rate - Deliverable Swaps	4	16.3	(0.0)	47.8	(0.0)	35.9	(0.0)
Interest Rate - Stirs	3	2.2	(0.5)	60.3	(13.0)	37.5	(8.1)
Interest Rate - US Treasury	5	5.9	(1.5)	64.0	(16.6)	30.2	(7.8)
Metals - Base	4	3.1	(0.0)	49.2	(0.3)	47.7	(0.3)
Metals - Precious	8	6.1	(0.1)	45.9	(1.1)	48.0	(1.1)

Table 4 filters down to the individual commodity level, selecting the most active commodities from various product classes. This table also provides a more detailed breakdown on automated versus manual trading for

a given commodity by dividing all trades into those where both, one or neither of the trade counterparties is automated (ATS–ATS, ATS–MAN, and MAN–MAN, respectively). Because this chart only includes the most liquid futures products, there is a weaker relationship between contract volumes and the level of automated trading. Also, even though these are the most active commodities within their product group, the frequency of trading between automated participants can be as low as 24 percent (gold) and as high as 72 percent (British pound).

Figure 1 (found at the end) expands on the information contained in Table 4 by showing, across the time period, the summarized percentages per participant-pair category across four actively traded products. For each product, the time series highlights periods of the sample where manual trading becomes much more prevalent relative to automated trading. Usually this happens during periods when market participants are rolling into new expirations through spread trades – each primary roll period usually lasts a few days.⁶ The frequency of rolls differs across contracts – where crude oil rolls on a monthly basis, the other three products roll quarterly. The figure also re-emphasizes the different levels of automated participation across the four selected commodities. Where the ATS–ATS pairing is by far the most frequent in Euro futures, ATS–ATS remains lower than ATS–MAN in crude oil.

The chart points to the likelihood that spread trading is, on a relative basis, a more manual activity. Table 5 provides a breakdown of activity for select products across the three participant-pair categories (ATS–ATS, ATS–MAN, and MAN–MAN) and across a secondary classification which separates all futures trades into outright (abbreviated as RO for “regular outright”) and spreads (SP). With this breakdown, it is clear that the Eurodollar contract is unique relative to peers. Though it is the most active futures product, it has a low automated participation rate relative to other active products. As indicated by Table 5, this is largely because Eurodollars have a similarly exceptional level of spread trading, where manual traders are more active. Omitting these trades, automated Eurodollar activity is noticeably higher, and much closer to others with similar trading activity.

In Table 5, data omissions in the trading activity of outright matched to spreads for some commodities arise because matching between outright and spread trades is not universally available. This ability is known as “implied spread” functionality, and is active for Treasuries, Eurodollars, and many physical commodities. It is not available in equity and FX futures, so the values for outright matched to spread trades for these underlyings are blank. Even controlling for this difference, it is clear that the use of spread trading widely differs across products, with a higher frequency of spreads in physical commodities (especially energy) and Eurodollars, as noted above. Many consumers of crude oil, and other physical commodity products, hedge short-term along with long-term anticipated energy needs, which can require trading in futures contracts which do not expire for months or years. For less

⁶Spread trades are one of the main distinguishing features of futures markets, relative to a market like equities. For futures, there is no perpetual underlying which can be held indefinitely; instead, each futures contract has a pre-specified expiration date, after which the contract no longer exists. Spread trades are trades which reference the relative price of two futures contracts, usually on the same product, with different expirations. For example, a participant could trade the December 2014 – December 2015 crude oil spread, which would simultaneously buy (sell) the first futures expiration and sell (buy) the same quantity of the second. Spread trades are most commonly used during the futures “roll” period when the shortest dated futures contract is close to expiration. If a market participant wants continued exposure to the commodity, the participant will need to exit its position in the expiring contract and enter into a similar position in a longer-dated contract.

Table 4: Manual Vs. Automated Trading Most Active Products

Notes: The chart summarizes participant activity by manual/automated type for the top three products by volume in each product group. This activity is reported as percentages of total volume in the case where both sides of the trade are automated (ATS-ATS), where only one side is (ATS-MAN), where neither side is (MAN-MAN). Source: CME transaction data, November 12, 2012 – October 31, 2014.

Product Group and Subgroup	Futures Product Name	Total Volume (mm)	ATS-ATS (%)	ATS-MAN (%)	MAN-MAN (%)
Agriculture - Grain and Oilseed	Corn	128	15.9	42.8	34.8
Agriculture - Grain and Oilseed	Soybean	93	15.2	42.5	36.3
Agriculture - Grain and Oilseed	Wheat	49	18.5	44.9	33.9
Energy - Crude Oil	Crude Oil	280	30.2	45.5	19.5
Energy - Natural Gas	Natural Gas (Henry Hub)	153	30.7	45.8	19.1
Energy - Refined Products	RBOB Gasoline Physical	67	24.6	43.7	23.0
Equities - US Index	E-mini S&P 500	866	43.4	42.6	13.8
Equities - US Index	E-mini NASDAQ 100	131	58.2	34.8	6.9
Equities - US Index	E-mini Dow (\$5)	73	60.9	33.6	5.4
FX - G10	Euro	110	64.5	28.9	4.4
FX - G10	Japanese Yen	76	68.8	25.0	3.5
FX - G10	British Pound	53	71.7	21.2	4.1
Interest Rate - Stirs	Eurodollar	1,125	36.1	49.3	12.5
Interest Rate - US Treasury	10 Yr Note	643	42.6	41.1	11.9
Interest Rate - US Treasury	5 Yr Note	360	50.5	32.8	8.6
Metals - Precious	Gold	85	24.0	43.9	25.5
Metals - Base	Copper	31	27.5	43.5	26.0
Metals - Precious	Silver	27	27.9	37.0	31.3

liquid contracts which have longer-dated expirations, it is common for trades to be executed as a spread to the front month, increasing spreads as a fraction of total trading. In those contracts where spread trades are more common, automated traders are more likely to be a larger component of spread activity. In crude oil, for instance, 77 percent of all spread to spread trades have an automated participant on the buy or sell side. For the S&P E-mini, where spread trading is uncommon, about 66 percent of spread to spread trades are executed between manual traders (manual participation in spread trading peaks across the identified commodities at 86 percent in silver). More generally, low trade interest and volume, as in other contexts, is often correlated with low automated participation.

Table 5: Volume Shares By Product and Trade Type

Notes: This table provides a further breakdown for the most active products found in Table 4. These products will be the primary focus for the rest of the study. For each product, percentages of total volume over the sample period are grouped by automated or manual category and by whether the trade matches regular outright (RO) or spreads (SP). Note the percentage of volume attributed to non-electronic trading is not included in this table. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Futures Product Name	ATS-ATS			ATS-MAN			MAN-MAN		
	RO-RO	RO-SP	SP-SP	RO-RO	RO-SP	SP-SP	RO-RO	RO-SP	SP-SP
E-mini S&P 500	43.1	-	0.3	40.6	-	2.0	9.5	-	4.3
E-mini NASDAQ 100	58.0	-	0.3	33.0	-	1.8	3.9	-	3.0
Crude Oil	11.1	2.5	16.6	15.1	4.5	25.9	4.4	2.1	13.0
Natural Gas (Henry Hub)	12.3	3.6	14.7	12.0	6.0	27.9	3.4	2.4	13.3
Eurodollar	11.7	6.0	18.3	17.1	8.1	24.2	5.2	2.4	4.8
10 Yr Note	41.4	0.3	1.1	36.3	0.4	4.5	8.1	0.1	3.6
5 Yr Note	49.2	0.3	1.2	26.8	0.3	5.6	3.2	0.1	5.2
Euro	64.3	-	0.2	27.0	-	1.9	2.3	-	2.1
Japanese Yen	68.4	-	0.3	23.0	-	2.1	1.6	-	1.9
Corn	10.1	1.9	3.9	17.7	6.9	18.2	9.2	6.1	19.5
Soybean	9.1	2.6	3.5	17.9	9.0	15.6	9.8	8.1	18.4
Gold	22.8	0.5	0.7	36.6	2.3	5.0	13.4	2.2	10.0
Silver	27.7	0.1	0.1	31.0	2.4	3.5	7.0	2.2	22.1

The introduction outlined common difficulties with identifying important subsets of automated traders. Table 6 considers one set of key participants, automated traders who trade large volumes in a given contract. Here, a large trader is defined as an account which contributes at least 0.5 percent to a given day’s volume; this definition allows for an account to be classified as “large” one day and “small” on other days, depending on where it sits relative to the 0.5 percent threshold. While there does exist variation across products, automated trading for large volume traders contributes the majority of total sample volume. However, for individual products like corn and soybeans, small traders contribute the majority share of volume; these contracts have a wide set of participants who, individually, execute small quantities, but are significant in aggregate. The Table also shows perhaps the more surprising fact that high activity traders are not equivalent to automated traders. There are many cases where high volume traders are manual, and where small volume traders are automated, emphasizing the diversity of trading strategies and activities across futures markets.

Table 6: Volume Shares By Product and Trader Group

Notes: Trading accounts are classified daily as “large volume trader” if they contribute at least 0.5 percent to total daily volume across all expirations. All accounts not meeting this criteria are grouped into the “small trader” category. Percentages represent automated and manual trading for the trader group. N captures the number of unique accounts observed in each category over the full sample period — note it is possible for an account to be included in both small and large totals. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Futures Product Name	Small Volume Trader			Large Volume Trader		
	N	ATS (%)	MAN (%)	N	ATS (%)	MAN (%)
E-mini S&P 500	126,675	19.2	29.3	469	45.5	5.8
E-mini NASDAQ 100	52,793	17.4	20.0	544	58.2	4.3
Crude Oil	65,410	18.4	30.5	538	34.6	11.8
Natural Gas (Henry Hub)	35,1813	15.1	24.5	649	38.5	17.5
Eurodollar	21,223	14.5	30.3	619	46.1	6.9
10 Yr Note	47,773	20.5	27.4	716	42.7	5.1
5 Yr Note	22,399	20.0	18.7	1,126	47.0	6.4
Euro	45,064	14.9	15.1	477	64.1	3.7
Japanese Yen	29,191	14.4	12.7	485	66.9	3.3
Corn	76,290	18.7	36.0	868	18.6	20.2
Soybean	59,739	19.2	42.2	815	17.2	15.4
Gold	52,675	19.0	31.8	630	26.9	15.8
Silver	23,934	17.0	24.4	696	29.3	25.4

Speed of Trading

Trading across all financial markets has gotten dramatically faster over the last decade, due to technological innovations at exchanges, better electronic order management at firms, and more sophisticated communication networks connecting the two. As noted by the SEC in recent analysis of equity market speed, just over 50 percent of market orders are canceled within half a second, approximately the speed of human reaction.⁷ The ability to enter or adjust orders quickly is useful for participants who provide liquidity and for those who take liquidity. Market makers, who provide liquidity, need to be able to quickly cancel or modify standing quotes during periods of directional price movement. Without this ability, the market maker could be overwhelmed by losses due to ‘adverse selection,’ cases where market moves leave the market maker with a position opposing the price direction. Market-takers, who remove liquidity from the order book and often trade on the basis of new information, are aided by being the first, or one of the first, participants to bring that new information into the market, resulting in gains from the subsequent price movement.

Table 7 summarizes the speed of a set of active products using one metric associated with execution speed, the time between order entry and execution. Here, order entry time is either the time when the order was initially entered, or, when modified, the time of the most recent modification. Note that this metric only takes into account a subset of orders. Submitted orders which are canceled without either partial or full execution do not show up in this sample set. Separate analysis of futures data has shown order to execution ratios are commonly as high as 10:1 even for liquid products, so this metric can often encompass only a fraction of total order book activity.

⁷For this, and related, analysis see “[Equity Market Speed Relative to Order Placement](#)” Data Highlight 2014-02, March 19, 2014.

Within this set of executed orders, the chart only summarizes times associated to the “passive” side of a trade — the side which provided liquidity for the execution. The “aggressive” side of a trade, because it executes against an already resting order, has little to no delay between order entry and execution. Note that there are a small fraction of trades, a subset of spread trades, where both sides of the trade are labeled passive; this leads some row sums in the chart to add to greater than 50 percent.

Across the selected commodities, a significant fraction of passive execution activity happens within one second of order entry or modification, similar to speeds seen in other large electronic markets. For instance, around 40 percent of executed passive orders are executed within one second in Euro and Yen futures contracts, along with around 20 percent in Treasury and S&P futures. As with the automated trader breakdown, physical commodities, with higher manual participation, are slower according to this speed metric: over 90 percent of passive executed corn orders remained in the book longer than one second, as are 85 percent of natural gas futures.

The chart also provides a breakdown of the relative frequency of manual versus automated orders within a time bucket. In one of the fastest markets, the E-mini S&P contract, the ratio of automated to manual trades executed within less than 100ms is just under 11 to 1; in contrast, the same ratio for orders executed between 10 seconds and 1 minute is less than 2 to 1. This disparity could be because automated, relative to manual, trades are traditionally closer to the top of the order book, and so experience faster execution, or it could be because automated participants cancel or modify standing orders faster than their manual counterparts (so executions must happen within an abbreviated window). For physical commodity contracts, the ratios per time interval are more consistent across buckets: for gold, 2.3 times as many orders executed within 100ms of order placement or modification are automated rather than manual. The same ratio for the 1–10sec interval is the negligibly different 1.925 times. There is also for physical commodities a greater balance in the liquidity provision by manual and automated traders. Up to 10 seconds, 3 times as much liquidity is provided by automated rather than manual participants in the E-mini contract; in gold, this ratio is just under 2.

Figure 2 extends this analysis by charting the level of daily liquidity provision by automated participants across the period, for a select set of commodities. Trends across individual commodities and periods are diverse; generally, but not always, the speed of markets relative to liquidity provision levels has been upward (as indicated by general increases and common moves in the speed distributions towards the higher frequency segments). The figure also highlights that, like Figure 1, the speed of liquidity provision is dependent on the frequency of spread activity in a market. During a futures roll, order-to-execution times increase, with spread quotes on average sitting in the book for a longer period; because of this speed reduction, the percentage of trading within the fastest buckets drops during contract roll periods. The charts also include a few non-roll periods where market speeds had non-negligible increases — for the 10 year Treasury note, market speeds increased both in June of 2013, during concerns about monetary tapering, and the fall of 2014, where large market moves were experienced on October 15. On a commodity-specific basis, the relative size of a given speed bucket differs, with the highest frequency bucket

Table 7: Time Between Order Placement and Execution for Liquidity Providers

Notes: Table shows percentage of total two-sided volume for 13 products by ATS vs. Manual and by the time between initial order placement or modification and execution. This includes all trades that designated as liquidity providing according to the aggressor indicator (non-aggressive side of the trade). A small subset of trades for active intra-commodity spreads have both sides marked as non-aggressive, which explains why the rows might sum to a number greater than 50 percent (e.g., Eurodollar row percentages sum to 54.5 percent). We assume each non-electronic trade has a single side marked as non-aggressive. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Futures Product Name	0 – 100 ms		100 ms – 500 ms		500 ms – 1 sec		1 sec – 10 sec		10 sec – 1 min		Remaining	
	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN
E-mini S&P 500	6.4	0.6	3.1	0.5	2.1	0.5	10.0	3.7	6.6	3.7		12.7
E-mini NASDAQ 100	8.9	0.8	5.5	0.6	3.5	0.5	11.8	2.9	5.0	2.4		8.0
Crude Oil	3.9	1.3	2.2	0.6	1.4	0.4	6.2	2.7	5.0	3.5		22.8
Natural Gas (Henry Hub)	3.2	1.2	1.4	0.5	1.0	0.3	5.8	2.4	5.4	3.6		26.3
Eurodollar	2.2	0.9	1.9	0.8	0.9	0.4	4.5	1.7	5.9	3.0		31.2
10 Yr Note	5.4	0.6	2.0	0.4	1.2	0.3	7.2	2.2	7.7	3.3		17.8
5 Yr Note	5.9	0.8	2.3	0.4	1.3	0.2	7.2	1.7	7.8	2.6		16.1
Euro	10.8	0.7	4.5	0.3	2.3	0.3	10.2	1.7	6.3	1.8		10.1
Japanese Yen	11.7	0.9	4.3	0.3	2.1	0.2	10.0	1.3	6.7	1.4		9.7
Corn	1.6	0.6	1.0	0.5	0.7	0.4	4.2	3.0	4.5	4.7		30.0
Soybean	2.0	1.1	1.4	0.7	1.0	0.6	5.3	4.2	4.2	5.6		26.6
Gold	4.2	1.8	3.1	1.0	1.9	0.7	7.7	4.0	3.7	3.4		16.6
Silver	4.1	1.4	2.3	0.5	1.6	0.4	7.6	2.9	4.9	3.6		20.1

(0–100ms) most common in the Euro contract, but the lowest frequency bucket (10s–60s) the most common in the 10Y for much of the charted period.

Figure 3 provides a comprehensive look at the relative speed of automated vs manual trading across four commodities. The charts show the cumulative percent of passive volume which executes within a specified period of time. The chart averages these speeds across all days in the sample, and includes a one-standard deviation band around this daily average. Across all the selected products, automated trading is faster than manual; half of automated volume in the E-mini contract is executed within 5 seconds, whereas only a quarter of manual volume trades at this speed. This roughly 2–1 ratio remains consistent across the three other products (40 versus 20 percent in crude oil, 60 versus 30 percent in the Euro). The widest disparity in speeds occurs in the crude oil market, where, as above, manual activity is more prevalent than in the other three markets. The difference in speed is wide enough that even a one-standard deviation faster manual day in any of the contracts is slower than a one-standard deviation slower automated day. Echoing Table 7, markets are not all equally fast. Where the fastest market, the Euro contract, executes half of automated volume within 4 seconds, it takes more than 15 seconds to reach a similar level of activity in the crude oil market.

Thus far the focus has been on the passive side of the trade, the cases where market makers are providing liquidity to the rest of the market. Figure 4 switches views and explores what percentage of aggressive trading within a trading day is sourced from automated participants. For these four products, the charts depict the percentage of total volume through the day associated with automated and manual aggressive trading, divided up into five minute intervals. Of particular interest are the spikes in ATS activity related to Crude Oil during the settlement windows at the end of the day and during the 10:30 AM time period when weekly crude oil inventory news is publicly disseminated. It is during these periods when the benefits of speed is often greatest, when information is quickly incorporated into the futures price. Similar trends in the other commodities are more muted, outside of an increase in automated participation at the electronic market open for the E-mini contract.

Finally, Table 8 addresses the last of the topics outlined in the introduction, the speed of position turnover. The Table compares the holding period of large automated vs. manual traders. To do this, we identify the “directional” and “non-directional” component of trading for each participant. The non-directional component is the minimum of their absolute buy and sell volume for a given product and contract expiration over the given period, with the remaining volume attributed to the directional component. These two volume measures are summed across all participants and then divided by the total traded volume. Here, the large traders are identified using the same method as Table 6. Across all products, automated trading accounts for the majority of non-directional trading. The level of non-directional trading is also always higher than that for manual traders — however, note that because directional trades are calculated on a contract basis, spread trading, common among manual traders, is considered directional even though it has a pseudo non-directional component. Over smaller intervals of time, the percentages of non-directional automated trading drop, but are still significant for products such as E-mini S&P

500 and NASDAQ 100. For the E-mini NASDAQ 100 contract, almost 60 percent of trading within 1 minute is non-directional, implying that positions opened by a participant have more than an even chance of being closed within the next minute. In some cases, the difference in the fraction of non-directional trading over the whole day is similar to that within a much shorter five-minute interval. Interestingly, the shortest manual turnover rates are seen in commodities, where manual activity, generally, is strongest.

Table 8: Non-Directional Trading of Large Volume Traders

Notes: This table explores the intraday trading activity of “large traders” – specifically the amount of non-directional trading in defined intervals of time. For example, for the full day column, we add up all buy and sell quantity by expiration and account. Then we label the minimum of buy and sell total quantity as the amount of non-directional trading. The same is done across smaller time intervals such as 1, 3 and 5 minute intervals. Source: CME transaction data, November 12, 2012 – October 31, 2014.

Futures Product Name	Total Large Volume Trader VM (%)							
	Full Day		1 Min		3 Min		5 Min	
	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN
E-mini S&P 500	86.7	8.1	57.7	3.7	69.3	5.0	73.5	5.5
E-mini NASDAQ 100	89.2	2.0	59.9	0.4	70.8	0.6	74.6	0.7
Crude Oil	68.8	17.6	39.5	4.5	47.2	6.5	50.6	7.7
Natural Gas (Henry Hub)	59.8	19.0	28.1	3.8	35.2	5.7	38.6	6.9
Eurodollar	75.4	6.8	22.9	0.8	30.2	1.3	34.5	1.5
10 Yr Note	86.3	5.8	42.0	0.8	56.1	1.4	62.4	1.7
5 Yr Note	82.0	3.4	32.5	0.4	45.9	0.6	52.3	0.8
Euro	90.9	3.0	49.1	0.6	61.4	0.9	66.5	1.0
Japanese Yen	90.5	1.8	43.4	0.2	56.1	0.3	61.5	0.4
Corn	42.1	33.6	19.8	6.2	25.2	9.3	27.9	11.2
Soybean	48.8	28.1	25.4	4.8	32.2	7.3	35.2	8.9
Gold	58.2	22.8	30.3	2.4	38.0	4.0	41.2	5.0
Silver	50.7	30.4	25.6	3.6	33.7	6.1	37.0	7.5

Conclusion

This research note has focused on the characteristics, and prevalence, of automated trading in some of the world’s largest and most active futures markets. This analysis makes use of trade level information on futures transactions, allowing for a highly granular view on automation within these markets, its effect on the speed of a given market, and on the relative levels of liquidity provision across products. Because U.S. futures markets are some of the deepest and fastest in the world, the benefits of automation have led to its increased use by a wide set of participants — in some product classes as high as 80 percent. These benefits appear to accrue both to those providing and those demanding liquidity. In addition to these expected trends, mirrored in many other asset classes, some aspects of current market activity are more surprising. The proliferation of algorithms has not led to the complete disappearance of traditional manual traders, especially in the trading of physical commodities. These manual traders can often be quite large as a percentage of trading volume — in the physical commodity space, 20 percent of volume by large accounts is executed through manual means. Also, since automated participation, in general, is related to the level of market


participation and the speed of that participation; slower markets, including the slower activity seen in spread trading across even active products, have much higher manual participation (often well over half).

The information presented in the tables and figures also raise a number of open questions such as how automated trading systems respond to liquidity shocks, the relationships between or prevalence of different types of automated trading (market makers, directional traders, those who use statistical arbitrage), the value of manual trading during roll periods, and the effects of scheduled news announcements. These questions will be explored in future research notes.


Figure 1: Daily Volume Percentages for ATS-ATS (Red), ATS-MAN (Blue), and MAN-MAN (Green)

Notes: For each commodity, total daily volume across all expirations is divided into trades where both sides result from an automated system (ATS-ATS), where one side is an automated system (ATS-MAN), neither side is an automated system (MAN-MAN), and non-electronic volume (not included in the figures below). Percentages are computed, then smoothed in SAS according to a noniterative smoothing spline transformation (Reinsch; 1967) with smoothing parameter of 5. Source: CME transaction data, November 12, 2012 – October 31, 2014.


(a) E-Mini S&P 500



(b) Crude Oil



(c) Euro



(d) 10 Yr Note






Figure 2: Daily Liquidity Provision of ATS by Time Groups

Notes: This figure summarizes the speed of passive order execution across four futures products. The time difference between passive order entry (or most recent modification) and trade execution is calculated for each transaction. These times are then grouped into five ranges: (1) 0 – 100 ms (Blue), (2) 101 ms – 500 ms (Red), (3) 501 ms – 1 s (Green), (4) 1.001 s – 10 s (Purple), and (5) 10.001 s – 60 s (Orange), and then charted as percentages of total volume for each trading day. The charts are smoothed in SAS according to a noniterative smoothing spline transformation (Reinsch; 1967) with smoothing parameter of 5. Source: CME transaction data, November 12, 2012 – October 31, 2014.


(a) E-Mini S&P 500



(b) Crude Oil



(c) Euro



(d) 10 Yr Note




Figure 3: Liquidity Provision Time Between Order Placement and Execution for ATS (Blue) and Manual (Red)

Notes: This figure summarizes the cumulative volume of passive execution across four futures contracts. For each transaction the time between order entry (or most recent modification) and execution was computed. These differences were rounded to the nearest 10 ms and ordered by increasing time. From this ordered set, the cumulative volume as a percent of total in both the ATS or MAN categories is charted. The bands show 1 standard deviation above and below the average across the sample period. Source: CME transaction data, November 12, 2012 – October 31, 2014.






Figure 4: Percentage of Aggressive Volume associated with ATS (Blue) and Manual (Red)

Notes: This figure shows the average percentage of aggressive volume associated with automated and manual trading, divided up into five minute intervals between 7 AM and 3 PM CT. Source: CME transaction data, November 12, 2012 – October 31, 2014.


(a) E-Mini S&P 500



(b) Crude Oil



(c) Euro



(d) 10 Yr Note

