Macro News Announcements and Automated Trading*

Richard Haynes† John S. Roberts‡

December 18, 2015

Introduction

A previous white paper by the authors considered the prevalence, and characteristics, of automated trading across all U.S. futures markets in recent years (See Haynes and Roberts [2015]). That paper quantified automated activity levels over time and across products, but did not generally isolate specific periods of time of more interest to market participants. In this paper, we continue this overview of automated trading by highlighting those periods which are commonly some of the most important in the process of price discovery and efficiency — periods of scheduled economic announcements. In our analysis we concentrate on one major macroeconomic announcement, the monthly Employment Situation Summary released by the Bureau of Labor Statistics.[1] This analysis will focus on trading activity around this announcement in products commonly used as bellwethers for general economic, and monetary, activity — the Chicago Mercantile Exchange's E-mini S&P 500 and 10 Year U.S. Treasury Note contracts, with somewhat greater focus on the latter.[2]

Scheduled economic announcements are commonly used in the finance literature to study how new information is incorporated into a market price by a diverse group of market participants. Because prices provide a guide to product value, based on the judgments of a large set of individuals, scheduled announcements are events where market participants can adjust this value as they simultaneously acquire, and react to, new information. Differentiation in the speeds of market participants, including the technology or analytical tools used, can differentiate how and when

*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.

[1]This monthly news release commonly occurs on the first Friday of the month at 8:30 a.m. EST. The reported number captures changes in total non-farm payroll employment. The report is a closely watched indicator of the overall health of the economy. Further details, along with historical reports, can be found at the BLS website: www.bls.gov.

[2]The E-mini is a futures contract which prices against the S&P 500 stock index. The 10 Year Treasury futures contract settles to the cheapest-to-deliver U.S. Treasury bond with remaining maturity between 6.5 and 10 years. Each of these products trade on the Chicago Mercantile Exchange & Chicago Board of Trade (CME Group). For specific product details, see www.cmegroup.com.

traders react to this public information. A slow trader who can correctly interpret how an economic number will affect the price may react too late to successfully trade on this interpretation. In contrast, a fast trader who has a flawed understanding of market dynamics may quickly gather a position, only to see the price eventually move in an opposing direction. One method to increase reaction speed is to automate some, or all, of the trading process. By doing this, and by incorporating automation within an accurate prediction model, one may be able to react more successfully, and more efficiently, to announcement information. In addition, for those firms who are liquidity providers, speed may allow them to reduce the possibility of adverse selection, especially prevalent during these volatile periods.

In this paper, we consider how automation may affect liquidity provision and price discovery during the periods before and after the scheduled announcement. We use transaction level data in the two noted markets over the last few years; these data include, as a trader-defined field, a binary which notes whether that trader uses automation to generate orders; this binary is known as the manual order entry indicator. More details about this binary flag can be found in Haynes and Roberts [2015]. For each news release date, we define an observation (or event) as the ten minute window surrounding the 8:30 a.m. release time, and consider only outright trades in the most active futures expiration during this time period. For each transaction, we have information about the trade quantity, time, and trade aggressor, along with the manual indicator.

Using these data we will consider a set of related topics:

- The level of automated liquidity provision and liquidity taking around scheduled announcements, and how it may change before or after the news release
- How reactions to news differs between manual and automated participants
- Whether prices react, or lead to, manual or automated activity in the period right after the news announcement

The rest of the paper is organized as follows. The next section provides a review of recent academic literature on trading activity around announcements and how the level of trading automation may affect this participation. Following this, we provide a general overview of activity around the announcement in the E-mini and 10 Year Treasury futures contracts, contrasting levels just before the announcement with that after. The following two sections consider liquidity provision and liquidity taking, dividing this activity between manual and automated participants. We finish with a more analytical section, including regressions which attempt to describe potential causal relationships between price changes in the post-announcement window and trading activity.

Literature Review

The effect of announcements on trading behavior and price movements has been studied in significant detail across a variety of markets. A number of recent papers have considered how these effects may have changed in the recent

period of increasing automation. Balduzzi et al. [2001], as we do below, looked at the effect of macroeconomic announcements on the Treasury market, and found not only that news explains most of the volatility after the announcement, but that the price adjustment, even in the early 90's, was concentrated within the first minute after release. They also find that levels of liquidity provision, and the cost of trading, return to pre-announcement levels within five to fifteen minutes after the news release. We find very similar results, with the period of re-pricing and liquidity return significantly reduced given the increased market speeds in recent years. Chaboud et al. [2014] examined the effect of automation in the FX market and found that automated participants tend to have correlated strategies post-announcement, although automated participants also generally seem to reduce market volatility during these periods. This volatility and price adjustment is often proportional to the unanticipated 'surprise' associated to the news announcement. The greater the deviation between announcement expectations and the actual announcement, the larger the product re-pricing (often an asymmetric effect, with 'bad' news experiencing greater reactions, see Andersen et al. [2002]).

Other research has looked specifically at the value of speed during announcement periods. A few of these examined the effect of releasing the results of the Michigan Consumer Sentiment survey to a select set of participants two seconds earlier than the public. One, Wu et al. [2014], found that this schedule translated to increased revenues for those traders who received the announcement results in the earlier release — confirming the value of getting, and responding to, new information prior to much of the marketplace. Another, Hu et al. [2013], using the same set of announcements, found that most of the price move occurred within the first ten percent of trades after the initial release, indicating that market prices can very rapidly reach a new equilibrium even when only a small subset of market participants are fully informed. The authors also find that the trading activity of this small set of participants, during this initial period, increased the general price efficiency of the examined contracts. More generally, Scholtus et al. [2014] find that a 0.3 second delay in response to an announcement corresponds to around a three percent reduction in returns, quantifying, for a selected set of cases, the value of incremental increases in speed. Below we will consider how position and price changes may adjust, in response to general market behavior, over similar time intervals.

General Overview

Scheduled announcements are market periods often balanced between levels of market uncertainty and certainty. Just prior to the announcement, positions held by market participants are subject to the risk that the announcement will move prices in an unfavorable direction. Just after the announcement, interpretations of the information may differ, resulting in unanticipated volatility or temporary price movements which can arise simply through liquidity demand. When interpretations of the information across participants are similar, movements to a new equilibrium price level may be quite rapid, relative to the speeds of individual traders; this can lead to very quick permanent

price movements, with very little deviation from trend during this period. Other, more ambiguous announcements, may lead to more equivocal behavior, with buyers and sellers simultaneously trying to move the price in opposing directions. As can be seen from price series of futures contracts on selected announcement dates, all of these forces commonly appear, to differing degrees across days (see Figure 1). Across all of these events, however, these forces are concentrated over relatively short periods, usually on the order of a few seconds, highlighting the high speed of trader response.

Traditionally, price movements and trading volume before announcements are abnormally low, given the desire to avoid adverse price movements in the following minutes. In contrast, activity after the announcement is significantly higher than average, as a multiplicity of traders express their views. Though traders' beliefs do not always coincide, price discovery usually happens very quickly, with the majority of permanent price changes restricted to the minutes, or in many cases seconds, just following the release.

Figure 1: Sample Price Paths Surrounding Announcements for 10 Year Treasury Note Futures

Notes: Each figure shows realized transaction prices, over ten minutes surrounding the announcement, in the most active expiration by volume. The figures include red reference lines representing the before and after five minute average price. Source: CME transaction data.

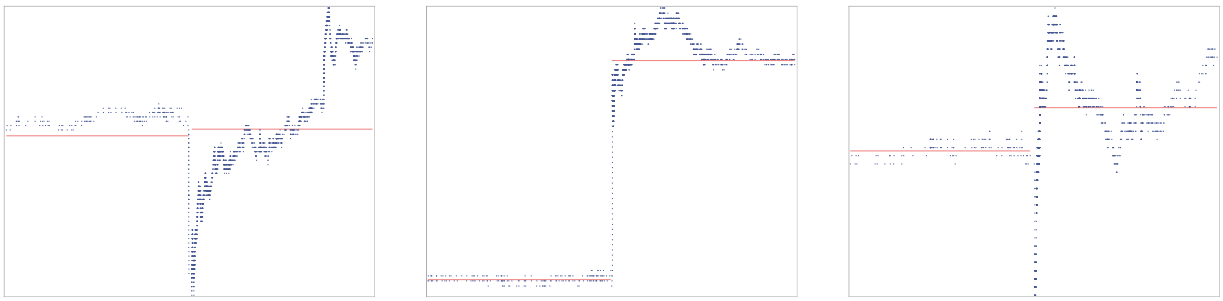

This difference in activity before versus after the announcement is clear across both of our examined contracts. Table 1 summarizes average activity during various periods around the announcement, isolating activity within the ten second and five minute windows straddling the announcement. This summary confirms the anticipated differences in volume and number of active participants for before versus after, with volume often increasing ten-fold in the post-announcement minutes.

Not surprisingly, this increase in volume and account activity is highly concentrated around the event itself. Given the reaction times made possible by electronic markets, many participants are able to respond to the new information with delays well less than a second, and time their orders accordingly. In the few seconds just after the announcement, trading volumes spike relative to the seconds just prior. Figure 2 provides a high-level summary of the level, and intensity, of response, by second, for the ten minute interval surrounding the public news event. While the charts display information from the 10 Year Treasury Note, similar patterns are observed for the E-mini S&P 500 product. Trading is subdued, relative to normal levels (based on Table 1), for all of the pre-announcement five

Table 1: Average Market Activity Before and After Announcements by Second

Notes: For each time interval, three groups are defined: PreNews represents the interval immediately prior to the news release; PostNews represents the interval immediately after the news release; NonEvent represents all days without a Jobs event. The ten second interval is defined as 8:29:50–8:30:00 and 8:30:00–8:30:10, for PreNews and PostNews, respectively. The five minute interval is defined as 8:25:00–8:30:00 and 8:30:00–8:35:00, for PreNews and PostNews, respectively. Source: CME transaction data.

	Ten Second Interval			Five Minute Interval		
Product Name	PreNews	PostNews	NonEvent	PreNews	PostNews	NonEvent
10 Year Treasury Note						
Volume	146	1,590	129	42	440	54
Active Participants	13	75	14	7	29	10
E-mini S&P 500						
Volume	48	817	62	26	225	24
Active Participants	13	88	15	8	31	9

minute period 2a, jumps in the immediate aftermath, and then, over the next few minutes, tapers to levels similar to, but often higher than, pre-announcement levels. The same is true for the average number of active accounts (Figure 2b). This pattern echoes findings in earlier papers and the general belief that announcements act as periods of price discovery where a broad set of market participants together express their trading interests.

Figure 2: Average Trading Volume and Number of Accounts by Second

Notes: The chart plots average trading volumes, and number of active accounts, in the most active 10 year Treasury futures contract for every second in the ten minutes surrounding the BLS jobs announcement. Averages are taken across the set of announcement dates from November of 2012 through September 2015. Source: CME transaction data.

(a) Trading Volume

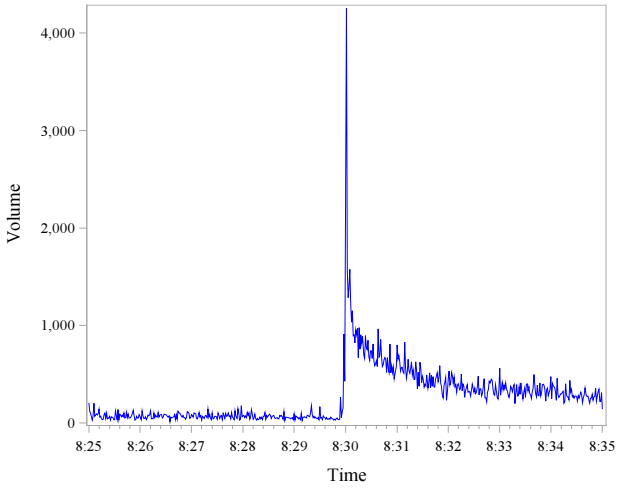

(b) Active Accounts

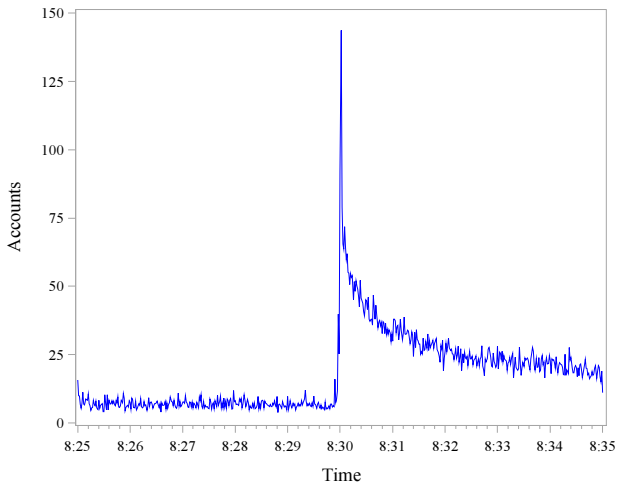

As will be clear below, though general patterns of behavior are consistent across announcements and products, implying similar responses by all participants to a given known event, individual accounts may have very different means of acting in response to information, and may act at very different speeds. In the next sections, we will look at two primary divisions across active accounts: liquidity providers vs. liquidity takers, and, as in the previous white

paper, manual (MAN) vs. automated (ATS) accounts. Activity by these groups differs significantly in the moments just before and just after the announcement. In brief, automated accounts, on a relative basis, are more active just after, versus just before, announcements, and generally seem to have much shorter holding periods for positions accumulated in response to the announcement. Manual accounts, instead, are more likely to have lower-frequency holding periods and strategies and, in certain cases, are responding to, rather than anticipating, price movements. We provide a brief discussion, and analysis, of all of these effects below.

Liquidity Provision around Announcements

The elevated interest in trading post-announcement requires both a market taker and a market maker — any participant interested in trading immediately after the announcement must trade against an order then sitting in the book. Either, or both, of these counterparties may be automated or manual. Because automation can provide the ability to respond to new information quickly, it may be the case that automation helps market makers more during the period after the announcement relative to before, when there are higher levels of price uncertainty. To identify the use of automation before and after public information, we divide market making activity by second, and monitor changes in this level over the selected ten minute period. To allow for cross-time comparisons and aggregations, we normalize by calculating these levels as a fraction of total volume and of total active accounts for the given time window. We find that liquidity provision by automated accounts often drops, sometimes significantly, just prior to the announcement and then quickly recovers, for both activity measures (see Figures 3a and 3b). The drop, and rebound, is of similar size across both products, although total automated activity is almost always lower in the E-mini S&P 500, where there is generally more participation by manual firms. Along with this reduction in automated market-making in the seconds preceding the announcement, general activity decreases, as we saw in the earlier table (See Table 1). This pull-back by automated participants fully reverses itself once the announcement has been made public. Participants who chose to reduce their liquidity provision, whether by quoting at smaller size or at wider bid-ask spreads for a period, often return during the very high volume period just after the news. These trends match patterns where there are relatively larger benefits of automation for market makers once information is known. By reacting to the new information, and strategically managing their inventory after the announcement, automated liquidity providers may be able to increase liquidity during the post-announcement period as they can more accurately avoid adverse selection.

Table 2 provides a further view into relative passive volume, broken down by manual or automated participation, over selected periods around the announcement. This table re-emphasizes the observation that manual liquidity provision is generally higher before the announcement than after as a fraction of volume, especially in the 10 Year Treasury contract. However, manual accounts can also be significant during certain periods of post-announcement activity. This is especially true in the case of the number of active accounts. Manual participation sometimes makes up well

Figure 3: Liquidity Provision by Automated Accounts for 10 Year Treasury (Blue) and E-Mini S&P (Red)

Notes: Focusing on the passive side of each electronic transaction, the volume and number of active accounts associated with automated order entry was calculated for each second to find the share of total second volume and number of accounts. These percentages are then smoothed in SAS according to a noniterative smoothing spline transformation (Reinsch; 1967) with smoothing parameter of ten. Source: CME transaction data.

(a) Trading Volume

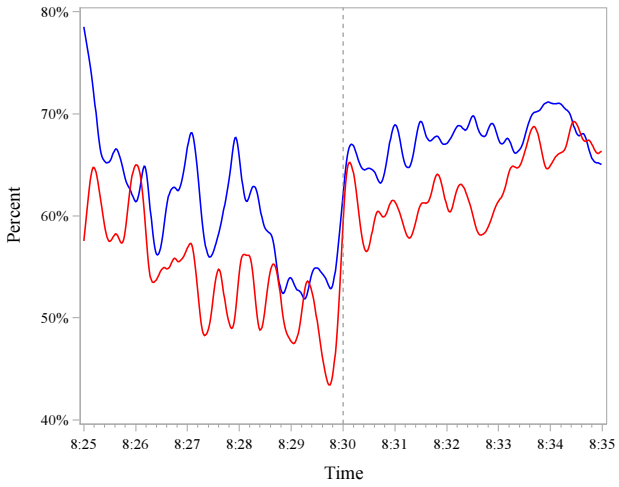

(b) Active Accounts

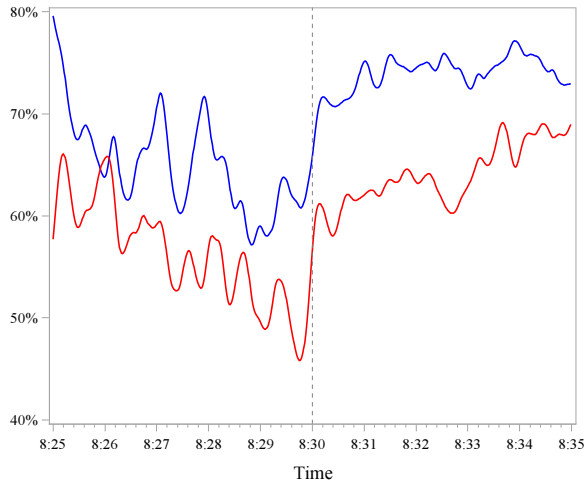

over 80% of all passive accounts within a half-minute period, even in cases where manual accounts are a minority of passive trading volume. Because of this, liquidity provision seems to be somewhat more concentrated across automated participants, with an average individual account providing liquidity across a larger number of trades. The balance between automated and manual participation does rapidly change just around the announcement event, though not always monotonically. Where automated and manual participation is reasonably balanced just before the announcement, automated levels drop slightly in the second coinciding with the event, but then grow in the subsequent minutes (again, most clearly in the 10 Year Treasury contract).

Table 2: Liquidity Provision by Order Entry Type

Notes: For the specified time group, passive electronic volume and number of unique accounts are divided by order entry type and percentages are then calculated. The time group excludes the second ending the interval, for example, the 8:30:00 – 8:30:01 time group only includes information from the initial second just after the announcement. Source: CME transaction data.

	10 Year Treasury (%)				E-Mini S&P 500 (%)			
	Volume		Accounts		Volume		Accounts	
Time Group	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN
8:29:50 – 8:30:00	51	49	33	67	52	48	27	73
8:30:00 – 8:30:01	46	54	40	60	37	63	20	80
8:30:01 – 8:30:10	57	43	29	71	55	45	14	86
8:30:10 – 8:30:30	63	37	32	68	55	45	15	85
8:30:30 – 8:31:00	61	39	32	68	53	47	15	85
8:31:00 – 8:32:00	64	36	30	70	56	44	14	86

The observed contribution from manual traders to liquidity provision immediately prior, and coincident, to the

public release of information is surprising given our intuition that technology advantages might aid automated traders during volatile periods. Actual trading behavior does not seem to always align well with this intuition. To better understand possible reasons for this, we make use of an audit trail time stamp which records either the order entry time or time of last modification, whichever is more recent. Using this, it is possible to calculate the resting time of passive orders until partial or full execution. Figure 4 shows the general level of order resting times dramatically changes in the seconds around the announcement. Where in the minutes prior to the announcement resting times are very short, these times jump significantly around 8:30. This often is a direct result of the large price moves around the announcement — orders which had been placed far from the prevailing price earlier in the day get executed as trades rapidly move towards those price levels. This phenomena seems to occur especially often for manual traders who may be holding a given position that they are interested in closing once a given profit level is hit.[3] If prices do move far enough as a result of the announcement, the standing manual order gets hit and the manual trader is aided by the given price volatility. In contrast, traders who hold less inventory and are primarily driven to capturing the bid/ask spread may need to more rapidly update, cancel or modify their orders during a price move, leading to no individual order staying in the order book beyond a relatively low time limit.

Figure 4: Order Resting Times by Order Entry Type (ATS=Blue, MAN=Red) for the 10 Year Treasury Product

Notes: Focusing on the passive side of each electronic transaction, the time difference between order entry and trade execution was calculated. Averages, by second, then are found by order entry type. Source: CME transaction data.

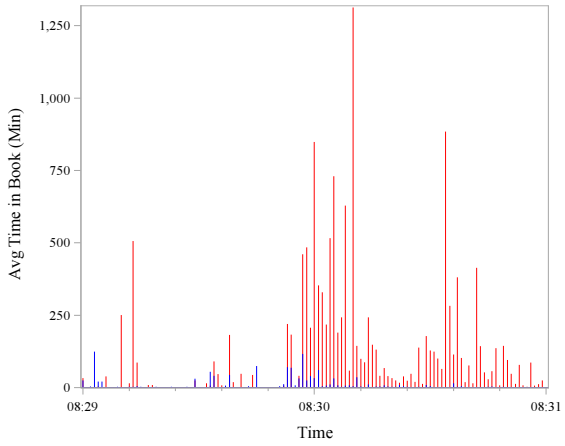

These findings are mirrored in a similar order resting time breakdown found in Figure 5, where we chart the percentage of executed passive orders with resting times over 30 seconds (left) or 5 minutes (right). Here, as before, resting times increase to their highest levels in the seconds after the announcement; in later seconds, manual resting times remain elevated, while automated resting times fall to levels much lower than those pre-announcement, especially in the 30 second resting-time window. These potential differing behaviors may point to reasons why,

[3]Of course, new positions can also arise as prices move toward a trader's willingness to acquire either a long or short position. Furthermore, we notice often longer order resting times are associated with the buying or selling of intra- or inter-product spreads (e.g., a spread between the 10 Year Treasury Note and the 5 Year Treasury Note). This figure indicates it is more common to see such activity associated with manual order entry — a topic for future research.

generally, manual liquidity provision rises just after the announcement — since these manual orders are sitting further away from the pre-announcement price, they are less prone to adverse selection, and thus are less in need of the benefits automation provides. This balance between speed and depth within the order book has been seen in other events within the futures market, such as activity on October 15th of last year in the Treasury market.[4]

Figure 5: Order Resting Times by Order Entry Type (ATS=Blue, MAN=Red) for the 10 Year Treasury Product, % of Trading Volume Greater Than Stated Time Difference

Notes: For each event and second, total quantity of liquidity provision is found by type (ATS or MAN). Then using the time difference between order entry time and trade time, we calculate the total quantity that is associated with time differences greater than or equal to the stated time difference in each figure below. For example, during the first second after news release, approximately 11 percent of ATS provided liquidity had a time difference of 5 minutes or greater (32 percent for MAN) — See 8:30:00 bar of Figure 5b. Source: CME transaction data.

(a) Thirty Seconds

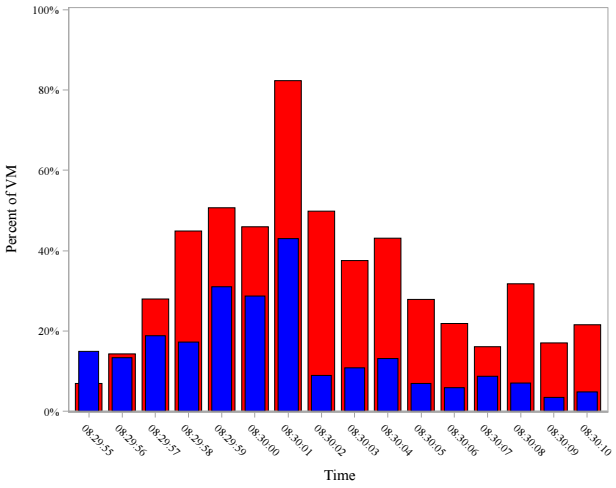

(b) Five Minutes

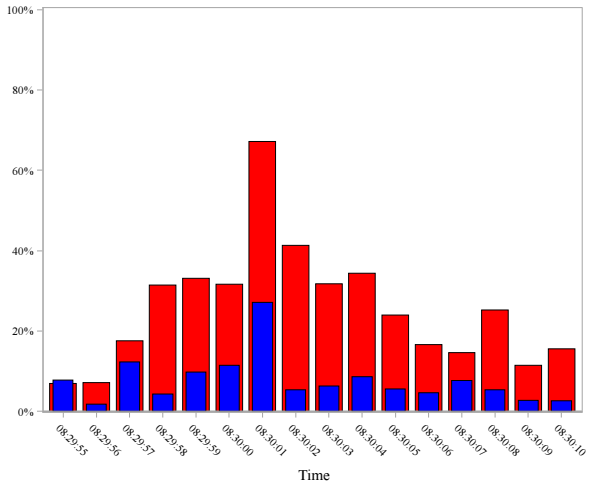

Liquidity Taking around Announcements

In the previous section, we highlighted differences in activity, and behavior, between automated and manual trading accounts acting as liquidity providers. Patterns in liquidity taking between these two groups are often similarly distinct. Speed, among other things, can allow an account to quickly respond to information and actively gather a profitable position. Table 3 provides an activity breakdown mirroring that of the prior section. Comparing this with Table 2, we see that in the minutes after the announcement, where automated accounts comprised the majority of liquidity provision, liquidity taking is often more balanced between both groups, dominated by automated accounts in the early seconds, but then often dominated by manual trading in the following minutes, especially in the E-mini contract. On a relative basis, as in the case of liquidity provision, automated accounts as a group exit the market just prior to the announcement, and take a bit of time, on the order of a second, to re-enter the market once

[4]For more on this activity, see the joint report: The U.S. Treasury Market on October 15, 2014 (2015).

information is known.

Table 3: Liquidity Taking by Order Entry Type

Notes: For the specified time group, non-passive electronic volume and number of unique accounts are divided by order entry type and percentages are then calculated. The time group excludes the second ending the interval, for example, the 8:30:00 – 8:30:01 time group only includes information from the initial second just after the announcement. Source: CME transaction data.

	10 Year Treasury (%)				E-Mini S&P 500 (%)			
	Volume		Accounts		Volume		Accounts	
Time Group	ATS	MAN	ATS	MAN	ATS	MAN	ATS	MAN
8:29:50 – 8:30:00	30	70	38	63	38	62	27	73
8:30:00 – 8:30:01	39	61	47	53	36	64	26	74
8:30:01 – 8:30:10	61	39	35	65	57	43	19	81
8:30:10 – 8:30:30	48	52	33	67	41	59	17	83
8:30:30 – 8:31:00	50	50	30	70	44	56	16	84
8:31:00 – 8:32:00	53	47	30	70	46	54	15	85

This aggressive activity is, on balance, in the direction of the price movement, with both aggressive automated and manual accounts buying as prices increase and selling as prices decrease (see Figure 6). Because aggressive activity can lead to price movements in the given direction (liquidity demand is subject to transaction costs, generally proportional to the size of the trade), this relationship is not that surprising, and seems to be similar across both groups. However, though the direction of trade is the same for automated and manual accounts, these groups differ in the length of time they hold these positions. Position holding periods for automated accounts, as a group, are relatively short, returning to a roughly flat position within the five minute window — coinciding with other analysis that many automated traders have short holding periods, often on the order of seconds or minutes.[5] In contrast to this, aggressive manual accounts continue to accumulate a position, again in the direction of the price movement, through the entire five minute period. Because of the time limit we place on our observations, and our concentration on just aggressive orders, these positions are likely closed over longer periods of time, and may be closed within the five-minute period using passive orders. However, with these caveats, the two categories of aggressive traders appear to be acting in a similar manner, but implement strategies over differing time scales.

Regression Analysis

To more formally analyze the patterns discussed in the previous sections, we incorporate price and position behavior into a vector autoregression (VAR), a common tool used when a set of variables may have complex, and interconnected, causal relationships. VARs provide an ability to look at the potential effect of a given variable on others in the same, or even delayed, periods. In this regression, we include five variables, each aggregated over half second periods: aggressive and passive position changes by automated and manual traders and price changes. As above, we concentrate on the 10 Year Treasury futures contract as a representative liquid contract affected by

[5]For an extensive literature review on academic research into automated and high-frequency trading, along with its effects, see the SEC's Equity market Structure Literature Review, Part II.

Figure 6: Liquidity Taking and Position Holdings by Order Entry Type with Price Movement

Notes: Figure shows the moving average, over the prior ten seconds, of aggressive positions as a fraction of total volume by automation class (blue and red lines). The average price (green) is a similar moving average, over the prior ten seconds, calculated relative to the average price in the five minutes before the news announcement. Prices and positions are normalized such that on those days with down price movements after the announcement, signs are reversed. Source: CME transaction data.

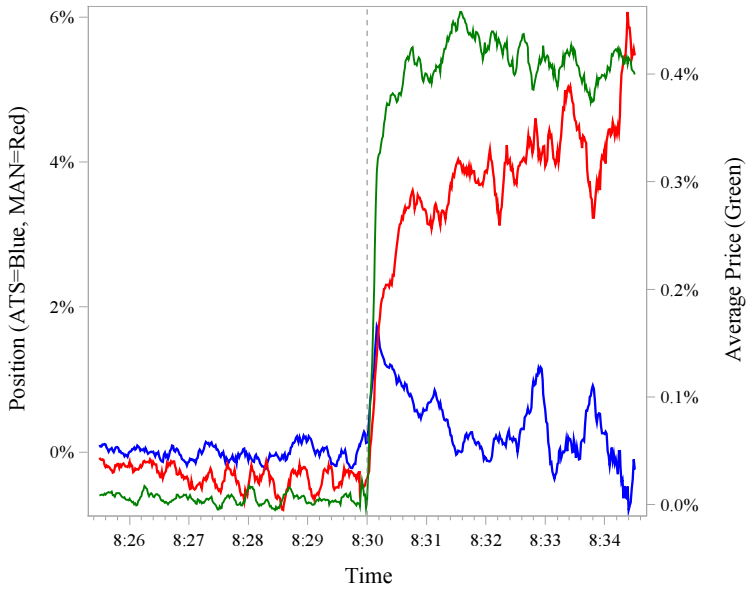

macroeconomic news. Position changes are defined such that if, in aggregate, automated accounts aggressively purchase ten contracts over a given half second period, they would be attributed a +10 in that time frame. All of these trades may have taken liquidity from automated accounts, in which case passive automated accounts would be assigned a -10 for the same period if this was their only activity. After constructing each individual change vector, all variables are scaled to have unit standard deviation and mean zero, in order to more easily compare effects across variables.

In building the VAR, we are interested in whether, or how, trading by different categories may translate into price changes around the announcement period, or whether accounts observe price changes and adjust their behavior in future periods accordingly. We begin with the assumption that effects should be larger in the post-announcement period, given informational effects, than in the pre-announcement period and thus concentrate on the latter.[6] We also assume that, because different accounts have different speeds, some will respond to prices, while others will be more coincident with prices. In order to potentially contrast activity during different periods of the post-announcement periods, we build impulse response charts for both the 8:30–8:31 one minute period (where there is greatest activity), and for the full 8:30–8:35 period.

Figures 7 and 8 summarize a selected subset of cumulative impulse response functions. As noted, we include a

[6]We tested pre-announcement activity using the same VAR specification and found few cross-variable effects, as expected, given that little new information, and thus price and position movement, would occur.

chart of responses during the one minute and during the five minutes just after the news announcement. These charts provide an overview of how various classes of accounts adjust their positions, through approximately a five second period (10 periods, each of half a second) after they observe an unexpected price movement. All of the position change categories see some level of response, though many at low levels of significance. In general, impulse response functions for automated versus manual traders are mirrors of each other, especially in the longer time period. Automated and manual market makers accumulate somewhat "offsetting" positions, with automated accounts slightly lengthening their positions, while manual traders gather short positions.[7] More varying, and interesting, dynamics happen for aggressive trading. As a response to a price increase, automated accounts actually aggressively sell over the course of the next few seconds, possibly closing positions that are now profitable given the price increase. Manual accounts, however, respond to price increases by progressively increasing their long positions over the next five seconds, peaking around 4 seconds after the initial price move, possibly anticipating further price increases.

Whether these anticipated price increases occur is often dependent on the time period. In the first minute after the announcement, prices may experience a high level of auto-correlation — price increases are followed by other increases — with, on average, almost all price changes permanent, representing the introduction of new information. This is less true, as we saw in the price charts in the initial sections, if we expand the region of interest to the full five minutes. A separate VAR analysis, seen in Figure 8, confirms this qualitative intuition. Here, price increases are usually followed by a gradual mean reversion, separating permanent from temporary price impacts. In both of these time windows, manual aggressive traders continue to accumulate positions. Automated aggressive accounts, however, trade in the direction of the mean reversion, selling as prices adjust downwards after the temporary liquidity shock.

Conclusions

In summary, we clearly find, as has been widely shown, that trading periods just after scheduled news announcements are both very active relative to normal times, and usually coincide with high market volatility as the new information gets processed. During this period, which often lasts on the order of a few seconds, a wide variety of accounts are active, on both the liquidity provision and liquidity demand sides. We find the use of automation, which can allow trading participants to make decisions and respond to market changes more quickly, to be higher during the period after the announcement, when the change in market prices and the forces of adverse selection are often similarly high. We also find that liquidity provision between manual and automated accounts can often differ during these periods of price change, with automated accounts reducing the average resting time of their orders, but manual

[7]Note that this discussion will use commonly used terms such as "going long" and "going short" in a relative context. Because the VAR analysis standardizes all variables to have mean 0 and standard deviation 1, every price and position change is relative to the average price/position change of that group, which has been assigned 0. This is close, but not necessarily equivalent, to absolute price and position changes.

accounts retaining orders in the book, often at price points that were further from the pre-announcement price levels. Finally, within a regression framework, we find persistent relationships between positions and prices for both automated and manual trading accounts; account activity, including positions gathered during trading, is often expressed at different speeds, with manual accounts commonly trading more slowly in response to market movements. In summary, automation does appear to help with inventory and price management during these known volatile periods, and seems to allow market participants to increase their levels of liquidity provision, even in turbulent markets, through its use.

Figure 7: VAR Impulse Response Curves: PostNews Period: 8:30 a.m. – 8:31 a.m. EST

Notes: The chart, focusing on activity during the first minute following news release, provides cumulative impulse response functions of price and position changes, by account category, representing responses to unexpected price changes. Source: CME transaction data.

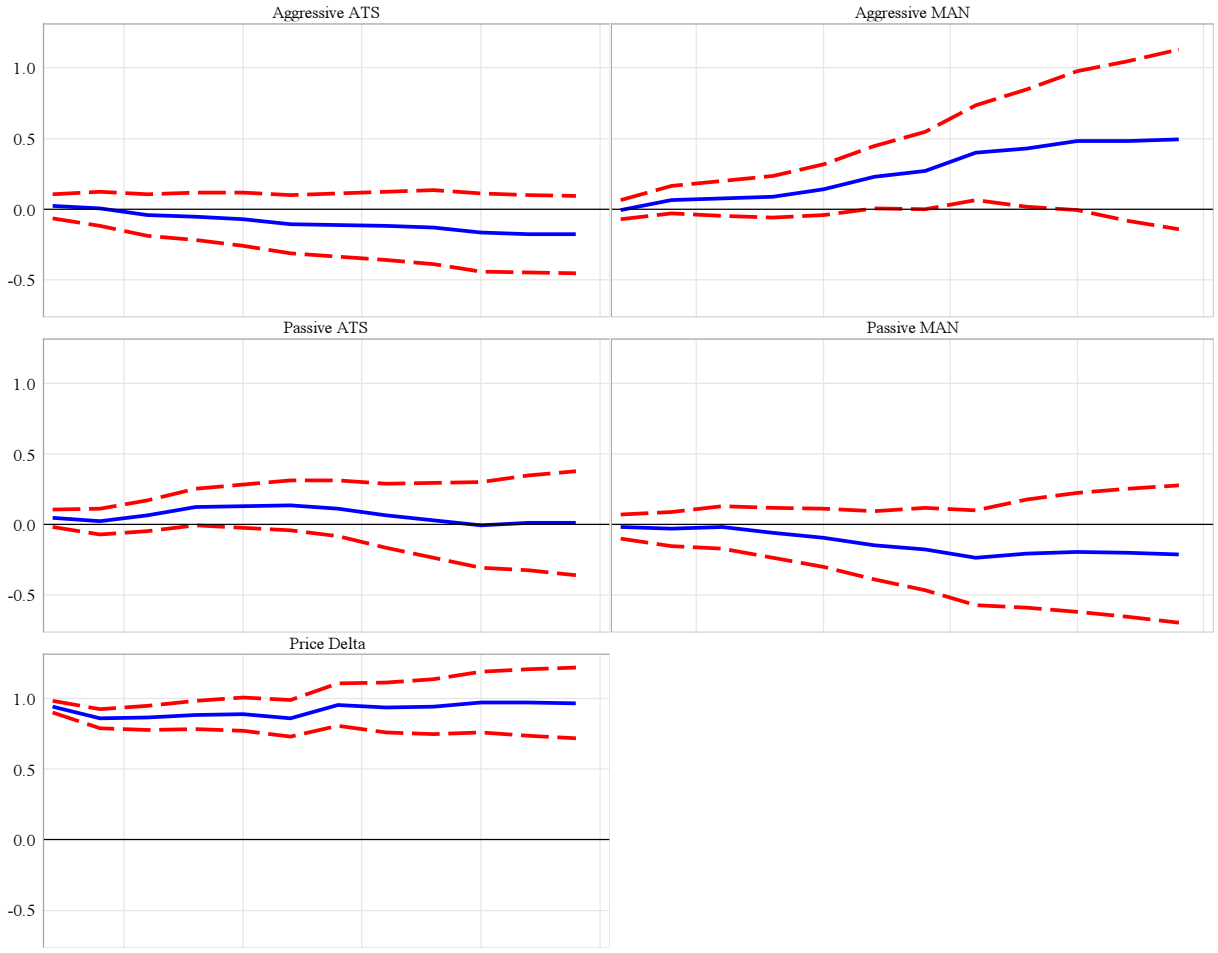

Figure 8: VAR Impulse Response Curves: PostNews Period: 8:30 a.m. – 8:35 a.m. EST

Notes: The chart, focusing on activity during the five minutes following news release, provides cumulative impulse response functions of price and position changes, by account category, representing responses to unexpected price changes. Source: CME transaction data.

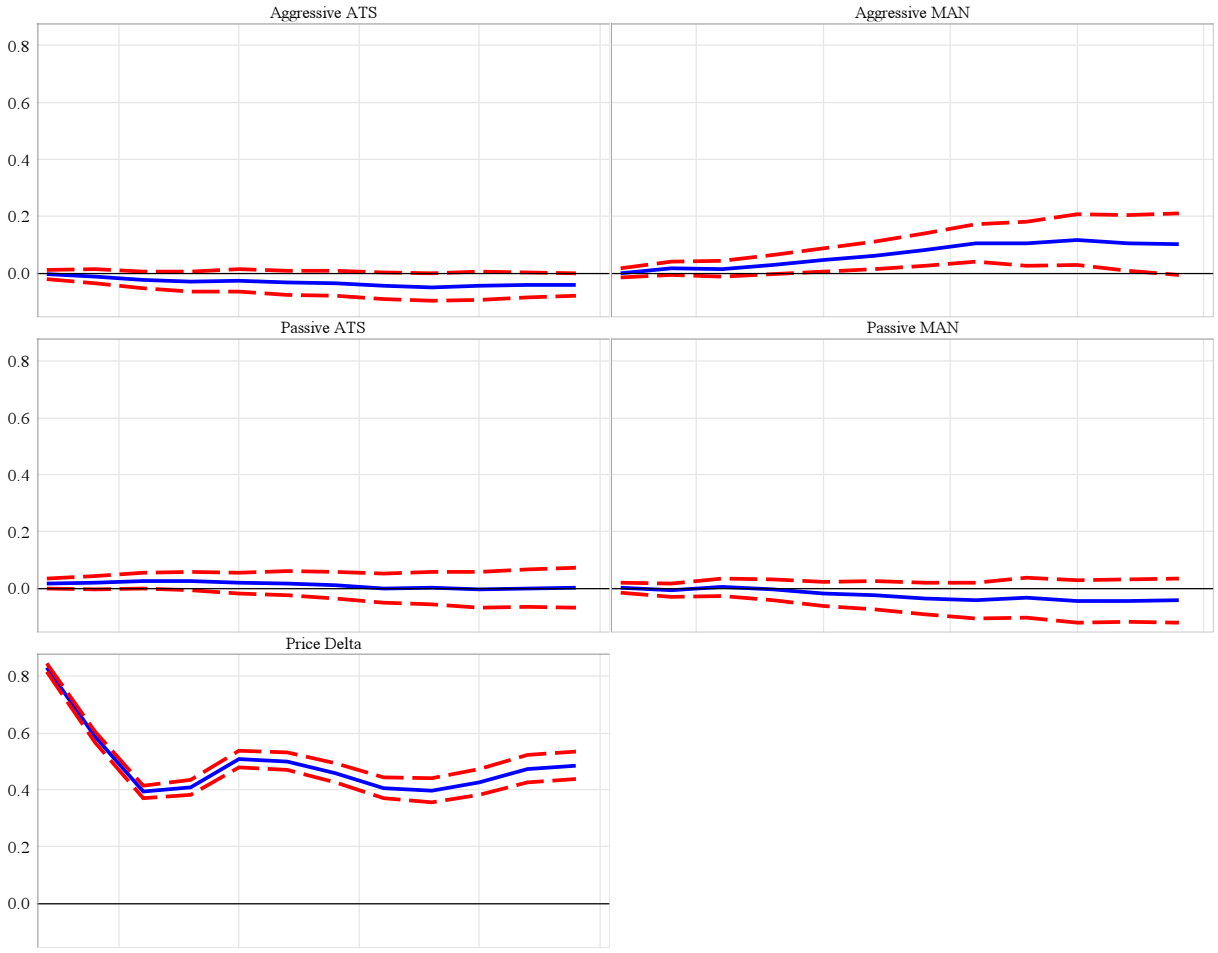

References

Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Clara Vega. Micro effects of macro announcements: Real-time price discovery in foreign exchange. Technical report, National Bureau of Economic Research, 2002.

Pierluigi Balduzzi, Edwin J Elton, and T Clifton Green. Economic news and bond prices: Evidence from the U.S. Treasury market. *Journal of financial and Quantitative analysis*, 36(04):523–543, 2001.

Alain P Chaboud, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega. Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5):2045–2084, 2014.

Richard Haynes and John S. Roberts. Automated trading in futures markets. *White paper, Office of the Chief Economist, Commodity Futures Trading Commission*, 2015.

Grace Xing Hu, Jun Pan, and Jiang Wang. Early peek advantage? *Available at SSRN 2361311*, 2013.

Martin Scholtus, Dick van Dijk, and Bart Frijns. Speed, algorithmic trading, and market quality around macroeconomic news announcements. *Journal of Banking & Finance*, 38:89–105, 2014.

Federal Reserve Bank of New York U.S. Securities U.S. Department of the Treasury, Board of Governors of the Federal Reserve System and U.S. Commodity Futures Trading Commission Exchange Commission. Joint staff report: The U.S. Treasury market on October 15, 2014. 2015.

Weishao Wu, Wenchien Liu, Sandy Suardi, and Chang Yuanchen. Informed high frequency trading with advance peek into the Michigan index of consumer sentiment. 2014.