

Intraday Momentum: The First Half-Hour Return Predicts the Last Half-Hour Return*

Lei Gao

Iowa State University

and

Yufeng Han

University of Colorado Denver

and

Sophia Zhengzi Li

Michigan State University

and

Guofu Zhou

Washington University in St. Louis[†]

First Draft: March, 2014

Current Version: December, 2015

*We are grateful to Phil Dybvig, Campbell Harvey, Steve Heston, Ryan Israelsen, Andrew Karolyi, Aditya Kaul, Naveen Khanna, Zhongjin Lu, Stefan Nagel, Matthew Ringgenberg, Ronnie Sadka, Robert Stambaugh, Jack Vogel, Mao Ye, seminar participants at Michigan State University, Shanghai Advanced Institute of Finance, Singapore Management University, Southwestern University of Finance and Economics, University of Missouri, and Washington University in St. Louis, and conference participants at the 2014 Tsinghua University Workshop, 2014 International Symposium on Financial Engineering and Risk Management, 2014 Financial Management Association meeting, 2015 Mid-Atlantic Research Conference in Finance, 2015 ITAM Finance Conference, and 2015 NFA Conference for helpful comments.

[†]Correspondence: Guofu Zhou, Olin School of Business, Washington University, St. Louis, MO 63130; e-mail: zhou@wustl.edu, phone: (314)935-6384.

Intraday Momentum: The First Half-Hour Return Predicts the Last Half-Hour Return

Abstract

Based on high frequency data of the S&P 500 ETF from 1993 – 2013, we document an intraday momentum pattern: the first half-hour return on the market predicts the last half-hour return. The predictability, both statistically and economically significant, is stronger on more volatile days, on higher volume days, on recession days, and on major macroeconomic news release days. This intraday momentum is also strong for ten other most actively traded domestic and international ETFs, and two major international equity index futures. Theoretically, the intraday momentum is consistent with the trading behavior of informed traders.

JEL Classification: G11, G14

Keywords: Predictability, Intraday, Momentum, Economic Value

Introduction

Since the seminal work of Jegadeesh and Titman (1993), it has been well-known that winners (losers) over the past six months to a year tend to continue to be winners (losers) over the next six months to a year. Griffin, Ji, and Martin (2003) show that momentum like this is common in global stock markets. In addition to this cross-section momentum, Moskowitz, Ooi, and Pedersen (2012) and Asness, Moskowitz, and Pedersen (2013) recently find evidence that time series momentum, where past returns of an asset positively predict its own future returns, is pervasive across asset classes such as equities, bonds, and currencies. To the best of our knowledge, however, almost all momentum studies are confined to the monthly or weekly frequency. An open question is whether there is any momentum at the intraday level. This question is of interest not only for examining the robustness of various momentum strategies, but also for understanding intraday market efficiency and the role played by daytraders including high-frequency traders.

In this paper, we provide the first study on intraday momentum, contributing in a unique way to the large literature of momentum studies.¹ Specifically, we find that the first half-hour return on the market significantly predicts the last half-hour return on the market.² We measure the market return by the actively traded S&P 500 ETF. The predictive R^2 of the first half-hour return on the last half-hour return is 1.6%, a level matching or exceeding a typical predictive R^2 at the monthly frequency (see, e.g., Rapach and Zhou, 2013). If the first half-hour return is combined with the twelfth half-hour return (the half-hour before the last half-hour), the R^2 increases further to 2.6%. We also find that predictability rises generally with volatility and volume. For instance, when the first half-hour volatility is high, the R^2 increases to 3.3% for the combined predictors. The predictability is stronger during recessions and on days with certain major economic news. Finally, we observe that intraday momentum is stronger on days when the first half-hour returns are positive than on days when the first half-hour returns are negative.

For out-of-sample (OOS) predictability, the R^2 is 1.2% using the first half-hour return as the only predictor, and 1.8% when this predictor is combined with the twelfth half-hour return predictor. Similar to the in-sample results, the OOS predictability is also greater than those typically found at the monthly frequency. In terms of economic significance, predictability based on either the first half-hour return alone or its combination with the twelfth half-hour return can generate certainty equivalent gains of 6.35% and 6.44% per annum, respectively, over ignoring the predictors for a mean-variance investor with a risk aversion of 5. In terms of market timing, the economic value is also substantial – the average return of the timing strategy using the sign of the first half-hour return is 6.67% per annum with a standard deviation of 6.19%. The Sharpe ratio is thus 1.08, which is remarkable compared to a level of 0.29 of a daily *Buy-and-Hold* strategy which delivers an average return of 6.04% per annum with a standard deviation of 20.57%. Moreover, the outperformance

¹For example, the latest number of Google citations of Jegadeesh and Titman (1993) is over 7360.

²In a recent study, Lou, Polk, and Skouras (2015) examine the intraday property of the cross-section momentum of stocks. In contrast, we study here the time-series momentum of the market.

remains significant even after accounting for transaction costs, which become increasingly lower due to quote decimalization in 2001 and advances in trading technology. Overall, the intraday momentum is both statistically and economically significant out of sample.

What are the economic forces that drive the intraday momentum? Theoretically, Admati and Pfleiderer (1988) show that informed traders will act strategically by timing their trades for high trading volume periods, which occur during the first and the last half-hours in our context due to the well-known intraday U-shaped trading volume pattern of the stock market. Hora (2006) demonstrates further that an optimal trading strategy is to trade rapidly at the beginning and at the end of the trading day, and trade more slowly in the middle of the day. Therefore, suppose there are good economic news in the first half-hour, informed traders are likely to bid up asset prices substantially. Then, in the last half-hour, their continued buying is likely to push the price further up, yielding our observed market intraday momentum. Moreover, this explanation is also consistent with the fact that the market intraday momentum is stronger when the first half-hour returns are positive than otherwise. In short, the trading behavior of informed traders explains the market intraday momentum.

The market intraday momentum is also consistent with the trading behavior of day-traders. Most major macroeconomic announcements, such as GDP and CPI, are released prior to 8:30am Eastern time, one hour before stock market trading starts. There is in addition various overnight news. Hence, a substantial rise in the first half-hour return is likely due to good economic news. In response to such a rise, some daytraders may go short to provide liquidity to the market, but they will unwind their positions later before the market closes. Shefrin and Statman (1985), Odean (1998), Locke and Mann (2000), Coval and Shumway (2005), and Haigh and List (2005) all suggest that daytraders can be subject to the disposition effect – they may be more reluctant to unwind losing positions than winning ones. Thus, as many of them are doing so during the last half-hour, their trading is likely to drive prices higher. The empirical evidence is consistent with this explanation. On a day when the first half-hour return is up substantially, the twelfth half-hour return is on average positive, making those who procrastinate have to unwind in the last half-hour. This also helps to understand the fact that the opening price on the following day is on average lower, which suggests that there is an adjustment of the price from the previous day's buying pressure in the last half-hour.

The intraday momentum is quite robust. It persists after accounting for reasonable transaction costs and market microstructure noises. Its economic value is significant for various risk aversion parameters and leverage constraints. Moreover, it is not limited to the S&P 500 ETF, but is also strong and significant for ten other most actively traded ETFs. These ETFs represent alternative stock indices, such as the Dow, the NASDAQ, and the Russell 2000. They also cover financial, real estate, bond and certain international equity indices. Interestingly, perhaps due to their lower liquidity, the out-of-sample predictability and the certainty equivalent gains on these ETFs are often greater than those on the S&P500 ETF. Furthermore, the intraday momentum exists in two major international equity index futures, the FTSE 100 and EuroStockXX 50. However, the intraday momentum does not

show up in major currency pairs or commodity futures. This is perhaps expected because, unlike the stock market, the daily open and close for currency and commodity futures are unclear because they are traded globally around the clock.

Our paper is related to the literature on intraday asset prices. Many of the existing studies have been focused on trading activity and volatility (see, e.g., Chordia, Roll, and Subrahmanyam, 2011; Corwin and Schultz, 2012). Heston, Korajczyk, and Sadka (2010) seem to be the only study that is closely related to ours. They find a striking pattern that returns on certain individual stocks tend to be persistent at the same half-hour intervals across trading days, and that this pattern can last for up to 40 trading days. In contrast to their study, we analyze intraday market momentum, namely, the predictability of the market’s first half-hour return for the market’s last half-hour return on the same day.

Our work is also related to the literature on price discovery. Barclay and Warner (1993), Chakravarty (2001), and Boehmer and Wu (2013) study how trading and traders of different types contribute to price discovery during a trading day and over longer horizons. Our paper by comparison suggests that the price discovery process can take at least a full trading day for the market to digest information, resulting in the intraday momentum.

The rest of the paper is organized as follows. Section 1 provides a description of the data. Section 2 documents the intraday momentum both in- and out-of-sample, and its property over volatility or volume regimes, and proposes two explanations. Section 3 provides an economic evaluation. Section 4 investigates its behavior over business cycles and news announcements. Section 5 examines the robustness of the results and Section 6 concludes.

1 Data

The intraday trading prices of the actively traded S&P 500 ETF (ticker SPY) are from the Trade and Quote database (TAQ) to compute half-hour returns. The sample period spans from February 1, 1993 through December 31, 2013. We exclude any trading days with fewer than 500 trades. For major news releases, we obtain the historical release dates of the Michigan Consumer Sentiment Index (MCSI) from the University of Michigan; the historical release dates of the GDP estimate from the Bureau of Economic Analysis; the historical release dates of the CPI from the Bureau of Labor Statistics; and the historical release dates of the Federal Open Market Committee (FOMC) minutes from the Federal Reserve Bank.³

Specifically, to examine the intraday return predictability, on any trading day t , we calculate the first half-hour return using previous day’s close price and the price at 10:00am Eastern time, and then every half-hour (30-minute) returns from 10:00am to 4:00pm Eastern

³The website for historical MCSI releases is <http://www.sca.isr.umich.edu/data-archive/mine.php>, for GDP releases is bea.gov/newsreleases/relsarchivegdp.htm, for Bureau of Labor Statistics announcements is www.bls.gov/bls/archived_sched.htm, and for FOMC minutes releases is www.federalreserve.gov/monetarypolicy/fomccalendars.htm.

time, a total of 13 observations per day, from

$$r_{j,t} = \frac{p_{j,t}}{p_{j-1,t}} - 1, \quad j = 1, \dots, 13, \quad (1)$$

where $p_{j,t}$ is the price at the j -th half-hour, and $p_{j-1,t}$ is the price at the previous half-hour, for $j = 1, \dots, 13$.⁴ Note that $p_{0,t}$ is the previous trading day's price at the 13th half-hour (4:00pm Eastern time). That is, we use the previous trading day's closing price as the starting price in calculating the first half-hour return on day t , i.e., $p_{0,t} = p_{13,t-1}$, so that the first half-hour return captures the impact of information released after the previous day's market close. To assess the impact of return volatility on return predictability, we also compute the volatility of the first half-hour return in two steps. First, we calculate the returns minute by minute within the first half-hour. Then, we compute the realized volatility using the estimated one-minute returns within the first half hour to obtain an estimate of the volatility of the first half-hour.

2 Intraday momentum

We first run predictive regressions to uncover the intraday momentum, and next investigate its out-of-sample performance. Then we examine the impact of volatility and volume on this momentum. Finally, we provide two intuitive explanations.

2.1 Predictive regressions

Consider first the simple predictive regression of the last half-hour return on the first half-hour return:

$$r_{13,t} = \alpha + \beta r_{1,t} + \epsilon_t, \quad t = 1, \dots, T, \quad (2)$$

where $r_{13,t}$ and $r_{1,t}$ are the last half-hour return and the first half-hour return on day t , respectively, and T is the total number of trading days in the sample.

The first column of Table 1 reports the results. The first half-hour return positively predicts the last half-hour return with a scaled (by 100) slope of 6.94, statistically significant at the 1% level, and an R^2 of 1.6%. Such a high predictive R^2 is impressive, as almost all typical predictors have lower R^2 's (see, e.g., Rapach and Zhou, 2013).

The twelfth half-hour (i.e., the second-to-last half-hour) may affect the last half-hour return too if there is a strong price persistence during the day. The second column of Table 1 reports the regression result using this predictor. It is clear that the twelfth half-hour return predicts the last half-hour return at the 1% significance level with an R^2 of 1.1%. We later show that this predictability largely comes from the recent financial crisis period,

⁴Similar results are obtained using log returns.

while that of the first half-hour return is always significant whether there is a crisis or not.

As r_1 or r_{12} predicts r_{13} individually, it is of interest to examine whether they can predict r_{13} jointly. The third column in Table 1 reports the predictive regression results using both predictors. Surprisingly, the slopes are barely changed from their individual regression values. Moreover, the R^2 of 2.6% is roughly equal to the sum of the individual R^2 's. The evidence suggests that r_1 and r_{12} are independent and complementary in forecasting the last half-hour return.

The standard monthly momentum strategy is known to have performed poorly during the recent financial crisis. How well the intraday momentum performs in this period is an interesting question. Panel B of Table 1 reports the predictive regression results from December 2, 2007, through June 30, 2009. The predictive power of r_1 in fact becomes stronger, with a larger slope of 13.6 and a higher R^2 of 4.1%. Moreover, the two predictors combined yield an amazingly high R^2 of 6.9%, rarely seen anywhere else. It may be noted that the predictive powers of r_1 and r_{12} are complementary during the crisis period too.

As the performance during the crisis period is so remarkable, a legitimate question is how the crisis affects the results of the whole sample period. Panel C of Table 1 addresses this question. Excluding those crisis days, performance clearly becomes much weaker. Although r_{12} is less significant, r_1 remains a powerful predictor of r_{13} with a sizable R^2 of 0.8%, comparable to many good predictors at the monthly frequency. The combined predictors yield a higher R^2 of 1.1%. Therefore, similar to studies on other trading strategies, although the predictability is time-varying due to, for example, the financial crisis, there is no doubt for the validity of intraday momentum over the entire sample period.

If the first and the twelfth half-hour returns can predict the last half-hour return, a natural question is whether any of the other ten half-hour returns can also predict r_{13} . To test the predictability of r_2, r_3, \dots , and r_{11} , we first examine if any of them used alone predicts r_{13} by performing a simple predictive regression analysis similar to Equation (2). Second, we examine if the explanatory power of r_1 and r_{12} on r_{13} remains after controlling for returns over other half-hour intervals by running a multiple regression that regresses r_{13} on r_1, r_2, \dots , and r_{12} simultaneously. To address the concern of data snooping, both simple and multiple regression analyses are performed not only for SPY but also for ten other most heavily traded index ETFs.⁵ Table IA.1 and Table IA.2 in the Internet Appendix report the results. Across all 11 ETFs, the predictability of r_1 is always statistically significant at the 1% level, and that of r_{12} is significant except for TLT. In contrast, none of the other 10 half-hour returns can significantly and consistently predict r_{13} across the board. In short, only the first and the twelfth half-hour returns can contribute to the intraday momentum.

2.2 Out-of-sample predictability

Our previous intraday momentum analysis is based on the entire sample (in-sample) estimation. While in-sample estimation is econometrically more efficient if regressions are stable

⁵Information on these index ETFs is detailed in Subsection 5.5 of Section 5.

over time, the financial crisis clearly destabilizes the estimation. At the monthly frequency, Welch and Goyal (2008) find that many macroeconomic predictors suffer from an instability problem, and their predictability largely vanishes once predictive regressions are estimated recursively out of sample (OOS). Thus, in-sample predictability does not necessarily imply OOS predictability.

To assess whether the intraday momentum persists out of sample, we run recursive regressions similar to other predictability studies at the monthly frequency. That is, to forecast return at any time t , we use data only up to time $t - 1$. Starting the regression using returns before January 3, 1998, we progressively add one more month of returns each time to form the OOS forecasts. Following Campbell and Thompson (2008), Rapach, Strauss, and Zhou (2010), Ferreira and Santa-Clara (2011), Henkel, Martin, and Nardari (2011), and Neely, Rapach, Tu, and Zhou (2014), among others, we use the OOS R^2 to measure the OOS predictability, defined as:

$$OOS R^2 = 1 - \frac{\sum_{t=1}^T (r_{13,t} - \hat{r}_{13,t})^2}{\sum_{t=1}^T (r_{13,t} - \bar{r}_{13,t})^2}, \quad (3)$$

where $\hat{r}_{13,t}$ is the forecasted last half-hour return from the predictive regression estimated through period $t - 1$, and $\bar{r}_{13,t}$ is the historical average forecast estimated from the sample mean through period $t - 1$. A positive *OOS* R^2 indicates that the predictive regression forecast beats the simple historical average.

Table 2 reports the results. When we use the first half-hour return alone, the *OOS* R^2 is 1.2%. When we use the twelfth half-hour return alone, the *OOS* R^2 is 0.7%. When we use both of them, the *OOS* R^2 achieves its highest value of 1.8%.⁶ The *OOS* R^2 's match or exceed those at the monthly frequency. As shown by Campbell and Thompson (2008) for monthly returns and confirmed later here, these levels of *OOS* R^2 are of substantial economic significance.

2.3 Volatility

Given that financial crisis is characterized by high volatility, earlier results during the crisis period are a special case of how intraday momentum performs under high volatility. In general, we can examine the impact of volatility by sorting all the trading days into three groups (terciles): low, medium, and high, according to the first half-hour volatility. For brevity, we consider the case of joint predictors of r_1 and r_{12} only.

Panel A of Table 3 reports the results. The predictability appears to be an increasing function of volatility. When the first half-hour volatility is low, the predictability is minimal with an R^2 of 0.6% and an insignificant coefficient for r_1 . At the intermediate volatility level, the R^2 rises to 1.0%, which is economically significant, and the coefficient of r_1 becomes highly

⁶Stronger results are obtained if we start the regression in later period. For example, the *OOS* R^2 is 2.08%, 1.19%, and 3.18%, respectively, if the regression is started after January 3, 2004.

significant. Finally, when the first half-hour volatility is high, the R^2 increases more than five times to as high as 3.3% compared to the low volatility case.

Overall, the intraday momentum seems highly related to volatility. The higher the volatility, the greater the predictability. This appears consistent with the theoretical model of Zhang (2006) that the greater the uncertainty, the stronger the persistence of a trend. In our context, the greater the volatility, the greater the likelihood that the first half-hour trend carries over to the last half-hour.

2.4 Explanations

Statistically, both the in- and out-of-sample analyses provide strong evidence on the intraday momentum. From an economic point of view, an interesting question is to find out what economic forces drive it. We provide below two explanations.

Our first explanation is based on the strategic trading of informed traders. Admati and Pfleiderer (1988) show theoretically that informed traders will time their trades for high trading volume periods. With a different preference specification, Hora (2006) also shows that an optimal trading strategy is to trade rapidly at the beginning and the end of the trading horizon, and to trade more slowly in the middle of the day. Figure 1A plots the average trading volume of the S&P 500 ETF every half-hour. Both the first and the last half hours have trading volume of close to 15 million shares, but the middle of the day has only about 5 million shares. The plot has a perfect U-shape, consistent with earlier findings about intraday trading activity (see, e.g., Jain and Joh, 1988). Now, according to the theories, given good economic news, informed traders are likely to trade more actively in the first half hour and thus bid up the price substantially. In the last half hour, they become active again and then their continued buying is likely to push the price up again. Figure 1B shows that the U-shape trading volume pattern is stronger on high volatility days, suggesting a stronger impact of informed trading as volatility rises. This is consistent with our earlier finding that intraday momentum is greater under greater volatility.

A direct assessment of the impact of volume on intraday momentum is given in Panel B of Table 3. Because trading volume has recently exhibited an upward trend largely because of substantially lower trading cost (Chordia, Roll, and Subrahmanyam, 2011), we need to control for the time trend effect in studying the volume and intraday momentum interaction. To do so, we first sort all trading days within each year into terciles based on the first half-hour trading volume, and then combine each volume tercile across all years to form the three volume groups. The predictive regression results in Panel B confirm that the intraday momentum is stronger when the first half-hour trading volume is higher. The R^2 increases from 1.1% when trading volume is low to 2.3% when trading volume is at an intermediate level, and then to 3.1% when trading volume is the highest.

Our second explanation is based on the trading behavior of daytraders. On a day when the first half-hour return is up substantially (e.g., due to overnight or early morning news), some traders may expect price reversion and go short. As they will almost surely unwind

to go flat before the market closes, some of them may wait to unwind in the last half-hour. Due to the disposition effect (see, e.g., Shefrin and Statman, 1985; Odean, 1998; Locke and Mann, 2000; Coval and Shumway, 2005; Haigh and List, 2005), they may be more reluctant to unwind losing positions than winning ones. On the other hand, on days with a substantial rise in price, the twelfth half-hour return is on average positive, making those who plan to unwind during this period wait to do so until the last half-hour. Therefore, there is likely even more unwinding of losing positions than usual in the last half-hour. Collectively, daytraders' buying is likely to push the last half-hour return higher than otherwise. Indeed, the opening price on the following day is on average lower, suggesting an adjustment of the price from the last half-hour buying pressure.

Both of our explanations provide an economic basis for the strong statistical evidence for the intraday momentum that the first half-hour return on the market predicts the last half-hour return on the market. Clearly, our explanations are limited in scope. Future research is called for to develop more models in understanding the trading motives, risk factors and the equilibrium factor risk premia.

3 Economic significance

In this section, to explore the economic significance of intraday momentum, we use the first half-hour and twelfth half-hour returns as timing signals either individually or jointly to examine the performance relative to a passive strategy that always holds the market (SPY) during the last half-hour. Then we use the predicted last half-hour returns from the OOS recursive predictive regressions to assess the certainty equivalent utility gains for a mean-variance investor.

3.1 Market timing

How well a predictor performs in market timing is a way to assess the value of the predictor. In our case, we use the first and twelfth half-hour returns as a timing signal to trade the market in the last half-hour. Specifically, we will take a long position in the market at the beginning of the last half-hour if the timing signal is positive, and take a short position otherwise. It is worth noting that the position (long or short) is closed at the market close on each trading day.

Consider first the use of the first half-hour return r_1 as the trading signal. Mathematically, the market timing strategy based on signal r_1 on day t will have a return in the last half-hour:

$$\eta(r_1) = \begin{cases} r_{13}, & \text{if } r_1 > 0; \\ -r_{13}, & \text{if } r_1 \leq 0. \end{cases} \quad (4)$$

The formula is clearly similar when using r_{12} as the timing signal.

When using both r_1 and r_{12} as the trading signal, we buy only if both returns are positive,

and sell when both are negative. Otherwise, we stay out of the market. Mathematically, the return is computed from

$$\eta(r_1, r_{12}) = \begin{cases} r_{13}, & \text{if } r_1 > 0 \ \& \ r_{12} > 0; \\ -r_{13}, & \text{if } r_1 \leq 0 \ \& \ r_{12} \leq 0; \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

3.1.1 Out-of-sample performance

Panel A of Table 4 reports summary statistics on returns generated from the three timing strategies. When we use the first half-hour return as the timing signal to trade in the last half hour, the average return is 6.67% on an annual basis.⁷ At first glance, this does not seem very high. To gauge the performance, we report three benchmark returns. The first is an *Always Long* strategy where we always take a long position in the market at the beginning of the last half-hour and close it at the market close. The first row in Panel B of Table 4 shows that the annualized average return of this strategy is only -1.11% . Hence, the timing strategy $\eta(r_1)$ outperforms this passive strategy substantially.

The second benchmark is a *Buy-and-Hold* strategy, where we simply take a long position in the market from the beginning of the sample, and hold it until the end of the whole sample period. The results are reported in the second row of Panel B. The average return is 6.04% per year, which is still below the average return delivered by the timing strategy, $\eta(r_1)$. Hence, 6.67% is remarkable, considering that we are in the market only for a half-hour each trading day instead of six and half-hours each day or all the time.

The third benchmark is a *Random Timing* strategy, where we use a random coin flip as the trade signal for the last half hour. The last row of Panel B shows that the average return is -0.68% per year. Therefore, the good performance of the timing strategy $\eta(r_1)$ is unlikely to be random.

Of course, we have to take risk into consideration. The standard deviation is 6.19% per annum for the timing strategy $\eta(r_1)$, resulting in a Sharpe ratio of 1.08. In contrast, the *Always Long* (*Random Timing*) strategy has a comparable standard deviation of 6.21% (5.79%), but a negative Sharpe ratio of -0.18 (-0.12). The long-term *Buy-and-Hold* strategy has a much higher standard deviation of 20.57%, and a much lower Sharpe ratio of 0.29. Note that the timing strategy $\eta(r_1)$ also enjoys a high positive skewness of 0.90 (versus -0.46 , -0.16 and -0.26 for the *Always Long*, *Buy-and-Hold* and *Random Timing* strategies, respectively) and a kurtosis of 15.65, suggesting that it often delivers high positive returns.

Note that the timing strategy trades only for the last half-hour even though we annualize the returns the same way as the daily return. However, since the timing strategy is exposed to market risk for only the last half-hour, its standard deviation is much lower and the Sharpe ratio is much higher than daily returns. As the Sharpe ratio is not very informative when

⁷Even though we are only in the market for the last half-hour, we still annualize the returns multiplying by a factor of 252 because we only trade once per day.

used to compare different strategies, we adopt another performance measure, the Modigliani-Modigliani measure (M2), which is related to the Sharpe ratio by

$$M2 = SRatio \times \sigma_b + r_f, \quad (6)$$

where $SRatio$ is the Sharpe ratio of the measured strategy, σ_b is the standard deviation of the benchmark portfolio, and r_f is the risk-free rate. Here we use the daily market return as the benchmark and assume the daily risk-free rate is zero. The economic interpretation of the M2 measure is that M2 is the average return of the measured strategy if the strategy is leveled up (down) to have the same volatility as the benchmark portfolio:

$$M2 = (\mu_s - r_f) \times \frac{\sigma_b}{\sigma_s} + r_f, \quad (7)$$

where μ_s and σ_s are the average return and standard deviation of the measured strategy. Table 4 shows that the M2 of the timing strategy $\eta(r_1)$ is 22.16% per annum, which suggests that this timing strategy would deliver an average return of 22.16% per annum if the timing strategy is leveled up to have the same risk (volatility) as the daily market returns (*Buy-and-Hold* strategy), which yields only 6.04% per annum.

Finally, we report the success rate, which is defined as the percentage of trading days with zero or positive returns. The success rate of the *Always Long* strategy is 50.42%, suggesting that the unconditional probability for the last half-hour returns is roughly 50-50. However, the success rate of the timing strategy $\eta(r_1)$ is higher at 54.37%.

Using the twelfth half-hour return as the timing signal yields similar but weaker results. The average return is about 1.77% per annum, Sharpe ratio is 0.29, skewness is 0.38, kurtosis is 15.73, and success rate is 50.93%. Overall, it still has a higher Sharpe ratio and a higher M2 measure than the *Always Long* benchmark.

Combining the two returns, r_1 and r_{12} , delivers improved performance over using only the twelfth half-hour return, but the performance is slightly weaker than using just the first half-hour return signal. For example, the average daily return is now 4.39% vs. 6.67% per annum, but the success rate is now much higher at an impressive value of 77.05%. This means that combining both r_1 and r_{12} does substantially improve the percentage of being right. Then, why does higher success rate yield lower average returns? The reason is that, when we combine the two signals, we take the long or short position only when both of them are positive or negative, which substantially reduces the number of days when we are in the market.⁸

⁸If we exclude the non-trading days with zero returns in the calculation, the strategy performs the best as expected, with an annualized average return of 8.85%, a standard deviation of 6.36%, and thus a Sharpe ratio of 1.39, a comparable skewness of 1.19, and a kurtosis of 18.30.

3.1.2 Impact of volatility or volume

We have observed in the in-sample predictive regression analysis that the intraday momentum is more pronounced on high volatility or volume days. To examine the impact of volatility on out-of-sample performance, we first sort all trading days into terciles based on the first half-hour volatility and report the out-of-sample timing results in Panel A of Table 5.

Overall, Panel A shows that timing strategies based on return predictability outperform the *Always Long* strategy under all scenarios, as is evident by higher average returns and Sharpe ratios. By looking at the impact of volatility, we find that the timing performance based on the first half-hour return is much better when the first half-hour volatility is higher. The average return per annum (and its t -statistic) of the $\eta(r_1)$ strategy rises substantially from 0.54% (0.43) in the low volatility group, to 4.75% (2.27) in the medium volatility group, and then to 14.73% (3.80) in the high volatility group. The Sharpe ratio (M2 measure) also rises from 0.18 (1.79% per annum) to 0.97 (15.48% per annum) and then to 1.63 (49.30% per annum). This enhanced out-of-sample performance of $\eta(r_1)$ on high volatility days is consistent with the better in-sample explanatory power of r_1 on high volatility days reported in Panel A of Table 3. Combining the first and twelfth half-hour returns as the timing signal confirms the positive interaction between the volatility and the predictability of the first half-hour return. Under the $\eta(r_1, r_{12})$ strategy, both the average return and the Sharpe ratio monotonically increase from the low to the high volatility groups.

Note that the first half-hour return predicts the last half-hour return both in sample and out of sample. If the predictability is due to the strategic trading of informed traders, as suggested by our first explanation, we would expect the intraday momentum effect to be stronger when the first half-hour trading volume is higher. To test this, we further form tercile volume groups, similar to Panel B of Table 3, run an out-of-sample timing performance analysis for days within each volume group, and report the results in Panel B of Table 5.

Comparing the three volume terciles in Panel B, we see that profitability of the $\eta(r_1)$ strategy improves both statistically and economically as the first half-hour trading volume increases. The average return per annum (and its t -statistic) of the $\eta(r_1)$ strategy increases from 1.67% (0.98) on low volume days to 6.46% (3.03) on medium volume days, and then further to a much higher level of 11.87% (3.23) on high volume days. The increase in the Sharpe ratio (M2 measure) from 0.42 (5.64% per annum) to 1.29 (23.93% per annum) and to 1.38 (37.67% per annum) of the $\eta(r_1)$ strategy also supports the implication that the first half-hour return predicts better on high trading volume days. Under the combined signal strategy of $\eta(r_1, r_{12})$, the average return rises from 2.10% per annum to 3.35% and then to 7.73% across the low, medium, and high volume terciles. All in all, these findings are consistent with the explanation that informed traders might time their trades for high volume periods such as the beginning and the end of the trading day, thus inducing a positive correlation between returns in the first and last half-hours.

3.2 Mean-variance portfolios

Instead of using only the signs to form timing strategies, here we use both the signs and magnitudes of the predictors to forecast the expected returns. Then we apply these expected returns to construct the optimal portfolio for a mean-variance investor who allocates funds between the market (SPY) and the risk-free asset (the Treasury T-bill).

The mean-variance efficient portfolio weights are given as

$$w_t = \frac{1}{\gamma} \frac{\hat{r}_{13,t+1}}{\hat{\sigma}_{13,t+1}^2}, \quad (8)$$

where $\hat{r}_{13,t+1}$ is the forecasted last half-hour return on day $t + 1$ conditional on information available at or before t and the predictor(s) at $t + 1$, and $\hat{\sigma}_{13,t+1}$ is the standard deviation of the last half-hour return, both of which are estimated from the recursive regression; and the relative risk aversion coefficient, γ , is set at 5. To be more realistic, we impose the portfolio constraint that weights on the risky asset must be between -0.5 and 1.5 , meaning that the investor is allowed to borrow or short 50% on margin. This will limit the potential economic gains from the usual unconstrained weights.⁹

Over the out-of-sample period, the realized utility is

$$U = \hat{\mu}_p - \frac{\gamma}{2} \hat{\sigma}_p^2, \quad (9)$$

where $\hat{\mu}_p$ and $\hat{\sigma}_p$ are computed based on the realized portfolio returns. In the out-of-sample forecasting literature, the historical average is usually the benchmark, and the certainty equivalent gain of predictability is computed from

$$CER = U_2 - U_1, \quad (10)$$

where U_2 is the realized utility of using the forecasted return $\hat{r}_{13,t+1}$, and U_1 is the realized utility of using the historical average mean forecast. From an economic perspective, CER can be interpreted as the gains of an investor who switches from believing in a random walk model of the intraday prices to believing in intraday momentum.

The results are reported in Table 6. Using the first half-hour returns to forecast the last half-hour returns yields an average return of 6.85% per annum, a standard deviation of 5.62% per annum, and thus a Sharpe ratio of 1.22, as well as large positive skewness and kurtosis. In sharp contrast, using historical average \bar{r}_{13} to predict the last half-hour return only generates an average return of 0.46% per annum, a standard deviation of 3.06% per annum, and hence a Sharpe ratio of merely 0.15. The CER using the first half-hour return is 6.35% per annum (the realized utility of using the historical average is only 0.46%), indicating sizable economic gains when investors switch from following a random walk model to following intraday momentum.

⁹The performance of the unrestricted portfolios is much stronger, which, although not reported for brevity, is indicated in Table IA.4 in the Internet Appendix.

When both the first and the twelfth half-hour returns are used to forecast the last half-hour returns, the portfolio delivers the best result, with an average return of 6.94% per annum, a Sharpe ratio of 1.13, and a CER of 6.44% per annum. Note that, unlike the case with market timing, using both predictors is slightly better than using the first half-hour return alone. This is because we are now always in the market. It is just that the allocation varies daily.

4 Macroeconomic events

In this section, we examine the performance of intraday momentum first over business cycles, and then on macroeconomic news releases.

4.1 Business cycles

We use the NBER dates for expansions and recessions to divide all trading days into these two types, and ask whether the intraday momentum effect interacts with the business cycle. We perform both in-sample predictive regression and out-of-sample timing performance for the two periods, and summarize the results in Table 7.

The comparison between these two periods suggests that intraday momentum has a more significant impact during recessions than expansions. Panel A shows that, during expansions, only the first half-hour return can predict the last half-hour return in sample. Albeit statistically significant, the predictability of r_1 is relatively weak, with an R^2 of 1.0% when using r_1 and r_{12} as joint predictors. During recessions, however, both the first and the twelfth half-hour returns are highly significant, and the R^2 increases more than six times to 6.6%. Such stronger predictability during recessions also translates into higher profits for market timing. For example, Panel B shows that, using the first and the twelfth half-hour returns as the timing signal (strategy $\eta(r_1, r_{12})$), the average return of the timing strategy in recessions is 16.79% per annum, seven times as high as 2.35%, the average return for the expansion periods. As a result, the Sharpe ratio is 2.10 in the recession periods, more than three times higher than the Sharpe ratio in the expansion period (0.66), despite the high volatility of the strategy (8.01% versus 3.57%). The performance of the timing strategy $\eta(r_1)$ also shows that the intraday momentum strategies perform better during recessions than during expansions.

4.2 News releases

Previously, we have found that intraday momentum is stronger on days with higher volatility or higher trading volume. One possible source of high volatility or trading volume may be the release of major economic news. It is hence of interest how news releases affect intraday momentum.

While there are many regular news releases, we here focus on four important news whose release times span across different time frames of the day. The first is the Michigan Consumer Sentiment Index (MCSI), released monthly at 10:00am. The next two are the major macroeconomic variables, the gross domestic product (GDP) and the consumer price index (CPI). Both of these are released monthly on pre-specified dates at 8:30am before the market opens, like most other macroeconomic news. The last is the minutes of the Federal Open Market Committee (FOMC), released regularly at 2:15pm about every six weeks. We analyze the impact of the news releases by dividing all the trading days into two groups: days with news releases, and days without.

Panel A of Table 8 reports the performance of intraday momentum for the two groups of trading days. On days without MCSI news, the R^2 is 2.6%. On days with MCSI releases, the R^2 more than doubles to 5.5%. That is, the intraday momentum becomes stronger. The same holds true when we compare the R^2 s on days without and on days with news announcements for GDP and CPI. These results seem to suggest that there is an information carryover effect of the news on market prices during the whole trading day.

The most astonishing result is for the releases of the FOMC minutes. While the no-release days have an R^2 of only 2.5%, the R^2 increases enormously to 11.0% on release days. There are two reasons why this result is astonishing. First, the R^2 is high by any standard, exceeding by far almost all predictors at the usual monthly frequency. Second, market participants seem to anticipate correctly in the first half-hour the message the Fed is going to send out to the market. Lucca and Moench (2015) find that pre-announcement excess equity returns account for sizable fractions of total realized stock returns, which is also a global phenomenon. Bernile, Hu, and Tang (2015) investigate market activity minutes prior to the release of the FOMC minutes. Unlike these studies, we focus on the intraday momentum. The high R^2 indicates that, even after the FOMC news release, there is a strong tendency of the market to continue the trend of the same direction anticipated in the first half-hour.

Will the higher R^2 s on the news release days imply greater economic gains? To answer this question, we examine the performance of the earlier market timing strategies on days with and without news release. Panel B of Table 8 reports only the results of using the first half-hour return, $\eta(r_1)$, for brevity. For the MCSI and CPI news, the gains are around three times the gains on the days without news releases. For the GDP news, the profits on release days are about twice as much. The greatest economic gains are delivered on the release days of the FOMC minutes. The annualized average return reaches a high level of 20.04%. This is close to four times the level on days without FOMC news. Overall, the performance of the intraday momentum is much stronger economically on the days with the four news releases.

5 Robustness

In this section, we examine the robustness of the intraday momentum on several dimensions. First, we analyze the intraday predictability conditional on the sign of the first

half-hour return. Second, we examine whether the gains of the intraday momentum can survive transaction costs. Third, we evaluate whether the intraday momentum is affected by microstructure noise. Further, we examine how the economic value measure may vary for various parameters and constraints on the mean-variance portfolio. Finally, we explore the evidence of intraday momentum on a set of the most actively traded ETFs, two popular international equity index futures, and other asset classes such as currencies and commodity prices.¹⁰

5.1 Conditional predictability

If either of our explanations holds, we would expect the intraday momentum to be concentrated mainly on days when the first half-hour returns are positive, and perhaps to be nonexistent when the first half-hour returns are negative. We test this implication by running predictive regressions conditioning on whether or not the first half-hour return is positive.

The results are reported in Table 9. During the whole sample period, the R^2 s for the three predictive regressions are 2.3%, 2.6%, and 4.5%, respectively, when the first half-hour return is positive. In sharp contrast, the R^2 s are only 0.5%, 0.3%, and 0.9% when the first half-hour return is negative. In addition, the first half-hour return, r_1 , is only marginally significant, and the twelfth half-hour return, r_{12} , is insignificant. An even greater difference is observed during the financial crisis period – the R^2 s increase to 4.5%, 11.0%, and 14.1%, respectively, for the three predictive regressions when r_1 is positive. On the other hand, the R^2 s are only 0.7%, 0.0%, and 0.9% when r_1 is negative, respectively, and both r_1 and r_{12} are insignificant. Finally, a similarly large difference is observed for periods excluding the financial crisis. For example, using r_1 as the predictor yields a R^2 of 1.1% when r_1 is positive compared to 0.1% when r_1 is negative. Neither r_1 nor r_{12} is significant conditional on r_1 being negative.

The results suggest that intraday momentum is a phenomenon specific to days when the first half-hour returns are positive, presumably because of good economic news, which is consistent with the two explanations we have proposed.

5.2 Transaction costs

What are the impacts of transaction costs on our results? With technological advancements and ever increasing competition in the financial industry, we have witnessed a significant decline in transaction costs over the past decades. This trend becomes even more evident after decimalization of quotations.

We examine the impact of transaction costs on the profitability of the intraday momentum using the market timing strategy as an example. To this end, we collect from the TAQ database the bid and ask prices at 3:30pm on each trading day and use the ask (bid) price

¹⁰Our study focuses on the intraday time-series momentum of the market or major indices. For the usual cross-section momentum, see Griffin, Ji, and Martin (2003), Schwert (2003) and references therein.

to calculate the last half-hour return if the market timing strategy takes a long (short) position.¹¹ Since the closing of the SPY is uniquely traded at the market clearing price for all the buys and sells, there will be no bid/ask spread effect for the price at 4:00pm.¹² Because of autoquotes of non-NYSE securities in the TAQ data before decimalization, we examine the effect of transaction costs only after decimalization (after July 1, 2001).¹³ The results are reported in Table 10.

Panel A of Table 10 shows that, using the first half-hour return as the timing signal, the average return reduces to 4.46% per annum, 2.47% lower than the average return before transaction costs, while the standard deviation remains the same at 6.10%. Nevertheless, the profits are still economically significant. Indeed, from the M2 measure, the strategy would yield an average return of 14.88% per annum if leveled up to have the same volatility as the daily returns. In contrast, the *Always Long* strategy which always invests in the market during the last half-hour yields an M2 of -2.45% per annum, and the daily market return (*Buy-and-Hold*) is 4.90% per annum for the same period. A slightly better result can be obtained when both the first and the twelfth half-hour returns are used to time the market; after adjusting for transaction costs, the average return is reduced by only 1.22% to 4.30% per annum with an M2 of 19.87%.

Figure 2 plots the time-series of the proportional spread after decimalization (after July 1, 2001). It shows clearly that the proportional spread narrowed after decimalization, and stabilized at around 1.2 basis point after 2005. To more closely capture the impact of transaction costs on future performance of the intraday momentum, we therefore consider the performance after January 1, 2005, reported in Panel B of Table 10. The average return of market timing using the first half-hour return is 6.52% after transaction costs compared with 7.96% before transaction costs. Similarly, the average return using both the first and twelfth half-hour returns is 4.74% after transaction costs versus 5.50% before transaction costs. Again, the leveraged average return (M2) is 20.77% and 20.82% per annum, respectively, much higher than the benchmark returns (-3.25% for the *Always Long* strategy and 6.75% for the *Buy-and-Hold* strategy).

5.3 Microstructure noise

Bid-ask bounce is known to induce negative autocorrelation, especially the first-order autocorrelation in high-frequency returns. If the bid-ask bounce effect is present in our data, it

¹¹We measure the bid and ask prices at 3:30pm using the median bid and ask prices at 3:30:00pm. If there is no quote at 3:30:00pm, we use the median bid and ask prices from the nearest previous second.

¹²We ignore the commission component of the transaction costs. At an online broker, such as Tradestation, an active individual investor may pay only \$4.99 commission for trading thousands of shares. The cost to active institutional investors can be even lower. In addition, some brokers even provide retail investors commission-free purchases and very low fees to sell.

¹³Autoquotes in the TAQ data are passive quotes by official dealers who are not making the market. Such quotes usually add a mechanical fraction on either side of the posted primary market quote, and hence will artificially inflate the quoted spread. The autoquotes issue is more severe in the pre-decimalization period, see Appendix B and Figure B-1 in Chordia, Roll, and Subrahmanyam (2001).

would indeed bias against our findings, which are based on returns formed from transaction prices. This is because the negative autocorrelation due to bid-ask bounce could attenuate the positive relation between r_1 and r_{13} and even more likely between r_{12} and r_{13} . To gauge this impact, we re-estimate the main predictive regressions in Table 1 using bid-to-bid, ask-to-ask, and midquote-to-midquote returns, and report the results in Panels B through D of Table IA.3 in the Internet Appendix. For completeness and to ease comparison, we also present the results using transaction price based returns (as in Table 1) in Panel A of the same table.¹⁴ As expected, the predictive power of r_{12} increases when returns are computed using bid, ask, or midquote prices over when returns are from transaction prices. For example, for the whole sample period regressions using only r_{12} as the predictor, the coefficient (t-statistic) of r_{12} increases from 11.94 (2.62) using transaction returns to 13.49 (2.88) using bid-to-bid returns, to 13.18 (2.80) using ask-to-ask returns, and to 13.65 (2.90) using midquote-to-midquote returns. The associated regression R^2 also increases from 1.1% in Panel A to 1.4% in Panel B, to 1.3% in Panel C, and to 1.4% in Panel D. The impact of bid-ask bounce on the predictive power of r_1 is minimal, however, as the estimated coefficient and t-statistic of r_1 stay largely the same across the four panels, and so does the R^2 . In short, the intraday momentum pattern cannot be induced by the bid-ask bounce but could actually be stronger after controlling for it.

5.4 Risk aversion and leverage

In Table IA.4 in the Internet Appendix, we examine the robustness of the out-of-sample mean-variance portfolio performance by varying the relative risk aversion coefficient, γ , and/or imposing different leverage restrictions on portfolio weights. For brevity, we consider only portfolios based on forecasts from using both the first and the twelfth half-hour returns. In Panel A, we keep $\gamma = 5$ and change the portfolio weight restrictions. The first alternative restriction is no-short sell and no-borrowing ($\psi_2 : 0 \leq w \leq 1.0$), which is more restrictive than the approach used in Table 6. Not surprisingly, the performance is poorer with an average return of 3.22% per annum but a Sharpe ratio of 0.82. The Sharpe ratio does not drop much because of the lower volatility of the portfolio. Relaxing the restriction by allowing shorting ($\psi_3 : -1.0 \leq w \leq 1.0$) increases the average return but also the volatility. In this case, the average return is around 7.35% per annum, CER is 6.61% per annum, and the Sharpe ratio is 1.26. Finally, we allow both shorting and borrowing ($\psi_4 : -1.0 \leq w \leq 2.0$), which delivers a much higher return (10.33% per annum), Sharpe ratio (1.19), and CER (9.55% per annum).

In Panel B, we set $\gamma = 2$ and impose various portfolio weight restrictions, and in Panel C, we allow γ to have a high value of 10. Overall, the results are very similar to Panel A where $\gamma = 5$. Of course, when no restriction is imposed, the average return and standard deviation are different for different γ as expected, and the lower γ is, the higher the average return and standard deviation are. But the Sharpe ratio should remain the same because

¹⁴The estimates in Panel A of Table IA.3 slightly differ from those in Table 1 because we here exclude days with fewer than one quote per half-hour to ensure the same sample across Panels A through D.

they are all on the same efficient frontier. Imposing portfolio restrictions, on the other hand, makes γ more or less irrelevant, and the portfolio performance is very close.

5.5 Other ETFs

Is the intraday momentum a special case for the S&P 500 ETF or a general phenomenon of the stock market? To address this question, we analyze the intraday returns of ten alternative ETFs.¹⁵ We choose the ten ETFs with highest average daily trading volume from their inception dates to December 31, 2013.¹⁶ Table 11 describes these ETFs. The asset classes are diverse. They include domestic alternative stock indices such as the Dow, the NASDAQ, and the Russell 2000 (DIA, QQQ, and IWM); international equity indices (EEM, FXI, EFA, VWO); two sector indices (XLF, IYR); and one bond index (TLT). If the intraday momentum found in SPY is also present in this diverse set of ETFs, it should lend more support to our trading behavior explanations.

We evaluate both the statistical and the economic significance of the intraday momentum in the same way as before. Table 12 reports the in-sample R^2 and out-of-sample performance measures for each ETF.¹⁷ We see a consistent pattern: the first half-hour return significantly predicts the last half-hour return. Moreover, utilizing such predictability generates substantial economic values. When the first half-hour return r_1 is used alone as a predictor, the in-sample R^2 ranges from 1.16% for DIA to 8.54% for EEM, and the out-of-sample R^2 is from 0.70% for QQQ to 6.53% for EEM. All the R^2 s strongly suggest that the first half-hour returns predict the last half-hour returns. In terms of economic value, the CER can be as high as 17.71% per annum for FXI, and many are greater than 10.0%. In comparison with the S&P 500 ETF, these ETFs are less liquid, so the price impact of the last half-hour trading is likely greater. This might help to explain their higher CERs in general. Adding r_{12} to r_1 as an additional predictor, we find a slight improvement over the single predictor r_1 , but the improvement is not uniform. In short, the results for various ETFs indicate a pervasive intraday momentum pattern in the stock market.

5.6 International Stock Futures

Another robustness analysis is to examine the intraday momentum in two other major international stock markets. To this end, we obtain from OneMarketData 30-minute intraday price and volume information for FTSE 100 index futures (FT) and EuroStockXX 50 index futures (XX) from July 1, 2003 to October 31, 2015.¹⁸ The trading hours of both futures have extended several times during the sample period. The current trading hour is from

¹⁵For the S&P 500, using futures data instead of the S&P 500 ETF produces similar results.

¹⁶We exclude a couple of heavily traded ETFs, which yield similar results, with inception dates later than 2005 and a few others to have a diverse and manageable set of ETFs.

¹⁷We delete trading days with fewer than 100 trades.

¹⁸Due to limited resources, those are the only international intraday data we have. We use the data from July 1, 2003 because, prior to this date, there is no volume information.

1:00am to 9:00pm local time for FT futures and from 7:50am to 10:00pm local time for XX futures, respectively.

However, trading activities are not evenly distributed over the entire trading day. Figure 3 plots the average trading volume every 30 minutes during the trading day for FT (Panel A) and XX (Panel B) futures, respectively. For FT futures, trading volume suddenly increases from almost zero and peaks at 8:30am (labeled as “1” in Panel A), similar to the first half-hour trading ending at 10:00am for SPY in the US stock market. Later in the afternoon, the trading volume increases again and peaks at 4:30pm (labeled as “17” in Panel A), similar to the last half-hour trading ending at 4:00pm for SPY. Between 8:30am and 4:30pm, the trading volume follows a U-shape, again, very much similar to the trading volume pattern of SPY. Similar pattern is observed for XX futures but at different time intervals. The trading volume drastically increases in the half-hour ending at 9:20am (labeled as “1” in Panel B), and then it peaks again at 5:50pm (labeled as “18” in Panel B). These trading volume patterns suggest that even though these index futures are traded in extended hours, traders/investors tend to trade with their regular working schedule. Therefore we designate 8:30am (9:20am) as the “first half hour”, and 4:30pm (5:50pm) as the “last half hour” for FT (XX) futures, and calculate the first half-hour return r_1 for FT (XX) futures from 4:30pm (5:50pm) of the previous trading day to 8:30am (9:20am) of the current trading day, and the last half-hour return from 4:00pm to 4:30pm (5:20pm to 5:50pm), denoted as r_{17} (r_{18}).

Table 13 reports the results. Panel A and B report the in-sample and out-of-sample predictive regression results, respectively, and Panel C reports the performance of mean-variance analysis. We denote the second to the last half-hour return for FT (XX) futures by r_{16} (r_{17}). For FT futures, r_{16} is not significant in the in-sample regression but significant (negative though) in the out-of-sample recursive regressions. However, the first half-hour return r_1 , is always significant and positive. Using r_1 as the only predictor, the in-sample R^2 is 1.90% and the out-of-sample R^2 is 1.87%, respectively. Similarly, for XX futures, r_1 is always significant and positive both in sample and out of sample. The second to the last half-hour return r_{17} is also significant and positive.

Panel C evidences the economic value of intraday momentum in FT and XX futures. For FT futures, the strategy that uses the first half-hour return to predict the last half-hour return delivers an average return of 1.84% and a Sharpe ratio of 0.43, compared to -0.72% and -0.29 of the strategy that uses the historical average (\bar{r}_{17}) to predict the last half-hour return. Similar results are obtained for XX futures. The intraday momentum strategy uses r_1 (r_1 and r_{17}) delivers an average return of 2.47% (3.03%) and a Sharpe ratio of 0.49 (0.61), compared to 0.28% and 0.05 of the strategy using the historical average (\bar{r}_{18}).¹⁹ Overall, it is quite interesting that the intraday market momentum is not only confined in the US, but also evident in the UK and in the Europe.

¹⁹The results are weaker than those for the US equity market probably due to the change of trading hours over time, which adds noise to the definition of the first and last half hour based on volume.

5.7 Currencies and Commodities

In this subsection, we further investigate the intraday momentum pattern beyond the stock and stock futures markets by examining nine major currencies and two major commodities.

The nine currencies are Australia, Canada, Euro, Japan, New Zealand, Norway, Sweden, Switzerland, and the United Kingdom, all of which are also examined by Moskowitz, Ooi, and Pedersen (2012). When these currencies are traded in the most liquid interbank cash market against the US dollar, they are quoted conventionally either in their own currencies or in US dollars: AUDUSD, USDCAD, EUROUS, USDJPY, NZDUSD, USDNOK, USDSEK, USDCHF, and GBPUSD. We obtain the intraday prices of these nine currency pairs from a major brokerage firm. Most of the pairs are available from November 11, 2004, through December 31, 2014, and the rest from January 6, 2005, through December 31, 2014. For commodities, the most liquid market is the futures market. We obtain intraday crude oil and gold futures prices from the same brokerage firm. The sample spans from September 1, 2005, through December 31, 2014. On each trading day, we use only the front-month contracts, i.e., the most traded and liquid ones.

Table 14 provides the results. For the currency pairs, the in-sample R^2 s are in general low and close to zero except for AUDUSD and USDJPY when using r_1 as the only predictor. R^2 s are improved substantially with the additional predictor r_{12} , suggesting strong autocorrelations between r_{12} and r_{13} in currency markets. However, for the out-of-sample tests, even lower and more negative R^2 s are observed. Again, this is especially true when r_1 is the only predictor. Results are marginally better when r_{12} is added to the regression. In addition, the CERs are small and even become negative in some cases. Similar results are obtained for the commodities. R^2 s are essentially zero when r_1 is used as the only predictor.

Overall, intraday momentum does not appear to exist in currency markets or commodity futures markets. These results are of no surprise, given the two explanations we have proposed, which critically depend on the structure of the stock market. In general, the majority of stock market participants can trade only when the exchanges are open from 9:30am to 4:00pm Eastern time, which helps generate the intraday momentum. Currency markets, however, trade 24 hours a day and 7 days a week. Therefore, traders do not have to wait until the market opens, or close their positions before the market closes. Similarly, even though commodity futures are still traded in the pit, electronic trading of commodity futures has become the dominant platform. Therefore, effectively, there are no open and close of the markets, and traders can trade continuously.

On the other hand, if we define the first and last half-hours for the currency pairs and commodity futures based on trading volume similar to the international stock index futures, we may find strong evidence for intraday momentum. Unfortunately, volume information is unavailable for the currency pairs and commodity futures. However, we do see some weak evidence of intraday momentum in the currency markets, particularly, for AUDUSD, EUROUS, USDJPY, and USDSEK. This predictability of the first half-hour return on the last half-hour return could be due to an artificial open and close of the markets. For example, a large proportion of currency traders work for prop trading desks of large US banks, and

most of their trades are submitted during regular working hours. We have observed similar pattern in the international stock index futures markets.

5.8 Data-snooping issues

Could our findings be caused by data snooping?²⁰ We argue that the intraday momentum pattern is strong and persistent, and so it is unlikely be explained by chance alone. First, the intraday momentum shows that r_1 is a powerful predictor. In Panel A of Table 1, the robust t -statistics of r_1 is 4.08, substantially exceeding the usual values of t -statistics between 2 and 3 seen in the return predictive regressions. Moreover, the in-sample R^2 of 1.6% is exceedingly high for an short-horizon return prediction problem. Such significant levels not only guard against the false discovery (see, e.g., Harvey, Liu, and Zhu, 2015), but also should bear less discount or “haircut” due to backtesting biases (see, e.g., Harvey and Yan Liu, 2015). Second, the performance of the intraday momentum is persistent throughout our sample, and, as summarized in Tables 1–8, the intraday momentum consistently emerges under vastly different market conditions characterized by financial crisis, volatility levels, trading volume levels, business cycles or macro news releases. Third, the intraday momentum is pervasive. As shown in Sections 5.5 and 5.6, the predicative power of the first half-hour return on the last half-hour return not only exists in SPY but also in a range of other most actively traded ETFs and two major international stock futures.²¹ In short, due to the plausible economic explanations and strong statistical evidence, the intraday momentum is likely to be a genuine phenomenon.

6 Conclusion

On the stock market intraday returns, we document in this paper that the first half-hour return on the market predicts the market return in the last half-hour. This intraday predictability is statistically significant both in- and out-of-sample. In terms of market timing and asset allocation, the economic gains of using the predictability are substantial. We also find that the market intraday momentum is stronger on high volatility days, high trading volume days, recession days, and important economic news (MCSI, GDP, CPI, FOMC) release days. Moreover, the intraday momentum is strong not only for the S&P 500 ETF, but also for ten of the most actively traded ETFs and two major international equity index futures. Finally, intraday momentum is more significant on days when the first half-hour returns are positive than otherwise. Theoretically, the market intraday momentum is consistent with the trading behavior of informed traders.

There are a number of open issues on intraday momentum. First, the documented em-

²⁰A related yet different concern is the overfitting bias, which is a result of using too many signals. Since the intraday momentum relies on a single or two signals, the overfitting issue is not relevant here.

²¹In fact, the first draft of this paper only considered SPY, and the robustness results on other ETFs and stock futures were found much later in responding to various helpful comments.

empirical facts in this paper call for new theoretical models of intraday trading to explain intraday momentum in a general equilibrium setting and to identify factors that determine its risk premium. Second, as trading costs become increasingly lower and trading execution becomes more automated, it is important to assess the asset pricing implications of intraday trading strategies and the associated implication for portfolio management. Thirdly, there is a huge literature on predictability at the monthly frequency, but it is unknown how intraday predictability imply monthly predictability. These are interesting topics for future research.


References

- Admati, A. R., and P. Pfleiderer. 1988. A theory of intraday patterns: Volume and price variability. *Review of Financial Studies* 1:3–40.
- Asness, C. S., T. J. Moskowitz, and L. H. Pedersen. 2013. Value and momentum everywhere. *Journal of Finance* 68:929–85.
- Barclay, M. J., and J. B. Warner. 1993. Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics* 34:281–305.
- Bernile, G., J. Hu, and Y. Tang. 2015. Can information be locked-up? Informed volume and price variability trading ahead of macro-news announcements. *Journal of Financial Economics* forthcoming.
- Boehmer, E., and J. Wu. 2013. Short selling and the price discovery process. *Review of Financial Studies* 26:287–322.
- Campbell, J. Y., and S. B. Thompson. 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21:1509–31.
- Chakravarty, S. 2001. Stealth-trading: Which traders' trades move stock prices? *Journal of Financial Economics* 61:289–307.
- Chordia, T., R. Roll, and A. Subrahmanyam. 2001. Market liquidity and trading activity. *Journal of Finance* 56:501–30.
- . 2011. Recent trends in trading activity and market quality. *Journal of Financial Economics* 101:243–63.
- Corwin, S. A., and P. Schultz. 2012. A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance* 67:719–60.
- Coval, J. D., and T. Shumway. 2005. Do behavioral biases affect prices? *Journal of Finance* 60:1–34.
- Ferreira, M. A., and P. Santa-Clara. 2011. Forecasting stock market returns: The sum of the parts is more than the whole. *Journal of Financial Economics* 100:514–37.
- Griffin, J. M., X. Ji, and J. S. Martin. 2003. Momentum investing and business cycle risk: Evidence from pole to pole. *Journal of Finance* 58:2515–47.
- Haigh, M. S., and J. A. List. 2005. Do professional traders exhibit myopic loss aversion? An experimental analysis. *Journal of Finance* 60:523–34.
- Harvey, C. R., Y. Liu, and H. Zhu. 2015. ...and the cross-section of expected returns. *Review of Financial Studies*, forthcoming.
- Harvey, C. R., and Yan Liu. 2015. Backtesting. Working paper, Duke University.

- Henkel, S. J., J. S. Martin, and F. Nardari. 2011. Time-varying short-horizon predictability. *Journal of Financial Economics* 99:560–80.
- Heston, S. L., R. A. Korajczyk, and R. Sadka. 2010. Intraday patterns in the cross-section of stock returns. *Journal of Finance* 65:1369–407.
- Hora, M. 2006. Tactical liquidity trading and intraday volume. Working paper, Credit Suisse Group.
- Jain, P. C., and G.-H. Joh. 1988. The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis* 23:269–83.
- Jegadeesh, N., and S. Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48:65–91.
- Locke, P. R., and S. C. Mann. 2000. Do professional traders exhibit loss realization aversion? Working paper, Texas Christian University.
- Lou, D., C. Polk, and S. Skouras. 2015. A tug of war: Overnight versus intraday expected returns. Working paper, London School of Economics and Athens University of Economics and Business.
- Lucca, D. O., and E. Moench. 2015. The pre-FOMC announcement drift. *Journal of Finance* 70:329–71.
- Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen. 2012. Time series momentum. *Journal of Financial Economics* 104:228–50.
- Neely, C. J., D. E. Rapach, J. Tu, and G. Zhou. 2014. Forecasting the equity risk premium: The role of technical indicators. *Management Science* 60:1772–91.
- Newey, W. K., and K. D. West. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55:703–8.
- Odean, T. 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53:1775–98.
- Rapach, D., and G. Zhou. 2013. Forecasting stock returns. In G. Elliott and A. Timmermann, eds., *Handbook of Economic Forecasting*, vol. 2A, chap. 6, 328–83. North-Holland, Amsterdam: Elsevier.
- Rapach, D. E., J. K. Strauss, and G. Zhou. 2010. Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies* 23:821–62.
- Schwert, G. W. 2003. Anomalies and market efficiency. In G. M. Constantinides, M. Harris, and R. M. Stulz, eds., *Handbook of the Economics of Finance*, vol. 1, chap. 15, 939–74. Amsterdam, Netherlands: Elsevier.

- Shefrin, H., and M. Statman. 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40:777–90.
- Welch, I., and A. Goyal. 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21:1455–508.
- Zhang, X. F. 2006. Information uncertainty and stock returns. *Journal of Finance* 61:105–37.

Panel A: Average 30-Minute Trading Volume



Panel B: Average 30-Minute Trading Volume Under High and Low Volatility




Figure 1: Average 30-Minute Trading Volume of SPY

For every 30-minute period from 9:30am to 4:00pm Eastern time, Panel A shows the average trading volume for SPY from February 1, 1993 through December 31, 2013. Each 30-minute period is labeled from 1 to 13 sequentially. Panel B plots the same 30-minute average trading volume on high volatility (top tercile) and low volatility (bottom tercile) days.

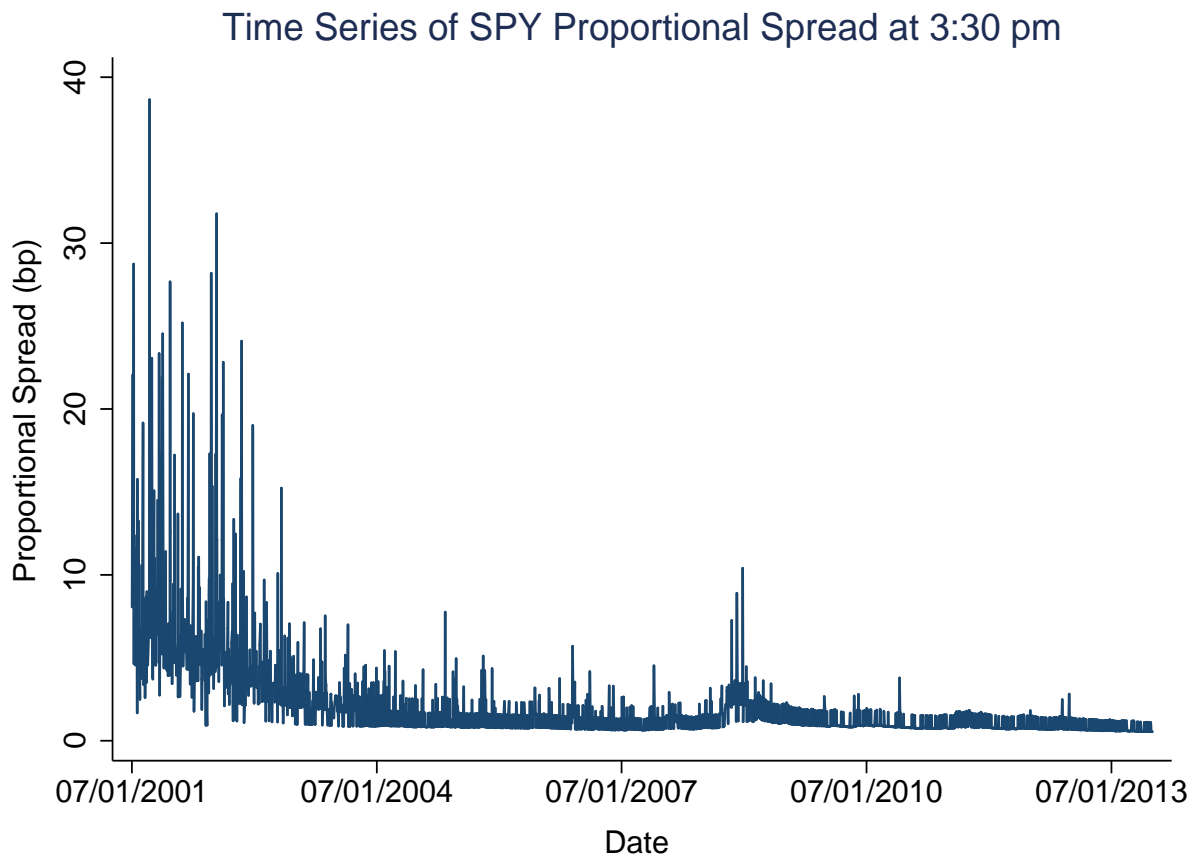



Figure 2: Time Series of Proportional Spread for SPY

This figure plots the proportional spread at 3:30pm on each trading day for SPY after decimalization (after July 1, 2001). The proportional spread is defined as $(Ask - Bid) / \text{Midquote}$, where the midquote price is the average of the bid and ask prices, $(Ask + Bid) / 2$.

Panel A: Average 30-Minute Trading Volume for FT



Panel B: Average 30-Minute Trading Volume for XX




Figure 3: Average 30-Minute Trading Volume of FT and XX index futures
 For every 30-minute period during trading hours, Panel A and B respectively show the average trading volume for FT and XX index futures from July 1, 2003 through October 31, 2015. Each 30-minute period is labeled sequentially, with “1” denoting the first half hour for both panels, and “17” (“18”) denoting the last half hour for Panel A (B).

Table 1: Predictability of the Last Half-Hour Returns

This table reports the results of regressing the last half-hour return (r_{13}) on the first half-hour return (r_1) and the twelfth half-hour return (r_{12}) of the day. The first half-hour return (r_1) is calculated from the closing of the previous trading day at 4:00pm to 10:00am Eastern time. Panels A, B, and C show results for three periods: the whole sample period, the financial crisis period from December 3, 2007, through June 30, 2009, and the periods excluding the financial crisis. The returns are annualized and in percentage, and the regression coefficients are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Predictor	r_1	r_{12}	r_1 and r_{12}	r_1	r_{12}	r_1 and r_{12}	r_1	r_{12}	r_1 and r_{12}
	Panel A			Panel B			Panel C		
	Whole Sample Period			Financial Crisis (12/2007–6/2009)			Excluding Financial Crisis		
Intercept	-1.63 (-1.16)	-1.33 (-0.94)	-1.82 (-1.28)	2.29 (0.29)	-1.66 (-0.20)	1.36 (0.17)	-1.63 (-1.25)	-1.25 (-0.97)	-1.72 (-1.31)
β_{r_1}	6.94*** (4.08)		6.81*** (4.14)	13.6*** (2.76)		13.2*** (2.88)	4.45*** (3.38)		4.40*** (3.36)
$\beta_{r_{12}}$		11.8*** (2.62)	11.4*** (2.60)		21.1* (1.95)	20.2** (1.99)		6.32* (1.88)	6.13* (1.83)
R^2 (%)	1.6	1.1	2.6	4.1	3.1	6.9	0.8	0.3	1.1

Table 2: Out-of-Sample Predictability

This table examines the out-of-sample predictability of the last half-hour return (r_{13}) by the first half-hour return (r_1) and the twelfth half-hour return (r_{12}) based on recursive estimations. The window of the estimation initially uses observations up to December 31, 1997, and progressively includes one more month of returns. The out-of-sample predictability is measured by the out-of-sample R-squared (OOS R^2):

$$OOS R^2 = 1 - \frac{\sum_{t=1}^T (r_{13,t} - \hat{r}_{13,t})^2}{\sum_{t=1}^T (r_{13,t} - \bar{r}_{13,t})^2},$$

where $\hat{r}_{13,t}$ is the forecasted last half-hour return from the predictive regression estimated through period $t - 1$, and $\bar{r}_{13,t}$ is the historical average return of the last half-hour estimated through period $t - 1$. The returns are annualized and in percentage, and the regression coefficients are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

	r_1	r_{12}	r_1 and r_{12}
β_{r_1}	4.51*** (29.5)		4.38*** (29.2)
$\beta_{r_{12}}$		6.88*** (22.8)	6.59*** (22.2)
OOS R^2 (%)	1.2	0.7	1.8

Table 3: Impact of Volatility or Volume

This table reports the results of regressing the last half-hour return (r_{13}) on the first half-hour return (r_1) and the twelfth half-hour return (r_{12}), under different levels of return volatility (Panel A) or trading volume (Panel B) in the first half-hour. The first half-hour volatility is estimated using one-minute returns within the first half-hour period, and then all the trading days are ranked into three terciles by their first half-hour volatility: low, medium, and high. For trading volume, we rank the trading days into low, medium, and high terciles by their first half-hour trading volume year by year to take into account the increasing trading volume over time, and then combine each volume tercile across all years to form three volume groups. The returns are annualized and in percentage, and the regression coefficients are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

	Panle A: Volatility			Panel B: Volume		
	Low	Medium	High	Low	Medium	High
Intercept	-2.18* (-1.76)	-3.07 (-1.51)	0.26 (0.07)	-4.36*** (-2.62)	1.22 (0.58)	-2.27 (-0.66)
β_{r_1}	2.34 (1.03)	5.40*** (2.93)	7.20*** (3.76)	4.32** (2.31)	7.22*** (3.32)	7.08*** (3.01)
$\beta_{r_{12}}$	8.81** (2.07)	8.39** (2.29)	12.7** (2.05)	10.1** (2.11)	6.16 (1.39)	13.7** (2.05)
R^2 (%)	0.6	1.0	3.3	1.1	2.3	3.1

Table 4: Out-of-Sample Market Timing

This table reports the economic value of timing the last half-hour market return using the first half-hour return, the twelfth half-hour return, or both jointly. For the timing strategy $\eta(r_1)$ ($\eta(r_{12})$), we use the sign of the first (twelfth) half-hour return as the timing signal – when the first (twelfth) half-hour return is positive (negative), we take a long (short) position in the market. When both returns are used jointly ($\eta(r_1, r_{12})$), we trade only when both returns have the same sign – long when both are positive and short when both are negative. The benchmark *Always Long* is to invest in the market during the last half-hour on each trading day, *Buy-and-Hold* is to buy and hold the market on a daily basis, and *Random Timing* is to time the last half-hour return using the toss of a coin as the signal. For each strategy, we report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, kurtosis, *M2* measure, and success rate (*Success*). The *M2* measure is estimated as the average return of the strategy with volatility leveled up to be the same as the volatility of the daily *Buy-and-Hold* strategy. The returns are annualized and in percentage. Newey and West (1987) robust *t*-statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Timing	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	M2(%)	Success(%)
Panel A: Market Timing							
$\eta(r_1)$	6.67*** (4.36)	6.19	1.08	0.90	15.65	22.16	54.37
$\eta(r_{12})$	1.77 (1.16)	6.20	0.29	0.38	15.73	5.88	50.93
$\eta(r_1, r_{12})$	4.39*** (3.96)	4.49	0.98	1.87	34.10	20.13	77.05
Panel B: Benchmark							
Always Long	-1.11 (-0.73)	6.21	-0.18	-0.46	15.73	-3.69	50.42
Buy-and-Hold	6.04 (1.19)	20.57	0.29	-0.16	6.61		
Random Timing	-0.68 (-0.54)	5.79	-0.12	-0.26	16.79	-2.23	

Table 5: Impact of Volatility or Volume on Out-of-Sample Timing Performance

This table reports the impact of the first half-hour volatility (Panel A) or trading volume (Panel B) on the economic value of timing the last half-hour market return using the first half-hour return (r_1), or the first half-hour return and the twelfth half-hour return (r_1 and r_{12}). The timing strategy is described in Table 4. We report the timing performance for three different levels of the first half-hour volatility or volume. For each strategy, we report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, and *M2* measure, which is the average return of the strategy with volatility leveled up to be the same as the volatility of the daily *Buy-and-Hold* strategy (not shown). The returns are annualized and in percentage. Newey and West (1987) robust *t*-statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Timing	Panel A: Volatility						Panel B: Volume					
	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	M2(%)	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	M2(%)
	Low Volatility						Low Volume					
Always Long	-2.04 (-1.62)	2.95	-0.69	-0.51	2.48	-6.80	-4.03** (-2.37)	3.98	-1.01	-0.78	6.08	-13.64
$\eta(r_1)$	0.54 (0.43)	2.95	0.18	-0.29	2.57	1.79	1.67 (0.98)	3.98	0.42	-0.54	6.30	5.64
$\eta(r_1, r_{12})$	0.97 (1.17)	1.93	0.50	0.12	5.87	4.94	2.10** (1.93)	2.53	0.83	1.08	13.25	11.14
	Medium Volatility						Medium Volume					
Always Long	-2.36 (-1.13)	4.89	-0.48	-0.25	2.83	-7.66	1.96 (0.92)	5.01	0.39	-0.02	3.94	7.23
$\eta(r_1)$	4.75** (2.27)	4.89	0.97	-0.14	2.91	15.48	6.46*** (3.03)	5.00	1.29	0.09	3.95	23.93
$\eta(r_1, r_{12})$	3.78*** (2.69)	3.28	1.15	0.79	9.07	18.32	3.35** (2.24)	3.50	0.96	0.74	14.09	17.68
	High Volatility						High Volume					
Always Long	1.05 (0.27)	9.10	0.12	-0.42	8.64	3.51	-1.29 (-0.35)	8.63	-0.15	-0.44	10.84	-4.08
$\eta(r_1)$	14.73*** (3.80)	9.06	1.63	0.76	8.50	49.30	11.87*** (3.23)	8.60	1.38	0.96	10.68	37.67
$\eta(r_1, r_{12})$	8.42*** (2.91)	6.77	1.24	1.44	17.62	37.75	7.73*** (2.80)	6.45	1.20	1.63	21.00	32.69

Table 6: Mean-Variance Portfolio Performance

This table reports the economic value of recursively predicting the last half-hour market return using the first half-hour return or combining with the twelfth half-hour return. We use the predicted returns to form a constrained mean-variance optimal portfolio for a mean-variance investor with a relative risk aversion of 5. Portfolio weights are restricted to between -0.5 and 1.5. For each strategy, we report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, kurtosis, and the certainty equivalent gain of return, *CER*, calculated as the difference in the certainty equivalent rate of return between the optimal mean-variance strategy and the benchmark (which uses the recursively estimated average returns of the last half hour returns instead of the forecasted last half-hour returns). The returns are annualized and in percentage. Newey and West (1987) robust *t*-statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Predictor	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	CER(%)
\bar{r}_{13}	0.46 (0.57)	3.06	0.15	0.48	18.05	0.46
$\beta_1 r_1$	6.85*** (4.55)	5.62	1.22	1.74	48.81	6.35
$\beta_1 r_1 + \beta_2 r_{12}$	6.94*** (4.23)	6.12	1.13	0.56	59.84	6.44

Table 7: Impact of the Business Cycle

This table reports the impact of business cycles on the predictability of the last half-hour return (r_{13}) by the first half-hour return (r_1) and the twelfth half-hour return (r_{12}). The expansion and recession periods are defined by the NBER. Panel A reports the results of the predictive regressions, while Panel B reports the results on the economic value of market timing. The timing strategy is described in Table 4. For each strategy, we report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, and kurtosis. The returns are annualized and in percentage, and the regression coefficients in Panel A are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Panel A: Predictive Regression					
Predictor	Intercept	β_{r_1}	$\beta_{r_{12}}$	$R^2(\%)$	
Expansion					
r_1	-2.34* (-1.76)	4.83*** (3.39)		0.9	
r_1 and r_{12}	-2.41* (-1.80)	4.80*** (3.39)	4.32 (1.26)	1.0	
Recession					
r_1	5.42 (0.89)	11.4*** (2.76)		3.2	
r_1 and r_{12}	4.79 (0.78)	11.0*** (2.87)	21.6** (2.30)	6.6	
Panel B: Market Timing Performance					
Timing	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis
Expansion					
Always Long	-1.73 (-1.29)	5.05	-0.34	-0.03	8.53
$\eta(r_1)$	4.63*** (3.44)	5.04	0.92	-0.13	8.61
$\eta(r_1, r_{12})$	2.35*** (2.46)	3.57	0.66	0.26	23.26
Recession					
Always Long	2.64 (0.37)	10.83	0.24	-0.65	8.10
$\eta(r_1)$	19.05*** (2.70)	10.77	1.77	1.13	7.75
$\eta(r_1, r_{12})$	16.79*** (3.19)	8.01	2.10	1.96	15.88

Table 8: Impact of Macro News Release

This table reports the impact of macro news releases on the predictability of the last half-hour market return. Panel A contrasts the results of regressing the last half-hour return (r_{13}) on the first and twelfth half-hour returns (r_1 and r_{12}) when there are macro news releases with those when there are no macro news releases. Panel B reports the profitability of timing the last half-hour market return using the first half-hour return, contrasting the days with certain macro news release with the days with no macro news release. The timing strategy is described in Table 4. We report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, and kurtosis. MCSI: Surveys of consumer confidence by University of Michigan release at 10:00 am Eastern time; GDP: monthly GDP estimate release at 8:30 am Eastern time; CPI: monthly release of CPI at 8:30 am Eastern time; FOMC: Federal Open Market Committee minutes release at 2:15 pm Eastern time. The returns are annualized and in percentage, and the regression coefficients in Panel A are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Panel A: Predictive Regression								
	No-Release	Release	No-Release	Release	No-Release	Release	No-Release	Release
	MCSI		GDP		CPI		FOMC	
Intercept	-1.70 (-1.15)	-7.16 (-1.21)	-1.72 (-1.17)	-6.75 (-0.94)	-1.93 (-1.31)	0.42 (0.06)	-1.49 (-1.03)	-12.6 (-1.61)
β_{r_1}	6.61*** (3.90)	14.4*** (3.40)	6.60*** (3.90)	11.7** (2.37)	6.63*** (3.90)	10.4* (1.95)	6.68*** (3.98)	14.4** (2.35)
$\beta_{r_{12}}$	11.9*** (2.64)	-5.51 (-0.48)	12.0*** (2.64)	-3.03 (-0.24)	11.4** (2.56)	11.7 (0.78)	10.9** (2.51)	34.1* (1.69)
R^2 (%)	2.6	5.5	2.7	3.0	2.5	5.0	2.5	11.0

Panel B: Impact of Macro News Release on Timing Performance							
	Macro News	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	
Non-Release	MCSI	6.05*** (3.83)	6.24	0.97	0.91	15.83	
Release	MCSI	19.09*** (3.41)	4.94	3.86	0.91	2.28	
Non-Release	GDP	6.28*** (4.01)	6.19	1.01	0.91	16.26	
Release	GDP	14.40** (2.08)	6.14	2.35	0.83	3.41	
Non-Release	CPI	6.10*** (3.88)	6.21	0.98	0.91	16.11	
Release	CPI	18.03*** (2.75)	5.80	3.11	0.90	3.84	
Non-Release	FOMC	6.24*** (4.01)	6.20	1.01	0.90	15.88	
Release	FOMC	20.04** (2.46)	5.84	3.43	1.07	7.22	

Table 9: Conditional Predictability

This table reports the results of regressing the last half-hour return (r_{13}) on the first half-hour return (r_1) and the twelfth half-hour return (r_{12}) of the day conditioned on the sign of the first half-hour return. Panels A, B, and C show results for three periods: the whole sample period, the financial crisis period from December 3, 2007, through June 30, 2009, and the periods excluding the financial crisis. The top panel reports the regression results when r_1 is positive, while the bottom panel reports the regression results when r_1 is negative. The returns are annualized and in percentage, and the regression coefficients are scaled by 100. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from February 1, 1993, through December 31, 2013, excluding days with fewer than 500 trades.

Predictor	Panel A			Panel B			Panel C		
	Whole Sample Period			Financial Crisis (12/2007–6/2009)			Excluding Financial Crisis		
	r_1	r_{12}	r_1 and r_{12}	r_1	r_{12}	r_1 and r_{12}	r_1	r_{12}	r_1 and r_{12}
When $r_1 > 0$									
Intercept	-8.85** (-2.52)	4.56** (2.41)	-8.47** (-2.50)	-14.2 (-0.82)	21.7* (1.82)	-14.4 (-0.91)	-5.70** (-2.16)	2.77 (1.64)	-5.62** (-2.15)
β_{r_1}	11.3*** (3.63)		10.5*** (3.58)	21.0** (2.15)		17.8** (2.02)	7.41*** (3.60)		7.19*** (3.58)
$\beta_{r_{12}}$		18.4*** (2.97)	17.2*** (2.85)		41.9*** (3.14)	39.5*** (2.98)		7.49 (1.52)	6.93 (1.44)
R^2 (%)	2.3	2.6	4.5	4.5	11.0	14.1	1.1	0.4	1.5
When $r_1 < 0$									
Intercept	-1.07 (-0.28)	-8.27*** (-3.44)	-0.83 (-0.21)	-5.07 (-0.25)	-26.2** (-1.97)	-2.82 (-0.13)	-2.77 (-0.87)	-5.90*** (-2.81)	-2.75 (-0.86)
β_{r_1}	5.72* (1.73)		5.90* (1.78)	9.74 (1.19)		10.5 (1.26)	2.73 (0.89)		2.79 (0.91)
$\beta_{r_{12}}$		6.60 (1.06)	6.93 (1.11)		7.78 (0.58)	8.98 (0.66)		5.45 (1.16)	5.53 (1.16)
R^2 (%)	0.5	0.3	0.9	0.7	0.0	0.9	0.1	0.2	0.3

Table 10: Market Timing with Transaction Costs

This table reports the economic value of timing the last half-hour market return using the first half-hour return or combining with the twelfth half-hour return, incorporating the transaction costs due to the bid–ask spread. The timing strategy is described in Table 4. The benchmark *Always Long* is to always invest in the market during the last half-hour on each trading day, and the benchmark *Buy-and-Hold* is to buy and hold the market on a daily basis. For each strategy, we report the average return (*Avg Ret*), standard deviation (*Std Dev*), Sharpe ratio (*SRatio*), skewness, kurtosis, and M2 measure, which is the average return of the strategy with volatility leveled up to be the same as the volatility of the daily *Buy-and-Hold* strategy. Panel A is for the period after decimalization (after July 1, 2001), and Panel B is for the period when the spread is stabilized (after January 1, 2005). The returns are annualized and in percentage. Newey and West (1987) robust *t*-statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively.

Timing	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	M2(%)
Panel A: After July 1, 2001						
$\eta(r_1)$	4.46*** (2.58)	6.10	0.73	1.21	19.82	14.88
$\eta(r_1, r_{12})$	4.30*** (3.44)	4.40	0.98	2.58	40.65	19.87
Always Long	-0.74 (-0.42)	6.12	-0.12	-0.53	20.06	-2.45
Buy-and-Hold	4.90 (0.85)	20.34	0.24	-0.17	8.07	
Panel B: After January 1, 2005						
$\eta(r_1)$	6.52*** (3.00)	6.51	1.00	1.42	20.48	20.77
$\eta(r_1, r_{12})$	4.74*** (3.01)	4.72	1.00	2.89	41.10	20.82
Always Long	-1.03 (-0.47)	6.54	-0.16	-0.54	20.78	-3.25
Buy-and-Hold	6.75 (0.98)	20.72	0.33	-0.26	9.78	

Table 11: Summary of Other ETFs

This table describes the ten index ETFs used for the robustness analysis in Table 12. These ETFs are the most heavily traded ETFs as measured by their average daily trading volume from their inception dates to December 31, 2013.

Symbol	Name	Inception
QQQ	Powershare NASDAQ 100	03/10/1999
XLF	Financial Select Sector SPDR	12/22/1998
IWM	iShares Russell 2000 ETF	05/26/2000
DIA	Dow Jones Industrial Average ETF	01/20/1998
EEM	iShares MSCI Emerging Markets ETF	04/11/2003
FXI	iShares China Large-Cap ETF	10/8/2004
EFA	iShares MSCI EAFE ETF	08/17/2001
VWO	Emerging Markets ETF	03/10/2005
IYR	iShares U.S. Real Estate ETF	06/19/2000
TLT	20+ Year Treasury Bond ETF	07/26/2002

Table 12: Out-of-Sample Portfolio Performance – Other ETFs

This table reports the average return (*Avg Ret*), standard deviation (*Std Dev*), in-sample R^2 (*INS R²*), out-of-sample R^2 (*OOS R²*), and *CER*, with the same analysis as Table 6 except replacing the market return by one of ten most heavily traded ETFs. All quantities are in percentage, and returns and standard deviations are annualized. Panel A reports the results using the first half-hour return (r_1) to forecast, and Panel B reports the results using both r_1 and r_{12} to forecast. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period for each ETF is from its inception date to December 31, 2013, excluding days with fewer than 100 trades.

Fund	Avg Ret(%)	Std Dev(%)	INS R^2 (%)	OOS R^2 (%)	CER(%)	Avg Ret(%)	Std Dev(%)	INS R^2 (%)	OOS R^2 (%)	CER(%)
	Panel A: $\beta_1 r_1$					Panel B: $\beta_1 r_1 + \beta_2 r_{12}$				
QQQ	7.75*** (3.65)	7.89	1.43	0.70	7.38	8.34*** (3.83)	8.08	2.26	0.50	7.96
XLF	12.04*** (4.36)	9.95	3.64	3.55	12.44	8.73*** (3.24)	9.70	4.37	2.19	9.13
IWM	11.72*** (5.18)	7.70	2.51	2.43	11.72	12.12*** (4.45)	9.26	4.53	3.81	12.09
DIA	3.46** (2.35)	5.69	1.16	1.03	4.16	4.63*** (2.79)	6.40	2.25	1.81	5.31
EEM	14.76*** (4.91)	9.01	8.54	6.53	14.69	18.46*** (6.01)	9.20	13.27	10.43	18.38
FXI	18.42*** (5.20)	10.17	7.80	5.90	17.71	15.98*** (4.35)	10.54	10.42	7.52	15.26
EFA	7.45*** (4.16)	5.82	3.53	1.90	7.18	6.53*** (3.69)	5.76	4.79	1.43	6.27
VWO	12.18*** (3.76)	8.72	5.72	4.39	12.12	13.61*** (4.15)	8.83	8.45	6.29	13.55
IYR	24.22*** (5.86)	12.29	5.29	4.60	14.98	29.80*** (6.43)	13.78	11.77	9.82	20.52
TLT	4.03*** (4.32)	2.89	1.77	1.65	2.26	4.50*** (5.14)	2.71	1.81	1.51	2.73

Table 13: Out-of-Sample Portfolio Performance – International

This table reports the results of using international index futures. Panel A reports the in-sample predictive regressions results, while Panel B reports the out-of-sample predictive rolling regression results. Panel C reports the performance of the mean-variance analysis. Returns and standard deviations are annualized. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively. The sample period is from July 1, 2003 through October 31, 2015.

Predictive Regression	Panel A: In-Sample			Panel B: Out-of-Sample		
	β_{r_1}	$\beta_{r_{16}}$	$R^2(\%)$	β_{r_1}	$\beta_{r_{16}}$	$R^2(\%)$
FTSE 100 Index Futures						
r_1	4.93*** (7.73)		1.90	4.38*** (25.15)		1.87
r_1 and r_{16}	4.93*** (7.71)	-0.01 (-0.01)	1.90	4.38*** (24.40)	-1.09*** (-4.41)	1.39
EuroStockXX 50 Index Futures						
r_1	3.02*** (5.17)		0.84	3.12*** (22.08)		0.67
r_1 and r_{17}	2.96*** (5.05)	2.75* (1.65)	0.93	3.05*** (21.56)	3.59*** (22.84)	0.60
Panel C: Mean-Variance Performance						
Variable	Avg Ret(%)	Std Dev(%)	SRatio	Skewness	Kurtosis	CER(%)
FTSE 100 Index Futures						
\bar{r}_{17}	-0.72 (-1.02)	2.48	-0.29	-1.54	19.43	-0.73
$\beta_1 r_1$	1.84 (1.49)	4.31	0.43	0.68	13.11	1.82
$\beta_1 r_1$ and $\beta_2 r_{16}$	1.67 (1.35)	4.29	0.39	0.79	11.69	1.65
EuroStockXX 50 Index Futures						
\bar{r}_{18}	0.28 (0.16)	5.97	0.05	-0.28	5.69	0.24
$\beta_1 r_1$	2.47* (1.72)	5.04	0.49	-0.09	9.42	2.45
$\beta_1 r_1$ and $\beta_2 r_{17}$	3.03** (2.14)	4.96	0.61	-0.09	9.76	3.00

Table 14: Out-of-Sample Portfolio Performance – Other Assets

This table reports the same analysis as in Table 12 except that the underlying asset is one of the 9 currency pairs, Australia, Euro, UK, New Zealand, Canada, Switzerland, Japan, Norway and Sweden versus the US dollar, and two commodity futures, crude oil and gold. All quantities are in percentage, and returns and standard deviations are annualized. Panel A reports the results using the first half-hour return (r_1) to forecast, and Panel B reports the results using both r_1 and r_{12} to forecast. Newey and West (1987) robust t -statistics are in parentheses, and significance at the 1%, 5%, or 10% level is given by an ***, an ** or an *, respectively.

Fund	Avg Ret(%)	Std Dev(%)	INS R^2 (%)	OOS R^2 (%)	CER(%)	Avg Ret(%)	Std Dev(%)	INS R^2 (%)	OOS R^2 (%)	CER(%)
	Panel A: $\beta_1 r_1$					Panel B: $\beta_1 r_1 + \beta_2 r_{12}$				
AUDUSD	1.34 (1.51)	2.45	0.69	0.31	1.57	3.13*** (3.02)	2.87	5.32	0.30	3.36
EUROUS	0.66 (1.43)	1.31	0.34	0.05	0.87	0.50 (1.07)	1.33	2.82	1.44	0.71
GBPUSD	0.46 (1.11)	1.19	0.14	0.02	0.24	0.46 (1.01)	1.30	2.39	0.45	0.23
NZDUSD	-0.34 (-0.35)	2.57	0.03	-0.24	-0.05	1.84 (1.54)	3.13	4.04	0.73	2.12
USDCAD	-1.26** (-2.19)	1.63	-0.01	-0.77	-1.12	-0.33 (-0.56)	1.71	0.18	-1.32	-0.19
USDCHF	0.43 (0.91)	0.94	0.20	-0.36	0.10	0.39 (0.74)	1.03	0.29	-0.31	0.05
USDJPY	0.75* (1.66)	1.30	0.82	0.28	0.86	0.66 (1.51)	1.25	1.69	0.16	0.77
USDNOK	0.65 (1.54)	0.88	0.04	0.01	0.14	0.40 (0.59)	1.41	1.88	0.25	-0.11
USDSEK	1.08* (1.66)	1.36	0.25	0.09	0.81	-0.09 (-0.13)	1.54	2.05	0.13	-0.36
OIL	1.05 (0.86)	3.10	0.01	-0.06	1.60	0.73 (0.62)	2.98	0.09	-0.78	1.28
GOLD	-0.20 (-0.15)	3.31	0.05	-0.20	0.71	2.54** (2.03)	3.07	3.33	2.82	3.45