Underreaction, Overreaction, and Dynamic Autocorrelation of Stock Returns

Hongye Guo*

December 16, 2019

Abstract

I document that in the US, the aggregate monthly stock returns correlate positively with past returns 2/3 of the time, and negatively 1/3 of the time. While the two arms of correlation are separately strong, they cancel with each other, leading to an average autocorrelation that is only weakly positive. I argue this pattern of aggregate return predictability will be generated if investors fail to see the time-varying autocorrelation structure of earnings news. In this model, investors act as if they have underreacted to past news 2/3 of the time, and overreacted to past news 1/3 of the time. I then look out-of-sample and find affirmative evidence in the cross section and the international stock markets. The paper shows that the traditional view on stock return autocorrelation misses important information, which is that it varies over time.

*Department of Finance, The Wharton School, University of Pennsylvania. Email: hoguo@wharton.upenn.edu. I thank my advisors, Jessica Wachter and Itamar Drechsler, for very helpful comments and guidance over the course of the project. I thank Cathy Schrand for very helpful comments on the accounting literature. I thank participants of the 2019 Yale Summer School in Behavioral Finance and especially to the organizer, Nick Barberis, for an excellent education on behavioral finance. I thank Nick Barberis, Jules van Binsbergen, John Campbell, Sylvain Catherine, Alice Chen, Vincent Glode, Craig Mackinlay, Michael Schwert, Nick Roussanov, Robert Stambaugh, and Xiao Han for helpful comments. All errors are mine.

1 Introduction

Autocorrelation of the aggregate stock market returns has been extensively studied. At the monthly frequency, the classic work is Poterba and Summers (1988). Using a variance ratio test, the authors show, among other things, that monthly market returns in the US have a small, insignificant positive autocorrelation over the horizon of 12 months. This weak autocorrelation can also be confirmed by regressing monthly stock returns on lagged returns over the past year. An important addition to this literature in the past decade is Moskowitz et al. (2012), which shows that in the international data, equity market index future excess returns exhibit strong positive autocorrleation over a look-back window of 12 months.

This paper first shows that these traditional views on stock market return autocorrelation miss an important feature, which is that it varies over time. I show that stock market returns correlate negatively with past returns when fresh earnings news comes out, and positively when old earnings news comes out. These positive and negative autocorrelation episodes correspond to the first half and the second half of the earnings reporting cycle, respectively. Given the stability of reporting cycles over time, they are set to be fixed months of the year within a country, and therefore could easily have been anticipated in advance.

Next, this paper shows similar empirical results in the cross section of industry returns [1]. Continuation in the cross section of the stock returns, as in Jegadeesh and Titman (1993), has also been extensively studied. Unlike the weak continuation found in the US aggregate market, momentum in excess stock returns is a much stronger and more robust effect (e.g., Asness et al. (2013)), and the industry component is shown to drive a large fraction of it (Moskowitz and Grinblatt (1999)). This paper shows

[1] For the record it also exists in the cross section of stock returns, though it seems to be operating mainly through the industry component. There could be, however, other components on which the pattern exists.

that the strength of the continuation in the cross section of industry returns also varies over time. Similar to the results found in the aggregate market, industry momentum is much stronger when old earnings news comes out, and virtually non-existent when fresh earnings news comes out. Moreover, I show that a similar pattern has been seen in country/territory momentum, and in country-/territory-industry momentum. The return predictability results of this paper are motivated by, but not constrained to, the time-series setting.

While the dynamic autocorrelation of stock returns is empirically interesting in its own right, it is also important to study the underlying reasons for it. In doing so, I connect with the behavioral finance literature and hope to contribute to it. In this literature, a large number of papers focus on the notions of under-, and over-reaction (Barberis (2018)). It is then natural to ask under what circumstances should we observe each. While this question is important, relatively few papers provide an answer, perhaps because it is not easy to construct one model that features both under- and over-reaction. Despite the challenge, Barberis et al. (1998) builds a model that successfully achieves exactly that. Starting with an earnings process that follows a random walk, it shows that if investors incorrectly believe that the autocorrelation structure of this earnings process is dynamic—specifically, follows a two-state regime-switching model featuring continuation and reversal—then they will overreact to news that seems to be in a sequence, and underreact to news that seems not.

This paper also speaks to this under-explored question. Contrary to Barberis et al. (1998), this paper relies on the earnings process to actually have time-varying autocorrelation. I first show empirically that the autocorrelation structure of the aggregate earnings news in the economy, as measured by the aggregate return on equity (ROE) change, is indeed dynamic in real, calendar time at the monthly frequency. Specifically, I show that while earnings news exhibit strong positive autocorrelation with past earnings news on average, such autocorrelation drops in the first half of the earnings reporting cycle. I then show in a stylized model that if investors incorrectly believe this autocorrelation is constant, they will exhibit underreaction when such autocorrelation is high, and overreaction when such autocorrelation is low. Overall, the theme is that investors act as if they have overreacted/underreacted to past news when fresh/old earnings news is coming out. These mechanisms will generate the aforementioned autocorrelation pattern in stock returns.

It is worth noting that while the broader logic behind such mechanisms is somewhat new in the literature, the paper is not the first to employ it. Specifically, Matthies (2018) finds that beliefs about covariance exhibit compression towards moderate values. He documents three pieces of supportive evidence: 1) natural gas and electricity futures exhibit moderate covariance despite persistent heterogeneity in the fundamental relation in the spot market; 2) macroeconomic forecasts made by professional forecasters exhibit predictable errors; and 3) participants in an experiment overestimate the stock market's low covariance with macroeconomic variables and compress covariances between individual stock returns towards moderate values. Behind Matthies (2018) and my paper is a particular bounded-rationality mechanism where investors' limited cognitive capacity prevents them from fully exploring the heterogeneity of a parameter, lending them to simply use a moderate representative value instead.

This paper also falls into the broader literature that studies the interaction between earnings announcements and stock returns (e.g., Beaver (1968), Bernard and Thomas (1990), Bernard and Thomas (1989)). An important piece of recent work in this area is Savor and Wilson (2016), which focuses on weekly stock returns. The authors first confirm that stocks have high returns on earnings announcement-week (as in Beaver (1968)), and additionally show that stocks that have high announcement week returns in the past are likely to have high announcement-week returns in the future. Among other things, the authors also show that early announcers earn higher returns than late announcers, and firms that are expected to announce in the near-term future have higher betas with respect to the announcing portfolios. Overall, the authors make a convincing case that earnings announcements of individual firms resolve systematic risks that have implications on the broader market.

Instead of the risks associated with earnings announcements, my paper focuses on the under- and over-reaction that are potentially related to them, as well as the resulting lead-lag relationship of stock returns. Also, instead of the returns to the portfolio that long the announcing firms and short the non-announcing firms, I focus on the aggregate market returns or the industry-level returns in excess of the market, neither of which strongly correlate with the spread between the announcing and non-announcing portfolio. In additional to these philosophical distinctions, specific difference in empirical results will be further discussed later in the empirical section.

This paper also relates to the broad literature studying the seasonality of stock returns, documented by Heston and Sadka (2008) and extended by Keloharju et al. (2016). This literature also studies the autocorrelation of stock returns, and makes the point that full-year lags have especially strong predictive power, which is a distinction of the independent variable. The main point of my paper, however, is that the predictive power of past returns is different according to the timing of the dependent variable. Philosophically, Heston and Sadka (2008) and Keloharju et al. (2016) are consistent with the notion of stationarity of stock returns, while my paper challenges it—specifically, the notion that the autocorrelation coefficients depend on displacement and not time. Again, specific distinctions will be further discussed in the empirical section.

The rest of the paper is structured as follows: section 2 motivates the analysis by demonstrating the dynamic autocorrelation structure of the aggregate market structure in the US. Section 3 provides the intuition behind those results, and substantiates those intuitions using fundamental data. Section 4 provides a simple stylized model with closed-form solutions that qualitatively illustrate the intuitions in section 3. Section 5 demonstrates dynamic autocorrelation structure in the cross section of 1) US industry returns, 2) global aggregate market returns, and 3) global industry returns. Section 6 demonstrates the forecasting usefulness of the US time-series results in real time. Section 7 concludes.

2 Dynamic Autocorrelation of the US Aggregate Market Returns

In this section, I first describe the earnings reporting cycle in the US and then demonstrate that the autocorrelation of the US aggregate market returns varies strongly with this cycle. For ease of expression, I first define three groups of months: "Group 1" contains January, April, July, and October, "Group 2" February, May, August, and November, and "Group 3" March, June, September, and December. In the US, I will call group 1 months the "newsy" months. The reason why they are called newsy is they are when *fresh* news on firm earnings come out the most *intensively*. The news is fresh because group 1 months immediately follow the end of the fiscal quarters, the majority of which are aligned with the calendar quarters. This is demonstrated in Table 1, which shows that in the US, about 85% of fiscal quarters end in the group 3 months. Moreover, Table 2 shows that in the US, about half of the firms report within one month after the end of a fiscal period. In fact, among all of the three types of months, most firms report in group 1 months. Therefore, group 1 months are when fresh news is reported intensively. The two features of freshness and intensity are the concrete meanings of the word "newsy." The bottom line is that in terms of the earnings news, in the US, group 1 months are the information relevant months.

Having said that, we look at Table 3, which reports results of the following monthly time-series regression that predicts the aggregate US stock market return: $mkt_t = \alpha +$

$\sum_{j=1}^{8} \beta_j mkt_{nm(t,j)} + \epsilon_t$. Here $mkt_{nm(t,j)}$ is the j th "newsy" month return strictly before the month t. Throughout my empirical analysis, returns on those newsy months are put on the right-hand side of the regressions. Figure 1 thoroughly illustrates how lagging is done on the regression: Suppose the dependent variable is return of November, then lag 1 newsy month (abbrv. lag 1nm) return is that of October, lag 2nm return is that of July and so on. As the dependent variable moves forward to December and January of the next year, the lagged newsy month returns on the right-hand side stay the same. However, when the dependent variable becomes return of February, the lag 1nm return will be moved forward by three months to January, as that is the most recent newsy month strictly before February.

Having clarified the specifics of the regressions we move to the results. Column 1 of Table 3 confirms the conventional view that the aggregate stock market exhibits only weak momentum with a look-back window of one year. Column 2 does the same regression, but only on the 1/3 of the sample where the dependent variables are returns of the newsy months. This column shows that in newsy months, returns are decidedly negatively correlated with past newsy-month returns. Column 3 does the regression for the rest of the sample, where the dependent variables are returns of non-newsy months. In those months, returns are positively correlated with past newsy-month returns. However, on average they cancel each other out, resulting in the weak unconditional autocorrelations shown in column 1.

The main empirical finding of this paper is that autocorrelation structure of stock returns varies by the timing of the dependent variable. Column 4 delivers this main point by showing the difference in the coefficients of columns 2 and 3. While the differences are not monotonic with lags, they are clearly all negative, and overall the effect seems stronger the smaller the lag. To evaluate the strength of the effect in different contexts, such as different historical periods, it will be useful to have one coefficient instead of eight. Throughout the empirical section, I will use the sum of the first four lags, or $\sum_{j=1}^{4} mkt_{nm(t,j)}$, as the flagship signal. Four lags and equal weighting are both choices I made. Since four lags correspond to the typical one-year look-back window of the various price momentum strategies, when I get to the cross section these choices will enable me to benchmark against those strategies and speak to when they work and don't work.

Table 4 focuses on the following regression: $mkt_t = \alpha + \beta \sum_{j=1}^{4} mkt_{nm(t,j)} + \epsilon_t$. Here $\sum_{j=1}^{4} mkt_{nm(t,j)}$ is the sum of the lag 1 to lag 4 newsy-month returns, the said flagship signal. Its coefficient here will roughly correspond to the average of the first four coefficients in Table 3. Column 1 shows the weak unconditional time-series momentum, pushed over the p-value cutoff of 10% by putting only the newsy-month returns on the right-hand side (regressing on the past 12-month return will result in a t-stat of 0.64). Columns 2 and 3 are the subsamples for newsy months and non-newsy-month returns. Again we see strong negative autocorrelation in newsy months and strong positive autocorrelations in non-newsy months. The difference in the coefficients, -0.233, is shown in the interaction term of column 4 with a t-stat of -4.87. These two values indicate that in the full CRSP sample of 1926-2019, the autocorrelation structure of US aggregate stock return is strongly time varying. Columns 5-7 show that these results are strong in the post-WWII period, first half, and second half of the sample, though the effect is stronger in the first half of the sample. Column 8 uses the alternative specification where group 2 months are considered part of the newsy "months." Here, the four newsy "months" of the year will be the four two-month periods of Jan+Feb, Apr+May, Jul+Aug, and Oct+Nov. While the results do seem to be present and in fact strong under this specification, column 8 seems substantially weaker than column 4 both economically and statistically. The return data seem to suggest that group 2 months should be considered non-newsy in the US.

3 Intuition and Earnings News Predictability

3.1 Intuition

In the past section we saw that in the US, the aggregate stock market returns have lower autocorrleation in newsy months, and higher autocorrelation in non-newsy months. In this section, we first provide some intuitions on the underlying reasons, and then substantiate the premises behind these intuitions using fundamental data.

Consider the following narrative: suppose investors see some great earnings-related news in April from the reporting firms, who are the early reporters for Q1. Then one may (correctly, as we will show in the next subsections) think that the news in the upcoming months is also great. In May and June that agent likely will indeed hear good news. This is because the reporting firms in those months are also reporting on Q1, and the news of those earnings reports shares a common time component—perhaps the first quarter is just a good time for everyone. After the streak of good news, investors may plausibly think the news will keep being good. However, in July, earnings of Q2, a different fiscal quarter, will be reported. The shared time component therefore disappears and news in July is less likely to resemble that in April. If the investors fail to anticipate this drop in earnings news autocorrelation, they will be disappointed in July. On the other hand, if they do not realize that the good news in May and June is "mechanically" driven by the shared time component, they may be positively surprised in those months.

Taking a step back from the narrative, the intuition here is that earnings news in the newsy months will have lower correlation with past earnings news, since reporters in those months are reporting on a new fiscal quarter. This makes the news in newsy months further away from past news in terms of fiscal time, holding constant the distance in calendar time. On the other hand, earnings news in non-newsy months will have higher correlation with past news. Imagine an investor trying to forecast the earnings news in the next months using past earnings news, and assume this investor treats the upcoming month as an average month, i.e., he or she makes no distinction between the newsy and non-newsy months in the forecasting practice. When past news has been good, in the upcoming newsy months this investor is likely to be disappointed, as if he or she has overreacted to past news. In the upcoming non-newsy months he/she is likely to see positive surprises, as if he/she has underreacted to past news.

The distinction between fiscal and calendar time is illustrated in Figure 2. For example, focusing on news in the lag 1 calendar quarter: when the forecast month is group 1, the average fiscal time between the news in the forecast month and that in the past calendar quarter is 1 fiscal quarter. When the forecast month is group 2, the average distance shortens to 2/3 fiscal quarter, and for group 3 it is only 1/3 fiscal quarter. Note this relationship holds for more than one lagged calendar quarter. The overarching message is that when investors are trying to forecasts the earnings news in the upcoming months, past information is more/less timely the later/earlier in the earnings reporting cycle.

3.2 US Fundamental News in Calendar Time

In this subsection we map the intuition in the previous subsection to fundamental data, and show that fundamental news indeed has lower autocorrelation earlier in the earnings reporting cycle.

First, we need a measure of the previously mentioned "earnings news," and ideally without using return information: It would be the most helpful to explain the pattern in stock returns without using stock returns. The measure I choose is the change in (trailing) aggregate ROE, which is aggregate quarterly earnings divided by aggregate book value of equity, less this ratio a year ago. Earnings news should be about earnings, and when aggregate earnings are high, the news is good. However, aggregate earnings is non-stationary, and needs to be scaled by something. I follow Vuolteenaho (2002) to use book value of equity, which could be thought of as a smoothed earnings measure over a long look-back window. This is because of the "clean accounting assumption," which holds reasonably well in reality (Campbell (2017)). Subtracting the ROE one year ago is a simple way to de-seasonalize and create a measure that is conceptually similar to earnings growth. This measure echoes with Ball and Sadka (2015), which highlights the importance of studying the aggregate earnings. However, this is not a perfect choice and reasonable people can use other measures. [2] Objections could be raised against this choice such as it is not compared against expectations data, or it is not as flexible as returns (not all news in the earnings report is captured by the earnings), or that subtracting the ROE four quarters ago is going to induce a mechanical negative autocorrelation in the fourth lag. However, despite those flaws, change in ROE does seem to be a simple and reasonable manifestation of earnings news, with which I can quantitatively illustrate the intuition in the previous section.

Having picked the specific earnings growth measure, I construct this measure for each calendar month among all the firms that report in the month. The ROE measure is the total earnings of the reporting firms divided by the total book value of equity. The change of ROE is computed as this month's ROE subtracted by that of 12 months ago. This computation reflects the earnings news investors receive in those calendar months.

Column 1 of Table 6 performs a simple multiple regression of change in ROE on the past three month's change in ROE, and the results are both simple and sensible: the coefficients are all significantly positive and monotonically declining with lags. However, columns 2-4 show that there are important complications underneath this simple and sensible result. These three columns split the regression in column 1 into

[2] Incidentally, the aggregate earnings growth measure itself will be badly behaved and cannot be used. This is because of the bad-divisor problem: aggregate earnings at the quarterly frequency can approach zero or go below it in the US and do so even more frequently globally.

three subsamples based on the months of the dependent variable: group 1, 2, and 3 months, where group 1 months are the newsy months (in the US) described earlier. Looking at the first row across the three columns, you see that in the newsy months, earnings news has a much lower first-order autocorrelation with the earnings news of the last month (0.005) compared to group 2 (0.622) and group 3 months (0.749). There are two reasons behind this pattern: first, in the newsy months firms are reporting on a new quarter, which is further away from past news in fiscal time—in fact, if you only look at one monthly lag it is guaranteed to be the same quarter. Second, in the newsy months, lag 1 month is the group 3 months where, in the US, a small fraction of firms report. The aggregate ROE changes therefore are substantially noisier than other times, which pulls the coefficients towards 0. This second point can be easily seen whenever group 3 months are on the right-hand side: lag 1 month for group 1 dependent variable (column 2), lag 2 month for column 3, and lag 3 month for column 4. They are the smallest coefficients in their respective rows. Columns 5-8 do similar regressions where the explanatory variable is the sum of the three lagged ROE changes. Here group 3 months account for 1/3 of the right-hand side across all columns. It is clear that in group 1 months earnings news has much lower loading on past earnings news than the average number in column 5. In the other two types of months, the loading is higher than average, though in group 2 months it is only slightly higher.

What if the investors do not consider the complications in column 2-4 and 6-8, and use only the simple and static results in columns 1 and 5? In that case, when past earnings news is good/bad, you would see negative/positive surprises in group 1 months, as if you had overreacted to the news. In group 2 and group 3 months you would see positive/negative surprises, as if you had underreacted to past news. Hence, not seeing the dynamic autocorrelation structure of earnings news can lead to dynamic autocorrelation structure in the market returns.

Note that in this framework it is not quite accurate to say that people underreact to news in the newsy months. If return continues over some horizons after the newsy months, one can characterize investors' initial reactions in the newsy months as underreactions. This indeed seems the case if one look at horizons that are less than two months, but in the third month the return reverses and the continuation weakens substantially. The main point of this paper is the *difference* in return autocorrelations between newsy and non-newsy months, not the sum. I have summarized the underlying mechanism as that people act as if they have overreacted/underreacted to past news in newsy/non-newsy months. Notice this is not a simple underreaction or overreaction story, but instead is a framework containing both features.

4 Stylized Model

In this section we build a stylized model delivering the intuition stated earlier in the section. The model is inspired by Guo and Wachter (2019). Consider an infinite-horizon discrete-time economy with risk-neutral investors. Let D_t denote the aggregate dividend at time t, and $d_t = \log D_t$. Assume that investors believe:

$$\Delta d_{t+1} = m \sum_{j=0}^{\infty} \rho^j \Delta d_{t-j} + u_{t+1} \tag{1}$$

$$u_t \overset{iid}{\sim} N(0, \sigma_u), \quad \forall t \tag{2}$$

And more generally, for all $i \geq 1$:

$$\Delta d_{t+i} = m \sum_{j=0}^{\infty} \rho^j \Delta d_{t+i-1-j} + u_{t+i} \tag{3}$$

$$u_t \overset{iid}{\sim} N(0, \sigma_u), \quad \forall t \tag{4}$$

In other words, the investors extrapolate an exponentially weighted moving average (EWMA) of past cash flow growth. Here ρ is the decay parameter in the EWMA, while m controls the degree of extrapolation. They lie between 0 and 1. Denote $x_t = \sum_{j=0}^{\infty} \rho^j \Delta d_{t-j}$, so that investors expect $\Delta d_{t+1} = m x_t + u_{t+1}$. Notice x_{t+1} can be recursively written as:

$$x_{t+1} = \Delta d_{t+1} + \rho x_t \tag{5}$$

This recursive relationship does not involve the investors' beliefs yet. Now given the investors' beliefs of future cash flow growth, it follows that they believe the following process of x going forward:

$$\begin{aligned} x_{t+1} &= \Delta d_{t+1} + \rho x_t & (6)\\ &= m x_t + u_{t+1} + \rho x_t & (7)\\ &= (m+\rho) x_t + u_{t+1} & (8) \end{aligned}$$

A similar relationship applies beyond period $t+1$. Notice $m+\rho$ needs to be less than 1 for the process of x_t to be stationary in the investors' minds.

While the investors use a constant extrapolation parameter m in their beliefs, in reality the process is driven by a dynamic parameter that takes values h and l in alternating periods, where $l < m < h$:

$$\Delta d_{t+1} = \begin{cases} h x_t + u_{t+1}, \text{ where } t \text{ is even} \\ l x_t + u_{t+1}, \text{ where } t \text{ is odd} \end{cases}$$

This reduced-form setup is motivated by the empirical ROE dynamics described in previous sections. While such dynamics are caused by heterogeneity in reporting lag among firms that share a common time component in earnings news, I do not model this particular mechanism. In principle, any mechanism that causes time-varying autocorrleation in earnings-related news can be mapped to this model. Given this cash flow process in reality, x_{t+1} actually follows the process:

$$x_{t+1} = \rho x_t + \Delta d_{t+1} = \begin{cases} (h+\rho)x_t + u_{t+1}, \text{ where } t \text{ is even} \\ (l+\rho)x_t + u_{t+1}, \text{ where } t \text{ is odd} \end{cases}$$

Having set up the investors' beliefs and how they deviate from reality, we now compute the equilibrium valuation ratio, which requires solely the beliefs, and the equilibrium equity returns, which require both the beliefs and the reality. Denote the current dividend on the aggregate market D_t. Let P_{nt} be the price of an equity strip that expires n periods away. Define:

$$F_n(x_t) = \frac{P_{nt}}{D_t} \tag{9}$$

We now show that $F_n(x_t)$ is indeed a function of x_t, our state variable representing past cash flow growth. Notice $F_n(x_t)$ must satisfy the following recursive relationship:

$$F_n(x_t) = E_t[\delta F_{n-1}(x_{t+1})\frac{D_{t+1}}{D_t}] \tag{10}$$

Where δ is the time discounting parameter of the investors. Conjecture $F_n(x_t) = e^{a_n + b_n x_t}$. Substitute the conjecture back into equation 10 and take the log of both sides:

$$a_n + b_n x_t = \log\delta + a_{n-1} + b_{n-1}(m+\rho)x_t + m x_t + \frac{1}{2}(b_{n-1}+1)^2\sigma_u^2 \tag{11}$$

Which leads to the following recursive relationship for a_n and b_n:

$$a_n = a_{n-1} + \log\delta + \frac{1}{2}(b_{n-1}+1)^2\sigma_u^2 \tag{12}$$

$$b_n = b_{n-1}(m+\rho) + m \tag{13}$$

Notice equation 13 along with the boundary condition of $b_0 = 0$ implies the solution:

$$b_n = \frac{1-(m+\rho)^n}{1-m-\rho}m \tag{14}$$

And a_n can be pinned down accordingly. Having solved for the valuation ratios of an equity strip that expire n periods away, we bring in the actual cash flow process to compute its return $R_{n,t+1}$. For even t, notice the log return needs to follow:

$$\begin{aligned}
\log(1+R_{n,t+1}) &= \log(\frac{F_{n-1}(x_{t+1})}{F_n(x_t)}\frac{D_{t+1}}{D_t}) &(15)\\
&= a_{n-1} - a_n + b_{n-1}x_{t+1} - b_n x_t + hx_t + u_{t+1} &(16)\\
&= a_{n-1} - a_n + (b_{n-1}+1)(hx_t + u_{t+1}) + (b_{n-1}\rho - b_n)x_t &(17)\\
&= a_{n-1} - a_n + (b_{n-1}+1)(h-m)x_t + (b_{n-1}+1)u_{t+1} &(18)
\end{aligned}$$

Similarly, for odd t, we have:

$$\log(1+R_{n,t+1}) = a_{n-1} - a_n + (b_{n-1}+1)(l-m)x_t + (b_{n-1}+1)u_{t+1} \tag{19}$$

Two points are worth noting: First, in returns there is an unpredictable component $(b_{n-1}+1)u_{t+1}$ that is completely driven by the unpredictable component in cash flow growth. Cash flow growth therefore will correlate positively with contemporaneous returns. Second, there is a predictable component that takes alternate signs. Hence, in the months where cash flow growth has high/low correlation with past growth, as represented by x_t, past cash flow growth positively/negatively forecasts the return. Given the contemporaneous correlation between return and cash flow growth, past returns would also positively/negatively forecast current return in the high/low correlation months.

It is important to note that this is a stylized model in which the "cash flow growth" can and should be given broader meaning if mapped to reality. Right now, x_t contains only past cash flow growth. In principle, any news that is revealed in the financial reports, e.g., launching of new products, acquisition of new assets, etc., can go into it. This news may have yet to affect cash flow, but investors may believe that they can affect future cash flow, potentially in the long run. If this news has a time component across firms, then it may contribute to the dynamic autocorrelation of stock returns without showing up in the dynamic autocorrelation pattern of earnings growth. The analysis of ROE is still a useful exercise for the theory, especially because it does not use return data. It is however important to note that the theory is not confined to fundamentals extrapolation, with the fundamentals narrowly defined as past cash flow growth.

5 Cross Sectional Analysis

5.1 US Cross Section

The main intuition behind the theory is the following: if a group of fundamentally connected stocks 1) have very similar fiscal period ends and 2) report progressively with sufficient heterogeneity of reporting lag, then earnings news will correlate much less strongly with past news when fresh news is coming out, i.e., in the newsy months. Failure to see this time-varying autocorrelation structure of earning news will result in more negative return autocorrelation in newsy months, and more positive return autocorrelation in the non-newsy months. Notice while it is very natural to apply this story to the aggregate stock market, it may additionally apply to the cross section of industry-level stock returns. Two random stocks in each industry-country are more connected with each other than two random stocks drawn from the same country. News and returns of each industry are often compared against those of the aggregate market via headlines like "Tech Leads the Market on Good Earnings".

In this section, I test whether industry excess return, or industry return subtracting the market return, exhibits similar dynamic autocorrelation structure to the aggregate stock market documented in the previous section. It is important to use industry excess return and not just the industry return to test this relationship. This is because the industry returns will all correlate very strongly with the aggregate market returns. Hence, using the industry return directly will make it difficult to see the additional effect on top of what is found for the aggregate market. Using the excess returns will help show this additional effect more directly.

Table 7 shows the regression results of industry excess returns on past newsy month industry excess returns. While the structure of this table looks similar to that in Table 3, they differ in two aspects: First, Table 7 is on an industry-month panel instead of a monthly time series. Second, Table 7 uses industry excess returns as opposed to the aggregate market return. What is used in Table 3 is what is being subtracted from the dependent and independent variables in Table 7.

Stocks data are taken from CRSP. In each cross section, I drop the small stocks, defined as those with market cap below the 10th percentile of NYSE. Industry-level returns are market cap-weighted averages. All regressions in Table 7 require that there be at least 10 stocks in the industry-month. This filter is applied because the theory requires a group of stocks to be in an industry. The empirical results are not sensitive to the particular choice of the cutoff or whether the filter is applied.

Column 1 shows that within a look-back window of about a year (so the first four lags), the cross section of industry stock returns exhibits positive price momentum on average. This contrasts the very tenuous momentum effect found on the aggregate market. Similar to the results on the aggregate market, this momentum effect is entirely concentrated in the non-newsy months, and even flips sign in the newsy months. This is shown in columns 2 and 3. Column 4 shows the difference between the coefficients in the newsy months and non-newsy months, and they seem to be undeniably negative within four lags, albeit being non-monotonic.

Table 7 uses the first digit of the stock's issuing company's four-digit SIC code as the industry classification variable. This is the least granular, or the least specific and most crude, classification. The first two, three, and four digits represent increasingly granular industry classification.[3] Table 8 examines whether the dynamic autocorrelation structure is robust to the choice of the industry classification variable. Here I will switch from individual lags to the sum of the first four lags, as before. Across Panels A to D, it seems that the autocorrelation structure of industry excess returns is strongly different across newsy months and non-newsy months. Hence, it seems that this relationship is not sensitive to the specific level of the industry classification.

The results in this section might seem to be related to those in Heston and Sadka (2008) and Keloharju et al. (2016), who show that the full-year lags, i.e., the 12, 24, 36, etc. monthly lags have unusually high predictive coefficients on stock-level excess returns. Full-year lags are indeed used in some of my regression results. Specifically, they are part of the RHS when the dependent variable are the group 3 months, like in column 2 of Table 7. However, observe that the point of my paper is that in group 1 months the return predictive coefficients are unusually low, exactly the opposite of the points made in the return seasonality literature. In other words, my results exist in spite of the loading on return seasonality, not because of it.

The results in this section may also seem to be related to Savor and Wilson (2016), who show that stocks' announcement week returns can be positively predicted by their past announcement week returns. Also, the authors show that earnings announcements are highly persistent, and therefore firms who have announced in January are likely

[3]Incidentally, the two-, three-, and four-digit SIC code classifications are called Major Group, Industry Group, and Industry. The one-digit SIC code is not often used. It is nonetheless a reasonable industry measure, as proximity in SIC code corresponds to similarity in the nature of the firms' business. In this paper we will be calling various levels of SIC "industry" in the generic sense.

to announce in April. Combining these two pieces of information, it may seem that an industry's excess returns 3/6/9/12 months ago might be driving part of my results. However, apart from the difference in return frequency (monthly versus weekly) and cross sectional unit (industry versus stock), observe again that the excess returns 3/6/9/12 months ago are used only when the dependent variables are group 1 months, and the effect goes against what is shown in Savor and Wilson (2016).

5.2 Global Aggregate Markets

5.2.1 Global Earnings Reporting Cycle

Having described the results in the US, we move to the global data. In addition to new data on stock returns, other countries have different earnings reporting cycles from the US and therefore can potentially provide a meaningful out-of-sample test of the theory. As Table 1 shows, globally, the fiscal quarter is also well aligned with the calendar quarter. Hence, news reported in group 1 months will remain the freshest in the global data. However, the reporting intensity is very different and can be used to generate some heterogeneity in the RHS variables in the global analysis. Table 9 shows that the smallest fraction of the firms is reporting in the group 1 months, among the three groups. This is especially true if you focus on the firms reporting relevant news, as in columns 5-6 and 7-8, where only 10% of the firms report in group 1 months. In contrast, group 2 months contain the largest fraction of firms reporting. Globally, it seems that both group 1 and group 2 months are viable candidates, as the former provide the freshest news and the latter come with the most intensive reporting. Group 3 months, however, do not seem to fall in either category. Columns 3 and 4 in Table 5 show that the median reporting lag is about two months globally. In other words, globally group 1 combined with group 2 months are the first half of the earnings reporting cycles. To enforce consistency across the US and the global analysis, I will therefore set the newsy "months" to be the combined periods of group 1 and group 2 months in the global data.

5.2.2 Global Fundamental News

Analysis of the earnings news predictability also supports the newsy status of group 2 months in the global data. Table 10 does similar regressions to those in Table 6 [4]. Columns 5-8 of this table show that earnings news has below average loading on past news in group 1 months and group 2 months, and above average loading in group 3 months. Note here the number of observations is small, as the data for reporting date do not go back very far. The difference between the coefficient in columns 5 and 8 is not significant given the short sample. While this evidence is weak, it seems consistent with the earnings reporting-lag results in their support of group 2 months' newsy status.

5.2.3 Return Predictability

Having described the earnings-related information globally, I move to the return predictability results. My country-/territory-level market returns come from Global Financial Data (GFD). All 50 non-US countries/territories for which GFD provides the equity return index are included. Returns are originally provided in USD and they are kept that way, since this is a simple way to deal with hyperinflationary episodes which will make local returns extreme and less informative because of the persistent inflation component in them. A number of return series go beyond 1926, but in this analysis we cut the sample at 1926 to be consistent with the US results[5]. The theory predicates

[4]Some may encourage me to do the computation on a country-by-country basis. However, at the country level even aggregate book value can get dangerously close to zero very frequently at the monthly level, due in part to the thin coverage of Compustat in certain non-US countries. This prompted me to aggregate the whole global economy, which effectively weights country-level ROE by their respective denominators. If the divisor is close to 0, the country-month will get a weight close to zero.

[5]After deleting some obviously problematic data in the UK in the 1600s and 1700s, inclusion of those early samples makes the results stronger.

on regular earnings reporting in the economy at the quarterly level [6]. While that does seem to be present in the US dating back to 1925 and potentially even earlier thanks to the requirements of the NYSE (Kraft et al. (2017)), outside the US the evidence seems less easy to find, at least not systematically in one place. Hence, I will stick to the same sample and subsamples used in the US analysis for consistency. Of course, if the data for a certain country do not go beyond these sample truncation cutoffs, the truncation will not be operative for that country.

Similar to the US cross section, it is also very important that we subtract the US aggregate market component from the global country-/territory-level returns to form the excess returns, since individual market returns are highly correlated globally. Different from the US results, most of Table 11 considers group 2 months as newsy months, as is motivated by the analysis of fundamental data earlier in the paper. This means that the newsy "months" each year will actually be four two-month periods. Results with the one-month newsy month specification are also reported and compared to the main results. In column 1, we observe a strong unconditional momentum effect. This reflects country-level momentum. Columns 2 and 3 then immediately show that this component is much stronger in the non-newsy months, here meaning the group 3 months. The results are robust to subsamples in 5-7. Column 8 considers the group 2 months as non-newsy months, which is the specification in the US regression. Here the interaction is insignificant, even though it has the correct sign. Comparing columns 4 and 8, we see the results are more tenuous in column 8. This illustrates that the international return data want to use the two-month newsy month specification.

Incidentally, this decline in significance from columns 4 to 8 mostly comes from the change on the right-hand side [7]. In other words, this is mostly because group 1 months

[6]Incidentally, it is not necessary that each firm in the economy reports quarterly. The story applies even if all firms in an economy report annually, as long as the fiscal year-end months are somewhat evenly distributed among March, June, September, and December.

[7]The reason I make this statement is that if you do the one-month NM specification regressions by dependent variable month group, then the coefficient for group 2 (.021 [t=1.86]) is about the average

are noisy measures of earnings news, which is not entirely surprising given the small fraction of firms that report in group 1 months globally, as in Table 9.

5.3 Global Cross Section

Similar to the analyses for the US, we look at the cross section of country-/territory-industry returns in excess of the corresponding country-/territory-level returns. Returns are mostly taken from Compustat Global, augmented with the Canadian stocks from Compustat North America. To be consistent with the global aggregate market results, we constrain the sample to the 50 countries/territories in the previous section. Returns are winsorized at the panel 0.5 and the 99.5 percentiles to deal with extreme values. All returns are in USD. In each country/territory-month, 10% of the smallest stocks are dropped, after which the returns are aggregated to the industry and market level using market cap weight. The difference of those two returns are then taken to arrive at the excess returns that we use. Here, we again require each country/territory-industry to have at least 10 stocks.

Table 12 shows similar patterns to those found in Table 8: Countries-/territories-industries exhibit momentum only in the non-newsy months. The results do not seem to be sensitive to the particular choice of industry variable.

6 Time-Series Implementation

6.1 No Look-Ahead Bias Estimation

Past research has cast doubts especially on the practicability of time-series strategies in the stock market. Specifically, Goyal and Welch (2008) show that in forecasting future stock market returns, predictors combined with expanding window coefficients

of that for group 1 (.009, [t=0.86]) and group 3 (.037 [t=3.26]). Hence from the perspective of the LHS it is about equally good to tie group 2 with group 1 as with group 3. See Table 14 in the appendix.

extracted without look-ahead bias fail to outperform the expanding window mean of past market returns. While Campbell and Thompson (2008) quickly show that imposing some simple and reasonable constraints on the regression estimation process will make the predictor-based approach clearly superior, it is nonetheless useful to make sure that the predictor in this paper can generate positive R^2 without involving coefficients estimated with look-ahead bias.

To investigate the R^2 generated with no look-ahead bias coefficients, I extend the CRSP aggregate market return series to 1871 using GFD data. We take valuation ratios data from Robert Shiller's website, and construct payout ratios and ROE series with data from Amit Goyal's website. These are used to generate the long-run return forecasts as in Campbell and Thompson (2008). While the estimation sample go back to 1872, the R^2 are evaluated starting from 1926, consistent with what is done in Campbell and Thompson (2008).

The monthly R^2 generated by various estimation methods are reported in Table 13. Coefficients are either constrained to be 1 or generated with simple expanding-window OLS estimations. No shrinkage procedure is applied at all. In method 1, we see that even the most naive method—combining signal with regression coefficients of returns on past signals and a freely estimated constant—generates an R^2 of 3.65%. This positive R^2 along with those generated with methods 2-7, should mitigate the concern that the strategy implied by the time-series results in this paper could not have been profitably employed.

6.2 Performance Concentration

Despite the positive R^2 reported in the previous section, one may still worry that the time-series result is driven by a handful of outliers, and is therefore not very robust. To mitigate this concern, I plot the rolling 24-month coefficients of future returns on the signal value used in methods 2-7 of Table 13 and no constant in Figure 3. The signal is the sum of the past four newsy month returns, subtracting its expanding-window mean, and sign flipped if the dependent variable is a newsy month. This signal is 1) designed to have a positive coefficient, achieved via the sign flip step, and 2) designed to have a long-run mean of zero, achieved via the demeaning step. Given 1), coefficients can be estimated without separating the newsy and non-newsy months. Given 2), the coefficients can be estimated without a constant. This no-constant specification ensures that the portfolio underlying the coefficients is formed relative to the signal value of 0, so that the look-ahead bias is only in the portfolio weight scaling, which is less problematic and eventually unavoidable.

Figure 3 conveys a few messages: First, the implicit strategy has been risky: there are many two-year episode with negative coefficient values, implying the strategy generated negative returns. Second, the strategy has been profitable on average: while the coefficients are sometimes negative, overall the mean does seem to be positive. Third, such profit is not driven by a handful of observations, but rather earned over multiple episodes. These results should mitigate the concern that the time-series results are driven by outliers.

It is important to realize that section 6 is about the implementation of the coefficient estimation procedure. It makes no statement on how much of these return predictability results will continue to exist in the future. Recent works explicitly making those statements include Mclean and Pontiff (2016) and Hou et al. (2018).

7 Conclusion

I show that stock returns have exhibited a dynamic autocorrelation structure. Specifically, they exhibit negative/positive autocorrelation when fresh/old earnings news comes out, as in the first/second half of the earnings reporting cycle. Such autocorrelation structure is found in four contexts: 1) Time series of the US aggregate market returns, 2) the cross section of US industry returns, 3) the cross section of global market returns, and 4) the cross section of global industry returns. I argue that such dynamic autocorrelation structure can be generated if people incorrectly think the autocorrelation structure of earnings news is static in real time. This explanation is supported by fundamental data alone, and can be qualitatively delivered in a stylized model featuring both underreaction and overreaction. The theme is that investors act as if they have overreacted/underreacted to past news when fresh/old news comes out. The key assumption behind this result is that they think in real, calendar time and fail to see that past news has time-varying predictability on future news.

References

Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013). Value and Momentum Everywhere. *The Journal of Finance*, 68(3):929–985.

Ball, R. and Sadka, G. (2015). Aggregate earnings and why they matter. *Journal of Accounting Literature*, 34:39–57.

Barberis, N. (2018). Psychology-Based Models of Asset Prices and Trading Volume. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 1, pages 79–175. North-Holland.

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.

Beaver, W. H. (1968). The Information Content of Annual Earnings Announcements. *Journal of Accounting Research*, 6:67–92.

Bernard, V. L. and Thomas, J. K. (1989). Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium? *Journal of Accounting Research*, 27:1–36.

Bernard, V. L. and Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics*, 13(4):305–340.

Campbell, J. (2017). *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press.

Campbell, J. and Thompson, S. B. (2008). Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies*, 21(4):1509–1531.

Goyal, A. and Welch, I. (2008). A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *The Review of Financial Studies*, 21(4):1455–1508.

Guo, H. and Wachter, J. A. (2019). 'Superstitious' Investors. SSRN Scholarly Paper ID 3245298, Social Science Research Network, Rochester, NY.

Heston, S. L. and Sadka, R. (2008). Seasonality in the cross-section of stock returns. *Journal of Financial Economics*, 87(2):418–445.

Hou, K., Xue, C., and Zhang, L. (2018). Replicating Anomalies. *The Review of Financial Studies*.

Jegadeesh, N. and Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance*, 48(1):65–91.

Keloharju, M., Linnainmaa, J. T., and Nyberg, P. (2016). Return Seasonalities. *The Journal of Finance*, 71(4):1557–1590.

Kraft, A. G., Vashishtha, R., and Venkatachalam, M. (2017). Frequent Financial Reporting and Managerial Myopia. *The Accounting Review*, 93(2):249–275.

Matthies, B. (2018). Biases in the Perception of Covariance. SSRN Scholarly Paper ID 3223227, Social Science Research Network, Rochester, NY.

Mclean, R. D. and Pontiff, J. (2016). Does Academic Research Destroy Stock Return Predictability? *The Journal of Finance*, 71(1):5–32.

Moskowitz, T. J. and Grinblatt, M. (1999). Do Industries Explain Momentum? *The Journal of Finance*, 54(4):1249–1290.

Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices: Evidence and Implications. *Journal of Financial Economics*, 22(1):27–59.

Savor, P. and Wilson, M. (2016). Earnings Announcements and Systematic Risk. *The Journal of Finance*, 71(1):83–138.

Vuolteenaho, T. (2002). What Drives Firm-Level Stock Returns? *The Journal of Finance*, 57(1):233–264.

Table 1

Fiscal Period End Month

	US		Global	
Group by Month	Count	Percent	Count	Percent
Group 1: Jan/Apr/Jul/Oct	75,278	9.09	17,261	2.03
Group 2: Feb/May/Aug/Nov	47,600	5.75	15,906	1.87
Group 3: Mar/Jun/Sep/Dec	705,235	85.16	817,392	96.1
Total	828,113	100	850,559	100

The table tabulates the number of company-quarter by 3 groups, where the groups are determined by the remainder of the quarter end month divided by 3. First two columns contains count and percentage on the full sample of US companies. The next two columns do the same tabulation on the full sample of Global companies.

Table 2

US Company-Quarter Count by Reporting Month

	All		0 mod 3 FQ end		Rpt Lag <= 91		Both filters	
Group by Month, Mod 3	Count	Pct	Count	Pct	Count	Pct	Count	Pct
Group 1: Jan/Apr/Jul/Oct	369,379	44.6	340,503	48.28	358,543	44.18	330,047	47.71
Group 2: Feb/May/Aug/Nov	350,961	42.38	314,499	44.59	347,812	42.86	312,791	45.22
group 3: Mar/Jun/Sep/Dec	107,773	13.01	50,233	7.12	105,183	12.96	48,893	7.07
Total	828,113	100	705,235	100	811,538	100	691,731	100

The table tabulates the number of company-quarter by 3 groups, where the groups are determined by the remainder of the reporting month divided by 3. First two columns contains count and percentage on the full sample. The next two columns do the same tabulation but only on the subsample where the fiscal quarter being reported end in Mar/Jun/Sep/Dec. The next two columns apply a filter requiring reporting lag to be less than 3 months. The last two columns apply both filters.

Table 3

US Time-series Return by Month Type

	(1) All	(2) NM	(3) Non-NM	(4) Difference
mkt_l1nm	0.087	-0.179**	0.220***	-0.400***
	[1.64]	[-2.42]	[3.86]	[-4.28]
mkt_l2nm	0.018	-0.182***	0.117**	-0.299***
	[0.43]	[-2.95]	[2.40]	[-3.81]
mkt_l3nm	0.045	0.005	0.064	-0.059
	[0.99]	[0.06]	[1.54]	[-0.59]
mkt_l4nm	0.019	-0.108	0.083**	-0.191**
	[0.47]	[-1.46]	[2.03]	[-2.27]
mkt_l5nm	-0.069*	-0.101	-0.054	-0.047
	[-1.89]	[-1.59]	[-1.35]	[-0.63]
mkt_l6nm	0.029	-0.012	0.050	-0.062
	[0.89]	[-0.22]	[1.23]	[-0.90]
mkt_l7nm	-0.055*	-0.169***	0.003	-0.172***
	[-1.73]	[-3.12]	[0.08]	[-2.62]
mkt_l8nm	-0.011	-0.080	0.023	-0.103
	[-0.30]	[-1.33]	[0.56]	[-1.42]
Const	0.008***	0.021***	0.002	0.020***
	[3.03]	[4.60]	[0.59]	[3.56]
N	1094	364	730	1094
R-sq	0.022	0.101	0.080	0.089

Column 1 of this table reports results from the following monthly time-series regressions: $mkt_t = \alpha + \sum_{j=1}^{8} \beta_j mkt_{t-jnm} + \epsilon_t$. Here mkt_{t-jnm} is the US aggregate market return in the jth newsy month (Jan, Apr, Jul, Oct) preceding the month t. Column 2 reports the same regression on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from this regression $mkt_t = \alpha + \sum_{j=1}^{8} \beta_j mkt_{t-jnm} + \sum_{j=1}^{8} \gamma_j mkt_{t-jnm} \times I_{t,nm} + \delta I_{t,nm} + \epsilon_t$. T-statistics computed with White standard errors are reported in square brackets.

Table 4

US Time-series Return Subsample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All	NM	Non-NM	All	Post-war	First Half	Second Half	2-Mon NM
mkt_t4nm	0.050*	-0.106***	0.127***	0.127***	0.078***	0.159***	0.089***	0.064***
	[1.92]	[-2.71]	[4.60]	[4.60]	[3.35]	[3.73]	[3.12]	[2.88]
$mkt\_t4nm \times I_{nm}$				-0.233***	-0.156***	-0.279***	-0.179***	-0.095***
				[-4.87]	[-3.41]	[-3.89]	[-3.08]	[-3.07]
I_{nm}				0.015***	0.010**	0.019***	0.011**	0.012**
				[3.30]	[2.41]	[2.60]	[2.03]	[2.56]
const	0.007***	0.017***	0.002	0.002	0.005**	-0.000	0.004	0.001
	[3.25]	[4.40]	[0.82]	[0.82]	[2.29]	[-0.06]	[1.48]	[0.24]
N	1106	368	738	1106	863	555	551	1105
R-sq	0.008	0.030	0.056	0.047	0.025	0.059	0.033	0.018

Column 1 of this table reports results from the following monthly time-series regressions: $mkt_t = \alpha + \beta \sum_{j=1}^{4} mkt_{t-jnm} + \epsilon_t$. Here mkt_{t-jnm} is the US aggregate market return in the jth newsy month (Jan, Apr, Jul, Oct) preceding the month t. Column 2 reports the same regression as in 1, but on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from a regression with additional interaction terms to that in column 1. Column 5-7 report results for the regression in column 4 on the subsamples of the post war period (1947-), the first half (1926-1972), and the second half (1973-2019). Column 8 reports results for the regression in column 4, where the four newsy "months" each year are set to the four two-month periods of Jan+Feb, Apr+May, Jul+Aug, and Oct+Nov. Note while titles of column 1-7 indicate subsamples, that for column 8 indicates a different regression specification. T-statistics computed with White standard errors are reported in square brackets.

Table 5
Statistics on Reporting Lags

	US		Global	
Stats	All	Rpt Lag <= 91	All	Rpt Lag <= 91
1%	11	11	22	20
5%	16	16	29	26
10%	18	18	37	30
25%	24	24	53	45
50%	33	32	74	57
75%	44	43	163	69
90%	59	55	529	79
95%	76	68	1019	85
99%	105	89	2446	90
Mean	37.0	35.2	221.6	56.2
Std. Dev.	22.5	15.9	447.3	17.4
Skewness	7.5	1.2	4.6	-0.1
Kurtosis	187.3	4.7	29.8	2.3
Obs	828,113	811,538	850,559	506,672

The table tabulates statistics on reporting lag, computed as the earnings announcement date of a company-fiscal period minus the the end date of that company-fiscal period. The first two columns are on the US firms, while the next two are on global firms. The computation are based on all instances of company-fiscal period and also with the filter requiring the reporting to be within 3 months of the fiscal period end.

Table 6

US Time-series Monthly ROE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All	Group 1	Group 2	Group 3	All	Group 1	Group 2	Group 3
Δroe_l1m	0.343***	0.005	0.622***	0.749***				
	[2.78]	[0.07]	[4.12]	[4.07]				
Δroe_l2m	0.263***	0.430**	0.125***	0.398**				
	[4.57]	[2.50]	[2.69]	[2.14]				
Δroe_l3m	0.151**	0.043	0.257***	0.025				
	[2.36]	[0.22]	[3.02]	[0.28]				
Δroe_t3m					0.253***	0.148***	0.273***	0.351***
					[6.65]	[3.83]	[7.72]	[4.44]
const	-0.000	-0.000	-0.000	0.000	-0.000	-0.000	-0.000	0.000
	[-0.24]	[-0.32]	[-0.35]	[0.30]	[-0.25]	[-0.40]	[-0.32]	[0.13]
N	555	185	185	185	555	185	185	185
R-sq	0.390	0.249	0.617	0.549	0.382	0.211	0.539	0.459

Column 1 of this table reports results from the following monthly level time-series regressions: $\Delta roe_t = \alpha + \sum_{j=1}^{3} \beta_j \Delta roe_{t-j} + \epsilon_t$. Here Δroe_t is the aggregate roe, or the aggregate net income divided by aggregate book value of equity, of firms reporting their earnings in month t. Column 2-4 report results from the same regression except on the subsamples. The subsamples are split according to the month of the dependent variable: Jan/Apr/Jul/Oct, Feb/May/Aug/Nov, and Mar/Jun/Sep/Dec. Column 5-8 report similar results from the following regression: $\Delta roe_t = \alpha + \beta \sum_{j=1}^{3} \Delta roe_{t-j} + \epsilon_t$. T-statistics computed with White standard errors are reported in square brackets.

Table 7
US Cross-sectional Return by Month Type

	(1)	(2)	(3)	(4)
	All	NM	Non-NM	Difference
exret_l1nm	0.038*	-0.033	0.074***	-0.106**
	[1.83]	[-0.96]	[2.90]	[-2.51]
exret_l2nm	-0.015	-0.071*	0.013	-0.084*
	[-0.73]	[-1.95]	[0.56]	[-1.94]
exret_l3nm	0.045**	0.024	0.055**	-0.031
	[2.07]	[0.60]	[2.24]	[-0.66]
exret_l4nm	0.038*	-0.027	0.070***	-0.097**
	[1.95]	[-0.82]	[2.89]	[-2.40]
exret_l5nm	-0.031*	-0.065**	-0.014	-0.051
	[-1.77]	[-2.24]	[-0.64]	[-1.43]
exret_l6nm	0.027*	0.035	0.023	0.013
	[1.69]	[1.37]	[1.12]	[0.39]
exret_l7nm	0.015	-0.009	0.028	-0.037
	[0.84]	[-0.35]	[1.17]	[-1.04]
exret_l8nm	0.007	0.029	-0.004	0.033
	[0.41]	[1.05]	[-0.20]	[0.97]
const	-0.000*	0.000	-0.000***	0.000**
	[-1.90]	[0.86]	[-2.79]	[2.34]
N	9,550	3,175	6,375	9,550
R-sq	0.008	0.014	0.017	0.016

Column 1 of this table reports results from the following month-sic1 panel regressions: $exret_{i,t} = \alpha + \sum_{j=1}^{8} \beta_j exret_{i,t-jnm} + \epsilon_{i,t}$. Here $exret_{i,nm(t,j)}$ is the value weighted average return of industry i in the jth newsy month (Jan, Apr, Jul, Oct) preceding the month t. Column 2 reports the same regression on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from this regression $exret_{i,t} = \alpha + \sum_{j=1}^{8} \beta_j exret_{i,t-jnm} + \sum_{j=1}^{8} \gamma_j exret_{i,nm(t,j)} \times I_{t,nm} + \delta I_{t,nm} + \epsilon_{i,t}$. Regressions are weighted by the market cap of industry i as of the month $t-1$, normalized to sum to 1 in each cross section. T-statistics computed with clustered standard errors by month are reported in square brackets.

Table 8

US Cross-sectional Return by Different Industry Measures

	(1) All	(2) NM	(3) Non-NM	(4) Difference
Panel A: SIC 1				
exret_t4nm	0.027**	-0.027	0.054***	-0.081***
	[2.57]	[-1.50]	[4.48]	[-3.74]
N	8,519	2,835	5,684	8,519
R-sq	0.003	0.003	0.012	0.009
Panel B: SIC 2				
exret_t4nm	0.021***	-0.012	0.038***	-0.050***
	[3.06]	[-1.01]	[4.65]	[-3.45]
N	39,430	13,127	26,303	39,430
R-sq	0.002	0.001	0.007	0.004
Panel C: SIC 3				
exret_t4nm	0.016***	-0.006	0.027***	-0.033***
	[2.62]	[-0.59]	[3.64]	[-2.61]
N	56,142	18,683	37,459	56,142
R-sq	0.001	0.000	0.004	0.002
Panel D: SIC 4				
exret_t4nm	0.018***	-0.003	0.028***	-0.030**
	[2.79]	[-0.25]	[3.49]	[-2.36]
N	45,672	15,218	30,454	45,672
R-sq	0.001	0.000	0.004	0.002

Column 1 of this table reports results from the following month-industry panel regressions: $exret_{i,t} = \alpha + \beta \sum_{j=1}^{4} exret_{i,t-jnm} + \epsilon_{i,t}$. Here $exret_{i,nm(t,j)}$ is the value weighted average return of industry i in the jth newsy month (Jan, Apr, Jul, Oct) preceding the month t. Column 2 reports the same regression on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from this regression $exret_{i,t} = \alpha + \beta \sum_{j=1}^{4} exret_{i,t-jnm} + \gamma \sum_{j=1}^{4} exret_{i,nm(t,j)} \times I_{t,nm} + \delta I_{t,nm} + \epsilon_{i,t}$. Panel A to D differ only in the industry variable used. Regressions are weighted by the market cap of industry i as of the month $t-1$, normalized to sum to 1 in each cross section. 10 stocks are required in each month-industry. T-statistics computed with clustered standard errors by month are reported in square brackets.

Table 9

Global Company-Quarter Count by Reporting Month Mod 3

	All		0 mod 3 FQ end		Rpt Lag <= 91		Both filters	
Group by Month, Mod 3	Count	Pct	Count	Pct	Count	Pct	Count	Pct
Group 1: Jan/Apr/Jul/Oct	196,889	23.15	184,704	22.6	62,856	12.41	56,089	11.4
Group 2: Feb/May/Aug/Nov	344,437	40.5	336,255	41.14	244,720	48.3	241,838	49.17
group 3: Mar/Jun/Sep/Dec	309,233	36.36	296,433	36.27	199,096	39.29	193,939	39.43
Total	850,559	100	817,392	100	506,672	100	491,866	100

The table tabulates the number of company-quarter by 3 groups, where the groups are determined by the remainder of the reporting month divided by 3. If a company reports earnings for a given fiscal quarter in October, this company-quarter belongs to group 1, as $10 \equiv 1 \mod 3$. First two columns contains count and percentage on the full sample. The next two columns do the same tabulation but only on the subsample of where quarter being reported end in Mar/Jun/Sep/Dec. The next two columns apply a filter requiring reporting lag to be less than 3 months. The last two columns apply both filters.

Table 10

Global Time-series Monthly ROE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All	Group 1	Group 2	Group 3	All	Group 1	Group 2	Group 3
Δroe_l1m	0.419***	-0.163	0.357*	0.847***				
	[3.45]	[-0.86]	[1.95]	[6.41]				
Δroe_l2m	0.192*	1.037***	-0.218	-0.002				
	[1.77]	[5.41]	[-1.30]	[-0.01]				
Δroe_l3m	0.069	-0.131	0.474**	0.084				
	[0.75]	[-0.86]	[2.66]	[0.84]				
Δroe_t3m					0.227***	0.216***	0.201***	0.271***
					[7.53]	[2.99]	[6.22]	[5.66]
const	-0.000	0.000	-0.001	-0.000	-0.001	-0.000	-0.001	-0.000
	[-0.61]	[0.26]	[-0.76]	[-0.04]	[-0.65]	[-0.01]	[-0.86]	[-0.38]
N	141	47	47	47	141	47	47	47
R-sq	0.347	0.448	0.469	0.526	0.316	0.269	0.339	0.363

Column 1 of this table reports results from the following monthly level time-series regressions: $\Delta roe_t = \alpha + \sum_{j=1}^{3} \beta_j \Delta roe_{t-j} + \epsilon_t$. Here Δroe_t is the aggregate roe, or the aggregate net income divided by aggregate book value of equity, of firms reporting their earnings in month t. Column 2-4 report results from the same regression except on the subsamples. The subsamples are split according to the month of the dependent variable: Jan/Apr/Jul/Oct, Feb/May/Aug/Nov, and Mar/Jun/Sep/Dec. Column 5-8 report similar results from the following regression: $\Delta roe_t = \alpha + \beta \sum_{j=1}^{3} \Delta roe_{t-j} + \epsilon_t$. T-statistics computed with White standard errors are reported in square brackets.

Table 11

Global Country/Territory Level Return Subsample

	(1) All	(2) NM	(3) Non-NM	(4) All	(5) Post-war	(6) First-Half	(7) Second-Half	(8) 1-Mon NM
$exmkt\_t4nm$	0.017***	0.005	0.040***	0.040***	0.041***	0.051**	0.044***	0.029***
	[3.22]	[0.78]	[4.58]	[4.59]	[4.39]	[2.32]	[4.15]	[3.63]
$exmkt\_t4nm \times I_{nm}$				-0.036***	-0.032***	-0.068**	-0.032**	-0.020
				[-3.30]	[-2.80]	[-2.57]	[-2.50]	[-1.47]
I_{nm}				-0.001	-0.001	-0.000	-0.002	0.004
				[-0.40]	[-0.19]	[-0.14]	[-0.71]	[1.62]
const	0.001	0.001	0.002	0.002	0.002	-0.000	0.003	-0.000
	[1.04]	[0.62]	[0.96]	[0.96]	[0.88]	[-0.14]	[1.28]	[-0.23]
N	29,605	19,723	9,882	29,605	26,793	12,209	22,143	29,657
R-sq	0.002	0.000	0.015	0.005	0.006	0.008	0.007	0.003

Column 1 of this table reports results from the following month-conutry/territory panel regressions: $exmkt_{i,t} = \alpha + \beta \sum_{j=1}^{4} exmkt_{i,nm(t,j)} + \epsilon_{i,t}$. Here $exmkt_{i,t} = mkt_{i,t} - mkt_{US,t}$, and $exmkt_{nm(t,j)}$ is the aggregate market return in the jth newsy "month" (Jan+Feb, Apr+May, Jul+Aug, Oct+Nov) preceding the month t for country/territory i. Column 2 reports the same regression as in 1, but on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from a regression with additional interaction terms to that in column 1. Column 5-7 report results for the regression in column 4 on the subsamples of the post war period (1947-), the first half (1926-1972), and the second half (1973-2019). Column 8 reports results for the regression in column 4, where the four newsy months each year are set to the four one-month periods of Jan, Apr, Jul, and Oct. Note while titles of column 1-7 indicate subsamples, that for column 8 indicates a different regression specification. T-statistics computed with clustered standard errors by month are reported in square brackets.

Table 12

Global Cross-sectional Return by Different Industry Measures

	(1) All	(2) NM	(3) Non-NM	(4) Difference
	Panel A: SIC 1			
$exret\_t4nm$	0.005	-0.002	0.019**	-0.021**
	[1.01]	[-0.32]	[2.47]	[-2.13]
N	82,197	54,756	27,441	82,197
R-sq	0.000	0.000	0.004	0.001
	Panel B: SIC 2			
$exret\_t4nm$	0.007	-0.000	0.021***	-0.021**
	[1.42]	[-0.05]	[2.95]	[-2.26]
N	149,775	99,761	50,014	149,775
R-sq	0.000	0.000	0.004	0.001
	Panel C: SIC 3			
$exret\_t4nm$	0.007	-0.000	0.023***	-0.023**
	[1.49]	[-0.07]	[3.08]	[-2.40]
N	144,370	96,237	48,133	144,370
R-sq	0.001	0.000	0.005	0.002
	Panel D: SIC 4			
$exret\_t4nm$	0.004	-0.004	0.021***	-0.025***
	[0.93]	[-0.72]	[2.93]	[-2.74]
N	136,484	90,918	45,566	136,484
R-sq	0.000	0.000	0.004	0.001

Column 1 of this table reports results from the following month-country/territory-industry panel regressions: $exret_{i,c,t} = \alpha + \beta \sum_{j=1}^{4} exret_{i,c,t-jnm} + \epsilon_{i,c,t}$. Here $exret_{i,c,nm(t,j)}$ is the value weighted average return of industry i of country/territory c in the jth newsy month (Jan, Apr, Jul, Oct) preceding the month t. Column 2 reports the same regression on the subsample where the dependent variable are returns of the newsy months. Column 3 is for the non-newsy months. Column 4 reports the difference between the coefficients in column 2 and 3, extracted from this regression $exret_{i,c,t} = \alpha + \beta \sum_{j=1}^{4} exret_{i,c,t-jnm} + \gamma \sum_{j=1}^{4} exret_{i,c,nm(t,j)} \times I_{t,nm} + \delta I_{t,nm} + \epsilon_{i,c,t}$. Panel A to D differ only in the industry variable used. Regressions are weighted by the market cap of industry i as of the month $t-1$, normalized to sum to 1 in each country/territory-month. 10 stocks are required in each month-country/territory-industry. T-statistics computed with clustered standard errors by month are reported in square brackets.

Table 13

Time Series R2 without Look Ahead Bias

Method	0	1	2	3	4	5	6	7
R^2	0.53%	3.65%	3.52%	3.61%	3.97%	3.78%	3.83%	4.33%

The R^2 in this table are calculated as $1 - \frac{\sum_{t=1}^{n}(r_t - \hat{r_t})^2}{\sum_{t=1}^{n}(r_t - \overline{r_t})^2}$, where $\overline{r_t}$ is the expanding window mean of past stock returns, and $\hat{r}_t$ is the forecast being evaluated. This R^2 is positive only when the forecast outperforms the expanding window mean of past stock returns. Method 0 comes solely from Campbell and Thompson (2008) and does not have anything to do with this paper. It is the valuation constraint + growth specification with fixed coefficients. Simple average is taken from the Dividend/Price, Earnings/Price, and Book-to-market ratios based forecasts. Method 1 uses the signal that is simply the sum of past four newsy month returns. The coefficients are extracted from simple expanding-window OLS of past returns on past signals, separately for newsy and non-newsy month dependent variables. The signal used in method 2-7 is the sum of past four newsy month returns, subtracting its expanding window mean, and sign flipped if the dependent variable is a newsy month. Method 2 uses the same coefficient estimation method as in method 1. Methods 3 replace the constant terms with the expanding window means of past newsy and non-newsy month returns. Method 4 replace the constant terms with the forecast in method 0. Method 5-7 are method 2-4 with the coefficients estimated on the combined sample of newsy and non-newsy months.

Figure 1: Timing of Independent and Dependent Variables in Return Forecasting Regressions, US

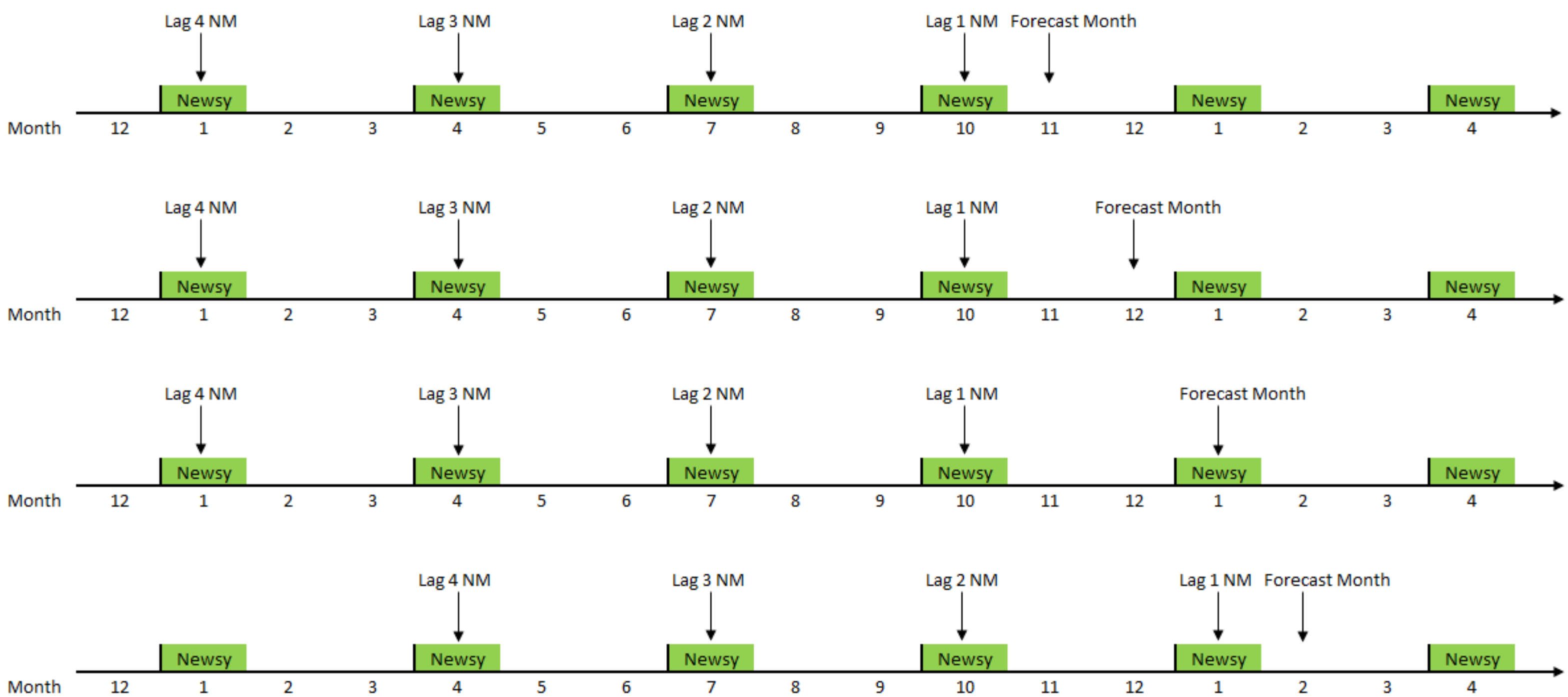

This figure shows how the independent variables in the US return forecasting regressions progress as the dependent variable move forward.

Figure 2: Distance in Calendar and Fiscal Time of Lagged Data in an Earnings Forecasting Setting

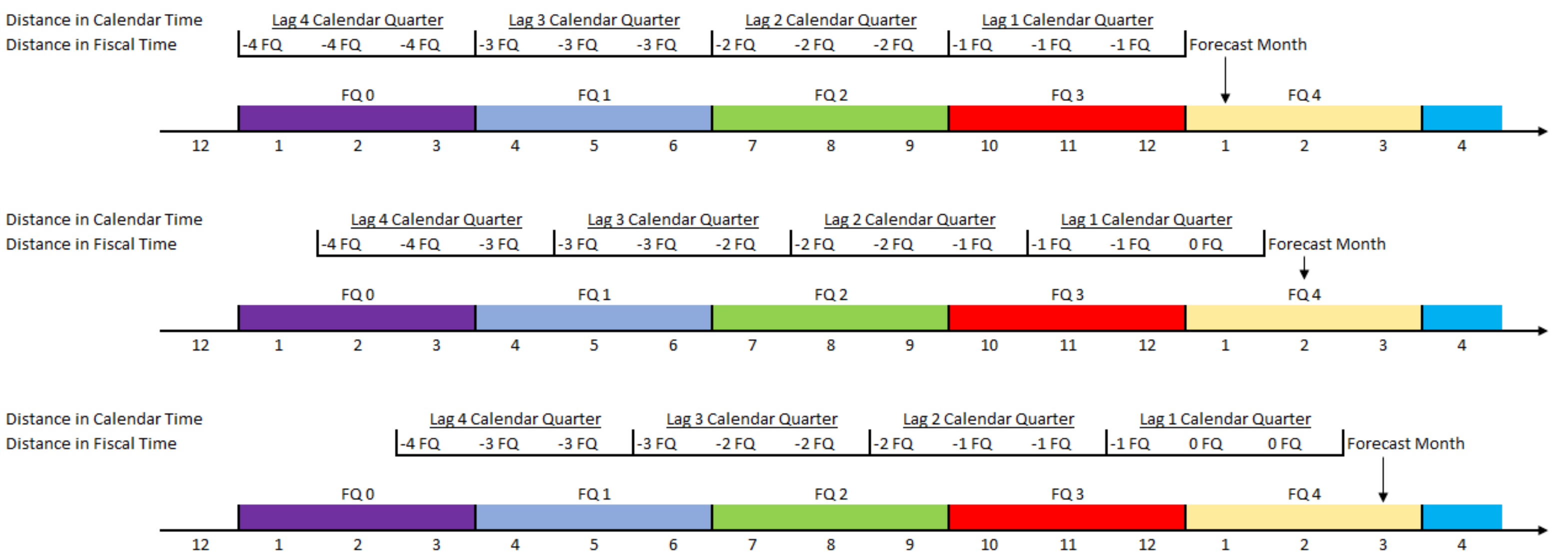

This figures depicts the dependent and independent variables used by a hypothetical investor trying to forecast the earnings news in each calendar month using past earnings news. It demonstrates that past earnings news contained over the same look-back window in terms of calendar time is actually further away in terms of fiscal time when trying to forecasts earnings news of the newsy months. Fiscal periods corresponding to earnings news reported in each calendar month are labeled and color coded.

Figure 3: Rolling 24-month Regression Coefficients, US Aggregate Market

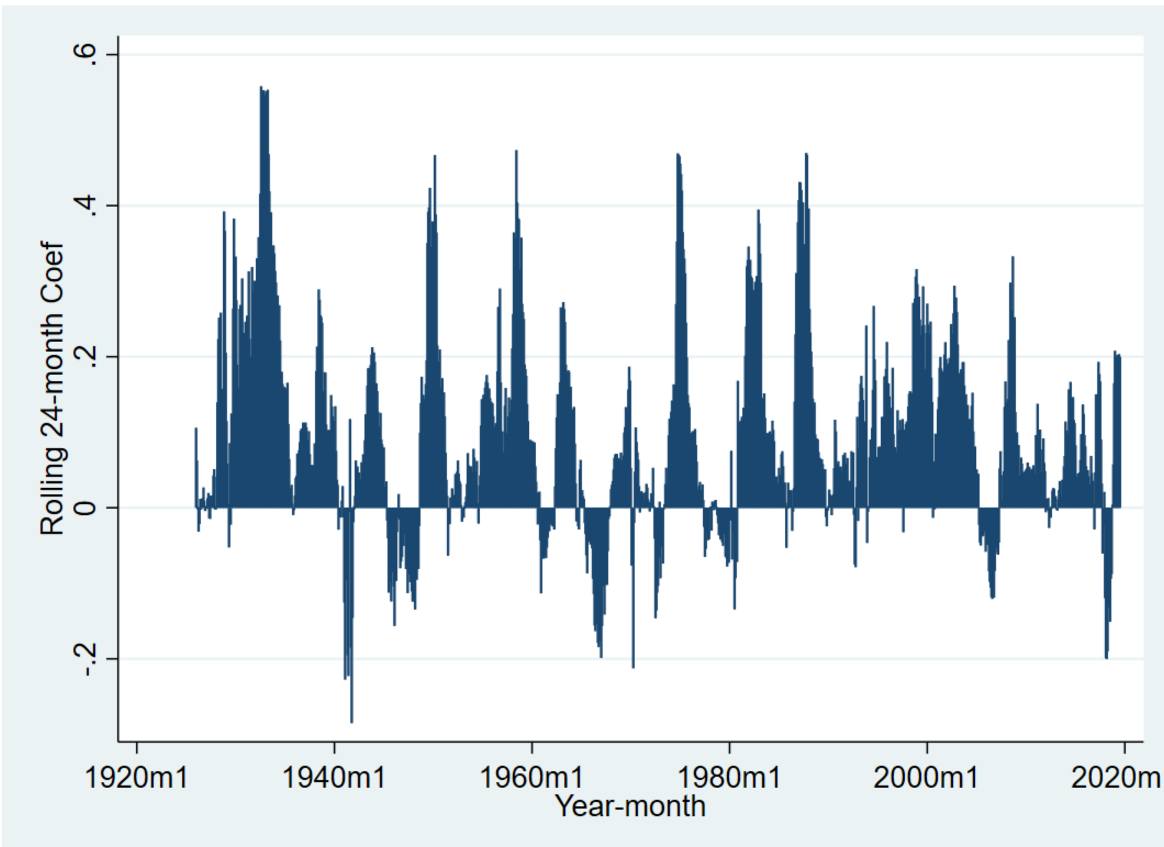

This figures shows the rolling 24-month regression coefficients of future US aggregate market returns on a signal and no constant. The signal is the sum of past four newsy month market returns, expanding window demeaned, and sign flipped for newsy month dependent variables.

A Appendix

Table 14 reports market return regressions with the dependent variables broken down into the 3 groups of months. Columns 1-3 are for the US where the dependent variables are outside the first calendar quarter of each year. Column 4-6 are where the dependent variables are within the first calendar quarter of each year. This motivation is that in the US, the earnings reporting for the last fiscal quarter of the fiscal years is substantially more slowly than that for other fiscal quarters. The median reporting lag in the first quarter of the US is about 41 days, which is about the same as the fastest reporting non-US country. This makes January in the US potentially less newsy than April, July, and October. The situation in the first quarter of the US is more similar to that in the global sample.

Comparing columns 1-3 against columns 4-6, we see that the dynamic autocorrelation in the US is much weaker in the first quarter than the others. Many reasons can be postulated. For instance, one can argue that January is not as newsy, and the effect of "reckoning" is not as strong. Hence the the coefficient in column 4 is not as negative as that in column 1. Similarly, coefficients in column 5 is less positive than that in column 2, potentially because of two reasons. First returns in January do not reflect earnings news as well. This reason is also potentially behind the discrepancy between the coefficient in column 4 and 6. Second, some of the "reckoning" that normally happens in the first month of the quarter happens in the second month, February, in the first quarter. Notice this situation in the first quarter of the US is similar to that in the global data. Consistent with this observation, patterns in column 4-6 are similar to those in column 7-9.

This consistency, however, should not be taken seriously at all because comparing regression coefficients across countries is a very dicey practice. Many factors other than the one you are interested in (here is the similarity/discrepancy of the reporting lag) can be driving the difference in regression coefficients. Just to name one example, difference in sample length alone can completely drive the results. Comparison within country, like that between column 1-3 and column 4-6, is potentially more reliable. Incidentally, globally, it is also true that the reporting lag is longer in the first quarter of the year, and that the pattern of dynamic autocorrelation is weaker.

Table 14

Time Series Return with One-month Newsy Month Specification

	US						Global		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	GP 1 no Jan	GP 2 no Feb	GP 3 no Mar	Jan	Feb	Mar	GP 1	GP 2	GP 3
mkt_t4nm	-0.130***	0.167***	0.122***	-0.016	0.056	0.086	0.009	0.021*	0.037***
	[-2.92]	[3.44]	[2.92]	[-0.23]	[1.05]	[1.14]	[0.82]	[1.86]	[3.26]
const	0.017***	0.003	0.000	0.016***	0.004	0.004	0.004*	-0.002	0.002
	[3.63]	[0.73]	[0.04]	[2.75]	[0.66]	[0.51]	[1.82]	[-1.08]	[0.81]
N	276	277	277	92	92	92	9,883	9,889	9,885
R-sq	0.043	0.082	0.054	0.001	0.016	0.026	0.000	0.002	0.007

This table reports results from the following monthly time-series regressions: $mkt_t = \alpha + \beta \sum_{j=1}^{4} mkt_{t-jnm} + \epsilon_t$. In this table, the newsy months are set to Jan, Apr, Jul, Oct regardless of the country. Here mkt_{t-jnm} is the market return in the jth newsy month preceding the month t. Column 1-6 are in the US, while column 7-9 are on the cross section global market returns. Column 1-3 are where the dependent variables are group 1-3 months outside the first quarter. Column 4-6 are where the dependent variables are returns of January, February, and March. Column 7-9 are where the dependent variables are group 1-3 months. T-statistics are reported in square brackets. In column 1-6 they are computed with White standard errors. In column 7-9 they are computed with standard errors clustered by month.