The Earnings Announcement Premium across an Earnings Season*

Roberto C. Gutierrez Jr.[a,†]
and
Cameron Pfiffer[b]

August 31, 2024

Abstract

Recent studies show that using firm-level earnings to systematically learn about a common factor induces a conditional risk premium on firms. We find that only the first half of firms reporting their earnings each quarter earn the well-documented earnings announcement return premium; the rest do not. Firms announcing late in the quarter provide no new incremental information about the factor and do not warrant a premium. A real-time Bayesian model of systematic learning finds that weekly reductions in economic uncertainty, achieved by observing the actual earnings flow, covaries positively with weekly announcement premiums. The more informative the week, the greater the announcement premium.

*We are grateful for comments from Jeremy Piger, Rob Ready, Youchang Wu, and seminar participants at the University of Oregon..

[a]Lundquist College of Business, University of Oregon (rcg@uoregon.edu)

[†]Corresponding author

[b]Graduate School of Business, Stanford University (cpfiffer@stanford.edu)

Quarterly earnings releases are important informational events that largely occur in clusters. The information flows of these "earnings seasons" are commonly used by agents to update their beliefs about industry trends and broad economic growth. Recent research has examined the bellwether aspect of earnings reports and its impact on stock pricing, which we build upon in this paper.

Savor and Wilson (2016) and Andrei et al. (2023) show that the use of a firm's earnings by investors to systematically learn about a market-wide common factor imparts a risk premium upon the firm at the event of its earnings release. The more systematically informative an earnings report is, the greater the risk premium. This systematic learning channel implies that the placement of a firm's earnings announcement within an earnings season is important. In particular, the well-documented finding of a stock-return premium for announcing earnings should depend critically upon when an earnings report falls inside a quarter.[1]

In this paper, we uncover a clear structure to the earnings announcement premium (EAP) that is remarkably consistent with a systematic learning channel. The systematic information content of earnings reports within each quarter should have a saturation point at which further reports reveal no incrementally new information about the common factor in earnings. At this point in the quarter, the EAP should disappear.

From 1995-2022, the first half of firms reporting their earnings in a quarter have a mean stock return premium around 0.40% in the week of their earnings announcement. In contrast, the second half of quarterly reporters display no return premium. Relatedly, Savor and Wilson (2016) show that the earlier an earnings report occurs in a quarter, the larger the announcement premium. We document that the contrast

[1] See, for example, Beaver (1968), Chari et al. (1988), Lamont and Frazzini (2007), Cohen et al. (2007), and Barber et al. (2013) for studies of the earnings announcement return premium.

between early and late announcements is starker than currently recognized. The announcement return premium is isolated, on average, to just the first three weeks of earnings seasons.

The finding that the earnings announcement premium exists only in the initial weeks of earnings seasons is robust and holds:

- in event time and calendar time,
- using CAPM abnormal returns, Fama-French five-factor model abnormal returns, and the return spread between announcers and non-announcers,
- within each half of the sample period,
- for a subsample of on-time announcements, and
- controlling for earnings surprises (SUE).

Both the on-time subsample and the tests controlling for differences in SUE across the initial and subsequent weeks of an earnings season are meant to address the concern that firms with higher earnings surprises tend to advance their announcements, while firms with lower earnings surprises tend to delay (e.g., Johnson and So, 2018).

We then consider whether three weeks into earnings seasons is a reasonable point for the systematic usefulness of firms' earnings reports to fade. To address this, we estimate a one-factor model of earnings from the perspective of a Bayesian investor in real-time. The investor uses the actual flow of earnings reports to update her beliefs each week on the common and firm-specific parameters of the model. Following Andrei et al. (2023), we define economic uncertainty to be the conditional variance of the common factor of earnings. The reduction in the posterior variance of the earnings factor is measured each week, and quantifies how much the investor learns from the earnings reports released during the week.

The earnings model indicates that three weeks does seem reasonable for the saturation point of systematic news. The economic uncertainty in the system declines on average by nearly 60% after observing earnings in week one, then a further 40%

decline after week two, and another 20% after week three. Further reductions are minimal. We then examine whether the estimates of the systematic learning temporally covary with the weekly earnings announcement premium. Reducing economic uncertainty by one standard deviation is associated with a increase in the EAP of 0.24 percentage points per week.

Lastly, we find that the volatility of residual returns from an unconditional CAPM is greater in the initial weeks of earnings seasons than the volatility in the subsequent weeks. We interpret this to mean that stock prices are updated more intensely during the first weeks, consistent with a conditional systematic risk that dissipates after the initial weeks.

Overall, our findings support systematic learning as a determinant of the earnings announcement premium. Researchers have offered other explanations for the premium that are not mutually exclusive, such as attention-driven price increases (Lamont and Frazzini, 2007) and the pricing of idiosyncratic earnings risk as a consequence of market frictions (Cohen et al., 2007; Barber et al., 2013; Johnson and So, 2018). However, we note that idiosyncratic risk remains abnormally high in the later weeks of earnings seasons when the EAP falls to zero. The decoupling of idiosyncratic risk and the EAP in the later weeks strains the second alternative.

In a related finding, Chan and Marsh (2022) show that the CAPM holds cross-sectionally during the first three days of earnings seasons in which a bulk of S&P 500 firms report their earnings (LEAD days). Their finding is consistent with systematic learning from earnings. Moreover, LEAD days appear to provide a better signal-to-noise ratio for estimating expected stock returns.

Overall, the findings in this paper emphasize the need for more research on the processing of earnings information in batches. Each quarter there are many days when over one hundred firms report earnings, and some days with over two hundred.

Examinations of how investors should optimally filter these information chunks and what price effects we should expect from such behavior seem fruitful.

1. Earnings announcement premium revisited

Our sample of earnings announcements comes from Compustat and ranges from 1995 through the first quarter of 2022. Dellavigna and Pollet (2009) find that earnings announcements dates before 1995 were recorded in Compustat (and I/B/E/S) with a much larger error rate, and with notably more errors on Friday announcements than on other days. Beginning in 1995, they find that recorded dates are accurate 96% of the time.

We retain only the U.S. firms having fiscal year-ends in December, which captures the vast majority of firms (e.g., see Noh et al. (2021)). Having firms that report their earnings for the same time periods facilitates the systematic signal extraction process we consider in this paper. Otherwise, two firms on different reporting cycles may announce their quarterly earnings in the same calendar week, but these earnings would only share one or two months in their respective reporting periods. Such cases would greatly complicate the earnings structure we model in section 2.

To isolate the most active weeks each quarter, we restrict the sample of earnings announcement dates to be at least 15 days and no more than 60 days after the end of each quarter. Announcements outside these windows are infrequent.

Stock returns are from CRSP. To reduce microstructure and illiquidity effects, we require a stock to have a price per share greater than $5 at the end of the prior quarter. Stocks are grouped into large (LG) and small (SM) stocks based on their market capitalizations at the end of the prior quarter. Large stocks have a market-capitalization above the median stock on NYSE, and small stocks are the remaining ones. CAPM betas are estimated from daily returns in the two quarters preceding the reporting period, with a minimum of 100 days of returns available. Betas and CAPM abnormal returns are winsorized within each quarter and each size category at the 1% and 99% levels.

Prior research referenced in the introduction has used daily, weekly, and monthly returns to examine the earnings announcement return premium. We choose to examine weekly returns for two reasons. First, grouping announcements by calendar weeks allows us to consider earnings seasons as batches of information. That is, we presume that investors are processing the available set of firms' news to extract a systematic signal, rather than learning from each firm's news in isolation. A model of a Bayesian investor's systematic learning is shown in section 2. Second, using weeks accommodates measurement error in the *expected* announcement dates, which can differ from the actual announcement dates. Firms that delay (advance) the release of their earnings announcement, relative to a date predicted by their reporting patterns, tend to have weaker (stronger) earnings news (e.g., Johnson and So, 2018). These adjustments in the timing of announcements tend to be a matter of days, not weeks. Nevertheless, we also consider an on-time sample of earnings announcements by relying on the methodology of Cohen et al. (2007).

The sample construction produces 286,893 earnings announcements from 11,445 firms, as measured by unique PERMNOs. The mean number of announcements per firm is 25, with a median of 14. The mean (median) number of firms per quarter is 2,632 (2,516).

Figure 1 displays the clustering of earnings announcements over the weeks of year. We measure weeks as Monday to Friday. The length of the earnings season in the first quarter is spread over seven weeks while the other quarters are essentially spread over five weeks. The slightly more diffuse season in the first quarters is due to fiscal-year reports requiring more time to prepare. The histogram also displays positive skewness within each quarter, meaning that the early reporters tend to cluster more in time than the later reporters.

We will examine how the earnings announcement return premium varies across earnings seasons using several approaches. First, we separate firms each quarter into the early half of reporters and the late half, and then average their respective returns in event time and in calendar time. We later examine returns week by week over a quarter.

1.1. Separating earnings seasons into halves

We sort firms each quarter by their announcement dates, from earliest to latest. Firms having announcement dates before or on the day the cumulative distribution of firms having reporting reaches 50% are labeled the "initial" reporters. The remaining half of the firms are labeled the "follow-on" reporters.

Since we are examining weekly returns, partitioning the firms based on a cumulative percentage cutoff very often produces one week each term which is a mixture of initial and follow-on reporters. This mixed week happens for 86% of the quarters in our sample. To disambiguate the initial reporters from the follow-ons, we drop these mixed weeks from the current analysis (but will include these mixed weeks later in the week-by-week analysis).

The initial-only reporting weeks are typically two weeks per quarter, with a 10^{th} percentile of one week and a 90^{th} percentile of three weeks. These initial weeks have a mean (median) of 456 (437) firms per week. As indicated by Figure 1, the distribution of the latter half of reporters is spread over more weeks. The follow-on reporting period lasts four weeks on average, with a 10^{th} percentile of three weeks and a 90^{th} percentile of five weeks. The mean (median) number of follow-on firms per week is 245 (221).

1.1.1. Event-time averaging

We estimate the earnings announcement return premium (EAP) in two ways. We primarily use an unconditional CAPM to measure the abnormal stock returns of the announcing firms, but we also measure the return of all announcing firms in a given week in excess of the return of all non-announcing firms that same week. We also consider robustness over time by splitting the sample period in half, examining the 1995-2008 and 2009-2022 subperiods.

Pooling all initial reporters over the sample period together, Table 1 shows their mean CAPM abnormal returns in the announcement week. Panel A shows that initial reporters earn a 0.42% premium per week, significant at the 1% level, over the full 1995-2022 period. Panel A uses the actual dates these initial firms announce their earnings.

Because firms have discretion over precisely when they announce, and because the decision to advance (delay) their announcement is associated with having (stronger) weaker earnings news (e.g., Johnson and So, 2018), we also consider the subset of firms that report their earnings on time, based on their reporting patterns. To determine on-time reporting, we follow the procedure of Cohen et al. (2007) and divide the 1995-2022 sample period into six blocks of four years and one block of three years. Within each of these blocks, we use the median trading day within each quarter of the year to estimate the expected announcement dates for each year-quarter in that time block. We require each firm to have the full three or four years of earnings data available to be considered for the on-time subsample. An earnings announcement is classified to be on-time when (i) the actual announcement date is within ± 1 day of the predicted date and (ii) the actual and predicted dates fall in the same week.[2]

Panel B of Table 1 examines the subset of on-time earnings releases. Throughout the various samples examined in the table, the number of on-time observations is roughly half of all earnings announcements. For the 1995-2022 period, we see a significant mean abnormal return of 0.39% in the announcement week, which is very close to the 0.42% across all announcers. For the two subperiods, the premia estimates for the on-time reporters are also very close to the respective all-announcers estimates. In short, the EAP for the initial reporters is robust to concerns about the discretionary timing of firms' earnings reports.

Table 2 shows the mean announcement return premium for the follow-on earnings reports. From 1995-2022 and for each subperiod, follow-on announcements do not display a return premium, neither in the full announcement sample in Panel A nor the on-time reports in Panel B.

The main takeaway from Tables 1 and 2, which is further supported in the next sections, is that the earnings announcement premium is a robust feature of only the initial reporters in each quarter. The follow-on earnings reports on average do not have an announcement premium. This is consistent with earnings announcement premia manifesting as compensation for systematic risk, whereby only the initial earnings reports each quarter provide signals about the state of the economy, while follow-on reports reveal no new systematic news (Savor and Wilson, 2016; Andrei et al., 2023).

Lastly, note that the fraction of on-time reporting is lower for the follow-on reports than for the initial reports, falling from about one half in Table 1 to about one third

[2] Using the four-year blocks, the median date is often calculated as a half day, for example, day 20.5. In these cases, we require both the rounded down and rounded up dates to fall in the same week as the actual announcement date.

in Table 2. This is consistent with the fact that large firms tend to report earlier than small firms and that large firms are more likely to report their earnings on time than small firms (e.g., Noh et al., 2021).[3] We turn now to examinations of EAP within large and small firms, respectively, in the next section.

1.1.2. Size-based analyses

Prior studies examine how the size of a firm affects the EAP, and with differing conclusions. Some find that the premium decreases with size (Ball and Kothari, 1991; Cohen et al., 2007). Others find that the premium increases with size (Lamont and Frazzini, 2007). Barber et al. (2013) find that the EAP for firms in the United States increases with size, while the EAP for firms in other countries decreases with size. In short, a size effect on the premium seems fragile across samples and methods.

To provide evidence on the size effect from the perspective of initial versus follow-on earnings announcements, we split the sample of announcements into two bins based on the market capitalization of the announcing firm's equity at the end of the quarter preceding each earnings announcement. Large firms (LG) are those having a market capitalization greater than the median value of NYSE stocks, and small firms (SM) are the remaining ones.[4]

Tables 1 and 2 show the mean abnormal returns for large and small stocks separately. In short, small stocks display the same robust pattern we noted earlier, whereby initial reporters earn an EAP while follow-on reporters do not. However, large firms — which comprise about 40% of the initial announcements and 20% of

[3] In 2002, the SEC began requiring larger firms to report earlier (CFR 17.240.12b-2). However, stability in earnings announcement dates based on firm size has been well documented for the decades preceding these regulations (Givoly and Palmon, 1982; Chambers and Penman, 1984). In this sense, the "accelerated filer" and "large accelerated filer" regulations do not appear to have disrupted common practices, but instead, codified elements of an existing structure within earnings seasons.

[4] We than Ken French for posting the decile breakpoints of market capitalization on his website.

the follow-on announcements — display some subperiod divergencies. Namely, large firms have an insignificant premium for initial reports in the early subperiod (Table 1), and a positive premium for follow-on reports in the later subperiod (Table 2).

First, note that large firms display the same patterns as small firms for the full sample period. Moreover, when we measure the EAP in calendar time as in section 1.2, these two subperiod deviations exhibit fragility.[5] Whether using either the CAPM to measure the premium or the returns of announcers in excess of non-announcers, each deviation fades. The two calendar-time analyses find that large-firm initial reporters have a significantly positive premium in the second subperiod (about 0.20% with t-statistics above 2.5) and that large-firm follow-on reporters have an insignificant premium in the first subperiod (between 0.06% and 0.18% with t-statistics below 1.59).

In sum, the placement of an earnings announcement within a quarter is more important to understanding the earnings announcement premium than the size of the firms making the announcements. Given this, subsequent analyses in this paper do not filter on the size of a firm.

1.2. Calendar-time averaging

Averaging return premia by calendar week first, and then across the weekly averages, is a common alternative measure of a return premium. This approach resembles the return performance of a real-time investor seeking exposure to the earnings announcement premium, and also provides a cleaner measure of standard errors, as cross-firm correlations are effectively purged.

[5] We use the same test design as Table 4 but examine only the large-firm announcements within each week having at least 200 firms of any size reporting their earnings, and additionally require at least 25 large-firm announcements in each week. This secondary filter on the number of large firms binds only in the follow-on weeks, where small firms dominate the earnings announcements.

However, an issue with the calendar-time approach is that some weeks have few earnings announcements, and hence, the returns of those firms receive relatively greater weight in the overall analysis. Specifically, the event-time weighting for firm i in calendar quarter q is $\frac{1}{N}$, where N is the number of earnings reports in the full sample. But the calendar-time weighting of the same firm's report is $\frac{1}{nT}$, where n is the number of firms reporting in quarter q and T is the number of calendar weeks in the full sample. Given that N is so much greater than T, filters on the number of firms in a given week must be considered.

For context, we examine the distributions of the number of firms per week for the initial earnings reports and for the follow-on reports. Recall that our definition of the follow-on portion of a earnings season is typically four weeks while the initial portion is typically two weeks. This implies that an initial week contains notably more announcements than a follow-on week. The 25th percentiles of the count of firms announcing per week are 194 firms for initial weeks and 26 firms for follow-on weeks. The median counts are 437 firms and 221 firms respectively. Therefore, it is the follow-on weeks where robustness may be an issue. To address this, we examine two filters, a 25-firm filter and a 200-firm filter, which are roughly the 25th and 50th percentiles of counts in the follow-on weeks.

The calendar-week approach also facilitates the use of alternative measures of the earnings announcement premium. In addition to using CAPM abnormal returns, we report the raw return of announcing firms in a given week in excess of the raw return of the non-announcing firms in that same week (e.g., Savor and Wilson, 2016; Barber et al., 2013). And, in untabulated results, we find that the five-factor model of Fama and French (2015) produces similar results.

Additionally, we winsorize CAPM abnormal returns (or the raw returns used in the second premium measure) at the 1st and 99th percentiles each calendar week, using the full cross section of announcer and non-announcer firms. And because the weekly mean returns display no serial correlation, we adjust standard errors only for heteroskedasticity.

Panel A of Table 3 uses the 25-firm requirement to select the weeks that comprise earnings seasons. With this specification, we find a significant return premium of 0.40% using the CAPM and 0.47% using the return spread between announcers and non-announcers ($A-NA$). Panel B uses the 200-firm requirement and finds premium estimates of similar magnitude to Panel A, which are also similar to those found in Table 1.

For the follow-on reports, the point estimates of the announcement premium are all negative, though only the $A - NA$ return spreads are statistically significant. However, at this point, we can readily conclude that the disparity between the initial and follow-on earnings announcements is evident in calendar time as well. Only the initial announcements of an earnings seasons receive the return premium that the literature has struggled to understand.

For additional perspective on the mixed evidence of a negative announcement premium for the follow-on earnings reports in Table 3, we examine the subperiod robustness in Table 4. We require at least 200 firms to announce their earnings for any week to be selected, as this specification seems more consistent with large clusters that would define an earnings seasons. Again, only the $A - NA$ return spread displays a statistically significant negative premium, and this occurs only in the early subperiod. Since 2008, the point estimates of the premium are zero for the follow-on reports.

Finally, note that Table 4 reveals a sizable premium for both measures and both subperiods. Hence, across all the tests we've undertaken, the one robust earnings announcement premium is the positive premium for the initial earnings reports of a quarter.

1.2.1. Announcement premium by week of the season

We now pivot from examining groupings of initial and follow-on earnings reporters. In this section, we examine the week-by-week evolution of the earnings announcement premium across earnings seasons. To be included in the following analysis, each week must have at least 200 firms reporting their earnings (though a 25-firm filter produces qualitatively similar findings).

The first week of each season often falls below 200 announcements per week, as the opening of an earnings season can be on any day of the week. Below shows the mean count of earnings reports per earnings-season week, as well as the number of weeks available after imposing the 200-firm filter (T). The mean announcement count

Week	Mean count	T
1	521	56
2	719	104
3	654	109
4	535	109
5	312	82
6	264	25

per week declines notably after week four.

Using these data, Figure 2 plots the mean announcement-week CAPM abnormal stock return by earnings-season week. The bars indicate a 95% confidence interval around each mean that is calculated from 10,000 resamplings of the data. The first three weeks of earnings seasons, and only the first three weeks, display a significantly positive return premium. These premia range from 0.61% in the first week to 0.33% in the third week. Hence, the main message of this paper is visible in this week-by-week setting as well: After a sufficient number of firms report earnings, the announcement return premium dissipates.[6]

1.2.2. The saturation point of learning

The prior findings are consistent with a systematic risk mechanism driving the earnings announcement premium. Andrei et al. (2023) show how the use of firm-level earnings to extract signals about a common factor in firms' payoffs can increase the priced risk and the risk premium of an announcing firm. In their model, investors seek to reduce their uncertainty about the economy, defined as the variance of the common factor.

The prior analysis finds that the earnings announcement premium dissipates after three weeks. Is three weeks consistent with systematic learning, and hence, systematic risk? In other words, how many weeks are required to learn about a common factor? Surely there is a saturation point at which observing more firms' earnings does not further reduce economic uncertainty. Is three weeks reasonable?

To investigate, we estimate a Bayesian model of learning about a common factor in firms' earnings. Following Andrei et al. (2023), we measure the reduction in the variance of the factor week by week using the actual flow of earnings reports. The model is detailed in section 2, but for now, we will simply note the average weekly reductions in uncertainty a Bayesian investor would have achieved during the 1995-2022 sample period.

Figure 4 plots the mean weekly reductions in uncertainty, as measured by UR_t and defined in section 2. The figure dovetails the earnings announcement premium

[6] This 200-firm specification of earnings-season weeks extends, at times, to a week seven. We omit the week seven mean from the figure because those weeks represent only 3.6% of the total weekly observations. The mean abnormal return in week seven is 0.30% with the largest 95% confidence interval that ranges from −0.15% to 0.74%.

results shown in Figure 2. The largest amount of learning from earnings occurs mostly in week one, where uncertainty falls 58%. Uncertainty falls further in week two by 40%, and then by 19% in week three. Figure 2 shows that the mean announcement premium becomes insignificant in week four, which is matched with a reduction in uncertainty of only 7%.

We are silent on the fair pricing of this economic uncertainty, and hence, whether the levels of uncertainty encountered in the model would require a risk premium throughout the first three weeks of earnings seasons. However, it is comforting to see that one approach to empirically modeling systematic learning from the actual flow of earnings reporting aligns with the structure of the earnings announcement premium over the season. In section 3, we will employ linear regressions to assess whether the reductions in uncertainty achieved from actual earnings flow covary with the weekly announcement premium. And, we will control for unexpected earnings (SUE).

1.3. Pre- and post-announcement returns

Our last exercise before introducing and estimating a model of systematic learning from earnings flow is to broaden the view of the announcement premium to pre- and post-announcement windows. Lamont and Frazzini (2007) and Barber et al. (2013) find a pre-announcement premium for United States and foreign stock returns, respectively, over days $[-10, -2]$ relative to the earnings announcement day. Lamont and Frazzini (2007) find a post-announcement premium over days $[2, 10]$. For the respective groupings of initial and follow-on earnings reports, we examine CAPM abnormal returns over weeks $[-2, -1]$ and $[1, 2]$.

As before, we conduct the testing in both event time and calendar time. In the event-time tests, we calculate holding-period abnormal returns by subtracting a stock's CAPM-based expected holding-period return for a given time window from the stock's realized holding-period return. The alternative is to sum the weekly abnormal returns over the given windows, but holding-period returns mitigate the biases that bid-ask spreads impart on the summing of returns (Blume and Stambaugh, 1983). For the calendar-time analysis, we require that at least 200 firms announce their earnings in week 0 for their pre- and post-announcement returns to be included.

Table 5 reports the event-time findings in Panel A and the calendar-time findings in Panel B. We see no significant abnormal return over weeks $[-2, -1]$. However, we see evidence in both panels of a positive premium over weeks $[1, 2]$ for the initial announcements only. Although not tabulated, we examine the robustness of the post-announcement premium over the two halves of the sample period and find no evidence of a premium in weeks $[1, 2]$ over the 2009-2022 subperiod, with neither the event-time nor the calendar-time analysis.

Given the lack of robustness across the two subperiods, we leave additional consideration of a post-announcement return premium for future research. We can offer a conjecture using a systematic learning channel though. The systematic spillover of earnings news might induce a smaller risk premium in non-announcing stocks whose prices are updated in the spillover (Ben-Rephael et al., 2021). However, because the only evidence of a premium we detect in Table 5 is for the initial reporters, the supposed spillover must affect only those stocks. Perhaps investors pay greater attention to early-reporting firms in general, or perhaps the spillover information is more precise and relevant for the valuations of the early reporters. Regardless, why this proposed spillover premium might fade in the more recent subperiod also seems a sticking point.

2. Earnings model

Earnings announcements may require a risk premium as investors extract systematic news from firms' earnings reports (Savor and Wilson, 2016; Andrei et al., 2023). The finding that the earnings announcement return premium exists only for the initial batches of announcements in an earnings seasons suggests that systematic risk may indeed be a mechanism for the premium.

To investigate further, we introduce a model of firm earnings growth that provides a real-time Bayesian investor's view of uncertainty regarding a common factor in earnings. We refer to uncertainty about this factor as the aggregate economic uncertainty in the system, as Andrei et al. (2023) do. The model generates a measure of this economic uncertainty to start each quarter. As earnings are observed, the uncertainty falls based on how informative newly released earnings reports happen to be. The magnitude by which economic uncertainty is reduced week by week from the actual flow of earnings announcements is then examined as a temporal driver of the earnings announcement premium.

We first describe the laws of motion of the economy that dictate the likelihood function, and then propose a measure of the weekly reductions in economic uncertainty. We assume that the structure of a firm's earnings growth is linear in parameters unknown to the investor, and that earnings growth, denoted X_t, is partly determined by a firm's own loadings on a one-dimensional latent process M_t. We make no claims about the interpretation of M_t other than its role as the common source of variation in earnings growth.

$$X_t = \mu + \beta M_t + \nu_t \tag{1}$$

$$M_t = \delta + \Phi M_{t-1} + \epsilon_t \tag{2}$$

where

- μ is an $N \times 1$ vector of firm-specific average earnings growth.
- β is an $N \times 1$ matrix of firm loadings on the latent process M_t. Denote β_i as firm i's loadings on M_t.
- ν_t is an $N \times 1$ vector of firm-specific mean-zero shocks with variance γ^2.
- δ is the unconditional mean value for the latent process, $E[M_t] = \delta$.
- ϕ is the AR(1) persistence parameter.
- ϵ_t is a white-noise shock to the latent process, with variance σ^2.

The above parameters and states are unknown to a rational Bayesian investor. Denote the parameter vector of the system with $\theta(t) = (\mu, \beta, \sigma, \gamma, \phi, \delta, \epsilon_{1:t})$ where the subscript 1:t denotes the set of data currently available to the real-time investor. We omit the idiosyncratic parameter ν_t.[7] The parameter vector has length $3N+2+(N+t)$.

Note that the latent factor shocks ϵ_t are treated as parameters of the model. We estimate all latent factor shocks to reconstruct the distribution of the latent process M_t in full. Doing this has the advantage of simplifying model estimation at the cost of some extra computational burden.

[7] Omitting "nuisance parameters" like ν_t is a common practice that speeds up inference since it reduces the number of parameters sampled. Note that, conditional on all other parameters, the distribution of X_t is Gaussian with mean $\mu + \beta M_t$ and variance γ. Notably, this allows a modeler to ignore ν_t as a parameter and instead use the variance of ν_t as a part of the likelihood.

Conditional upon observing the realized earnings growth of firms each quarter, the investor updates her beliefs about the joint density of the parameters $P(\theta \mid X_{1:t})$, where $X_{1:t}$ is the filtration $X_{1:t} = \{X_1, X_2, \ldots, X_t\}$. Bayes' rule gives

$$P(\theta \mid X_{1:t}) = \frac{P(X_{1:t} \mid \theta)P(\theta)}{P(X_{1:t})}$$

The explicit posterior given above is unlikely to be analytically tractable due to the complexity in calculating the denominator $P(X_{1:t})$. However, as is common in numerical Bayesian inference, we can state that the posterior is proportional to the likelihood times the prior:

$$P(\theta \mid X_{1:t}) \propto P(X_{1:t} \mid \theta)P(\theta)$$

Recasting posterior inference as a proportionality problem is what allows statements about the empirical distribution of parameters in the model, and, more broadly, is a simplification that permits the use of many modern Bayesian inference methods.[8]

The first component of the posterior, $P(X_{1:t} \mid \theta)$, is referred to as the likelihood. The likelihood function is proportional to the Gaussian PDF centered at $\mu + \beta M_t$ for firms announcing in time t, with variance γ:

$$P(X_{1:t} \mid \theta) \propto \prod_{s=1}^{t} \mathcal{N}(X_s \mid \mu + \beta M_t, \gamma^2 I)$$

The second component of the posterior, the prior $P(\theta)$, is characterized as follows.

- Firm means: $\mu \sim \mathcal{N}(\hat{\mu}, I)$. We set each firm's mean earnings growth to the average of that firm's earnings growth in the first five years of the firm's existence, or for younger firms, to the average earnings growth across all firms from

[8]See Gelman et al. (2013) for a rigorous mathematical treatment of contemporary Bayesian methods, or McElreath (2018) for a more casual text.

1980-1985. This vector is denoted $\hat{\mu}$. The prior variance is set to one, which assumes that about 10% of the 15.52% aggregate earnings growth variance in the sample is due to prior uncertainty in average earnings growth.

- Earnings betas: $\beta_i \sim \mathcal{N}(1, 0.1I)$ for each firm i. We center the prior mean of β_i at one, to predispose firms to having a positive exposure to aggregate earnings growth shocks. The prior variance of 0.1 reflects a relatively strong belief that the model should not permit too many firms to have a cashflow beta of zero.[9]

- Firm idiosyncratic variances, γ_i^2, follow an Inverse Gamma prior, where each firm has a different distribution. To set the priors on firm idiosyncratic risk, we estimate a firm's earnings growth variance in the first five years of the sample period and set the α and β hyperparameters of the Inverse Gamma distribution to the pair that makes the mean of the distribution equal to half the firm's estimated earnings growth variance. Firms entering later in the sample begin with a prior equal to the aggregate earnings growth variance. The choice of prior is equivalent to assuming that half of a firm's observed earnings growth variance is idiosyncratic.

- Latent parameter mean: $\delta \sim \mathcal{N}(0, 0.5)$. We assume that, on average, aggregate earnings growth shocks are close to zero in line with Kothari et al. (2006).

- Latent variance: $\sigma^2 \sim$ Inverse Gamma$(10, 143.83)$. We estimate the variance of all firm's earnings growth in the first five years of the sample period, and set the α and β hyperparameters of the Inverse Gamma distribution to the pair that sets the variance of the latent shock equal to the empirical variance.

[9]If the prior variance on β_i is too large, the posterior distribution for β_i could be centered tightly around zero, and the latent factor could be removed entirely as a source of variation. Earlier versions of the model used a variance of 1, and many estimates of β_i were tightly centered around zero. This had the unfortunate byproduct of inflating firm idiosyncratic risk γ_i^2 as well.

- Latent process persistence: $\phi \sim \mathcal{N}(0, 0.2) \in [-1, 1]$. We assume that earnings growth persistence is on average centered around zero to permit negative autocorrelation, though empirically there is some evidence that persistence is as high as 0.7 (Kothari et al., 2006). We truncate the distribution between -1 and 1 to rule out AR(1) processes that are not covariance stationary.

- Latent process shocks: $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$. We assume latent process shocks are zero on average, with a hierarchical prior parameter of σ.

2.1. Measuring reductions in uncertainty

After observing a firm's earnings growth, the model infers the value of the parameters governing the latent factor M_t and the history of the common shocks ϵ_t. Firms differ on the ability of their earnings to reduce uncertainty. For example, a firm with a large idiosyncratic earnings growth variance γ_i^2 relative to the volatility of aggregate earnings would offer a noisy signal about the current common shock.

Suppose a Bayesian investor were attempting to form a nowcast of the unobserved latent process M_t after observing firm earnings growth $X_{i,t}$. The Bayesian investor's forecast of economic uncertainty is $\mathrm{Var}[M_t \mid X_{i,t}]$. Lemma 1 derives this forecast and its variance.

Lemma 1. The variance of the latent process M_t conditional upon observing the earnings growth $X_{i,t}$ of an announcing firm j is[10]

$$\mathrm{Var}[M_t \mid X_{j,t}] = E_{\hat{\theta}}\big[\mathrm{Var}[M_t \mid X_{i,t}, \hat{\theta}]\big] + \mathrm{Var}_{\hat{\theta}}\big[E[M_t \mid X_{i,t}, \hat{\theta}]\big]$$

[10]This follows from the law of iterated expectations: $\mathrm{Var}[X] = E\big[\mathrm{Var}[X \mid Y]\big] + \mathrm{Var}\big[E[X \mid Y]\big]$

where

$$\mathrm{Var}[M_t \mid X_{i,t}, \hat{\theta}] = \frac{\sigma^2}{1-\phi^2}\left(\frac{\beta_i^2\sigma^2}{\gamma_i^2(1-\phi^2)} + 1\right)^{-1} \tag{3}$$

$$E[M_t \mid X_{i,t}, \hat{\theta}] = \delta + \sum_{s=1}^{t-1}\phi^s(\delta + \epsilon_{t-s}) + \left(1 + \frac{\gamma_i^2(1-\phi^2)}{\beta_i\sigma^2}\right)^{-1}\left(\beta_i\sum_{s=1}^{t-1}\phi^s\epsilon_{t-s} + \nu_{j,t}\right) \tag{4}$$

Lemma 1 is straightforward to compute numerically, once equipped with posterior samples of $\hat{\theta}^k$. Simply computing Equations 3 and 4 for all posterior samples $\hat{\theta}^k$ and calculating the sum of the sample average (of Equation 3) and variance (of Equation 4) yields the appropriate posterior $\mathrm{Var}[M_t \mid X_{j,t}]$.

We construct a measure that succinctly describes the information content that earnings provide about the latent factor. We denote the reduction in economic uncertainty after observing firm i's earnings growth as

$$UR_{i,t} = 1 - \frac{\mathrm{Var}_t[M_t \mid X_{i,t}]}{\mathrm{Var}_t[M_t]} \tag{5}$$

The denominator is the measure of uncertainty before observing the earnings growth for firm i. The ratio on the righthand side is bounded between 0 and 1, with a greater ratio indicating a weaker signal about common earnings. Subtracting from one means that larger values of uncertainty reduction, $\mathrm{UR}_{i,t}$, indicate greater reductions in economic uncertainty.

We have described the measure for firm i, but we can easily view the Bayesian updating and the measure of uncertainty reduction for any set of firms. And we will focus in this paper on the estimate of $\mathrm{UR}_{i,t}$ achieved by observing the set of firms that announce their earnings in any given week.

2.2. Estimation overview

The sample of announcements used in the Bayesian estimation of the earnings model differs from the data sample in the prior section in a few ways. By starting this sample in 1980, we allow for a 15-year burn-in period before testing the model's output on the sample of announcement returns. Having such a long burn-in diminishes the impact of our choices for prior beliefs. Secondly, we exclude firms with a market capitalization of equity below $500 million. This reduces the number of firms in the system, and hence the computational demands. Additionally, we do not exclude earnings that are announced before day 15 of a quarter or after day 60. This allows the Bayesian investor to learn from any earnings released outside of earnings seasons.

We collect quarterly earnings reports from 1980 through the first quarter of 2022, which includes 1,535 firms. The sample is drawn from the CRSP/Compustat Merged Database using all domestic firms with a stock ownership code of zero (STKO), to exclude ADRs. We use only firms with fiscal quarters ending in March, June, September, and December to simplify the modeling exercise. Firm-quarters without income before extraordinary items (IBQ), common shares outstanding (CSHOQ), quarterly end price (PRCCD), or total assets (ATQ) are excluded, as are firms with a market cap below $500 million at quarter end.

We calculate seasonally-adjusted earnings growth $X_{i,t}$ using the method in Savor and Wilson (2016) and Kothari et al. (2006). Firm earnings growth is calculated as the difference between a firm's earnings (IBQ) and the firm's earnings four quarters prior, scaled by the firm's market cap in the previous quarter. Firms are thus required to have at least four quarters of observations before they can be included in the sample. We multiply earnings growth by 100 to reduce the risk of numerical errors that can arise with small numbers. All parameters should thus be interpreted in terms of percentage points (i.e. $\mu_i = 1.5$ indicates an average earnings growth of 1.5%).

The model is written in the probabilistic programming language Turing.jl. We perform Markov Chain Monte Carlo (MCMC) sampling using the No U-Turn Sampler (NUTS), a Hamiltonian Monte Carlo method. NUTS is considered to be an easy-to-use MCMC algorithm that scales well to high dimensional problems (Hoffman and Gelman, 2014). NUTS obviates much of the difficulties of Bayesian inference for high-dimensional problems like the one we have presented. The Appendix provides details of the estimation techniques, an overview of some model diagnostics, and examples of the time-series evolution of firm-level parameters.

Lastly, before moving to tests of how the weekly reductions in economic uncertainty covary with the earnings announcement premium, readers may like to see the level of aggregate uncertainty in the system as the Bayesian investor starts each quarter. Figure 3 plots $\sqrt{Var(M_t)}$ each quarter before any earnings reports are observed. Although not shown, uncertainty starts at over 30% in 1985 and then stabilizes near 20% around 1990. The late 1990s are a time of relatively lower uncertainty, and outside of the third quarter of 2015, uncertainty is largely between 18% to 20%.

3. Systematic learning and the announcement premium

In this section, we examine whether the reduction in economic uncertainty arising from the weekly flow of earnings news covaries with the earnings announcement premium. To the extent that the announcement premium reflects a reward for bearing systematic risk, weeks that convey more systematic information should have a greater announcement premium.

Estimates of the announcement premium are from section 1.2.1. UR_t is the measure each week's systematic news, and is the incremental weekly reduction in the posterior variance of the common earnings factor. The Bayesian investor updates her views of the full set of the model's parameters each week after observing all the earnings news of the week.

A regression framework will also afford us the opportunity to control for the level of the earnings surprises that are realized each week. As noted earlier, the literature recognizes a tendency for firms to advance (delay) their announcements when their earnings surprises are stronger (weaker). Therefore, we estimate standardized unexpected earnings (SUE) for each earnings announcement as done by Livnat and Mendenhall (2006).[11] We require IBES to have two analysts' forecasts to calculate SUE, and we winsorize SUE each quarter at the 1^{st} and 99^{th} percentiles.

When controlling for SUE, we restrict the analysis to the subset of earnings announcements having a measure of SUE available. Roughly 60% of announcements can be paired with a SUE, so we investigate several filters for the minimum number of firms announcing their earnings each week, namely 25, 100, and 200 firms respectively. While we expressed a preference for the 200-firm requirement in the prior section, the data loss when requiring 200 announcements to have a SUE available each week is substantial, as we will discuss.

Lastly, we also investigate whether the earnings announcement premium covaries with VIX. Andrei et al. (2023) employ VIX as a proxy for the level of economic uncertainty, and hence, investors' demand for earnings information. They find that investors pay more attention to earnings on high-VIX days, and that earnings response coefficients are also greater on these days.

[11] We thank Freda Song Drechsler for providing the python code to estimate SUE on her website, https://www.fredasongdrechsler.com/data-crunching/pead.

All variables in the regression models of Table 6 are standardized. The regression residuals display mild serial correlation in columns (1) and (3), having a significant AR1 coefficient near 0.18, while the residuals of columns (2) and (4) have a lower and insignificant AR1 coefficient. The Newey-West HAC model with one lag is used to estimate standard errors in all models of the table. Using lags of three or six weeks does not affect the inferences.

Column (1) of Table 6 reports the results for time-series regressions of the announcement-week mean abnormal returns on UR_t. Requiring 25, 100, or 200 earnings reports in each week has little effect on the results. The coefficient in each panel is positive at a one-percent level of significance. A one standard deviation reduction in economic uncertainty is associated with a 0.15 increase in the standard deviation of the announcement return premium in Panel A, and increases just slightly in the other panels. The economic magnitude of these coefficients are very stable across the panels, implying a 0.22% to 0.24% increase in the return premium when UR_t increases one standard deviation.

For column (2), each earnings announcement is required to have a measure of SUE. Panels A and B find a statistically positive relation between the weekly means of SUE and the announcement return premium. However, when requiring 200 announcements to have SUE, the number of weeks available falls from 503 in column (1) to 328 in column (2). The loss of power hinders the ability to detect a relation between SUE and the announcement premium. Hence, we will examine subperiod variation in the regression findings of Table 6 using a 100-firm requirement in the next table.

But, before that, note that VIX displays no explanatory power for the announcement premium in any specifications of column (3). We have also examined the weekly change in VIX and a high-VIX dummy set to one when VIX is above its 75th percentile, and zero otherwise. These alternative specifications display no covariance with the announcement premium either. This VIX result perhaps not so surprising after observing the changes in the earnings announcement premium that occur within an earnings seasons (e.g., Figure 2). VIX simply does not display within-season dynamics to match this. However, UR_t does. VIX may track variation in the return premium at lower frequencies, which we leave for future research to consider.

Column (4) shows results for multiple regressions with all three explanatory variables. Panels A and B reveal that the temporal comovement between UR_t and the announcement premium is largely orthogonal to temporal variation in SUE. Again, Panel C seems to suffer from power deficiencies, as SUE has no effect either when requiring at least 200 firms each week.

We examine subperiod robustness in Table 7 using the 100-firm requirement on the number of announcers each calendar week. Column (1) detects evidence of UR_t covarying with the earnings announcement premium in each subperiod, albeit with a 10% significance level in the latter subperiod.[12] After controlling for SUE and VIX, column (4) finds no statistical evidence of comovement between UR_t and the premium. But again, requiring SUE for each announcement reduces the number of weeks available for analysis. Note that even SUE is not robust within both the subperiods.

Overall, we find results that are consistent with the central message of Andrei et al. (2023). When the systematic information provided by earnings reports is greater, the risk premium for the announcing firms is also greater.

[12]The 25-firm filter also reveals a covariance in each subperiod. But the first subperiod in that case has the 10% significance level while the second subperiod has the 5% significance level.

4. Price updating during earnings seasons

The spillover of systematic information from earnings reports to the stocks of non-announcing firms underpins our interpretations of the findings thus far, as the spillover creates the conditional systematic risk exposures of the announcing firms. Mapping the pathways of the information spillovers is nontrivial, and we leave mapping attempts for future research.[13] Instead, we focus here on how intensely prices are being updated within the initial and follow-on weeks of earnings seasons respectively.

We expect the information content of earnings to be greater during the initial weeks of earnings seasons than during follow-on weeks because initial announcements should convey both idiosyncratic and systematic news, while follow-on announcements should convey only idiosyncratic news. Following prior research, we utilize abnormal idiosyncratic risk to gauge the information content of earnings. The applicability of the measure comes more from the "abnormal" aspect rather than the "idiosyncratic" one, as we explain shortly.

Abnormal volatility of daily residual returns (AIVOL) is calculated similarly to the design of Barber et al. (2013), aggregated to the weekly frequency. To create an event-week estimate of idiosyncratic risk for each stock, daily CAPM abnormal returns are squared and summed over the days of the given week. This event-week variance is scaled by the estimate of weekly residual variance obtained from the regressions to estimate the betas in section 1. Lastly, we take the square root of this ratio of event-week variance to prior-period variance, and then subtract one so that an AIVOL of

[13] Ben-Rephael et al. (2021) provide evidence on how firm-level earnings news spills over to other stock prices, using (institutional) investors' firm-level search activity on Bloomberg. They do not focus their analysis on the sizes of the spillover networks nor how networks change over earnings seasons, as we would require here.

zero means no abnormal volatility. AIVOL is winsorized each calendar week at the 1% and 99% levels.

Systematic learning from earnings induces conditional increases in the systematic risk of firms announcing their earnings (Savor and Wilson, 2016; Andrei et al., 2023). Using unconditional beta estimates to measure AIVOL implies that any conditional changes in systematic risk are embedded in the measure. This interpretation of AIVOL pairs with our use of unconditional betas to estimate the announcement return premium in earlier sections. The mean abnormal return of an unconditional CAPM, we presume, reflects the conditional risk premium associated with announcing earnings.[14]

Because AIVOL displays marked skewness, we examine both the mean and median values of AIVOL in calendar time. And we bootstrap the distribution of each test statistic using 10,000 resamplings. Panel A of Figure 5 displays the mean of the weekly AIVOL means, while Panel B displays the median of the weekly AIVOL medians. To be considered an earnings announcement week in this calendar-time analysis, a week must have at least 200 announcers.

In both panels, we see that AIVOL increases during announcement weeks, consistent with prior research. We can also see that AIVOL falls notably from the initial weeks to the follow-on weeks. In panel A, the mean AIVOL declines on average from 37% to 31%, a 16 percent relative decline from initial to follow-on weeks. In panel B, the median AIVOL declines on average from 14% per year to 4%, a 71 percent relative decline. The bootstraps reject the null hypothesis that AIVOL is equal across

[14]Viewing CAPM alpha and AIVOL as capturing a component of priced systematic risk seems a simpler interpretation to us for an announcement premium than having market frictions induce the pricing of truly idiosyncratic risk (e.g., Merton, 1987; Cohen et al., 2007). Nevertheless, priced idiosyncratic risk is inconsistent with our finding that the earnings announcement premium does not exist in the follow-on weeks of earnings seasons, as abnormal idiosyncratic risk remains high in those weeks too.

the initial and follow-on weeks at the 5% level in Panel A, and the 1% level in Panel B. In sum, these findings are consistent with the initial weeks of earnings seasons revealing more value-relevant information than the follow-on weeks.

Interestingly, the variability in SUE across initial and follow-on weeks is opposite to the pattern of AIVOL. The mean of the weekly cross-sectional standard deviations of SUE is 0.0054 during initial weeks, and rises to 0.0080 during the follow-on weeks. So, dispersion in earnings surprises is not driving the greater AIVOL in the initial weeks.

For completeness, Figure 5 also shows AIVOL for the non-announcing firms within each week type. There is no statistical difference between AIVOL for the non-announcers in the initial weeks and the follow-on weeks in either panel. In fact, the average AIVOL for the non-announcing firms during earnings weeks is not statistically different from AIVOL for these firms in weeks outside of earnings seasons. From these findings, we conclude that the spillover of earnings information from announcers to non-announcers is of similar magnitude to the flow of non-earnings information that generally moves stock prices.

Overall, this section adds to the literature by showing that the value-relevant information flows during the initial weeks of earnings seasons are greater than those during the follow-on weeks. The incremental difference is consistent with a systematic component of earnings information existing only for the initial announcements of an earnings season.

5. Conclusion

Savor and Wilson (2016) and Andrei et al. (2023) show that the use of firm-level earnings by investors to learn about a common factor in earnings induces a conditional risk premium into the stock prices of the firms whose earnings investors are using to systematically learn about the non-announcing firms. Their insights offer a novel explanation for the earnings announcement premium. We test the ability of a systematic learning channel to explain the premium in several ways.

First, we document that the announcement premium only exists in the initial weeks of each quarter's earnings season. This is consistent with a systematic learning channel because the systematic information of earnings reports must eventually become redundant after many hundreds of firms have reported earnings. Secondly, we measure a real-time Bayesian investor's view of the variance of the common earnings factor. This posterior variance captures economic uncertainty in the system. Weekly reductions in this uncertainty covary with the weekly time-series of earnings announcement premiums, controlling for the level of SUE . These findings strongly support a systematic learning channel to the announcement premium.

Moreover, our findings imply the importance of observing the placement of an earnings report within the season when studying the informativeness and efficiency of stock prices. Larger risk premiums for the initial announcements can, for example, disrupt measurements of the price responses to earnings.

References

Andrei, Daniel, Henry Friedman, and N. Bugra Ozel, 2023, Economic uncertainty and investor attention, *Journal of Financial Economics* 149, 179–217. [page 1], [page 2], [page 8], [page 14], [page 17], [page 25], [page 27], [page 29], [page 30]

Ball, Ray, and S. P. Kothari, 1991, Security returns around earnings announcements, *The Accounting Review* 66, 718–738. [page 9]

Barber, Brad M., Emmanuel T. De George, Reuven Lehavy, and Brett Trueman, 2013, The earnings announcement premium around the globe, *Journal of Financial Economics* 108, 118–138. [page 1], [page 3], [page 9], [page 11], [page 15], [page 28]

Beaver, William H, 1968, The information content of annual earnings announcements, *Journal of accounting research* 67–92. [page 1]

Ben-Rephael, Azi, Bruce I. Carlin, Zhi Da, and Ryan D. Israelsen, 2021, Information Consumption and Asset Pricing, *The Journal of Finance* 76, 357–394. [page 16], [page 28]

Betancourt, Michael, 2021, A Short Review of Ergodicity and Convergence of Markov chain Monte Carlo Estimators. [page 35]

Blume, Marshall E., and Robert F. Stambaugh, 1983, Biases in computed returns: An application to the size effect, *Journal of Financial Economics* 12, 387–404. [page 16]

Chambers, Anne E, and Stephen H Penman, 1984, Timeliness of reporting and the stock price reaction to earnings announcements, *Journal of accounting research* 21–47. [page 9]

Chan, Kam Fong, and Terry Marsh, 2022, Asset pricing on earnings announcement days, *Journal of Financial Economics* 144, 1022–1042. [page 3]

Chari, Varadarajan V, Ravi Jagannathan, and Aharon R Ofer, 1988, Seasonalities in security returns: The case of earnings announcements, *Journal of Financial Economics* 21, 101–121. [page 1]

Cohen, Daniel A, Aiyesha Dey, Thomas Z Lys, and Shyam V Sunder, 2007, Earnings announcement premia and the limits to arbitrage, *Journal of Accounting and Economics* 43, 153–180. [page 1], [page 3], [page 5], [page 7], [page 9], [page 29]

Dellavigna, Stefano, and Joshua M. Pollet, 2009, Investor Inattention and Friday Earnings Announcements, *The Journal of Finance* 64, 709–749. [page 4]

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22. [page 11]

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin, 2013, *Bayesian Data Analysis* (CRC Press). [page 19], [page 35], [page 36]

Gelman, Andrew, and Donald B. Rubin, 1992, Inference from Iterative Simulation Using Multiple Sequences, *Statistical Science* 7. [page 36]

Givoly, Dan, and Dan Palmon, 1982, Timeliness of annual earnings announcements: Some empirical evidence, *Accounting review* 486–508. [page 9]

Hoffman, Matthew D, and Andrew Gelman, 2014, The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo, *Journal of Machine Learning Research* 15, 31. [page 24], [page 34], [page 35]

Johnson, Travis L, and Eric C So, 2018, Time will tell: Information in the timing of scheduled earnings news, *Journal of Financial and Quantitative Analysis* 53, 2431–2464. [page 2], [page 3], [page 5], [page 7]

Kothari, S P, Jonathan Lewellen, and Jerold B Warner, 2006, Stock returns, aggregate earnings surprises, and behavioral finance, *Journal of Financial Economics* 32. [page 20], [page 21], [page 23]

Lamont, Owen A, and Andrea Frazzini, 2007, The earnings announcement premium and trading volume, *NBER working paper* . [page 1], [page 3], [page 9], [page 15]

Livnat, Joshua, and Richard R. Mendenhall, 2006, Comparing the Post–Earnings Announcement Drift for Surprises Calculated from Analyst and Time Series Forecasts, *Journal of Accounting Research* 44, 177–205. [page 25]

McElreath, Richard, 2018, *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (CRC Press). [page 19]

Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *The Journal of Finance* 42, 483–510. [page 29]

Noh, Suzie, Eric C. So, and Rodrigo S. Verdi, 2021, Calendar rotations: A new approach for studying the impact of timing using earnings announcements, *Journal of Financial Economics* 140, 865–893. [page 4], [page 9]

Savor, Pavel, and Mungo Wilson, 2016, Earnings Announcements and Systematic Risk: Earnings Announcements and Systematic Risk, *The Journal of Finance* 71, 83–138. [page 1], [page 8], [page 11], [page 17], [page 23], [page 29], [page 30]

Vehtari, Aki, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner, 2021, Rank-Normalization, Folding, and Localization: An Improved $\hat{R}$ for Assessing Convergence of MCMC (with Discussion), *Bayesian Analysis* 16, 667–718. [page 36]

6. Appendix

Monte Carlo Markov Chain (MCMC) methods vary greatly by the specific algorithm employed, but in general (most) methods follow a common procedure:

1. Initialize model parameters θ_0 to some value (typically random draws from the prior).

2. Propose a new parameterization $\tilde{\theta}$, chosen stochastically by a MCMC kernel $K(\theta_0)$. In this case, this kernel is the NUTS.

3. Accept the proposed $\tilde{\theta}$ with some probability given by the kernel $K(\theta, \theta_0)$. If accepted, set the draw to $\theta_1 = \tilde{\theta}$. If rejected, set $\theta_1 = \theta_0$.

4. Repeat steps 2-4 sequentially until a predetermined number of samples are drawn.

The kernel $K(\cdot)$ is replaceable with numerous MCMC algorithms. For example, a common MCMC kernel is a random-walk Metropolis-Hastings, which proposes a new parameterization by perturbing the old parameterization with a Gaussian shock, i.e. $\tilde{\theta} \sim \mathcal{N}(\theta_{n-1}, \Sigma)$ for a fixed proposal variance Σ. The new proposal $\tilde{\theta}$ is accepted with probability $\alpha = P(\tilde{\theta})/P(\theta_{n-1})$, where $P(\tilde{\theta})$ is the joint density of the model (prior probability of $\tilde{\theta}$ multiplied by the likelihood evaluated at $\tilde{\theta}$). The kernel we employ, NUTS, is a significantly more complex statistical construction that treats parameters as if they were high-energy particles with momentum, traversing the probability space of the problem. We refer the reader to Hoffman and Gelman (2014) for detailed information about NUTS.

For each quarter, we generate a vector of earnings growth for firms announcing in the quarter. We then estimate a new model for each new quarter of data, where each model incorporates all past quarterly data. For example, the 16th model uses 16 quarters' worth of earnings growth data for all firms that announced. The 17th model has 17 quarters of data available, for example. Each model contains 4,000 samples (1,000 across 4 chains set to random initial values), with an additional 1,000 samples per chain at the beginning of sampling for burn-in and adaptation. The 1,000 burn-in samples are discarded after sampling is completed as they are generally non-representative of the true posterior.

The MCMC kernel, NUTS, is designed to have few degrees of freedom for the researcher. The two sampling parameters we modify are the target acceptance rate and the maximum tree depth. The target acceptance rate is the proportion of proposed samples that the sampler should accept. Typically, NUTS handles higher acceptance rates by making smaller proposals that are close to the previous sample. We set the target acceptance rate to 80% in accordance with the suggestion in Hoffman and Gelman (2014). The other parameter, maximum tree depth, limits the number of points used by NUTS to find a new sample. Higher limits for tree depth permit more significant jumps from sample to sample with higher acceptance rates; but due to computational constraints, we select a slightly lower tree depth of 5 than the default of 7. Lowering the maximum tree depth does not make the samples less accurate or incorrect asymptotically, but it does mean that the sampler will generally converge slower on average.

Under certain conditions, MCMC methods asymptotically guarantee that all parameters are drawn from the posterior joint density $P(\theta \mid X_{1:t})$.[15] In particular, the distribution of the sampled values θ^n approximate the distribution $P(\theta \mid X_{1:t})$ as the

[15] See Betancourt (2021) for a brief overview of finite-sample convergence in MCMC samples, or Gelman et al. (2013) for a more comprehensive text on MCMC in general.

number of samples reach infinity. In the absence of asymptotic draws, however, statisticians must rely on various diagnostics to indicate that a Markov chain has reached a stationary point. Various measures have been proposed, such as the Gelman-Rubin diagnostic of Gelman and Rubin (1992). The Gelman-Rubin diagnostic, commonly referred to as $\hat{R}$, tests whether there is sufficient evidence to conclude that the first half of a Markov chain has the same mean as the second half. In an equilibrium state, a Markov chain should have the same mean at every point in the sample.[16] A second diagnostic is to conduct a visual inspection of trace plots for samples to verify that the samples appear stationary.

To verify that the quarterly models are likely to be valid MCMC samples (in the sense that their sampling distribution approximates the true posterior density), we visually check trace plots for a sample of parameters. A common diagnostic technique in Bayesian inference is to verify that the trace plot of a parameter displays no obvious autocorrelation and appears to follow a stationary distribution. Figure 6 demonstrates a randomly selected example of a trace and histogram plots for the parameters of IBM using 163 quarters of earnings growth. The trace plots on the left are good examples of a "caterpillar plot", wherein there seems to be little autocorrelation or trend in the trace. A review of 100 trace plots from different chains indicates that we can proceed under the assumption that the samples are directly drawn from a sufficiently close approximation of $P(\theta|X_{1:t})$.

An additional diagnostic tool employed in Bayesian methods is $\hat{R}$ (Gelman and Rubin (1992)). Gelman et al. (2013) notes that an $\hat{R}$ less than 1.1 is generally considered acceptable, though there is no hard-and-fast statistical criterion for whether

[16] Vehtari et al. (2021) suggests that $\hat{R}$ is an inadequate convergence diagnostic when the chain has a heavy tail or when sample variance is different across chains. The model is a linear model with Gaussian and inverse-gamma priors, which does not have the complex geometry of models studied in Vehtari et al. (2021).

a chain has converged. We calculate that the maximum $\hat{R}$ across all parameters and models is 1.042, indicating that all parameters are drawn from a stationary distribution.[17] The result of these tests allows us to proceed with the analysis, comfortable that sampling error due to chains that have not converged is unlikely to be driving the results.

All variables in the model have distributional (at a single point in time) and time-series interpretation (distributions across time). Figure 7 shows the evolution of firm-specific characteristics for the Walt Disney company from 1984 to 2021. Firm parameters vary meaningfully over the sample period. As an example, note that Panel B of Figure 7 demonstrates a steady decline in Disney's loading on the common factor β_i, though as of 2020q2, this trend reverses and suggests a large increase in the cashflow beta. A significant proportion of Disney's earnings comes from theatrical film releases and theme park attendance, both of which declined substantially during the onset of the COVID-19 pandemic.

[17]It is not surprising that the chains demonstrate as good a performance as they do. The model is linear, has simple priors, and uses a relatively flat hierarchical structure.

Table 1 Announcement-week mean CAPM abnormal returns of the *initial* earnings announcements each quarter. Panel A examines all announcements, while Panel B examines only the subset of announcements that are deemed to be on-time. Standard errors are clustered by firm and by calendar quarter. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively.

Period	Firms	AR (%)	t-statistic	N
A. All announcements				
1995-2022	All	0.42***	4.41	103,691
	LG	0.16**	2.36	40,491
	SM	0.59***	4.56	63,200
1995-2008	All	0.52***	3.54	58,390
	LG	0.21*	1.91	21,417
	SM	0.70***	3.61	36,973
2009-2022	All	0.30***	2.73	45,301
	LG	0.11	1.39	19,074
	SM	0.43***	2.92	26,227
B. On-Time announcements				
1995-2022	All	0.39***	4.44	50,124
	LG	0.17**	2.42	22,304
	SM	0.56***	4.57	27,820
1995-2008	All	0.47***	3.50	25,931
	LG	0.22**	2.04	11,169
	SM	0.67***	3.47	14,762
2009-2022	All	0.29***	2.78	24,193
	LG	0.13	1.34	11,135
	SM	0.43***	3.09	13,058

Table 2 Announcement-week mean CAPM abnormal returns of the *follow-on* announcements each quarter. Panel A examines all announcements, while Panel B examines only the subset of announcements that are deemed to be on-time. Standard errors are clustered by firm and by calendar quarter. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively.

Period	Firms	AR (%)	t-statistic	N
A. All announcements				
1995-2002	All	−0.13	−1.31	111,152
	LG	0.04	0.47	22,516
	SM	−0.17	−1.53	88,636
1995-2008	All	−0.13	−0.92	66,040
	LG	0.22**	2.04	11,647
	SM	−0.20	−1.33	54,391
2009-2022	All	−0.13	−0.98	45,112
	LG	−0.16	−1.54	10,869
	SM	−0.13	−0.78	34,243
B. On-Time announcements				
1995-2022	All	0.07	0.66	38,062
	LG	0.07	0.76	9,004
	SM	0.07	0.54	29,058
1995-2008	All	0.14	1.00	18,865
	LG	0.26*	1.72	3,935
	SM	0.11	0.65	14,930
2009-2022	All	0.00	0.02	19,197
	LG	−0.07	−0.61	5,069
	SM	0.03	0.15	14,128

Table 3 Announcement-week mean abnormal returns in calendar time. Panel A shows results when requiring at least 25 announcements in each calendar week. Panel B shows results when requiring at least 200 announcements in each week. The abnormal return (AR) is measured either using the CAPM or the raw return of announcing firms in excess of the raw return of non-announcing firms in the same week (A−NA). Standard errors are adjusted for heteroskedasticity using White's (1980) method. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively. T is the number of calendar weeks meeting the respective filter in each panel.

Measure	Rank	AR(%)	t-statistic	T
A. Requiring at least 25 reporters each week				
CAPM	Initial	0.40***	4.64	215
	Follow-On	−0.16	−1.63	343
A−NA	Initial	0.47***	8.76	215
	Follow-On	−0.19***	−2.94	343
B. Requiring at least 200 reporters each week				
CAPM	Initial	0.49***	5.41	169
	Follow-On	−0.04	−0.41	240
A−NA	Initial	0.51***	8.70	169
	Follow-On	−0.10*	−1.80	240

Table 4 Announcement-week mean abnormal returns in calendar time. We require at least 200 reporters in each week. Panel A shows results for the early subperiod, and Panel B the later. The abnormal return (AR) is measured either using the CAPM or the raw return of announcing firms in excess of the raw return of non-announcing firms in the same week (A−NA). Standard errors are adjusted for heteroskedasticity using White's (1980) method. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively. T is the number of calendar weeks meeting the respective filter in each panel.

Measure	Rank	AR(%)	t-statistic	T
A. 1995-2008				
CAPM	Initial	0.64***	4.37	82
	Follow-On	−0.05	−0.47	148
A−NA	Initial	0.51***	6.32	82
	Follow-On	−0.16**	−2.57	148
B. 2009-2022				
CAPM	Initial	0.36***	3.26	87
	Follow-On	−0.01	−0.06	92
A−NA	Initial	0.50***	6.02	87
	Follow-On	0.00	0.05	92

Table 5 Mean abnormal returns over pre-announcement and post-announcement weeks. Panel A averages CAPM abnormal holding-period returns in event time, clustering the standard errors by firm and calendar quarter. Panel B averages CAPM weekly abnormal returns in calendar time requiring at least 200 reporters in each week, adjusting the standard errors for heteroskedasticity using White's (1980) method. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively. N is the number of observations used in Panel A. T is the number of calendar weeks meeting the respective filter in Panel B.

Window	Rank	AR(%)	t-statistic	N/T
A. Event time				
$[-2, -1]$	Initial	−0.03	−0.21	103,691
	Follow-On	−0.26	−1.50	111,152
$[1, 2]$	Initial	0.29**	2.05	103,691
	Follow-On	−0.07	−0.38	111,152
B. Calendar time				
$[-2, -1]$	Initial	−0.01	−0.14	278
	Follow-On	−0.08	−1.31	350
$[1, 2]$	Initial	0.15**	2.40	278
	Follow-On	0.04	0.56	350

Table 6 Weekly calendar-time regressions of mean announcement-week CAPM abnormal returns on UR_t and control variables. UR_t is the incremental reduction in uncertainty achieved each week by observing the announced earnings. Panel A requires that at least 25 firms report earnings in a given week; panels B and C require at least 100 and 200 firms respectively. Columns 2 and 4 require each announcing firm to have an estimate of SUE. Newey-West standard errors with one lag are in parentheses. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively. T is the number of calendar weeks available.

	Mean Abnormal Return			
	(1)	(2)	(3)	(4)
A. Requiring at least 25 reporters each week				
UR_t	0.153***			0.103**
	(0.047)			(0.043)
Mean SUE_t		0.159***		0.136**
		(0.061)		(0.059)
VIX_{t-1}			−0.066	0.012
			(0.075)	(0.055)
R-squared	0.024	0.025	0.004	0.035
T	652	575	652	575
B. Requiring at least 100 reporters each week				
UR_t	0.174***			0.099**
	(0.046)			(0.050)
Mean SUE_t		0.220***		0.226***
		(0.060)		(0.054)
VIX_{t-1}			0.013	0.073
			(0.066)	(0.064)
R-squared	0.030	0.048	0.000	0.063
T	562	470	562	470
C. Requiring at least 200 reporters each week				
UR_t	0.177***			0.075
	(0.047)			(0.054)
Mean SUE_t		0.116		0.119
		(0.082)		(0.080)
VIX_{t-1}			0.015	0.132*
			(0.073)	(0.075)
R-squared	0.031	0.013	0.000	0.036
T	503	328	503	328

Table 7 Weekly calendar-time regressions of mean announcement-week CAPM abnormal returns on UR_t and control variables. At least 100 firms must announce earnings in a given week. UR_t is the incremental reduction in uncertainty achieved each week by observing the announced earnings. Panel A reports results for the first subperiod, and Panel B the later subperiod. Columns 2 and 4 require each announcing firm to have an estimate of SUE. Newey-West standard errors with one lag are in parentheses. Significance at the 1%, 5%, and 10% levels are denoted with ***, **, and *, respectively. T is the number of calendar weeks available.

	Mean Abnormal Return			
	(1)	(2)	(3)	(4)
A. 1995-2008				
UR_t	0.232***			0.102
	(0.062)			(0.081)
Mean SUE		0.231*		0.216
		(0.136)		(0.143)
VIX_{t-1}			−0.042	0.019
			(0.086)	(0.074)
R-squared	0.054	0.053	0.002	0.064
T	308	234	308	234
B. 2009-2022				
UR_t	0.116*			0.096
	(0.070)			(0.065)
Mean SUE		0.257***		0.269***
		(0.068)		(0.056)
VIX_{t-1}			0.090	0.172**
			(0.084)	(0.081)
R-squared	0.014	0.066	0.008	0.101
T	254	236	254	236

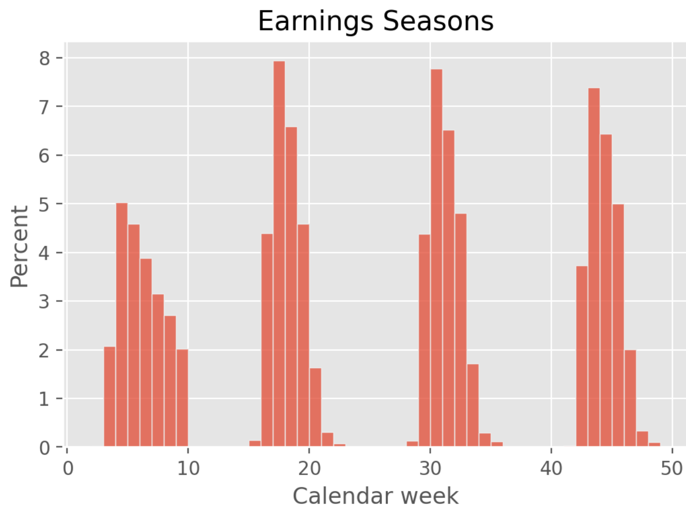

Figure 1. Histogram of earnings announcements by calendar week of the year.

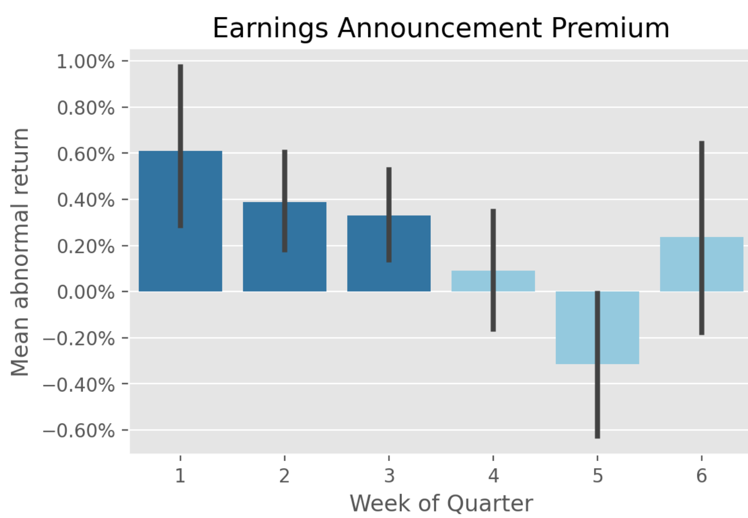

Figure 2. Mean CAPM abnormal stock returns of firms announcing in each week of a quarter. A given week must have at least 200 reporting firms to be included. The bars indicate bootstrapped 95% confidence intervals. Dark blue weeks represent means that statistically differ from zero.

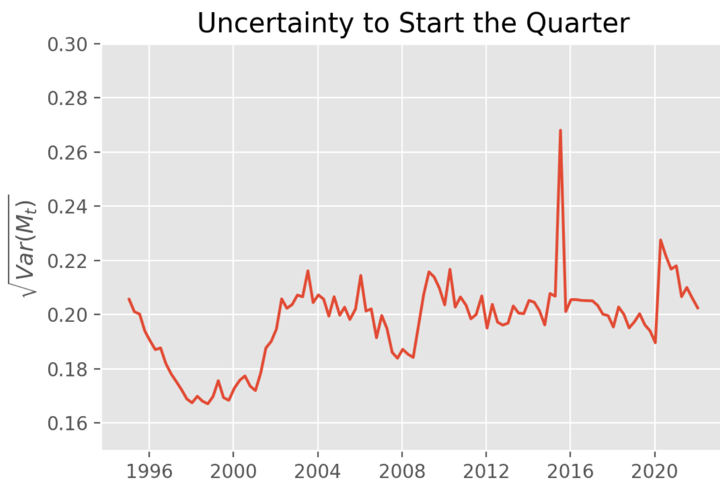

Figure 3. Level of economic uncertainty to start each quarter. The square root of $VAR(M_t)$ before any earnings are observed in a quarter is shown.

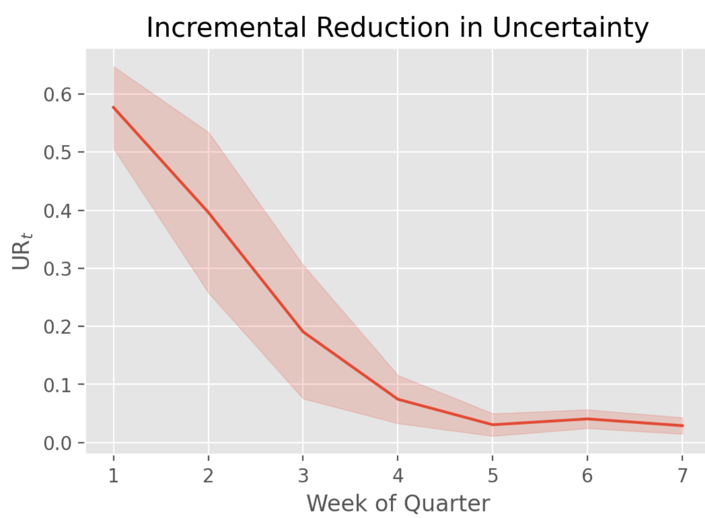

Figure 4. Reduction in economic uncertainty each week of a quarter. The reductions capture the systematic information provided by the actual earnings reports each week. The ribbon is plus/minus one standard deviation. A given week must have at least 200 reporting firms to be included.

Figure 5. Abnormal IVOL as a fraction of normal IVOL from the prior two quarters. Daily CAPM residual returns are squared and summed each week, then scaled by idiosyncratic volatility measured in the prior two quarters, and lastly, one is subtracted from this ratio to produce abnormal IVOL. Average AIVOL is reported for the Initial and Follow-On weeks of earnings seasons, respectively, as well as the remaining calendar weeks that are outside of those two periods. Stocks are separated into announcing firms (red) in a given week and non-announcing firms (blue). Panel A reports the means of weekly means of AIVOL. Panel B reports the median of weekly medians of AIVOL.

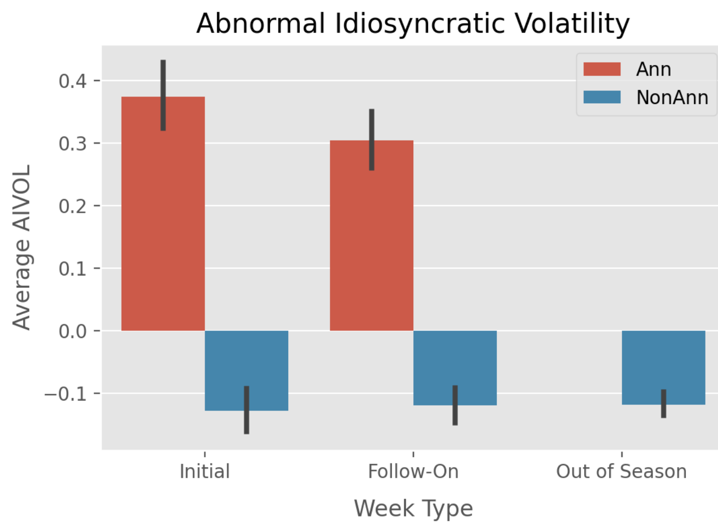

(a) Mean of calendar-time means.

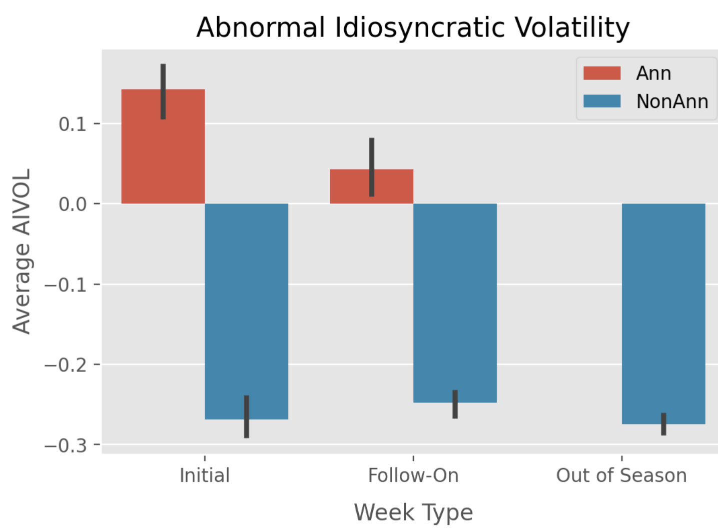

(b) Median of calendar-time medians.

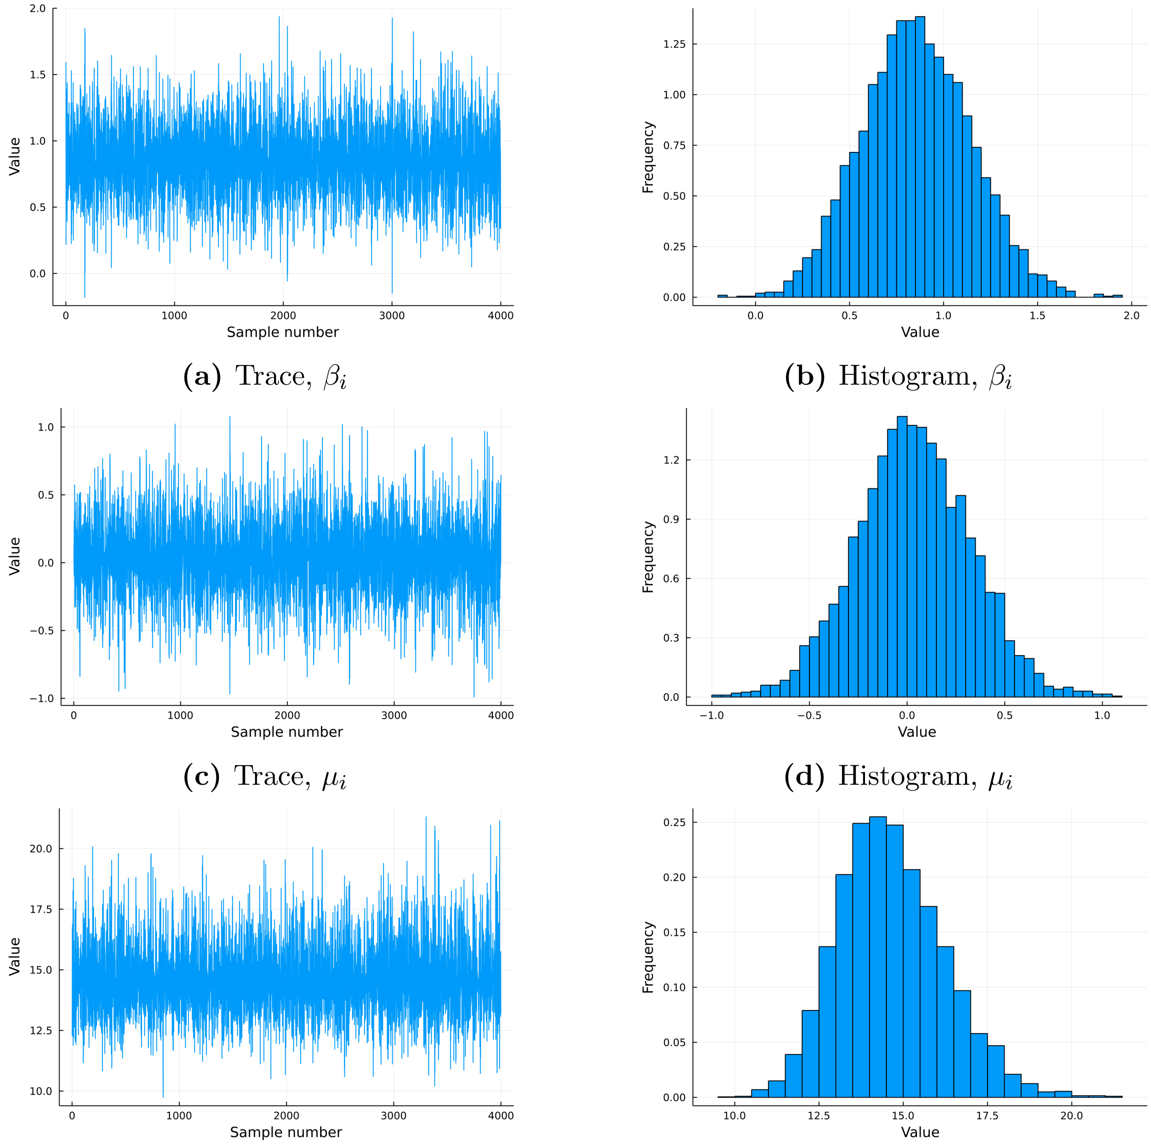

(a) Trace, β_i **(b)** Histogram, β_i

(c) Trace, μ_i **(d)** Histogram, μ_i

(e) Trace, γ_i **(f)** Histogram, γ_i

Figure 6. Trace and histogram plots for parameters estimated for IBM using data from 1984q1 to 1994q4.

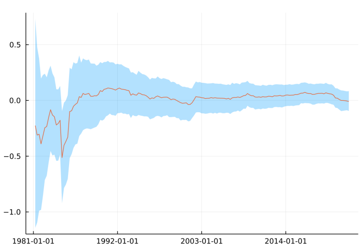

(a) Idiosyncratic mean μ

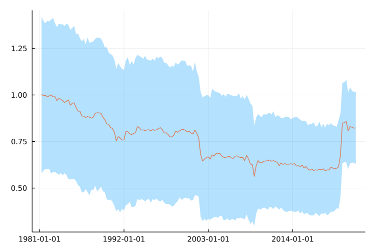

(b) Common factor loading β

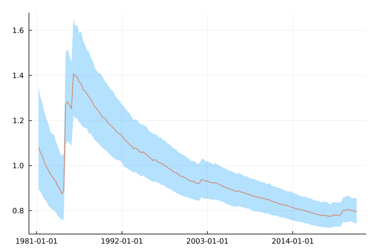

(c) Idiosyncratic variance γ

Figure 7. Median estimates of parameters (bounded by 10% and 90% quantiles) estimates of parameters for the Walt Disney Company.