PARAMETER UNCERTAINTY, CASHFLOW BETAS, AND EARNINGS ANNOUNCEMENT PREMIA

by
CAMERON S. PFIFFER

A DISSERTATION

Presented to the Department of Finance
and the Division of Graduate Studies of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctorate of Philosophy

September 2022

DISSERTATION APPROVAL PAGE

Student: Cameron S. Pfiffer

Title: Parameter Uncertainty, Cashflow Betas, and Earnings Announcement Premia

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Finance by:

Roberto Gutierrez	Chair
Robert Ready	Core Member
Youchang Wu	Core Member
Jeremy Piger	Institutional Representative

and

Krista Chronister	Vice Provost for Graduate Studies

Original approval signatures are on file with the University of Oregon Division of Graduate Studies.

Degree awarded September 2022.

DISSERTATION ABSTRACT

Cameron S. Pfiffer

Doctor of Philosophy

Department of Finance

September 2022

Title: Parameter Uncertainty, Cashflow Betas, and Earnings Announcement Premia

I apply Bayesian methods to estimate parameters describing the relationship between firm earnings and unobserved common earnings shocks. I estimate a firm's Bayesian cash-flow beta, which measures the comovement between firm earnings and a latent aggregate earnings factor, along with estimating the uncertainty about the firm's cash-flow beta. Firms with high parameter uncertainty have higher expected stock returns and lower stock price reactions to earnings, consistent with investors' rational learning in the presence of parameter uncertainty. A novel measure summarizes the capacity of a firm's earnings news to convey information about the macroeconomic state and reveals that earnings responses and announcement risk premia increase with a firm's informativeness. The most informative firms tend to announce earlier in earnings seasons.

ACKNOWLEDGEMENTS

I am extraordinarily grateful for the work of my advisor, Ro Gutierrez, who went to great lengths to help me achieve my personal and academic goals. Ro has been a superlative advisor with the marvelous ability to challenge my basic assumptions in an informative and open dialogue. I thank my committee members Robert Ready, Youchang Wu, and Jeremy Piger for their expertise, availability, kindness, and especially for their willingness to help me through uncommon circumstances. I thank Shoshana Vasserman for affording me a bright new world to explore when I needed it most, and for putting on a masterclass on intelligence, rigor, curiosity, and compassion.

I acknowledge the support of all the past and present faculty members of the University of Oregon's Department of Finance who have provided assistance to me during my time at the University. In alphabetical particular order: Ioannis Branikas, Maria Chaderina, John Chalmers, Diane Del Guercio, Brandon Julio, Stephen McKeon, Phil Romero, Albert Sheen, and Jay Wang. I also thank my past and current PhD student peers for their companionship and comments during my studies: Gretchen Gamrat, Arash Dayani, Donghyeok Jang, Wensong Zhong, Yuwen Yuan, Sina Davoudi, Wendi Wu, and Xuanyu Bai.

I am grateful for the assistance of the various communities I have been involved with. I acknowledge the support and kindness of the individuals who founded the Microstructure Exchange with me, and who treated me as an equal: Björn Hagströmer, Andriy Shkilko, Katya Malinova, and Andreas Park. I am grateful to Hong Ge and the Turing.jl team for inspiring me to study Bayesian methods, with particular thanks to David Widmann for his help in making me a better engineer.

I acknowledge the emotional, social, and academic support provided by the broader economics and statistics community. Specific and pointed suggestions were given by Brad Ross, David Childers. I thank various individuals treating me with respect and kindness: Toni Whited, Andrew Baker, Luke Stein, Tony Cookson, Ashvin Ghandi, Akhil Rao, Demitri Pananos, Daniel Bergstresser, Arpit Gupta, Benn Eifert, Anne Laski, and John Horton. I further thank uncountable individuals in the community for their small acts of kindness, inspiration, and humor.

I acknowledge the support and resources of Stanford University's computational cluster, Sherlock, Café Roma for cheap, delicious coffee and baked goods, and the piano located in the Erb Memorial Union for allowing me to practice. Lastly, I thank the members of the Oregon Association of Rowers for providing me an outlet and teaching me to row, a skill I now hold dear.

DEDICATION

I dedicate this dissertation to everyone who helped me in my educational journey. Tammy Schrader, who was so patient with me and my inability to complete any mathematics assignment. Maria Wickwire, who thought I could do whatever I put my mind to. My fifth grade teacher Patrick Webb, who saw that I need to live to thrive. My undergraduate teacher Curtis Bacon, who inspired me to continue to graduate school. I dedicate this dissertation to every educator or authority figure who believed in my capabilities even when I did not.

I dedicate this work to all the people who have made me a better person or supported my personal development. I thank my brother, Quinlan Pfiffer, who continues to show me how to live, how to make time for joy, and how to push yourself. I thank Patricia Levi for providing me with a vision of a prismatic future as we spin around the loop for a little while. I thank Elizabeth Koonce for her kindness, spirit, and bottomless support.

Above all, I dedicate this dissertation to my mother, Lynnette Pfiffer. She demonstrated an unparalleled certainty of character during challenging years. My mother gave my brother and I the gift of her personal values. Her only wish for us was that we be Renaissance men — a wish that inspired us to be skilled, artistic, and kind, while specifying no path or objective other than that we explore our curiosity with an open heart. We were left with the space to attend our inner and outer lives, to build what we could, and, above all, to enjoy the act of learning and doing. The generality, openness, and singularity of purpose afforded to us by the desire that we pursue a breadth of knowledge ensures that we have whatever we need to thrive. I am grateful for my mother inspiring me to find joy in infinite improvement.

So far, I think I've done a pretty good job as Renaissance man. How do you think I did, ma?

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1

INTRODUCTION

Firm earnings announcements are important events for market participants. Quarterly earnings news contains information about the firm profits, business activities, and challenges. For the econometrician, a cross-section of quarterly earnings announcements provide a rich data source of a firm's exposure to common macroeconomic factors. Earnings accrued during the same quarter that co-move between firms is indicative of shared exposure to the business cycle. Importantly, investors do not know with certainty whether there is a latent earnings component, what form it takes, or how firm earnings interact with the latent factor. Investors can only infer commonality in earnings growth through the data they can observe. Investors face a learning problem where they are attempting to understand firm and common factors simultaneously. The purpose of this paper is to quantify how much can be learned from earnings announcements once it is acknowledged that the underlying commonality in earnings may not be observed with certainty.

I propose a novel Bayesian model of learning through earnings growth and use my model to address several topics of interest to the financial economist. I derive a high-dimensional Bayesian statistical model of the economy where no parameter is known to an investor – the true state of the economy is unobservable, and can only be inferred through observed firm earnings growth. Investors in the real economy face a learning problem whereby they jointly estimate the parameters that govern the factor process and the parameters that govern firm-level earnings shocks. My modelling approach is unique and permits me to study the uncertainty of a firm's exposure to aggregate earnings growth, and grants me a wealth of information about the joint distribution of the economy, such as average firm earnings growth, firm systematic risk exposure, and aggregate earnings growth risk.

A primary output of my model is the cashflow beta, which captures the linear exposure of a firm's earnings growth to unobserved common earnings growth shocks. The canonical cashflow beta measures the growth in raw earnings as a function of aggregate earnings growth. The cashflow beta is distinct from the CAPM beta which measures the sensitivity of an asset's return to the market's return, in that the CAPM beta includes both a cashflow component and a discount-rate component. Estimation of the cashflow beta is a non-trivial exercise. Methods include vector autoregressions (Campbell and Vuolteenaho 2004), fitting ARIMA models on consumption growth (Da 2009), or estimating principal components (Ball, Sadka, and Sadka 2009). In each case, researchers typically proxy for aggregate earnings growth with one or more measures, or estimate aggregate earnings growth through some model. I estimate a cashflow beta that differs from previous estimates, in that it is a cashflow beta on an uncertain and unobserved latent factor. Directly estimating a cashflow beta partially frees my cashflow beta from the specification issues in Campbell and Vuolteenaho (2004), as highlighted by Chen and Zhao (2009).

Commonality in asset payoffs has a well-studied history and in some sense functions as the primary causal driver of asset prices. Asset prices vary with the exposure of the asset's payoffs to systematic risk. Assets that tend to provide low cashflows when aggregate cashflows are low should have higher expected returns, as cashflows arrive precisely when those cashflows are the least desirable. However, asset pricing research has shown that there are many facets of systematic risk. To capture the timing, size, and risk of cashflows and their arrival times, researchers have described an asset's systematic risk exposure in several forms. The original Sharpe-Lintner CAPM market beta, for example, has been decomposed into a discount rate beta and a cashflow beta (Campbell and Vuolteenaho 2004). Da (2009) further decomposes the cashflow beta into a covariance measure, describing the covariance between cashflows and contemporaneous aggregate consumption, and into a duration measure that describes the temporal relationship between current cashflows and future cashflow growth. The cashflow beta and its associated characteristics has generally been found to drive cross-sectional asset prices. Campbell and Vuolteenaho (2004), Da (2009), Ball, Sadka, and Tseng (2021), and Bansal, Dittmar, and Lundblad (2005) all find positive risk premia associated with the cashflow beta[1]. I do not find a positive risk price using my cashflow beta estimate. Perhaps quarterly firm earnings are insufficient for estimating relevant cashflow risk, or perhaps there are model limitations that need to be addressed. I discuss alternative modelling approaches to consider.

A novel contribution of my paper is examining the effects of uncertainty about the cashflow beta. There may be insufficient evidence at a point in time to conclude the precise point estimate of a cashflow beta, a phenomenon known as parameter uncertainty or estimation risk. Parameter uncertainty is suggested to play a role in the traditional CAPM-style market beta (Handa and Linn 1993; Kumar et al. 2008), but there does not currently exist a paper studying the role of the cashflow beta and its associated parameter uncertainty. The lack of existing literature on parameter uncertainty as it relates to the cashflow beta is surprising, considering that the literature on parameter uncertainty in the CAPM beta highlights significant effects in expected returns. Insofar as the CAPM market beta is considered to be composed of some measure of a cashflow beta (Campbell and Vuolteenaho 2004), it stands to reason that parameter uncertainty in the cashflow beta should influence expected returns just as it does in the CAPM beta. I find, surprisingly,

[1]Positive risk prices for the cashflow beta is not always the case. As shown in Ellahie (2021), there are inconsistent cashflow beta prices associated with different measures of aggregate earnings – some cashflow beta prices are positive, some are negative.

that the cashflow beta is not a priced risk in my sample of stock returns, but that there seems to be partial evidence that uncertainty about the cashflow beta is. I highlight several reasons why my unintuitive finding could arise, including model misspecification, lack of time variation in beta, or unidentifiability in the latent factor. I provide guidance for future analysis in how to better model a cashflow beta and its associated parameter uncertainty.

The cashflow beta also provides for a natural extension into research questions other than cross-sectional asset prices. If a firm's cashflows are in part determined by a shared but unobserved common shock (through the cashflow beta), then firm earnings news should function as a signal about that common shock. News from different firms should have information content that varies in quality – firms that have large betas and low idiosyncratic cashflow volatility are more precise signals about the common shock, while firms with low betas or high idiosyncratic volatility are noisy signals that reveal little of the aggregate earnings shock.

I provide an additional set of tests motivated by the informational content of firm cashflow signals. First, I study the unconditional earnings announcement premia accrued by firms with different characteristics. Numerous papers have highlighted the existince of elevated abnormal returns across earnings announcements without conditioning on earnings surprise. The literature highlights two general mechanisms for the earnings announcement premia. Evidence suggests that (a) risk increases during an earnings announcement, or (b) that risk is unchanged, but the market premium is detectable only during announcement weeks. In case (a), Savor and Wilson (2016) and Patton and Verardo (2012) highlight that announcing firms earn a premium over non-announcing firms because firm earnings are noisy signals about aggregate cashflow news. Savor and Wilson (2016) provide a theoretical model of this premium, their model does not permit any form of heterogeneity in the characteristics of announcing firms. For case (b), Chan and Marsh (2022) shows that the return-to-beta association only appears to hold during earnings announcement periods. In either case, the premium accrued seems to be a function of systematic risk exposure and asset informativeness. I find evidence in support of both explanations for earnings announcement premia, and further find evidence that heterogeneity in firm characteristics plays a significant role in the earnings announcement premia. Firms that are more informative accrue the majority of the earnings announcement premium, as do firms with larger cashflow betas. My findings suggest that a model of announcement premia that permits variation in firm characteristics would better inform financial economists about why risk appears to be elevated across earnings announcements, and to which firm characteristics that risk is attributable to.

Second, I explore the *response* to firm earnings news as a function of the cashflow beta and its uncertainty. Earnings news constitute an information revelation to the market. Investors can revise their beliefs about firm characteristics after observing an earnings report. News from firms with higher cashflow betas (holding idiosyncratic risk constant) implies that a larger proportion of the firm's news is systematic, and thus there is a lower informational content about the idiosyncratic elements of a firm. When much of the news of a firm has little to do with the firm itself, we should expect reduced earnings responses. Similarly, if parameter uncertainty is high, the ability of an investor to attribute observed earnings news to idiosyncratic or common shocks is diminished – the same unit of news observed by an uncertain firm and a certain firm is processed differently. Schmalz and Zhuk (2019) and Beyer and Smith (2021) both predict that higher cashflow beta uncertainty implies lower responses to unexpected earnings in normal economic conditions. Consistent with the literature's predictions, I find that earnings responses are lower for firms with high cashflow betas and with high parameter uncertainty.

Third, I produce a sufficient statistic that describes the signal-to-noise ratio of a firm's earnings news about common cashflow shocks. Informative firms (high signal-to-noise ratios) are commonly referred to as "bellwethers" (Bonsall, Bozanic, and Fischer 2013). News from bellwether firms conveys high-quality information about the broader industry or macroeconomy in which it operates, and commensurately draw more attention from analysts (Hameed et al. 2015). I derive a novel, parametric, continuous measure of firm bellwetherness, and use the measure to demonstrate that bellwether firms have different return behavior during earnings announcements, relative to non-bellwether firms. My statistic, the variance ratio, summarizes the posterior joint density of a firm's parameters and its ability to convey macroeconomic news. I show that my variance ratio is qualitatively distinct from an alternative and simpler measure of cashflow informativeness like R^2 of a regression of aggregate earnings on firm earnings, a statistic commonly employed in the bellwether literature (see Hameed et al. (2015) for one example). I show that the most informative firms, as measured by the variance ratio, have higher earnings responses, higher earnings announcement premia, and tend to announce earlier.

Fourth, I provide novel evidence on the realized ordering of firm earnings announcements. I show that, on average, firms announce earnings in a way that appears to be relatively optimal. Suppose a social planner were to assign earnings announcement order with the goal to reveal the most information about the aggregate earnings shock as early as possible. My model predicts which firms should generally tend to announce earlier under the social planner's problem, and I find evidence that more informative firms tend to announce earlier. I contribute to the literature on the effects of earnings announcement timing and socially optimal earnings disclosure, and highlight future areas of research.

The paper proceeds as follows. Section 2 outlines the model I use and derives measures for informativeness and ambiguity. Section 3 summarizes the existing literature and formulates predictions that can be made on each topic I consider. Section 4 produces the estimated quantities and details summary statistics about the MCMC estimation routine. Section 5 applies the estimated parameters from my model to each topic. Section 6 discusses the implications of my modelling choices and outlines future research areas. Section 7 details an alternative model specification. Section 8 concludes.

Chapter 2

MODEL

In this section, I propose a novel model of firm earnings growth that is simplistic enough to form a basic economic intuition, yet complex enough to require a joint estimation process. As is common with contemporary Bayesian models, I focus first on describing the laws of motion of the economy that dictate the likelihood function. I focus next on my choice of priors. The section concludes with a discussion of the advantages and disadvantages of my model.

I assume that the structure of firm earnings growth is linear in parameters unknown to the investor. I assume that earnings growth, denoted X_t, is partly determined by a firm's own loadings on a one-dimensional latent process M_t. Furthermore, I make no claims about the interpretation on M_t, other than that it is a driver of common sources of variation in earnings growth.

$$X_t = \mu + \beta M_t + \nu_t \tag{2.1}$$

$$M_t = \delta + \Phi M_{t-1} + \epsilon_t \tag{2.2}$$

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.

All the above parameters and latent states are unknown to a rational Bayesian investor. Denote the parameter vector of the system with $\theta(t) = (\mu, \beta, \sigma, \gamma, \phi, \delta, \epsilon_{1:t})$. I omit the idiosyncratic parameter ν_t.[1] The parameter vector has length $3N + 2 + (N + t)$.

[1] 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.

Note that the latent factor shocks ϵ_t are treated as parameters of the model. I 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.

Each quarter t, the investor observes the realized earnings growth of all firms announcing in that quarter. Conditional upon observing earnings growth, the investor updates their 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, I 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 me to make any statements about the empirical distribution of parameters in my model, and, more broadly, is a simplification that permits the use of many modern Bayesian inference methods.[2]

The first component of the posterior, $P(X_{1:t} \mid \theta)$, is referred to as the likelihood. My likelihood function is proportional to the Gaussian PDF centered at the $\mu + \beta M_t$ 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 $P(\theta)$, was chosen to predispose my estimates towards nonzero cashflow betas.[3]

- Firm means: $\mu \sim \mathcal{N}(\hat{\mu}, I)$. I set firm mean earnings growth to the average of that firm's earnings growth in first five years if the firm was in existence at that time, or to the average earnings growth across all firms between 1980-1985 if a firm enters after 1985 (this vector is denoted $\hat{\mu}$). The prior variance is 1, which assumes that about 10% of the 15.52% aggregate earnings growth variance in my sample is due to prior uncertainty in average earnings growth.

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

[3]Selecting priors for a high-dimensional model is nontrivial, and the computational burden of my model means that it is infeasible to conduct prior sensitivity analyses.

- Cashflow betas: $\beta_i \sim \mathcal{N}(1, 0.1I) \forall i$. I center the prior mean on β_i at one, with the belief that firms are *on average* positively exposed to aggregate earnings growth shocks. The prior variance of 0.01 reflects my relatively strong belief that the model should not permit too many firms to have a cashflow beta of zero.[4]

- Firm idiosyncratic variances γ_i^2 follow an Inverse Gamma prior, where each firm has a different distribution. To set my priors on firm idiosyncratic risk, I estimate a firm's earnings growth variance between 1980 and 1985 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 that did not exist between 1980 and 1985 begin the sample with a prior equal to the aggregate earnings growth variance. My choice of prior is equivalent to assuming that half of a firm's observed earnings growth variance is idiosyncratic.

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

- Latent variance: $\sigma^2 \sim \text{InverseGamma}(10, 143.83)$. I estimate the variance of all firm's earnings growth between 1980 and 1985, 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.

- Latent process persistence: $\phi \sim \mathcal{N}(0, 0.2) \in [-1, 1]$. I 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, Lewellen, and Warner 2006). I truncate my distribution between -1 and 1 to rule out AR(1) processes that are not covariance stationary.

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

2.1 Model discussion

Existing models of cashflow betas, such as Campbell and Vuolteenaho (2004), employ a vector autoregression (VAR) to model cashflow betas. Campbell and Vuolteenaho (2004) use firm responses to macroeconomic conditions measured by (1) excess market returns, (2) the yield spread between long- and short-term bonds, (3) the market's smoothed price-to-earnings ratio, and (4) the small-stock value spread. In the context of my model, this

[4] 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.

is roughly equivalent to replacing the latent shocks with data, rather than estimating the latent shocks directly. There are several distinguishing features of my model relative to Campbell and Vuolteenaho (2004) that I contribute to the study of asset pricing.

First, I demonstrate how to examine cashflow betas in the context of a Bayesian model when the latent factors are unknown. A firm's latent factor loading β_i is important to understand jointly with the latent factor, as we do not generally know the latent factor and instead use rough proxies (as in Campbell and Vuolteenaho (2004)). Unfortunately, estimating the latent factor jointly with the firm's loading on that factor can be challenging to do analytically – the posterior product distribution of $\beta_i M_t$ is challenging to derive analytically even with very simple priors[5]. My model has the advantage of permitting the statistician to understand a firm's loading on an uncertain factor in a way that is new to the literature. When the latent factor is estimated in a way that permits the econometrician to be uncertain in its exact form, latent factor uncertainty can appropriately propagate to estimates of a firm's beta.

Previous papers have modeled latent factors as principal components and used firm loadings on principal components (Ball, Sadka, and Sadka 2009). My method is distinct from principal component analysis (PCA) in that my posterior densities are not strictly linear in the data, as in PCA. The posterior densities in my model are jointly determined and thus capable of describing the higher moments of my parameters in an internally consistent way, i.e. the posteriors are simply functions of the likelihood and of the prior. Additionally, PCA does not generally permit a degree of uncertainty in either the structure of the principal components, nor in the weightings on the principal components. Thus, the simple form of principal components analysis[6] is not able to speak to parameter uncertainty in the same way as is my model.

Second, to my knowledge, I am the first to estimate a model that parameterically decomposes firm earnings into common and idiosyncratic components. Prior studies generally proxy for commonality in earnings using ratios of R^2 (Brown and Kimbrough 2011), whereas my paper provides an explicit parametric form of the economy. Additionally, my model permits an intuitive understanding of the *uncertainty* a statistician should have in that decomposition – if the posterior joint samples of a firm's common factor loading β_i are highly negatively correlated with a firm's idiosyncratic average earnings growth μ_i, then it is difficult for the statistician to infer whether a firm's observed performance comes from μ_i or from β_i. [7]

[5] For example, if M_t was strictly Gaussian (which it is not, due to the AR parameter ϕ), the resulting distribution requires a nontrivial evaluation of modified Bessel functions (Cui et al. 2016).

[6] I distinguish the "simple form" of PCA from more sophisticated methods like probabilistic PCA (Tipping and Bishop 1999).

[7] I find a negative posterior correlation between μ_i and β_i, particularly for large firms that are tightly linked to aggregate earnings shocks. The reason this occurs is that, when a firm's observed earnings closely mimic the aggregate earnings growth series, the model is not able to distinguish between earnings that are because of a firm's idiosyncratic average performance μ_i

Third, my model permits partial time-variation in cashflow parameters[8], in that I re-estimate the model each quarter using new data. The model permits the statistician to change their entire understanding of a firm's parameters based on all the data available at the time. I opt for "constant" parameters in the laws of motion of my economy, but the estimates of those parameters are continuously updated each quarter.

My model has a weakness that is important to consider when interpreting my results. In particular, firm idiosyncratic performance is homoscedastic, while the latent factors are heteroskedastic. Recall the laws of motion

$$X_t = \mu + \beta M_t + \nu_t$$
$$M_t = \delta + \Phi M_{t-1} + \epsilon_t$$

with $\nu_{i,t} \sim \mathcal{N}(0, \gamma_i^2)$ and $\epsilon_{i,t} \sim \mathcal{N}(0, \sigma^2)$. Firm earnings growth is thus a function of (a) their own characteristics (average growth μ_i and stochastic shock ϵ_t) and (b) some loading (denoted with β_i) on an autoregressive process M_t.

Astute readers will note that both M_t and X_t have constant prior variances throughout time. Latent factor shocks ϵ_t have variance σ^2 and firm earnings growth has variance γ_i^2. In frequentist approaches to econometrics, it is extremely common to handle heteroscedasticity through clustering or robust standard errors, and the absence of an obvious heteroskedasticity-robust estimator may strike the frequentist reader as flawed at best and incorrect at worst.

In practice, however, note that σ^2 is a *hierarchical* parameter, and that the actual shocks ϵ_t have their own posterior distribution that is estimated using σ^2 as a prior. The posterior distribution for two different shocks ϵ_r and ϵ_s need not have the same *posterior* variance, even if they share the same *prior* variance. Few of the latent factor shocks I estimate have identical posterior variance, which addresses any concerns about the heteroscedasticity problem in full for the latent process and in part for the firm processes. Using the latest model where all 165 quarters are available, 2008q2 has a posterior standard deviation of 0.24, nearly twice that of the 2019q2's posterior standard deviation of 0.13.

Unfortunately, the issue still remains that the model I estimate treats $\nu_{i,t}$ as a parameter with a constant variance – firms are not permitted to have higher residual uncertainty

and the firm's exposure to the aggregate earnings process β_i. The "exchangeability" of μ_i and β_i for these firms is essentially a lack of identifiability quantified through the posterior correlation in μ_i and β_i.

[8] It is not clear that time-variation in parameters is critical to my empirical tests in Section 5. Campbell and Vuolteenaho (2004) find (notably) that firm cashflow and discount rate betas have significantly different characteristics between the first and second half of their sample period. Fortunately, Campbell and Vuolteenaho (2004) also replicate their cross-sectional tests using a rolling estimator, and find that constant and time varying parameters both produce comparable prices for cashflow beta risk.

at different time periods, since γ_i^2 is constant. An outcome of my modelling choice is that firm idiosyncratic variances are biased upwards to rationalize observed earnings growth in cases where a firm's earnings growth history is discordant with the estimated latent factor.

While I do not use γ_i^2 directly in my tests, an upwards-biased idiosyncratic variance can have the effect of attenuating the expectation of β_i towards the prior mean of 1, since a high idiosyncratic variance rationalizes a large set of observed earnings growth. The model can then reduce the joint density of a parameter space by pushing β_i closer to 1. This is not a definite outcome in my estimation and depends on the relative strength of the priors on γ_i^2, β_i, and on the degree of observed commonality in earnings growth.

My model is a novel approach to statistical modelling in contemporary economics. I can reduce a complex economic system to a manageable level, while maintaining the uncertainty and complexity of the system. Further, I demonstrate how one might apply contemporary Bayesian methods to financial economics. The subsequent section details the estimation of my model.

Chapter 3

TOPICS

This section details several topics in financial economics that I can address. Section 3.1 details existing research on the cashflow beta, and provides several tests to determine whether the cashflow beta I estimate measures exposure to systematic risk. Section 3.2 highlights common market responses to a firm's earnings announcements and outlines tests using my model outputs. Section 3.3 describes the role of firm informativeness and details several hypotheses about observable characteristics for highly informative firms.

3.1 Cashflow beta

The primary motivation of this paper is the study of the asset pricing effects of a firm's exposure to systematic risk. Early works in asset pricing theorized that some measure of an asset's covariance (beta) with consumption should be priced (Rubinstein 1976; Lucas 1978; Breeden 1979). This style of model, commonly referred to as the conditional CAPM or CCAPM, theorizes that differences in asset prices should be driven by differences in the covariance between an asset's returns and aggregate consumption (the consumption beta). Early tests of the CCAPM did not find that the consumption beta explained variation in asset price (Campbell 1996; Cochrane 1996).

Tests of the CCAPM varied in their approaches to attempt to explain the failure of the unconditional CCAPM to explain variation in returns. One branch proposes that consumption is simply measured poorly and better proxies for consumption growth provide consumption betas that better explain prices (Savov 2011; Kroencke 2017). Others estimate a time-varying CCAPM (Lettau and Ludvigson 2001). Some tackle the failure of the CCAPM by applying different preferences (Campbell and Cochrane 1999).

The focus of this paper is about one such extension intended to address failures of the CCAPM, the cashflow beta. The cashflow beta differs from the consumption beta in that it describes the covariance of an asset's cashflow with consumption, rather than the covariance of asset returns and consumption. The seminal work of Bansal and Yaron (2004) ties asset prices to their cashflow beta. A later work Bansal, Dittmar, and Lundblad (2005) notes that "[Under Epstein-Zin preferences t]he cash flow betas and standard consumption betas differ and do not provide the same information as risk premia." The cashflow beta is a quantity with a solid theoretical grounding and some empirical suport.

A variety of empirical tests find strong support for a positive cashflow beta premium (Da 2009; Campbell and Vuolteenaho 2004; Botshekan, Kraeussl, and Lucas 2012; Ball, Sadka, and Sadka 2009). However, a large porportion of these papers were called into question by Chen and Zhao (2009), who note that the cashflow betas estimated by any paper using the return decomposition approach of Campbell and Shiller (1988) could be identifying severely missepcified cashflow betas. Chen and Zhao (2009) shows that the common approach of first estimating discount rate news and then treating the residual effect as cashflow news can produce severe misspecification errors. Discount rate news is difficult to estimate and using proxies to capture discount rate news effect can leave large and potentially biased measurement errors in cashflow news, which produces biased and nonsensical cashflow betas. Papers with questionable results due to their use of the Campbell and Shiller (1988) include Campbell and Vuolteenaho (2004), Botshekan, Kraeussl, and Lucas (2012) and Bansal, Dittmar, and Lundblad (2005). Papers untouched by the Chen and Zhao (2009) critique include Da (2009) and Ball, Sadka, and Sadka (2009), but it remains the case that much of the evidence for a cashflow beta has been called into question for econometric purposes.

My key contribution is an alternative approach to modelling commonality in earnings growth. Rather than select a particular proxy for earnings growth, I allow the data, guided by my model, to estimate a latent earnings growth factor directly. Estimating the latent factor by parametric fit rather than by an alternative method like PCA (discussed in Section 6.1) allows my model to better interpret past earnings shocks in addition to current ones. For example, earnings news today that is inconsistent with a non-persistent latent factor can revise the posterior distribution of the common shock that appeared years previous.

My cashflow beta is not dissimilar to the ones estimated by Da (2009) or Ellahie (2021), who project observed or forecasted earnings growth on aggregate earnings growth. The Bayesian model I employ has a less obvious intuitive interpretation than the linear projection in that the posterior distribution has a non-zero weighting towards the prior density of my joint parameters θ. I select priors in my model to tilt the estimation towards the belief that the single latent factor generally influences many firms by setting the variance of the prior density of β_i to 0.1. The tightness of the prior on β_i means that the model must be presented with large amounts of data to reduce the latent factor to meaninglessness.

The interpretation of the cashflow beta, however, is not complicated. The parameter β_i is the channel the model has available to it to determine a firm's exposure to a common shock that is shared by all firms in the sample. The cashflow beta should be thought of as an uncertain parameter governing the firm's loading on an uncertain first common component of aggregate earnings, though the common factor I find is not a maximal variance-reducing component, but an uncertain, evidence-weighted composite informed by firm earnings. Ball, Sadka, and Sadka (2009) shows that firm loadings on the principal components of aggregate earnings are priced.

To understand this intuition better, let us walk through a simplified version of the way the model estimates a cashflow beta. If all firms tend to have highly volatile and persistent earnings growth shocks that move together, the model will assign more weight to the posterior distribution where firms have non-zero betas. If there is evidence of common shocks that are differentially shared by firms, then the model should down weight the density of β_i for firms that are less exposed and up weight the density of β_i for firms that are more exposed. The posterior density of the cashflow beta $P(\beta_i \mid X_{1:t})$ thus measures a firm's exposure to common sources of earnings growth.

My second contribution is my estimation of parameter uncertainty about the cashflow beta. Cashflow risk is thought to have several characteristics, of which the cashflow beta is just one. The cashflow beta is a measure of the covariance of asset cashflows with aggregate cashflows. I suggest that parameter uncertainty should play a role in asset prices just as the true value of the cashflow beta.

Parameter uncertainty is defined as uncertainty an investor faces in the parameters that govern the data generating process for an asset's characteristics, such as cash flows or returns. The effect of parameter uncertainty (also called estimation risk in older works) on asset prices is fairly well studied, though the literature has focused on restricted models of Bayesian learning. A notable early entry in the study of parameter uncertainty is Klein and Bawa (1977), which finds that parameter uncertainty should generally produce mean-variance portfolios that are disjoint from the traditional full-information portfolio of the CAPM.

Clarkson, Guedes, and Thompson (1996) summarizes the literature by noting that "...in a CAPM framework, the literature has shown that differential parameter uncertainty generally has a systematic component and should be priced to some degree". Handa and Linn (1993) expands on this in more detail, and notes that parameter uncertainty about idiosyncratic components (such as a firm's alpha) tends to function *as if* the asset has an elevated correlation with the stochastic discount factor. Investors attribute more systematic risk to assets for which they have less information, which reduces the prices for low-information assets. Differentially informed assets (low- and high-information) can arise due to different asset age, structural breaks in parameters, or a spin-off that generates a new firm.

Investors in Handa and Linn's model attribute more systematic risk to low-information firms because their parameters are more uncertain. This (rational) attribution of risk results in lowered prices for low-information firms and higher prices for high-information firms. Notably, Handa and Linn's model does not require correlated priors – investors in their model know an asset's idiosyncratic risk, and the only source of deviation from a full-knowledge APT portfolio arises because of common estimation risk in the underlying factor betas.

The effects of parameter uncertainty on prices can be non-linear, and it does not appear to be the case that higher parameter uncertainty should unambiguously increase or decrease asset prices. More recent papers reinforce the findings that parameter uncertainty causes the naïve CAPM tangency portfolio be suboptimal (Kan and Zhou 2007). Cvitanić et al. (2006) and Xia (2001) suggest that hedging estimation risk can be a significant proportion of risk asset demand, and that investment horizons can play a large role in portfolio construction. Importantly, in these papers, the degree to which parameter uncertainty shifts asset demand is a function of asset cross-covariance and not simply absolute parameter uncertainty. My paper only tests the simpler predicted form of parameter uncertainty of Handa and Linn (1993), which predicts that higher uncertainty yields higher expected returns.

I contribute parameter uncertainty in the systematic proportion of cashflow risk to a growing body of work that study alternative aspects of cashflow risk. Da (2009) models the *duration* of asset cashflows, a measure of whether cashflows arrive in the near or far term. Botshekan, Kraeussl, and Lucas (2012) studies differential cashflow betas in good and bad times, and finds that the bad-times cashflow beta has a higher risk price than does the good-times cashflow beta. Li and Zhang (2017) proposes cashflow betas with short- and long-term exposure to aggregate consumption growth.

To test my predictions, I propose two forms of tests. First, I estimate Fama-MacBeth cross-sectional regressions, controlling for the Fama-French 3-factor betas. I include the cashflow beta and its uncertainty. If the cashflow beta or my measure of parameter uncertainty has a risk price associated with it, the cross-sectional regressions will show a significant coefficient in the cross-sectional regressions. Importantly, my proposed tests should control for the Fama-French 3-factor betas of Fama and French (1993) to ensure that my test is not simply a proxy for a redundant but known risk factor.

Second, a common test when there is a hypothesized risk factor is to construct portfolios sorted by some characteristic and evaluate whether the ex-post returns of the portfolios differ across characteristics. Typically, researchers calculate abnormal returns for these characteristic portfolios and test whether the abnormal returns are statistically different across the high and low portfolios[1]. I provide univariate sorts on the cashflow beta and uncertainty, as well as a bivariate sort on the cashflow and its beta to highlight any conditional pricing effects.

My cashflow beta β_i should have an associated risk price as long as exposure to aggre-

[1]An alternative approach when performing portfolio tests of risk factors is to explicitly test the monotonicity in portfolio returns, sorted by the quartile of the characteristic. Patton and Timmermann (2010) provides tests on the assumption that portfolio returns should be monotone in the sorting characteristic.

gate earnings growth is a systematic risk, which is generally agreed to be the case (Campbell and Vuolteenaho 2004). Section 5.1 does not find evidence that my cashflow beta is a priced factor, but that the parameter uncertainty is priced. I discuss my findings and their implications in Section 6.2 and provide an alternative model specification for the cashflow beta in Section 7.

3.2 Earnings news

The cashflow beta in my model permits an investor to infer the contemporaneous latent state, although imperfectly. A firm's observed earnings growth $X_{i,t}$ is a function of idiosyncratic factors like μ_i and γ_i^2, but also of the cashflow beta-adjusted latent factor $\beta_i M_t$. Several papers provide theoretical or empirical evidence that the degree of a firm's covariance with the latent factor can induce observable effects in returns during earnings announcements.

First, a vein of research focuses on the earnings announcement premium, which describes abnormal returns that accrue during firm earnings announcements. The earnings premium occurs without conditioning on earnings surprise. Two general strands of research hypothesize why the premium may accrue. The first strand of literature suggests that stocks become more risky during earnings announcements due to rational learning (Savor and Wilson 2016; Patton and Timmermann 2010). The second strand posits that risk does not increase during earning announcement, but that the risk premium associated with a firm's market beta is better estimated during earnings announcement dates (Chan and Marsh 2022). In either case, the premium exists and seems to be a robust empirical feature.

Savor and Wilson (2016) show that a portfolio long on earnings announcers and short non-announcers has positive alpha. Patton and Verardo (2012) produce a similar finding. Savor and Wilson motivate the risk premium they find through a simple model of cross-learning, where information about a common cashflow shock can only be conveyed imperfectly through earnings news. Firms who announce in batches with other firms have a higher covariance with the common factor because their earnings news when they announce represents a larger proportion of beliefs about the aggregate shock than during other times. Savor and Wilson's model assumes that all firms are identical in characteristics, so their model cannot make any implications about how the premium they highlight could accrue to firms with different characteristics. As providing such a model is beyond the scope of my paper, I attempt to provide empirical evidence that the Savor and Wilson premium may vary with firm characteristics.

The form of heterogeneity I focus on is the cashflow beta and its parameter uncertainty.

Firms with higher cashflow betas are *ceteris paribus* clearer signals of aggregate earnings shocks, and thus it could be the case that the partial announcement effect of Savor and Wilson is amplified for these firms. Parameter uncertainty has a less clear role in earnings announcement premia, as it can arguably increase or decrease earnings announcement premia. On one hand, parameter uncertainty could *increase* earnings announcement premia for the same reason that parameter uncertainty in Handa and Linn (1993) functions as systematic risk – when investors do not *know* how much of firm earnings is systematic, they act *as if* the firm is more systematic than it may truly be. On the other hand, ex-ante parameter uncertainty may reduce earnings announcement premia relative to other announcers by adjusting uncertain firm weights in the announcer portfolio's signal about common earnings. If two firms announce earnings in the same week, but one firm has a highly uncertain cashflow beta, a rational response by market investors is to apply more weight to the firm with a certain cashflow beta.

To test earnings announcement premia, I perform Fama-MacBeth cross-sectional regressions with a dummy variable for a firm's announcement week. The predictions in Savor and Wilson (2016) are specifically on the premium for a portfolio of announcing firms relative to non-announcing firms, so any test I conduct should attempt to capture the same premium. To capture any heterogeneity in premium, I interact the announcement week dummy with a measure I believe should capture some degree of differential exposure to risk premia, such as the cashflow beta or parameter uncertainty. Finding a positive loading on any of the interaction terms implies a risk premium that accrues differentially to firms with a given characteristic. There is reason to believe that the earnings announcement tests may better describe the risk premium associated with my cashflow beta than does the general cross-sectional tests described in Section 3.1, as there seems to be evidence that premia are better estimated (and potentially only relevant) during earnings announcements (Chan and Marsh 2022).

My second set of tests applies the cashflow beta and its parameter uncertainty to earnings responses. Firm earnings announcements contain an expected component that is typically priced ahead of the earnings release. Earnings announcements also contain a surprise component that can be used by investors to revise their beliefs about a firm's parameters. Prior works quantify the degree of price response using the earnings response coefficient, which is the amount of abnormal returns that occur during an earnings announcement per unit of earnings surprise (Hirshleifer, Lim, and Teoh 2009).

Schmalz and Zhuk (2019) study earnings response coefficients in a Bayesian model similar to my own, though their model has some known structural parameters. Their primary finding is that earnings responses are higher in bad times, since earnings news is an unambiguous signal about firm quality. Schmalz and Zhuk derive an expression for a firm's earnings response coefficient (ERC) as a function of the relative parameter uncertainty about a firm's cashflow beta. All else equal, their expression predicts that higher betas and higher parameter uncertainty should imply lower earnings responses.

Testing earnings responses is relatively straightforward. First, I calculate the cumulative abnormal returns (CAR) on the day of and the day following an earnings announcement. Second, I regress the CAR on a firm's surprise unexpected earnings, or SUE, and a parameter I expect to influence the cumulative abnormal returns interacted with SUE. The construction of SUE is described in Section 5.2. If the interaction between the parameter of interest (the cashflow beta or parameter uncertainty) and a firm's surprise unexpected earnings is significant, it implies that market responses to the same unit of news differs for firms with different characteristics.

In Section 5.2, I provide tests of earnings response coefficients that vary with parameter uncertainty and with the cashflow beta, finding strong evidence in support of predictions drawn from Schmalz and Zhuk's model. I find that earnings responses are decreasing in (a) the cashflow beta and (b) parameter uncertainty. Parameter uncertainty appears to play a larger role in earnings responses than does the level of the cashflow beta.

Schmalz and Zhuk (2019) test earnings responses in good and bad economic periods, since their claim is largely that parameter uncertainty declines in bad times and thus earnings responses increase. I provide similar tests to see if decreasing parameter uncertainty corresponds with higher earnings responses segmented by good and bad periods. I do not find statistically significant evidence that earnings responses vary between good and bad times.

Overall, I note that the literature generally predicts higher earnings response coefficients for firms with lower uncertainty and lower cashflow betas, but that the role of the cashflow beta is unclear as it relates to earnings announcement premia. Sections 5.2 provides empirical tests of my predictions. I find evidence that firms with higher parameter uncertainty accrue larger earnings announcement premia, but have lower earnings response coefficients. I find no premium associated with the cashflow beta, but I do find that higher cashflow betas reduce earnings response coefficients.

3.3 Firm informativeness

My model permits me to define the signal-to-noise ratio of a firm about common earnings shocks. Firms with precise signals are often referred to as "bellwethers". A bellwether is typically defined in the financial economics literature as a firm that acts as a highly informative signal about other firms. Hameed et al. (2015) defines a bellwether as a firm whose fundamentals best predict the fundamentals of peer firms. They find that revisions in analyst forecasts about bellwether firms causes large price revisions in peer firms. Bonsall, Bozanic, and Fischer (2013) shows that bellwethers provide timely information about industry and market shocks. Hann, Kim, and Zheng (2019) shows that earnings announcements from bellwethers reduce option implied volatility in peer firms, suggesting that bellwether firms reduce uncertainty in non-bellwether firms, as measured by implied volatility.

In the spirit of these papers, I explicitly derive a new statistical measure of bellwetherness that succinctly communicates the ability of a firm's earnings to reveal aggregate earnings news. I make two broad predictions about bellwether firms in my sample.

First, bellwether firms should generally have higher earnings response coefficients. Bellwether firms have high systematic cashflow exposure (high betas) and low idiosyncratic volatility by their very definition. Further, bellwether firms draw more analyst attention (Hameed et al. 2015) which should correspond to elevated earnings response coefficients (Hirshleifer, Lim, and Teoh 2009). I expect that earnings responses for highly informative firms should be larger, because each unit of news they produce is both more informative and more economically meaningful.

Second, highly informative firms may accrue higher risk premia due to elevated market risk from a learning channel. Informative firms are those with (a) high betas, (b) low parameter uncertainty, and (c) low idiosyncratic risk. In the context of Savor and Wilson (2016), bellwether firms are those that are the clearest signals about aggregate earnings news. Announcing bellwether firms thus may have a larger earnings premium than other firms because they have high systematic exposure. Alternatively, these firms may have lower premia because they are very precise signals. As a counterfactual example, Savor and Wilson (2016) note that no risk premium would be earned if firm earnings fully revealed the common earnings news. Savor and Wilson (2016) is silent on modeling cross-sectional variation in the premium, so again I provide empirical results to motivate future analysis.

I derive a measure of firm informativeness to test my two predictions. Recall that my model focuses on firms' exposure to an uncertain latent factor. The latent factor is never known or revealed to investors, so they face a learning problem where observed firm earnings growth is used to infer the value of the parameter governing the latent factor M_t, its parameters, and the history of the shocks ϵ. The information value of any firm's earnings growth to the latent shock varies with firm characteristics. For example, if a firm has a large idiosyncratic earnings growth variance γ_i^2 relative to the volatility of aggregate earnings, that firm's earnings are high-noise and cannot be used to precisely determine the value of the current latent factor shock.

I define the variance ratio $\mathrm{VR}_{i,t}$, which summarizes a firm's ability to precisely and reliably inform investors about aggregate growth shocks. $\mathrm{VR}_{i,t}$ can be thought of as the degree of “bellwetherness”, that is, how much observing the firm’s actual earnings can reduce residual uncertainty about aggregate earnings growth shocks. 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 is $\text{Var}[M_t \mid X_{i,t}]$, with forecast variance $\text{Var}[M_t \mid X_{i,t}]$. 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[2]

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

where

$$\text{Var}[M_t \mid X_{i,t}, \hat{\theta}] = \frac{\sigma^2}{1-\phi^2}\left(\frac{\beta_i^2\sigma^2}{\gamma_i^2(1-\phi^2)} + 1\right)^{-1} \tag{3.1}$$

$$E[M_t \mid X_{i,t}, \hat{\theta}] = \delta + \sum_{s=1}^{t-1}\phi^s(\delta + \epsilon_{t-s}) + \left(1 + \frac{\gamma_i^2(1-\phi^2)}{\beta_i\sigma^2}\right)^{-1}\left(\beta_i\sum_{s=1}^{t-1}\phi^s\epsilon_{t-s} + \nu_{j,t}\right) \tag{3.2}$$

□

Lemma 1 is straightforward to compute numerically, once equipped with posterior samples of $\hat{\theta}^k$. Simply computing Equations 3.1 and 3.2 for all posterior samples $\hat{\theta}^k$ and calculating the sum of the sample average (of Equation 3.1) and variance (of Equation 3.2) yields the appropriate posterior $\text{Var}[M_t \mid X_{j,t}]$.

$\text{Var}[M_t \mid X_{j,t}]$ is not a scale-free measure, since the model has limited data available towards the beginning of the sample and variances can be elevated. I construct a scale-free measure that succinctly describes the information content a firm provides about the latent factor. I denote the amount of variance about the latent process M_t reduced after observing firm i’s earnings growth ($\text{VR}_{i,t}$) as

$$\text{VR}_{i,t} = \frac{\text{Var}_t[M_t \mid X_{i,t}]}{\text{Var}_t[M_t]} \tag{3.3}$$

The ratio $\text{VR}_{i,t}$ ratio is bounded between 0 and 1. The intuition of Equation 3.3 is that a

[2]This follows from the law of iterated expectations:

$$\text{Var}[X] = E\big[\text{Var}[X \mid Y]\big] + \text{Var}\big[E[X \mid Y]\big]$$

low ratio indicates that, on average, firm i contains precise information about the current value of M_t. The variance ratio is a sufficient statistic for the nowcasting capabilities of a firm's earnings, since (a) it is always possible to construct an unbiased estimate of M_t (in Equation 3.2) and thus signals can be discriminated only by their precision $\mathrm{Var}[M_t \mid X_{j,t}]$. Values of $\mathrm{VR}_{i,t}$ close to 1 indicate that the firm's earnings announcement provides little new information about the common factor. Note that the variance ratio $\mathrm{VR}_{i,t}$ summarizes both average variance and the conditional expectation, which addresses both the accuracy and location of an investor's nowcast.

I highlight the variance ratio because it is a reasonably effective way of reducing a high-dimensional parameter space into a statistic that describes a firm's informativeness. The variance ratio includes parameter uncertainty in addition to cross-learning from firms, making it a uniquely informative statistic about the power of a firm's realized earnings to forecast aggregate earnings.

The estimate $\mathrm{VR}_{i,t}$ has a shortcoming, in that it is a counterfactual evaluation of a world where firm i is the first firm to announce in a quarter. The firm with the lowest variance ratio in a quarter is the firm that an investor would choose to announce first, since the announcement carries the most information. A more rich measure of the informational content of a firm's earnings content would be to integrate across all possible announcement orders, though this is computationally infeasible as it would require the computation of 1535! variance ratios. Section 3.3.1 details what my model can and cannot say about optimal ordering of firms.

To summarize, I predict that bellwether firms, firms with low variance ratios $\mathrm{VR}_{i,t}$, should have higher earnings response coefficients. I am unclear on the role of bellwetherness and earnings announcement premia. In Section 5.3 I find evidence for elevated earnings responses but lower earnings announcement premia.

3.3.1 Earnings announcement order

The previous section provides a sufficient statistic for the amount of information a firm reveals if it were to go first in each earnings season. One question is whether *observed* earnings season orderings are consistent with an informationally optimal ordering. My model is capable of providing some information about the general earnings order observed and whether my model would predict a similar ordering.

Suppose a social planner were selecting the order in which firms should announce their earnings to the market. Assume that the social planner's objective is to maximize the amount of information transmitted as early as is possible. The first firm the social planner would select is the firm with the lowest variance ratio $\mathrm{VR}_{i,t}$, since this is the firm that most reduces uncertainty about the latent factor.

For the second and higher firms, the order is not simple to calculate numerically since it requires conditioning on the set of all firms who have already announced. A brute-force algorithm to compute optimal ordering would be to calculate a sequence of variance ratios:

$$\begin{aligned}
\mathrm{VR}^1_{i,t} &= \frac{\mathrm{Var}_t[M_t \mid X_{i,t}]}{\mathrm{Var}_t[M_t]} \\
\mathrm{VR}^2_{i,t} &= \frac{\mathrm{Var}_t[M_t \mid X_{i,t}, X_{\iota(1),t}]}{\mathrm{Var}_t[M_t \mid X_{\iota(1),t}]} \\
&\dots \\
\mathrm{VR}^N_{i,t} &= \frac{\mathrm{Var}_t[M_t \mid X_{i,t}, X_{\iota(1),\iota(2),\dots,\iota(N-1),t}]}{\mathrm{Var}_t[M_t \mid X_{\iota(1),\iota(2),\dots,\iota(N-1),t}]}
\end{aligned}$$

where $\iota(j) = \arg\min_i \mathrm{VR}^j_{i,t}$ indicates the firm with the minimum variance ratio in position j. The firm order vector $\iota = [\iota(1), \iota(2), \dots, \iota(N)]$ is then informationally optimal (by construction) in that each firm's earning news provides the most information to investors.

Unfortunately, calculating the optimal ordering is computationally infeasible as it would require the calculating a total of 1535! variance ratios. A more reasonable first-pass at optimal ordering would be to organize firms by their variance ratios, assuming they were all the first to announce. Releasing the earnings news of these firms first should generally correspond with large information events happening towards the beginning of the earnings season.

Few of the above papers make theoretical predictions about exactly where firms should land in the earnings announcement order. The only paper I am aware of on the topic is Guttman, Kremer, and Skrzypacz (2014), which models a manager's decision of when to release information early or late. They show that the manager should release information later rather than later, unless the manager is myopic and cares about early-period prices. Guttman, Kremer, and Skrzypacz (2014) points to the basics of disclosure theory and suggests that managers in a non-competitive environment would prefer to release later. However, Guttman, Kremer, and Skrzypacz (2014) does not address the effects of information competition in earnings disclosure. If there were two managers in the model, would the late-announcement preference still hold? This isn't clear, and so it remains an open question as to what order we should expect to observe in competitive markets.

It is not clear what we should expect to see by comparing this pseudo-optimal ordering to observed earnings announcement orderings, since there are many factors that have little to do with a hypothetical social planner's decision. Further, most existing papers on earnings announcement order tend to focus on marginal deviations from scheduled earnings. Kim, Pierce, and Yeung (2021) finds that firms announcing good news tend to delay their earnings news, and that the delay is unrelated to earnings management. In a similar vein, Johnson and So (2018) shows that firms who *schedule* earlier tend to announce worse news. Noh, So, and Verdi (2021) shows that firms who announce earlier (through a quasi-exogeneous shock to ordering) have higher earnings announcement premia and greater media attention. Savor and Wilson (2016) shows that early announcers have higher earnings announcement premia. Frederickson and Zolotoy (2015) finds that investors choosing between the earnings news of firms who announce on the same day tend to focus on more "visible" firms, defined in their study as firms with higher media attention, advertising spending, analyst coverage, and size.

Section 5.4 provides basic statistical tests showing that highly informative firms tend to announce first, even when controlling for size and book-to-market. My finding provides some validation to the model in that there is no explicit requirement in my model that early announcing firms also be informative, but observed non-model earnings announcement orders seem to coincide with a high-information equilibrium. The next section provides statistical evidence for the topics described above.

Chapter 4

ESTIMATION

4.1 Model estimation

My sample consists of quarterly earnings releases for 1,535 firms between January 1984 and December 2021. Firms are selected from the CRSP/Compustat dataset. I draw all domestic firms with a stock ownership code of zero (STKO) to exclude ADRs. I use only firms with fiscal quarters ending in March, June, September, and December to simplify my modelling 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.

I calculate seasonally-adjusted earnings growth $X_{i,t}$ using the method in Savor and Wilson (2016) and Kothari, Lewellen, and Warner (2006). Firm earnings growth is calculated as the difference between a firm's earnings (IBQ) and the firm's earnings four quarter prior, rescaled 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 my sample. I 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%).

My model is written in the probabilistic programming language Turing.jl. I 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 I have presented. As NUTS (and Markov chain Monte Carlo in general) may not be familiar to traditional financial economists, I include a brief description of the functionality of MCMC.

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 my 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 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 I 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. I refer the reader to Hoffman and Gelman (2014) for detailed information about NUTS, as understanding how it functions is both outside the scope of this paper and noncritical to understanding my results.

For each quarter, I generate a vector of earnings growth for firms announcing in the quarter. I 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.

My MCMC kernel, NUTS, is designed to have few degrees of freedom for the researcher. The two sampling parameters I 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. I set my 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 evaluated 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 I select a slightly lower tree depth of 5 than the default of 7. Lowering the maximum tree depth does not make my samples less accurate or incorrect asymptotically, but it does mean that my sampler will generally converge slower on average.

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

[1] 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.

ples 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[2]. A second diagnostic is to conduct a visual inspection of trace plots for samples to verify that the samples appear stationary. I provide diagnostics in the next section.

4.2 Estimation results

I estimate a total of 165 probabilistic models. Each of the 165 models incorporates a new quarter's earnings growth for each firm that announced earnings. The 165th model thus incorporates all earnings growth disclosures from 1984Q1-2021Q4 (165 total quarters).

To verify that each of the 165 models are likely to be valid MCMC samples (in the sense that their sampling distribution approximate the true posterior density), I 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 4.1 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. After reviewing 100 trace plots from different chains, I can proceed with the analysis under the assumption that my 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 a chain has converged. I calculate that the maximum $\hat{R}$ across all parameters and models is 1.042, indicating that all parameters are drawn from a stationary distribution[3]. The result of these tests allow me to proceed with my analysis, comfortable that sampling error due to chains that have not converged is unlikely to be driving my results.

[2] 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. My 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).

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

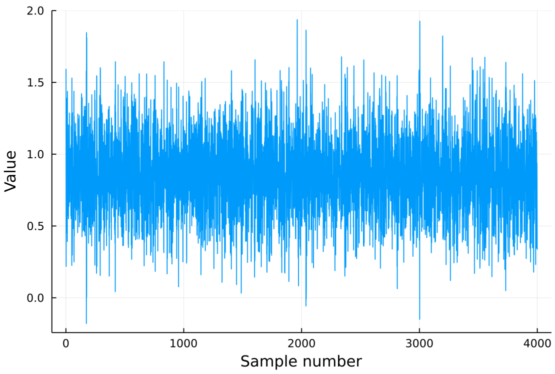

(a) Trace, β_i

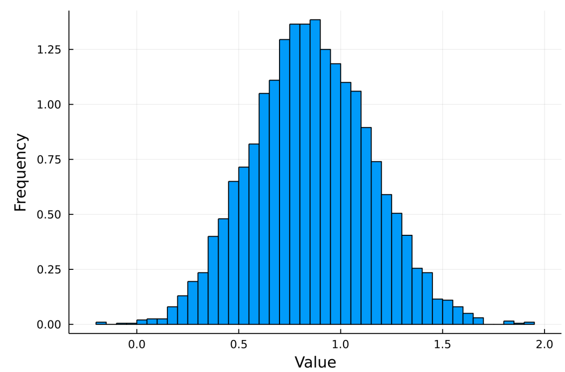

(b) Histogram, β_i

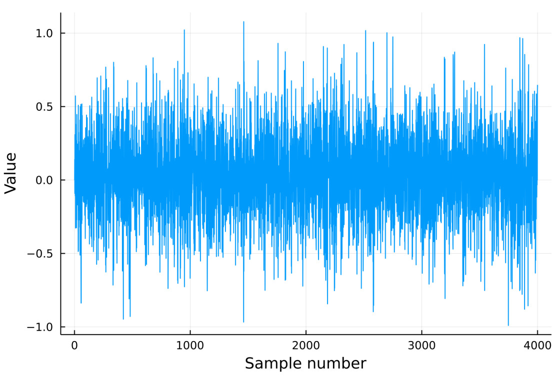

(c) Trace, μ_i

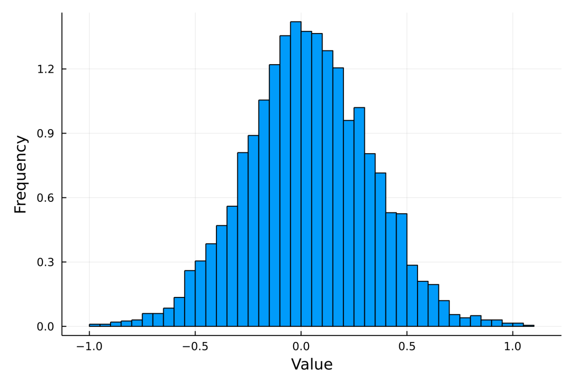

(d) Histogram, μ_i

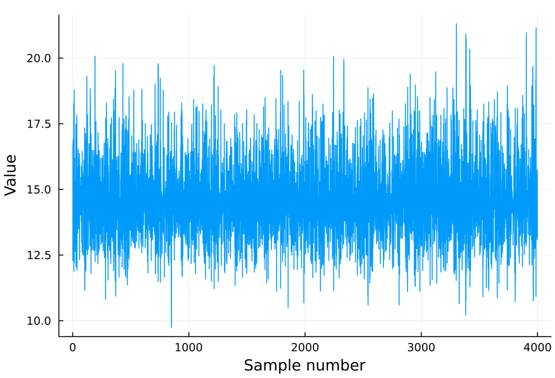

(e) Trace, γ_i

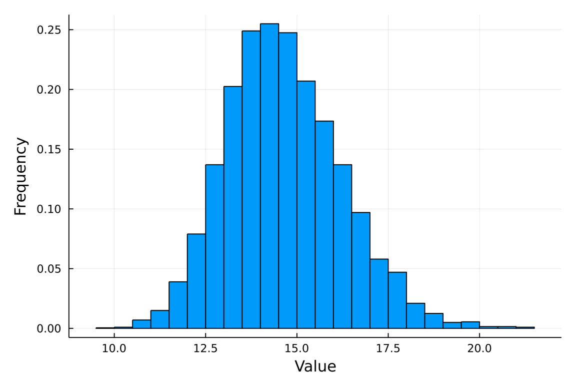

(f) Histogram, γ_i

Figure 4.1: Trace and histogram plots for parameters estimated for IBM using data from 1984q1 to 1994q4.

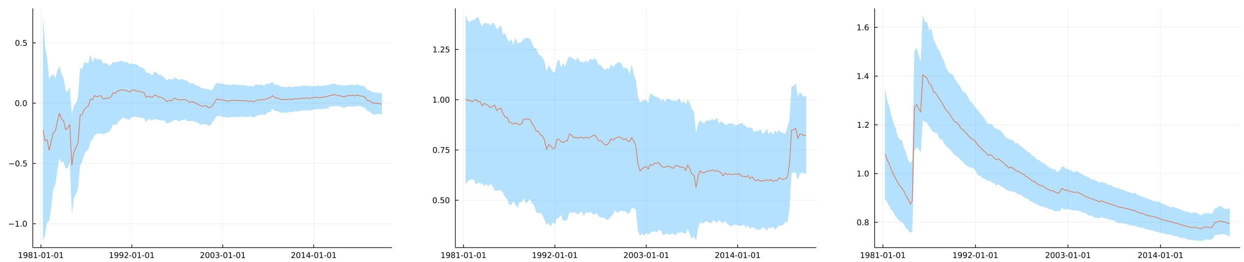

(a) Idiosyncratic mean μ (b) Common factor loading β (c) Idiosyncratic variance γ

Figure 4.2: Median estimates of parameters (bounded by 10% and 90% quantiles) estimates of parameters for the Walt Disney Company.

All variables in my model have distributional (at a single point in time) and time-series interpretation (distributions across time). Figure 4.4 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 4.2 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.

Next, I explore the determinants of the latent factor, and the cross-sectional distribution of firm parameters.

4.2.1 Latent factor results

Since the latent factor plays a key role in my analysis, I extract all posterior estimates that govern the latent factor to reconstruct it. Figure 4.3 plots the latent factor using the latest model, estimated with data from 1984-2021. I note that the latent process closely tracks recessions, particularly the 2001, 2008, and 2020 recessions. The precision of the latent factor is quite high, in that the posterior standard deviation is small. My model is relatively confident in the shape and timing of the latent factor.

Figure 4.2 displays the time series posteriors for the parameters that govern the latent factor M_t. Each point in time is the filtered estimate of the parameter in each panel, using data up to and including the quarter on the x-axis. Aggregate earnings growth shocks have become more persistent since roughly 2008. Note that the standard deviation of latent factor shocks σ appears to have declined from 3.7 on average to 1.25. While this may appear alarming to some readers, I note that this is more a function of the posterior variance of σ contracting as more information is revealed.

The unconditional mean aggregate earnings growth δ begins the sample in 1985 with a

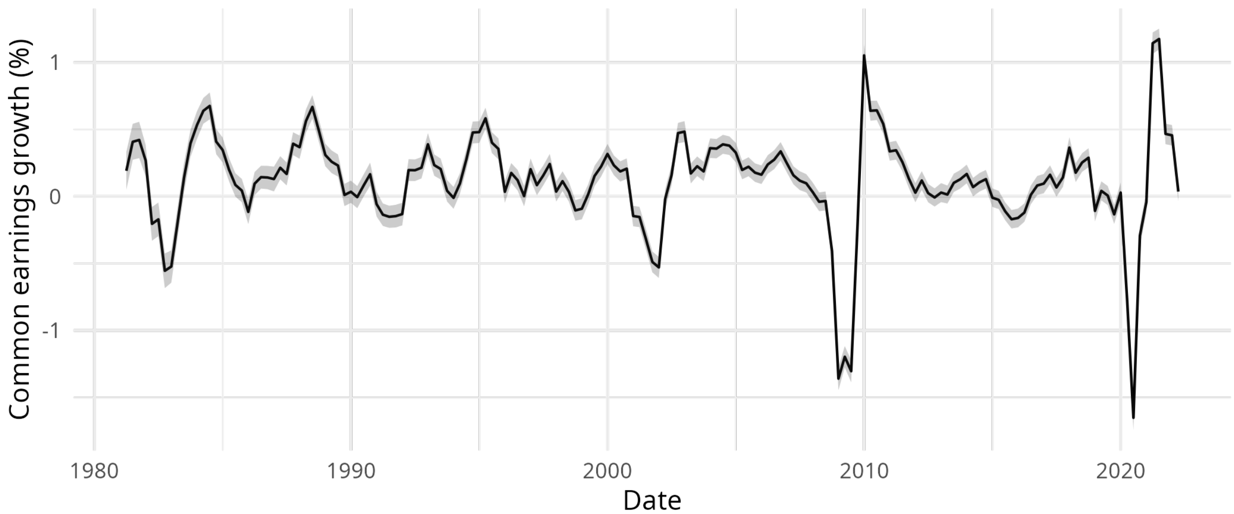

Figure 4.3: The estimated series of the latent factor M_t, using all 165 quarters of data. The thick blue line is the mean value $E[M_t \mid X_{1:165}]$, while the shaded ribbon represents the mean ±1 standard deviation $\sqrt{\text{Var}[M_t \mid X_{1:165}]}$.

median posterior belief of -0.5bps and increases to 10bps at the end of the sample. The 5% highest posterior density[4] for δ as of December 31st, 2021 is $[-0.125, 0.282]$. A frequentist interpretation is that δ is not significantly different from zero at the end of the sample using a 5% significance level, though I would caution against thinking this way – there is a large mass of δ that is positive (84% of the posterior samples are greater than zero), suggesting that δ is highly likely to be small but non-zero. My findings are qualitatively and quantitatively similar to Kothari, Lewellen, and Warner (2006), who find an unconditional mean of 0.14.

Beliefs about the latent factor persistence parameter ϕ begins near a median of 0, but asymptotes towards a median value of 0.225 with a 5% highest posterior density of $[-0.088, 0.555]$ as of the end of 2021. As of 2021, the posterior probability that $\phi > 0$ is 91%. Aggregate earnings growth persistence suggest that shocks are positively autocorrelated to some extent. A non-zero value of ϕ has the benefit of an explicit time-series learning component in the economy, since knowing the current quarter's shock has some predictive power over the next period's shock.

My posterior estimates of systematic variance (σ) declined steadily since 1985, starting near a median of 14% and declining towards a posterior median of 1.66%. The 5% highest posterior density at the end of the sample is relatively tight around $[1.363, 2.063]$. A firm with $\beta_i = 1$ would thus expect to have median earnings growth with a variance of 1.66% per quarter, assuming an idiosyncratic shock variance of zero.

To put the systematic shocks σ into context, note that the median idiosyncratic earn-

[4]The highest posterior density is similar to a standard confidence interval for unimodal posteriors.

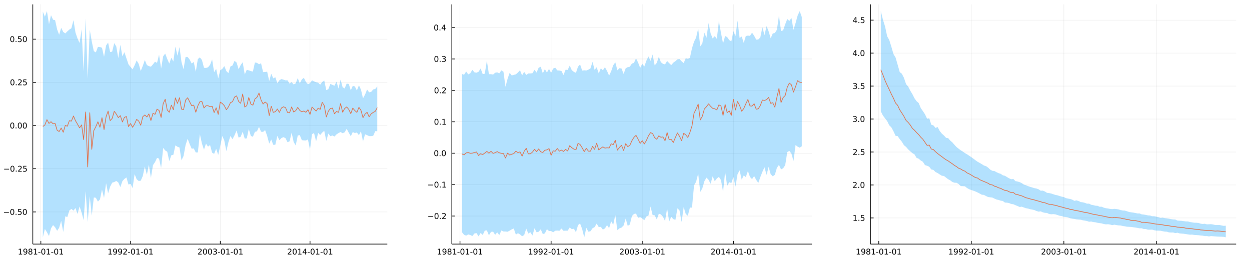

(a) Latent factor mean, δ (b) Latent factor persistence ϕ (c) Systematic variance σ

Figure 4.4: Median estimates of parameters (bounded by 10% and 90% quantiles) governing the latent process, given by the identity $M_t = \delta + \phi M_{t-1} + \epsilon_t$ where $\epsilon_t \sim N(0, \sigma)$.

ings growth γ_i is 5.23% as of 2021. I estimate that the median firm with $\beta_i = 1$ derives 24% of their earnings growth innovation from systematic exposure. This is roughly consistent with the explanatory power of the first principal component in Ball, Sadka, and Sadka (2009). It is half the average proportion of macroeconomic information in Bonsall, Bozanic, and Fischer (2013), though their use of a semipartial R^2 makes it difficult to compare to my measure and to the measure in Ball, Sadka, and Sadka (2009).

4.2.2 Firm results

Summarizing the estimated firm posteriors for μ_i, γ_i, and β_i is difficult due to the large number of these parameters, so I have opted to calculate the expected posterior values for all firms in each quarter, and display their cross-sectional dispersion by decade. Figure 4.5 displays the results.

Expected growth rates μ_i have declined from about 20bps per year in 1980 to about 10bps per year as of 2020. Expected cashflow betas β_i have dispersed substantially by the 2010s and 2020s, and firms are much more spread out around $E[\beta_i] = 1$. Lastly, firm idiosyncratic earnings variance γ_i has increased from the earlier decades. The distributions are also not particularly smooth, in that there are several "modes". In the 2010s, for example, idiosyncratic volatility has a modal point near 2% standard deviation in quarterly earnings growth, and a much smaller mode at about 0.5%. In all cases, it is clear that the cross-sectional distributions have moved away from the prior-dominated density of the 1980s.

The variance ratio $\text{VR}_{i,t}$ is calculable from the posterior draws. Figure 4.6 portrays the results for each firm overlaid on top of each other, to convey a sense of scale and location. I standardize the variance ratio in each quarter by subtracting the mean and dividing the difference by the standard deviation. Firms below zero in this panel are firms that are more informative on average about aggregate earnings growth. The distribution of vari-

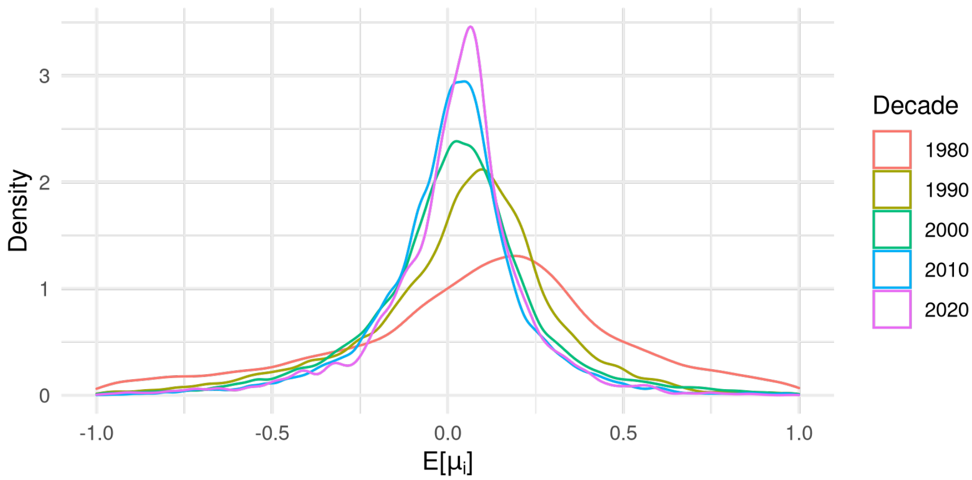

(a) μ_i

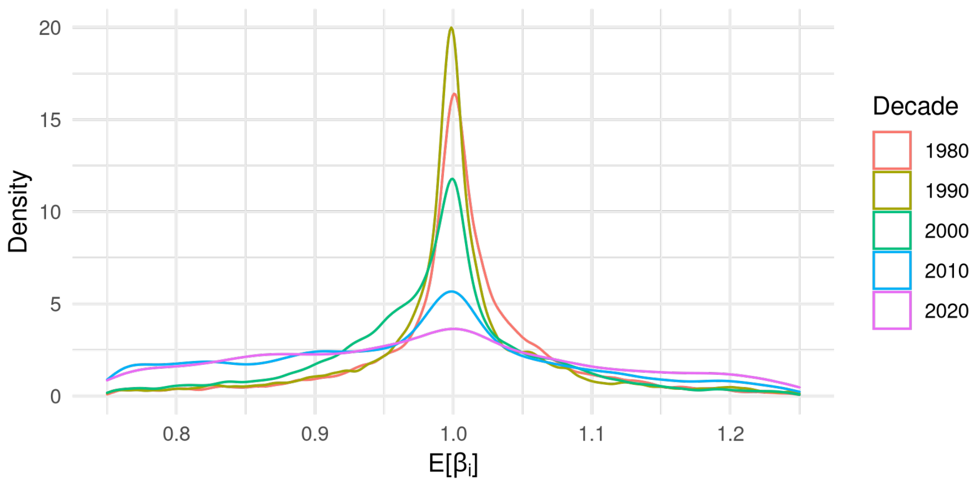

(b) β_i

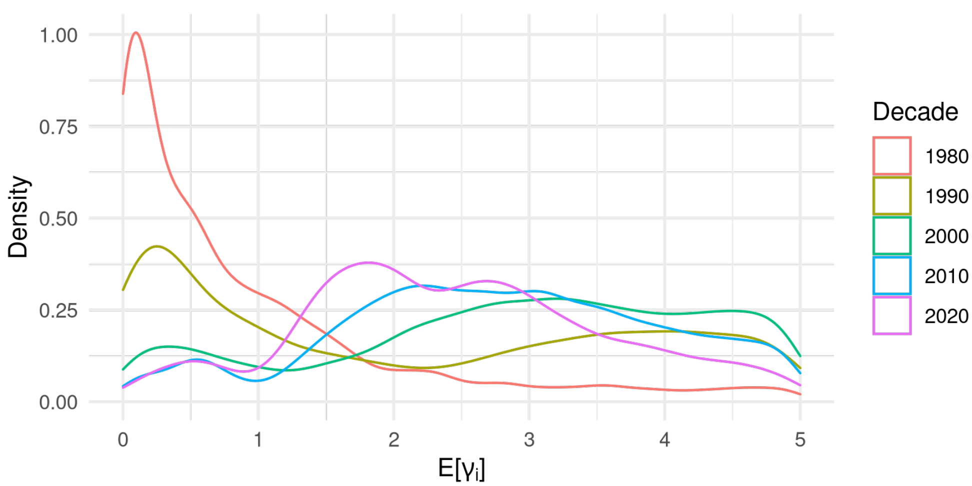

(c) γ_i

Figure 4.5: Cross-sectional dispersion in firm parameter estimates, grouped by decade.

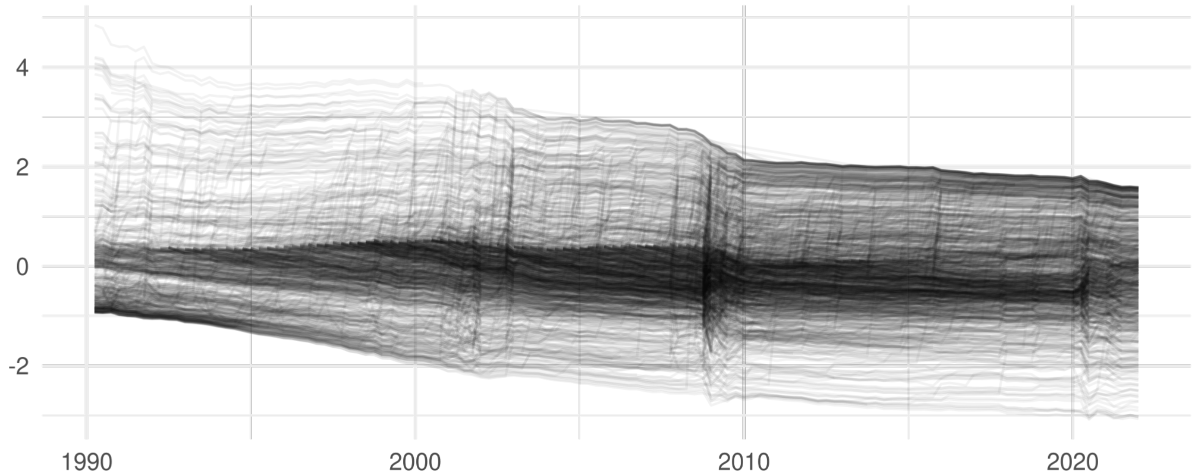

Figure 4.6: Firm variance ratios $\mathrm{VR}_{i,t}$

ance ratios is fat tailed, with a larger mass of firms having low variance ratios pre-2000. This shifts after 2000, with the bulk of firms being above the mean variance ratio. The 2008 crisis induced severe cross-sectional dispersion in variance ratios, and introduced a regime change in the average dispersion. The COVID-19 pandemic had a similar outcome on cross-sectional dispersion.

In this section, I outlined the mechanics of estimating my model. I then provided a brief overview of the results and the estimated parameters. The next section summarizes the existing literature on a set of topics that my model is capable of addressing.

Chapter 5

EMPIRICAL RESULTS

In this section, I perform several tests to determine whether the measures I construct have any statistical power for (a) earnings announcement premia and (b) pricing the cross-section of returns. All empirical tests begin at 1990 to eliminate highly volatile parameter estimates during 1980-1989, when my Bayesian model was relatively low-data and heavily prior driven.[1] I standardize the posterior variance of beta $\sqrt{\text{Var}[\beta_i]}$ to make the results more numerically interpretable.

Table 5.1 presents an overview of my sample, grouped by industry. The additional columns contain the parameters I estimate, discussed below. The top three industries are chemicals, depository institutions, and business services. These three industries collectively compose 24% of my sample. The largest industry by average market capitalization is miscellaneous retail. I calculate the average earnings growth $X_{i,t}$ for each firm as well – the fastest growing industry in my sample is stone, clay, glass, and concrete products. The quickest contracting industry is amusement and recreation services.

Table 5.1 highlights the average estimated parameters for my model, as well as the corresponding parameter uncertainty. Average betas are generally near one with some attenuation towards zero. Growth rates are near zero (as in Kothari, Lewellen, and Warner (2006)) and tend to have large posterior standard deviations, suggesting most unconditional firm earnings growth is due to aggregate earnings growth. Average idiosyncratic earnings growth suggests that between 1-3% of firm earnings growth is due to idiosyncratic risk.

I calculate the measures $\text{VR}_{i,t}$, $E[\beta_i]$, and $\sqrt{\text{Var}[\beta_i]}$ directly from my posterior samples in each quarter. I assign the above measures to cumulative abnormal returns in the quarter following to avoid bias, meaning all results are interpretable as the effect of a particular posterior quantity on CAR in the following quarter. The variance ratio $\text{VR}_{i,t}$, cashflow beta $E_t[\beta_i]$, and parameter uncertainty $\sqrt{\text{Var}[\beta_i]}$ are normalized to make estimates comparable across quarters. I subtract the quarterly cross-sectional mean from each observation and divide by the cross-firm standard deviation.

Table 5.2 highlights correlations between estimated quantities in my model. Many of the posterior variances are highly correlated. Firms with very uncertain cashflow betas also tend to have uncertain mean growth rates. The variance ratio is highly positively correlated with the posterior variance of firm parameters $\text{Var}_t[\mu_i]$, $\text{Var}_t[\beta_i]$, and $E_t[\gamma_i]$. There

[1] The exclusion of the 1980s does not change the qualitative results, but it does change the quantitative results as there are several extreme posterior estimates I generate during the 1980s.

Industry	Obs	Firms	Mkt cap	$X_{i,t}$	$E[\mu]$	$E[\beta]$	$E[\gamma]$	$\mathrm{VR}_{i,t}$
Chemicals and Allied Products	9748	130	20,648.4	0.06	-0.05	0.92	1.51	-0.00
Other	9718	134	13,406.9	0.23	0.05	0.96	1.68	-0.01
Business Services	9206	129	22,051.84	0.22	0.07	0.91	1.67	0.01
Depository Institutions	8910	124	13,552.6	0.20	0.15	0.97	1.59	-0.07
Insurance Carriers	6950	89	10,343.58	0.35	0.07	0.99	1.89	0.29
Electric, Gas and Sanitary Services	6910	78	9,924.22	0.05	0.05	0.90	1.48	-0.09
Electronic & Other Electrical Equipment & Components	6635	82	17,748.89	0.08	-0.04	1.03	1.72	0.02
Industrial and Commercial Machinery and Computer Equipment	6181	81	7,268.638	0.13	-0.06	1.02	1.66	-0.06
Oil and Gas Extraction	5760	85	10,147.25	-0.16	-0.09	1.02	1.95	0.52
Measuring, Photographic, Medical, & Optical Goods, & Clocks	5430	67	11,180.76	0.08	0.03	0.90	1.38	-0.31
Transportation Equipment	3844	48	14,176.23	0.45	0.03	1.06	1.95	0.14
Communications	3419	49	31,944.12	0.60	-0.03	0.95	1.89	0.11
Holding and Other Investment Offices	3123	40	8,139.993	0.29	0.06	0.93	1.67	0.08
Food and Kindred Products	2665	34	20,375.96	0.11	0.06	0.84	1.25	-0.39
Security & Commodity Brokers, Dealers, Exchanges & Services	2532	40	13,313.46	0.18	0.10	0.98	1.55	-0.18
Wholesale Trade - Durable Goods	1691	20	4,577.746	0.10	0.01	0.98	1.33	-0.58
Paper and Allied Products	1627	22	10,754.32	0.22	-0.07	1.02	1.49	-0.41
Primary Metal Industries	1611	22	4,306.214	-0.12	-0.45	1.12	2.34	0.59
Health Services	1587	20	8,334.007	0.20	0.09	0.90	1.64	0.20
Eating and Drinking Places	1522	22	11,925.99	0.06	0.07	0.86	1.30	-0.37
Metal Mining	1462	20	12,599.25	-0.07	-0.06	1.01	1.91	0.43
Petroleum Refining and Related Industries	1409	22	60,385.21	0.12	-0.10	1.05	1.65	-0.01
Fabricated Metal Products	1396	15	5,578.4	0.09	0.01	1.05	1.44	-0.34
Transportation by Air	1214	17	5,484.36	2.31	0.15	1.03	3.16	0.69
Motor Freight Transportation	1163	16	8,913.742	0.11	0.03	0.91	1.38	-0.32
Amusement and Recreation Services	1112	16	6,362.425	-0.49	-0.01	1.02	2.05	0.29
Engineering, Accounting, Research, and Management Services	1075	16	3,527.093	0.06	-0.01	0.93	1.57	0.05
Water Transportation	926	15	2,831.631	-0.21	-0.06	0.98	1.72	0.20
Wholesale Trade - Nondurable Goods	888	11	9,574.891	0.02	0.04	0.85	1.42	-0.29
Nondepository Credit Institutions	877	14	33,127.64	0.39	0.21	0.94	1.44	-0.28
Rubber and Miscellaneous Plastic Products	870	12	3,093.69	-0.05	-0.06	0.97	1.64	0.09
Apparel, Finished Products from Fabrics & Similar Materials	851	11	5,477.718	0.09	0.06	0.91	1.35	-0.42
Printing, Publishing and Allied Industries	789	10	3,287.056	0.08	-0.04	0.93	1.23	-0.54
Insurance Agents, Brokers and Service	696	9	11,165.12	0.12	0.12	0.75	1.20	-0.32
Lumber and Wood Products, Except Furniture	634	8	5,180.051	0.15	-0.07	1.09	1.43	-0.45
Railroad Transportation	616	7	19,115.12	0.15	-0.00	0.99	1.52	-0.06

Table 5.1: Overview of quarterly data for firms in my sample by industry.

	$E[\mu_i]$	$\text{Var}[\mu_i]$	$E[\beta_i]$	$\text{Var}[\beta_i]$	$E[\gamma_i]$	$\text{Var}[\gamma_i]$	log(mcap)	$\text{VR}_{i,t}$
$E[\mu_i]$	1.000							
$\text{Var}[\mu_i]$	-0.027	1.000						
$E[\beta_i]$	-0.153	0.295	1.000					
$\text{Var}[\beta_i]$	-0.047	0.606	0.542	1.000				
$E[\gamma_i]$	0.018	0.352	0.070	0.138	1.000			
$\text{Var}[\gamma_i]$	0.023	0.121	0.011	0.031	0.890	1.000		
log(mcap)	0.119	-0.361	-0.179	-0.371	-0.009	0.009	1.000	
$\text{VR}_{i,t}$	-0.070	0.783	0.035	0.527	0.328	0.099	-0.194	1

Table 5.2: Correlation matrix of regressors.

is little correlation between the variance ratio and the level of the cashflow beta, and understandably so – the ability of firm earnings to convey macroeconomic news is primarily a function of parameter uncertainty and fundamental uncertainty γ^2. This is largely consistent with Schmalz and Zhuk (2019) who show that information transmission about aggregate earnings is only a function of posterior variance, not posterior expectations.

5.1 Cross-sectional results

In this section, I place my measures of systematic risk uncertainty and informativeness into the canonical literature on cross-sectional asset pricing, and whether the factors I construct represent priced sources of risk. Section 3.2 provides further context and literature review on this topic area. I predicted that firms with higher cashflow beta parameter uncertainty $\text{Var}_t[\beta_i]$ should have higher cross-sectional returns due to increased systematic risk exposure.

To test the cross-sectional implications of my parameter estimates, I perform the standard regressions of Fama and MacBeth (1973) using several specifications. The cashflow beta $E_t[\beta_i]$ and $\sqrt{\text{Var}_t[\beta_i]}$ are cross-sectionally normalized by subtracting each variable from the average in a week and divided by the week's standard deviation of each variable in order to focus on the effects of the relative position of a firm's characteristics, since the posterior I draw from moves substantially between the beginning and end of the sample. Section 6.1 provides additional context on the scale and potential unidentifiability of β_i. I use Newey-West standard errors with three lags. I gather weekly returns for all firm-weeks in my sample, and estimate Fama-French 3-factor betas[2] for time t using weeks in $[t-120, t-2]$, requiring at least 60 weeks of returns. My cross-section contains 1,266,809 firm-week observations after eliminating firm-weeks where FF3 betas were not estimable.

Table 5.3 presents my results. The first specification regresses weekly returns on only a firm's cashflow beta $E[\beta_i]$, without controlling for the Fama-French 3-factor coefficients. I

[2]Data on the Fama-French factors is gathered from the Ken French data library.

Table 5.3: Fama-MacBeth regressions of weekly returns (in percent) on covariates. Standard errors appear in parentheses. $E[\beta_i]$ is a firm's expected cashflow beta. $\sqrt{\text{Var}[\beta_i]}$ is the standard deviation of a firm's posterior cashflow beta. Standard errors are Newey-West using three lags.

	Returns						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$E[\beta_i]$	0.006		−0.003	−0.011		−0.009	
	(0.012)		(0.007)	(0.008)		(0.010)	
$\sqrt{\text{Var}[\beta_i]}$		0.017*		0.013	0.006	0.016*	
		(0.010)		(0.009)	(0.007)	(0.009)	
$\text{VR}_{i,t}$							0.004
							(0.007)
$E[\beta_i] \times \sqrt{\text{Var}[\beta_i]}$						0.002	
						(0.002)	
β_i^{HML}			0.010	0.009	0.008	0.008	0.008
			(0.033)	(0.033)	(0.034)	(0.033)	(0.034)
β_i^{SMB}			0.048**	0.047**	0.046**	0.047**	0.040*
			(0.020)	(0.021)	(0.021)	(0.021)	(0.021)
β_i^{MKT}			0.029	0.026	0.027	0.026	0.037
			(0.039)	(0.039)	(0.040)	(0.039)	(0.039)
Observations	1,266,809	1,266,809	1,266,809	1,266,809	1,266,809	1,266,809	1,242,481
R-squared	0.278	0.275	0.347	0.349	0.347	0.350	0.351
Notes:	***p < .01; **p < .05; *p < .1						

find no support that my cashflow beta estimate has any substantive effect on cross-sectional asset prices. Column 1, 3, 4, and 6 all highlight no significant effect of a cashflow beta on asset returns. In fact, the point estimates seem to suggest a *negative* price, though they are insignificant – a negative cashflow beta price suggests that my model's cashflow estimate does not capture the same type of cashflow covariance risk measured in papers like Campbell and Vuolteenaho (2004) or Da (2009), both of whom find positive and significant cashflow betas.

The lack of a risk price associated with my cashflow beta could simply be the result of low statistical power. Chan and Marsh (2022) shows that the expected return predictions of the CAPM only hold during earnings announcement periods. My cashflow beta is not the same as the market beta they estimate, but Section 5.2 highlights that the cashflow beta does seem to be a priced risk during an earnings announcement week, consistent with Chan and Marsh (2022).

I find (limited) support for a risk price associated with the parameter uncertainty. Column (2) regresses cross-sectional returns on only the parameter uncertainty measure, and finds a positive risk price of 0.017. The economic magnitude is small. The median firm should expect an extra 0.5bps of weekly returns due to their parameter uncertainty, which annualizes to 29.7bps per annum. Firms with higher parameter uncertainty at the 80th percentile of parameter uncertainty have high returns by about 0.9bps per week, or 46.9bps per annum.

I test for any effect that omitting the cashflow beta might have on the risk price associated with parameter uncertainty. In Column (4), I include the cashflow beta with no interaction between the beta and its uncertainty. I find that the risk price is numerically similar to Column (2), but that it is no longer significant. Similarly, controlling for the Fama-French betas in Column (5) removes the effect of parameter uncertainty.

However, in Column (6), I permit an interaction between the beta and parameter uncertainty, and the effect I detected in Column (2) becomes significant again. $E_t[\beta_i]$ and $\text{Var}_t[\beta_i]$ are heavily correlated. My findings suggest that variation in parameter uncertainty provides incremental explanatory power once the correlated component between $E_t[\beta_i]$ and $\text{Var}_t[\beta_i]$ is removed by the interaction term. Importantly, the magnitude in this specification is quantiatively similar to the one in Column (2), where parameter uncertainty was the only regressor. Column (6) summarizes the main result of this section by providing evidence that parameter uncertainty in the cashflow beta explains variation in cross-sectional returns. My finding suggests a nonlinearity in parameter uncertainty, consistent with more complicated hedging-based models of parameter uncertainty (Cvitanić et al. 2006).

The interpretation in Column (6) may seem unintuitive since the main effect on parameter uncertainty is weakly significant, while the interaction is not. Note that both $E_t[\beta_i]$ and $\text{Var}_t[\beta_i]$ are normalized variables, which implies that the coefficient on $\text{Var}_t[\beta_i]$ indicates the effect of parameter uncertainty on returns is different from zero when $E_t[\beta_i]$ is at its cross-sectional mean. This leads to a follow-up question: why are no coefficients in Column (4) significant, where the interaction term is omitted? There are several explanations. First, it may simply be that the non-interacted specification is misspecified, and that the effect of parameter uncertainty is only locally identified once $E_t[\beta_i]$ is fixed. Second, as noted by Jun and Pinkse (2009), the addition of the "irrelevant" interaction term $E_t[\beta_i] \times \text{Var}_t[\beta_i]$ can improve the efficiency of the estimator.

The remaining columns show no effect on weekly returns. The variance ratio $\text{VR}_{i,t}$ does not suggest any difference in weekly returns. I suggest that perhaps a firm's informativeness about quarterly cashflows is insufficient to describe variation in weekly returns. Section 5.2 highlights that the informativeness measure seems to explain cross-sectional stock return variation for a firm's announcement week, precisely when the informativeness of firm earnings news is most valuable.

A common test of a cross-sectional pricing story is to provide portfolio sorts, as popularized by Fama and French (1993). To provide an analysis of the dual role of the cashflow beta and its associated parameter uncertainty, I provide two sets of portfolios. I construct univariate sorts of firms into quintiles based on their cashflow beta and on their parameter

	$\text{Var}_t[\beta_i]$			$E_t[\beta_i]$	
Quintile	Raw	Abnormal	Quintile	Raw	Abnormal
1 (Low)	0.220	-0.0230	1 (Low)	0.223	-0.0201
	(3.79)	(-1.47)		(4.23)	(-1.25)
2	0.230	-0.0331	2	0.244	-0.0304
	(3.47)	(-2.08)		(3.97)	(-1.90)
3	0.244	-0.0333	3	0.264	-0.0425
	(3.47)	(-2.00)		(3.83)	(-2.30)
4	0.262	-0.0348	4	0.268	-0.0314
	(3.52)	(-1.81)		(3.60)	(-1.50)
5 (High)	0.273	-0.0299	5 (High)	0.245	-0.0307
	(3.41)	(-1.50)		(3.16)	(-1.27)
5-1	0.041	-0.007	5-1	0.021	-0.011
	(0.44)	(-0.27)		(0.23)	(-0.36)

Table 5.4: Average weekly returns in percent grouped by the firm's quintile of $E_t[\beta_i]$ and by the quintile of a firm's parameter uncertainty, $\text{Var}_t[\beta_i]$. Abnormal returns are the residuals of the Fama-French 3-factor model.

uncertainty. I calculate abnormal returns as the residuals from constructing the Fama-French 3-factor betas on the portfolio's excess returns.

Table 5.4 highlights the average returns of univariate sorts on cashflow beta and on parameter uncertainty. Abnormal returns for portfolios sorted on the cashflow beta or on parameter uncertainty do not seem to be statistically different between the high and low portfolios on either axis. The limited support for my statistical results in the cross-sectional regressions is not surprising, given that the economic magnitude of parameter uncertainty on returns is small. Further, portfolio tests tend can be relatively low power as compared to regression tests when the number of cross-sectional observations is small, as in my sample (Cattaneo et al. 2020).

Note that several of the portfolios have abnormal returns that are significantly negative. This is in part the result of my sample selection, which tilts more towards firms with low book-to-market ratios (Zhang 2013). Similar results hold for untabulated analyses on high/low large/small portfolios provided by Ken French.

I conclude this section by noting that the cross-sectional regressions indicate that the cashflow beta is not priced, but that the parameter uncertainty is. However, portfolio sorts suggest that neither the cashflow beta nor its uncertainty are priced cross-sectionally, though portfolio sorts are likely to be relatively low-power given my small cross-section. I cannot conclude that there is *no* risk price associated with either attribute, but I can say that I do not find sufficient evidence.

5.2 Earnings announcements

Section 3.2 summarizes the current literature on earnings announcement premia and earnings response coefficients as it relates to parameter uncertainty. I contend that that firms with an increase in systematic parameter uncertainty should have higher earnings response coefficients.

I construct cumulative abnormal returns (CAR) for each firm-earnings announcement pair. I calculate abnormal returns for each day by regressing returns from $[t-120, t-14]$ (requiring at least 60 days' worth of returns) on the Fama-French three factors (size, value, and market). Abnormal returns are then the residual of the raw return on day t after removing the projected FF3 effects. Cumulative abnormal returns $CAR(s, k)$ are the sum of all abnormal returns between $[t+s, t+k]$. This is standard practice, see Kothari and Warner (1997) for one of many examples.

Lastly, I calculate a measure of raw earnings news, standardized unexpected earnings (SUE). I measure SUE as the difference between actual earnings and median forecasted earnings, scaled by the previous quarters' ending price. I require a firm to have at 5 analysts in the 90 days preceding a firm's scheduled earnings announcement. This is similar to the method used by Patton and Verardo (2012), among many others. My definition of SUE is written as

$$\text{SUE}_{i,t} = \frac{\text{EPS}_{i,t} - E[\text{EPS}_{i,t}]}{p_{i,t-1}}$$

I perform two types of tests. First, I test for earnings response coefficients, which describe the abnormal returns to a stock as a function of unexpected earnings news (SUE). Second, I test for earnings announcement premia, which are abnormal returns that accrue to firms regardless of their unexpected earnings. The earnings announcement premium tests are cross-sectional, while the earnings response coefficients examine within-firm variation in earnings response due to changing parameter beliefs.

In Section 3.2, I summarized the literature on parameter uncertainty, and suggested that firms with higher total parameter uncertainty might have higher earnings announcement premia due to a lower signal-to-noise ratio as per Savor and Wilson (2016). I also suggested that firms with increased parameter uncertainty should have higher earnings response coefficients – intuitively, if an investor is highly uncertain about a firm's parameters, any new information is weighted higher than it would be than if the firm had lower parameter uncertainty.

Table 5.5 tests these hypotheses. I find that when a firm's parameter uncertainty increases, they have *lower* earnings responses. Alternatively stated, this implies large earnings responses when parameter uncertainty is low. Column (1) indicates that earnings

Table 5.5: Regression of cumulative abnormal return (in percent) during an earnings announcement, CAR(0,1). Subsamples are by NBER recession indicators. Abnormal returns are daily, adjusted for Fama-French 3-factor betas estimated using returns from $[t-120, t-14]$, requiring at least 60 days of returns. SUE is the actual EPS less the median analyst estimate, scaled by the previous quarter's price. Standard errors are clustered by firm and by quarter. Fixed effects are by firm and by quarter.

	CAR(0,1)		
	(1)	(2)	(3)
$E_t[\beta_i]$	0.154***	0.147***	0.159***
	(0.048)	(0.049)	(0.049)
$\sqrt{\text{Var}_t[\beta_i]}$	0.129**	0.107*	0.127**
	(0.063)	(0.061)	(0.063)
$E[\beta_i] \times$ SUE		-12.239^{***}	-10.781^{**}
		(4.566)	(4.922)
$\sqrt{\text{Var}[\beta_i]} \times$ SUE	-96.173^{***}		-95.072^{***}
	(24.117)		(23.317)
SUE	102.068***	61.235***	107.900***
	(22.004)	(12.333)	(23.029)
Observations	79,919	79,919	79,919
R-squared	0.040	0.037	0.040
Adjusted R-squared	0.019	0.016	0.020
Notes:	***p < .01; **p < .05; *p < .1		

responses are reduced when firms have high parameter uncertainty (large $\sqrt{\text{Var}[\beta_i]}$). My evidence is consistent with the model in Schmalz and Zhuk (2019), who show that greater parameter uncertainty about the cashflow beta should yield lower earnings response coefficients.

Columns (2) and (3) demonstrate that higher betas also reduce a firm's earnings response, though the effect is smaller than the effect of parameter uncertainty. This finding is not surprising. Earnings news about a firm with a high cashflow beta should yield a weaker response, since more of the firm's earnings news is driven by common factors, meaning that stock returns incorporate more of the firm's earnings news. Belief revision about the firm's parameters should be relatively small, since the information conveyed by earnings news is driven more by the aggregate earnings shock.

The previous tests focus on earnings response coefficients, but do not address risk premia that accrue to firms with different parameterizations. To test risk premia, I conduct Fama-MacBeth regressions using dummies for firm-weeks that are announcement periods. The test of earnings announcement premia in Savor and Wilson (2016) is fundamentally a cross-sectional test, which I attempt to replicate through a regression framework rather than a portfolio sort.

The primary variable of interest is Announcement week interacted with the variance ratio $\text{VR}_{i,t}$, the cashflow beta $E_t[\beta_i]$, and parameter uncertainty $\text{Var}_t[\beta_i]$. My hypothe-

Table 5.6: Fama-MacBeth regressions of weekly returns (in percent) on announcement week indicators. $E[\beta_i]$ is a firm's expected cashflow beta. $\sqrt{\text{Var}[\beta_i]}$ is the standard deviation of a firm's posterior cashflow beta. $\text{VR}_{i,t}$ is firm i's variance ratio. Standard errors appear in parentheses and are Newey-West adjusted, using three lags.

	Weekly Returns			
	(1)	(2)	(3)	(4)
Announcement week	0.159*	0.063	0.093	0.077
	(0.088)	(0.117)	(0.078)	(0.157)
Announcement week $\times \text{VR}_{i,t}$		−0.221**		
		(0.112)		
Announcement week $\times E[\beta_i]$			0.626*	
			(0.379)	
Announcement week $\times \sqrt{\text{Var}[\beta_i]}$				0.167
				(0.281)
$\text{VR}_{i,t}$		−0.0001		
		(0.010)		
$E[\beta_i]$			0.004	
			(0.010)	
$\sqrt{\text{Var}[\beta_i]}$				0.004
				(0.009)
β_i^{MKT}	0.074*	0.080*	0.073*	0.072
	(0.043)	(0.045)	(0.043)	(0.044)
β_i^{HML}	−0.046	−0.049	−0.048	−0.048
	(0.033)	(0.034)	(0.032)	(0.032)
β_i^{SMB}	0.053**	0.055**	0.051**	0.050**
	(0.025)	(0.026)	(0.025)	(0.025)
Observations	666,573	653,789	666,573	666,573
R-squared	0.334	0.339	0.337	0.336
Notes:	***p < .01; **p < .05; *p < .1			

ses have concluded that, in general, it is difficult to determine the *a priori* expectation whether any of these parameters should increase or decrease premia. The closest available work that studies announcement premia as functions of my estimates is Savor and Wilson (2016), which unfortunately does not permit heterogeneity in firms. I provide these estimates to readers for further research on the effects of informativeness, cashflow beta, and parameter uncertainty on earnings announcement premia.

To ensure that the cross-sectional results are not driven by announcement weeks with a few announcers. I require a week to have at least thirty announcers and thirty non-announcers. This restriction reduces my sample size by nearly half. Table 5.6 presents the results of my cross-sectional regressions. Column (1) reinforces the Savor and Wilson (2016) finding that announcers accrue premia relative to non-announcers – announcers earn a premia of approximately 16bps during announcement weeks that is unexplained by the Fama-French three-factor model. Annualized, this 8.61%, which is similar in magnitude to the 9.9% reported by Savor and Wilson (2016).

Column (2) shows that more informative firms earn higher premia during their announcement weeks relative to less informed firms. The magnitude of the elevated premium is similar to the magnitude of the unconditional premium, which provides evidence that heterogeneity in the signal-to-noise quality of firm earnings about aggregate cashflow growth is an important feature of the earnings announcement premium. A 20th percentile informative firm has 13bps higher weekly returns for an annualized return of approximately 7.24%. The median firm has an announcement informativeness premium of 2.45bps which annualizes to 1.28%. I show that much of the premium that accrues goes to the most informative firms, and that my estimates seem to suggest that almost all the Savor and Wilson (2016) announcement premium is driven by firm informativeness.

Column (3) tests the effect of the cashflow beta on announcement premia. I show that higher cashflow betas correspond to higher earnings announcement premia during a firm's earnings announcement. The cashflow beta effect subsumes the unconditional earnings announcement premium just as in Column (2). The median cashflow beta firm accrues an additional 9.1bps during their quarterly announcement week, or 4.84% in annualized returns. Firms with relatively high cashflow betas in the 80% percentile accrue 21.62bps, or nearly 11.88% per annum. Column (4) shows no effect of parameter uncertainty.

The results in Table 5.6 suggest that much of the results in Savor and Wilson (2016) are driven primarily by heterogeneity in firm characteristics, and that a disproportionate amount of earnings premia accrue to firms that are informative or have high cashflow betas. As the model in Savor and Wilson (2016) does not permit any form of heterogeneity, my results suggest that additional modelling of announcement premia with a language for heterogeneity would be a fruitful topic for future research.

To conclude this section, I find that earnings responses for firms with high parameter uncertainty are lower, and that earnings responses for firms with higher cashflow betas are lower. I find that informative firms earn higher announcement premia, as do firms with larger cashflow betas.

5.3 Firm informativeness and earnings response

The variance ratio $\text{VR}_{i,t}$ summarizes the ability of a firm's earnings to precisely communicate aggregate earnings news. The previous section focuses on the effects of systematic parameter uncertainty. This section focuses on how responsive markets are to informative and uninformative firm news, as measured by their variance ratio. In each quarter, I divide firm variance ratios into quintiles. The method for calculating a firm's variance ratio is described in Equation 3.3. I use the same cumulative abnormal returns, controls, fixed

VR Quintile	Days since QE	Std. SUE	Median market cap	Avg. market cap	Avg. $E_t[\beta_i]$	Avg. $\text{Var}_t[\beta_i]$	Avg. $\sqrt{E_t[\gamma_i^2]}$
1 (informative)	26.07	0.005	7,296.2	23,559.4	0.979	0.068	1.05
2	27.72	0.006	4,137.1	13,870.5	0.928	0.086	1.30
3	29.20	0.007	3,011.3	11,917.7	0.933	0.091	1.44
4	30.17	0.011	2,607.2	11,106.7	0.958	0.093	1.65
5 (uninformative)	29.56	0.024	3,276.5	10,724.3	1.006	0.098	2.83

Table 5.7: Summary statistics of firms in my sample, grouped by the quintile of their variance ratio.

effects, and standard error clustering as in Section 5.2.

Table 5.7 collects summary statistics sorted by their variance ratio quintile, with quintile 1 being the "most informative" in the language of my model. The most informative firms are also generally the largest, though the size effect does not have a clear direction in the median for quintiles 2-5. The standard deviation of SUE is monotone increasing with the variance ratio quintile. Informative firms have the smallest variance in earnings surprise. This is in part an output of my model – highly predictable firms have low posterior variance in each of their parameters because the model has frequent, repeated, and consistent evidence about a firm's parameters. The motivation for the variance ratio is that low variance ratios indicate firms that are predictable and low-noise, which is consistent with the correlations in Table 5.7.

Variance ratios move throughout time, and different firms and industries shift between levels of informativeness to the market. Figure 5.1 highlights the time variation in informativeness by averaging the standardized variance ratio for each industry-quarter, filtered to the top-ten largest industries. Recall that a lower variance ratio indicates higher informativeness, so industries towards the bottom of the figure are more informative in the language of my model.

Depository institutions were among the least informative industries at the beginning of my sample, but towards the end are the most informative of the large industries. Curiously, depository institutions trended towards more informative continuously until the 2008 banking crisis. Insurance carriers are continuously among the least informative industries. Utilities tended towards highly informative, but seem to have lost their informative role for much of the 2000s (except for a high-informativeness period during the mid-2010s).

Next, I take my firm-level estimates of variance ratios to the data. Table 5.8 presents the results. All specifications use the standardized variance ratio $\text{VR}_{i,t}$, which is *negative* for informative firms. Column (1) indicates that firms who become more informative have higher earnings response coefficients. Since the variance ratio is standardized, the economic conclusion one should take away is that a one-standard-deviation decrease (meaning

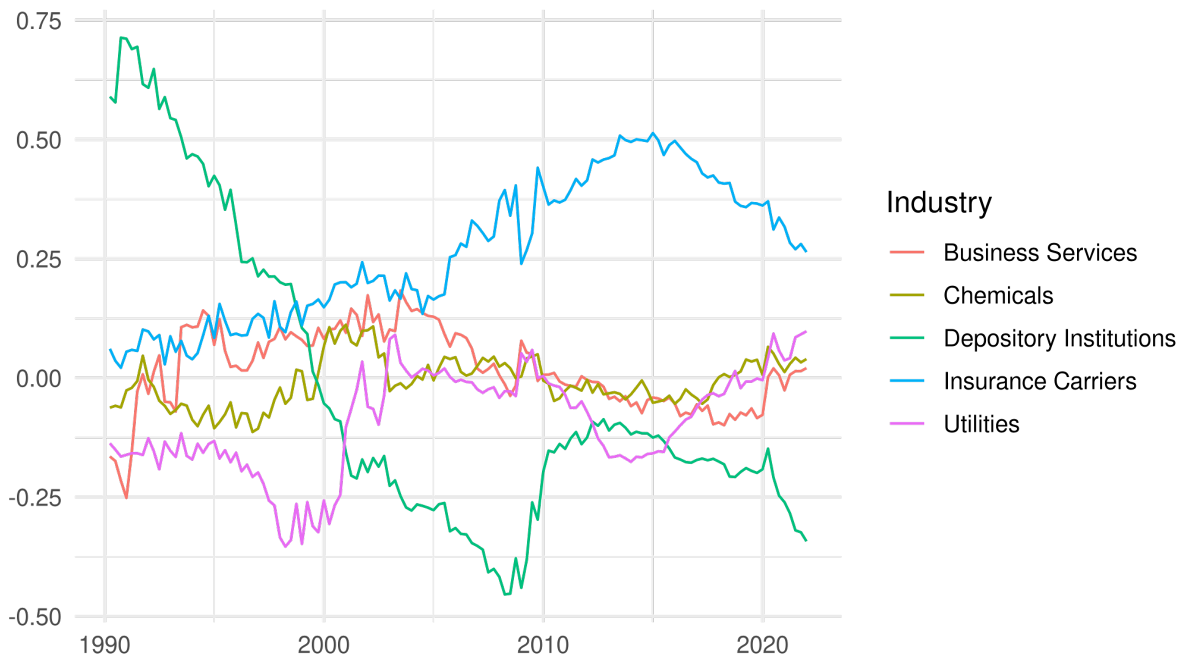

Figure 5.1: Average variance ratios $\text{VR}_{i,t}$ by industry, filtered to the top five largest industries in my sample. A lower variance ratio means the firm is more informative.

more informative) in the variance ratio means that, on average, the earnings response coefficient is 132.3, as compared to the median informative firm with an earnings response coefficient of 98.99.

My finding is consistent with the model of Schmalz and Zhuk (2019). In Lemma 5, for example, they demonstrate that when a firm's idiosyncratic variance or parameter uncertainty is high, earnings response coefficients are generally decreasing. Beyer and Smith (2021) produce a similar result. Firms should have higher earnings response coefficients when they are more informative.

I address the concern that the variance ratio provides little incremental information above a simpler measure like the R^2 of a regression of aggregate earnings on firm earnings. The R^2 would measure the capacity of a firm's cashflow news to convey information about the aggregate shock. I calculate the aggregate earnings growth measure of Kothari, Lewellen, and Warner (2006), defined as the year-over-year difference in the sum of firm earnings, rescaled by the total market capitalization of all firms in the prior quarter.

I estimate a regression for each firm and quarter t. In each regression, the data are an expanding window using data from 1 to t. I extract the R^2 statistic to compare its interpretation to that of my variance ratio $\text{VR}_{i,t}$. I multiply the R^2 by -1 to make its interpretation comparable to the variance ratio, since more negative variance ratios are more

Table 5.8: Regression of cumulative abnormal return (in percent) across a firm's earnings announcement, testing earnings responses as a function of firm informativeness. Abnormal returns are daily, adjusted for Fama-French 3-factor betas estimated using returns from $[t-120, t-14]$, requiring at least 30 days of returns. SUE is the actual EPS less the median analyst estimate, scaled by the previous quarter's price. Standard errors are clustered by firm and by quarter. Fixed effects are by firm and by quarter.

	CAR(0,1)	
	(1)	(2)
$\text{VR}_{i,t}$	0.116^{*}	
	(0.060)	
$\text{VR}_{i,t} \times$ SUE	-34.316^{***}	
	(7.794)	
SUE	97.987^{***}	58.977^{***}
	(20.701)	(15.980)
$-R^2_{i,t}$		0.141
		(0.558)
$-R^2_{i,t} \times$ SUE		46.295^{*}
		(26.200)
Observations	79,919	51,715
R-squared	0.041	0.052
Adjusted R-squared	0.020	0.023
Notes:	***p < .01; **p < .05; *p < .1	

informative. Similarly, $-R^2_{i,t}$ is more negative when firms are less informative. I find that the correlation between $-R^2_{i,t}$ and $\text{VR}_{i,t}$ is small, only 0.04.

Columns (2) perform the same tests, but this time with the $R^2_{i,t}$ calculated from projecting aggregate earnings growth on firm earnings growth. Notably, the $R^2_{i,t}$ measure does not reproduce the qualitative or quantitative results as when the variance ratio is used. The only significant coefficient related to the $-R^2_{i,t}$ measure is on the earnings response coefficient, which is of the opposite sign to the variance ratio tests – firms with high $R^2_{i,t}$ informativeness have *lower* earnings responses, whereas the variance ratio predicts *higher* earnings responses. I cannot say which of these estimates is more "correct", only that they correlate with different outcomes.

I find that the most informative firms are also those with the strongest earnings response coefficients. Markets respond more strongly to earnings news when firms are informative, consistent with rational learning about systematic risk. Further, I find evidence that my variance ratio provides quantitatively and qualitatively different interpretations of earnings response coefficients.

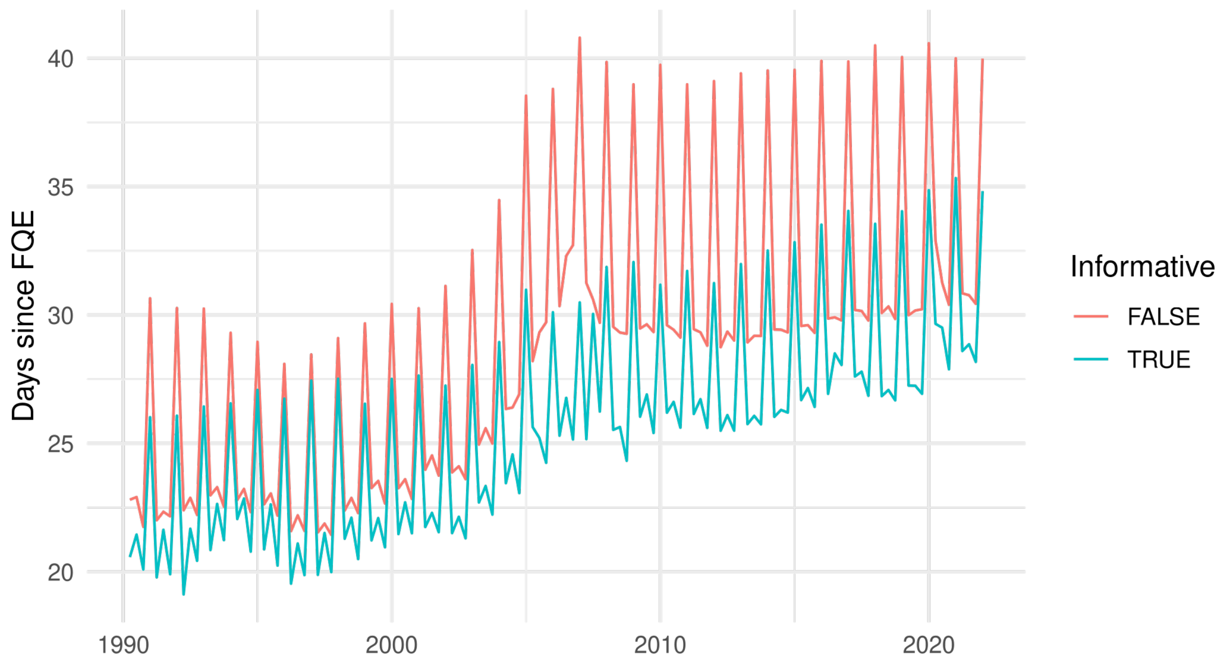

Figure 5.2: Average days-to-announcement from the end of firm fiscal quarters. *Informative* here is defined as a firm's quarterly variance ratio $\text{VR}_{i,t}$ being in the lowest quintile.

5.4 Earnings announcement order

Section 3.3.1 highlights the existing literature on earnings announcement order, and highlights that it remains an open question whether we should expect actual earnings orderings to be high-information first, low-information first, or some mix of both. To provide some guidance for future research, I estimate whether firms I predict to be high information announce earlier or later in my sample. I also examine whether high-information firms tend to announce earnings by themselves or as part of a cohort, since cohort information releases are more information-rich environments.

Using my data on earnings announcements from Section 5.2, I calculate the number of days between a fiscal quarter end and a firm's actual earnings announcement. Figure 5.2 graphs average days to fiscal quarter end by firm informativeness. I group all firms in the first quintile of their variance ratio per quarter and calculate the average number of days a firm takes to announce. The graph demonstrates a clear pattern of informative firms announcing earlier on average. The spikes in the graph represent additional delays in earnings announcements at year-end, when firms typically release an annual report.

I perform Fama-MacBeth cross-sectional regressions of the number of days since a firm announced earnings on their variance ratio $\text{VR}_{i,t}$, market capitalization, and book-to-

Table 5.9: Cross-sectional regressions of the difference between a firm's earnings date and the firm's fiscal quarter end in days. Columns (1) and (2) use the log number of days between a firm's quarter-end and the firm's earnings announcement date. $\text{VR}_{i,t}$ is the firm's variance ratio, cross-sectionally standardized.

	log(Days since FQE)		Days since FQE	
	(1)	(2)	(3)	(4)
$\text{VR}_{i,t}$	0.030***	0.052***	0.829***	1.445***
	(0.009)	(0.007)	(0.248)	(0.243)
BOOK_TO_MARKET	0.001		−0.043	
	(0.003)		(0.084)	
LOG(MCAP)	−0.033***		−1.042***	
	(0.006)		(0.177)	
CONSTANT	0.014	0.020**	0.176	0.686***
	(0.011)	(0.008)	(0.308)	(0.225)
Observations	79,653	115,011	79,653	115,011
R-squared	0.731	0.622	0.693	0.400
Notes:	***p < .01; **p < .05; *p < .1			

market ratios. I demean the number of days in each quarter to reduce the effect of differential average days-to-announce during the year-end. The specifications using market capitalization and book-to-market generally have a smaller sample size due to limited availability in accounting figures for some firms in my sample. I provide specifications controlling for market capitalization and book-to-market to demonstrate that variance ratio provides incremental explanatory power above size and book-to-market, which can proxy for financial distress.

Table 5.9 presents the regression results. As before, I use fixed effects and clustered standard errors by firm and by quarter.I find that, on average, informative firms announce significantly earlier than uninformative firms. All columns of Table 5.9 show positive and significant effects of the variance ratio on the number of days between the quarter end and the firm's earnings announcement. Recall that informative firms have *negative* variance ratios, so a variance ratio for a firm with a -1 standard deviation of informativeness of -1 would announce 0.829 days *earlier* (using Column 3) than firms with average informativeness. 0.829 days may seem a small effect, but Noh, So, and Verdi (2021) builds their work on a sample where nearly 70% of variation in days-to-announce is ± 1 day.

The results in Table 5.9 may provide some explanation of the results in Sections 5.3, which shows that earnings responses are higher for informative firms. It is possible that much of the result of those sections is driven by informative firms announcing earlier and with fewer co-announcers, though a more rigorous economic model of disclosure order may be needed.

Overall, I provide novel evidence that informative firms with highly predictable earnings tend to announce ahead of other firms. The existing literature is too sparse to place my findings into an appropriate context, but it is my hope that my empirical findings motivate further theoretical research to determine precisely why informative firms announce ahead of other firms.

Chapter 6

DISCUSSION

In this section, I provide discussion about my model, the assumptions used, and the results those assumptions may have on my results. Section 6.1 discusses the role of the latent factor M_t. Section 6.2 discusses the lack of time variation in the cashflow beta β_i, and provides alternative estimation procedures for the cashflow beta. Section 6.3 discusses the role of the Gaussian likelihood in my model. Section 6.4 analyzes my priors. Finally, Section 6.5 suggests additional research areas to which my model could be applied.

6.1 The latent factor

This section discusses the role the latent factor plays in my analyses. A primary goal of this paper is to quantify the uncertainty of a firm's cashflow beta with respect to an unknown latent factor. I propose a parametric description of the common earnings growth factor, and allow the data to inform the parameterization. Importantly, my latent factor is jointly estimated with firm loadings on that factor to better capture the difficulty an investor faces in interpreting earnings news.

The first concern is why there is residual posterior uncertainty given the size of the cross-section. One would expect that $\mathrm{Var}_t[M_t \mid X_{1:t}]$ would fall to zero as $N \to \infty$, but in fact there is a positive posterior standard deviation (between 0.06 and 0.3 in different time periods and across model estimates). For example, the model I estimate for the fourth quarter of 1984 estimates a shock of 0.11 ± 0.06, where 0.06 is the posterior standard deviation. As of the 2022 model, the estimate of the latent shock in 1984q4 has an estimate of 0.35 ± 0.09, an increase of 1/3 in standard deviation. Why then does the parameter uncertainty in the latent factors not fall to zero?

As with any estimator in a finite sample regime, my Bayesian estimate of M_t has uncertainty just as would a frequentist estimator of M_t. My estimator likely yields a slower convergence rate relative to a frequentist estimator since (a) the size of my dataset is small relative to the size of the parameter space, and (b) my parameter space grows with the cross-section – each firm added to my sample requires the estimation of β_i, γ_i, and μ_i, so simply growing the cross-section and holding the number of time periods T fixed will not resolve posterior variance in M_t. Exact convergence rates are difficult to derive for even simple Bayesian models, let alone models such as the one in this paper that do not have an obvious closed-form posterior.

Inference on the latent factors can be a difficult problem in general. In a Bayesian

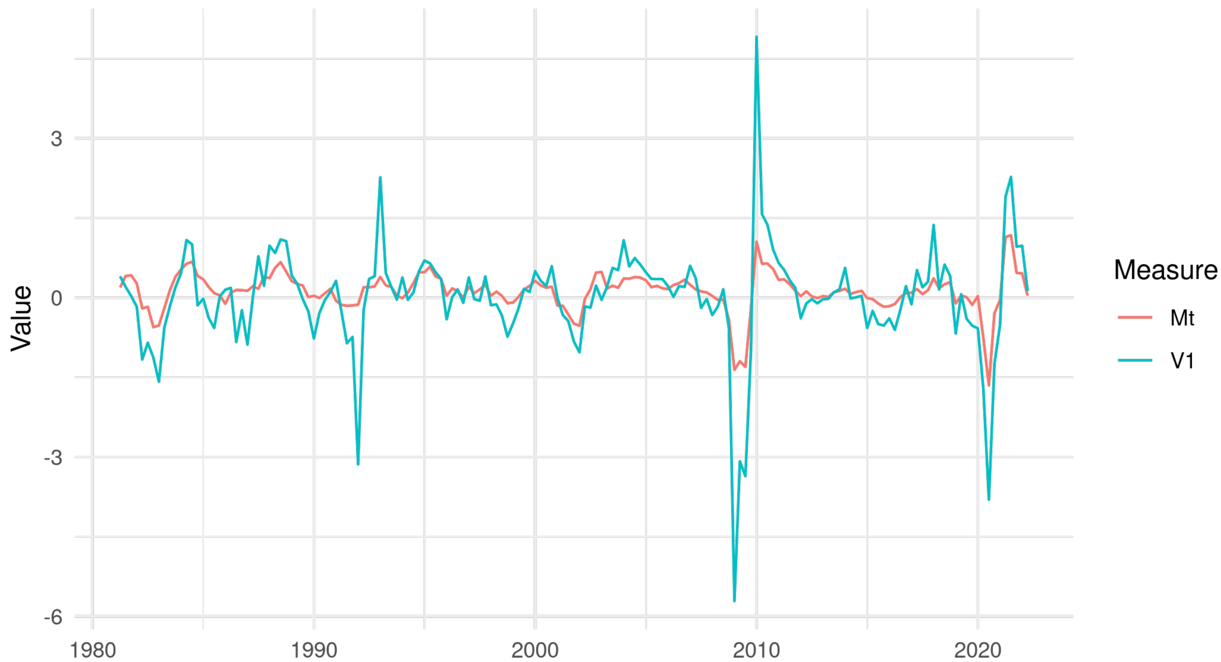

Figure 6.1: Comparison of the first principal component of earnings growth, compared to the latent process M_t.

framework, latent parameters need not always asymptotically converge (in time periods T) to the Bernstein-von Mises maximum likelihood estimate (Bontemps 2011; Li, Yu, and Zeng 2020), particularly for a fixed cross-section N. Latent parameters in my model are "partially consistent", as so named in Neyman and Scott (1948). Partially consistent variables have a finite amount of data attributable to them and may not even have a consistent maximum likelihood estimate, and similarly may not have an asymptotically consistent Bayesian estimator.

While such issues are important to consider, the primary benefit of Bayesian methods is their application to finite samples. We do not have a set of infinite time periods nor an infinite cross-section. Whether the Bernstein-von Mises theorem holds or not is important but not the key consideration – the focus of this paper is on a proposed representative learning process faced by an investor who does not live an infinite number of time periods, nor observe an infinite number of assets. The investor faces finite sample uncertainty consistently, but may not have a consistent estimator. What they do have is an uncertain guess about how cashflows share common cashflow risk that may not be fully resolved.

A second issue to consider of the jointly-estimated latent factor is that it may be unidentified, since the variance of the latent factor σ_2 and each firm's cashflow beta β_i are jointly estimated. In a traditional frequentist paradigm, estimating these two parameters together with maximum likelihood estimation is a fool's errand. It is less clear to what is identified under my joint Bayesian framework – the priors constrain σ^2 and β_i to a particular region, but the constraint imposed by the priors could have unanticipated effects on the full joint density $P(\theta(t) \mid X_{1:t})$. The magnitude of the latent factor shock is inversely related to the average magnitude of the cashflow betas in my model, and so only the ratio can be identified. A common practice is to fix the variance of the latent factor $\text{Var}[M_t]$ to 1, and allow the betas to be identified through a normalization.

Importantly, the parameter ϕ must also be modified to identify σ^2. Note that the variance of an AR(1) process is given by

$$\text{Var}[M_t] = \frac{\sigma^2}{1 - \phi^2}$$

In order to fix $\text{Var}[M_t] = 1$, any pair (σ^2, ϕ) must satisfy the identity $\sigma^2 = 1 - \phi^2$. The identity above means that σ^2 can be removed as a parameter since it is given as a direct function of ϕ. By maintaining ϕ as a parameter sampled from the domain $[-1, 1]$, $sigma^2$ is a parameter with a prior density defined only through a functional transform of the prior density on ϕ.

An alternative specification is to estimate a latent process via principal component analysis (PCA). Under this specification, one would estimate the latent factor, fix the latent process to the estimated factors, and allow only the firm's weightings on the latent process to fluctuate. Doing so is akin to assuming that the econometrician knows the latent factor with certainty but the factor loadings to be uncertain.

I provide visual evidence of the principal component alternative. Since I have an unbalanced panel with many missing values (as firms enter and exit at different times), I use the imputation method of Bai and Ng (2021) to produce a principal component. I standardize and demean all firm earnings growth $X_{i,t}$, and estimate the first principal component using all 165 quarters of data. Since principal components are identified up to a sign, I multiply the factor by -1 to match the peaks and troughs of M_t. Figure 6.1 demonstrates that the first principal component is a close fit to the estimated latent process $E[M_t]$. Quantitatively, the correlation between the two measures is 0.875.

The principal component factor is much more volatile than is M_t. A key reason is that the latent factor is a "predictability-weighted" latent process, in that common components are shared by firms that have earnings growth well described by my model. Firms with earnings that appear stationary and Gaussian across my sample end up molding the shape of M_t more than do other firms. Principal components analysis does not do this – it attempts to find a common factor that reduces the variance across all firms, without explicitly modelling firm earnings growth.

The principal components alternative could be an acceptable alternative to the joint latent factor distribution, though there are several shortcomings of the approach. First, principal components are a non-parametric method – there are no parameters that govern a system estimated by PCA. It is an ex-post rationalization of observed data without provision of a meaningful time-series structure. Certainly, this has a role in explaining realized common variation (as in Ball, Sadka, and Sadka (2009)), but the lack of an economic structure provided by PCA limits its use as an economic tool. For example, there is no longer a sense in which observed earnings growth connects to forecastable quantities as under my baseline model. Investors cannot make (noisy) predictions about future earnings, and so the value of learning is reduced. Much of the results in the learning literature are functions of explicit parametric forms for earnings or returns, and a non-parametric method simply cannot provide an analogue to a valuable source of economic intuition.

Second, the latent factor does not have an associated level of uncertainty that can propagate through to the factor coefficients in a tractable way. This may be positive or negative depending on your perspective. In some sense it is a good thing, since removing the latent factor's uncertainty alleviates some concern with the prior's role in determining the size and scale of β_i and σ^2. In a less positive sense, the key analysis of this paper is to provide a model of earnings growth that intuitively maps onto the learning process of an investor. Investors cannot know the latent factor any more than they can know the latent factor loadings, and one must expect uncertainty in the latent factor to drive up uncertainty in the beta coefficients. Shutting down the uncertainty of the latent space does not permit the flow of uncertainty from an uncertain space to the firm-level parameters.

To conclude, the latent factor could be replaced by a principal component estimate. However, further study is needed to determine whether the qualitative outcomes match those using the latent process M_t. I defer to future research further work on the estimation method of the latent factor and the associated parameter uncertainty in cashflow betas.

6.2 Time variation in the cashflow beta

A restriction imposed by my model is that the cashflow beta β_i is time invariant. Assuming a fixed beta is akin to assuming that a firm has the same risk level throughout its lifetime. A motivator for my modelling choice is to capture a firm's long-run average exposure to systematic risk. Exposure to long-run risk is generally considered a key component of the cashflow beta (Bansal and Yaron 2004).

Such a restriction may be too aggressive, particularly since researchers typically permit time-varying market or discount rate betas. Ghysels (1998), for example, highlights significant structural breaks in CAPM-style betas. Zhang (2005) micro-founds one such mechanism for time-variation in betas by noting that value firms (high book-to-market ratio) are simply riskier in bad times, due to such firms having more assets-in-place and fewer growth options. Gomes, Kogan, and Zhang (2003) provides a similar motivation for time-variation in betas.

Time variation is commonly detected in cashflow betas as well, though explicit modelling of time-variation in the cashflow beta remains more limited than the work done for market betas. Botshekan, Kraeussl, and Lucas (2012) performs their up-down cashflow-discount rate beta decomposition on time-varying betas estimated on 60-month rolling windows. Campbell, Polk, and Vuolteenaho (2010) and Campbell and Vuolteenaho (2004) both highlight striking differences in their cashflow beta estimates after they separate their sample into two time periods. Da (2009) provides time-varying versions of cashflow covariance and duration, the paper's two key attributes of cashflow sensitivity.

Formal economic models of time-variation in cashflow betas are sparse, though appear capable of reconciling anomalous outcomes in asset pricing. Li and Zhang (2017) expand the long-run consumption risk model of Bansal and Yaron (2004) to include time-varying cashflow betas on short- and long-term cashflow risk, and find significant explanatory power for momentum and value portfolio returns.

The fixed beta of my model is at odds with much of the prior literature, and indeed it may be a primary cause as to why my results on the cashflow beta $E_t[\beta_i]$ are not as clear as might be expected. It may even play a role in the partial significance of parameter uncertainty $\text{Var}_t[\beta_i]$, in that parameter uncertainty may function as a catch-all variable for firms predisposed to structural breaks in their cashflow betas. In this case, my measurement of parameter uncertainty may not necessarily describe the quantity it is intended to capture.

A potential solution to time-variation in betas is to estimate posterior densities in sequences on earnings growth, rather than on expanding windows. Using an iterative procedure would remove the restriction that the beta be fixed in the likelihood, but would maintain its Bayesian interpretation. Cashflow betas would then be more flexible and reactive to current data. However, there are important considerations in how the model would be estimated in sequence.

The priors of the model must be updated in sequence. In the current version of the model estimation, priors are fixed at the beginning of the sample, but the expanding window means that the role of the prior deteriorates as more data are used. In a sequential estimation procedure, using the same prior constructed at the beginning of the sample is not consistent with a Bayesian methodology – the posterior at time t should act as the posterior in time $t + 1$.

Rather than the current posterior $P(\theta(t) \mid X_{1:t})$, consider the alternative time-varying joint density

$$P\big(\theta(t) \mid X_t\big) \propto P\big(X_t \mid \theta(t)\big) P\big(\theta(t) \mid \theta(t-1)\big) \tag{6.1}$$

Note that the prior $P\big(\theta(t) \mid \theta(t-1)\big)$ is the density of the parameterization $\theta(t)$ evaluated at the posterior density using the set of earnings growth X_{t-1}. The likelihood of the vector X_t is then

$$P\big(X_t \mid \theta(t)\big) \propto \mathcal{N}(X_t \mid \mu + \beta M_t, \gamma^2 I) \tag{6.2}$$

Note that this is a unit-observation model, in that only a single quarter is exposed to the posterior at a time. I do not allow for a rolling window of earnings growth. A rolling window is inappropriate in my model, since overlapping observations between quarters will result in excess "certainty" of the posterior. For example, consider the time t posterior estimated using earnings growth between $t-k$ and t. The prior, estimated at $t-1$, uses earnings growth between $t-k-1$ and $t-1$, meaning that the overlapping data available to the time t posterior includes k samples that are already included in the posterior at time $t-1$. If the k observations are included in the posterior estimation for t, the resulting density will be overfitted on replicated data.

Estimating the sequential density is simple given an analytic posterior, but unfortunately my model does not permit a set of tractable closed-form posterior densities. I sample from the posterior density using Markov chain Monte Carlo instead. Markov chain Monte Carlo is not a sequential method, in that the posterior samples cannot easily or accurately be used as inputs in a subsequent prior. I propose three alternative estimation routines that must be performed to permit the sequential updating of the posterior density $P\big(\theta(t) \mid X_t\big)$.

First, rather than use NUTS as the inference algorithm, one could instead use a sequential Monte Carlo (SMC) method. SMC is an algorithm that performs inference on a *series* of distributions much like the ones in Equation 6.1 (Chopin 2002; Chopin 2004). Such an approach is technically uncomplicated with probabilistic programming, and may be the simplest solution in general.

Second, one could maintain NUTS as the sampler and approximate each posterior by fitting a distribution to the marginal posterior density. For example, each β_i would have a Gaussian prior $P(\beta_i \mid X_{t-1})$, with the prior mean $E_{t-1}[\beta_i]$ and prior variance $\mathrm{Var}_{t-1}[\beta_i]$. The key downside of this posterior-fitting approach is that the co-movement structure between parameters is lost – for example, if β_i and β_j are highly positively correlated, fitting priors on the marginal densities only will lose this information and require the model to learn the co-movement from the new observation X_t. The priors in this case are not incorrect in any sense, but they are less informative in that a large amount of valuable information about the relationships between parameters is lost.

Third, the model could be re-written to accommodate time-variation in beta directly. I defer this conversation to Section 7, which provides a detailed analysis of alternative model specifications that can address time-variation in the cashflow beta.

I suggest that future research in the probabilistic estimation of cashflow betas utilize a sequential estimation method, whether that is sequential Monte Carlo, priors set by fitting known distributions to posterior estimates, or an alternative model specification. The method applied in this paper may not be capturing an appropriate cashflow beta, since as the model fixes the cashflow to be constant throughout the 40-year sample period.

6.3 The Normality assumption

The use of the Gaussian likelihood may play a large role in my analyses. This section describes the role of the likelihood in my model and the influence it may have on my posterior density. The likelihood is Gaussian in my model due to the assumption that idiosyncratic shocks ν_t are Gaussian. Recall the likelihood is specified

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

for mean $\mu + \beta M_s$ and variance-covariance matrix $\gamma^2 I$.

A key reason that I selected the Gaussian distribution for the likelihood is that it has a simple conditional expectation. I use the conditional expectation of a multivariate Gaussian in my construction of the variance ratio to derive the amount of variance reduction attributable to a given firm's earnings. This is the primary reason for selecting a Gaussian shock. However, it is important to consider the role that the Gaussian likelihood plays in my posterior estimates, and to attempt to understand in which ways it may distort my empirical tests.

The impact of using Gaussian errors ϵ_t can have unintuitive results on the posterior density. Müller (2013) highlights that misspecified likelihoods can cause certain obvservations to be weighted more strongly than others, though the degree and direction of this weight adjustment is difficult to determine in general. My model is high-dimensionsal and minor misspecifications could cause my posterior estimates to be pushed into regions that may not necessarily provide sufficient coverage of the "true" parameterization of my model. Unfortunately, the size of my model makes it difficult to determine ex-ante what role the

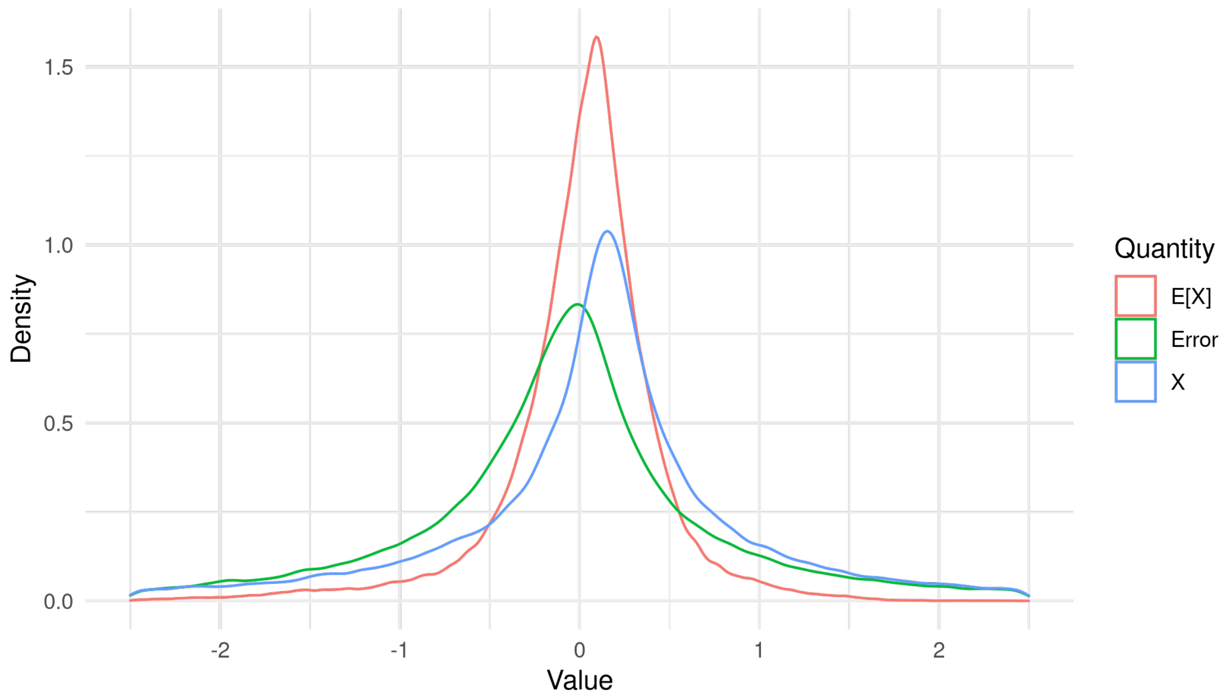

Figure 6.2: Density of the posterior expectation $E_t[X_{i,t}]$, realized earnings growth $X_{i,t}$, and the model error $E_t[X_{i,t}] - X_{i,t}$ starting from 1990.

Gaussian likelihood may play.

What I can highlight is that my model seems to capture an element of non-normality in its estimation, which may address some concerns. Figure 6.2 highlights the difference between the model-predicted earnings growth $E_t[X_{i,t}]$ and the observed earnings growth $X_{i,t}$. Note that the posterior expectation has much thinner tails and muted skewness, while the observed earnings growth in blue demonstrates significant fat tails and moderate negative skewness. The forecast error in green displays similar fat tails but is mean zero. The form of non-normality captured in the model prediction, $E_t[X_{i,t}]$, is non-normal through excess kurtosis, but is only minorly skewed. One could say that the thin tails in $E_t[X_{i,t}]$ are not correctly describing earnings growth, and that the posterior density is biased towards zero to accomdate the skewness and fat tails of observed earnings growth $X_{i,t}$.

However, it seems that the variance ratio $\text{VR}_{i,t}$ describes firms that are the most prone to this non-normality. To test this, I group all firms in my sample by their variance ratio quintile, pool their forecast errors, and calculate Jarque-Bera normality test statistics on the pooled model error $E_t[X_{i,t}] - X_{i,t}$. Table 6.1 presents the results. Notably, the most "informative" firms in my sample, those in quintile 1, have the *highest* deviations from normality in my sample. The variance of model errors is similarly much lower, though they are significantly more skewed and have a higher kurtosis than do firms in other samples. The informative firms have the lowest parameter uncertainty as well. The model detects,

$VR_{i,t}$ Quintile	Jarque-Bera	Mean	Sd. dev.	Skew	Kurtosis	Var[β_i]	$VR_{i,t}$
1	39099.981	-0.015	0.962	0.290	9.402	-1.259	-1.031
2	30881.801	-0.037	0.986	0.179	8.711	-0.507	-0.033
3	16047.247	-0.044	1.111	0.138	7.138	-0.113	0.201
4	5621.551	-0.118	1.360	0.112	5.495	0.355	0.312
5	444.687	-0.156	1.738	0.081	3.749	1.485	0.525

Table 6.1: Sample normality statistics on the model error $E_t[X_{i,t}] - X_{i,t}$, by variance ratio quintile. The first column is the Jarque-Bera normality test statistic. Other columns describe the mean, standard deviation, skewness, and kurtosis for model errors. The last two columns detail average variance ratios and parameter uncertainty.

perhaps correctly, that the firms with the highest skewness are those that are most informative about the macroeconomic state, and at the same time is able to precisely estimate their cashflow betas.

This could be consistent with the model allocating too much weight in the likelihood to informative firms with non-Gaussian earnings. Such an overweighting would also significantly move the posterior density of the latent factor M_t, which could cause unintuitive outcomes on the posterior density of other firms in my sample. The firms I find to be the more informative about the latent economy are those with the *most* non-Gaussian earnings growth. Informative firms have highly skewed model errors, and significantly fatter tails, suggesting that the likelihood weighting overweights firms with non-Gaussian earnings.

The answer as to whether the Normal likelihood is an issue for my estimates is difficult to determine without a formal Monte Carlo analysis or more detailed econometric analysis of my model. Further, the answer depends on the object of study. If the objective of the analysis is to construct a variance ratio or other measure of informativeness, then it is not obvious that the Gaussian likelihood is an issue: it is tractable, easy to use, and yields a variance ratio that captures skewness in earnings growth, and appears to produce estimates that are consistent with theory and existing literature.

6.4 The role of the prior

In this section, I discuss the role my selection of priors may play in determining my results. Priors play an important role in Bayesian analysis. Selecting appropriate priors in even simple models can have outsized effects on the posterior densities. In more complex models like the one applied in this paper, priors and their relative strength can influence the posterior in unintuitive ways. For example, consider a piece of evidence X_t that is less consistent with a common source of variation in earnings growth. My selection a low-variance prior on the cashflow beta β_i means that the posterior may adjust the latent factor variance downwards instead of reducing the betas if the prior variance on σ^2 is relatively low. Further, the adjustment of σ^2 in response to the tightness in the cashflow beta prior can influence the location and scale of the densities for the latent factor mean δ, and even idiosyncratic means μ.

In short, a detailed prior analysis is a nontrivial exercise. The size and computational burden of my model make it difficult to conduct the formal prior sensitivity tests common in Bayesian analyses, so instead I provide some context from a modeler's perspective by commenting on the priors of relevant parameters that I believe play a significant role in the results I obtain. The priors I believe play a significant role are those on cashflow betas β_i, idiosyncratic volatility σ^2, firm idiosyncratic volatility γ_i^2, and the latent persistence parameter ϕ.

The prior on the cashflow beta, β_i, is homogeneous across firms $\beta_i \sim \mathcal{N}(1, 0.1I) \quad \forall \quad i$. The cashflow beta prior is a simple one intended to predispose my posterior towards a common structure in firm earnings growth. Early posterior estimates of the model in this paper with a larger standard deviation of 1 resulted in a model that was too permissive – the posterior estimates concluded that there was very little evidence of a common variance structure in earnings growth. Such a finding strongly disagreed with my personal priors and with the existing literature (Ball, Sadka, and Sadka 2009).

The cashflow beta prior I applied is also overly simplistic in that it is homogeneous for all firms. There is little reason to expect that a large and old firm that sells consumer goods should have the cashflow risk exposure as a software firm. Indeed, one of the primary contributions of Li and Zhang (2017) is the evidence that differential exposure to firm short- and long-term cashflow risk appears differently in firms sorted on size, value, and momentum.

A more well-reasoned prior would be to assign a relatively lower prior for beta to firms with lower book-to-market ratios following the evidence in Campbell and Vuolteenaho (2004). To implement the differential priors in my model, could calculate the rational difference between the firm's estimated beta as a proportion of the sample average cashflow beta. Using Campbell and Vuolteenaho, this would imply an average beta of 0.55 for small growth assets, 1.18 for small value assets, 0.27 for large growth, and 1 for large value. Permitting heterogeneous priors in cashflow betas should allow for a more informed posterior that better incorporates the existing body of work on cashflow beta estimation.

The second prior to consider is the prior on common risk volatility. I choose a prior that yields an average variance equal to the observed earnings growth variance across all firms between 1980 and 1985. The prior thus assumes all cross-sectional variance is due only to systematic innovations, though this is not a particularly strong assumption for a large cross-section. Indeed, the principal component analysis conducted in Section 6.1

suggest the variance of the first principal component has a standard deviation of roughly 14%, which is close in magnitude to the empirical variance in earnings growth in 1980-1985 of 15.98%.

The joint selection of priors on σ^2 and β_i could be problematic, in particular since the magnitude of the latent factor is not formally identified without an additional source of data[1]. The choice of $E[\beta_i] = 1$ and $E[\sigma^2] = 15.98$ is equivalent to an alternative prior(in terms of expected common risk exposure) where $E[\beta_i] = 3$ and $E[\sigma^2] = 15.98/3 = 5.33$. The choice of prior location where beta is centered at one is roughly equivalent to one where it is centered at 3, as long as the other variables are scaled accordingly. I discuss in Section 7 an alternative specification that removes this issue by fixing $\sigma^2 = 1$ and allow the cashflow betas to dictate the magnitude of the latent factor shocks.

My choice of prior on σ^2, β_i, and γ_i^2 is intended to suggest that the common factor provides nearly half of the observed variance for firm i, since the priors on γ_i^2 is the result of fitting an inverse Gamma distribution to half of a firm's 1980-1985 variance if the firm existed. If the firm did not exist and entered much later in the sample, it will enter with the prior assumption that it share the same variance structure as a firm from 1980-1985.

Just as with the priors on firm betas, there is reason to expect significant heterogeneity in idiosyncratic risk γ_i^2. Firms in financial distress, small firms and firms in new or novel industries should have significantly less predictable returns. Firms with low posterior variance ratios are firms with significantly lower idiosyncratic earnings volatility, and similarly low earnings surprise volatility.

Fortunately, my priors handle this heterogeneity in part, as long as a firm's earnings variance in 1980-1985 is illustrative of the firm's idiosyncratic volatility. For firms that enter after 1985, there is zero heterogeneity in priors on idiosyncratic variance. Such a non-informative prior is perhaps too simplistic – a better prior would be to permit the average idiosyncratic variance to vary by industry, which is typically known quite early on in a firm's lifetime.

The final prior I consider is the persistence of the latent factor, ϕ. The persistence parameter ends up having a significant influence on the size of the systematic risk a firm is exposed to. The (conditional) variance of a firm's earnings growth attributable to the latent factor is $\beta_i^2\sigma^2/(1-\phi^2)$. For posterior samples of ϕ that are zero (for example), this is simply $\beta_i^2\sigma^2$. But as ϕ approaches 1 in absolute value, the firm's exposure to common risk increases. The posterior mean estimate $E_{165}[\phi] \approx 0.44$ suggests that the magnitude of a firm's total variance is nearly 25% higher than would be expected by just β_i and σ^2.

The role of the persistence parameter in determining the total amount of common risk

[1]One method of rectifying the unidentifiability of the latent factor is to include a component in the likelihood for the latent factor. For example, one could assume that the aggregate earnings measure of Kothari, Lewellen, and Warner (2006) is simply M_t measured with some unit of noise, which would permit the latent variance σ^2 to be matched to observed data.

faced by a firm is important, and the prior may not appropriately reflect this risk. Recall that my prior was specified as $\phi \sim \mathcal{N}(0, 0.2) \in [-1, 1]$, which has a standard deviation of roughly 0.45. This prior is quite wide relative to its domain, and may have resulted in a much larger posterior variance for all parameters in my model. A tighter but positive prior, perhaps $\phi \sim \mathcal{N}(2, 0.05) \in [-1, 1]$, would allow the model to focus more on the parameter of interest, β_i.

The remaining parameters are less critical to the discussion in that they are intended to capture residual variation in earnings growth, such as firm- or common-level earnings trends. I conclude this section by noting that my priors are reasonable, but could utilize more *a priori* information about the fundamental processes that govern earnings growth.

6.5 Additional topics

In this section, I consider further analyses that could be studied with my model. The model I study provides a rich description of firm earnings growth, though only parts of the model output are used. I generate estimates of numerous quantities that are not focal points of my analysis. For example, I estimate average firm earnings growth μ_i and idiosyncratic risk γ_i^2, though those quantities are generally marginalized in my computation of the cashflow beta $E_t[\beta_i]$, parameter uncertainty $\text{Var}_t[\beta_i]$, and the variance ratio $\text{VR}_{i,t}$.

Further study is needed on the informational effects of earnings announcement order. It is striking that my findings that firms with low variance ratios tend to announce early, even using the crude approximation to the true conditional informativeness of a firm's earnings news given by the variance ratio. Recall that the variance ratio is computed through an implicit counterfactual assuming that a firm were to go first, without accounting for preceding earnings news from other firms. The value of the variance ratio as computed becomes less useful for firms that empirically announce later, since announcements from late announcers have a significant amount of conditioning information available.

Explicitly studying the full earnings announcement order would be a novel application of the model I propose. Guttman, Kremer, and Skrzypacz (2014) studies the role of disclosure timing, but in a noncompetitive environment. Suppose firms were characterized only by the signal-to-noise ratio their earnings news provides about common earnings growth. Firms are allowed to partially determine their announcement order. Should we expect firms with high signal-to-noise ratios (the variance ratio in my model) to go first? There is limited theoretical support for such a claim, but providing some guidance clarify the extent to which my findings on the variance ratio are expected under a rigorous analytic framework.

Using the model in this paper, a researcher could model observed and counterfactual orders and compare the informational efficiencies from the perspective of a representative investor. How efficient is the current earnings announcement order? Does the existing order resolve uncertainty early or late, and how quantitatively efficient is that order relative to the first-best optimal order? Questions like these cannot be resolved without a measure of the informational efficiency of a firm's earnings news and a well-grounded theoretical model of earnings announcement ordering.

An additional expanded study built upon my model would focus more closely on the learning process throughout the earnings season. Estimating a new model in each earnings week (rather than in each quarter) to better track the evolution of beliefs about firm parameters throughout the earnings season. One question that could be addressed is the effects of belief revision. For example, a firm that announces bad news early during a quarter may have their poor performance attributed to idiosyncratic factors. If it becomes apparent that the common earnings shock is negative during the earnings season, one might expect beliefs about the firms' parameters to revise as more firms announce. Firms that experience belief revision should have some form of reversal in returns following their announcement as the proportion of their earnings news believed to come from idiosyncratic factors decreases.

Analyzing belief revision throughout the earnings season is straightforward with my model, in that calculating belief revision can be done by calculating a prior-posterior distance measure like KL divergence or the Wasserstein metric. Large belief revisions for firm i that occur after another firm j announces should cause proportional revisions in returns. Zhang (2006) finds that increased uncertainty during an information release corresponds to greater under-reaction to firm news. Zhang's findings could be augmented by explicitly calculated measures of uncertainty and of signal-to-noise ratios. Providing explicit posterior distributions for parameters and determining if prices respond to directly estimated belief revision can provide evidence about the persistent discussion on the relative influence of rational learning and behavioral biases (Brav and Heaton 2002).

Lastly, my model as implications for the predictive capacity of earnings news. Much of the analysis in this paper focuses on the ability of a firm's earnings to reveal the contemporaneous common earnings shock. However, there is an as-yet unused component of the model that can be applied to the forecasting of future earnings shocks. The parametric form of M_t (when $\phi \neq 0$) permits some measure of forecastability, in that understanding the current level of M_t provides a more accurate prediction of the level in M_{t+1}. The same firms who provide accurate nowcasts (low variance ratios) may also be the same firms who provide accurate forecasts, and it is worth evaluating to what degree these firms' earnings news predicts future common shocks and future firm earnings growth. Focusing on forecasting would allow a researcher to test alternative specifications in the specification for the latent factor. Comparing various time series models (ARIMA, ARMA, etc.) could provide valuable insight into which latent factor permits firms news to convey the most earnings news.

To conclude this section, I note that my model has many promising applications to the study of information revelation and learning. Extending my model to a more rich analysis of earnings announcement order would be a fruitful topic for future study due to the limited existing theoretical and empirical literature. My model produces estimates of earnings growth that can be well-applied to the rational learning literature. Lastly, I suggest a direction of study for understanding the ability of firm earnings news to forecast aggregate earnings.

Chapter 7

ALTERNATIVE MODEL SPECIFICATION

This section considers an alternative model specification after considering the discussion in Section 6. The two key issues are (a) time variation in the cashflow beta β_i, and (b) fixing the latent factor variance to provide identification of β_i. Fixing the variance of the latent factor is simple, as doing so amounts to setting σ^2 to 1.

The more challenging modelling problem is to permit time-variation in firm cashflow betas. Estimating the model in a one-step filtering approach as in Section 6.2 can address time-variation in betas without modifying the likelihood of the model. However, one could improve the economic interpretability of the model through the inclusion of an explicit form for a time-varying beta.

What form would such a beta take? The literature supports time-variation in cashflow betas, but at the same time also suggests that firms may have different exposure to short- and long-term cashflow risk. The seminal work of Bansal and Yaron (2004) focuses on exposure to long-term cashflow risks. Li and Zhang (2017) expands Bansal and Yaron (2004) to include firm exposure to short-term shocks in consumption growth. Da (2009) motivates his study through heterogeneous asset exposure to cashflows that arrive at different times. My model should attempt to capture the same exposure, and to permit the analysis of parameter uncertainty in addition to known effects due to the level of the cashflow beta.

An improved model would permit variation in firm betas within the model rather than modifying the estimation procedure, as described in Section 6.2. Any model extension should permit the ease of estimation afforded by Hamiltonian Monte Carlo methods due to its desirable high-dimensional convergence properties. Finally, the model should permit some notion of short- and long-term exposure to common cashflow risk.

Consider the modified laws of motion

$$X_t = \mu + \beta_t M_t + \nu_t \tag{7.1}$$

$$\beta_{i,t} = \rho_i \beta_{i,t-1} + (1 - \rho_i)\bar{\beta}_i + \xi_{i,t} \tag{7.2}$$

$$M_t = \delta + \phi M_{t-1} + \epsilon_t \tag{7.3}$$

where $\text{Var}[\epsilon_t] = 1 - \phi^2$, $\text{Var}[\nu_t] = \gamma^2 I$, and $\text{Var}[\xi_t] = \varsigma_i^2 I$. Note that σ^2 is no longer a parameter to be inferred but a function of ϕ – this addresses a major criticism in the model I estimate, and permits better identification of firm-level cashflow betas $\beta_{i,t}$.

The new component of the model is the time-series process for firm i's beta in Equation

7.2. Under this process, a firm's beta varies with time, and has a short- and long-term component. The short-term cashflow beta is given by its current level $\beta_{i,t}$. The long-term cashflow beta $\bar{\beta}_i$, which measures the unconditional average cashflow beta, captures the firm's exposure to long-term systematic risk. The autocorrelation coefficient $\rho_i \in [-1, 1]$ measures the degree of persistence in deviations from average $\bar{\beta}_i$. The variance of beta shocks $\xi_{i,t}$ is ς_i^2.

My model is similar to a simplified version of the one in Li and Zhang (2017), which incorporates an AR(1) process in the cashflow beta. Li and Zhang (2017) estimate their model using the simulated method of moments, and find small exposure to short-run risk but large and persistent exposure to long-run risk. In Equation 7.2, this would be equivalent to a small ρ_i, the parameter that governs short-term risk persistence.

To estimate this model, one could estimate $\beta_{i,t}$ using a conditional filter (if $\xi_{i,t}$ were assumed Gaussian, this could be the Kalman filter) during each MCMC step. Doing so would mean that the only new parameters added to the model would be ρ_i, $\bar{\beta}_i$, and the variance term ς_i^2. For a cross-section of N firms, this yields $2N$ new parameters, since the term β_i in my model is replaced by $\bar{\beta}_i$.

Regarding the priors of this model, I would incorporate the priors discussed in Section 6.4, with several additions to accommodate the alternative specification for the latent factor. The prior for firm means μ_i would use the same mean from the pre-sample, given by that firm's average earnings growth between 1980 and 1985. However, for firms that entered later in the sample, I would use the average earnings growth for the firm's industry using data from 1980 to the quarter prior to the firm's entrance. For all firms, I would set the prior variance to the standard deviation of the earnings growth in the years prior to the firms entrance.

The prior on idiosyncratic risk γ_i would be similar to its current form. I would estimate the variance of a firm's earnings growth between 1980-1985, or from 1980-$t-1$ for entry date t, and fit an inverse Gamma distribution to yield an average variance equal to half of the firm's observed earnings variance. Maintaining the prior that half of a firm's observed risk comes from idiosyncratic risk is a semi-informed prior that permits the data to sway the posterior estimates towards the evidence.

Regarding the new law of motion for β_t, I would set ρ_i to a truncated Gaussian with mean of 0.5 and prior standard deviation 0.4. The bounds would be $[-1, 1]$, though I find it unlikely that any persistence in beta would be negative. The prior distribution 0.5 ± 0.4 would allow the prior to cover relatively high persistence as estimated by Li and Zhang (2017), but also allow my model to find low persistence if there is sufficient evidence to do so.

The long-term beta $\bar{\beta}_i$ would be set to predispose the model towards non-zero cashflow exposure as in the current prior specification. However, the prior mean and variance must be adjusted to match the scale of observed aggregate earnings growth, since the latent factor variance σ^2 is now normalized to one. I suggest an average cross-sectional beta of 15.98, with the assumption that $\bar{\beta}_{i,t}$ explains the total cross-sectional aggregate variance in my sample from 1980-1985. To adjust for heterogeneity in firms, I suggest multiplying 15.98 by the firm's corresponding characteristic (noted in Section 6.4) when the firm first enters the sample. For a large growth firm, this would yield an average $\bar{\beta}_i = 15.98 \times 0.27 = 5.395$. A small value firm would have an average $\bar{\beta}_i = 15.98 \times 1.18 = 23.576$. Firms not in the extreme quintiles would have the corresponding scalar calculated from Campbell and Vuolteenaho (2004). The variance of the cashflow beta should be correspondingly widened – I suggest that, since my model is relatively new, the priors should be widened to allow betas to fall ±50% within one standard-deviation. For a large growth firm, the prior would be $\mathcal{N}(5.395, 2.697^2)$.

The prior on the variance of beta, ς_i^2, should be relatively small compared to aggregate earnings growth. I believe it highly unlikely that "true" cashflow betas should increase or decrease between quarters by more than 5 units with a one-standard-deviation shock. For a firm with a prior beta of 10, for example, a $\varsigma_i = 10$ implies that shocks to the beta process double a firm's exposure to systematic risk. If ρ_i is highly persistent (close to 1), this would result in large and unstable processes for firm betas. I recommend the prior InverseGamma$(5, 100)$ as an example of a fairly wide prior with an average standard deviation of 5. I do not have more informed priors about whether firms of different sizes, value, or industry should have differential priors, in part because Li and Zhang (2017) does not provide estimates of their shock volatility sorted by characteristics. For this parameter, I would prefer to permit the data speak for itself.

Lastly, I would not change the priors on the latent factor priors for δ, ϕ, and latent factor shocks ϵ_t. I would maintain the prior $\delta \sim \mathcal{N}(0, 0.5)$ for the mean. The prior on latent factor persistence would remain a truncated Gaussian $\phi \sim \mathcal{N}(0, 0.2) \in [-1, 1]$. Lastly, the latent factor innovations would only differ in that they are no longer hierarchical, since their standard deviation is normalized to one, i.e. $\epsilon_t \sim \mathcal{N}(0, 1) \forall t$.

The model specified above can address shortcomings in the currently estimated model by providing greater identifiability and allowing for time variation in firm cashflow betas. Further, the subdivision of firm cashflow betas into a time-varying form $\beta_{i,t}$ and the long-run mean $\bar{\beta}_{i,t}$ permits the additional study of the relative parameter uncertainty investors face in the two forms of exposure to cashflow risk.

Chapter 8

CONCLUSION

I propose a model of Bayesian learning through earnings growth and use my model to address several topics of interest to the financial economist. I directly estimate the joint density of an economy where no parameter is known and study how the levels, uncertainty, and time variation in beliefs influences cross-sectional variation in returns as well as firm price responses during earnings seasons.

I find evidence that investors respond to earnings news consistent with rational learning in the presence of parameter uncertainty. I find that firms with higher parameter uncertainty have lower earnings responses. Similarly, I show that firms with earnings that provide precise information about aggregate earnings growth accrue greater earnings announcement premia than do other firms, and have larger stock price reactions to unexpected earnings news. My findings regarding earnings announcements are consistent with existing theoretical and empirical work in rational learning and parameter uncertainty, and suggest that my methodology may have a promising application in directly estimating parameter of interest to researchers studying the effects of learning, uncertainty, and the behavior of rational Bayesian investors.

However, I do not find robust support that my cashflow beta or its parameter uncertainty explains variation in stock returns in non-announcement weeks. I do find a positive risk price associated with the cashflow beta, but only during a firm's announcement week, consistent with a literature that indicates that risk prices are more salient during announcement weeks. My results on parameter uncertainty about the cashflow beta seems to suggest a positive risk price associated with parameter uncertainty, though I highlight a potential nonlinearity in the pricing role of parameter uncertainty.

I provide a discussion on the shortcomings of my model and why it may be the case that I find a cross-sectional risk price associated with my cashflow beta only during announcement weeks. I note that there is a potential identifiability issue in my cashflow betas, since they are not pinned down in levels due to the joint estimation of the latent factor variance. I highlight solutions and alternative model specifications that may allow my cashflow beta to better capture asset cashflow risk, including sequential Monte Carlo estimation, as well as a model that explicitly permits time-variation in cashflow betas.

Lastly, I contribute to the discussion and measurement of bellwether firms by deriving a measure that summarizes a firm's ability to communicate common earnings shocks to the market. I find that these bellwether firms are larger on average and have larger earnings response coefficients. Further, bellwether firms accrue larger earnings announcement premia than uninformative firms. I show that bellwether firms tend to announce earlier than other firms. Curiously, my model makes no requirements that informative firms announce first, so the finding that informative firms usually go earlier is novel. I highlight that additional economic modelling is required to understand the competitive dynamics of firm earnings announcement ordering. My results on informative firms is consistent with rational learning and information choice. My findings imply that the information content of firm payoff signals provides an important source of variation in stock returns during information events such as earnings announcements.

My paper contributes more generally to the Bayesian modelling of fundamental asset risk. I estimate a model and provide guidance for best practices in future models by analyzing the role played by my prior selection and the form of the likelihood. I provide alternative estimation methods and model definitions that permit time variation in cashflow betas to guide future analyses on the effect of learning under parameter uncertainty.

BIBLIOGRAPHY

Bai, Jushan and Serena Ng (2021). "Matrix Completion, Counterfactuals, and Factor Analysis of Missing Data". *Journal of the American Statistical Association* 116.536, pp. 1746–1763.

Ball, Ray, Gil Sadka, and Ronnie Sadka (2009). "Aggregate Earnings and Asset Prices". *Journal of Accounting Research* 47.5, pp. 1097–1133.

Ball, Ray, Gil Sadka, and Ayung Tseng (2021). "Using Accounting Earnings and Aggregate Economic Indicators to Estimate Firm-Level Systematic Risk". *Review of Accounting Studies*.

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad (2005). "Consumption, Dividends, and the Cross Section of Equity Returns". *The Journal of Finance* 60.4, pp. 1639–1672.

Bansal, Ravi and Amir Yaron (2004). "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles". *The Journal of Finance* 59.4, pp. 1481–1509.

Betancourt, Michael (2021). *A Short Review of Ergodicity and Convergence of Markov Chain Monte Carlo Estimators*.

Beyer, Anne and Kevin C. Smith (2021). "Learning about Risk-Factor Exposures from Earnings: Implications for Asset Pricing and Manipulation". *Journal of Accounting and Economics* 72.1, p. 101404.

Bonsall, Samuel B., Zahn Bozanic, and Paul E.. Fischer (2013). "What Do Management Earnings Forecasts Convey About the Macroeconomy?" *Journal of Accounting Research* 51.2, pp. 225–266.

Bontemps, Dominique (2011). "Bernstein–von Mises Theorems for Gaussian Regression with Increasing Number of Regressors". *The Annals of Statistics* 39.5, pp. 2557–2584.

Botshekan, Mahmoud, Roman Kraeussl, and Andre Lucas (2012). "Cash Flow and Discount Rate Risk in Up and Down Markets: What Is Actually Priced?" *Journal of Financial and Quantitative Analysis* 47.6, pp. 1279–1301.

Brav, Alon and J. B. Heaton (2002). "Competing Theories of Financial Anomalies". *The Review of Financial Studies* 15.2, pp. 575–606.

Breeden, Douglas T. (1979). "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities". *Journal of Financial Economics* 7.3, pp. 265–296.

Brown, Nerissa C. and Michael D. Kimbrough (2011). "Intangible Investment and the Importance of Firm-Specific Factors in the Determination of Earnings". *Review of Accounting Studies* 16.3, pp. 539–573.

Campbell, John Y. (1996). "Understanding Risk and Return". *The Journal of Political Economy* 104.2, pp. 298–345.

Campbell, John Y. and John H. Cochrane (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior". *Journal of Political Economy* 107.2, pp. 205–251.

Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho (2010). "Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns". *The Review of Financial Studies* 23.1, pp. 305–344.

Campbell, John Y. and Robert J. Shiller (1988). "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors". *The Review of Financial Studies* 1.3, pp. 195–228.

Campbell, John Y. and Tuomo Vuolteenaho (2004). "Bad Beta, Good Beta". *American Economic Review* 94.5, pp. 1249–1275.

Cattaneo, Matias D. et al. (2020). "Characteristic-Sorted Portfolios: Estimation and Inference". *The Review of Economics and Statistics* 102.3, pp. 531–551.

Chan, Kam Fong and Terry Marsh (2022). "Asset Pricing on Earnings Announcement Days". *Journal of Financial Economics* 144.3, pp. 1022–1042.

Chen, Long and Xinlei Zhao (2009). “Return Decomposition”. *The Review of Financial Studies* 22.12, pp. 5213–5249.

Chopin, Nicolas (2002). “A Sequential Particle Filter Method for Static Models”. *Biometrika* 89.3, pp. 539–552.

— (2004). “Central Limit Theorem for Sequential Monte Carlo Methods and Its Application to Bayesian Inference”. *The Annals of Statistics* 32.6, pp. 2385–2411.

Clarkson, Pete, Jose Guedes, and Rex Thompson (1996). “On the Diversification, Observability, and Measurement of Estimation Risk”. *The Journal of Financial and Quantitative Analysis* 31.1, pp. 69–84.

Cochrane, John H. (1996). “A Cross-Sectional Test of an Investment-Based Asset Pricing Model”. *Journal of Political Economy* 104.3, pp. 572–621.

Cui, Guolong et al. (2016). “Exact Distribution for the Product of Two Correlated Gaussian Random Variables”. *IEEE Signal Processing Letters* 23.11, pp. 1662–1666.

Cvitanić, Jakša et al. (2006). “Dynamic Portfolio Choice with Parameter Uncertainty and the Economic Value of Analysts’ Recommendations”. *The Review of Financial Studies* 19.4, pp. 1113–1156.

Da, Zhi (2009). “Cash Flow, Consumption Risk, and the Cross-section of Stock Returns”. *The Journal of Finance* 64.2, pp. 923–956.

Ellahie, Atif (2021). “Earnings Beta”. *Review of Accounting Studies* 26.1, pp. 81–122.

Fama, Eugene F. and Kenneth R. French (1993). “Common Risk Factors in the Returns on Stocks and Bonds”. *Journal of Financial Economics* 33.1, pp. 3–56.

Fama, Eugene F. and James D. MacBeth (1973). “Risk, Return, and Equilibrium: Empirical Tests”. *Journal of Political Economy* 81.3, pp. 607–636.

Frederickson, James R. and Leon Zolotoy (2015). “Competing Earnings Announcements: Which Announcement Do Investors Process First?” *The Accounting Review* 91.2, pp. 441–462.

Gelman, Andrew and Donald B. Rubin (1992). “Inference from Iterative Simulation Using Multiple Sequences”. *Statistical Science* 7.4.

Gelman, Andrew et al. (2013). *Bayesian Data Analysis*. CRC Press.

Ghysels, Eric (1998). “On Stable Factor Structures in the Pricing of Risk: Do Time-Varying Betas Help or Hurt?” *The Journal of Finance* 53.2, pp. 549–573.

Gomes, Joao, Leonid Kogan, and Lu Zhang (2003). “Equilibrium Cross Section of Returns”. *Journal of Political Economy* 111.4, pp. 693–732.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz (2014). “Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure”. *American Economic Review* 104.8, pp. 2400–2420.

Hameed, Allaudeen et al. (2015). “Information, Analysts, and Stock Return Comovement”. *The Review of Financial Studies* 28.11, pp. 3153–3187.

Handa, Puneet and Scott C. Linn (1993). “Arbitrage Pricing with Estimation Risk”. *The Journal of Financial and Quantitative Analysis* 28.1, pp. 81–100.

Hann, Rebecca N., Heedong Kim, and Yue Zheng (2019). “Intra-Industry Information Transfers: Evidence from Changes in Implied Volatility around Earnings Announcements”. *Review of Accounting Studies* 24.3, pp. 927–971.

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh (2009). “Driven to Distraction: Extraneous Events and Underreaction to Earnings News”. *The Journal of Finance* 64.5, pp. 2289–2325.

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, p. 31.

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* 53.6, pp. 2431–2464.

Jun, Sung Jae and Joris Pinkse (2009). “ADDING REGRESSORS TO OBTAIN EFFICIENCY”. *Econometric Theory* 25.1, pp. 298–301.

Kan, Raymond and Guofu Zhou (2007). “Optimal Portfolio Choice with Parameter Uncertainty”. *Journal of Financial and Quantitative Analysis* 42.3, pp. 621–656.

Kim, Mark P., Spencer R. Pierce, and Ira Yeung (2021). “Why Firms Announce Good News Late: Earnings Management and Financial Reporting Timeliness”. *Contemporary Accounting Research* 38.4, pp. 2691–2722.

Klein, Roger W. and Vijay S. Bawa (1977). “The Effect of Limited Information and Estimation Risk on Optimal Portfolio Diversification”. *Journal of Financial Economics* 5.1, pp. 89–111.

Kothari, S P, Jonathan Lewellen, and Jerold B Warner (2006). “Stock Returns, Aggregate Earnings Surprises, and Behavioral Finance”. *Journal of Financial Economics*, p. 32.

Kothari, S. P. and Jerold B. Warner (1997). “Measuring Long-Horizon Security Price Performance”. *Journal of Financial Economics* 43.3, pp. 301–339.

Kroencke, Tim A. (2017). “Asset Pricing without Garbage”. *The Journal of Finance* 72.1, pp. 47–98.

Kumar, Praveen et al. (2008). “Estimation Risk, Information, and the Conditional CAPM: Theory and Evidence:” *The Review of Financial Studies* 21.3, pp. 1037–1075.

Lettau, Martin and Sydney Ludvigson (2001). “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying”. *Journal of Political Economy* 109.6, pp. 1238–1287.

Li, Jun and Harold H. Zhang (2017). “Short-Run and Long-Run Consumption Risks, Dividend Processes, and Asset Returns”. *The Review of Financial Studies* 30.2, pp. 588–630.

Li, Yong, Jun Yu, and Tao Zeng (2020). “Deviance Information Criterion for Latent Variable Models and Misspecified Models”. *Journal of Econometrics* 216.2, pp. 450–493.

Lucas, Robert E. (1978). “Asset Prices in an Exchange Economy”. *Econometrica* 46.6, pp. 1429–1445.

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

Müller, Ulrich K. (2013). “Risk of Bayesian Inference in Misspecified Models, and the Sandwich Covariance Matrix”. *Econometrica* 81.5, pp. 1805–1849.

Neyman, J. and Elizabeth L. Scott (1948). “Consistent Estimates Based on Partially Consistent Observations”. *Econometrica* 16.1, pp. 1–32.

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* 140.3, pp. 865–893.

Patton, Andrew J. and Allan Timmermann (2010). “Monotonicity in Asset Returns: New Tests with Applications to the Term Structure, the CAPM, and Portfolio Sorts”. *Journal of Financial Economics* 98.3, pp. 605–625.

Patton, Andrew J. and Michela Verardo (2012). “Does Beta Move with News? Firm-Specific Information Flows and Learning about Profitability”. *The Review of Financial Studies* 25.9, pp. 2789–2839.

Rubinstein, Mark (1976). “The Valuation of Uncertain Income Streams and the Pricing of Options”. *The Bell Journal of Economics* 7.2, p. 407.

Savor, Pavel and Mungo Wilson (2016). “Earnings Announcements and Systematic Risk: Earnings Announcements and Systematic Risk”. *The Journal of Finance* 71.1, pp. 83–138.

Savov, Alexi (2011). "Asset Pricing with Garbage". *The Journal of Finance* 66.1, pp. 177–201.

Schmalz, Martin C and Sergey Zhuk (2019). "Revealing Downturns". *The Review of Financial Studies* 32.1, pp. 338–373.

Tipping, Michael E. and Christopher M. Bishop (1999). "Probabilistic Principal Component Analysis". *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61.3, pp. 611–622.

Vehtari, Aki et al. (2021). "Rank-Normalization, Folding, and Localization: An Improved Rˆ for Assessing Convergence of MCMC (with Discussion)". *Bayesian Analysis* 16.2, pp. 667–718.

Xia, Yihong (2001). "Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation". *The Journal of Finance* 56.1, pp. 205–246.

Zhang, Lu (2005). "The Value Premium". *The Journal of Finance* 60.1, pp. 67–103.

Zhang, X. Frank (2006). "Information Uncertainty and Stock Returns". *The Journal of Finance* 61.1, pp. 105–137.

Zhang, Xiao-Jun (2013). "Book-to-Market Ratio and Skewness of Stock Returns". *The Accounting Review* 88.6, pp. 2213–2240.