Risk exposures from risk disclosures: What they said and how they said it

Rahul Mazumder
MIT
rahulmaz@mit.edu

Seth Pruitt
ASU
sjpruitt@asu.edu

Landon J. Ross
Tulane University
lross7@tulane.edu

March 20, 2025

Abstract

We extract information from 10-K risk disclosures: topic models measure *what* is discussed, and context models measure *how* it is discussed. We find that both contain significant predictive information about future aggregate risk exposures, even controlling for (structured) firm characteristics, and that this information is both economically valuable and statistically distinct. We present evidence that management's discussion helps to predict some future corporate actions, which in turn can alter the firm's exposure to aggregate risk.

Keywords: text analysis, conditional asset pricing, factor model, word embedding, group lasso
JEL Codes: C23, C53, G11, G12, G17

1 Introduction

This paper uses forward-looking text from the risk disclosure section of 10-K filings to predict firms' aggregate risk exposure.[1] We use two methods of extracting data from text, and show that they offer distinct predictive information. The first method uses *topic* models—successfully used by recent work in Cong et al. (2020), Bybee et al. (Forthcoming), Lopez-Lira (2023), amongst others—to measure what concepts are discussed. The second method uses text embedding to measure the context in which concepts are discussed—this approach has been less widely used in the finance literature. We find that both context and topics significantly predict future betas—they are complementary information sources. Therefore, agents obtain useful information from firm managers' choice of both *what* to say and *how* to say it.

Topic models measure groups of words chosen and their proportion of existence in one document relative to other documents. They measure *what* is discussed by a document. One might ask of a topic model: How much is an exchange-rate topic discussed? We find that managers' choice of topics helps predict future beta with respect to many tradable and nontradable factors, although notably not to market beta.

Contexts—how a keyword is described by the surrounding text—are measured by text-embedding models. They aim to map words to numeric vectors where operations affect recognizable dimensions of meaning.[2] One might ask of a text-embedding model: How is the discussion of exchange rates in one document different than the discussion in another document? We find that managers' choice of context helps predict future exposure to all the aggregate risks we consider.

[1] From the required Section 1A section presenting management's views on the risks facing the firm.

[2] A classic example is to consider the words *queen*, *king*, *woman*, and *man*. A useful text embedding could be one that maps these to vectors where $queen - woman + man = king$. That is, these vectors would identify that *queen* and *king* are similar in terms of being royal persons, and different in gender as identified by the words *woman* and *man*.

Topics and contexts deliver distinct and important predictive information for aggregate risk exposures. In forecast combination regressions, both topic and context predictive variables are statistically significant much of the time. When not the case, it is the context-based forecast that remains statistically significant. In particular, this is true for investment risk exposure, which is interesting because investment decisions, of all the portfolio-sorting variables we considered, seem the most affected by management's decisions, of which they might have private forward-looking information.

Importantly, the aforementioned results hold while controlling for the predictive content in current beta. In robustness checks we also consider the predictive content of other market- and accounting-based firm characteristics, many of which Kelly et al. (2021) find are significant predictors of future beta, and our results continue to hold.

Ours is the first paper to predict future exposures to familiar aggregate risks using both topic- and context-modeling of 10-K disclosures. The application is important and natural. Item 1A of a firm's 10-K must be a disclosure of the firm's relationship to risks that may affect future performance. Such text is naturally informative, if anything, about the importance of future risks both systematic and idiosyncratic. Our focus on aggregate risk exposure comes from two motivations. First, we contribute to the finance literature's abundant analysis of aggregate risks. Second, many important agents (investors, firm managers, risk managers, etc.) use factor models to make consequential choices. Therefore, we are showing how to extract important information about which both researchers and practitioners care.

The fact that we use *both* topic and context modeling is significant. Gentzkow et al. (2019) note that textual data use in the finance literature is on the rise, alongside the concomitant complication of mapping text to numerical quantities. Topic and context modeling represent two distinct means of such mappings. Our read of the respective literature suggests that the former is becoming common in finance and economics, while the latter is a main focus for computer science applications. Our results give concrete empirical evidence that *both*

methods of extracting data from text are valuable. Moreover, since we find they are often complementary, they should be used jointly by future research.

To aid in our results' interpretation, we present a few additional analysis. First, we estimate whether or not the text-based beta forecasts also *predict* future firm characteristics (e.g. investment), which in turn are *contemporaneously* related to the firm's beta. That is, can our text-based beta forecasts be seen as forecasting future firm conditions of which the managers might be better informed? We find the answer to be yes. Second, we provide examples of 10-K text snippets that most align with the beta forecasts, and show how their natural interpretation is being used to derive predictive content, in order to increase confidence that our method is not an inexplicable black box. Finally, we measure the economic significance of the text in two ways. One, we calculate the information's certainty equivalent for a Bayesian mean-variance investor, finding it is about 28 basis points per annum. Two, we calculate factor-based expected returns and find a nontrivial degree of out-of-sample predictability. Altogether, these results clarify some channels through which managers reveal important information to the market.

Our paper builds on Loughran and McDonald (2011), which shows that context is crucial to effectively using financial documents for economic research. Loughran and McDonald (2011) show variation in a word's context is a potential source of significant measurement error when constructing text-based proxy variables for financial concepts. Much of the subsequent text-analytic literature uses carefully crafted dictionaries whose words are unlikely to occur in unintended contexts to reduce proxy variables' measurement error, e.g., Buehlmaier and Whited (2018), Garcia and Norli (2012), and Hoberg and Maksimovic (2015). We depart from Loughran and McDonald (2011) by viewing variation in a word's context information about firm-level variation in financial concepts.

Our results also contribute to a growing literature using firms' risk disclosures to study accounting, economic, and financial questions. Cohen et al. (2020) shows changes in the

length of firms' risk disclosures predicts changes in firms' future earnings and other financial variables. Campbell et al. (2014) show that firms' risk disclosures accurately reflect pre-disclosure proxy variables for firms' risks using dictionary methods and a topic model. Tetlock et al. (2008) shows that negative fundamentals-related news predicts lower future earnings and returns. Florackis et al. (2022) uses firms' risk disclosures to document a cybersecurity-related risk factor in the cross-section of stock returns. Dyer et al. (2017) uses a topic model to document the effect of SEC- and FASB-related reporting requirements on risk factors' contents, length, specificity, and informativeness. Kim et al. (2024) use large-language models to analyze earnings-call transcripts to predict future stock price volatility.

The plan for the paper is as follows. The next section details the data we use, both structured (already numeric) and unstructured (text requiring mapping to numbers). Section 3 explains how topic and context models distinctly extract data from the text, and the familiar forecasting methods we employ to make our case. After that we report our results demonstrating the significance of text-based information in Section 4. Section 5 provides further understanding of how text is mapping to useful information about future risk exposures. We then conclude.

2 Data

This section reports the data we use for the paper's results. First we describe the stock-market and macroeconomic data used to calculate betas. Then we detail the sample of text taken from Form 10-K filings' Item 1A.

Firms' aggregate-risk exposures are crucial—betas are used for a variety of financial applications. Anyone using the CAPM wants to know market beta: an investor to calculate expected returns, a corporate manager for capital-budgeting decisions, or a risk manager to simplify estimation of stocks' covariance matrix. Anyone instead using a multifactor model

would require betas on the additional portfolio-based factors as well, and the same is true for nontradable factor models.

Firm betas are widely used—and in many applications, the beta we would like to know *pertains to the future*. It is the future aggregate covariance of the firm's return that is important. While this beta is about the future covariance of returns, it is known today in accordance with standard no-arbitrage theory. Were betas static, then this distinction would not matter: but a host of papers (see Ferson and Harvey, 1999; Kelly et al., 2019, amongst others) point to the importance of beta dynamics. Therefore, market participants are motivated to estimate the betas that will govern future returns, and it is this point that guides our main research question.

We measure realized betas using simple regressions. The year t and firm i estimate of $\beta_{i,t}$ with respect to some factor is

$$r_{i,t}^d = \alpha_{i,t} + \beta_{i,t} f_t^d + e_{i,t}^d \tag{1}$$

where $r_{i,t}^d$ is the daily return of stock i on day d in year t and f_t^d is the daily return of some factor on day d in year t. The $\beta_{i,t}$ parameter is the beta of stock i with respect to the given factor during year t. We allow the intercept to freely vary across i and t and its estimate is not used in our analysis.

If agents are rational, then this realized beta is noisy realization of the beta investors expected before year t. The estimated $\beta_{i,t}$ will surely have some sampling error. So long as this is uncorrelated with information before year t, the measurement error satisfies classical assumptions. Therefore, using the $\beta_{i,t}$ estimates as dependent variables (instead of the unobservable true beta) will only reduce the R^2 of the forecasting regressions we run. Nevertheless, we want a reasonably small amount of measurement error in our dependent variables. With this in mind, we restrict ourselves to $\beta_{i,t}$ estimates coming from daily data

with an abundance of available observations. If a firm has less than 230 daily observations in year t, then the firm is omitted from the sample. We start year t on the first business day in January and end on the last business day in December.

Our decision to use daily data provides virtually no limitation on what tradable factors we can consider, since portfolio returns are easily available at the daily frequency. However, there are a host of nontradable factors that are observed only at a far lower frequency. For instance, macroeconomic series like unemployment and inflation are available only monthly, whereas consumption and gross domestic product are available only quarterly. This renders such nontradable factors unsuitable for our analysis, because we'd expect $\beta_{i,t}$ estimates to be unduly noisy due to the paucity of new observations during year t. Nevertheless, we would like to consider *at least some* macroeconomic factors that experience tells us that firm managers consider. The requirement for them to be included is that they have meaningful variation at the daily frequency.

We start by following Fama and French (2015) in choosing nontradable factors. We use daily excess returns on the market (Mkt-RF), the size portfolio (SMB), the value portfolio (HML), the investment portfolio (CMA), and the profitability portfolio (RMW). Of course, a host of other portfolios' returns are available daily. Due to its familiarity, we further consider the momentum factor (Jegadeesh and Titman, 1993). All of these data are available from Ken French's website.

As our three nontradable factors, we choose three macroeconomic series that reasonably affect firm performance and are available at the daily frequency—we download them from the Federal Reserve Economic Database (FRED), whose mnemonics we give. The first is the exchange rate which, given the global nature of commerce and supply chains, can have important effects on firms' performance. Since we focus on U.S. stock market returns, we use the exchange rate of the USD to a broad basket of foreign currencies constructed by the Federal Reserve (DTWEXBGS). The second nontradable factor is the credit spread, which

can have differential effects on the borrowing costs for firms of different creditworthiness. We measure the credit spread using the daily return for Moody's seasoned BAA corporate bond index minus the return for Moody's seasoned AAA corporate bond index (BAA-AAA). Our final macroeconomic factor is the term spread, which can have differential effects on the borrowing costs for firms borrowing from their bank's floating-rate facility versus firms able to issue corporate bonds at various maturities. This is measured by the daily return for the 10-year US treasury constant maturity index minus daily return for the 3-month US treasury constant maturity index (T10Y3M).

The construction of our Item 1A text sample begins with all Form 10Ks filed at the Securities and Exchange Commission (SEC) between the years 2006 and 2022. Our text sample begins in 2006 because this is the first year Item 1A is generally present in Form 10K filings. Within this sample period, there are several factors that determine the specific sample of Item 1A texts we use for the paper's results. First, we only use Item 1A texts from Form 10K filings where we can successfully extract an Item 1A text with high confidence. Second, we only use Item 1A texts when we can successfully link the Item 1A text to a CRSP/Compustat observation using the associated Form 10K's filing date and Central Index Key (CIK). Third, we drop Item 1A texts from the sample whose word counts are above the sample's 99th percentile and below the sample's 1st percentile because these texts contents is typically not the Item 1A portion of a filing. More detail is given in Appendix A.1.

3 Empirical Methods

Using lagged beta to forecast is simple and needs no description. However, our methods of extracting information from 10-K text are much more involved, as has been discussed in Gentzkow et al. (2019) and elsewhere. Our two methods are topic- and context-modeling,

which now we describe.

3.1 Topic model

We use Latent Dirichlet allocation (LDA) to factorize word counts into a collection of latent topics following Blei et al. (2003). This procedure is widely used, and so detailed explanation is relegated to Appendix A.2.

We use the LDA estimates of documents' topic mixtures, i.e. $\theta_{i,t} \in R^{25}$ for firm i and year t, as firm characteristics—the dimension of twenty-five is chosen following Lopez-Lira (2023). Since LDA represents documents as relatively low-dimensional features we use a standard ordinary least squares regression to estimate the relationship between firms' betas and Item 1A text topic mixtures:

$$\beta_{i,t+1} = a + \theta_{i,t}^T \phi + e_{i,t+1}. \tag{2}$$

The left-hand side variable is a firm's beta with respect to some factor. The parameter a is an intercept. The vector $\phi \in R^{25}$ contains the model's slopes. Scalar $e_{i,t+1}$ is an error term.

3.2 Context Model

Our main method of extracting context from 10-K text is to use text embedding methods that are very common in computer-science and statistical applications, and gaining traction in the finance literature. We relegate to Appendix A.3 detailed description of the procedure that finds the average context for each keyword w in a 10-K document. The end result of the context model is $x_{i,t}^w$, the context vector for word w in the 10-K for firm i in year t.

We use a group lasso (glasso) specification to estimate the relationship between our embedding for Item 1A texts and firms' aggregate risk exposures. A few factors make the glasso an appropriate specification. First, our text embedding produces approximately 18,000 firm

characteristics (7 embedding dimensions, multiplied by approximately 2500 keywords) to represent Item 1A texts' information about future returns. Since the text embedding is high-dimension, we consider some manner of regularization natural to ensure our empirical estimates plausibly generalize to new data. Second, we are equally interested in producing empirically useful descriptions of firms' aggregate risk exposures and interpreting our empirical estimates of firms' aggregate risk exposures. Since one of our aims is interpretation, we consider sparsity an appropriate form of regularization because sparse empirical solutions are relatively amenable to inspection. Third, the natural unit of sparsity for our text embedding is at the level of the context vector for one word, which contains more than one variable. So, group sparsity, where each group contains the context vector for one word, is the appropriate form of sparsity to consider.

The linear model associated with the glasso specification is

$$\beta_{i,t+1} = a + \sum_{w} x_{i,t}^{w^T} \phi_w + e_{i,t+1} \tag{3}$$

where a is an intercept, $x_{i,t}^{w}$ is the context vector for word w, $\phi_w \in R^7$ is the slopes for components of the word w context vector, and $e_{i,t+1}$ is an error term. We estimate the linear model's slopes via the regularized least squares problem

$$\hat{a}, \hat{\phi}_1, \ldots, \hat{\phi}_W = \underset{a,\phi_1,\ldots,\phi_W}{\arg\min} \ \frac{1}{2} \sum_{i,t} \left(a + \sum_{w} x_{i,t}^{w^T} \phi_w - \beta_{i,t+1} \right)^2 + \lambda \sum_{w} \|\phi_w\|_2. \tag{4}$$

The minimization problem's first term is the standard least squares term. The problem's second term is a glasso penalty which shrinks each context vector's slopes towards zero. The λ is a scaling parameter controlling the magnitude of the glasso penalty. We use five-fold cross-validation and a grid search to estimate this parameter. Note that the penalization term sums across words w, and each summand is the sum-of-squares. Yuan and Lin (2006) show

that this produces two regularization effects. Within a chosen word, the context slopes ϕ_w are L_2-regularized, as in ridge regression. Across words, we get selection as in L_0-regularization and L_1-regularization. Hence, we get sparsity in the word (w) dimension, as desired for interpretability, as well as shrinkage for (hopefully) better out-of-sample performance.

3.3 Out-of-sample analysis

We take seriously the potential for overfitting. The large dimension of our empirical analysis stems from the large dimensionality inherent in text information (see Gentzkow et al., 2019). Therefore, an important concern is the possibility of overfitting when estimating the text-embedding mapping and the beta-forecasting parameters. We have detailed how those parameters are estimated by a variety of dimension-reduction techniques, and these constitute our first safeguard against overfitting. Our second, important safeguard is sample-splitting. We estimate those text-related parameters using only the first half of the sample, 2006–2014. Then we evaluate their performance only on the second half of the sample, 2015–2022. Therefore our main results are out-of-sample, which is crucial to delivering conclusions that are not driven by overfitting.

4 Results

We start by showing the out-of-sample predictive performance of forecasts based on past betas, topics, or contexts. Text information, of both types, are significant predictors of future risk exposure. These conclusions are robust to controlling for a wealth of structured firm-characteristic data.

4.1 Out-of-sample performance

Table 1 reports out-of-sample R^2 of forecasting realized beta, using forecasts errors from the 2015–2022 sample. That is, the statistic is constructed as

$$1 - \left(\sum_{2015-2022} (\beta_{i,t} - b_{i,t-1})^2 \right) / \left(\sum_{2015-2022} \left(\beta_{i,t} - \overline{b}_i \right)^2 \right).$$

We have abused notation and use "2015-2022" to denote the set of firm-year observations in the 2015–2022. In the numerator we forecast $\beta_{i,t}$ using a prediction $b_{i,t-1}$ based on year $t-1$ information—the forecast comes from year $t-1$ beta, topics, or contexts. In the denominator, the R^2 uses as its benchmark $\overline{b}_i$ the firm's beta on the 2006–2014 sample, or if the firm does not exist during 2006–2014 we use the average across firms. This benchmark forecast seems imminently reasonable, as it is simply the long-run average beta the agent saw during 2006–2014 (or else the average across firms). Note that this benchmark allows each firm i to have a different benchmark forecast, somewhat analogous to if one constructed a panel R^2 by using forecast errors from firm fixed effects as the benchmark for the denominator.[3] We prefer this statistic as it embodies the idea that different firms have different betas, and so our R^2 is positive only if the predictions $b_{i,t-1}$ improve upon this basic fact.

Our first important result is that only contexts provide positive out-of-sample R^2 for market beta. The 7.2% is considerably strong—for context, Kelly et al. (2021) found an *in-sample* R^2 of 25.6% using thirty-six firm characteristics over a longer period. Strikingly, lagged beta gives negative forecast performance, implying it is usually a worse predictor of market beta than the longer-run average beta for that firm. Putting these observations together, we see evidence that market beta is dynamic, but its own past value is too noisy

[3]If we instead take every firm's benchmark forecast as simply the average beta in 2006–2014 across firms, this serves to increase all the R^2s relative to what we report—hence our statistic is a conservative approach. Additionally, for the market beta it could make sense to take a value of 1 as the benchmark for firms born after 2015; but for other factors there is no such sensible assumption. Therefore we use the average across firms for all factor betas.

Table 1: Out-of-sample performance

Notes– Reports the out-of-sample R^2 of forecasting realized beta, using forecasts errors from the 2015–2022 sample. All parameters are estimated on the 2006–2014 sample. The realized beta for each is estimated from a univariate regression of daily stock returns on daily factor realizations, as described in the text. The R^2 uses as its benchmark the firm's beta on the 2006-2014 sample, and if the firm does not exist during that period we use the average across firms.

Factor	Lagged Beta	Topic	Context
Panel A: Portfolios			
Market	−0.054	−0.034	0.072
Size	0.156	0.016	0.082
Value	0.294	0.151	0.159
Investment	0.140	0.292	0.320
Profitability	0.599	0.103	0.104
Momentum	−0.584	0.230	0.236
Panel B: Macroeconomic			
Exchange Rate	−0.169	0.112	0.098
Credit Spread	−0.304	0.015	0.028
Term Spread	−0.165	0.006	0.038

to provide a good forecast. On the other hand, management's risk disclosures can provide valuable risk-exposure information that backward-looking market information does not.

Topics provide negative forecast performance for market beta. This could follow because management's information about the firm's exposure to the market, an aggregation of everything that is going in all firms, is not related to their choice of topics. We see a great variety of topics discussed in 10-Ks (we will see this below), but a particular filing's choice of topic tells us nothing about their future market beta. Yet, *the context of how* management describes topical keywords does reveal insight into market beta. This is a first glance at our second important result, that topic and context models reveal complementary and distinct information.

While only contexts are useful in forecasting market beta, there are several other aggregate risks for which both lagged betas and topics also perform well. Continuing with

the portfolio factors, we see that lagged beta delivers a solid 15.6% for size exposure. The choice of topics delivers a modest 1.6%, indicating that management's choice of topics is only slightly connected to their future size beta. Meanwhile, the contexts chosen by management deliver 8.2% of the variation in future size beta, meaning that their disclosures do provide useful information to investors if extracted properly.

When forecasting value beta, lagged beta is once again a good forecaster, delivering R^2 of 29.4%. For value risk, both topic and context information are very similarly useful and deliver R^2s of 15.1% and 15.9%, respectively. Qualitatively similar, lagged beta provides a very strong prediction of future profitability beta with an out-of-sample R^2 of 60%, while topics and contexts are moderately useful with R^2s of 10.3-10.4%. Hence, value and profitability exposures are predicted quite well by their own recent past values, while text-based information is less useful.

Meanwhile investment risk has a flipped result: text-based information is *more* informative of future exposure than its own lag. Whereas 14% of the out-of-sample variation in investment beta is explained by its lagged value, text-based information explains more than twice as much: topic-based forecasts deliver a R^2 of 29% and context-based forecasts deliver 32%.

Finally, momentum risk provides a stark contrast between backward-looking-market and forward-looking-text information. Its lagged value provides extremely negative forecasting power for momentum beta, with an out-of-sample R^2 of -58%. Nonetheless, there is significant information about future momentum beta coming from 10-K risk exposures. Both topic and context information deliver an out-of-sample R^2 of around 23-24%. The contrast here highlights the novel richness of forward-looking text-based information extracted by both topic and context models. Daniel and Moskowitz (2016) show that momentum strategies have experienced significant crashes, including during our 2006-2014 training sample. Kelly et al. (2021) show that other firm characteristics are useful for identifying large shifts in a

firm's momentum exposures and avoiding crashes, and here we show a similar story for 10-K text information.

It is noteworthy that management's disclosures provide the most predictive information for exposure to investment risk. Of all the portfolios considered, this factor has the most to do with management's own choices. The investment factor is constructed as a spread portfolio between firms with conservative (low) and aggressive (high) investment ratios. Fama and French (2015) measure this using the firm's asset growth rate, which is exactly the type of financial ratio management's decisions can directly affect. Our results suggest that firm managers have forward-looking information about their future investment decisions, that they reveal it in 10-K disclosures, and it has significant power for understanding their future investment-risk exposure. This fact has clear value for investors viewing investment risk as an important source of returns' first and second conditional moments.

Turning to momentum, we see very poor performance in lagged beta. As mentioned, this is in comparison to a benchmark of the firm's average beta. It turns out that these average momentum betas are not very disperse amongst stocks—over a longer term, firms have a similar exposure to momentum, in keeping with that portfolio's great deal of turnover. Hence, the benchmark for all stocks is relatively similar, and in part it is the excessive cross-sectional variation of lagged momentum betas that contributes to the negative R^2. In turn, this excessive volatility does not forecast anything, and therefore must come from unforecastable shocks. One might have taken this to suggest that expected momentum betas are not time-varying—but the text-based forecasts argue against this. There is year $t-1$ information that usefully predicts future momentum exposure: that is, its conditional expectation is dynamic. Management's choice of topics and their context reveal that this forward-looking information exists.

Panel B of Table 1 reports results for nontradable factors. Across the board, lagged beta provides negative forecasting power. Meanwhile, text-based forecasts both provide

moderately strong forecast performance for exchange-rate beta, and more mild performance for credit-spread and term-spread betas. Overall, the out-of-sample R^2 for macroeconomic factors are lower than for portfolio factors, consistent with the well-known observation that nontradable factors can be somewhat weak (e.g. Giglio and Xiu, 2021). Nevertheless, forward-looking text information uniformly provides better predictions than backward-looking market information.

4.2 Marginal forecast significance

To investigate whether or not the information in lagged betas and text is complementary, now we estimate forecast combinations. If one forecast is statistically significant while controlling for other forecasts, this is evidence that the former provides important predictive content. Since we are running these regressions over the 2015–2022 sample, the next set of results involve an in-sample aspect. But note that the construction of the forecasts—in particular, the text-extraction parameters—continue to be estimated on the 2006–2014 sample, to protect against possible overfitting. Moreover, the regressions we run have at most three regressors and hence there are no high-dimensionality concerns. The R^2s reported in Table 2 are not directly comparable to the out-of-sample R^2s reported in Table 1.

Panel A of Table 2 reports the market-beta results for several combinations of the context, lagged-beta, and topic forecasts. The first three columns essentially provide Mincer-Zarnowitz regressions of each forecast. For the joint null hypothesis that the intercept equals 0 and the slope equals 1, only the topic-based forecasts fail to reject. The context-based forecasts reject the Mincer-Zarnowitz null because the slope is 0.90 and significantly different than 1, not because the intercept is far from 0.

Meanwhile, the lagged-beta forecasts reject the Mincer-Zarnowitz null because the slope is 0.54 and the intercept is 0.51, and both are far from the hypothesized values. In fact,

Table 2: Marginal forecast significance

Notes– From regressions over 2015–2022 of future beta on out-of-sample forecasts coming from contexts, lagged betas, or topics. Standard errors are clustered by firm, and t-statistics in parentheses. We denote 5% significance by *, 1% significance by **, and 0.1% significance by ***. The number of observations for each panel is more than 10,300.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Panel A: Market			
Context	0.903***			0.527***	0.781***		0.454***
	(34.03)			(23.26)	(23.90)		(17.93)
Lagged Beta		0.542***		0.410***		0.474***	0.406***
		(41.85)		(28.72)		(34.82)	(28.24)
Topic			1.011***		0.284***	0.532***	0.179***
			(21.76)		(5.34)	(17.00)	(5.12)
Intercept	0.0464	0.507***	-0.0852	0.0301	-0.147**	-0.0506	-0.0915**
	(1.47)	(31.25)	(-1.54)	(1.45)	(-2.98)	(-1.61)	(-2.89)
R^2	0.217	0.286	0.126	0.343	0.223	0.316	0.345
				Panel B: Size			
Context	0.830***			0.532***	0.701***		0.450***
	(37.82)			(28.85)	(25.08)		(20.77)
Lagged Beta		0.517***		0.367***		0.437***	0.360***
		(47.15)		(33.66)		(39.52)	(32.79)
Topic			1.049***		0.339***	0.609***	0.231***
			(26.73)		(7.49)	(22.31)	(7.20)
Intercept	-0.181***	0.481***	-0.481***	-0.111***	-0.470***	-0.281***	-0.309***
	(-5.68)	(33.37)	(-8.68)	(-4.90)	(-9.75)	(-8.16)	(-9.04)
R^2	0.212	0.236	0.127	0.304	0.220	0.274	0.307
				Panel C: Value			
Context	0.973***			0.696***	0.563***		0.403***
	(48.48)			(39.27)	(14.30)		(12.27)
Lagged Beta		0.468***		0.272***		0.280***	0.258***
		(52.26)		(30.23)		(30.93)	(28.85)
Topic			1.169***		0.570***	0.829***	0.428***
			(46.37)		(12.48)	(39.69)	(11.41)
Intercept	-0.626***	0.0902***	-0.777***	-0.433***	-0.750***	-0.537***	-0.537***
	(-34.04)	(10.63)	(-35.28)	(-28.21)	(-36.48)	(-29.85)	(-30.92)
R^2	0.266	0.223	0.260	0.320	0.281	0.318	0.328

continued on next page

Table 2 – continued from previous

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Panel D: Investment			
Context	0.752***			0.424***	0.812***		0.471***
	(17.62)			(13.21)	(12.55)		(11.28)
Lagged Beta		0.439***		0.381***		0.409***	0.381***
		(41.21)		(32.40)		(32.71)	(32.98)
Topic			0.821***		-0.112	0.426***	-0.0882
			(8.38)		(-0.84)	(6.47)	(-1.02)
Intercept	-0.146***	-0.172***	-0.120***	-0.0485***	-0.165***	-0.0318*	-0.0638***
	(-10.81)	(-15.41)	(-4.59)	(-4.96)	(-6.44)	(-2.03)	(-3.92)
R^2	0.102	0.219	0.056	0.247	0.102	0.233	0.247
				Panel E: Profitability			
Context	0.631***			0.301***	0.525***		0.254***
	(21.98)			(18.80)	(14.55)		(11.40)
Lagged Beta		0.511***		0.463***		0.480***	0.462***
		(48.19)		(46.23)		(47.50)	(46.33)
Topic			0.763***		0.222***	0.347***	0.101***
			(18.39)		(4.68)	(16.59)	(3.52)
Intercept	0.729***	-0.115***	0.964***	0.392***	0.934***	0.483***	0.487***
	(16.32)	(-14.52)	(14.37)	(15.45)	(14.96)	(13.90)	(14.23)
R^2	0.096	0.284	0.065	0.303	0.099	0.296	0.304
				Panel F: Momentum			
Context	0.620***			0.671***	0.558***		0.598***
	(13.07)			(14.31)	(6.03)		(6.37)
Lagged Beta		-0.0235*		-0.0506***		-0.0453***	-0.0509***
		(-2.18)		(-4.94)		(-4.32)	(-4.95)
Topic			0.531***		0.0718	0.572***	0.0853
			(12.32)		(0.86)	(13.42)	(1.00)
Intercept	-0.170***	-0.168***	-0.170***	-0.170***	-0.170***	-0.170***	-0.170***
	(-23.91)	(-21.89)	(-23.63)	(-22.93)	(-23.94)	(-22.73)	(-22.95)
R^2	0.016	0.001	0.012	0.019	0.016	0.015	0.019

continued on next page

Table 2 – continued from previous

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Panel G: Exchange Rate			
Context	0.595***			0.413***	0.398***		0.277***
	(23.27)			(19.99)	(9.81)		(8.69)
Lagged Beta		0.337***		0.293***		0.298***	0.290***
		(37.99)		(34.42)		(34.06)	(33.75)
Topic			0.734***		0.328***	0.506***	0.230***
			(22.28)		(6.72)	(19.20)	(5.86)
Intercept	-0.00205***	-0.00757***	-0.000116	-0.00222***	-0.000230	-0.000889*	-0.000947**
	(-5.88)	(-74.01)	(-0.26)	(-8.14)	(-0.55)	(-2.57)	(-2.82)
R^2	0.064	0.119	0.058	0.148	0.069	0.145	0.150
				Panel H: Credit Spread			
Context	1.316***			0.991***	0.802***		0.594***
	(23.51)			(20.92)	(9.98)		(8.82)
Lagged Beta		0.292***		0.234***		0.237***	0.223***
		(28.71)		(25.36)		(24.74)	(24.20)
Topic			2.084***		1.242***	1.590***	0.996***
			(27.39)		(10.98)	(24.39)	(10.34)
Intercept	0.00199	0.0416***	-0.0322***	0.00125	-0.0304***	-0.0256***	-0.0247***
	(0.79)	(46.43)	(-9.73)	(0.60)	(-9.62)	(-9.39)	(-9.25)
R^2	0.059	0.070	0.059	0.101	0.071	0.102	0.108
				Panel I: Term Spread			
Context	1.116***			0.659***	0.812***		0.466***
	(26.44)			(19.46)	(12.76)		(10.26)
Lagged Beta		0.459***		0.371***		0.392***	0.361***
		(36.05)		(28.20)		(27.66)	(26.58)
Topic			1.578***		0.722***	0.925***	0.485***
			(25.45)		(7.85)	(18.31)	(7.16)
Intercept	-0.0133***	0.0349***	-0.0503***	-0.00995***	-0.0468***	-0.0332***	-0.0325***
	(-3.97)	(26.27)	(-10.44)	(-4.24)	(-10.60)	(-9.78)	(-9.86)
R^2	0.101	0.157	0.081	0.187	0.111	0.182	0.191

these qualities of the lagged-beta forecast led the out-of-sample R^2 to be negative in Table 1.

Based on the Mincer-Zarnowitz test we'd say that the topic-based forecasts are systematically unbiased—but that null hypothesis is not the last word on economic significance. Looking at the R^2s, though the topic-based forecasts are unbiased they also are the least powerful: the context-based and lagged-beta forecasts provide about double the predictive content, proving their value.

Column 4 shows that context remains strongly significant (at the 0.1% level) even when controlling for lagged beta. The context slope drops from 0.90 to 0.53 and the lagged-beta slope drops from 0.54 to 0.41. Hence, context appears to contain some backward-looking information that is captured in lagged beta—controlling for the latter suggests that about 60% ($\approx 0.53/0.90$) of the context information is forward-looking. Column 6 reports a similar finding for topic. Column 5 includes just the two text-based forecasts, and both remain significant, with the topic slope falling by about 75% but the context slope falling by only 15%. Therefore we find two important results: first, text-based forecasts provide valuable forward-looking information even while controlling for backward-looking market information; second, topic and context are complementary to each other.

The final column reports that all three forecasts are statistically significant when evaluated together—the clearest sign yet that topics, contexts, and lagged betas provide distinct predictive content. Since we are regressing onto fitted values, it is somewhat natural that the slope coefficients for the eventual forecast (what we'll call the fitted value in column 7) almost sum up to 1. Roughly speaking, about two-fifths of the eventual forecast comes from both context and lagged betas, and one-fifth from topic.[4] The R^2 in column 7 shows that these distinct information sources explain 34.5% of the variation in market betas, an

[4]If anything this is generous to topic, as the context and lagged-beta R^2s are similar implying their fitted values are about equally variable. A formal variance decomposition would need to take this into account—we opt to instead speak in rough terms.

economically significant amount as we'll later quantify.

Panel B shows the results for size betas. Once again, all three forecasts enter the combination significantly and the R^2 is an economically-significant 30.7%. As for market beta, from the slope coefficients we see that context is the biggest part of the eventual forecast, with lagged beta and topic following. Panel C shows that all three forecasts are significant, with the two text-based forecasts providing similar amounts to the eventual forecast whose R^2 is 32.8%. Jumping to profitability, panel E column 7 shows a similar R^2 of 30.4%, but this time lagged beta is responsible for most of the eventual forecast, with context and topic coming next in that order.

As highlighted in the previous section, results for investment and momentum betas show differences. For investment betas, panel D column 4 shows us that topic information is basically nested in context information. That is, in the bivariate regression we see that the context is strongly significant while topic is not. While both text-based forecasts are significant while controlling for backward-looking lagged-beta information, we see evidence that context is extracting the relevant predictive information for future investment beta. This shines through in column 7 where context and lagged beta are strongly significant, but topic is not.

Meanwhile, momentum betas provide more nuanced differences. Notably, while lagged beta is significant in Panel F column 2, it enters with a negative sign—this holds true even in column 7. Meanwhile, topic is significant on its own (column 3) but is driven out by context (columns 5 and 7). Notably, the context slopes barely change between columns 1, 5, and 7. This indicates that context uniquely contains the 10-K's forward-looking predictive information, and it is virtually uncorrelated with lagged beta's backward-looking information.

Panels G–I of Table 2 show that lagged beta and the text-based forecasts are all statistically significant. For these macroeconomic risks, topic and context are uniformly com-

plementary. A modest amount of their informativeness is correlated with lagged beta, as the slope coefficients vary with the latter's inclusion. But they consistently are important contributors to the eventual forecast in column 7, where R^2s range respectably from 11% to 19%. Therefore, text-based information is useful for predicting future nontradable-factor risk exposures.

A broad takeaway from Table 2 is that context and lagged beta are always significant, complementary predictors of future risk exposures.[5] Most often, but not always, topic is significant as well. For most factor risks, backward-looking market information is enhanced by both types of forward-looking text information we extracted.

4.3 Firm characteristic information

So far we have used lagged beta to capture structured, non-textual information that agents have available to forecast future risk exposure. Yet Kelly et al. (2021) gives us reason to believe that other firm characteristics may also be valuable.

Table 3 shows that our main results continue to hold with respect to this larger structured information source. Comparing to column 7 of Table 2, we see the same pattern of topic and context statistical significance. In fact, in numerous cases the context and topic coefficient estimates are negligibly changed.

4.4 Statistical selection of information

We have explored two types of structured information (lagged beta and firm characteristics) and two types of textual information (topic and context). We have found that these information sources provide forecasts that are often complementary to one another, using

[5]We have also calculated standard errors clustered by year—while this serves to lower the t-statistics nearly every qualitative conclusion remains. But we note that issues with small numbers of clusters (we have 7 when clustering by year) are well known.

Table 3: Marginal forecast significance, robustness to firm-characteristic forecasts

Notes– From regressions over 2015–2022 of future beta on out-of-sample forecasts coming from contexts, firm characteristics, or topics. Standard errors are clustered by firm, and t-statistics in parentheses. We denote 5% significance by *, 1% significance by **, and 0.1% significance by ***. The number of observations for each column is more than 10,300.

	Market	Size	Value	Investment	Profitability	Momentum
Context	0.592***	0.491***	0.389***	0.694***	0.331***	0.354***
	(18.84)	(17.41)	(10.13)	(11.05)	(9.42)	(4.46)
Characteristic	0.534***	0.563***	0.535***	0.526***	0.478***	1.192***
	(19.27)	(14.84)	(12.78)	(4.89)	(12.49)	(24.32)
Topic	0.248***	0.294***	0.518***	-0.230	0.195***	-0.225**
	(5.52)	(7.43)	(11.95)	(-1.75)	(4.49)	(-3.03)
Intercept	-0.514***	-0.892***	-0.989***	-0.0581	1.392***	-0.173***
	(-11.07)	(-17.25)	(-37.27)	(-1.73)	(19.51)	(-28.66)
R^2	0.286	0.292	0.322	0.137	0.139	0.121
	Exchange Rate	Credit Spread	Term Spread			
Context	0.331***	0.594***	0.651***			
	(8.37)	(7.23)	(9.86)			
Characteristic	0.320***	1.197***	0.701***			
	(9.33)	(14.14)	(10.88)			
Topic	0.282***	1.143***	0.685***			
	(5.96)	(10.65)	(7.73)			
Intercept	0.00264***	-0.0701***	-0.0871***			
	(5.10)	(-17.11)	(-15.92)			
R^2	0.080	0.095	0.132			

forecast-combination regressions. We now provide another means of understanding the value of these different information sets. We use the glasso with varying levels of regularization and ask: what type of information is chosen? If we allow the glasso to only select 10 information groups, which are the ones it picks up? How does this change as we allow more groups to be selected? Answering these questions tell us which types of information are selected by purely statistical methods.

The striking conclusion from Figure 1 is that context is the earliest and most prevalent information statistically selected. While it is true there is only one lagged beta variable to be chosen, there are dozens of topics and firm characteristics that could be chosen. Uniformly, context variables constitute the majority of groups selected for every factor and every total number of active groups. Visually this is apparent by the prevalence of blue in Figure 1, denoting the number of context groups chosen.



Figure 1: Information selected by glasso

Notes– How many groups of each information type are chosen. The types are context vectors, firm characteristics, lagged beta, and LDA topics. Going horizontally is the number of groups allowed to be active, from 10 to 100 by 10. The vertical axis measures how many groups of each type are chosen, as the height of the bar of each color.

Let us consider in detail the top-right panel for market beta. When ten groups are allowed active, seven of them are context variables—two firm characteristics and the lagged betas compose the remaining three. Only when forty groups are allowed active does a single topic get selected—no more are ever chosen. By time one hundred groups are allowed active, only five firm characteristics are included, alongside lagged beta and the lone topic. The remaining 93 active groups are context variables. For every factor considered, context information comprises more than ninety percent of what the glasso eventually selects.

Scanning across the other panels, it is only the topic, firm characteristic, and lagged beta selection that really varies across factor betas. Topic groups are never selected for size, investment, profitability, and momentum betas. For value and investment betas, the first ten groups selected are purely context.

Turning to nontradable factors in the last row, the prominence of context information is unchanged. Exchange-rate betas never choose topic information, term-spread betas choose a single topic, and credit-spread betas choose two topics. Despite there being thirty-six firm characteristics available, never more than handful are selected.

There are a few ways to interpret Figure 1 in light of results we've seen before. Recall that forecasts based on topics and firm characteristics are very often statistically significant in forecast-combination regressions. Since the glasso chooses so few of them, it is natural to suppose that those valuable topic and firm-characteristics forecasts were driven by a small number of topics and firm characteristics, respectively. Meanwhile, it is perhaps the case that context variables have a more complicated relationship to one another, such that a greater number of them are consistently chosen as we loosen the regularization. One could say that the value of topic and firm-characteristic information tends to rest in just a few variables. At the same time, the broad pattern of selection indicates that when glasso can choose a new variable to include in beta forecasting, it almost always chooses context information. This suggests that context information is quite rich.

5 Interpretation

In this section we further explore the information provided by 10-K text—for the sake of exposition, we focus on the information extracted by the context model. First, we present simple reduced-form regressions which indicate how 10-K contexts are forward-looking via managerial decisions or business-condition predictions. Second, then illustrate particular pieces of 10-K text that represent this forward-looking information. Finally, we calculate two measures of economic significance: a certainty equivalent for a Bayesian mean-variance investor, and out-of-sample predictability results for individual stocks.

5.1 Context and future firm characteristics

We do not believe that management simply *writes* into existence its future aggregate beta. Instead, our view is that management's risk disclosures are revealing their information about future managerial decisions or business conditions, and that *these* are signicant drivers of the realized covariance the firm experiences. With this in mind, we run simple reduced-form regressions to suggest why the text information is predictive. For simplicity and since it was statistically significant for all them, we focus on context information and the five Fama and French (2015) factor betas; as above, we consider firm characteristics from Kelly et al. (2021).

To be specific, the two types of regressions we run are referred to as *contemporaneous* and *predictive* and can be expressed as follows. Let $\hat{\beta}^f_{t|t-1}$ be the context-based forecast of factor f β^f_t using context at $t-1$ and c^f_t be a firm characteristic related to factor f (measured as the average over the year). Then the *contemporaneous* regression is

$$\beta^f_t = a_0 + a_1 c^f_t$$

Table 4: Context forecasts and future firm characteristics

Notes– All regressions have more than 6,200 observations and employ standard errors clustered by firm. In the *Predictive* Context column we report the point estimate and t-statistic of forecasting the firm characteristic named in that row using the context-based beta forecast for the factor whose column in the *Contemporaneous* panel is not blank. The *Contemporaneous* panel reports regressions of the year t realized beta of the factor in the column label (e.g. Value) on the year t characteristic (e.g. Book/Market). All predictive regressions are univariate; the Profitability contemporaneous regression is bivariate while the remaining are univariate. Standard errors are clusted by firm.

	Predictive	*Contemporaneous*				
	Context	Market	Size	Value	Investment	Profitability
Book Leverage	0.294	0.003				
	(2.21)	(4.93)				
Earnings/Assets	−0.059		−1.959			
	(−8.12)		(−15.09)			
Book/Market	0.392			0.335		
	(15.74)			(3.11)		
Asset growth	−0.095				−0.374	
	(−8.90)				(−2.33)	
Profitability	1.910					0.0002
	(1.67)					(2.76)
Earnings/Assets	0.065					3.576
	(11.05)					(15.52)
R^2		0.002	0.107	0.092	0.029	0.131

and the *predictive* regression is

$$c_t^f = a_2 + a_3 \hat{\beta}_{t|t-1}^f.$$

The first regression tells us if the characteristic is related to realized beta contemporaneously. The second regression tells us if the context-based forecast of future beta also predicts the future characteristic value.

According to classic corporate theory, a firm increases its measured (levered) market beta by increasing its leverage. Perhaps this is one of the few financial-managerial decisions that drives the CAPM. Table 4's first row provides supporting evidence—first look at the column labeled Market. There we see that book leverage during year t is significantly and

positively related to the realized market beta during year t, with a t-statistic of 4.93, in a *contemporaneous* regression. Hence, a firm's increased leverage explains a part of the increase in market beta. And the first column of that row (labeled Context) shows that context is *predictive* of future leverage. From managerial risk disclosures we get forward-looking information of the firm's future indebtedness, which of course carries significant effects for firms.

Market capitalization is not something managers can directly control, yet it should (at some point) be connected to firm performance. As a measure of performance, earnings over assets gives a sense of how productively the firm is employing its capital, and this productivity should result in firm growth. The second row shows a strongly negative relationship between between earnings/assets and size beta ($t = -15.09$), which is consistent with this idea because size exposure is smaller for bigger firms. Meanwhile, context strongly predicts future earnings/assets ($t = -8.12$). Hence, size exposure is somewhat explained by its contemporaneous asset productivity, which in turn is predicted by risk-disclosure text.

The story for value and investment exposures is more direct. The book-to-market ratio explains value beta ($t = 3.11$) and asset growth explains investment beta ($t = -2.33$) with the expected signs. And the context of firm disclosures significantly forecast the future book-to-market ($t = 15.74$) and asset growth ($t = -8.90$). In the case of the former, we think this has more to do with management's inside information about the firm's future valuation, since they cannot directly control stock prices. In the case of the latter, context may reveal management's investment plans which manifests as as asset growth.

Context does not forecast future (gross) profitability as strongly as the other characteristics we've considered ($t = 1.67$), though profitability is contemporaneously related to profitability-risk exposure ($t = 2.76$) as one expects. Because of this we considered earnings/assets as well, because it too is a firm productivity measure and it worked significantly

in the size-exposure case.[6] We see that context robustly forecasts earnings/assets ($t = 11.05$) which itself is contemporaneously related to profitability exposure ($t = 15.52$).[7] Therefore risk disclosures are more informative of future bottom-line profit measures, and these relate strongly to aggregate risk exposure.

The straightforward story presented here is that 10-K disclosures are forward looking because managers have useful information about future business conditions or future corporate plans. Future research might explore related questions, such as: which type of information is more precise? (we might expect the latter); or, which type is more useful to investors? (does a firm's business-conditions forecast spillover to information about related firms?); or, does management have incentives to reveal some plans but not others?

5.2 Text driving beta predictability

How has the context model extracted predictive data? To provide intelligible examples, we take two steps—details are relegated to Appendix A.4 and here we provide a simple summary. First, we identify which words' contexts are most important to the forecast of beta. This is done by considering the glasso solution path, and arranging words by the order of their inclusion in the model. Second, we map actual 10-K text snippets whose embedding most closely aligns to the direction given by the word's predictive coefficients $\hat{f}_w$, using cosine similarity. For the sake of exposition, we focus on investment betas because Table 1 shows that text information is particularly predictive thereof.

The texts indicate that the firm is negatively exposed to the investment factor, as the cosine similarity scores are close to -1. Recall that CMA is short firms *aggressively* investing, and we can see that the text relate to such cases. *Patents* is the most important word for

[6] Hence the contemporaneous regression for profitability beta is run with both profitability and earnings/assets as regressors.

[7] Context predicts earnings/assets differently in the size and profitability cases—the estimates are -0.059 and 0.065, respectively. This is because the forecast of size beta is a different linear combination of context variables than the forecast of profitability beta.

Table 5: Text driving investment betas

Notes– The top seven words chosen by the glasso model are listed in the first column, in descending order. We select text from 10-Ks whose magnitude of cosine similarity with $\hat{f}_w$ is large. The cosine similarity measure is reported in the second column, and the third column gives the text snippet.

Word	Cosine	Text
patents	−0.995	We do not currently have any issued patents in any foreign jurisdictions that cover our most advanced product candidates.
technologies	−0.994	These technologies can be expected to require refinements and there is the risk that advanced new technologies will be introduced.
delays	−0.982	This could interfere with product sales or cause interruptions of or delays in our research and development programs.
volatile	−0.990	The price of aircraft fuel is unpredictable and has been increasingly volatile over the past few years.
damages	−0.991	Defects and errors in our software products could result in claims for substantial damages against us.
equipment	−0.991	The risk of a lack of available suppliers or equipment may be heightened as a result of market and economic conditions.
engineering	−0.989	Our manufacturing operations focus on highly engineered components which require extensive engineering and research and development time.

investment beta forecasts, and the similar text is from a firm identifying a risk that their most advanced product candidates may not have global protection. *Technologies* follows in importance, and the text depicts a firm with technological advancements that could be displaced (see Kazemi, 2022, for more on displacement risk and its relation to intangible investment). Further down the list of importance, *equipment* aligns with a firm discussing the risk of being unable to invest in equipment due to macroeconomic conditions. All of these pieces of text describe firms that are tied to aggressive investment, either as producers or consumers of investment goods. Hence, we get a broad idea of how forward-looking text in a firm's 10-K labels its future activities as tied to aggregate investing conditions.

5.3 Economic significance

We provide two measures of the economic significance of context-based beta forecasts. The first quantifies the predictive information's importance to a Bayesian investor. The second calculates factor-based expected returns and reports their out-of-sample accuracy.

5.3.1 Bayesian mean-variance investor

One way to measure the economic significance of risk-disclosure information is by calculating a certainty equivalent, answering the question "What value do investors ascribe to this text data?" A modestly challenging aspect to this question is that the reduction in uncertainty occurs for *parameters* employed by the investor. That is, a typical certainty-equivalent calculation takes a risky payoff and calculates the riskless payoff yielding the same utility. But in our setup it is uncertainty *about the perceived risk* that is at play. Therefore, we adopt a Bayesian perspective, to jointly capture both the uncertainty about the payoff (unaffected by the text information) and the uncertainty about the model parameters (what the text information *does* affect).

To ease exposition, we put the details in Appendix A.5 and here give an overview of the problem. Agents have mean-variance preferences and use an APT-consistent factor model, à la Stambaugh (1983), to form expectations of future returns—our twist is that the agent is Bayesian and uncertain of the β parameters determining conditional expected returns. Employing conjugate priors, the typical mean-variance solution obtains, albeit with the agent viewing returns as given by a multivariate non-central t-distribution. Intuitively, the perceived variance of returns is increasing in the variance of investors' prior on β—hence Bayesian mean-variance investors dislike parameter uncertainty.

Let $MV^*(rf, \tilde{\kappa})$ denote the agent's optimized utility as a function of the risk-free rate rf

and hyperparameter $\tilde{\kappa}$. The certainty equivalent calculation solves for s in

$$MV^*(rf, \tilde{\kappa}_1) = MV^*(rf + s, \tilde{\kappa}_2), \quad \text{with } \tilde{\kappa}_1 > \tilde{\kappa}_2, s \geq 0, \tag{5}$$

where $\tilde{\kappa}_1, \tilde{\kappa}_2$ govern prior-distribution precision (of the beta parameters) in the case of excluding text information ($\tilde{\kappa}_2$) or including text information ($\tilde{\kappa}_1$). Then s is the non-negative increase in the risk-free rate that compensates the investor for excluding text information. That is, s is the Bayesian certainty-equivalent value of that information.

Using Fama and French (2015) data for our sample period, a moderately high risk aversion of $\gamma = 10$, and simplifying assumptions detailed in the appendix, we solve (5) for s. We find that the certainty equivalent is 28 basis points per annum. With the average risk-free rate around 14 basis points per annum, this says that investors require a risk-free rate twice as high to compensate for losing valuable information.[8] Even when simply judging the certainty equivalent at face value, we view 28 basis points as a reasonable yet economically-significant value of the text information we've extracted.

5.3.2 Stock predictability

Our main focus is the prediction of systematic risk exposures, which is useful for a variety of applications including risk management and asset allocation. In the standard mean-variance framework, the value of accurate loading predictions flows through more accurate assessments of returns' covariance matrix and expected returns, and the last section shows that this accuracy is economically valuable. We now focus solely on the expected return.

Under no-arbitrage assumptions, $\mathbb{E}_t(r_{i,t+1}) = \beta_{i,t}\lambda_t$; so we need only add an estimate of the factor's risk premium to produce an expected return from our beta forecast. As tradable

[8] Of course, risk-free rates at the end of our sample are much higher than this average. We have found that a higher risk-free rate serves to increase the certainty equivalent calculation, all else equal, so believe our 28 basis point conclusion is conservative.

Table 6: Stock return predictability

Notes– Out-of-sample R^2 (in percentage) of stock-return forecasts using factor-based expected returns, using either lagged betas or context-forecasted betas.

	Lagged Betas	Context
Market	−0.75	0.23
Size	−2.84	0.12
Value	−4.29	0.58
Profitability	0.83	0.42
Investment	−2.83	0.90
Momentum	−7.23	−0.18

factors' conditional mean is not our object of study, we take as our simplest possible estimate their average value over the training sample.[9]

Table 6 shows that text-based beta forecasts lead to out-of-sample return predictability. Lagged betas do not provide useful return predictions, except for profitability-factor-based expected returns, which is natural given the good profitability beta results in Table 1. Meanwhile, text-based betas lead to useful return predictions for all factors, except momentum. Consistent with text's greater accuracy in forecasting investment betas, the investment-factor-based expected returns are most accurate and achieve a 0.9% out-of-sample R^2. Broadly speaking, the takeaway is that text helps to predict firms' future exposure to aggregate risks, which in turn can help forecast future returns.

6 Conclusion

This paper uses the risk disclosure text of 10-Ks to predict future aggregate risk exposures. We find notable success, even controlling for structured (firm characterstic) data that agents already have in hand. As an accurate measure of beta is valuable to any agent using a factor

[9]As a benchmark for the out-of-sample R^2, we use each stock's average beta's over the training sample.

model to understand expected returns, cost of capital, or covariance—this shows that text provides significant information for actual investors. Topics and contexts provide significant information that is distinct, and economically valuable according to two metrics we provide. We show that these forecasts also predict future firm conditions, which are themselves contemporaneously related to aggregate risk exposure, providing evidence that managers' private information about future business conditions are revealed in the disclosures.

References

Blei, David M, Andrew Y Ng, and Michael I Jordan (2003) "Latent dirichlet allocation," *Journal of machine Learning research*, Vol. 3, No. Jan, pp. 993–1022.

Buehlmaier, Matthias MM and Toni M Whited (2018) "Are financial constraints priced? Evidence from textual analysis," *The Review of Financial Studies*, Vol. 31, No. 7, pp. 2693–2728.

Bybee, Leland, Bryan Kelly, Asaf Manela, and Dacheng Xiu (Forthcoming) "Business News and Business Cycles," *Journal of Finance*.

Campbell, John L, Hsinchun Chen, Dan S Dhaliwal, Hsin-min Lu, and Logan B Steele (2014) "The information content of mandatory risk factor disclosures in corporate filings," *Review of Accounting Studies*, Vol. 19, pp. 396–455.

Cohen, Lauren, Christopher Malloy, and Quoc Nguyen (2020) "Lazy prices," *The Journal of Finance*, Vol. 75, No. 3, pp. 1371–1415.

Cong, William, Tangyuan Liang, and Xiao Zhang (2020) "Textual Factors: A Scalable, Interpretable, and Data-driven Approach to Analyzing Unstructured Information,"Technical report, Cornell University.

Daniel, Kent and Tobias J. Moskowitz (2016) "Momentum crashes," *Journal of Financial Economics*, Vol. 122, No. 2, pp. 221–247.

Dyer, Travis, Mark Lang, and Lorien Stice-Lawrence (2017) "The evolution of 10-K textual disclosure: Evidence from Latent Dirichlet Allocation," *Journal of Accounting and Economics*, Vol. 64, No. 2-3, pp. 221–245.

Fama, Eugene F and Kenneth R French (2015) "A five-factor asset pricing model," *Journal of financial economics*, Vol. 116, No. 1, pp. 1–22.

Ferson, Wayne E. and Campbell Harvey (1999) "Conditioning Variables and the Cross Section of Stock Returns," *Journal of Finance*, Vol. 54, No. 4, pp. 1325–1360.

Florackis, Chris, Christodoulos Louca, Roni Michaely, and Michael Weber (2022) "Cybersecurity Risk," *The Review of Financial Studies*, Vol. 36, No. 1, pp. 351–407, URL: `https://doi.org/10.1093/rfs/hhac024`, DOI: 10.1093/rfs/hhac024.

Garcia, Diego and Øyvind Norli (2012) "Geographic dispersion and stock returns," *Journal of Financial Economics*, Vol. 106, No. 3, pp. 547–565.

Gavish, Matan and David L Donoho (2014) "The optimal hard threshold for singular values is 4/*sqrt*{3}," *IEEE Transactions on Information Theory*, Vol. 60, No. 8, pp. 5040–5053.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019) "Text as Data," *Journal of Economic Literature*, Vol. 57, No. 3, pp. 535–574.

Giglio, Stefano and Dacheng Xiu (2021) "Asset Pricing with Omitted Factors," *Journal of Political Economy*, Vol. 129, No. 7, pp. 1947 – 1990.

Hoberg, Gerard and Vojislav Maksimovic (2015) "Redefining financial constraints: A text-based analysis," *The Review of Financial Studies*, Vol. 28, No. 5, pp. 1312–1352.

Hoffman, Matthew, Francis Bach, and David Blei (2010) "Online learning for latent dirichlet allocation," *advances in neural information processing systems*, Vol. 23.

Jegadeesh, Narasimhan and Sheridan Titman (1993) "Returns to buying winners and selling losers: Implications for stock market efficiency," *The Journal of finance*, Vol. 48, No. 1, pp. 65–91.

Kazemi, Maziar (2022) "Intangible Investment, Displacement Risk, and the Value Discount,"Technical report, Arizona State University.

Kelly, Bryan T, Tobias J Moskowitz, and Seth Pruitt (2021) "Understanding momentum and reversal," *Journal of financial economics*, Vol. 140, No. 3, pp. 726–743.

Kelly, Bryan T., Seth Pruitt, and Yinan Su (2019) "Characteristics are covariances: A unified model of risk and return," *Journal of Financial Economics*, Vol. 134, No. 3, pp. 501–524, URL: `https://www.sciencedirect.com/science/article/pii/S0304405X19301151`, DOI: https://doi.org/10.1016/j.jfineco.2019.05.001.

Khodak, Mikhail, Nikunj Saunshi, Yingyu Liang, Tengyu Ma, Brandon M Stewart, and Sanjeev Arora (2018) "A La Carte Embedding: Cheap but Effective Induction of Semantic Feature Vectors," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12–22.

Kim, Alex, Maximilian Muhn, and Valeri V. Nikolaev (2024) "From Transcripts to Insights: Uncovering Corporate Risks Using Generative AI,"Technical report, University of Chicago.

Lopez-Lira, Alejandro (2023) "Risk Factors that Matter: Textual Analysis of Risk Disclosures for the Cross-Section of Returns,"Technical report, University of Florida.

Loughran, Tim and Bill McDonald (2011) "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks," *The Journal of finance*, Vol. 66, No. 1, pp. 35–65.

Mikolov, Tomas, Edouard Grave, Piotr Bojanowski, Christian Puhrsch, and Armand Joulin (2018) "Advances in Pre-Training Distributed Word Representations," in *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.

Stambaugh, Robert F. (1983) "Arbitrage pricing with information," *Journal of Financial Economics*, Vol. 12, No. 3, pp. 357–369, URL: https://ideas.repec.org/a/eee/jfinec/v12y1983i3p357-369.html.

Tetlock, Paul C, Maytal Saar-Tsechansky, and Sofus Macskassy (2008) "More than words: Quantifying language to measure firms' fundamentals," *The journal of finance*, Vol. 63, No. 3, pp. 1437–1467.

Yuan, Ming and Yi Lin (2006) "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Vol. 68, No. 1, pp. 49–67.

A Appendix

A.1 Text data

This appendix summarizes our procedure for extracting and processing Item 1A sections from Form 10-K documents. We first discard plain text 10-K filings from the sample. We discard plain text filings because they make a negligible fraction of the sample, occur only in the first few years of the paper's sample, and have difficult to identity Item 1A sections. Next, we use two programs to extract Item 1A from the sample of HTML 10-K documents.

The first program we use to extract Item 1A sections from Form 10-K filings has three subroutines. The first subroutine identifies the table of contents at the beginning of a Form 10-K. The routine finds a document's table of contents by searching for anchor tags within the document's HTML whose text is similar to the phrase "table of contents." The routine then locates the HTML element with an id attribute matching the id given by the document's links to its table of contents. The second subroutine parses a Form 10K document's table of contents into individual sections. The subroutine first finds the beginning of the Item 1A section by searching for a section title similar to "Item 1A. Risk Factors." Then, the subroutine finds the section immediately after the document's risk factors section. Next, the subroutine finds links to the beginning of each section's contents within both sections' table of contents entries. The third subroutine extracts the Item 1A text from the Form 10-K document using the table of contents links to the Item 1A section and the section immediately after the Item 1A section. The subroutine extracts all of the text between the link to the Item 1A section and the link to the next section.

The second program we use to extract Item 1A sections from Form 10-K filings has three subroutines. The first subroutine directly searches the text of a Form 10-K section for phrases similar to "Item 1A" at the beginning of an HTML tag. When the subroutine finds multiple matches for this condition, the subroutine assumes the last match is the beginning of the Item 1A section. This assumption works well because practically all other uses of the phrase "Item 1A" within a Form 10-K occur in the Management Discussion and Analysis portion of the Form 10-K, which is earlier in a Form 10-K than the risk factors section. The second subroutine searches for the occurrence of the phrase "Item" at the beginning of an HTML tag after the start of the Item 1A section's text. This second subroutine usually matches several different positions in the text. We assume the first match is the conclusion of the Item 1A section and the beginning of the filing's next section. The third subroutine extracts the Item 1A section from the filing using the section beginning and section conclusion identified by the previous two routines. The third routine also discards sections with less than ten words. The purpose of this last condition is to discard Item 1A sections where the text is "Not required" or another similar phrase.

We perform additional processing after extracting Item 1A sections from Form 10-K documents. The purpose of this further processing is to standardize the documents' words. We first remove all HTML-related tags from the texts. We also remove anchor tags' text because practically all links in risk factors sections are table of contents links. We convert

Unicode characters to their closest ASCII equivalents. We remove page numbers. Last, we remove punctuation and other non-alphanumeric characters.

A.2 Topic model

Latent Dirichlet allocation (LDA) is a Bayesian approach to factoring discrete data, here word counts, into a collection of latent topics (Blei et al., 2003). LDA assumes the following specification for generating a document's words from K latent topics:

1. Choose $N \sim \text{Poisson}(\xi)$.
2. Choose $\theta \sim \text{Dirichlet}(\alpha)$.
3. For each of the N words w_n in the document:
   (a) Choose a topic $z_n \sim \text{Multinomial}(\theta)$.
   (b) Choose a word w_n from multinomial distribution $p(w_n|z_n, B)$.

Variable N represents the number of words present in a document and is drawn from a Poisson distribution with parameter ξ. Vector $\theta \in R^K$ represents the distribution of K latent topics that generate a document's words. Vector θ is drawn from a Dirichlet distribution with parameter $\alpha \in R^K$. The expression z_n is a categorical variable representing the topic generating the word at position n in a document. The variable z_n is drawn from a multinomial distribution whose parameter θ controls the topic mixture for the document. Last, word w_n at position n in the document is drawn from a topic z_n specific multinomial distribution over all of the words occurring in the document sample. The parameter B is a vector specifying the probability each word in the sample is generated by topic z_n. The primary user-chosen parameter for LDA is the number of latent documents generating the text sample. We use an LDA specification with 25 latent topics because Lopez-Lira (2023) finds LDA factorizations of Item 1A texts with 25 topics have greater topic coherence than LDA factorizations with other numbers of topics. We use an online variational Bayes algorithm from Hoffman et al. (2010) to estimate the LDA specification's three unknown parameters: the parameter α controlling the distribution of topics within the Item 1A sample, the parameter θ controlling the distribution of topics within each document, and the parameter B controlling each topic's distribution over words present in the Item 1A sample.

A.3 Context model

Given the possible lack of familiarity to the reader, we start with simplified example of context. Suppose companies A and P both mention the term *exchange*. Company A says, "The Company uses derivative instruments, such as foreign currency forward and option contracts, to hedge certain exposures to fluctuations in foreign currency exchange rates." On the other hand, company P says, "Fluctuations in exchange rates, including as a result of currency controls or other currency exchange restrictions have had, and may continue to have, an adverse impact on our business, financial condition and results of operations." Both companies mention exchange-rate risk, but using different contexts: company A states it is actively hedging its exposure, while company P discusses adverse impacts. One might

interpret the different contexts as implying (their respective management teams believe) company A has *less* future exchange-rate exposure than company P.

There is empirical support for this interpretation, it turns out. The above companies are Apple and PepsiCo, and the sentences are from the text of their respective FY2019 10-K's Section 1A. In those filings, each company mentions "exchange rate" five times—at a simplified level, the prevalence of the exchange-rate *topic* is similar in both documents.[10] During 2019 both Apple's and PepsiCo's exchange-rate betas are insignificant (in fact, current exchange-rate beta does not predict future beta, according to results below). Thereafter during the turbulence of the COVID-19 pandemic, PepsiCo's realized beta becomes strongly significant ($t = 3.8$) while Apple's remains insignificant. While current beta and the exchange-rate topic *are not* predictive, our simple interpretation of the exchange-rate *context is* predictive. Our text-embedding approach essentially systematizes the illustration into a machine-learning method that is suitable for analyzing big data.

We represent the context of word w in the year t Item 1A text for firm i with the vector $x_{i,t}^{w} \in R^{L}$. The definition of the context vector $x_{i,t}^{w}$ is

$$x_{i,t}^{w} = A \frac{1}{|C_{i,t}^{w}|} \sum_{c_{i,t}^{j} \in C_{i,t}^{w}} \sum_{w_{i,t}^{k} \in c_{i,t}^{j}} v^{w_{i,t}^{k}}. \tag{A.1}$$

Let's unpack the context vector's construction from right to left. The right-most variable is $v^{w_{i,t}^{k}} \in R^{M}$. Vector $v^{w_{i,t}^{k}}$ is a pre-trained vector for the word $w_{i,t}^{k}$ at position k of the Item 1A text for firm i in year t. The next symbol is $c_{i,t}^{j}$, which represents the context of the word at position j in the Item 1A text for firm i and year t. The context of word $w_{i,t}^{j}$ is the set $c_{i,t}^{j}$ of words within the K-word neighborhood of word j, i.e.

$$c_{i,t}^{j} = \{w^{k} \in W_{i,t} : |k - j| \leq K\} \tag{A.2}$$

where $|k - j|$ is the distance between words $w_{i,t}^{k}$ and $w_{i,t}^{j}$ within a document and K controls the context's width. We are now able to interpret the construction's right-most summation expression;

$$\sum_{w_{i,t}^{k} \in c_{i,t}^{j}} v^{w_{i,t}^{k}} \tag{A.3}$$

sums pre-trained word vectors for words in the context of word w at position j in the Item 1A text for firm i in year t.

The same word often occurs at several different positions within a document. In this case, a word will be associated with more than one context within a document. We use the expression $C_{i,t}^{w}$ to represent the set of all contexts for word w within the Item 1A text for

[10] It turns out that topics are useful to forecasting exchange rate beta, on average. But these two documents have a similar predictive weight on informative topics—in other words, from a topic-model perspective these two observations should forecast the same. Meanwhile, the context model successfully classifies these two statements differently.

firm i and year t with the definition

$$C_{i,t}^{w} = \{c_{i,t}^{j} : w_{i,t}^{j} \in W_{i,t}, w_{i,t}^{j} = w\} \tag{A.4}$$

where the $w_{i,t}^{j} \in W_{i,t}$ restricts the set's contents only to contexts within the Item 1A text for firm i and year t and condition $w_{i,t}^{j} = w$ restricts the set's contents to contexts with keyword w. Additionally, we use the expression $|C_{i,t}^{w}|$ to represent the number of contexts for word w in an Item 1A text. We can now interpret more of the context vector's construction. The expression

$$\frac{1}{|C_{i,t}^{w}|} \sum_{c_{i,t}^{j} \in C_{i,t}^{w}} \sum_{w_{i,t}^{k} \in c_{i,t}^{j}} v^{w_{i,t}^{k}} \tag{A.5}$$

is the average vector representation for the context of word w in the Item 1A text for firm i and year t.

The last undefined expression in the context vectors' construction is the matrix $A \in R^{L \times M}$. The matrix A serves two purposes. The first purpose of A is to adapt the pre-trained word vectors to better represent words meanings' within Item 1A texts. Pretrained word vectors are estimated using very large text samples taken from a variety of domains. The matrix A tilts the pre-trained vectors' meanings away from their original multiple-domain setting and towards the paper's finance-specific domain. The second purpose of matrix A is to dimension reduction. The standard dimension of pretrained word vectors is 300 (Mikolov et al., 2018). The matrix A maps high-dimensional pre-trained word vectors into lower-dimensional vectors for the paper's specific task.

The paper's context vector construction requires specifying several details for the paper's results. We use the fasttext word vectors from Mikolov et al. (2018) as pre-trained word vectors. We use a context window size of $K = 20$. We estimate the matrix A using a modified form of the specification from Khodak et al. (2018), which we describe in appendix A.3.1. We report results for context vectors with seven dimensions. We determine seven components are sufficient for the context vectors using a test from Gavish and Donoho (2014), which we also describe in appendix A.3.1. We estimate context vectors for all nouns that occur in at least 5% of the Item 1A texts and that are not stop words. We restrict our sample of contexts to nouns because risks are typically nouns.

A.3.1 A Matrix Estimation

This section describes our estimate of the matrix A from equation (A.1). We use a methodology similar to the one proposed in Khodak et al. (2018) with an additional step that reduces the context vectors' dimension. The matrix A is the solution to the multivariate least squares regression problem

$$V = UA + E \tag{A.6}$$

$$V \in R^{W \times K} \quad U \in R^{W \times D} \quad A \in R^{D \times K} \quad E \in R^{W \times D}. \tag{A.7}$$

The matrix V contains vector representations for the W words in the paper's vocabulary taken from the FastText embedding after a dimension reduction step. Let $\tilde{V} \in R^{W \times D}$ be the matrix where each row is the FastText vector representation of some word occurring in the Item 1A sample. Let the singular value decomposition of the matrix $\tilde{V}$ be

$$\tilde{V} = FGH^T \tag{A.8}$$

where we use F, G, H to represent the singular value decomposition of $\tilde{V}$ to avoid collisions with the paper's other notation. The matrix V is the seven largest left singular vectors of $\tilde{V}$ scaled by the seven largest singular values of $\tilde{V}$:

$$V = \sum_{j=1,\ldots,7} F_{:,j} G_{j,j}. \tag{A.9}$$

We used cross-validation to select the number of left singular vectors we use to construct V.

The matrix U contains context vectors for the W words in the paper's vocabulary. Row w of the matrix U is the vector $u_w \in R^D$ with the definition

$$u_w = \frac{1}{|C^w|} \sum_{c^j_{i,t} \in C^w} \sum_{w^k_{i,t}} v^{w^k_{i,t}}. \tag{A.10}$$

The set C^w contains the context for every occurrence of the word w in our sample of Item 1A texts; i.e., $C^w = \cup_{i,t} C^w_{i,t}$. Note that C^w does not contain duplicate contexts because C^w is a set. So, each context occurs only once in the set C^w even though the context may belong to several $C^w_{i,t}$ sets.

A.4 Text driving beta predictions

A.4.1 Context selection order

Which words' contexts are most important for explaining variation in firms' future factor betas? We answer this question by studying the solution path for the glasso regressions of betas onto word contexts. The glasso solution path traces the contexts assigned nonzero coefficients as we vary the glasso penalty. The early portion of the solution path has a large penalty and assigns few words' contexts nonzero coefficients. The glasso model gives contexts nonzero coefficients along the solution path when their inclusion produces the greatest reduction in residuals relative to the remaining contexts with zero coefficients. Thus, contexts given nonzero coefficients early on the solution path explain more variation in betas than contexts activated later and are natural candidates for the contexts most important for explaining future betas.

Table A.1 reports the order word contexts receive nonzero coefficients along the glasso solution path for the Market, Size, and Value factor regressions. The "Context" columns report the words with contexts assigned nonzero coefficients early on the solution path. The "Entry Order" columns report the order contexts are assigned nonzero coefficients. E.g.,

the context in the first row of table A.1 for the market factor is "retail." The entry order for "retail" is one. This means the context for the word "retail" is the first context assigned nonzero coefficients by the glasso model for firm's market betas. The subsequent rows present the second, third, fourth, etc. contexts chosen by the glasso model for explaining firms' future betas. Occasionally, two or more contexts have the same entry order. This occurs when the penalty parameter interval over which the contexts' coefficients are sequentially assigned nonzero values is relatively small and difficult to numerically locate. Table A.2 reports solution path results for the investment, profitability, and momentum factors. Table A.3 reports solution path results for the exchange rate, credit spread, and term structure factors.

Table A.1: Context Glasso Solution Path

Notes– Reports the order contexts are assigned non-zero coefficients along the glasso solution path for the Market, Size, and Value factors.

Market		Size		Value	
Context	Entry Order	Context	Entry Order	Context	Entry Order
retail	1	shares	1	loans	1
volatile	2	ratings	2	intellectual	2
borrowings	3	billion	3	deposits	3
covenants	4	million	4	banks	4
lenders	5	directors	5	collateral	5
industries	6	common	6	values	5
losses	7	retail	6	license	6
domestic	8	covenants	7	patent	7
loans	9	executive	8	loan	8
finance	10	quarter	9	acceptance	9
care	11	investors	10	mortgage	10
downturn	12	senior	11	allowance	11
net	13	pension	12	borrowings	12
distribution	14	volume	13	distribution	13
warranty	15	stockholder	14	borrowers	14
quarter	16	finance	15	research	15
China	17	rating	16	Reserve	16
raise	18	natural	17	licenses	17
utility	19	warranty	17	Drug	17
reserves	20	raise	17	deterioration	18
bank	21	distribution	17	asset	18
properties	21	allowance	18	savings	19
steel	22	key	19	downturn	19
certificate	23	downturn	20	options	19
stockholder	23	care	21	care	20
investors	23	Stock	22	lenders	20
agencies	24	volatile	23	applications	20
production	25	commodities	24	labeling	21
holders	26	profitability	25	expansion	22
economy	26	transmission	25	infringement	23

Table A.2: Context Glasso Solution Path

Notes– Reports the order contexts are assigned non-zero coefficients along the glasso solution path for the Investment, Profitability, and Momentum factors.

Investment		Profitability		Momentum	
Context	Entry Order	Context	Entry Order	Context	Entry Order
patents	1	equity	1	loans	1
technologies	2	securities	1	intellectual	2
delays	3	shares	2	patents	3
volatile	4	retail	3	banks	4
damages	5	loans	4	deposits	5
equipment	6	raise	5	acceptance	6
engineering	7	investors	6	license	7
countries	7	million	7	values	8
find	8	directors	8	patent	9
shares	9	weather	9	borrowers	10
stockholders	9	loan	10	research	11
contract	10	holders	11	distribution	12
issuance	11	losses	12	equivalents	13
certificate	12	allowance	13	Drug	13
initial	13	volatile	14	March	14
quantities	14	Bank	15	collateral	15
officers	15	fuel	16	allowance	16
rapid	16	Corporation	17	loan	17
pension	17	deposits	18	greenhouse	18
exploration	18	convertible	19	issuance	19
encounter	19	industries	20	shares	20
analysts	20	fund	21	convertible	20
savings	21	finance	22	approval	20
ratings	22	pension	23	analysts	21
royalty	23	components	24	suppliers	22
complex	23	certificate	25	credit	23
life	24	distribution	26	borrowings	23
rating	25	warranty	27	labeling	24
stockholder	26	utility	28	licenses	24
retail	27	quarter	28	sensitive	25

Table A.3: Context Glasso Solution Path

Notes– Reports the order contexts are assigned non-zero coefficients along the glasso solution path for the Exchange Rate, Credit Spread, and Term Spread factors.

Exchange Rates		Credit Spread		Term Spread	
Context	Entry Order	Context	Entry Order	Context	Entry Order
production	1	retail	1	retail	1
gas	2	utility	2	borrowings	2
covenants	3	care	3	components	3
equipment	3	Bank	4	domestic	4
exploration	4	domestic	5	care	5
properties	5	bank	6	quarter	6
volatile	6	components	6	utility	7
retail	7	borrowings	7	industries	8
dollar	8	loan	8	shares	9
currency	9	quarter	9	China	9
facility	9	pension	10	default	10
joint	9	currency	11	volatile	11
reserves	10	weather	12	lenders	12
engineering	11	economy	13	economy	13
steel	12	distribution	14	distribution	14
industries	13	China	15	allowance	15
indebtedness	14	volatile	15	bank	16
default	15	allowance	16	currency	16
distribution	16	default	17	utilities	17
domestic	17	currencies	18	Bank	18
drilling	18	fuel	18	downturn	19
lenders	19	industries	19	approval	20
countries	19	utilities	20	warranty	21
currencies	20	valuation	21	pension	22
air	20	Health	22	currencies	23
accidents	21	transmission	23	goodwill	23
care	22	approval	24	transmission	23
publicity	22	shares	25	Health	23
financing	23	degree	26	agencies	24
marketing	23	deterioration	27	gross	25

Overall, the early contexts selected by the factors are consistent with economic intuition about the variables prior literature finds relevant to firms' covariances and expected returns. Almost all of the glasso models select the retail context relatively early, which is consistent with the prominent role of consumption in theories of asset prices. Contexts relevant to firm leverage, such as borrowings, covenants, lenders, loan, banks, ratings, default, credit, and indebtedness, frequently occur across the solution paths. Contexts relevant to equity financing, governance, and corporate control also appear: distribution, certificate, stockholder, shares, director, common, executive, senior, officers, equity, and issuance. Macroeconomic subjects are frequent across all solution paths: domestic, downturn, China, production, steel, economy, commodities, countries, weather, fuel, greenhouse, dollar, gas, dollar, currency. Financial intermediation relevant terminology like bank, collateral, deposits, lending, and mortgage appear for several factors. Many of the remaining terms are relevant to describing tangible and intangible assets: intellectual, license, patent, research, infringement, technologies, engineering, equipment, properties, reserves, and goodwill. Broadly speaking, these words contexts are plausible sources of information about firms' covariances with the economy's significant risks.

Focusing on the Market factor, early contexts such as "retail," "finance," and "downturn" underscore the significance of consumption, financial conditions, and general economic conditions for market risk. For the size factor, contexts of "billion" and "million" are natural for distinguishing large and small firms. Other terms like "directors," "executive," denote corporate governance risks often associated with small firms, and "ratings," "covenants," and "investors" are consistent with financing related explanations for risk premia associated with smaller firms. The presence of several research-related terms early in the value and investment path is consistent with previous literature noting small research-intensive firms are typically growth stocks with elevated levels of investment. For the exchange rate path, most terms have a natural correspondence with international trade and foreign currency. Contexts selected for the credit spread model broadly capture information about firms' credit risk. Contexts selected for the term spread contain information about the effects of broad economic movements on firms, e.g. "retail," "borrowings," "domestic," "economy," and "downturn."

A.4.2 Context sentences

Next, we look at the information present in individual contexts about firms' future betas. Since we use embeddings that approximately preserve meaning under linear operations, we interpret a context's coefficients as a direction in the context's embedding where variation in meaning aligns with variation in betas. If an observation is close to the coefficients' line in the embedding space, then much of the observation's meaning is relevant to future betas, and the observation is a natural candidate for making sense of the coefficients' meaning. Likewise, observations far from the coefficients' line through the embedding space contain little information about betas, and make poor candidates for making sense of variation in contexts' meaning aligned with betas.

We need to be careful, however, with the notion of proximity we use to judge observations' distance from the line formed by a context's slopes. We use cosine similarity between a context's observations and slopes because it orders observations by the proportion of their meaning shared with a context's slopes. We do not use dot products or other projection concepts to measure distance when interpreting context coefficients because these concepts depend on the observations' norms, which, while relevant to measuring magnitudes, would place observations with large norms but minimal alignment close to a context's line through the embedding.

Table A.4 reports sentences with the greatest absolute cosine similarity to the first ten contexts slopes from the glasso model for market factor betas. To simplify the sentences' presentation and interpretation, the sentences shown in table A.4 are taken from the subset of risk disclosure sentences with one independent clause and less than twenty words. We also report only sentences from the in-sample period in table A.4 so that we are certain the sentences' content was available during model estimation. We report similar results for sentences from the out-of-sample period in table yada yada. Cosine calculations for all sentence, context, and factor triples are available on the authors' websites. We use absolute cosine similarity to identify sentences most aligned with the contexts' slopes. However, we report the sentence's cosine similarity, not absolute cosine similarity, in table A.4 so that the results also report the direction of the sentences' effects on betas.

The sentences in table A.4 generally have a clear economic relationship with firm betas. The sentences aligned with the "borrowing" context's slopes describe standard risks associated with debt-related financing. The second sentence, in particular, connects firm risk back to aggregate interest rates, i.e. an increase in aggregate interest rates would increase the firm's market beta. The sentences for the "convenants" context also describe standard risks and constraints associated with debt and credit financing. Both sentences within the "convenants" context indicate the marginal effect of covenants on firm betas is positive. The "domestic" context's sentences ties firms' betas to their reported exposures to domestic and global economic conditions. Both of the "domestic" sentences would have a positive marginal effect on a firm's beta. Like the previous financing related contexts, the "finance" context indicates constraints on firm's financing opportunities have a positive effect on the market betas. The "industries" context relates firms' dependence on specific industries' business conditions to increases in market betas. The "lenders" context sentences indicate financing constraints have a positive effect on firms' market betas. Interestingly, unlike the previous financing-related contexts, the sentences for the "loans" context describe firm revenue associated with issuing loans with Loan-writing business operations increasing a firm's market beta. Sentences aligned with the "lossess" context indicate reporting scenarios where a firm incurs losses have a positive effect on market betas. The sentences for the "retail" context indicate general exposure to retail operations, after accounting for the other contexts, reduces firms' market betas. Last, the "volatile" context sentences indicate general price volatility associated with a firm's physical and financial inputs have a positive effect on firm's market beta.

Table A.4: Cosine Similarity Between Group Slopes and Sentences for Market Factor

Notes– Reports sentences with greatest absolute cosine similarity to slopes for first ten word contexts selected by glasso model of firms' market factor betas.

Context	Cosine	Sentence
borrowings	1.00	Any additional borrowings would further increase the amount of our indebtedness and the associated risks.
	0.99	All of our borrowings under the Revolving Facility are at variable rates of interest and expose us to interest rate risk.
covenants	1.00	Our credit facility contains operating covenants that may restrict our business and financing activities.
	1.00	The covenants in our outstanding indebtedness impose restrictions that may limit our operating and financial flexibility.
domestic	0.99	We are subject to the risks arising from adverse changes in domestic and global economic and political conditions.
	0.99	Our business is subject to risks arising from adverse changes in domestic and global economic conditions.
finance	0.99	Our ability to finance our ongoing operations may not always be possible solely from internal sources of capital and liquidity.
	0.99	We may not have access in the future to sufficient funding to finance desired growth.
industries	1.00	Substantially all of our customers operate in the energy or mining industries.
	1.00	Our operating results can be adversely affected by the general cyclical pattern of certain industries in which we operate.
lenders	0.99	The lenders can unilaterally adjust the borrowing base and therefore the borrowings permitted to be outstanding under our credit agreement.
	0.99	From time-to-time we may require consents or waivers from our lenders to permit actions that are prohibited by our credit facility.
loans	1.00	Significant increases in prevailing interest rates may significantly affect demand for loans and the value of bank collateral.
	1.00	These factors have resulted in a decline in the market value of many mortgage loans and related securities.
losses	0.99	Our loss reserves are based on estimates and may be inadequate to cover our actual losses.
	0.99	Any of these acquisition risks could result in unexpected losses or expenses and thereby reduce the expected benefits of the acquisition.
retail	−1.00	European regulatory officials are considering multiple proposed regulations intended to improve retail investor protection.
	−1.00	A general slowdown in the retail sector would result in additional pricing and marketing support pressures on Housewares.
volatile	0.99	Prices for oil and gas historically have been extremely volatile and are expected to continue to be volatile.
	0.99	Certain of our deposits and other funding sources may be volatile and impact our liquidity.

A.5 Bayesian Certainty Equivalent

Consider an investor with mean-variance preferences and information—they solve the problem

$$\max_{w\in\mathbb{R}^N} \mathbb{E}\left(\left[\begin{array}{c} 1-w'\iota \\ w \end{array}\right]'\left[\begin{array}{c} rf \\ r \end{array}\right]\middle| rf,\mathcal{I}\right) - \frac{\gamma}{2}\mathbb{V}\left(\left[\begin{array}{c} 1-w'\iota \\ w \end{array}\right]'\left[\begin{array}{c} rf \\ r \end{array}\right]\middle| rf,\mathcal{I}\right) \tag{A.11}$$

where risk aversion $\gamma > 0$, risky returns $r \in \mathbb{R}^N$, and risk-free $rf \geq 0$. We are emphasizing that $\{rf, \mathcal{I}\}$ is the investor's information set, because later we will alter both on our way to calculating a certainty equivalent accounting for parameter uncertainty—we call this a *Bayesian certainty equivalent*.

We focus solely on uncertainty about the parameters $\beta \in \mathbb{R}^{N\times K}$ that appear in a factor model, following standard APT assumptions as in Stambaugh (1983). That is, agents believe returns' mean is given by $\beta\lambda$ and covariance is given by $\beta\Delta_f\beta' + \Delta_\epsilon$. For simplicity we assume $\lambda \in \mathbb{R}^K, \Delta_f \in \mathbb{R}^{K\times K}, \Delta_\epsilon \in \mathbb{R}^{N\times N}$ are known ($\Delta_f, \Delta_\epsilon$ are positive definite with the latter diagonal).

We assume $r|\mathbb{E}(r), \mathbb{V}(r) \sim N(\mathbb{E}(r), \mathbb{V}(r))$ and therefore use the conjugate, normal-inverse-Wishart prior to express uncertainty about $\mathbb{E}(r)$ and $\mathbb{V}(r)$. A wrinkle in our formulation is that we wish to express uncertainty about β and this parameter appears *in both* the mean and variance. We have not seen this particular challenge tackled before, so to simplify analysis we represent beta uncertainty solely via the hyperparameter governing the conditional distribution of returns' mean.[11] That is, we assume (with abuse of notation) that

$$\mathbb{E}(r)|\beta_0, \kappa, \Delta_f, \Delta_\epsilon \sim N\left(\beta_0\lambda, \frac{1}{\kappa}\left[\beta_0\Delta_f\beta_0' + \Delta_\epsilon\right]\right) \tag{A.12}$$

$$\mathbb{V}(r)|\beta_0, \upsilon, \Delta_f, \Delta_\epsilon \sim W^{-1}\left(\beta_0\Delta_f\beta_0' + \Delta_\epsilon, \upsilon\right) \tag{A.13}$$

where $W^{-1}(\cdot)$ is the inverse-Wishart distribution. Note that by $\mathbb{E}(r)$ we mean $\beta\lambda$. Given these assumptions, we have

$$r \sim t_{\upsilon-N+1}\left(\beta_0\lambda, \frac{\kappa+1}{\kappa(\upsilon-N+1)}\left[\beta_0\Delta_f\beta_0' + \Delta_\epsilon\right]\right) \tag{A.14}$$

which is the multivarate non-central t-distribution—this is the distribution that the parameter-uncertain investor perceives as driving returns.

The usual mean-variance solution holds where the risky-asset portfolio is

$$w^* = \frac{1}{\gamma}\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r)$$

and using this alongside the mean and variance of the multivariate t given in (A.14), we find

[11] Future work can explore the viability of this assumption, in a fully-formed analysis

that the optimized utility can be written

$$\begin{aligned}
MV^* &= (w^*)'\mathbb{E}(r) + (1 - \iota' w^*) rf - \frac{\gamma}{2}(w^*)'\mathbb{V}(r) w^* \\
&= (\frac{1}{\gamma}\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r))'\mathbb{E}(r) + (1 - \iota'\frac{1}{\gamma}\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r)) rf - \frac{\gamma}{2}(\frac{1}{\gamma}\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r))'\mathbb{V}(r)\frac{1}{\gamma}\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r) \\
&= rf + \left(\frac{1}{2}\mathbb{E}(r) - rf\iota\right)'\left[\mathbb{V}(r)\right]^{-1}\mathbb{E}(r)\frac{1}{\gamma} \\
&= rf + \frac{\tilde{\upsilon}\tilde{\kappa}}{\gamma}\left(\frac{1}{2}\beta_0\lambda - rf\iota\right)'\left[\beta_0\Delta_f\beta_0' + \Delta_\epsilon\right]^{-1}\beta_0\lambda \qquad \text{(A.15)}
\end{aligned}$$

for $\tilde{\upsilon} \equiv \upsilon - N - 1$ and $\tilde{\kappa} \equiv \kappa(\kappa+1)^{-1}$.

Now we can simply state our Bayesian certainty equivalent calculation. Write $MV^*(rf, \tilde{\kappa})$ because we will hold constant the other parameters in (A.15). We find the s solving the equation we repeat from the main text

$$MV^*(rf, \tilde{\kappa}_1) = MV^*(rf + s, \tilde{\kappa}_2), \quad \text{with } \tilde{\kappa}_1 > \tilde{\kappa}_2, s \geq 0 \qquad (5)$$

for some given $\tilde{\kappa}_1, \tilde{\kappa}_2$. Note that $\tilde{\kappa}$ is an increasing transformation of κ: therefore their increase represent an increase in the precision of the prior distribution by (A.12). Note that s is a non-negative *increase* to the risk-free rate the agent receives. Intuitively, this risk-free-rate bump is positive when the smaller $\tilde{\kappa}_2$ means that the agent is facing *more* parameter uncertainty, which flows into the perceived variance of returns by (A.14) which the agent dislikes. Therefore we call s the *Bayesian certainty equivalent* for the increase in parameter precision (and therefore the agent's posterior return precision) implied by increasing $\tilde{\kappa}_2$ to $\tilde{\kappa}_1$.

Generally (A.15) requires stock-specific values owing to the presence of β_0 and Δ_ϵ. To empirically ground our calculations as simply as possible, we now assume that the $N = K$ and the assets available to the investor are the K factors themselves. Of course, literally these means that $\beta_0 = I$ and it appears on the surface that we've assumed away β_0 uncertainty entirely. However this is not an accurate depiction of what is going on in reality, because in the real world investors need to mimic the factors using the available assets, and *these* mimicking weights require β_0. Our conjecture is that a fully-formed analysis would show that β uncertainty translates into "something like" uncertainty about the factors in the large N limit, but we leave this analysis for future work. With our assumption, we can represent

$$MV^*(rf, \tilde{\kappa}) = rf + \tilde{\kappa}\frac{\tilde{\upsilon}}{\gamma}\left(\frac{1}{2}\lambda - rf\iota\right)'\Delta_f^{-1}\lambda \qquad \text{(A.16)}$$

Plugging in the mean and covariance of the Fama and French (2015) factors as well as the average rf from in Ken French's data over our sample period, we set $\gamma = 10$ and $\tilde{\upsilon} = 20$, we find s by solving (5).