Asset Pricing with News

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der KIT-Fakultät für Wirtschaftswissenschaften des
Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

M. Sc. Viktoria K. Klaus

Tag der mdl. Prüfung:	17. Dezember 2024
Referent:	TT-Prof. Dr. Julian Thimme
Korreferent:	Prof. Dr. Sven Klingler
Karlsruhe 2024	

Acknowledgements

I am deeply grateful to Julian Thimme for supervising my dissertation and enabling me to shape my own PhD journey, to Talina Sondershaus and Kathrin Hackenberg for the inspiring discussions and their guidance, to Sven Klingler for taking on the role of co-advisor and sharing my sense of humor, to Renee Adams for supporting women in finance, to Tim Knopp for all the technical support, to Fanchen Meng for being a wonderful collegue, to Christoph Schiller for broadening my horizon, and to participants of multiple research seminars at KIT, multiple annual meetings of the DGF, the WiMFEH Workshop at DIW in Berlin (2023), the Workshop on Sustainability and Aging at ZEW in Mannheim (2023), the EEA Conference in Barcelona (2023), and the AFA Meetings in San Antonio, TX (2024) for helpful comments and advice.

Furthermore, I gratefully acknowledge financial support from the PhD scholarship of the Konrad Adenauer Stiftung.

To my family and Martin Mehl: Thank you for everything.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Asset pricing is all about news. Every new piece of information entering the market can potentially change perceptions of asset risks. Thus, news concern the central objective of asset pricing theory: the correct pricing of assets according to their inherent risk. Within the risk-return ratio, the risk component cannot be measured without error in many places. The use of quantitative analyses of textual news data opens up new possibilities to improve the approximation of central risk factors in asset valuation. Hence, this thesis extracts information from news texts to investigate these risks. A prominent advantage of news texts is, that they are rich of narrative. That is, news convey interpretable information, refined by verbal explanations of complex context. By studying information disseminated in newspaper articles, researchers in finance can profit from journalists' expertise in deducing and elaborating the causes of events affecting markets (see Bybee et al., 2023). As news reflect concerns of readers (see Mullainathan and Shleifer, 2005), it is likely that investors seek exactly this sort of information, giving newspaper publishers a strong motivation to provide timely and pertinent details for market participants.

Earliest work on textual analysis of news in a financial context originates from Tetlock (2007, 2011), as well as Tetlock et al. (2008) and García (2013) measuring interactions between the stock market and the media.[1] A main hurdle for working with texts has long been the large data size which requires powerful computers. With computational power becoming less costly and feasibility of algorithms for textual analyses increasing, the growing of this strain

[1] A comprehensive literature review on handling and usage of textual data can be found in Gentzkow et al. (2019) and Loughran and Mcdonald (2016).

of literature gained traction over the last years. In asset pricing, the availability of news data is used to challenge established market models. Liu and Matthies (2022) test the long-run risk model by implementing a news-based consumption proxy and Bybee et al. (2023) as well as Aleti and Bollerslev (2024) estimate a Stochastic Discount Factor (SDF) with the help of news. Bybee et al. (2024) and Baker et al. (2016) take more of a high-level approach and use newspaper articles to describe the state of the economy. A recent application for news usage is addressing climate change concerns. Engle et al. (2020) and Ardia et al. (2023) develop news indices, that capture public attention towards climate change. Eventually, the authors - as well as further subsequent researchers - analyze the existence of a "greenium", i.e. a risk premium related to pollution levels of firms, with the help of these attention measures.

The aim of this dissertation is to incorporate news into classical challenges within the field of asset pricing and thereby offering guidance to a better understanding of markets and investors. In particular, throughout the thesis, we first introduce a new econometric method to lift macroeconomic time series to a higher frequency and provide comprehensive data of monthly time series.[2] Second, we develop novel news topic based asset pricing factors that capture fundamental risks of stocks. We exploit our new approaches in empirical studies to reevaluate and gain further insights on stock markets. Third and last, we use public attention to study trading behavior on another significant asset market: bank loans to low rated firms.

A general prerequisite for empirical models to work well is a large sample size. Hence, the validation of theoretical economic models developed in the late 20th century often hinges on the availability of data. While for example in asset pricing, returns on financial assets are available at high frequencies (up to milliseconds), macroeconomic data are only available at low frequencies. To overcome this frequency mismatch, returns are usually aggregated to lower frequencies causing the sample size to be significantly reduced. An appealing, but more complex option to solve the data mismatch is to lift macroeconomic variables to higher frequencies. Exploiting news data assumes that news media report timely and relevant information on the state of the economy. Machine learning algorithms can then leverage the large text data and filter out relevant information to proxy high-frequency economic variables. These generated time series provide significantly more information on economic developments for researchers as well as practitioners and policy advisers, and can then be used to test and improve empirical models,

[2] I have kept the pronoun "we" in the following chapters, as it is used in the working paper versions, to show that their basis is joint work.

for example in asset pricing.

A different strain of asset pricing models originates rather from empirical evidence than from theory and thus only offers little economic foundation. For these factor based models, a key issue is understanding the economic significance of common risk factors. Precisely, what fundamental risks do factors like SMB, HML, RMW, and CMA represent?[3] Answering this question can shed light on the economic developments that investors prioritize when taking portfolio decisions. These developments typically include macroeconomic events such as inflation, GDP changes, or political disturbances. Fama and French's (2015) return based factors likely capture various types of these fundamental shocks. Under the assumption that news are timely incorporated in investment decisions (see Fedyk, 2024), using news data can provide useful context to disentangle these underlying risks and understand the economic nature of return movements.

Nevertheless, news are not only relevant to improve comprehension of fundamental asset pricing models. As news reflect public attention, they are also a valuable tool to investigate investor trading behavior. More specifically, understanding investors' attention towards climate change is crucial in navigating the role of finance on the journey towards a low carbon society. If investors have no restrictions and can always buy securities from polluting companies divestments by banks or mutual funds, which do face regulatory constraints, might have minimal impact overall. This makes the leveraged loan market particularly noteworthy, as it serves as main funding source for lower-rated or private firms. So systematic withdrawal from polluting firms in this sector could severely hinder their operations. As the leading investors in the leveraged loan market and a key financing source for lower-rated firms, Collateralized Loan Obligations (CLOs) play a significant role. CLO investments are less transparent compared to those by institutional investors in the bond or equity markets, and CLO investors have fewer chances to withdraw their funds during the CLO's operation. Therefore, CLOs face less pressure to divest from polluting firms, allowing them to benefit from divestments of other investors. News data enable us to identify periods of high and low public attention towards climate change and lets us analyze, whether CLOs acquire loans from polluting firms during times of heightened climate change awareness, when other investors are more likely to withdraw.

[3]Besides a market factor, SMB (small-minus-big), HML (high-minus-low), RMW (robust-minus-weak), and CMA (conservative-minus-aggressive) are established portfolio based asset pricing factors by Fama and French (1993, 2015) and commonly used by researchers as well as practitioners.

1.2 Structure of the Dissertation

This thesis applies news data to different questions within the field of asset pricing. As presented in Figure 1.1, the dissertation begins with more general theory and ends with an analysis of a specific investor type. Precisely, the structure is as follows:

In Chapter 2, which is based on the working paper Klaus and Thimme (2024a), we develop a method to lift economic time series, such as output or consumption growth, from a low frequency, typically annual, to higher frequencies, such as monthly or even higher. Our method is based on the idea that the frequencies of certain word combinations used in newspaper articles correlate with the time series under consideration. Our approach identifies the most important word combinations through penalized regression on the low-frequency data and "predicts" the time series variation at the higher frequency. It accounts for characteristics of the time series, such as time aggregation and seasonal adjustment. We apply our method to a large number of low-frequency economic time series and find plausible patterns in the predicted high-frequency counterparts. The text dataset we use in our application includes the full text of all issues of the New York Times (NYT) over the past hundred years. By replicating two classic studies from consumption-based asset pricing, we illustrate the potential of our method for future research in economics.

Based on the working paper Klaus and Thimme (2024b) Chapter 3 investigates the economic foundation of the commonly used asset pricing factors SMB, HML, RMW, and CMA by Fama and French (1993, 2015). With the help of human coders, we use news coverage of stock market developments on days of large changes in factor returns to determine the economic topic behind. Our results show that macroeconomic news, monetary policy news, and news about the earnings of individual companies and industries are the most common causes of high factor returns for all four factors. From constructed topic-specific factors, we document plausible differences between the economic drivers of asset pricing factors. While HML is linked to macroeconomic news CMA is tied to news about commodities. The interpretation of major drivers of SMB and RMW is more complex: SMB correlates with exchange rate news and a sentiment-driven factor, and RMW is mainly influenced by firm-specific news. These findings suggest that both risk-based and behavioral channels contribute to explanations of factor risk premia.

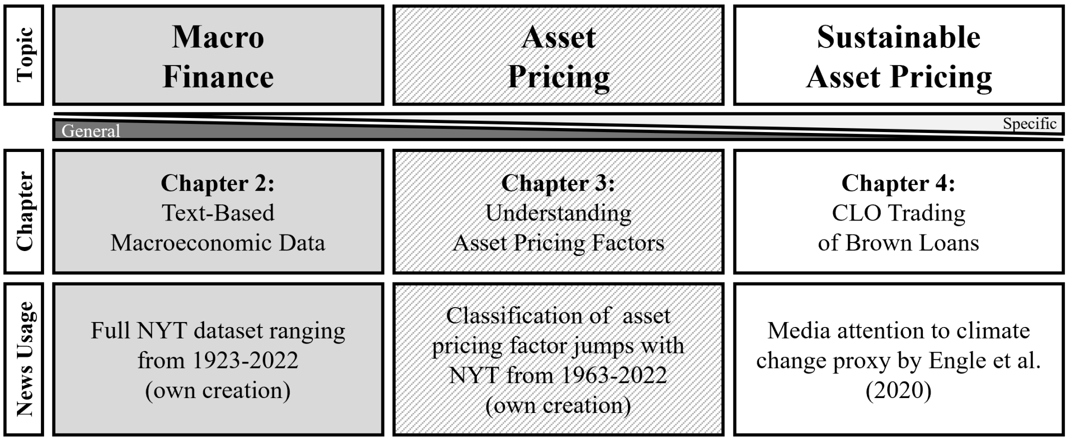

Figure 1.1: **Structure of this dissertation.**

Chapter 4 is based on the working paper Hackenberg et al. (2024)[4] and uses news indices to approximate public attention to climate change. We investigate the tremendously growing asset class of leveraged loans. Specifically, this chapter shows that CLOs, the main investors in leveraged loans, take advantage of heightened media attention to climate change. We document that loans from firms in carbon-intensive industries trade at a discount during these times. CLOs exploit these price discounts by increasing their investments in carbon-intensive industries. Notably, commitments to environmental considerations, such as signing the Principles of Responsible Investing (PRI), do not change this investment strategy. Instead, bank affiliated CLO managers particularly increase their investments, suggesting that banks use CLOs to offload (climate related) risks to improve the looks of their balance sheets. Hence, our findings highlight that CLOs are an investor class that purchases from those divesting from brown industries, indicating a potential conflict between responsible investment commitments and actual investment practices.

Chapter 5 summarizes the main results of the three chapters and concludes this dissertation with a brief outlook into future research opportunities.

[4]An earlier version of this paper was part of the dissertation by Hackenberg (2022). However, the current version differs significantly from this earlier draft as the analyses have been severely extended and the structure, the data sets used, and the writing have been completely revised.

Chapter 2

Text-based Macro Data and Asset Prices

2.1 Introduction

Economic theories should, like all scientific theories, be testable and falsifiable, i.e., conclusively rejectable by the data (Popper, 1968). In practical applications, this is often complicated by the fact that the economic variables, between which a theory hypothesizes a structural and testable relationship, are only observable at low frequency. Examples of important but only low-frequency (of sufficient quality) observable variables include output (Gross Domestic Product), consumption expenditures, many price indices, and annual economic data based on financial statements of companies. The resulting problems are twofold: Firstly, it is impossible to test high-frequency implications of economic models. Secondly, economic time series are often only available since the early 20th century or later, meaning that the samples are necessarily small. In this Chapter 2, we suggest and apply a text-based method to lift low-frequency time series data to higher frequencies.

The key idea underlying our approach is that newspaper articles contain up-to-date information relevant for capturing the current state of the economy (see Bybee et al., 2024). Our approach relies on the baseline assumption that the frequency of specific word combinations found in newspaper articles have a correlation with the time series being analyzed. More specifically, we suggest a two-step process to "predict" high frequency within-year variation of time series that are originally only available at low frequencies:[1] First, we employ penalized regression techniques, such as support vector regressions (SVR) and least absolute shrinkage and selection

[1] In the following sections, we will use the terms "predict" or "extrapolate" with the understanding that neither refers to actual out-of-sample prediction, but rather to filling in gaps within time periods by bringing the time series to a higher frequency.

operator (LASSO) regressions, to estimate relationships between annual economic indicators and word frequencies. These techniques help manage the high dimensionality of the word frequency data by shrinking (and selecting) regression coefficients, thereby ensuring a focus on the most relevant predictors. Second, our method uses the estimated coefficients from the first step to derive higher frequency fitted values.

Our method accounts for characteristics of the underlying macroeconomic time series. In particular, we show how to deal with time-aggregated annual time series (e.g., quantities measured as annual average over annual average) by generalizing the time-aggregation formula by Hall (1988) to arbitrary frequencies and applying it to word frequencies. Additionally, our approach accounts for leap years and seasonal adjustments, ensuring the consistency and validity of extrapolated series across different time periods. Together, these steps result in a comprehensive method that constructs consistent monthly time series from annual series by leveraging textual data.

For our application, we obtain full text data from the NYT historical archive from 1923 onwards. This makes the data set used in our work one of the largest among recent text-based studies in economics. For comparison, studies by Manela and Moreira (2017) incorporate only article abstracts, while Bybee et al. (2024) include full articles but only between 1984 and 2017, and Liu and Matthies (2022) only use articles which contain certain terms related to economic growth. Building the data set consists of several steps: The articles between 1923 and 1981 are available as photos, which have to be translated into machine readable text using Optical Character Recognition (OCR) software. Since this step can produce imperfect results, we employ algorithms to correct spelling errors (see García, 2013). After further extensive data cleaning, we count occurrences of words and word combinations of two, so-called bigrams.

Using this large matrix of bigram frequencies, we run an SVR and LASSO regression of annual economic time series on annual word frequencies. We then use the estimated slope coefficients together with the monthly word frequencies to create monthly proxies of multiple macroeconomic growth rates. Our goal is not to achieve strong out-of-sample performance but to precisely capture the within-year variations. Hence, our setting requires a different approach than the typical three-way data split over time into training, validation, and test sample. Instead, we tune the hyperparameters by considering time series that are observable on the monthly frequency (such as, e.g., industrial production) and aggregate them to the annual frequency. We

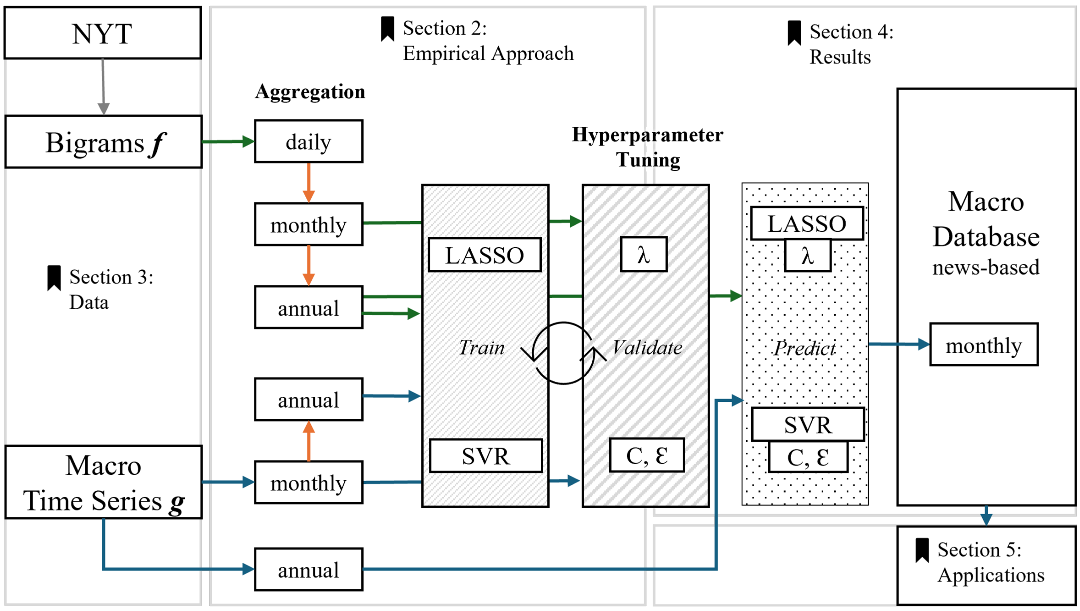

Figure 2.1: **Workflow.** This figure depicts the methodological workflow and structure of this chapter. Orange lines indicate aggregation. Green lines indicate input of bigram frequency data and blue indicates macroeconomic time series.

then apply our extrapolation method and investigate for which hyperparameters we achieve the best fit to the true actual within-year variation.

Figure 2.1 shows how our approach can be thought of in terms of the train-validate-test split procedure in standard machine learning studies: Intuitively, we can think of the annualized monthly time series, together with the annual word frequency data as the training data. Using standard performance metrics such as R^2, we then validate the accuracy of the predicted monthly text-based time series and tune our model hyperparameters. The observed monthly time series, together with the monthly word frequencies, can thus be thought of as validation data. The test data set consists of monthly time series of further time series. Since our approach is supposed to also work well on time series not observable on the monthly frequency, it is crucial that the optimal choice of hyperparameters is to a certain degree identical across time series.

Interestingly, we find that the best performance on the monthly validation data is achieved with a very mild regularization, meaning that the annual fitted regression models match the annual training data very closely (with R^2s close to one). This result raises the question whether our models suffer from overfitting. In standard machine learning applications, overfitted models are prone to only explain the training data while delivering poor results when confronted with new, unknown data for predictions. Our approach, however, is conceptually different, since all events we try to "predict" are already part of the training data set. A higher model complexity turns out to be helpful in finding a sufficient number of word combinations related to these events. Our findings second Kelly et al. (2024) who praise the "virtue of complexity" in prediction tasks for the US stock market, and Chen and Dim (2023), who argue that data-mined portfolio strategies perform well and are methodologically sound if detected in a consistent fashion. In a similar vein, Didisheim et al. (2024) show that richer models with thousands of parameters perform better out-of-sample in asset pricing settings.

In addition to the good time series fit on the monthly time series, our text-based approach also has the advantage that the selected word combinations can be interpreted economically. We look at the bigrams with the largest absolute coefficients and consistently find plausible terms. As an example, the bigram with the highest weight for the Consumer Price Index (CPI) is "price,increase." It is important to note that we do not predefine terms, as done by Baker et al. (2016) or Liu and Matthies (2022), but rather let the data speak for itself, i.e., allowing the algorithm to select from thousands of word pairs those that best explain the time series variation at the annual level.

After fixing the hyperparameters of SVR and LASSO regression models, we apply our method to further time series, that are only observable on the annual frequency. The resulting "high" (monthly) frequency data can then be used to test (macro-)economic models. We reconsider two classic applications from macrofinance to exemplify the potential use of our monthly time series. First, we estimate the risk aversion coefficient from consumption and asset return data, in the context of the Consumption-Capital Asset Pricing Model (Consumption-CAPM) (Lucas, 1978), with a Generalized Method of Moments estimator (Hansen, 1982). Second, we follow Hall (1988) and estimate the Intertemporal Elasticity of Substitution in consumption (IES) from a regression of consumption growth on interest rates.

Contribution to the Literature. The idea to use newspapers for text-mining purposes in the field of macroeconomics has gained traction in recent years. For instance, Baker et al. (2016) count the frequency of certain word combinations, such as "uncertain + economy + white house" in newspaper articles to construct a measure of economic policy uncertainty. The paper closest to our idea is Manela and Moreira (2017). They aim to provide a long history of data on the volatility index (VIX), a measure which is originally calculated from options and becomes only available in 1990. They use an SVR of the VIX on the relative frequencies of the words used in front-page articles of the Wall Street Journal (WSJ) over the sample starting 1990 and, then, use the fitted value from that regression to extrapolate the VIX back to 1890. While their approach is fully agnostic about the relation between words and uncertainty, they find that the frequencies of words like war and tax are very informative about the volatility of financial assets.

Our methodological approach adds to current literature on the textual analysis of news data (see Gentzkow et al. (2019) for an overview). In contrast to many studies that use topic modeling approaches (see e.g. Bybee et al., 2023, 2024; Adämmer and Schüssler, 2020), we use regularization to handle the high-dimensionality of textual data. We follow an agnostic approach towards word choices and, just as Manela and Moreira (2017), "let the data speak." The only restriction we put on bigrams in order to enter our news corpus is a minimum occurrence of 3600 times over the whole sample, spanning almost 100 years. This sets us apart from Fisher et al. (2022) who create a daily measure of attention to macroeconomic news from NYT and WSJ data from 1980 to 2020. For author defined word lists that are set out to capture attention to each category of news, the authors count the number of articles per day that relate to each category. For example, for U.S. output growth, they use "gross domestic product", "GDP", "gross national product", and "GNP". While their goal is to capture attention towards macroeconomic indicators, we aim to directly measure changes in the respective variable.

We further contribute to literature on consumption based asset pricing, i.e. Savov (2011), Kroencke (2017), and Pukthuanthong et al. (2023), who use alternative data as proxy for consumption. The latter use daily US household shopping data to create a new measure for consumption. Applied to an asset pricing setting the daily consumption measure results in lower risk aversion which solves the equity premium puzzle. Liu and Matthies (2022) predict consumption growth using only WSJ articles that relate to economic growth until 2013 and apply the prediction in a long-run risk asset pricing model. A growing literature on the "virtue of complexity" in economic models (see Kelly et al., 2024; Didisheim et al., 2024) challenges the conventional econometric notion of the superiority of simplistic models with only few free parameters. They show that richer models can perform better out-of-sample if applied accurately.

In economics, knowing the time lag in publication that the elaborate measuring process imposes, there is also a large literature on nowcasting of economic data (see e.g. Bok et al., 2018) which relates to our work. Although researchers also estimate high-frequency data, our approach is not meant to predict the future, but rather utilize historic data to fill in gaps. For a literature review on economic nowcasting see Stundziene et al. (2024).

2.2 Empirical Approach

2.2.1 General Approach

The idea underlying our approach is that the frequency of certain words or word combinations used in the newspaper is correlated with the economic time series we are studying. This is the standard assumption in studies that follow a bag-of-words approach. We denote the frequency of word (combination) w in year y, month m, day d as $f^w_{y,m,d}$. Aggregated frequencies are denoted with a shorter subscript: $f^w_{y,m} = \sum_d f^w_{y,m,d}$ and $f^w_y = \sum_{m,d} f^w_{y,m,d}$ denote the frequencies of word combination w in month y, m and in year y, respectively. For simplicity, we refer to years and months, as in our application, we convert annual time series to a monthly frequency. However, any periods for low frequency (y) and high frequency (m) are conceivable, as long as the high frequency does not exceed that of the text data used for the application (in our case, daily).

The time series to be extrapolated is denoted as $(g_y)_y$ and g_y denotes the realization of the time series in year y. We use the same subscripts as above to denote the realizations of the time series over different periods of time, such as months or days. Fitted values, that means estimated time series, corresponding to the annual time series $(g_y)_y$ used in the estimation, are denoted as $(\hat{g}_y)_y$, $(\hat{g}_{y,m})_{y,m}$, and $(\hat{g}_{y,m,d})_{y,m,d}$.

In the remainder of Section 2.2, we will discuss necessary refinements to our approach, but the basic idea is very simple. The approach proceeds in two steps:

1. We estimate the coefficients b^w in the following linear regression equation using a penalized regression method:

$$g_y = a + \sum_w b^w f^w_y + \eta_y, \tag{2.1}$$

2. To extrapolate g to a higher frequency, we then use the fitted value

$$\hat{g}_{y,m} = \frac{a}{12} + \sum_w b^w f^w_{y,m} \tag{2.2}$$

For this to work, the considered time series must be stationary and have the property that the latent monthly time series can be aggregated by summing to the observed annual time series. If the time series has a unit root, as e.g. when considering asset prices or a stock index, we take the logarithm of the time series and consider the first differences. Logarithmic transformation is important because gross returns, while stationary, are not additive. But there can be other reasons why time series on the high frequency do not aggregate to the low frequency equivalent by taking sums.

First, the time series might be stationary but not quantify a change from a previous point in time. This would be the case, e.g., with an interest rate or a ratio such as the ratio of labor income to asset holdings. In this case, one could regress the given time series on the word frequencies $f^w_{y,m,d}$ on the day of measurement (typically a key date at the end of the year) and then extrapolate to all other days (or days at the end of every month). While this approach is conceptually correct, it is not advisable for practical reasons: daily word frequencies are highly volatile, so it is recommended to work with partially aggregated data even at the higher frequency (at least on a weekly, better on the monthly frequency).

The (better) alternative is to also consider first differences for stationary time series. In this way, the regression design described above is both conceptually correct (since first differences are again additive and can be aggregated from the high to the low frequency) and practical when dealing with volatile daily word frequencies. As we will discuss later, it is optimal to basically perfectly match the annual time series in the penalized regression (i.e., to choose the penalization such that many word combinations receive a positive weight), so using the first differences instead of the original time series does not compromise quality: the original time series at the lower frequency can always be recovered by summing.

Second, the time series might be time-aggregated, i.e., the individual realizations of the level variables could be sums or averages over a longer period. This is very often the case with important macroeconomic time series. For example, industrial production is measured over an entire month, quarter, or year, rather than at a specific point during or at the end of such a period. As a result, while the level is additive (the sum of industrial production across the months of a year equals that of the year), it is not stationary. The monthly logarithmic growth rates, however, do not sum to the annual growth rates. These are given by the logarithmic production in one year minus the logarithmic production in the previous year.

Other examples of time-aggregated macroeconomic variables include consumption, output, imports and exports, and government spending. Although statistics like consumer price indices are conceptually point-in-time measurements, they are practically measured over a period and must therefore also be treated as time-aggregated. Moreover, there are some cases that are neither clearly point-in-time nor time-aggregated. For instance, the Bureau of Labor Statistics determines its unemployment rate through household interviews during the middle week of a month. Other variables are hybrid, such as the price-dividend ratio, where the price of a security is measured at a specific point in time, but the dividends usually reflect those of the past year.

While these hybrid variables are difficult to extrapolate, there is a simple variant of the approach presented in Equations (2.1) and (2.2) for commonly time-aggregated variables. The exact functioning of this variant is explained in the following Section 2.2.2.

2.2.2 Time-Aggregation

Denote by X_y a time-aggregated level variable, e.g. industrial production, over the year y and by $X_{y,m}$ the monthly analogue, so that $X_y = \sum_m X_{y,m}$. The annual time-aggregated growth rate in year y is given by

$$g_y = \log(X_y) - \log(X_{y-1}).$$

As described above, the problem for our approach is that g_y does not correspond to the sum of the time-aggregated monthly growth rates $g_{y,m} = \log(X_{y,m}) - \log(X_{y,m-1})$ over the twelve months of year y. Nevertheless, annual time-aggregated growth rates can approximately be expressed as a weighted sum of monthly time-aggregated growth rates. Hall (1988) and Patton and Timmermann (2011) use the following formula, which is based on a first-order Taylor approximation of the log:

$$g_y \approx \sum_{m=2}^{12} \frac{m-1}{12} g_{y-1,m} + \sum_{m=1}^{12} \frac{12-m+1}{12} g_{y,m}.$$

More generally, this formula can be used for all pairs of high and low frequencies to aggregate time-aggregated log growth rates at the high frequency $g_{l,h}$ to those of the low frequency g_l:

$$g_l \approx \sum_{h=2}^{H} \frac{h-1}{H} g_{l-1,h} + \sum_{h=1}^{H} \frac{H-h+1}{H} g_{l,h}. \tag{2.3}$$

Here, H represents the number of time units of the high frequency within one time unit of the low frequency. The case where the number H depends on the instance of the low frequency unit (e.g., leap years or months with varying numbers of days) is discussed in Section 2.2.3. A derivation of Equation (2.3) is presented in Appendix A.1.

Our approach is motivated by Hall (1988), who wants to estimate the intertemporal elasticity of substitution in consumption. He assumes a monthly decision interval, so that the standard consumption-based asset pricing model implies a linear relation between monthly household consumption growth and monthly interest rates (see Section 2.5.2). Since consumption is only available on the annual frequency, he aggregates monthly interest rates to an annual time-aggregated series, using Equation (2.3), and then regresses the available annual time-aggregated consumption growth time series on this time series constructed from monthly interest rates. Note that the constructed time series does not have a direct economic interpretation. However, since the time aggregation formula in Equation (2.3) is a linear operation, and linear regression is also a linear operation, the regression coefficients of the two time-aggregated time series are the same as if one were to regress the monthly time series against each other. This only holds in expectation and asymptotically, as the sample size is obviously reduced through aggregation.

When extrapolating time-aggregated economic time series, we follow Hall (1988) and also time-aggregate our independent variables, in our case the word frequencies. More precisely, in step 1 of the approach presented in Section 2.2.1, we regress on time-aggregated annual frequencies $\tilde{f}_y^w$ instead of f_y^w, and in step 2 we calculate fitted values by using time-aggregated monthly frequencies $\tilde{f}^w_{y,m}$ instead of $f^w_{y,m}$. Here, $\tilde{f}^w_y$ and $\tilde{f}^w_{y,m}$ are given by[2]

$$\tilde{f}^w_y := \sum_{d=2}^{365} \frac{d-1}{365} f^w_{y-1,d} + \sum_{d=1}^{365} \frac{365-d+1}{365} f^w_{y,d} \tag{2.4}$$

$$\tilde{f}^w_{y,m} := \sum_{d=2}^{30} \frac{d-1}{30} f^w_{y,m-1,d} + \sum_{d=1}^{30} \frac{30-d+1}{30} f^w_{y,m,d} \tag{2.5}$$

In Equations (2.4) and (2.5), we have assumed 365 calendar days per year and 30 calendar days per month. This, of course, is not always correct. How the formula needs to be adjusted depends on whether the original time series is seasonally adjusted and whether the extrapolated time series is meant to remove the mechanical seasonality caused by the varying number of days in a month. We will discuss this topic in the following Section 2.2.3.

2.2.3 Leap Years and Seasonal Adjustments

Equation (2.3) holds only if the low-frequency period always contains the same number (namely H) of high-frequency subperiods. This applies in the example from Hall (1988), since a year always has twelve months. However, in our case, adjustments are necessary because months and years can have different numbers of days. The derivation in Appendix A.1 naturally generalizes to

$$g_l \approx \sum_{h=2}^{H_{l-1}} \frac{h-1}{H_{l-1}} g_{l-1,h} + \sum_{h=1}^{H_l} \frac{H_l-h+1}{H_l} g_{l,h} + \log(H_l) - \log(H_{l-1}), \tag{2.6}$$

where H_{l-1} represents the number of high-frequency subperiods in low-frequency period $l-1$ and H_l represents the equivalent for period l.

In our application, we work with annual time series. Macroeconomic time series are typically leap year adjusted, but the word frequencies have to be aggregated with care. If the annual value to be extrapolated is measurable at a specific point in time, then in leap years we multiply

[2]Equations (2.4) and (2.5) are with abuse of notation. $f^w_{y,d}$ with d running from 1 or 2 to 365 is supposed to denote the word frequencies of the days in a year. We dropped the explicit mention of the month in the subscripts. $f^w_{y,m-1,d}$ is not defined properly, if month m is January. We then mean by $m-1$ the December of the previous year.

the sum of the daily word frequencies by the correction factor 365/366, i.e., in general,

$$f_y^w = \frac{365}{H_y} \sum_{d=1}^{H_y} f_{y,d}^w, \tag{2.7}$$

where H_y denotes the number of calendar days in year y.

If the annual statistic to be extrapolated is time-aggregated, we correct for the difference of the logarithmic terms at the end of Equation (2.6) and multiply the sums of the daily word frequencies by correction terms:

$$\tilde{f}_y^w = \frac{365-1}{H_{y-1}-1} \sum_{d=2}^{H_{y-1}} \frac{d-1}{H_{y-1}} f_{y-1,d}^w + \frac{365+1}{H_y+1} \sum_{d=1}^{H_y} \frac{H_y-d+1}{H_y} f_{y,d}^w. \tag{2.8}$$

The correction terms result from the usual summation formulas

$$\sum_{h=2}^{H} \frac{h-1}{H} = \frac{1}{2}(H-1), \text{ and } \sum_{h=1}^{H} \frac{H-h+1}{H} = \frac{1}{2}(H+1).$$

In Step 1 of our approach from Section 2.2.1, we regress the annual time series on the thus aggregated and corrected word frequencies $(f_y^w)_w$ and $(\tilde{f}_y^w)_w$ to estimate the coefficients a and b^w. In the second step, we extrapolate using these coefficients on a monthly frequency. Since our extrapolated monthly time series should at least be adjusted for the mechanical effects that different months have different numbers of days, the monthly word frequencies must also be corrected. More specifically, we again use Equation (2.2) to estimate $\hat{g}_{y,m}$, but with adjusted monthly word frequencies:

$$f_{y,m}^w = \frac{(365/12)}{H_m} \sum_{d=1}^{H_m} f_{y,m,d}^w, \tag{2.9}$$

$$\tilde{f}_{y,m}^w = \frac{(365/12)-1}{H_{m-1}-1} \sum_{d=2}^{H_{m-1}} \frac{d-1}{H_{m-1}} f_{y,m-1,d}^w + \frac{(365/12)+1}{H_m+1} \sum_{d=1}^{H_m} \frac{H_m-d+1}{H_m} f_{y,m,d}^w, \tag{2.10}$$

where H_m denotes the actual number of days in month m.

2.2.4 Penalized Regression Approaches

The linear regression in Equation (2.1) involves thousands of independent variables on the right-hand side in our application, making a standard OLS regression infeasible. Instead, we use two different penalized regression techniques, both of which modify the usual least squares objective function by adding penalty terms for large regression coefficients.

First, we use SVR. For a formal introduction to the method, see Smola and Schölkopf (2004). Intuitively, an SVR seeks coefficients b^w such that the fitted value of the regression does not deviate more than a tolerance value ε from the true value. If the deviation between the fitted value and the true function value for an observation d exceeds the tolerance by ξ_d, this results in a penalty term. The objective function is given by

$$J((b^w)_w) = \frac{1}{2}\sum_w (b^w)^2 + C\sum_d \xi_d, \tag{2.11}$$

where the hyperparameter C determines the trade-off between regularization, meaning the goal of achieving small coefficients b^w in magnitude, and a good model fit. SVR shrinks all coefficients toward 0, but it does not result in only a small number of coefficients being different from zero. Manela and Moreira (2017) use the method for backward extrapolation of the VIX.

Second, we use LASSO regressions. Here, the usual least squares objective function is augmented by the constraint that the sum of the absolute regression coefficients must not exceed a regularization coefficient λ. This implies the objective function

$$J((b^w)_w) = \frac{1}{Y}\left(g_y - (a + \sum_w b^w f_y^w)\right)^2 + \lambda|b^w|.$$

Compared to SVR, LASSO has the property that only a small number (depending on the regularization parameter λ) of word frequencies have a non-trivial coefficient b^w. This makes the fitted value potentially more sensitive to random, non-systematic fluctuations of individual word frequencies. However, the selected word combinations are not as numerous, and as a result, the monthly time series is more interpretable.

Since more volatile word frequencies ceteris paribus exhibit smaller regression coefficients and are preferentially selected by LASSO, we need to standardize the word frequencies. To do this, we divide each annual word frequency time series $(f_y^w)_y$ or $(\tilde{f}_y^w)_y$ by its standard deviation:

$$\sigma^w = \left(\frac{1}{Y-1}\sum_{y=1}^{Y}\left(\mathrm{f}_y^w - \frac{1}{Y}\sum_{y=1}^{Y}\mathrm{f}_y^w\right)^2\right)^{\frac{1}{2}}, \tag{2.12}$$

where Y denotes the number of years in the sample, and $\mathrm{f} = f$ or $\mathrm{f} = \tilde{f}$, depending on whether g is point-in-time or time-aggregated.

When estimating the time series, the monthly word frequencies $(f^w_{y,m})_{y,m}$ or $(\tilde{f}^w_{y,m})_{y,m}$ must then be divided by the same annual standard deviation to ensure that the resulting monthly time series is consistent.

Our approach can be summarized as follows:

1. To estimate the coefficients a and $(b^w)_w$, we perform an SVR with hyperparameters C and ε or a LASSO regression with hyperparameter λ, corresponding to the regression equation

$$g_y = a + \sum_w b^w F^w_y + \eta_y. \tag{2.13}$$

Here, $F^w_y = f^w_y/\sigma^w$ if g is point-in-time, and $F^w_y = \tilde{f}^w_y/\sigma^w$ if g is time-aggregated. f^w_y and $\tilde{f}^w_y$ are defined in Equations (2.7) and (2.8), and σ^w in Equation (2.12).

2. To extrapolate g to a monthly frequency, we use the fitted value

$$\hat{g}_{y,m} = \frac{a}{12} + \sum_w b^w F^w_{y,m}. \tag{2.14}$$

Here, $F^w_{y,m} = f^w_{y,m}/\sigma^w$ if g is point-in-time, and $F^w_{y,m} = \tilde{f}^w_{y,m}/\sigma^w$ if g is time-aggregated. $f^w_{y,m}$ and $\tilde{f}^w_{y,m}$ are defined in Equations (2.9) and (2.10), and σ^w in Equation (2.12).

In order to assure comparability across macroeconomic time series, we also divide each macro time series by its standard deviation. This is necessary as we want to select hyperparameters that fit well for all macroeconomic time series. We can later multiply with the standard deviation to return to a non-standardized time series. The hyperparameters used in step 1 are crucial for the estimated coefficients. How we choose them will be discussed in the next Section 2.2.5.

2.2.5 Hyperparameter Tuning Process

In standard applications of penalized regressions, the dataset is divided into training, validation, and testing subsets. The hyperparameters are selected in such a way that the model's performance on the training dataset transfers as well as possible to the validation dataset. A statistic, such as R^2 or the sum of squared errors, is used to measure the quality of the model on the validation dataset, and this statistic is maximized by optimally choosing the hyperparameters. The test dataset is then used to investigate the true out-of-sample performance of the model. The same statistic as before can be applied for this purpose. Ultimately, the goal is to generate a model with good out-of-sample performance.

However, the goal of our approach is different. We are not aiming for good out-of-sample performance, but rather for accurately capturing the intra-annual variation of an annual time series. A word frequency that shows strong explanatory power for variation in an economic variable in a given sample does not necessarily need to do so in the future. Therefore, a 3-way split of our annual sample is not suitable for our application.

Instead, we use time series that are also available at a high frequency to select the hyperparameters. We aggregate these to an annual frequency and perform our extrapolation approach, described in Equations (2.13) and (2.14). The estimated monthly time series can then be compared with the actual time series, and we select the hyperparameters that maximize the monthly R^2. We then check whether the same hyperparameters also perform well for other time series that are available at a monthly frequency.

To use the usual terminology, we divide the set of available monthly observable time series into two groups. One group serves as the training and validation dataset. The training dataset is the time series aggregated to an annual frequency, and the validation dataset is the original monthly time series. The monthly time series of the other group then serve as test datasets. We then use the same hyperparameters to extrapolate the time series that are only available at a low frequency to a high frequency.

We use the monthly R^2 as the performance metric to be optimized, because the goal of our approach is to generate the most accurate possible monthly time series. It is calculated as

$$R^2_{\text{monthly}} = 1 - \frac{\widehat{\text{Var}}(\xi_{y,m})}{\widehat{\text{Var}}(g_{y,m})}, \tag{2.15}$$

where $\xi_{y,m}$ is the residual from an OLS regression of the true monthly values on the predicted monthly values, that means

$$g_{y,m} = \alpha + \beta \hat{g}_{y,m} + \xi_{y,m}.$$

The monthly R^2_{monthly} quantifies the share of the variation in the true monthly time series that can be explained by the predicted time series. It is partly in-sample and partly out-of-sample, because the annual variation is used to train the algorithm, while the monthly variation within the years is not used when training the algorithm.

We also report the annual R^2, which is calculated as

$$R^2_{\text{annual}} = 1 - \frac{\widehat{\text{Var}}(g_y - \hat{g}_y)}{\widehat{\text{Var}}(g_y)} \tag{2.16}$$

It can be interpreted as an in-sample R^2.

Since the monthly R^2 is also affected by the overall fit of the annual variation (which is in-sample), we also compute a measure that is supposed to quantify within-year fit. We call it the *Within-*R^2. It is calculated as

$$R^2_{\text{within}} = 1 - \frac{\widehat{\text{Var}}(\zeta_{y,m})}{\widehat{\text{Var}}(\zeta^{FE}_{y,m})}, \tag{2.17}$$

where $\zeta_{y,m}$ are the residuals from a regression of true monthly values on the predicted monthly values, including year fixed effects:

$$g_{y,m} = \beta_{\text{within}}\hat{g}_{y,m} + \sum_y \gamma_y \mathbf{1}_y + \zeta_{y,m} \tag{2.18}$$

and $\zeta_{y,m}$ are the residuals from a regression of true monthly values on year fixed effects only:

$$g_{y,m} = \sum_y \gamma^{FE}_y \mathbf{1}_y + \zeta^{FE}_{y,m} \tag{2.19}$$

The *Within-*R^2 is purely out-of-sample because the annual variation used to train the algorithm is taken out by the year fixed effects. We are also interested if the coefficient β_{within} is close to one and perform t-tests to formally test the hypothesis if $\beta_{\text{within}} = 1$.

2.3 Data

2.3.1 Words and Word Combinations

Our text data set consists of the full text corpus of the NYT from 1923 to mid-2022. These data are directly downloaded from the NYT website. Before 1981 newspapers are not available in a machine readable format, but only as images. In order to still utilize these data we translate these into text via OCR. This process is erroneous and commands the following corrections before we can start counting word and word combination occurrences.

First, we convert all text into lower case characters. This is crucial as words need to be written in the exact same way throughout the text to count occurrences properly. Second, we delete hyphens before a line break, indicating that a word is divided in order to fit into the given line space. Third, we split contractions, e.g. "aren't" becomes "are not". Fourth, we employ spelling corrections on each word using SymSpell[3]. Its dictionary-based approach increases the accuracy of the word correction. For example, due to bad image quality an "i" can be translated into an "f" turning the word "economic" into "economfc". SymSpell is able to identify the unknown word and correct the mistake. Fifth, we replace remaining hyphens with an empty space. Sixth, we remove all numbers and punctuation. Seventh, we delete unknown words (words that are not in the English dictionary) as well as words that consist of less than two characters. Eighth, we drop stop words because they do not convey relevant information (see Gentzkow et al., 2019). At last, we reduce all words to their word stem. This is useful to cover different grammatical forms of words. For example, "increase", "increasing", and "increases", all carrying the same meaning, fall with our approach under the same stemmed word "increas".

We follow a classic bag-of-words approach which leads us to count word frequencies for each day. After the cleaning process explained above, a single issue of the NYT contains on average about 84,000 words, so our one-hundred-year daily dataset comprises about 3 billion words. The number of distinct words amounts to 45,770.

In order to better access the scope of information provided by the text, we use bigrams (phrases consisting of the word stems of two key words, such as e.g. "price increas" instead of "price" and "increas" separately). Using bigrams instead of words largely increases the number

[3]Detailed information on the procedures of SymSpell by Wolf Garbe can be found at `https://github.com/wolfgarbe/SymSpell`.

of unique occurrences from thousands to millions. Ideally, we want our model to allocate large coefficients to bigrams that occur frequently and are therefore informative throughout time. Hence, we limit our data set to only include bigrams that on average occur at least three times per month, thus 3600 times in the whole sample.[4] This leaves us with 142,126 unique bigrams. Some economic time series are only available later within the century. We additionally drop bigrams that do not occur at least once in the time period where the economic time series is available.

The NYT is published in principle seven days a week for our entire sample. However, some issues are missing. Specifically, during a strike in the fall of 1978, no newspaper was published. Since our method relies on daily word frequencies, we replace word frequencies for missing days by the average of 30 days before and 30 days after the occurrence.

At last, we normalize the bigrams by dividing each day's bigram count by the sum of bigrams on that day to account for different newspaper lengths. Even though this may initially seem confusing, as we have always spoken of absolute word frequencies in Chapter 2.2, it is completely irrelevant for our method. We just sum the daily relative word frequencies using Equations (2.9) and (2.10) to calculate monthly (and annual) aggregated word frequencies. The use of relative word frequencies is better than absolute ones because the overall length of the NYT has significantly decreased over time. As a result, absolute word frequencies exhibit a negative trend in the cross-sectional average of bigrams.

2.3.2 Macroeconomic Time Series

We select the following macroeconomic time series, which are available on a monthly frequency, to analyze the optimal choice of the hyperparameters. All of them are downloaded from the website of the St. Louis Fed at `fred.stlouisfed.org/series`.

Industrial Production. The data comes as a seasonally adjusted monthly index. To aggregate the time series to the annual frequency, we total the monthly index values of a calendar year and calculate the logarithmic growth rates from one year to another. Thus, the annual time series is a time-aggregated series, which is consistent with the nature of the monthly time

[4] We also test a lower threshold of 1000 occurrences and find similar results. Thus, the exact choice of the threshold does not seem crucial for our results.

series, also being time aggregated. We additionally use sub-indices of industrial production, i.e. manufacturing, energy, and consumer goods.

Employment Data. We calculate the year-to-year logarithmic growth rate for the annual totals of the unemployed and the employed people as well as the employed women in the US. We calculate the sum over the twelve monthly quantities of (un)employed persons/women and then calculate the year-over-year growth rate in these sums. Note that an aggregated quantity value in the annual time series cannot directly be interpreted as the number of, e.g., unemployed in a calendar year. Rather, it is to be interpreted as the total number of person-months being unemployed in that year.

Price Indices. The monthly seasonally adjusted consumer price index is aggregated to an annual time series by taking the averages of the monthly index values within a year. Thus, the annual price index time series quantifies average consumer prices within the years. We then calculate log growth rates, again leading to a time-aggregated annual time series. We follow the same approach for the producer price index.

Dividends. We retrieve daily aggregated dividends from Compustat. For each day, total dividends are calculated as the average of the following 90 days. We follow that procedure, since dividends are typically paid out at the end of a fiscal quarter, even though earned throughout the period. To calculate seasonally adjusted dividends we apply an X-13 ARIMA filter. Daily dividends are then summed up within calendar years. The annual time series used in our analysis is the log growth rate of these annual dividend payments, which is time-aggregated.

Sales. We retrieve monthly sales data from Compustat. Following the same logic as for dividends, we compute for each month sales as the average reported sales over the subsequent three months. In order to calculate real seasonally adjusted sales values we again use the monthly CPI and the X-13 ARIMA filter. The annual growth rate in sales is again a time-aggregated series.

Confidence Index. We download the time series as a seasonally adjusted monthly index. To aggregate the time series to the annual frequency, we average the index values of the twelve months in a calendar year and then calculate the year-over-year growth rate in these averages. The time series is time-aggregated.

Table 2.1 shows summary statistics and the respective samples. All our samples end in 2019, since some time series, such as unemployment, show extreme spikes during the Covid pandemic in 2020. This leads to the situation that algorithms primarily pick bigrams related to Covid, which, however, are not at all informative about the time series variation outside of the Covid period. The problem that single extreme events have the potential to contaminate the bigram choice requires rolling-window variants are other "conditional" versions of our algorithm. We leave their development for future research.

Table 2.1: **Summary statistics of macro data.** This table reports summary statistics of annual macroeconomic log growth rates used in this study to tune the hyperparameters for LASSO and support vector regressions.

	Name	start	end	count	mean	std	min	25%	50%	75%	max
1	Production Index	1924	2019	96	0.030	0.085	-0.248	0.008	0.035	0.077	0.233
2	Consumer Price Index	1948	2019	72	0.034	0.027	-0.010	0.016	0.028	0.041	0.127
3	Unemployment	1949	2019	71	0.014	0.190	-0.455	-0.091	-0.034	0.062	0.656
4	Dividends	1927	2019	93	0.064	0.118	-0.520	0.038	0.071	0.107	0.431
5	Producer Price Index	1930	2019	90	0.028	0.059	-0.168	0.002	0.023	0.055	0.206
6	Employment	1940	2019	80	0.020	0.026	-0.044	0.008	0.020	0.031	0.122
7	Women Employment	1965	2019	55	0.025	0.020	-0.028	0.014	0.022	0.040	0.077
8	Confidence	1961	2019	59	0.000	0.010	-0.029	-0.005	0.001	0.006	0.026
9	PI - Manufacturing	1924	2019	96	0.029	0.092	-0.264	0.004	0.037	0.083	0.245
10	PI - Energy	1955	2019	65	0.019	0.035	-0.077	-0.000	0.017	0.044	0.119
11	PI - Consumer Goods	1940	2019	80	0.026	0.044	-0.078	0.004	0.024	0.044	0.181
12	Sales	1963	2019	57	0.074	0.068	-0.149	0.040	0.073	0.111	0.236

2.4 Text-based Macro Time Series

2.4.1 Hyperparameter Tuning

We perform an extensive grid search to select the optimal hyperparameters on our training data, i.e. the monthly macro time series aggregated to annual frequency, listed in Table 2.1. We then calculate the metrics described in Section 2.2.5 to evaluate model fit. Figure 2.2 plots R^2_{monthly} depending on λ in the LASSO regression. For the *Production Index* we find a maximum of R^2_{monthly} at $\lambda = 0.078$ and for the *Consumers Price Index* the metric peaks at $\lambda = 0.082$. For *Unemployment*, R^2_{monthly} plateaus at a similar value, with the maximum being located at about

0.09. For the other time series as well, we find maxima near the value 0.08, and between 0.08 and the actual maximum, R^2_{monthly} changes only marginally. We therefore choose 0.08 as the penalty parameter in the LASSO regression.

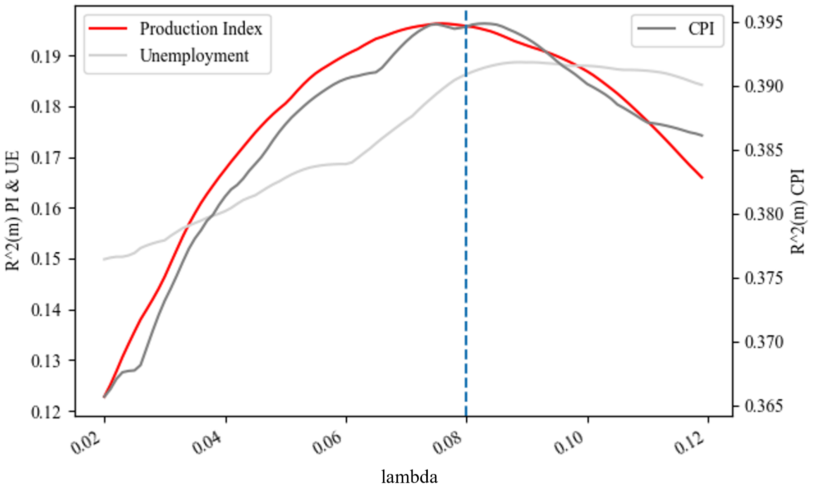

Figure 2.2: **LASSO hyperparameter tuning.** The figure shows R^2_{monthly} as a function of the tuning parameter λ in the LASSO regression. The red line shows R^2_{monthly} for the *Industrial Production* index, the dark gray for the *CPI*, and the light gray for *Unemployment*. The left y-axis shows the level of R^2_{monthly} for *Industrial Production* and *Unemployment* and on the right for the *CPI*.

Figure 2.3 shows results of the same analysis for the two tuning parameters C and ε in the SVR, holding the respective other tuning parameter fixed at the value selected later. In our analysis of the optimal tuning parameters, we actually perform a grid search on a two-dimensional grid but only show conditional line plots here for better interpretability. For the three series mentioned above, we find R^2_{monthly} decreasing in ε and increasing in C. However, all R^2_{monthly}s plateau at around $C = 0.00025$ and $\varepsilon = 0.01$, so that further increasing C or decreasing ε does not improve the fit on the monthly frequency materially. We thus select $C = 0.00025$ and $\varepsilon = 0.01$ for our application.

The patterns of the two line graphs in Figure 2.3 indicate that our method benefits from very weak regularization. In SVR, a small ε means that only very small deviations of the fitted value from the true function value are accepted without penalty. A large C penalizes deviations beyond ε particularly severely (see Equation (2.11)). The goal of choosing the coefficients b^w to be small is comparatively less important. As we discuss below, this leads to a de facto perfect fit in the regression of the annual time series on the annual word frequencies. A similar picture emerges in the LASSO regressions: The penalty parameter λ is chosen so low that a very large number of non-zero coefficients are always selected, often more than 50% of the sample size. Here, too, the explanatory power on the annual training datasets is consistently very high.

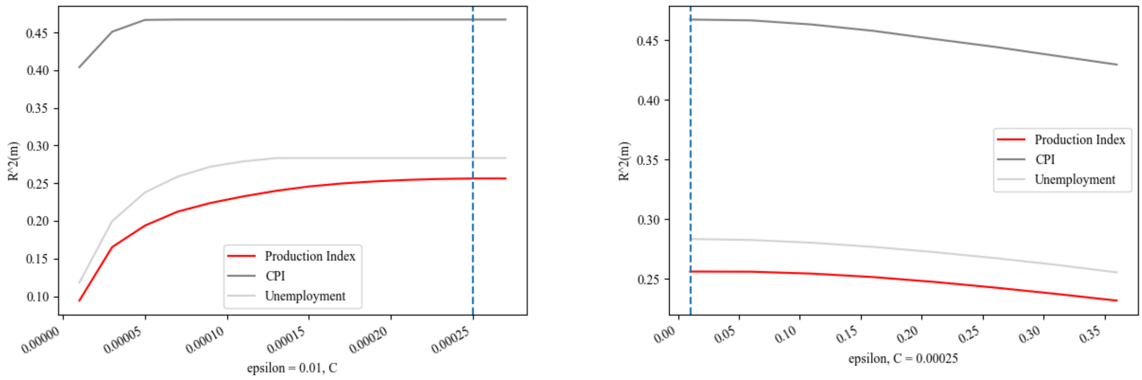

(a) R^2_{monthly} as a function of C.

(b) R^2_{monthly} as a function of ε.

Figure 2.3: **SVR hyperparameter tuning.** The figure shows R^2_{monthly} as a function of the tuning parameters C (graph a) and ε (graph b) in the SVR regression. The red line shows R^2_{monthly} for the *Industrial Production* index, the dark gray for the *CPI*, and the light gray for *Unemployment*.

This result raises the question whether our models suffer from overfitting. Usually, rich models with many parameters are seen with a certain dose of skepticism because overfitted models are prone to only explain the training data while delivering poor results when confronted with new, unknown data for predictions. However, our setting differs from other commonly used time series applications as we explicitly refrain from predicting "unknown" future events. All events we try to "predict" with our models are already part of the training data set. An extremely close fit to the training data is advantageous to identify (possibly a large number of) relevant bigrams on an annual level that then have the potential to also closely match the within-year timing of the events.

Our results are in line with recent arguments by Kelly et al. (2022, 2024), who praise the "virtue of complexity" in certain prediction tasks, in their case, the prediction of asset returns. They emphasize that "simple" models – models in which the number of observations largely exceeds the number of features – understate the complexity of the true functional relationship to be identified. In a similar vein, Didisheim et al. (2024) show that richer models with thousands of parameters perform better out-of-sample in asset pricing settings, contrary to conventional econometric wisdom.

Table 2.2: **LASSO and SVR on the training data.** Performance statistics of LASSO and SVR regressions. R^2_{annual}, R^2_{monthly}, R^2_{within} and β_{within} are defined in Equations (2.15) to (2.19). # non-zero denotes the number of non-zero coefficients in the LASSO regression, $corr_{\text{monthly}}$ the correlation coefficient between $(\hat{g}_{y,m})_{y,m}$ and $(g_{y,m})_{y,m}$, and t is the t-statistic of the test of the hypothesis $\beta_{\text{within}} = 1$. In the LASSO regressions we use $\lambda = 0.08$. In the SVR we use $C = 0.01$ and $\varepsilon = 0.00025$.

LASSO	Time Series	# non-zero	R^2_{annual}	R^2_{monthly}	$corr_{\text{monthly}}$	R^2_{within}	β_{within}	t
1	Production Index	51	0.927	0.196	0.444	0.083	0.888	-0.696
2	CPI	34	0.968	0.395	0.630	0.051	0.472	-3.961
3	Unemployment	49	0.964	0.186	0.433	0.068	0.666	-1.960
4	PPI	44	0.948	0.266	0.519	0.097	0.971	-0.132
5	PI - Manufactoring	52	0.925	0.182	0.427	0.075	0.841	-0.949
6	PI - Energy	39	0.972	0.036	0.201	0.008	0.417	-3.953
7	PI - Consumer Goods	40	0.926	0.043	0.253	0.006	0.238	-2.532
8	Dividends	35	0.924	0.064	0.261	0.020	0.570	-3.032
9	Employment	43	0.951	0.282	0.531	0.033	0.483	-4.864
10	Women Employment	35	0.981	0.361	0.619	0.009	0.149	-9.268
11	Confidence Index	39	0.970	0.062	0.316	0.013	0.253	-7.029
12	Sales	42	0.970	0.114	0.351	0.017	0.387	-4.780

SVR	Time Series	R^2_{annual}	R^2_{monthly}	$corr_{\text{monthly}}$	R^2_{within}	β_{within}	t
1	Production Index	1.0	0.256	0.506	0.133	1.034	0.231
2	CPI	1.0	0.467	0.684	0.094	0.919	-0.397
3	Unemployment	1.0	0.283	0.532	0.159	1.039	0.312
4	PPI	1.0	0.266	0.517	0.074	0.870	-0.490
5	PI - Manufactoring	1.0	0.243	0.493	0.128	1.064	0.407
6	PI - Energy	1.0	0.062	0.249	0.020	0.918	-0.287
7	PI - Consumer Goods	1.0	0.178	0.421	0.085	1.011	0.076
8	Dividends	1.0	0.091	0.303	0.026	0.861	-0.456
9	Employment	1.0	0.381	0.618	0.109	0.858	-1.186
10	Women Employment	1.0	0.565	0.752	0.109	0.814	-1.289
11	Confidence Index	1.0	0.221	0.470	0.082	0.891	-0.733
12	Sales	1.0	0.176	0.421	0.039	0.698	-1.709

Table 2.2 shows performance statistics on the annual datasets, as well as on the monthly datasets. As mentioned above, the R^2_{annual} is virtually equal to 1 for SVR and very high (between 0.9 and 1) for LASSO regressions on the training dataset. The most important statistic for our application is R^2_{monthly}, because we are interested in explaining as much as possible of the monthly variation. Table 2.2 also reports correlation coefficients between the predicted time series $(\hat{g}_{y,m})_{y,m}$ and the actual time series $(g_{y,m})_{y,m}$. We find that the correlation varies between 0.2 (for the *Energy Price Index*) and 0.63 (for the *Consumers Price Index*) in the case of LASSO and between 0.25 (again for the *Energy Price Index*) and 0.75 (for *Women Employment*) in the case of SVR. With the exception of the *Producer Price Index*, where the performance for SVR and LASSO is about the same, SVR yields monthly predictions that are consistently more strongly correlated with the actual time series than LASSO.

Since positive $R^2_{monthly}$s could be entirely caused by the predictors matching the (in-sample) annual variation, we also investigate if the predictors also explain some of the *within-year* variation in the monthly time series. We find that all R^2_{within} are positive and, for SVR, average at about 10%. These findings provide empirical support for the hypothesis that word frequencies that are correlated on the annual frequency with economic time series, also correlate with the monthly counterparts. Reassuringly, we also find that the regression coefficients β_{within} are close to 1 and, in case of SVR, we cannot reject the hypothesis that $\beta_{\text{within}} = 1$ for any of the twelve time series. LASSO statistics are similar, however, with lower R^2_{within} and β_{within} further away from one.

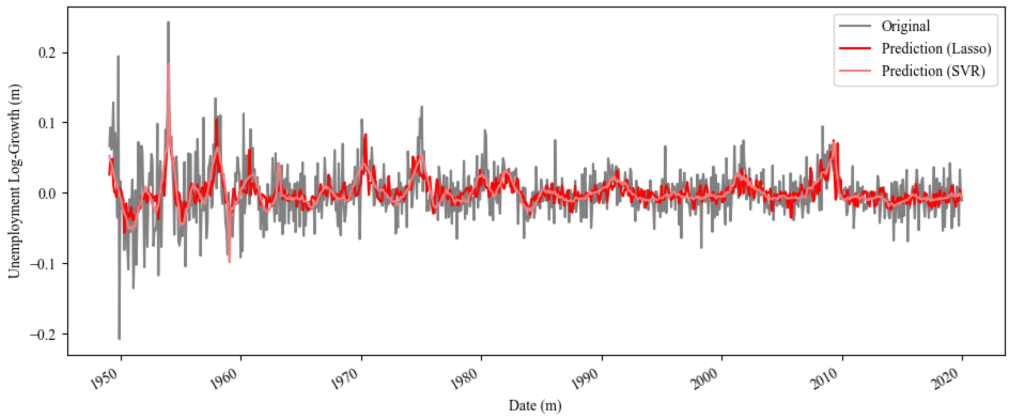

Figure 2.4: **Text-based monthly unemployment.** Time series plot of *Unemployment*. The original monthly time series is plotted in gray, the time series predicted by LASSO is plotted in red, and the one predicted by SVR is plotted in orange.

Figure 2.4 shows the two predicted monthly time series (in red for LASSO and orange for SVR), together with the original time series (in gray) for the log growth rate of *Unemployment*. Both predicted time series capture the major swings in the original time series, which is not surprising however, given the good in-sample fit. We find that the time series predicted by LASSO exhibits a high within-year volatility, relative to SVR, reflecting the high within-year volatility of the original time series. This is, again, not surprising, as LASSO selects only a small number of bigrams, which themselves have volatile frequency time series. The time series predicted by SVR is more strongly auto-correlated. However, the higher monthly- and within-R^2 of the SVM predictor indicates that the strong within-year variation of the LASSO predictor does not always accurately capture the within-year variation of the original time series. In

Section 2.4.2, we analyze more thoroughly, if the predictors capture specific events well on the monthly frequency. Still, one may be interested in producing a time series with moments, including auto-correlation, that are similar to those of the original time series.

We also combine the predicted time series of LASSO and SVR by weighting each with a factor of 0.5 and summing them up.[5] For the resulting new monthly mixed time series we again compute performance statistics and present them in Table 2.3. Most of the time series exhibit a R^2_{monthly} close to the value for SVR. Intriguingly, for *PPI* and *Dividends*, the ensemble approach delivers even higher R^2_{monthly} than LASSO or SVR individually. Keeping the monthly fit stable while at the same time increasing the within-year variation makes the ensemble approach our preferred approach for the applications in the upcoming sections.

Table 2.3: **Weighted results of SVR and LASSO regressions.** This table reports model parameters and results of LASSO and SVR regressions from Table 2.2 equally weighted (50% / 50%) with key statistics. For equations see Chapter 2.2.5. $coef$ are coefficients of the within-model regression and t are the respective t-test statistics with $\beta_{within} = 1$.

MIX	Time Series	R^2_{annual}	R^2_{monthly}	$corr_{\text{monthly}}$	R^2_{within}	β_{within}	t
1	Production Index	0.986	0.249	0.505	0.133	1.193	1.156
2	CPI	0.993	0.465	0.682	0.091	0.870	-0.670
3	Unemployment	0.993	0.262	0.514	0.132	1.036	0.268
4	PPI	0.989	0.289	0.541	0.111	1.201	0.764
5	PI - Manufactoring	0.986	0.237	0.492	0.127	1.212	1.224
6	PI - Energy	0.994	0.059	0.244	0.017	0.815	-0.878
7	PI - Consumer Goods	0.985	0.136	0.369	0.043	0.769	-0.736
8	Dividends	0.985	0.093	0.306	0.031	0.945	-0.220
9	Employment	0.990	0.360	0.602	0.087	0.887	-0.887
10	Women Employment	0.996	0.513	0.716	0.049	0.499	-3.218
11	Confidence Index	0.994	0.180	0.425	0.046	0.636	-2.545
12	Sales	0.993	0.166	0.408	0.034	0.677	-2.154

2.4.2 Plausibility of Results

Besides analyzing the quality of the predicted monthly time series on statistical grounds, we also investigate the plausibility based on economic arguments. To do so, we plot each original monthly time series together with its text-based version (using the ensemble approach) in Figure 2.5. Ideally, the text-based data (red) should match the original data (grey). Indeed, we observe large spikes in the original data that match exactly with the timing of spikes in the predicted time series. For example, a soaring *unemployment* rate in 1954 is spot on replicated, as well

[5] Combining multiple econometric estimation methods can significantly improve results. For a literature review on ensemble approaches in financial time series estimations see Albuquerque et al. (2022).

as a drop in growth rates of the $Production\ index$ in the 1940s or a rise in PPI, again in the 1940s. Further evidences come from the timing of recessions, which matches well the predicted macro time series.

Table 2.3 shows that the monthly fit for *PI - Energy* lacks behind the other time series, followed by *Dividends* and *PI - Consumer Goods.* A potential reason for the difference could be autocorrelation. All three time series exhibit negative or close to zero serial correlation. News can be assumed persistent concerning major topics, making our model prediction's fit to positively autocorrelated data more favorable. Still, as autocorrelation on the annual frequency is not a reliable predictor for autocorrelation on the monthly frequency, deciding whether the monthly news-based prediction of a time series with our approach works well needs to be evaluated case by case.

Besides analyzing the time series behavior of the predicted series, our text-based approach has the advantage that we can give an economic interpretation to the predicted series by analyzing the bigrams that have been assigned large coefficients from the LASSO or SVR algorithm. Remember that we let the algorithm decide which bigrams to pick, making our approach a priori agnostic about bigram choices.

Table 2.4 shows bigrams with the largest positive and negative coefficients, for LASSO and SVR, for the three times series *Production Index*, *CPI* and *Unemployment.* A declining *Production Index* connects to bigrams such as "short,sell", "increas,unemploy", and "present,busi", see LASSO results in Table 2.4. Increased short-selling activities regularly gained public interest in the US, especially in times of troubled stock markets, making it a very plausible negative macro indicator. The same is true for increasing unemployment, which points at challenges for corporations and therefore reduced production. While clearly context-related, "present business" is more difficult to interpret directly as a negative bigram, since it does not indicate a positive or negative signal for economic growth in general and industrial production in particular.

Results for SVR are more difficult to interpret, apart from "short,sell," which we again find among the bigrams with the largest coefficients. The difference between SVR and LASSO becomes apparent, when comparing the magnitude of the coefficients. LASSO coefficients are about 100 times bigger in magnitude than SVR coefficients. This indicates that the top SVR bigrams must be interpreted with care since there is a very large number of bigrams with small weights, making the occurrence of difficult-to-interpret terms more likely - they could be due

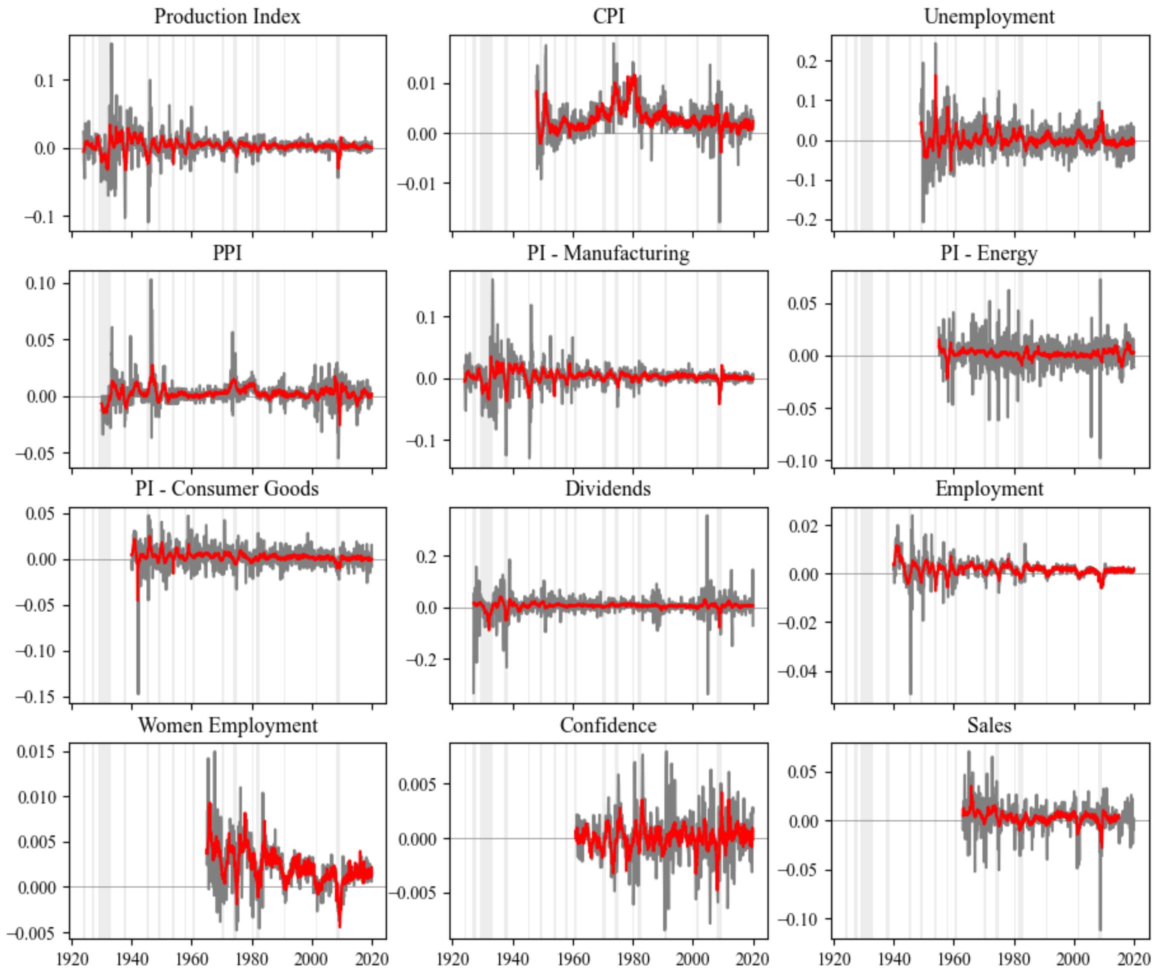

Figure 2.5: **Text-based vs. regular macro time series.** This Figure shows monthly text-based macro time series (log growth rates) in red and the measured time series in grey. Grey shaded areas are US recessions according to NBER definition.

to random in-sample correlation without being economically meaningful - and less important, since their impact is masked by the variation in hundreds of other bigram frequencies. We thus put more focus on the interpretations of the bigrams selected by LASSO.

For *CPI* the most important bigram with a positive coefficient is "price,increas," pointing at high inflation. Other terms on the negative side are puzzling as e.g. "girl,go." But with less than half of absolute coefficient size, results are clearly driven by bigrams with positive coefficients. Interestingly, bigram which are assigned large coefficients by SVR are related to oil as "gasoline,ration" or "suppli,price." For *Unemployment*, we find plausible bigrams such as "rise,unemploy" among those with the highest absolute weights.

Table 2.4: **Bigram frequencies.** This table reports bigrams with the five highest and lowest coefficients in lasso and support vector regressions. Minimum occurrences overall years must be 3600 times. $freq$ is the overall number of occurrences in the data. $coef$ is the coefficient of the Lasso/Support Vector regression. Top five bigrams for the remaining macro data can be found in Appendix Table A.2. We provide additional context examples from the NYT in Appendix Table A.1 for a selection of bigrams that are less intutive.

Production Index	bigram (+)	coef (+)	freq (+)	bigram (-)	coef (-)	freq (-)
LASSO						
1	ahead,year	0.106872	4371	short,sell	-0.162886	5828
2	state,ticket	0.092784	4873	never,worn	-0.154956	4770
3	end,end	0.092291	18983	increas,unemploy	-0.121399	4936
4	advertis,caution	0.082868	6018	present,busi	-0.086433	4681
5	avail,everi	0.069508	6928	influenc,peopl	-0.081708	3620
SVR						
1	candid,unit	0.001549	5657	island,open	-0.001500	7279
2	state,ticket	0.001524	4873	young,young	-0.001454	4029
3	candid,state	0.001509	7413	chines,armi	-0.001434	3839
4	candid,repres	0.001462	3729	intern,settlement	-0.001419	7702
5	vote,governor	0.001351	4640	short,sell	-0.001372	5828

CPI	bigram (+)	coef (+)	freq (+)	bigram (-)	coef (-)	freq (-)
LASSO						
1	price,increas	0.168063	62271	girl,go	-0.068375	7517
2	organ,individu	0.154819	6396	end,tuesday	-0.018867	3700
3	call,product	0.105025	3645	team,match	-0.009613	7455
4	treasur,island	0.095580	5778	without,addit	-0.006844	6833
5	market,organ	0.085202	4092	includ,power	-0.005276	5311
SVR						
1	treasur,island	0.000519	5778	citi,better	-0.000450	3649
2	mover,system	0.000467	3683	convent,visitor	-0.000423	5212
3	monday,monday	0.000448	6596	girl,go	-0.000413	7517
.	.	.	.	.	.	.
7	gasolin,ration	0.000412	4792	time,weather	-0.000388	11202
8	suppli,price	0.000412	3642	vic,sade	-0.000387	4015

Unemployment	bigram (+)	coef (+)	freq (+)	bigram (-)	coef (-)	freq (-)
LASSO						
1	pool,gym	0.259468	8194	seventh,fleet	-0.088779	4162
2	offer,low	0.129566	8870	also,arrang	-0.064519	6777
3	news,farm	0.129443	4637	glass,menageri	-0.051456	4119
4	rise,unemploy	0.120603	4903	citi,set	-0.047276	6371
5	product,declin	0.106865	5164	green,park	-0.046784	6398
SVR						
1	schuyler,west	0.001012	14510	right,day	-0.000789	4546
2	news,farm	0.000952	4637	old,vic	-0.000784	5390
3	homelik,hotel	0.000916	4925	seventh,fleet	-0.000773	4162
4	pool,gym	0.000855	8194	pioneer,new	-0.000740	3699
5	avail,next	0.000825	4645	annul,associ	-0.000738	4456

Table A.2 in Appendix A.2 shows bigrams with large coefficients also for the other nine time series. For *PPI*, the most important bigrams are "rise,price" (with a positive coefficient) and "fall,price" (with a negative coefficient). While "rise,unemploy" has a high positive coefficient for *Unemployment*, and gets a large negative coefficient for *Employment*. For *Dividends*, we find “increas,dividend” as the bigram with the seventh largest coefficient and for *sales*, we find “product,declin” and “econom,slowdown” among those with the most negative coefficients.

Many more examples of plausible bigram choices can be found in the Appendix Table A.2. Even though there are also many dubious bigrams with large coefficients, we conclude that the penalized regression is able to select economically meaningful bigrams, which correlate with low and, plausibly, also high frequency variation in the considered time series.

2.4.3 New Monthly Time Series

We use the Macrohistory Database by Jordà et al. (2017) to retrieve annual macroeconomic data.[6] Table 2.5 presents the list of all variables used. It also shows if and how we calculate growth rates or first differences, and if we treat the time series as point-in-time measured or time-aggregated. The database also includes indicators and informational data (such as ISO country codes), which we ignore. We focus on extrapolating data from Jordà et al. (2017), as their data represent major macroeconomic time series. We refrain from estimating text-based return data from Jordà et al. (2019), see Table A.3 in Appendix A.2 for an overview, as returns are easily available on high frequencies. Additional results from Jordà et al. (2019, 2021) can be found in the Appendix Table A.2 and Figure A.1.

In contrast to the series analyzed in Section 2.4.2, for most of the time series considered here no reliable monthly data are available to evaluate our results. Hence, we rely on a twofold validation process: (1) visual inspection of the predicted monthly time series and (2) interpretation of bigrams with large coefficients.

The predicted monthly text-based time series are shown in Figure 2.6. We observe that monthly growth rates fluctuate in line with expectations. For example, in the monthly *Gross Domestic Product (GDP)* time series, we find growth rates turning positive right after the Great Depression in April 1933. It also shows negative export rates during recession months.

We again interpret bigrams that got assigned large relative coefficients in our model. Again, we mainly interpret LASSO bigrams as they carry more weight than SVR bigram coefficients. Top five bigrams by time series and regression model type can be found in Table 2.4 in Appndix A.2. For *Wage* the LASSO regression allocates high negative weight to “wage,cut”. As public interest is high, wage cuts by large corporations often become public in due time after the

[6]The database is openly accessible at `www.macrohistory.net/database/`.

Table 2.5: **JST Marcohistory Database variables.** This table reports variables used from Jordà et al. (2017, 2019, 2021). *Type* differentiates variables between "PIT" (point-in-time), "TAG" (time aggregated), and "DUM" (dummy). *Stat* defines how the series is transformed to become stationary with "LNR" (log returns) or "DIF" (first difference). Variables not used for our analysis are marked with N/A.

Variable	Description	Type	Stat
year	Year	N/A	N/A
country	Country	N/A	N/A
iso	ISO 3-letter code	N/A	N/A
ifs	IFS 3-number country-code	N/A	N/A
pop	Population	PIT	DIF
rgdpmad	Real GDP per capita (PPP, 1990 Int$, Maddison)	TAG	LNR
rgdpbarro	Real GDP per capita (index, 2005=100)	TAG	LNR
rconsbarro	Real consumption per capita (index, 2006=100)	TAG	LNR
gdp	GDP (nominal, local currency)	TAG	LNR
iy	Investment-to-GDP ratio	TAG	LNR
cpi	Consumer prices (index, 1990=100)	TAG	LNR
ca	Current account (nominal, local currency)	TAG	N/A
imports	Imports (nominal, local currency)	TAG	LNR
exports	Exports (nominal, local currency)	TAG	LNR
narrowm	Narrow money (nominal, local currency)	PIT	LNR
money	Broad money (nominal, local currency)	PIT	LNR
stir	Short-term interest rate (nominal, percent per year)	PIT	DIF
ltrate	Long-term interest rate (nominal, percent per year)	PIT	DIF
hpnom	House prices (nominal index, 1990=100)	TAG	LNR
unemp	Unemployment rate (percent)	TAG	LNR
wage	Wages (index, 1990= 100)	TAG	LNR
debtgdp	Public debt-to-GDP ratio	TAG	LNR
revenue	Government revenues (nominal, local currency)	TAG	LNR
expenditure	Government expenditure (nominal, local currency)	TAG	LNR
xrusd	USD exchange rate (local currency/USD)	N/A	N/A
tloans	Total loans to non-financial private sector (nominal, local currency)	PIT	LNR
tmort	Mortgage loans to non-financial private sector (nominal, local currency)	PIT	LNR
thh	Total loans to households (nominal, local currency)	PIT	LNR
tbus	Total loans to business (nominal, local currency)	PIT	LNR
bdebt	Corporate debt (nominal, local currency)	PIT	LNR
crisisJST	Systemic financial crises (0-1 dummy); included since R5	DUM	N/A
crisisJST_old	Systemic financial crises (0-1 dummy); as coded in all prior releases (R1 – R4)	DUM	N/A
peg	Peg dummy	DUM	N/A
peg_strict	Strict peg dummy	DUM	N/A
peg_type	Peg type (BASE, PEG, FLOAT)	DUM	N/A
peg_base	Peg base (GBR, USA, DEU, HYBRID, NA)	N/A	N/A
JSTtrilemmaIV	JST trilemma instrument (raw base rate changes)	N/A	N/A
housing_tr	Housing total return, nominal. r[t] = [[p[t] + d[t]] / p[t-1]] - 1	TAG	DIF
housing_capgain_ipolated	1 if housing capital gains and total returns interpolated e.g. wartime	N/A	N/A
housing_capgain	Housing capital gain, nominal. cg[t] = [p[t] / p[t-1]] - 1	TAG	DIF
housing_rent_rtn	Housing rental return. dp_rtn[t] = rent[t]/p[t-1]	TAG	DIF
housing_rent_yd	Housing rental yield. dp[t] = rent[t]/p[t]	TAG	DIF
lev	Banks, capital ratio (in %)	PIT	DIF
ltd	Banks, loans-to-deposits ratio (in %)	PIT	DIF
noncore	Banks, noncore funding ratio (in %)	PIT	DIF

decisions to cut wages. Among the bigrams with positive coefficients there are e.g. "increas,cost" or "product,manag." The former could appear in the context of rising worker costs of firms when wages are increasing. For example, the NYT writes on August 9, 1947: "Higher wages do not necessarily **increase costs** except perhaps in the case of marginal industries. Efficient industries can afford higher wages precisely because their unit costs are low."

Employment is generally a big concern of people and therefore frequently covered in news. Full employment is considered a strong positive indicator for an economy. The bigram "full,employ" gets allocated a high positive weight for *Total loans to households*. On the negative coefficient side, the apparent role of banks in loan granting is reflected by terms as "pro-

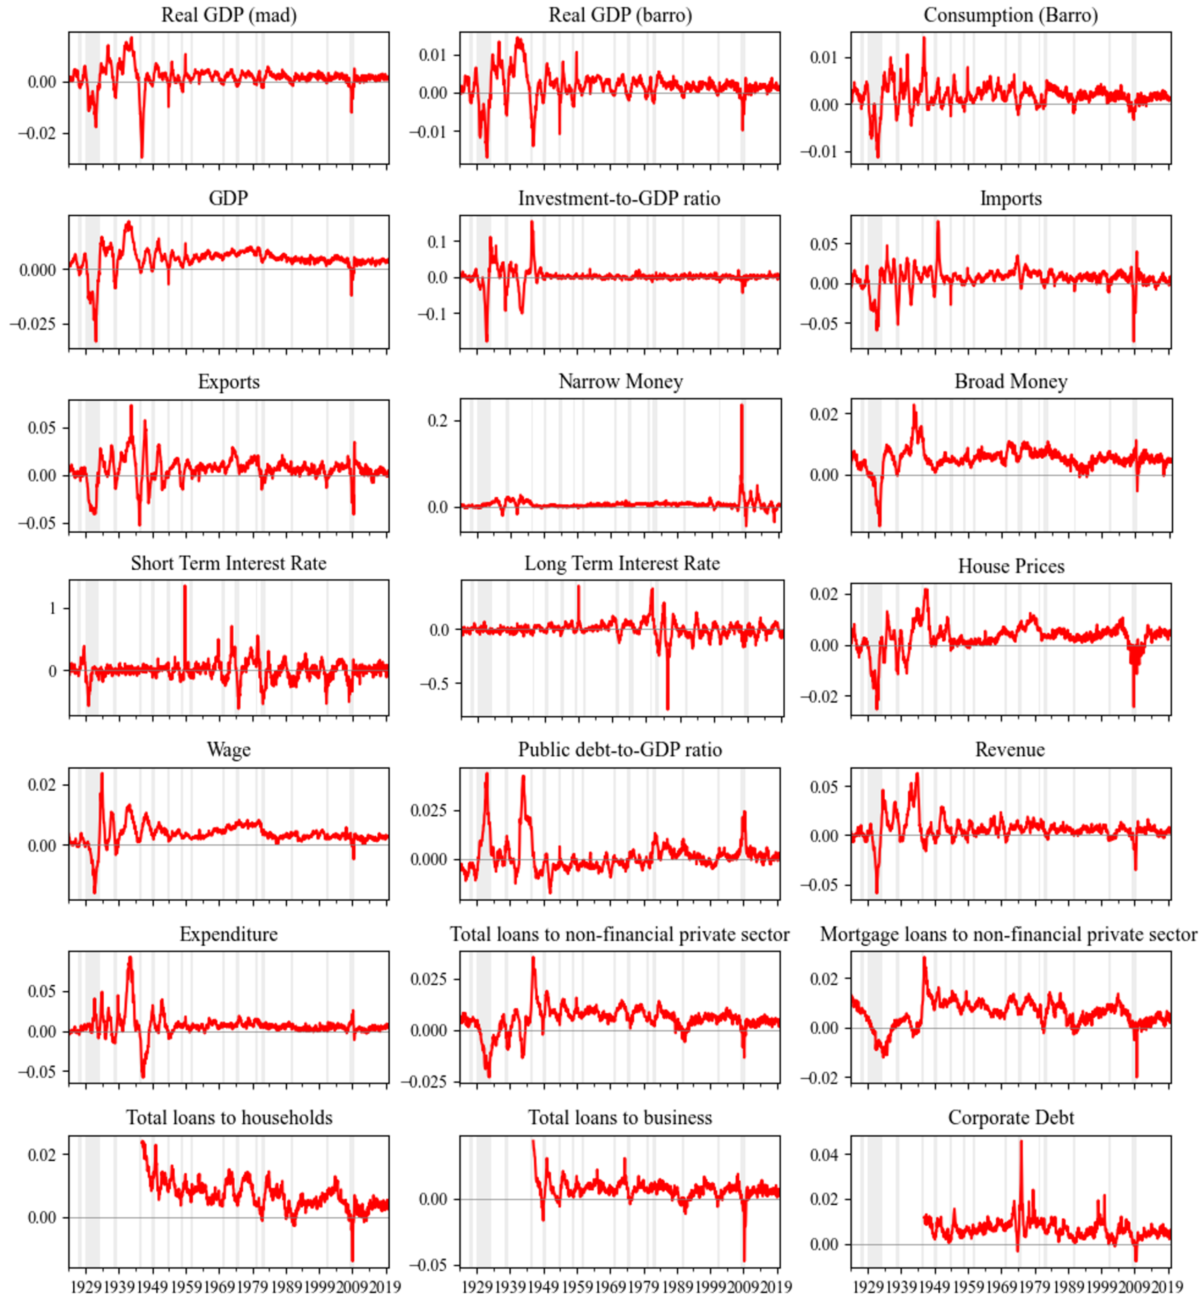

Figure 2.6: **Text-based macro time series.** This figure shows monthly text-based macro time series. For details see Table 2.5. Grey shaded areas are US recessions according to NBER definition.

gram,expand" or "bank,charg". *Short Term Interest Rates* are driven by bigrams related to the monetary policy by central banks': Bigrams with negative coefficients include "declin,first" and "point,cut", and those with positive coefficients include "point,increase."

Another relevant topic is international trade activity, measured by the National Current Account which is set up of *Imports* and *Exports*. When US *Exports* rise, we find heightened usage of bigrams such as “fee,char” or “order,larg”. Both bigrams belong in the context of trade activity, as e.g. fees charged for transportation at boarders are crucial in trading decisions. The NYT titles on November 23, 1966: “Big Sale to Chile Cuts Contracts – Indonesia Also Hints at Plan to *Order Large* Shipment”. For increasing *Imports* “put,full”, “use,great”, or “receiv,paymen” receive large coefficients.

Nevertheless, some macro data also deliver bigrams that are less (or not at all) intuitive at first sight. Applying our method to *Broad Money* yields the following top five positively connoted bigrams: “american,land”, “john,chamberlain”, “servic,local”, “car,purchas”, and “minist,inform”. While the latter three appear somewhat reasonable, finding the artist John Chamberlain as a major driver to increase money supply is surprising. However, after researching the NYT, we find the mentions of John Chamberlain probably refers to a political and economical journalist with the same name, writing for the NYT, WSJ, and other newspapers. Similarly, economic interpretations might also be hidden behind other dubious bigrams, which only reveal themselves upon closer inspection.

2.5 Applications

2.5.1 The Consumption-CAPM

The Consumption-CAPM, first suggested by Lucas (1978), assumes a representative agent with time-additive isoelastic preferences. Solving the optimization problem of the agent, the consumption-savings problem, results in the central testable implication of the model, known as the Consumption Euler Equation:

$$E\left[\beta \exp(-\gamma g_{C,t+1}) R_{t+1}\right] = 1, \tag{2.20}$$

where g_C denotes consumption growth and β and γ are preference parameters quantifying the representative investor's subjective time preference rate and her risk aversion, respectively. R denotes the gross return on an arbitrary asset. Looking at return differences yields the following version of the Euler Equation:

$$E\left[\beta \exp(-\gamma g_{C,t+1}) R^e_{t+1}\right] = 0, \tag{2.21}$$

where R^e denotes an asset's excess return relative to the risk-free rate. The central implication of the model follows from solving the equation for $E[R^e]$: The expected excess return on any asset only depends on the asset return's covariance with consumption growth.

Empirical investigations of the model usually result in unrealistically high estimates of the risk aversion coefficient γ, which is at odds with arguments put forward by Mehra and Prescott (1985) and others. This phenomenon is know as the *equity premium puzzle*. While the exact parameter estimates depend on the choice of the consumption measure and the test asset, we follow the standard approach and use consumption of non-durable goods and services and the return on the aggregate US stock market portfolio in our estimation.

As a benchmark, we use annual consumption data of the US between 1930 and 2019 from the Bureau of Economic Analyses. We aggregate the two quantity indices of consumed non-durable goods and services by Fisher aggregation, using the respective price indices, and divide by the US population to get per capita consumption levels. We then calculate annual growth rates. The resulting time series is time-aggregated. Kroencke (2017) points out that the standard test, using the consumption-growth time series in combination with stock returns, which are point-in-time, is inconsistent and can lead to biased results.

We compare these benchmark results to an estimation using monthly data, generated via our text-based approach. Figure 2.7 shows the monthly text-based consumption time series. We find phases of low consumption growth rates coincide with US recessions, which is exactly what we would expect. Recall that we train the model on annual growth rates, so major swings should be correctly covered. However, we not only match major swings but correctly predict recession months supporting the plausibility of the text-based monthly time series. For example, the beginning of the first Oil Crisis of 1973 is correctly predicted by our method. Furthermore, our text-based time series captures the drastic plunge in household consumption in October 2008, the onset of the Great Financial Crisis.

In addition, the bigrams with major regression weights are economically intuitive. For example, bigrams with the largest positive weights are 'project approv', 'citi bill', and 'receiv payment', and those with the largest negative weights are 'advertis may', 'short sell', and 'increase rate'. Approval of e.g. governmental project may boost consumption spending as well as receiving payments. "citi bill" might relate to discussed citizenship bills in the post WWII phase or to positions of the Citizen Union to bills. For the negatives, "short sell" is a common indicator of time series models for declines during the Great Financial Crisis.

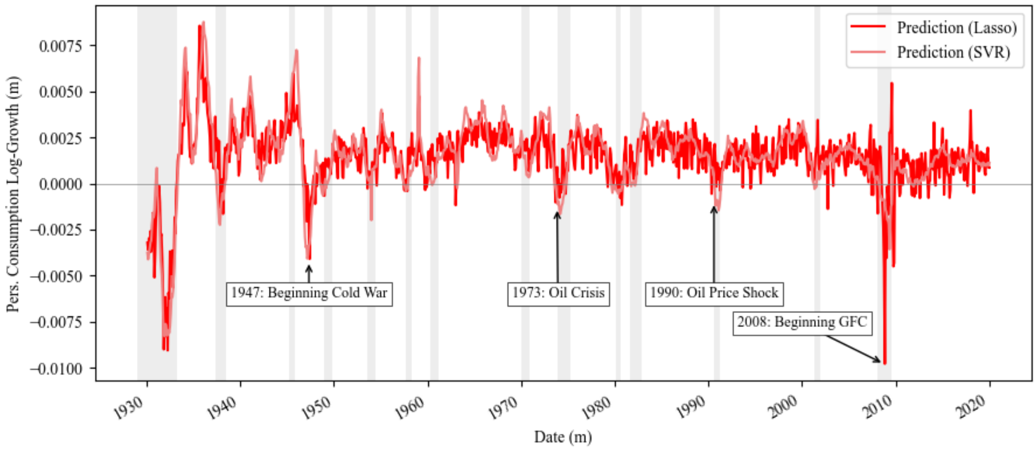

Figure 2.7: **Consumption.** This is a time series plot of personal consumption expenditures (non-durable goods and services). Grey shaded areas are US recessions according to NBER definition.

We follow the empirical design of Kroencke (2017) and only estimate γ under the assumption that $\beta = 0.95$. More precisely, we run a Generalized Method of Moments (GMM) estimation, using the sample analogue of Equation (2.21) as moment condition. Column 1 in Table 2.6 shows the usual result: The risk aversion coefficient is close to 40 and significantly larger than 10.

Table 2.6: **Estimation of the CCAPM.** Estimates of the risk aversion coefficient γ from a GMM estimation using Equation (2.21) as moment condition. Standard errors are shown in parentheses. The excess return on the aggregate stock market relative to the risk-free rate is used as a test asset return and downloaded from Kenneth French's website. In columns 1 and 3, we use annual consumption of nondurable goods and services as a consumption proxy. In columns 2 and 4, we use monthly text-based consumption. For the estimates in columns 1 and 2, we use the end-of-period timing convention, while for those in columns 3 and 4, we use the beginning-of-period timing convention. The

	end-of-period		beginning-of-period	
	annual	monthly	annual	monthly
γ	39.92	244.46	21.90	225.85
	(13.41)	(94.63)	(8.93)	(87.50)

Column 2 shows the estimate of γ when re-estimating the model using our monthly consumption time series. Here, we extrapolate annual quantity and price indices and aggregate these to a monthly consumption growth time series. We find that the use of the monthly text-based measure does not solve the equity premium puzzle. In contrast, the estimate is much larger than with annual data.

As pointed out earlier, the time-aggregation bias can lead to biases when using annual time-aggregated consumption growth in conjunction point-in-time measures asset returns. Using monthly consumption growth mitigates this issues, but there is still time-aggregation bias, since the monthly consumption growth measure is also time-aggregated. Some studies (see, e.g., Yogo, 2006) argue, that a "beginning-of-period" timing convention, i.e., aligning consumption growth of year $t+1$ over year t with returns over year t (instead of $t+1$) could be accurate. While one approach is conceptually as false as the other, beginning-of-period timing alludes to the idea that the agent decides upon all consumption expenditures in a month or year at the beginning of the period.

Columns 3 and 4 show estimates of γ when the estimation is performed using a beginning-of-period timing convention. The results are very similar to those in columns 1 and 2, so the message does not change.

Due to the joint-hypothesis nature of the empirical tests, interpretations of the estimation results are difficult. A high estimate of the risk aversion coefficient could on the one hand indicate that the model framework is excessively simplified. Especially the notion that only contemporaneous changes to consumption matter for asset prices, a feature of time-additive models, may be false. Extensions with recursive preferences have the feature that also news concerning future consumption growth and uncertainty about consumption growth are priced (see Epstein and Zin, 1989; Bansal and Yaron, 2004). The fact the annual consumption works "better" than monthly consumption growth, could point to the fact, that long-term consumption growth is more accurate that contemporaneous consumption changes.

Alternatively, consumption of non-durables and services may not be an appropriate measure of the crucial drivers of the utility function of the marginal investor. For example, Yogo (2006) emphasizes the important role of durable good consumption, which is often omitted from the analysis because the timing of expenditure for and consumption of these goods does not coincide. Dittmar et al. (2024) point out that various nondurables and services are not substitutable for one another, meaning that the different types of consumption appear independently in the utility function. They emphasize the important role of energy consumption. Ait-Sahalia et al. (2004) examine the role of luxury goods and suggests that these are more elastic compared to basic goods and thus better suited for such analyses. Savov (2011) and Kroencke (2017) argue that consumption data is excessively filtered, leading to the loss of important variation, and propose alternative measures of consumption. Malloy et al. (2009) and Vissing-Jørgensen (2002) argues that the average consumer does not match the marginal investor and suggests survey-based consumption of stockholders as a better measure.

2.5.2 Elasticity of Intertemporal Substitution in Consumption

Substituting a risk-free asset return in the conditional version of Equation (2.20) and taking the logarithm yields

$$r^f_{t+1} = -\log(E_t[\beta \exp(-\gamma g_{C,t+1})]) \approx -\log(\beta) + \gamma E_t[g_{C,t+1}], \tag{2.22}$$

Where r_f denotes the logarithmic risk free rate. In the time-additive isoelastic model, γ represents the risk-aversion coefficient and, at the same time, the inverse of the intertemporal elasticity of substitution (IES) in consumption, see the discussion in Hall (1988) and Epstein and Zin (1989). Hall (1988) formulates Equation (2.22) as a regression equation

$$g_{C,t+1} = const + \gamma^{-1} r^f_{t+1} + \epsilon_{t+1} \tag{2.23}$$

to estimate the IES γ^{-1}. Intuitively, the IES quantifies the willingness of the representative consumer to intertemporally shift consumption in response to a varying investment opportunity set, quantified here by the risk-free interest rate. As discussed in Section 2.2.2, Hall (1988) uses annual data and accounts for the fact that annual consumption growth is time-aggregated while interest rates are point-in-time. To overcome this problem, he time-aggregates interest rates using Equation (2.3).

We replicate his study using annual data between 1930 and 2019 and, as a comparison, use our monthly text-based consumption time series and run the regression using monthly time series. While the two designs have to give the same results in large samples, additional within-year common variation in monthly consumption growth and interest rates could be masked on the small sample, leading to small sample point estimates, deviating from those of the larger monthly sample.

Table 2.7 shows that the results are very consistent across the two data frequencies. We estimate an IES of 0.0138 on the annual sample and 0.0134 on the monthly sample. Our estimates are more precise when using the monthly sample, since the sample size increases from 90 to 1080. However, we take our monthly text-based consumption measure as given and do not account for estimation uncertainty in the extrapolation here.

Table 2.7: **Estimation of the IES.** We estimate γ^{-1} using the condition in Equation (2.23), using an Ordinary Least Squares (OLS) regression. Standard errors are shown in parentheses. r_f is the logarithmic interest rate downloaded from Kenneth French's website. Consumption growth g_C is the logarithmic growth rate of consumption of non-durable goods and services. Column 1 shows estimation results when using annual data, and column 2 shows estimation results when using montly data. The sample ranges from 1930 to 2019.

	annual	monthly
const	0.0177 (0.0032)	0.0014 (0.0001)
γ^{-1}	0.0138 (0.0719)	0.0134 (0.0209)

Our findings resemble those of Hall (1988), who argues that "all estimates [of the EIS] are small. Most of them are quite precise, supporting the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero." Whether the IES is actually that small remains a matter of debate. Simplifying assumptions are necessary to derive Equation (2.23), including a first-order Taylor approximation in Equation (2.22). However, further assumptions, e.g. regarding the choice of proxies for consumption and return data, and the underlying utility function, are also necessary and may be overly simplistic. Numerous studies estimate the IES based on more general assumptions (see Thimme, 2017, for a survey of the literature). Our result can only rule out that Hall's results came about as an artifact of the annual data frequency.

2.6 Conclusion

In this Chapter 2 we developed a method to lift economic time series from lower to higher frequencies. The approach exploits the correlation of frequencies of select bigrams in NYT articles to changes in macroeconomic indicators in order to increase the frequencies of these indicators from annual to monthly. Using the penalized regression models SVR and LASSO, relevant bigrams are identified on the low-frequency data. Our approach shows how to account for characteristics of the respective time series, such as time aggregation and seasonal adjustment. The models then predict time series variation at the higher frequency. They prove capable of producing higher frequency data for general applications. Thus, our method represents a novel tool for researchers requiring higher frequency data for calibrations and empirical tests of economic models.

A potential problem with our approach becomes evident in Section 2.5: The extrapolated monthly time series exhibit estimation uncertainty. If they are assumed to be observable and used as such in empirical estimations and tests, the resulting standard errors can be too small. Future research could focus on developing a method to account for the uncertainty in the monthly predicted time series. Additionally, it might be of interest to define a rolling-window or otherwise "conditional" version of our approach to account for the different epochs and their particularities, especially with regard to the vocabulary used in newspapers during those times.

Chapter 3

Understanding Asset Pricing Factors

3.1 Introduction

The asset pricing factors SMB, HML, RMW, and CMA are omnipresent in financial market research and practice. They were designed with the goal of capturing a large chunk of the common time-series variation in stock returns, as the Arbitrage Pricing Theory (Ross, 1976) implies that a linear combination of such factors represents the stochastic discount factor. In other words, exposures to these four factors do a good job of explaining cross-sectional differences in average stock returns. Indeed, the pricing performance of the Fama and French (1993, 2015) model is decent for many portfolios. However, it remains unclear what fundamental economic shocks originally drive the innovations in the factors. In a sense, the model is context-free, or rather economics-free. While this may be acceptable for practitioners as long as the model works well empirically, economists seek to understand the fundamental causes behind the common variations of stocks along the boundaries of factor portfolios.

To explore this, we examine a catalog of days with high absolute factor returns. We sometimes refer to these movements as jumps, though they may also involve steady price shifts throughout the day. To understand the economic drivers of these movements, we analyze newspaper articles from the following day and systematically categorize the realized factor returns. We find that macroeconomic news, monetary policy news, and news about the earnings of individual companies and industries are the most common causes of high factor returns for all four factors. This is consistent with the findings of Baker et al. (2021, BBDS henceforth), who identified these three topics as major drivers of the market factor. To understand the differences between the factors, we construct topic-specific factors as the principal components of the returns on a set of auxiliary assets on days with news from a particular category, and we analyze how these fundamental factors correlate with the four factors SMB, HML, RMW, and CMA.

The macro factor we construct is almost perfectly correlated with HML. This result supports an extensive theoretical literature that interprets the value premium as compensation for systematic risk inherent in macroeconomic shocks (see, e.g., Dittmar and Lundblad, 2017; Kogan and Papanikolaou, 2014; Vassalou, 2003). CMA is most strongly correlated with our commodities factor, which essentially reflects shocks in oil supply and prices. This supports theoretical arguments by Dittmar et al. (2024) and Gao et al. (2022).

For SMB and RMW, the interpretation is less clear. SMB is, on the one hand, strongly correlated with the exchange rate factor (see Francis et al., 2008; Starks and Wei, 2003), which makes sense insofar as exchange rate risks are harder for small firms to hedge than for large firms. This is a key unifying, cross-industry characteristic of small firms and explains part of their common variation. Additionally, SMB is also strongly correlated with the "Unknown" factor, which is constructed from the principal component on days when journalists cannot identify a clear reason for the large movements in the stock market and often cite psychological factors (flight-to-quality, etc.). RMW is predominantly affected by shocks in individual companies or industries, which are represented in the long or short portfolio of the RMW factor. Examples of this include news about regulatory changes (often affecting software stocks) and sovereign military events (affecting defense stocks).

Our results are interesting in light of the recent discussion on whether the existence of factors points to rational explanations or arises from mispricing. Kozak et al. (2018) argue that mispricing in stocks can persist if the affected stocks exhibit common variation, as arbitrageurs would otherwise eliminate it. The existence of this common variation imposes a significant risk on the arbitrageurs' books, which only becomes systematic when the arbitrageur takes the position. Building on this idea, Daniel et al. (2020) construct mispricing factors, including one based on earnings announcements. These are naturally connected to profitability anomalies, as they are typically based on earnings relative to book values. Bali et al. (2023), Böll et al. (2024), and Frey (2023) also categorize many profitability anomalies as mispricing-based. Therefore, it is not surprising that we do not find a consistent economic channel for innovations in RMW.

Methodologically, our study is a direct extension of the study by BBDS, which focuses on the economic causes behind large returns of the aggregate market portfolio. Just like BBDS, we recruit several students who, under our supervision, search newspapers for information on the causes of stock market movements and categorize each day with extreme factor returns into one of 16 categories. In comparison to BBDS, who examine international markets, we only focus on the US stock market. Furthermore, we use only historical editions of the NYT for our analysis, while the coders of BBDS use several newspapers, predominantly relying on the WSJ in the US. As a sanity check, we had our students code days from the BBDS dataset and found widely consistent results.

Our dataset is available as supplementary material to this article and is a direct extension of the BBDS dataset. In addition to the categorizations made by human coders, it also includes codings by ChatGPT 3.5. We discuss the differences between the two and, similar to BBDS, find that the human coders produced higher-quality categorizations.

Both ChatGPT and the student coders were trained with a coders guide, which we adopted from BBDS with very few exceptions. Specifically, we use the same 16 category definitions, which has the advantage that our dataset directly extends that of BBDS. Another crucial advantage of the categories used by BBDS for us is that many of them correspond directly to economic channels that are considered candidates for explaining factor risk premia in the theoretical asset pricing literature.

One of the most prominent theories to explain factor risk premiums is macroeconomic risks. Bansal and Yaron (2004) and Dittmar and Lundblad (2017) emphasize the relevance of innovations in expected consumption. Vassalou (2003) looks at the role of news about Gross Domestic Product. Kogan and Papanikolaou (2014) advocate the importance of technology shocks. All of these pieces of news fall under the “Macro” category in our classification. We find that macroeconomic news can be particularly relevant for the economic grounding of the HML factor.

Exchange rate risks, according to Francis et al. (2008) and Starks and Wei (2003), are important drivers of returns in the cross-section. These are categorized under our “Exchange Rates” category. Our findings suggest that these risks are particularly relevant for the SMB factor. Dittmar et al. (2024) and Gao et al. (2022) find that energy consumption risks and oil price risks can help interpret return differences in the cross-section. The results in this Chapter 3 confirm the importance of this channel, especially for the CMA factor.

Other significant channels have been discussed in the asset pricing literature. Belo et al.

(2013) identify government spending as a key driver of returns in the cross-section; and "Government Spending" is categorized separately in our classification. An extensive literature is dedicated to the influence of monetary policy news on stock prices, see for example Ai and Bansal (2018), Rigobon and Sack (2004), and Wachter and Zhu (2022), which are coded under our "Monetary" category.[1] The intermediary asset pricing literature (He and Krishnamurthy, 2013; Adrian et al., 2014) highlights the importance of the risk-bearing capacity of financial intermediaries and their regulatory constraints. News about regulation is captured as a separate category in our classification. Importantly, our results cannot rule out that these channels are important drivers of factor risk premiums. Our approach merely identifies the individual topical factors that are most strongly correlated with the four characteristic-based factors.

While the channels discussed above are all risk-based, it could a priori also be the case that the factor risk premiums are behavioral phenomena. Kozak et al. (2018) explain that factor structures can also be found in an economy where all return differences in the cross-section are based on mispricing. In this case, innovations in a factor would be driven by firm-specific news (in our case, the "Corporate" category) or by news primarily affecting specific industries (such as our "Sovereign Military" category). Our results suggest that this could be the case for the RMW factor.

This chapter contributes to the literature on text-based interpretations of asset prices (see Loughran and Mcdonald, 2016; Gentzkow et al., 2019). More specifically, our paper contributes to research on news data used for asset pricing questions. For example Kerssenfischer and Schmeling (2024) build an extensive news database starting in 2002 and find that innovations through news account for up to 35% of price movements of stocks and bonds in the US and Euro area. Pettenuzzo et al. (2020) analyze daily dividend announcements and use these information to explain stock returns as well as jumps.

Closest to our paper is the work of Aleti and Bollerslev (2024) who estimate a stochastic discount factor (SDF) and use Dow Jones Newswires Archive to analyze innovations in high frequency. Moreover, Bybee et al. (2023) use media attention to 180 topics as input in IPCA to construct "narrative risk factors" and Jeon et al. (2022) investigate news directly related to large shifts in prices of individual stocks. Lopez-Lira (2019) quantifies firms' risk exposures by using

[1] There is also a large empirical literature on the impact of scheduled monetary policy announcements on asset price, see, e.g., Ben-Rephael et al. (2021), Bernanke and Kuttner (2005), Gürkaynak et al. (2005), Kroencke et al. (2021), Patton and Verardo (2012), Savor and Wilson (2013, 2014), and Schmeling and Wagner (2016).

topic analysis of 10-K reports. In contrast to these papers, our work focuses on the economic drivers behind characteristic-based factors of the Fama and French (1993, 2015) model, heavily used in financial research and industry.

Earlier research has tried other approaches to study the underlying economic forces of the Fama and French (1993) factors, most prominently Hahn and Lee (2006), Petkova (2006), Vassalou (2003), and link SMB to aggregate default risk and HML to economic growth expectations and the slope of the term structure. Aretz et al. (2010) find macroeconomic indicators to be influential for the size, book-to-market as well as the momentum factor.

The chapter continues as follows: Section 3.2 explains how we compile the data set. Section 3.3 provides a descriptive analysis of the data set and checks its plausibility. In Section 3.4, we construct topical factors from our data set and analyzes their relations with the four characteristic-based factors. Section 3.5 concludes.

3.2 Compiling the Data Set

3.2.1 Large Innovations in Characteristic-Based Factors

Methodically, we follow BBDS to create comparability between our and their datasets. BBDS analyze news about stock market movements on days when the aggregated market portfolio exhibits returns exceeding 2.5% in absolute terms. We extend this analysis to days when the four characteristic-based factors SMB, HML, RMW, and CMA exhibit large absolute returns. Thus, our dataset can be understood as a direct extension of the dataset by BBDS.

The cutoff of 2.5% is not necessarily appropriate for the four characteristic-based factors, since they are less volatile. The days categorized by BBDS include about 3.5% of the trading days in the sample period between 1900 and 2020. Since the time series for RMW and CMA only begin in 1963, our sample period is significantly shorter, covering only about 15,000 trading days. Looking at the 500 trading days with the largest absolute return per factor leads us to a similar coverage ratio of about 3.3%. The corresponding percentiles in the distributions of absolute factor returns vary slightly across the four factors, reflecting their different variances. The most volatile factor is HML, for which we set the cutoff at 1.4%. The days on which the absolute HML return exceeds 1.4% account for about 20% of the absolute factor variation and 50% of the quadratic variation. For SMB, we set the cutoff at 1.21%, for RMW at 0.96%, and for CMA at 0.89%, resulting in a share of 40-50% of the quadratic variation for each factor.[2] Figure 3.1 shows the four factor time series and highlights days with large innovations in each factor.

Although there are approximately 2,000 factor innovations to analyze in total, the actual number of days to be analyzed is lower because the factors often exhibit large innovations on the same days. Therefore, when considering only the number of days on which at least one factor shows a high absolute return, we need to examine 1,395 days. Additionally, on some of these days, the aggregated market return exceeds 2.5% in absolute value, meaning that BBDS have already categorized these days. This leaves us with 1,170 days to categorize.

3.2.2 Identifying Relevant Articles

Just like BBDS, we recruit a group of student assistants, whom we will refer to as **coders**, to help us analyze the factor innovations. Their first task is to identify newspaper articles that provide meaningful information about the causes of stock market movements. To do this, each student receives a randomly selected list of days with large innovations, identified as described in Section 3.2.1. We do not provide any information about which factors exhibited innovations on the respective days, nor whether these innovations were positive or negative.

One difference from the analysis by BBDS is that our coders use the NYT instead of the WSJ as a source. Nimark and Pitschner (2019) suggest that although different news outlets may have individual thematic priorities when reporting on crucial events, they tend to harmonize their narratives, which results in a cohesive news environment. The NYT is the second-largest U.S. newspaper with a substantial business section which includes stock market news. The format of this daily stock market section varies over time ("Market Place", "The Markets", "DealBook") but ultimately reports on the previous day's market events.

To find an appropriate article, the coders search the NYT of the day following the day with the large factor innovation for stock market related key words as e.g. "stocks" or "Wall Street". Usually a range of possible articles then appears and must be screened by the coders. In rare cases, it can happen that the search yields no results due to missing textual readability of articles

[2]Note that there are slightly more than 500 jumps per category as there are multiple observations with the same (up to one basis point) factor returns. We choose thresholds to cover **at least** 500 jumps per factor. For exact numbers see Table 3.1.

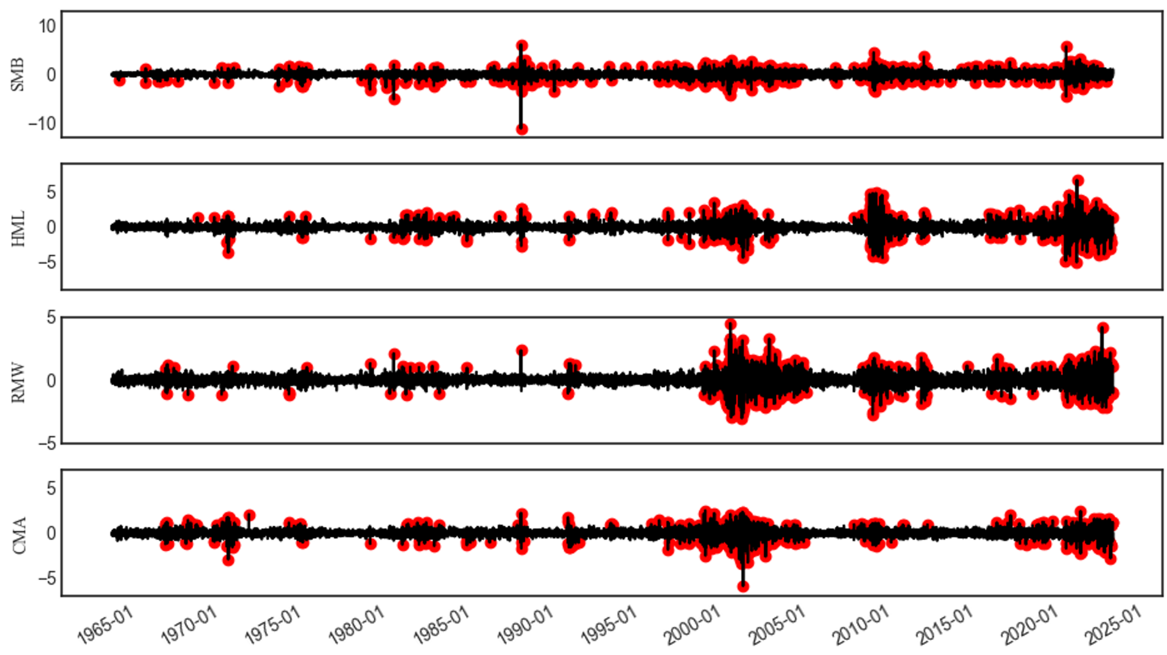

Figure 3.1: **Daily factor returns.** Time series of daily returns on characteristic-based factors. Red dots mark returns on days considered in our analysis.

and meta information.[3] Coders then must manually screen the newspaper for relevant articles. As a last resort, the coder categorizes the day as "No Article", indicating that no suitable article could be found.

3.2.3 Categorization by Human Coders

To assign categories to the identified articles, we follow BBDS and have our coders read and categorize the articles. Alternatively, we use ChatGPT (see Section 3.2.4). We train our coders with a comprehensive guidebook, which aligns almost entirely with the guide by BBDS, except for a few specific differences. These differences cover aspects that are important for identifying stock movements that affect not the overall market but only certain stocks. An excerpt from coders' guide is presented in Appendix B.1 and the complete coder's guide is available upon request.

We use the same 16 categories as BBDS. Our guide provides coders with many examples of these categories and information on how to tackle more complicated set ups. Additionally,

[3]Starting from 1981, the NYT provides fully machine-readable texts. Before that, the full articles are available as photographs, and only titles, abstracts, and keywords are available as machine-readable (and searchable) text. However, even this information is missing for some articles.

every coder has the option to seek advice from one of the authors of this chapter supervising the task in case unforeseen difficulties come up. To verify that our coders are up to the task and that the categorization using the NYT is consistent with that using the WSJ, we have our coders first categorize about 20 days that were also categorized by BBDS. We find quite consistent categorization results.

The 16 categories are:

1. *Commodities*
2. *Corporate* earnings and outlook
3. *Elections* and political transitions
4. *Ex*change *rate* policy and capital controls
5. *Foreign* stock markets
6. *Gov*ernment *spend*ing
7. International *trade* policy
8. *Macro*economic news and outlook
9. *Monetary* policy and central banking
10. *Other* non-policy
11. *Other pol*icy
12. *Reg*ulation
13. *Sov*ereign *mil*itary and security
14. *Taxes*
15. *Terror*ist attacks and large-scale violence by non-state actors
16. *Unknown* and no explanation

We henceforth use the parts in italics print of the above category names as abbreviations.

Some articles cite multiple reasons for changes in stock prices. In these cases, along with the primary category, we allow our coders to report a secondary category. If the article does not clearly indicate which category should take priority, we follow the order in which the reasons are mentioned.

Our dataset, which includes more recent events such as the Covid pandemic, comprises many factor innovations in 2020 (see Figure 3.1). After 2018, the NYT replaced its “The Market” section with “The DealBook,” which provides less detailed analysis of the reasons behind stock movements on a given day. Despite this, we still find articles discussing events that could potentially influence stock movements, even if these articles do not directly connect the events to the stock market. To capture this information, we have introduced an “own evaluation” checkbox for coders.

Just as BBDS, we further gather information on "journalist confidence" to capture the level of certainty with which the article author determines the reasons for stock price movements. Moreover, coders have to evaluate the "ease of coding," i.e., the difficulty of finding the right category using a three point scale. Additionally, the coders must provide a link to the article, its headline, and a key passage that justifies their categorization. This data is available in the supplementary material for this article.

We have each article categorized by at least two coders. This allows us to validate whether the selected categories are consistent. If the primary category coded by the two coders differs, the corresponding data is reviewed and categorized by a third coder. In cases where the third coder faces significant difficulty in coding (e.g., when both initial coders use the "own evaluation" option), we default to "no article" coding. This applies to 18 days with large factor innovations in our sample.

Coders initially agree on the primary category for about half of the days in our sample. For the other half, we often see a reversal of the primary and secondary categories. We create an *initial disagreement dummy* to label days where the initially coded primary categories differ across coders. This dummy is strongly negatively correlated with "journalist confidence" and "ease of coding," and strongly positively correlated with "own evaluation." More precisely, the disagreement rate of initial coders is 66.1% (72.8%) when journalist confidence (ease of coding) is low (difficult), and 42.2% (37.2%) when it is high (easy). In cases where one of the initial coders used the "own evaluation" option, the disagreement rate was 68.5%. These findings suggest that categorization is difficult for the coders when the journalists themselves are unable to clearly identify the origin of a major innovation in the stock market.

3.2.4 Categorization by ChatGPT

In addition to our human coders (henceforth referred to as *HI* - Human Intelligence), we use version 3.5 of ChatGPT (henceforth referred to as *AI* - Artificial Intelligence) to categorize days with large factor innovations.

The prompt includes the NYT article identified by our human coders, as described in Section 3.2.2, along with the jump category description from the coders' guide and one example per category. We instruct the AI to provide a category, as well as classifications for ease of coding and journalist confidence. To better understand the classification choice, we also request the key passage of the coded article. The full prompt is provided in Appendix B.2.1. The dataset accompanying this chapter includes both HI and AI categorizations for each day with large factor innovations.

The AI categorization aligns with the HI primary category on 42.39% of the days. A detailed analysis of similarities and differences, along with a discussion of possible reasons and conclusions, is provided in Section 3.3.3.

3.3 Descriptive Analysis of the Data Set

3.3.1 Importance of Different News Categories

Table 3.1 shows the absolute number of days that exhibit large factor innovations, categorized by human coders (columns 1 to 5) and by ChatGPT (columns 6 to 10). Column 5 (10) displays the total number of days in the respective category in our sample, while columns 1-4 (6-9) break down the days based on when the respective four characteristic-based factors show high absolute returns. Columns 5 and 10 are not given by the sum of the previous four columns, as factors often show high absolute returns on the same days. The days coded by BBDS are equally integrated into HI and AI.

The three most frequently selected categories by the human coders are, by far, *Macro*, *Corporate*, and *Monetary*. This is true not only for the total numbers but also equally for all four factors. This result aligns with the findings of BBDS, who also coded these three categories most frequently (at least in the post-war era, see their Table 1), and remains true even when we exclude the days coded by BBDS. The fourth most frequently coded category by human coders (again, for all four factors) is *Other* (non-policy). This category is often coded when the reason for changing stock prices described by journalists is "bargain hunting," i.e., stocks have been falling for a while, and suddenly prices rally as investors find the prices fair again, leading them to buy. Other reasons described include panic selling to meet margin calls or other psychological factors driving investors' trading behavior. In BBDS, the category *Other* ranks fifth, behind *Govspend*, which in our results comes in fifth. Overall, our findings are very consistent with those of BBDS. While this outcome was not necessarily expected, since the characteristic-based factors could spike for different reasons than the overall market, the result is reassuring.

Table 3.1: **Absolute number of days in categories.** Absolute number of days with large factor innovations by coded category. Columns 1-5 refer to the categorization by human coders and columns 6-10 to the categorization by ChatGPT. Categorizations of BBDS are included in both data sets.

	Human intelligence					Artificial intelligence				
	SMB	**HML**	**RMW**	**CMA**	Total	**SMB**	**HML**	**RMW**	**CMA**	Total
Commodities	17	14	10	14	41	40	34	31	42	101
Corporate	73	93	136	141	300	51	61	75	76	181
Elections	10	8	10	7	24	4	0	1	1	6
Exrate	6	0	0	2	8	4	1	2	2	7
Foreign	6	4	2	4	12	96	73	117	132	292
Govspend	25	32	16	8	54	17	24	15	9	38
Macro	129	113	115	101	329	142	130	107	105	351
Monetary	79	69	67	79	203	35	45	34	35	99
Other	44	28	30	30	89	9	8	7	3	14
Otherpol	9	11	3	4	16	7	10	7	4	17
Reg	5	2	5	3	9	1	1	1	2	3
Sovmil	6	10	19	19	36	1	4	7	6	11
Taxes	5	3	4	6	13	1	2	2	2	4
Terror	1	0	2	2	4	0	0	1	2	2
Trade	4	5	2	4	12	1	0	1	1	3
Unknown	20	27	15	25	53	30	26	28	27	111
No Article	73	81	69	54	192	73	81	69	54	192
Total	512	500	505	503	1395	512	500	505	503	1395

With AI coding, *Macro* is again the most frequently selected category. However, the categories *Corporate*, *Monetary*, and *Govspend* are chosen significantly less often, while *Commodities* and *Foreign* are selected much more frequently by ChatGPT compared to the human coders. A detailed analysis of these differences follows in Section 3.3.3.

To identify differences between the four characteristic-based factors, it is easier to compare the relative proportions of the coded categories. Figure 3.2 shows these for HI (Graph A) and AI (Graph B).

Overall, the relative frequency of the coded categories across the four factors is very similar. In Graph A, we can see that for all four factors, *macro*economic and *monetary* news dominate more than a third of all days with large factor innovations. Together with *corporate* news, they account for more than half of the days. However, there is also a significant difference between the days: The proportion of days with *corporate* news is significantly higher for RWM and CMA, at around 28%, compared to SMB and HML, each at about 15%. In contrast, the categories *Govspend* and *NoArticle* appear somewhat more frequently for SMB and HML. Apart from that, the differences are rather small.

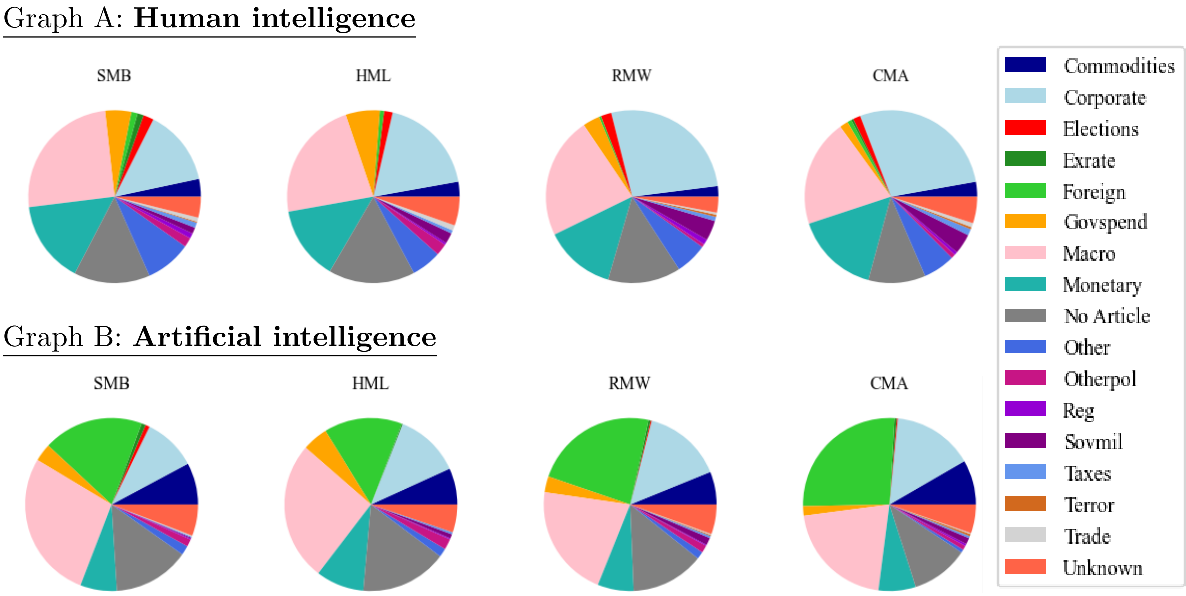

Figure 3.2: **Dominant category per factor (HI).** This figure shows the percentage of each category human coders chose as main reason for stock movements on the respective factor jump days. The whole sample covers 1395 unique jump days.

Graph B shows the categories coded by ChatGPT and presents a similar overall picture. As previously mentioned, the categories *Foreign* and *Commodities* appear more frequently here, while *Corporate* appears less often. *Foreign* is coded particularly often on days with large innovations in RMW and CMA. In comparison, according to ChatGPT, SMB and HML are influenced somewhat more frequently by *Macro* news. Altogether, however, we again find little variation between the factors.

Figure B.2 in the Appendix shows the relative shares of different news categories in the quadratic variation of the four factors. All in all, it also confirms that the differences between the factors are not particularly pronounced, and the observable differences are very similar to those described above.

The key takeaway from this initial analysis is that the differences in the economic origins of variations between the four factors are not revealed solely by looking at the frequency of the news categories that influence them. This is plausible since all the factors are long-short portfolios of the same stocks. If a piece of news affects a specific group of stocks, it will affect exactly one factor only in the case where the membership of that group is exactly orthogonal to the membership in the long and short portfolios of all other factors. In such a case, we could conclude that the factor is strongly influenced by this type of news and could interpret the factor risk premium accordingly. However, it is not surprising that we find that different types of news can influence all the factors alike.

This is consistent with the fact that the factors often move together. Table 3.2 shows how often a factor has an absolute return exceeding one standard deviation on days with a large innovation in another factor, as coded in our sample. We divide these observations into returns with the same sign and with opposite sign than the coded factor innovation. We find that the factors very often exhibit relatively high absolute returns on the same days, while the direction of these movements can vary significantly. This aligns with the notion that the characteristic-based factors are linear combinations of various fundamental factors.

Table 3.2: **Relation between factor innovations.** Number of days where a factor (in columns) have absolute returns exceeding one standard deviation, given a factor exhibits a large innovation, as defined in Section 3.2.1. Columns 1,4,7, and 10 show the number of days with returns with the same signs as the factor innovations. Columns 2,5,8, and 11 show the number of days with returns within one-standard deviation bands. Columns 3,6,9, and 12 show the number of days with returns with opposite signs as the factor innovations.

	SMB			HML			RMW			CMA		
	same	~ 0	**opp**	**same**	~ 0	**opp**	**same**	~ 0	**opp**	**same**	~ 0	**opp**
SMB	512	0	0	179	242	91	60	243	209	134	273	105
HML	186	220	94	500	0	0	219	166	115	321	153	26
RMW	67	221	217	239	165	101	505	0	0	259	148	98
CMA	129	263	111	398	90	15	273	147	83	503	0	0

Our findings highlight the need for a more differentiated analysis. In Section 3.4, we construct "topical factors" by calculating the principal components of underlying securities on days with news from individual categories. This allows us to determine which securities consistently move on these days and how these movements relate to the four characteristic-based factors.

3.3.2 Plausibility of Categorization

Scheduled events often have a strong impact on stock prices and possibly on the four characteristic-based factors. In this section, we examine whether the factors frequently show large innovations on such days and whether categories corresponding to the themes of the events were coded in these cases. We first look at announcements and election dates. Subsequently, we check conversely whether, on days when *ExRate* or *Commodities* are coded, there were large movements in the FX or crude oil markets. These are important plausibility checks for our codings.

Announcements. We obtain announcement dates between 1970 and 2022 from the economic calendar at `www.investing.com`, a source also used by Bloomberg and many other researchers (see, e.g., Kerssenfischer and Schmeling, 2024). We use major news announcements (highest importance ranking) of the following seven categories: Balance, Bonds, Central Banks, Confidence Index, Economic Activity, Employment and Inflation. *Balance* includes news on crude oil inventories. *Bonds* subsumes events concerning e.g 10-year or 30-year note auctions. *Central Banks* includes days with Fed interest rate decisions, testimonies or FOMC meeting minutes. These announcements are only available starting from 1982. *Confidence Index* are the release dates of the monthly CB Confidence Index. *Economic Activity* encompasses the release of the following indices: Chicago PMI, Philadelphia Fed Manufacturing Index, Existing Home Sales, ISM Manufacturing PMI, New Home Sales. *Employment* entails initial jobless claims, unemployment rate and non-farm payrolls. *Inflation* includes release days of ISM Manufacturing Prices, CPI, Core CPI, and Core PCE.

Table 3.3 provides an overview of the sampled announcement days. It is noticeable that *Central Bank* announcements are relatively more frequently associated with large absolute factor innovations, while announcements about *Economic Activity* and *Employment* occur significantly more often in absolute terms and therefore lead to large factor innovations on more days in absolute terms as well. Looking at the different factors, it is noticeable that HML shows large innovations on announcement days much more frequently than the other three factors.

Table 3.3: **Summary statistics announcements.** Frequency of scheduled announcements and days coinciding with days with large factor innovations. *# ann* denotes the absolute number of announcements, *w/ jump* denotes the number of announcement days which coincide with days with large factor innovations, *share* is w/ jump divided by *# ann*, and the remaining columns show the absolute frequencies of the respective factor jumps on announcement days.

	# ann	w/ jump	share	SMB	HML	RMW	CMA
Balance	1081	156	14.43%	59	70	53	43
Bonds	365	53	14.52%	17	29	17	18
Central Banks	786	134	17.05%	53	66	49	37
Confidence Index	661	54	8.17%	20	25	15	23
Economic Activity	1984	271	13.66%	105	124	82	82
Employment	2105	267	12.68%	96	130	92	84
Inflation	1315	150	11.41%	52	80	47	44

Of the 1,395 days with large factor innovations, 541 days (38.8%) coincide with announcement days. Figure 3.3 shows which categories were selected by the human coders on these days. Figure B.3 in the Appendix shows the same information for ChatGPT's coding. In both graphs, it must be noted that multiple announcements often occur on the same days, so it is not always possible to correctly identify the cause of a stock market reaction. Additionally, an innovation in a factor can also be due to another event, independent of the announcement, but taking place on the same day as the announcement.

Despite these limitations, we see that on most days with announcements from six of the seven types, the *Macro* category was coded. This is plausible, as major releases of economic indicators are expected drivers of stock prices and should be categorized as *Macro* news. The exception is *Central Bank* announcements, where the *Monetary* category was most frequently coded. Both apply to human and AI codings.

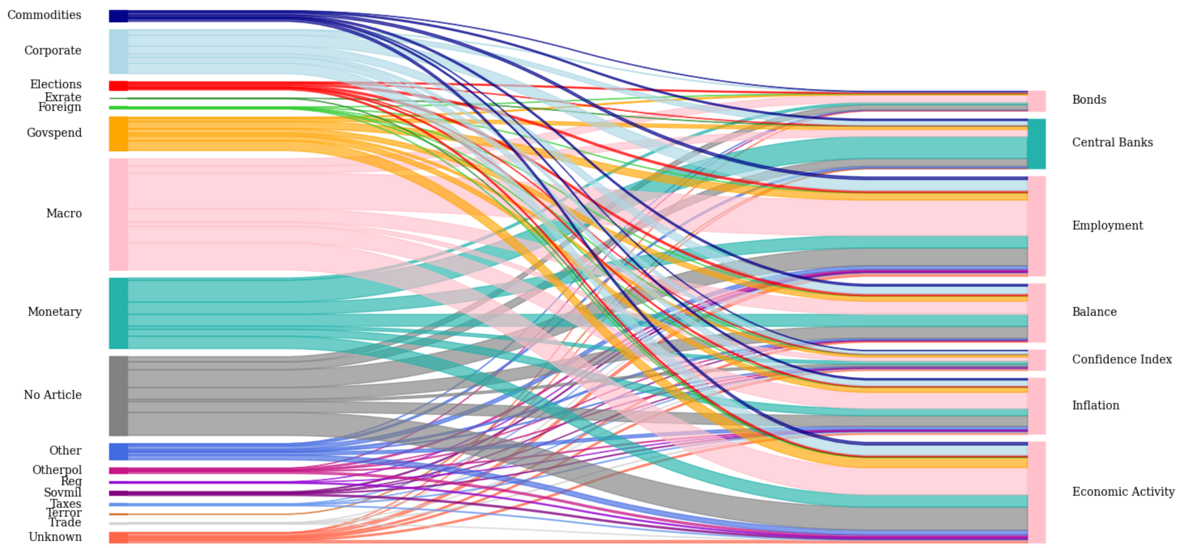

Figure 3.3: **Announcements (HI).** This figure shows the scheduled announcement categories (right hand side) as well as our human coded categories on announcement days (left hand side). Colors on the right hand side indicate the major contributor (pink = macro, green = monetary). Announcement dates are retrieved from the economic calendar from `https://www.investing.com/economic-calendar/`.

Intriguingly, 113 out of 192 days (59%) on which no article was found that could provide information on the cause of a factor innovation are announcement days. This suggests that the factors reacted to the news contained in the announcements on these days, but the news was not significant enough to warrant coverage in the NYT. Following this intuition, we can assume that many innovations with *NoArticle* codings can be attributed to macroeconomic news.

Elections. We look at the dates of 15 presidential elections and 14 midterm elections in the United States between 1964 and 2020. On five of the 15 presidential election dates, there were large innovations in one of the factors. One of these days was categorized by BBDS as *Elections*. The remaining four were also categorized as *Elections* days by the human coders. Recent examples in our sample are the elections of Donald Trump in 2016 and Joe Biden in 2020. ChatGPT categorized these days as *Unknown* (2016) and *Corporate* (2020), and the two remaining days as *Macro*.

Of the 14 midterm elections, four days appear in our sample. One of these was categorized by BBDS as *Corporate*, and one was categorized as *Monetary* by both the human coders and ChatGPT. The two remaining days were categorized as *Elections* by the human coders and as *Foreign* and *Macro* by ChatGPT. Overall, these results are reassuring regarding the coding quality of the human coders, while ChatGPT's coding here appears somewhat questionable. Although it is not clear whether every factor innovation on election days can actually be attributed to the elections, it is rather unlikely that the elections, as suggested by the AI, did not critically influence the factors in any case.

Exchange Rates. We consider daily logarithmic returns on an equally weighted portfolio of four currencies, relative to the US Dollar, namely the German Mark (from January 1999 the Euro), the French Franc (from January 1999 the Euro), the Swiss Franc, and the Japanese Yen. Columns 1 and 2 of Table 3.4 show the standard deviations of the return on the respective days (i.e., the square root of the squared deviations from the full sample mean of the return).

For the days categorized by human coders, we can observe that the standard deviation on days coded as *Exrate* is significantly higher than for all other categories and more than three times as high as on days when none of the factors show large fluctuations. Other topics with increased standard deviations in exchange rates are *Govspend*, *Elections*, *Monetary*, and *Macro*, which also seems reasonable. Overall, these results are reassuring concerning the performance by human coders. In contrast, when coded by ChatGPT, no increased standard deviation in exchange rates on *Exrate* days is observed. We explore the differences between HI and AI coding in more detail in Section 3.3.3.

Oil Prices. In Columns 3 and 4 of Table 3.4, we see, analogous to the FX returns in Columns 1 and 2, the standard deviations of the daily changes in the logarithmic West Texas Intermediate (WTI) oil price. Similar to the FX returns, we observe that the returns on days categorized by

Table 3.4: **Standard deviation of exchange rates and oil prices**. Square root of squared difference between returns on days which are categorized in respective category and full sample average returns. FX returns refer to the logarithmic daily returns on an equal weighted portfolio of German Mark (Euro from 1999), French Franc (Euro from 1999), Swiss Franc, and Japanese Yen. The sample is January 1971 till December 2022 and the data are downloaded from the website of the Federal Reserve Bank of St. Louis. West Texas Intermediate (WTI) returns refer to daily logarithmic oil price changes. Oil prices are Cushing, Oklahoma WTI Spot Prices, downloaded from the website of the Energy Information Administration.

	FX returns		WTI returns	
	HI	AI	HI	AI
Commodities	0.50	0.54	7.97	5.29
Corporate	0.46	0.48	3.86	4.73
Elections	0.59	0.92	3.14	4.57
Exrate	1.13	0.31	1.72	3.51
Foreign	0.30	0.41	2.38	3.09
Govspend	0.74	0.85	5.93	6.14
Macro	0.55	0.58	17.29	16.67
Monetary	0.58	0.53	3.35	3.44
NoArticle	0.34	0.34	4.27	4.27
Other	0.41	0.42	5.23	12.12
OtherPol	0.39	0.35	8.88	4.20
Reg	0.58	0.92	9.15	15.37
SovMil	0.54	0.69	10.60	14.47
Taxes	0.44	0.53	8.25	5.56
Terror	0.27	0.35	8.81	12.23
Trade	0.38	0.70	3.26	0.31
Unknown	0.45	0.52	2.60	3.11
No Jump	0.34	0.34	3.38	3.38

human coders as *Commodities* are more than twice as volatile as on days without significant factor innovations. In the coding by artificial intelligence, the effect is noticeably weaker.

Overall, however, the oil price moves just as or even more strongly on days categorized otherwise. This is because the oil price is also heavily influenced by other news, such as growth prospects (*Macro*). A good example is April 20, 2020, when the WTI oil price ended in negative territory. However, this day was influenced by negative growth forecasts and was categorized as *Macro* by both human coders and ChatGPT. Another example is the outbreak of military conflicts (*SovMil*), which also frequently have a noticeable impact on oil prices.

3.3.3 Comparison of AI and HI

An interesting result from Sections 3.3.1 and 3.3.2 was the pronounced differences between the codings of human coders and ChatGPT. This raises the question of whether the human coders or ChatGPT might be making systematic coding errors. To investigate this, we closely examine ChatGPT's codings, particularly in cases where they differ from human codings. Figure 3.4 shows how often different categories were selected in various combinations by human coders and ChatGPT.

When ChatGPT coded categories such as *Macro*, *Corporate*, or *Monetary*, there was generally strong agreement with the human coders. However, there are numerous differences in the codings. Three aspects stand out in particular:

First, ChatGPT categorizes more than twice as many days as days with news about *Commodities* compared to the human coders. When reading the articles categorized exclusively by ChatGPT as commodity news, it becomes clear that although commodities are mentioned, they are usually not the source of innovation in stock prices. For one, many articles discuss changes in stock, bond, and commodity prices, especially gold, in response to other types of news. Additionally, previous day's returns on stocks in various sectors, including gold and silver mining stocks, are often reported. In other cases, ChatGPT appears to misinterpret the context of words (example: "...market was fueled by hopes..."). In such cases, ChatGPT tends to categorize the article under *Commodities*, suggesting that the AI struggles to differentiate causal relationships from mere mentions of a topic or even a term. Appendix B.2.2 discusses an article where this error occurred.

Second, ChatGPT selects the category *Foreign* stock markets for about 20% of all examined days, a category that is rarely chosen by human coders. This category is meant to be used only when innovations in the U.S. stock market are solely due to events on foreign stock markets. When reading the relevant articles, it is often completely unclear why they were categorized this way by the AI, as foreign stock markets are frequently not mentioned at all (for examples, see Appendix B.2.2). Here, we can only speculate that the term "foreign" was not sufficiently taken into account.

Third, ChatGPT selects the *Monetary* category much less frequently than the human coders, often favoring the categories *Foreign* stock markets (see the discussion above) and *Macro*

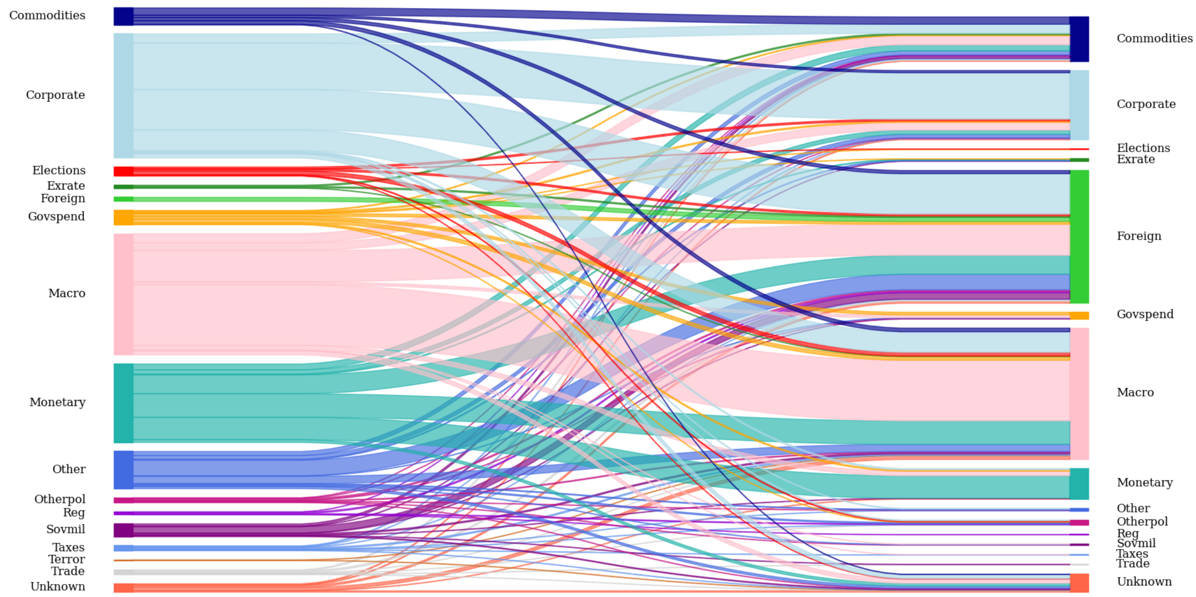

Figure 3.4: **Categorizations by HI and AI.** Human coder categorizations are on the left side of the chart, ChatGPT categorizations for the same jump dates are on the right side of the chart. Codings taken from BBDS are excluded from the graph.

instead. In some cases, it is indeed difficult to distinguish between *Monetary* and *Macro*. However, there are also articles categorized as *Macro* by ChatGPT, where innovations in the stock market are clearly attributed to monetary policy decisions by the Fed or ECB. Examples can again be found in Appendix B.2.2.

Even though human categorizations can never be error-free and algorithmic categorizations allow for better reproducibility, the problems outlined above in ChatGPT's categorization have led us to use the human-coded data as our benchmark dataset. Nonetheless, as a robustness check, we will also refer to the AI-coded dataset in Section 3.4. Newer and future versions of ChatGPT or other large language models will likely generate better categorization results than the ones we currently have.

3.3.4 Factor Innovations during Crises Episodes

Figure 3.5 shows the temporal distribution of factor innovations. We observe that factors tend to exhibit particularly high absolute returns when market volatility is generally elevated due to various crises. Specifically, during the periods 2000-2001, 2008-2009, and 2020-2022, we find a high number of large absolute factor innovations, which we attribute to the Dot-Com Crisis, the Global Financial Crisis, and the COVID Crisis. We now examine whether the different crisis episodes differ in terms of the causes of the factor innovations. In addition to the three episodes just mentioned, we also look at the crises of the 1970s and 1980s, such as the two oil crises and the Flash Crash of 1987.

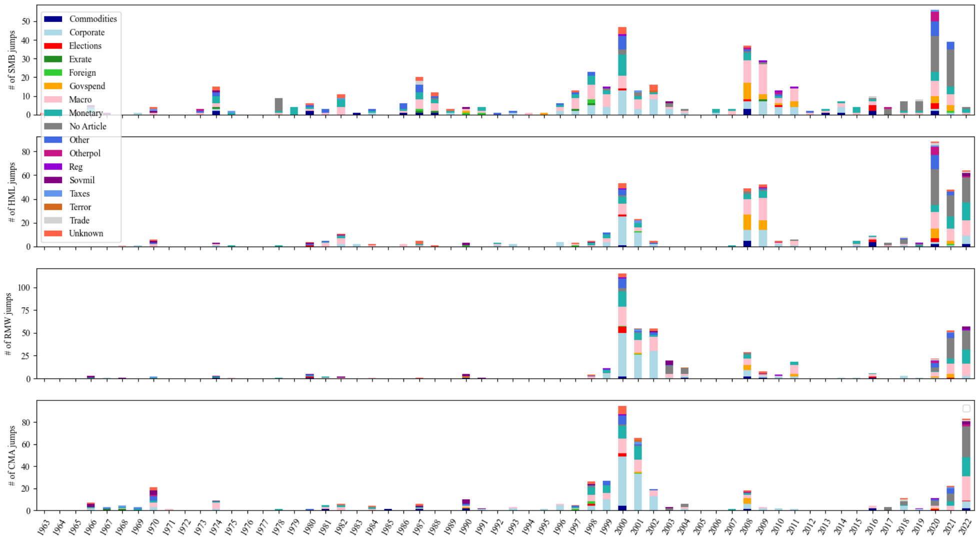

Figure 3.5: **Factor jumps with categories.** This figure shows the frequency of dominant category occurrences between 1963 and 2022. The four subplots show jumps of (1) SMB, (2) HML, (3) RMW, and (4) CMA. Colors indicate the categorical driver coded by HI. Number of observations per factor: SMB: 512, HML: 500, RMW: 505, CMA: 503.

We see that the crises of the 1970s and 1980s primarily affected the SMB and CMA factors, while RMW in particular exhibits hardly any major innovations before 2000. In contrast, CMA shows few major innovations during the financial crisis. During the Dot-Com Crisis, we observe large innovations across all four factors, with RMW standing out particularly. About half of all days with major innovations in RMW fall within this crisis episode. The COVID episode, in turn, leads to numerous spikes in all four factors.

An important follow-up question is whether the temporal distribution of large innovations significantly influences the relative frequency of the categories of the different factors. This would be the case, for example, if a factor spikes particularly often during a specific crisis, and this can be attributed to a particular type of news. Figure 3.6 shows the relative distribution of the categories in the respective time windows. We do indeed find very strong differences between the crisis episodes. Most notably, during the Dot-Com Crisis, over 50% of the days with large innovations can be attributed to *Corporate* news. This explains the high relative share of

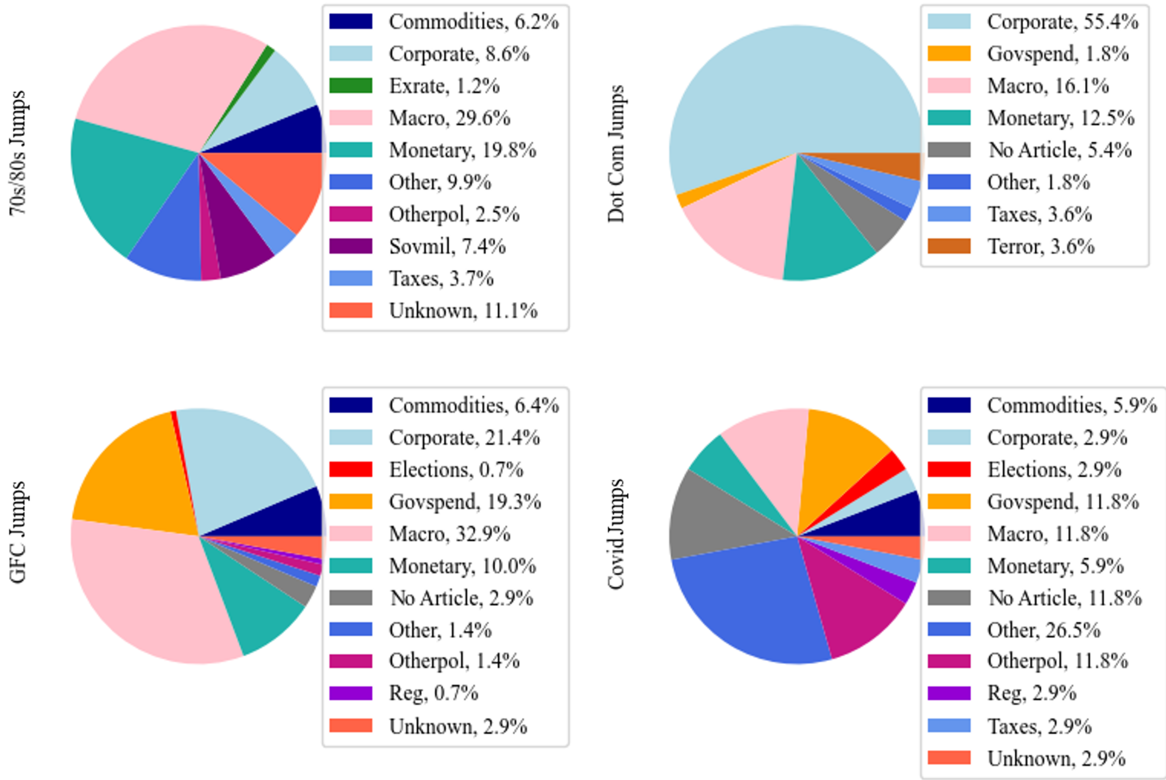

Figure 3.6: **Categories during crises.** This figure shows the percentage of each category human coders chose as main reason for stock movements on the respective factor jump days. Crises are defined as periods of economic stress, i.e. NBER recession periods.

Corporate news for RMW innovations, which we documented in Section 3.3.1, as RMW exhibits high absolute returns, especially during the Dot-Com Crisis.

Macro and monetary news are frequently coded categories in the other crises. During the financial crisis, news about government spending also played an important role (recall the Economic Stimulus Act and other measures). Similarly, *Govspend* news played a key role during the COVID crisis. In 2020 and 2021, numerous COVID-19 relief laws were passed, aimed at determining how government-provided funds should be used to support businesses and households. During the COVID crisis, many days were also coded as *Other*, *Otherpol*, or *NoArticle*. In fact, *Otherpol* almost exclusively occurs during the COVID episode, as this category includes news about COVID policies (such as quarantine rules) that strongly influenced the markets during this phase. The high number of *NoArticle* codings is also due to the fact that during this period, the "The Market" section of the NYT was replaced by "The Dealbook," and stock market reporting became less transparent or less speculative from then on.

3.3.5 Up- and Downside Innovations

In this section, we analyze the direction of the innovations and answer the question of whether certain types of news consistently influence factors in a particular direction. Figures 3.7 and 3.8 show histograms with negative (positive) factor innovations on the negative (positive) vertical axis. The charts also display this information separately for positive, negative, and (nearly) neutral market movements. We will evaluate this information further below, but first, we can look at the sums of the respective three histograms (in the positive and negative areas) for each factor.

Overall, it is noticeable that none of the dominant categories lead to exclusively positive or negative reactions in any of the factors. For example, macro news affects all four factors both positively on some days and negatively on others. This could be because macro news (like all other types of news) can be both positive and negative for the overall economy. Alternatively, it could be that, for example, value stocks react to a particular type of news, while growth stocks are less affected by that news. Since assessing the news is difficult, we use the simultaneous reaction of the overall market as a gauge. If the aggregate stock market reacts positively to news, we interpret this as positive. If this is accompanied by a positive reaction from HML, we can conclude that, primarily, stocks on the long side of the factor—i.e., value stocks—had positive returns, while the returns on growth stocks were less positive.

No clear pattern emerges for SMB. The factor shows innovations both in the same direction as the market and in the opposite direction. From the graph, we can only see that positive monetary policy news tends to have a negative impact on SMB. The same applies to macro news, except during the financial crisis, where we find an exactly opposite pattern.

For HML, we find that positive news from all categories is associated with negative factor innovations, and negative news with positive factor innovations. This applies to all categories but with two notable exceptions: during the financial crisis and in 2020, the first year of the COVID crisis, HML behaves exactly the opposite. We can conclude that overall, positive (negative) news normally influences growth stocks more positively (negatively), with the exception of the aforementioned periods, during which value stocks were primarily affected.

For RMW and CMA, we similarly find that they consistently react in the opposite direction to the market. This applies to all news categories and time periods without exception.

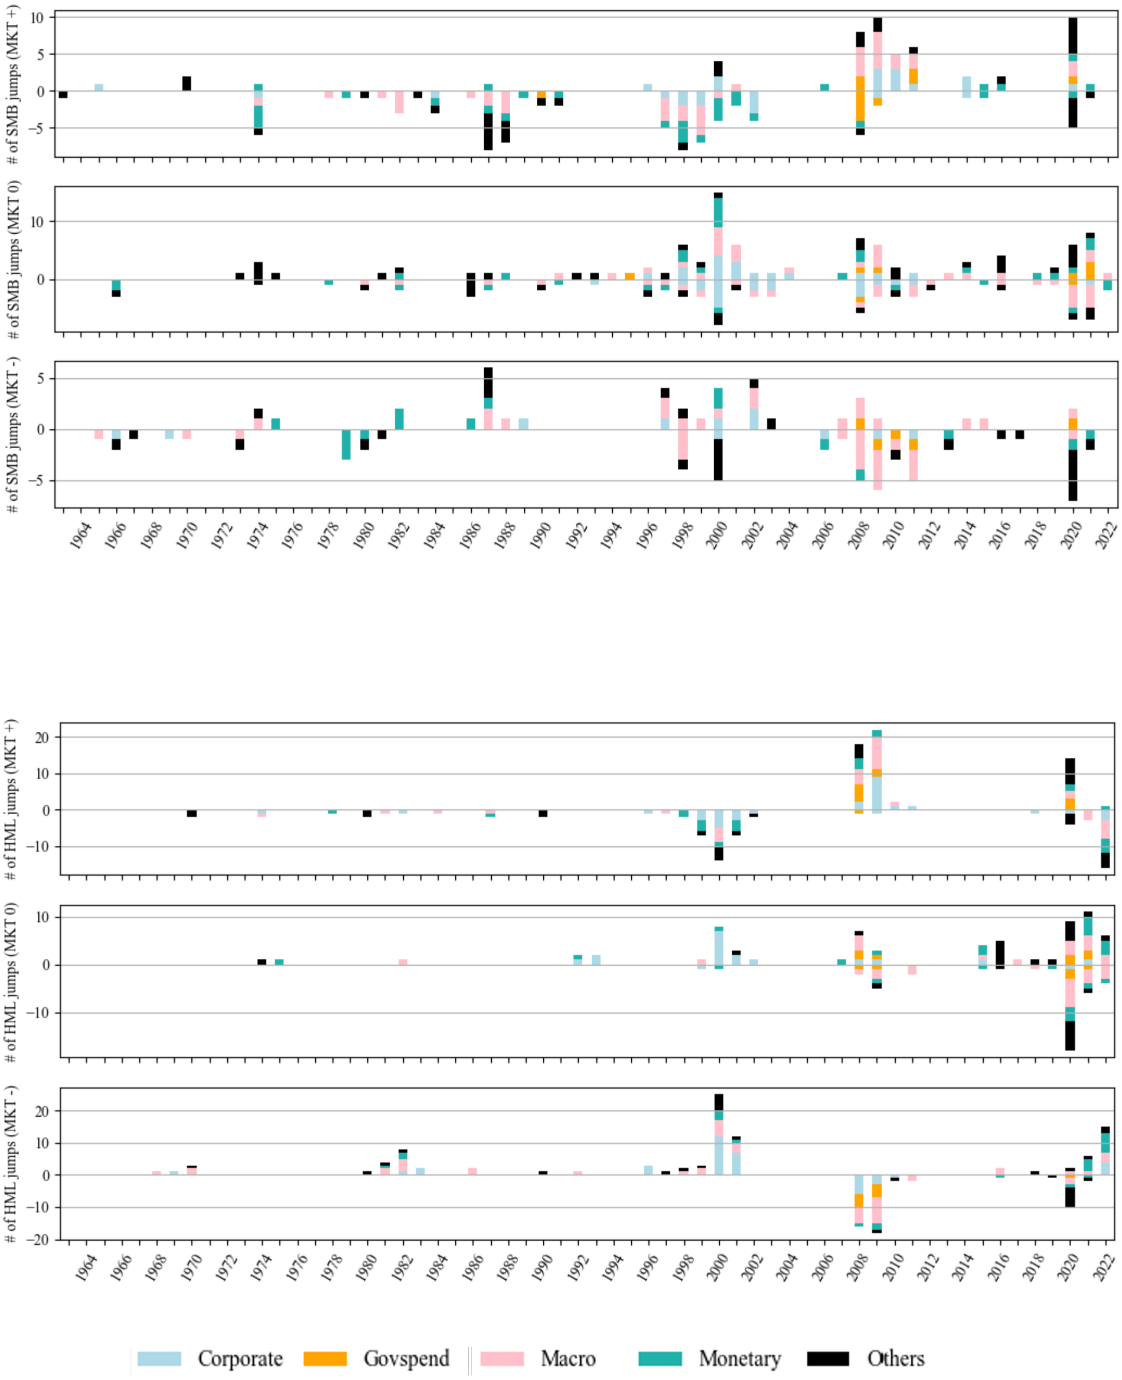

Figure 3.7: **SMB and HML jumps with categories.** This figure shows the number of SMB (top three plots) and HML (bottom three plots) jumps per year between 1963 and 2022. "Others" (black) encompasses all remaining categories. For either factor, the bar chart on top filters factor jumps when the market factor (MKT-RF) is larger than one standard deviation. The bar chart on the bottom filters jumps when MKT-RF is negative and the absolute value is larger than one standard deviation. The chart in the middle shows factor jumps when the market is moving less than one standard deviation.

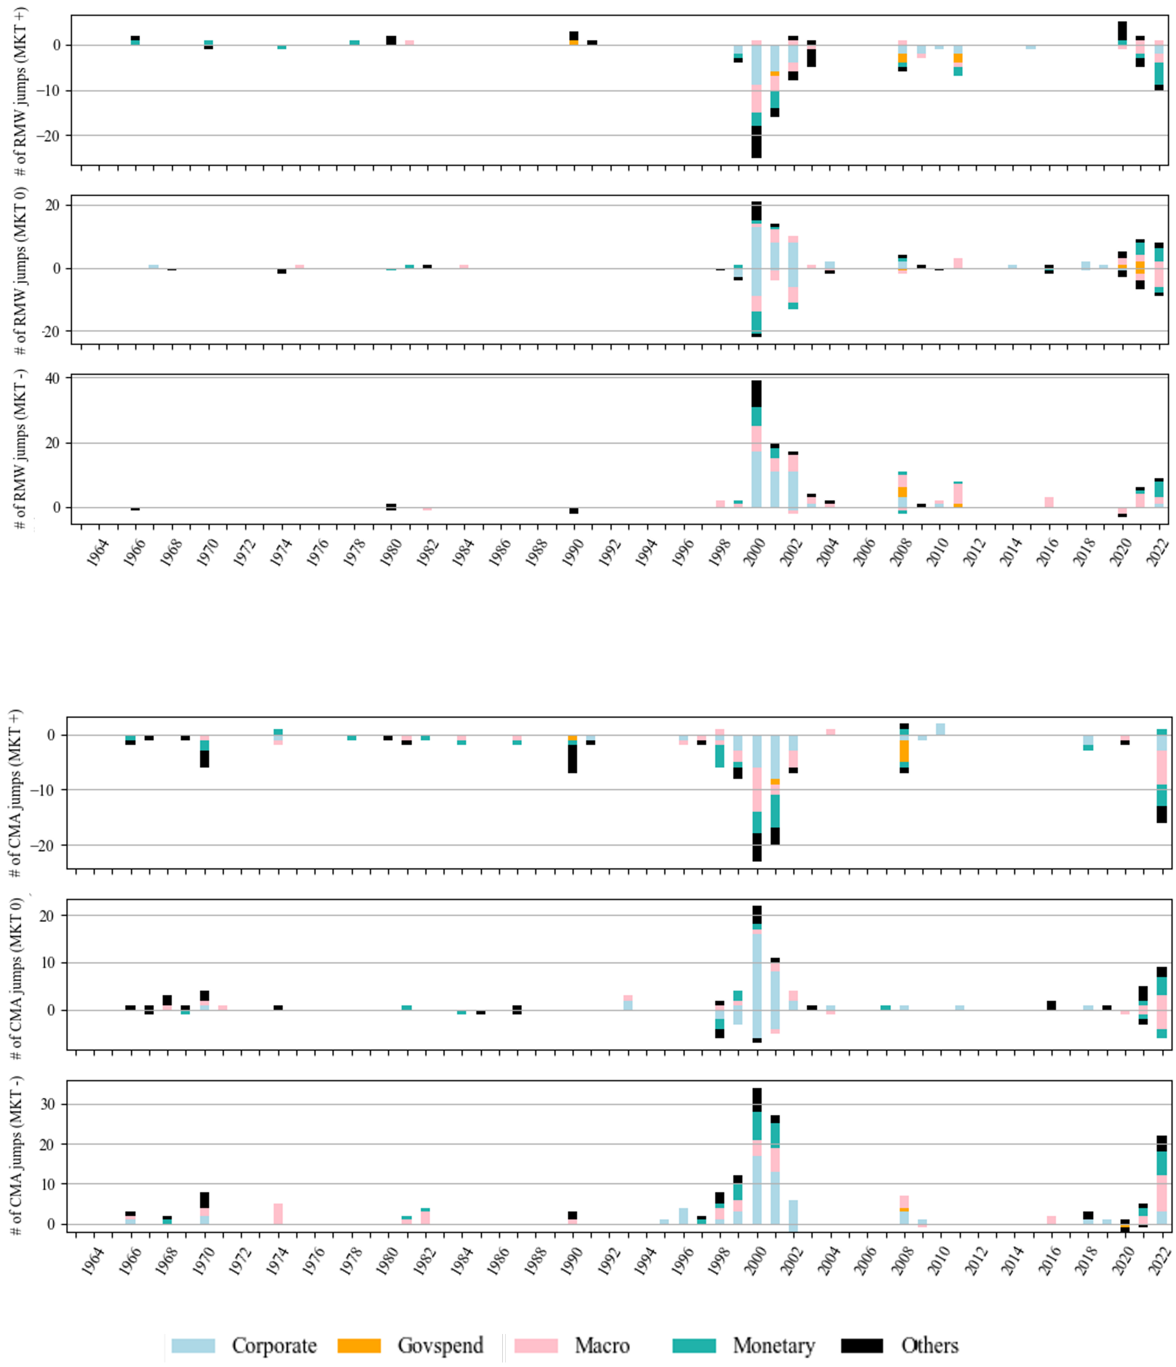

Figure 3.8: **RMW and CMA jumps with categories.** This figure shows the number of RMW (top three plots) and CMA (bottom three plots) jumps per year between 1963 and 2022. "Others" (black) encompasses all remaining categories. For either factor, the bar chart on top filters factor jumps when the market factor (MKT-RF) is larger than one standard deviation. The bar chart on the bottom filters jumps when MKT-RF is negative and the absolute value is larger than one standard deviation. The chart in the middle shows factor jumps when the market is moving less than one standard deviation.

This result suggests that weak and aggressive stocks react more strongly to news than robust and conservative stocks and are primarily responsible for the pronounced innovations in the factors. Overall, the results for these two factors and for HML (at least outside the two discussed episodes) confirm the usual findings and the motivation for the four-factor model as a whole: Stocks in the short leg of the respective factor are more strongly positively correlated with market fluctuations and still have lower average returns than their counterparts in the long leg. This is true even when conditioned on days with large factor innovations or specific types of news.

3.4 Topical Factors

3.4.1 Construction

So far, our analyses are largely descriptive. In order to deeper investigate the relation of characteristic-based factors and economic topics, in this section, we develop "topical factors" pertinent to the aforementioned categories. This allows us to determine which assets consistently move on topic related days and then understand how these movements are connected to the four characteristic-based factors. We select a set of auxiliary assets and scrutinize the performance of these assets on days coinciding with news events within the respective categories. The factors are then constructed as weighted sums of those auxiliary asset returns that are most sensitive to the respective types of news.

To elaborate, for a given news category, say *Elections*, we identify the set $T_{\text{Elections}}$ comprising days characterized by significant factor innovations and news pertaining to elections. Subsequently, we examine the principal components of the daily returns generated by the auxiliary assets on the days within $T_{\text{Elections}}$. These principle components are linear combinations of the auxiliary asset returns themselves and we use the weights of the auxiliary assets to also construct the principle components on days not included in $T_{\text{Elections}}$. On days not within $T_{\text{Elections}}$, where news related to elections is absent, the behavior of the auxiliary assets may not exhibit any discernible patterns linked to the aforementioned principal components. Thus, the factors, by design, exert significant influence on returns specifically during days associated with topic-related news events but may not exert comparable influence on other days.

Unsurprisingly, we observe that the first principal component on news days across all categories demonstrates nearly identical loadings on all auxiliary assets. This observation remains consistent irrespective of the selection of auxiliary assets and underscores the prominence of the market factor in elucidating variations in returns across all types of news days. Consequently, it is more interesting to explore the second principal components on news days, while always acknowledging that any model incorporating topical factors should invariably incorporate the market factor as well.

The intuition behind our approach is straight-forward: Imagine that there was exactly one auxiliary asset that reacted to *Elections* news and all the other auxiliary assets were completely insensitive to these news. The returns on the latter were then driven by the market return plus asset-specific low-volatility noise components on these days and would not have pronounced weights in the principal components we back out. The *Elections* factor would thus be given by the return difference between the one auxiliary asset loading on *Elections* news and the market factor. Our approach generalizes that idea, also considering all possible portfolios formed from the auxiliary assets, instead of single assets only.

We examine two sets of auxiliary assets: the first comprises eight portfolios, constituted by the respective long and short positions of the four factors under consideration. Data comes from Kenneth French's data repository,[4] adhering to the portfolio construction methodology outlined therein. In particular, the portfolio of, for example, growth stocks is constructed controlling for size, i.e., averaging across returns on portfolios double-sorted by market capitalization and book-to-market ratio along the size dimension.

The simplicity in interpreting the weights used to construct the topical factors is a key advantage of this small set of auxiliary assets. Importantly, the characteristics-based factors are linear combinations of these portfolios themselves, with weights, for example, of the form $(1, -1, 0, 0, 0, 0, 0, 0)'$ in case of the size factor. Thus, we can directly compare the weights of the topical factors with these "pre-defined" weights.

However, a limitation of this choice of auxiliary assets is that the 17 factors derived from them can by construction only encompass an eight-dimensional subspace within the return space. While this limitation is acceptable given our primary focus on the four characteristic-based factors, it is conceivable that the 17 topical factors collectively perfectly explicate the four factors, as they originate from the same portfolio returns. We have to take this fact into account

[4] Retrieved from `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

when interpreting the topical factors and their relations with the characteristics-based factors in Section 3.4.2.

To address these concerns, we additionally incorporate a second set of auxiliary assets, comprising 40 portfolios constructed in a similar manner as the aforementioned ones. These portfolios stem from 10×10 double sorts based on size and one of the remaining three characteristics: book-to-market ratio, operating profitability, and investments, sourced again from Kenneth French's repository.

3.4.2 Relation of Topical Factors with Characteristics-Based Factors

Table 3.5 exhibits the weightings of the returns associated with the eight auxiliary assets employed in the construction of the topical factors. We standardize the signs of the loadings so that the loading of the small size portfolio is always positive. Noteworthy patterns emerge from this analysis. Some topical factors predominantly load on portfolios representing small and large firms (e.g., *Unknown*), while others exhibit loading tendencies towards high and low book-to-market ratio firms (e.g., *Macro*). The loadings on portfolios defining profitability and investment factors tend to be of lower magnitude compared to size and value factors. Specifically, factors displaying relatively strong loadings on portfolios categorized by robust, weak, conservative, or aggressive stocks also tend to exhibit strong loadings on high and/or low book-to-market ratio stocks.

In general, opposite signs are observed for portfolios on opposing sides of the factors, particularly for pronounced loadings. For instance, the *Foreign* factor loads positively on small, high, robust, and conservative portfolios, and negatively on the remaining four portfolios. This pattern is consistent with expectations considering that the market factor, roughly corresponding to the average of the two opposing sides, is taken out. However, not all factors show the same patterns. For example, factors such as *Commodities* and *Govspend* display stronger loadings on robust stocks compared to weak stocks, while *Corporate*, *Elections*, and *Monetary* exhibit an inverse pattern.

Subsequently, we proceed to analyze the topical factors themselves. Due to substantial time series correlations between the four characteristic-based factors (and within the portfolios constituting them) reliance solely on the weights from Table 3.5 is insufficient. A topical factor may exhibit a near-perfect correlation with a characteristics-based factor, even in the absence

Table 3.5: **Weights of the eight auxiliary assets in the topical factors.**

	small	big	high	low	robust	weak	cons.	aggr.
Commodities	0.05	-0.07	0.71	-0.51	-0.19	0.01	0.19	-0.39
Corporate	0.11	0.01	0.79	-0.40	0.05	-0.25	0.17	-0.33
Elections	0.41	-0.25	0.72	-0.35	0.01	-0.18	0.13	-0.29
Exrate	0.66	-0.59	0.31	-0.16	-0.14	0.19	0.18	-0.12
Foreign	0.58	-0.46	0.52	-0.30	0.12	-0.11	0.19	-0.18
Govspend	0.28	-0.33	0.73	-0.41	-0.25	0.03	-0.09	-0.22
Macro	0.15	-0.05	0.78	-0.46	0.03	-0.19	0.13	-0.33
Monetary	0.27	-0.14	0.72	-0.39	0.02	-0.17	0.27	-0.37
NoArticle	0.20	-0.08	0.73	-0.47	0.07	-0.27	0.10	-0.33
Other	0.22	-0.09	0.71	-0.44	0.12	-0.30	0.08	-0.37
OtherPol	0.16	-0.17	0.77	-0.49	-0.09	-0.05	-0.08	-0.31
Reg	0.15	-0.46	-0.32	0.13	-0.45	0.54	-0.21	0.32
SovMil	0.02	0.10	0.45	-0.31	-0.43	0.36	0.47	-0.40
Taxes	0.17	-0.03	0.71	-0.36	0.13	-0.31	0.31	-0.35
Terror	0.50	-0.49	0.37	-0.31	-0.25	0.18	0.28	-0.33
Trade	0.02	0.01	0.71	-0.54	0.07	-0.17	0.11	-0.40
Unknown	0.58	-0.70	-0.20	0.09	-0.10	0.26	-0.17	0.12

of weights aligning precisely with the predefined weights discussed earlier.

Correlations between the characteristics-based factors and the constructed topical factors are presented in Table 3.6. Notably, the size factor (SMB) demonstrates a time series correlation of 88.71% with the *Unknown* factor. This result is somewhat unsatisfactory, as small and big stocks apparently move characteristically in opposite directions on days when the origins of the innovations are unclear. The second most informative topical factor for SMB is the *Exrate* factor, which is, however, considerably less correlated in the time series, with a correlation of only 74.74%. *Exrate* is also the most informative factor when using the alternative set of test assets and *Unknown* is the most informative one when using the categorizations by ChatGPT. Going back to the baseline case, a linear combination of *Unknown* and *Exrate* factors exhibits a correlation of 96.89 with SMB (not reported in Table 3.6). This is not surprising given the loadings shown in Table 3.5: both factors load similarly on small and big, but with opposite signs on high, low, conservative, and aggressive.

Moreover, the value factor (HML) exhibits strong correlation with the *Macro* factor, registering a time series correlation of 97.84%. This finding is extremely robust across specifications.

Table 3.6: **Correlations between characteristics-based and constructed topical factors**. Correlations are in %.

	SMB	HML	RMW	CMA
Commodities	10.46	87.24	10.53	72.82
Corporate	17.90	94.02	10.88	57.22
Elections	49.72	82.86	-6.71	43.66
Exrate	74.74	31.86	-42.17	3.86
Foreign	63.11	51.96	-20.38	16.80
Govspend	34.10	81.81	2.05	59.48
Macro	21.64	97.84	11.49	63.47
Monetary	37.08	87.76	-2.11	53.31
NoArticle	20.24	95.53	21.88	67.90
Other	19.02	93.17	26.62	68.81
OtherPol	18.51	85.99	14.16	64.30
Reg	37.71	-44.00	-45.27	-26.31
SovMil	22.21	64.85	-44.12	41.48
Taxes	23.94	83.41	7.66	49.35
Terror	72.89	70.55	-27.75	55.33
Trade	1.83	88.60	26.64	71.45
Unknown	88.71	-20.06	-38.98	-18.20

A similar pattern is observed with the *NoArticle* factor (correlation of 95.53%) and the *Other* factor (correlation of 93.17%). These two factors, however, are also highly correlated with the *Macro* factor (correlations of 98.78% and 89.25%, respectively). This could indicate that stock returns on days with stock market news where the origin is difficult to categorize (*Other*), or without articles (*NoArticle*), are governed by macro news as well. This conclusion aligns well with our finding in Section 3.3.2 that days with *Unknown* news are often macroeconomic announcement days.

We do not find comparably strong results for the RMW factor. The topical factor with the highest correlation with RMW is the *Reg* factor, with a correlation of -45.27%. When adding *SovMil*, the correlation of RMW with a linear combination of the two increases to 88.43% (not reported in Table 3.6).

For CMA, we find correlations of 72.82% and 71.45% with the *Commodities* and *Trade* factors, respectively. Combining the two does not add any explanatory power, since they are almost perfectly correlated with a correlation coefficient of 98.28%. A combination of the *Commodities* and the *OtherPol* factors increases the correlation to 87.85%.

Adding more topical factors further increases the explanatory power for the characteristics-based factors. However, the limitations of our approach make an interpretation of such findings challenging. While the most influential factors are consistent across sets of auxiliary assets for SMB and HML, and the pairs *Reg* and *SovMil* (for RMW) and *Commodities* and *OtherPol* (for CMA) are again among the most informative pairs when using the alternative set of auxiliary assets, we do not find consistent triples of topical factors to explain the characteristics-based factors. It is thus important to think about economic reasons for relations between topical factors and characteristics-based ones.

3.4.3 Discussion

Small-Minus-Big. The fact that the *Unknown* factor is most strongly correlated with SMB is interesting, as it could indicate that innovations in only big or only small stocks are driven by sentiment or institutional frictions rather than attributable to a concrete piece of news arriving on the respective days. The day with the biggest spike in the *Unknown* factor is October 20, 1987, the day after Black Monday. While all stock prices plummeted sharply on October 19, the market recovered on the 20th. However, as the WSJ writes: "The rally was limited only to blue-chip highly liquid stocks, which were aggressively sought out by investors and many corporations through share buy-back programs. But there was no relief for many of the smaller-capitalization issues, such as those traded in the over-the-counter market and on the American Stock Exchange." Indeed, the price of the big portfolio rose by 2.46% that day, while the small portfolio lost 8.73%. An analyst is quoted as saying: "if it's not a quality name, investors don't want it." A similar 'flight-to-quality' was observed on October 16, 1989, when large stocks rose by 2.09% and small stocks fell by 1.45%.

We also discuss the second strongest (or the strongest, on the alternative set of test assets) factor, namely *Exrate*. Several classic studies suggest that exchange rate risk should be a priced risk factor (see Frenkel, 1981; Froot and Klemperer, 1989; Rogoff, 1996). Dumas and Solnik (1995), De Santis and Gerard (1998), and Carrieri et al. (2006) find in empirical studies that the aggregated US stock market includes an exchange rate risk premium. Francis et al. (2008) find an exchange rate risk premium in the cross-section of US industries. When we look at the cross-section of individual firms, Jorion (1990) initially finds few firms with quantitatively relevant exchange rate risk. However, Starks and Wei (2003) point out that nonzero exchange rate exposures will only be observed under certain conditions, namely in firms that have difficulty handling short-term cash flow fluctuations and can easily fall into financial distress because of this. They show that this characteristic is particularly related to firm size, but also to other characteristics that are cross-sectionally correlated with firm size, such as short-term leverage, availability of internal funds, costs of underinvestment, and product specialization. Apergis et al. (2011) also show that exchange rate risk is empirically more relevant for small firms than for big firms.

In essence, an important cause of common variation among small companies could be that, unlike large companies, they have difficulty effectively managing exchange rate risks. If these companies either generate a significant portion of their revenue from foreign business or import primary goods from abroad, an exchange rate shock can have far-reaching consequences for them. This also aligns with some of the shocks identified in our dataset as exchange rate shocks. For example, the NYT wrote on December 15, 1967, "Continued frantic trading in European gold markets touched off a number of rumors, ranging from possible devaluation of the French franc to devaluation of the United States dollar." The small portfolio suffered a loss of 0.75% on that day.

In contrast, the NYT wrote on November 17, 2009: "Wall Street traders latched onto signs on Monday that the dollar would continue to remain weak," citing comments from Ben Bernanke and some analysts. On that day, the market portfolio rose by 1.54%, but the small portfolio (2.72%) benefited particularly, relative to the large portfolio (1.42%). This suggests that small companies, in particular, benefit from favorable export conditions due to a weak dollar.

High-Minus-Low. The connection between the *Macro* factor and HML is not surprising given the extensive literature on the topic. Kogan and Papanikolaou (2014) emphasize the influence of investment-specific technology shocks as significant drivers of the economic cycle. Crucially, these shocks affect growth opportunities and assets in place differently. This explains the difference in risk premia between value and growth stocks. Empirical support for this mechanism comes from Liew and Vassalou (2000), who show that the characteristic-based factor HML has predictive power for future GDP growth. Cooper et al. (2022) empirically reconsider a global version of the classic model of Chen et al. (1986) and find that exposures to macroeconomic risk factors have explanatory power for the return differences between value and growth stocks. The factors include the term spread, a well-known predictor of business cycles (see Harvey, 1988; Chen, 1991; Estrella and Hardouvelis, 1991), the default spread (see Friedman and Kuttner, 1998; Gilchrist et al., 2009) and industrial production. Our results confirm this mechanism by showing that HML essentially behaves like our constructed macro factor, which is composed of principal components on days with news about macroeconomic conditions, industrial production, manufacturing activity, etc.

Robust-Minus-Weak. When we look at the correlation between RMW and the *Reg* factor, we see that news regarding the development of the antitrust lawsuit against Microsoft on November 8, 1999, and April 3, 2000, primarily led to significant returns in tech companies, which are mostly found in the weak profitability portfolio. Here again, the impression arises that RMW innovations are often driven by news regarding individual companies or at least specific sectors. The same is true for *SovMil* news, which, again, often affect certain industries which happen to be in the robust portfolio. For example, on January 4, 1968, robust stocks fell by 0.75% while weak stocks gained by 0.41% upon news on peace negotiations in Vietnam. The NYT wrote about that day: "Many of the glamour stocks, which include those engaged in aerospace and defense work as well as in the electronics industry, fell sharply, apparently hurt by the reports of peace feelers." These and other *Reg* and *SovMil* news also led to strong and cross-sectionally varying returns in growth vs. value and conservative vs. aggressive stocks, which is why the *Reg* and *SovMil* factors are not correlated more strongly with RMW than -45.27% and -44.12%.

Conservative-Minus-Aggressive. The topical factor most strongly correlated with CMA is the *Commodities* factor. It is important to know that *Commodities* news almost exclusively consists of news about oil prices or oil production volumes by OPEC countries. The NYT frequently emphasizes that different sectors respond to this news. For example, they wrote about March 8, 2000: "Airlines benefited from a sharp decline in the price of oil after Saudi Arabia and Iran sent a strong signal of significant production increases. Continental; Delta Air Lines; AMR, parent of American; and US Airways all gained three points or more." On this day, the robust portfolio fell by 0.27%, while the weak portfolio rose by 1.24%. About November 30, 2016, they wrote, "Oil stocks climbed after OPEC nations, which collectively produce more than one-third of the world's oil, agreed to trim production for the first time in eight years." and additionally, "Higher oil prices mean more revenue for companies that extract or sell oil, and energy companies made big gains on Wednesday. Exxon Mobil picked up \$1.40, or 1.6 percent, to \$87.30, as Chevron rose \$2.22, or 2 percent, to \$111.56." On this day, the robust portfolio rose by 1.16%, while the weak portfolio fell by 0.83%.

The oil and airline industries are perfect examples of sectors that react differently to oil supply shocks. However, almost all companies, to varying degrees, use oil or more generally energy goods as input factors for production. Dittmar et al. (2024) emphasize the central role of energy consumption as a factor for understanding the cross-sectional variation in stock returns in general, and particularly those related to investments. They show that more energy-intensive companies invest less than those that are less energy-intensive. The central economic channel is that energy goods are non-substitutable for consumers, and therefore oil supply shocks directly affect the marginal utility of households. As a result, companies that are particularly sensitive to these shocks have to pay high risk premia.

3.5 Discussion and conclusion

We examine a large number of days with high absolute returns in the four asset pricing factors SMB, HML, RMW, and CMA and seek to understand their economic origins. For that purpose, we assemble a team of student research assistants and read newspapers from the following day to systematically categorize the realized factor returns. Consistent with BBDS, macroeconomic news, monetary policy news, and news about the earnings of individual companies and industries are the most common causes of high factor returns for all four factors. To understand the differences between the factors, we construct topic-specific factors as the principal components of the returns on a set of auxiliary assets on days with news from a particular category, and we analyze how these fundamental factors correlate with the four factors.

Our *Macro* factor is almost perfectly correlated with HML, suggesting that the value premium is (also, see the discussion below) a compensation for macroeconomic risk. This result supports an extensive theoretical literature that interprets the value premium accordingly. CMA is most strongly correlated with our *Commodities* factor, which essentially reflects shocks in oil supply and prices. This supports theoretical arguments by Dittmar et al. (2024) and Gao et al. (2022).

For SMB and RMW, the interpretation is less clear. SMB is, on the one hand, strongly correlated with the *Exrate* factor, suggesting that a part of the common variation in the returns on small firms is driven by news on the foreign exchange market. However, SMB is also strongly correlated with the *Unknown* factor, which rather suggests behavioral origins of SMB returns (flight-to-quality, etc.). RMW is predominantly affected by shocks in individual companies or industries, which are represented in the long or short portfolio of the RMW factor. Examples of this include news about regulatory changes (often affecting software stocks) and sovereign military events (affecting defense stocks).

Viewed from a higher perspective, our results suggest that the value factor HML and the investment factor CMA can be interpreted more in terms of risk-based explanations, while our results for the profitability factor RMW lean more towards behavioral explanations. For the size factor SMB, both channels seem to play a role. However, for all four factors, it is quite possible that many different economic channels, including both risk-based and behavioral ones, play a part. If all four factors represent differently weighted linear combinations of various fundamental factors, numerous channels might be relevant in their interpretation. Our approach in Section 3.4 only identifies the fundamental factor most strongly correlated with the characteristic-based factors.

Despite this limitation, due to the strong correlation between HML and *Macro*, it is very unlikely that HML is not causally related to macroeconomic risks. Therefore, this chapter lends support to a large body of theoretical literature explaining factor risk premiums. Since *Macro* is a very broad category, future research should more precisely analyze which types of macroeconomic risks are relevant for which stocks. Moreover, it will be important to better understand which institutional frictions play a significant role in the profitability premium.

Chapter 4

Exploiting Media Attention to Climate Change: CLO Trading of Brown Loans

4.1 Introduction

Institutional investors are increasingly compelled to consider the environmental impact of their investments and reduce exposure to polluting companies. These environmental considerations limit the breadth of their investment opportunities, potentially leading to inferior returns (e.g., Pedersen et al., 2021). The challenge of balancing environmental impact with the need to generate competitive returns for stakeholders becomes even more severe during periods of heightened media attention to climate change. When society pays closer attention to climate change, the pressure to divest from polluting firms further increases and potentially leads to lower prices for securities issued by polluting firms. Although a large literature (reviewed below) shows that institutional investors divest from polluting firms, little is known about the investors who take the other side of these divestments. If these other investors are unconstrained and always able to purchase securities issued by polluting firms, divestments by the likes of banks or mutual funds may have little overall effect.

In this context, the leveraged loan market is an important case to consider. Leveraged loans are the primary source of financing for lower-rated or private companies, and systematic divestment from polluting firms in this market can therefore significantly reduce the operations of polluting companies. Collateralized Loan Obligations (CLOs) are the largest investors in the leveraged loan market and the most important source of financing for lower-rated companies. CLOs are actively managed investment vehicles that raise debt and equity from investors by pooling cash flows of leveraged loans into different tranches. CLO investments are less transparent than those of institutional investors in bond or equity markets, and CLO stakeholders have limited opportunities to withdraw their investments during the CLO's lifetime. Hence, CLOs are likely under less pressure to divest from polluting firms than other stakeholders, enabling them to capitalize on divestments made by other investors.

Given this critical role of CLOs, we investigate if they purchase loans from polluting firms during periods of elevated attention to climate change, when other investors are likely to divest. Utilizing leveraged loan transactions conducted by CLOs, we first show that loans of firms in carbon-intensive industries trade at a discount during periods of increased media attention to climate change. This discount is in line with investors reducing their exposure to polluting firms when society pays closer attention to climate change. Our main finding is that CLOs exploit the heightened media attention to climate change and the resulting price discounts by increasing their investments in carbon-intensive industries. This investment behavior is still present for CLO managers who signed the Principles of Responsible Investment (PRI) and are thereby committed to considering environmental factors in their investments. We further show that CLOs that have a bank affiliated manager show a particularly strong increase in investments in carbon-intensive industries when attention to climate change is high. Hence, CLOs are one investor class purchasing from institutional investors who divest from brown industries.

To examine CLOs' trading activities during times of heightened media attention to climate change, we first need to identify episodes of heightened media attention. Our preferred proxy is the Crimson Hexagon Negative News Index (CHNeg index) developed by Engle et al. (2020). This index captures negative sentiment around the keyword "climate change" in a wide array of media outlets. We define months of heightened media attention as periods when the CHNeg index is above its 80% quantile. The main challenge for our empirical analysis is to distinguish polluting firms from other firms. Because over 85% of the borrowers in our sample are private companies with virtually no public information on their operations, our preferred approach distinguishing polluting firms (henceforth "brown" firms) from other ("non-brown") firms is to rely on the companies' industry classifications. For our main tests, we rely on the Moody's industry classifications in our data and use the rankings provided by Bolton and Kacperczyk (2021) to categorize borrowers into brown and non-brown. In additional tests, we match our borrower sample to Refinitiv-Eikon, obtain more granular industry classifications, and rank borrowers based on either the total carbon emissions or the emission intensity of the borrowers in that granular industry.

Combining our measure of media attention with the definition of brown firms, we then arrange our empirical analysis around four hypotheses. Our first hypothesis and the starting point of our investigation is that loans to brown borrowers trade at a discount when media attention to climate change is elevated. We test this hypothesis using the leveraged loan transactions executed by CLOs. In line with our hypothesis, leveraged loans to issuers in brown industries trade at a discount during times of heightened media attention to climate change. The price difference between brown and non-brown loans during these times exceeds 1% of the average loan price and ranges from \$0.915 to \$1.025, depending on the specification. This result is robust to adding loan-level controls, such as maturity and traded amount, as well as granular CLO-level controls.

Our second hypothesis, and the main focus of this chapter, is that CLOs exploit elevated media attention to climate change by increasing their net purchases of loans from brown borrowers. This hypothesis is motivated by two institutional features unique to CLOs: (i) CLOs resemble closed-end mutual funds in that they face little risk of investor withdrawals before their maturity date and (ii) CLO portfolio holdings are relatively opaque as they comprise mainly loans to private companies. We therefore argue that CLOs face less pressure to divest from polluting firms than other institutional investors, such as banks or mutual funds.

To test this hypothesis, we examine the trading behavior of CLOs over time. For each CLO, we construct a variable that captures the net purchases in different industries and over time. Specifically, for CLO i, industry r, and time period t, we define $\textit{Net Purchases}_{i,r,t}$ as the notional amount of all loans in industry r purchased by CLO i minus the notional amount sold. We normalize this measure with the total portfolio holdings of CLO i at time $t-1$. This approach gives multiple time series for each CLO in our sample and allows us to control for any unobservable factors affecting the CLO trading behavior by controlling for CLO-time.

In line with our hypothesis, we find that CLOs increase their net purchases of brown loans in periods of heightened media attention to climate risk. This increase corresponds to 0.055% on a monthly basis and is statistically significant at a 1% level. To illustrate the economic magnitude of this effect, we note that the total average monthly net purchase is 0.212%. Hence, the increase in net purchase during times of elevated attention to climate change is equivalent to 25% of the normal CLO activities. To gain additional insights, we repeat our analysis by directly comparing the net purchases of loans from issuers in brown industries to a subgroup of loans from issuers in industries with the lowest carbon emissions, which we classify as "green" loans. Using this direct comparison substantially increases the economic magnitude of the increase in net purchases, which now ranges from 37.8% to 41.0% of the average total net purchases.

We next repeat our analysis with the more granular Refinitiv-Eikon classifications. Instead of using Moody's industry classifications, we group the issuers into quintiles based on either the average carbon emissions or the emission intensity of their granular industries. We then construct *Net Purchases* for each CLO-month and each emission quintile. We find statistically and economically similar results. During periods of heightened media attention to climate change, CLOs increase their net purchases for loans from issuers with higher carbon emissions or greater emission intensities.

Our third hypothesis is that CLOs whose managers committed to considering environmental aspects in their investments do not exploit heightened media attention to climate change. To test this hypothesis, we examine if CLOs whose managers signed the PRI behave differently from other CLOs. The PRI are an initiative by the United Nations (UN) in which signatories commit to, among other things, considering environmental factors in their investment decisions. Modifying our previous analysis, we now add interaction terms controlling for whether the CLO manager signed the PRI. Contrary to our expectations, we find that managers who signed the PRI do not behave differently from managers who did not sign the PRI.

Our final hypothesis is that CLOs under bank-affiliated managers exploit media attention to climate change to a larger extent. One purpose of CLOs is to help banks conduct regulatory arbitrage. As explained by Cordell et al. (2023), holding CLO tranches instead of outright positions in leveraged loans can lower banks' regulatory capital requirements. In a similar spirit, it is plausible that holding CLO tranches instead of outright positions in brown loans helps banks veil their investments in brown industries. Confirming this hypothesis, we find that bank-affiliated CLOs increase their investments in brown industries even more during times of heightened media attention to climate change.

We conclude by examining the holdings of loan mutual funds, which are the second largest participants in the leveraged loan market. In contrast to CLOs, loan mutual funds are open-ended mutual funds in which investors can withdraw their capital on a daily basis. In addition, loan mutual funds are more visible to regulators and the public because their main investors are retail investors. Hence, investors in loan mutual funds could pressure the fund managers to divest from polluting firms or industries. We therefore expect that loan mutual funds sell when CLOs are buying.

We have access to loan fund holdings during the six months surrounding the Paris Agreement of December 2015 and investigate changes in CLO and loan mutual fund investments around this event. While this time period is a narrow focus, it has the advantage that the Paris agreement is arguably an exogenous shock to media attention to climate change.[1] We follow previous studies (e.g., Degryse et al., 2023; Delis et al., 2024; Ehlers et al., 2021; Müller and Sfrappini, 2021; Reghezza et al., 2022) and estimate a difference-in-differences model around the Paris agreement. We find that for CLOs, net purchases of high-emission industries increased compared to net purchases of loans from other industries in the month following the Paris Agreement. By contrast, loan mutual fund investments into high-emission industries decreased during the exact same period.

Contributions to the Literature. Our findings contribute to three streams of literature. First, we contribute to the extensive literature examining the influence of climate risk on asset prices. So far, this literature mainly focuses on equity and bond markets. Bolton and Kacperczyk (2021), Bolton and Kacperczyk (2023), Görgen et al. (2020), Hsu et al. (2023), among others, show that stocks of firms with higher carbon emissions generate higher returns, suggesting that equity investors demand compensation for exposure to carbon emission risk. Sautner et al. (2023) find that shocks to climate change exposure correlate negatively with firm valuations. Huynh and Xia (2021) and Huynh and Xia (2023) find that climate change news and environmental disasters affect corporate bond prices. To the best of our knowledge, this is the first paper to study the link between media attention to climate change and leveraged loan prices.

Second, we contribute to the literature linking concerns about climate change with institutional investor behavior. As discussed by Krueger et al. (2020) and Bolton and Kacperczyk (2021), among others, institutional investors are concerned about climate risks. These concerns lead mutual funds to re-direct their investments toward greener companies (Ceccarelli et al., 2024) and banks to incorporate carbon risks in their loan pricing (Delis et al., 2024). Kacperczyk and Peydro (2021) show that banks reduce their lending to firms with high carbon emissions. These divestments lead brown firms to reduce their leverage (e.g., Ginglinger and Moreau, 2023

[1] One concern with our previous analysis might be that the results are clouded by confounding events. For example, when governments decide to pass regulations to curb climate change, awareness of climate change rises while, at the same time, the regulation might change the trading environment.

and Nguyen and Phan, 2020). Although climate concerns affect institutional investors, it is not clear if banks committing to "responsible investments" behave differently from other banks (e.g., Giannetti et al., 2023 and Ehlers et al., 2021). Additionally, Müller et al. (2024) show that banks use securitization to sell brown loans to CLOs. We contribute to this literature by examining the investment behavior of CLOs during times of heightened media attention to climate risk. To the best of our knowledge, this is the first paper highlighting a scenario in which institutional investors invest more aggressively in polluting industries.

Finally, we contribute to the growing literature on the role of CLOs in financial markets, as previously studied by, among others, Benmelech et al. (2012), Bord and Santos (2015), Liebscher and Mählmann (2017), Peristiani and Santos (2019), Loumioti and Vasvari (2019), Nicolai (2020), Fabozzi et al. (2021), Bhardwaj et al. (2022), Elkamhi and Nozawa (2022), Fleckenstein (2022), Kundu (2022), Cordell et al. (2023), Elkamhi et al. (2023), Griffin and Nickerson (2023), Kundu (2024), and Emin et al. (2024). Fabozzi et al. (2021) highlight that more active CLOs outperform less active ones. Cordell et al. (2023) claim that CLOs exist for the purpose that traditional financial intermediaries can mitigate regulatory constraints and benefit from regulatory arbitrage. We contribute to this strand of literature by highlighting that CLOs tilt their portfolios toward browner industries when media attention to climate change is elevated.

4.2 Background and Hypotheses

In this section, we describe the institutional backdrop of our analysis. Building on this background, we link media attention to climate change and CLO trading behavior, deriving four testable hypotheses.

4.2.1 Institutional Background

We focus our study on CLOs, which are structured finance products that operate as special purpose vehicles (SPVs). CLOs issue equity and debt tranches with different seniority to finance their investments, which serve as collateral for the debt tranches. As is common for structured finance products, investment losses affect the debt holders according to a waterfall structure: The equity tranche absorbs the first losses, followed by the more junior debt tranches. In contrast to other structured finance products, such as Collateralized Debt Obligations (CDOs), CLOs increased substantially in popularity after the global financial crisis in 2007 and 2008. As discussed by Cordell et al. (2023), the outstanding volumes of CLOs have nearly tripled since the global financial crisis, exceeding $750 billion at the end of 2019. This sheer size makes CLOs an important player in the leveraged loan market. Leveraged loans are syndicated loans from firms with lower credit credit quality, and CLOs are the predominant investor in this market. According to Kundu (2022), CLOs purchase up to 75% of the leveraged loan issuances in the primary market. Hence, CLOs are the most important financing source for risky firms.

In the context of our analysis, the main feature of CLOs is that the manager can rebalance the collateral portfolio because they are active investors in the leveraged loan market. This rebalancing is subject to four types of constraints. First, diversification constraints require the CLO manager to spread investments across different borrowers and industries. Second, collateral quality constraints require that the share of loans with credit ratings below triple-C and the share of defaulted loans in the collateral portfolio do not exceed a pre-specified threshold. Third, the CLO is subject to regular performance tests, such as over-collateralization (OC) tests (e.g., Fabozzi et al., 2021) to ensure the safety of senior investors' cash-flows. Finally, active rebalancing is typically limited in the amortization period, which starts approximately five years after the CLO is issued. During this period, the CLO gradually repays its debt and equity tranches instead of reinvesting in new loans. We further discuss the life cycle of a CLO in the Appendix (Figure C.1).

Despite these constraints on active rebalancing, CLO managers face a limited risk of investor withdrawals. The debt tranches of the CLO have a fixed maturity and therefore resemble the debt of a closed-end mutual fund. The only potential risk of investor withdrawals is that the equity investors can "call the deal" and force the manager to unwind her positions and repay the debt tranches (starting with the most senior ones). The main CLO equity investors are hedge funds and the decision to call a deal is typically driven by the profitability of the equity tranche.

In summary, CLOs invest in loans from lower-rated companies and face limited risk of investor redemptions. This makes CLOs less constrained by environmental considerations compared to other institutional investors. Unlike banks or mutual funds, CLOs have relatively opaque collateral portfolios, investing almost exclusively in loans from lower-rated, often private, companies. As a result, CLO investors have limited insight into the environmental impact of these investments. Additionally, even if CLO investors had a full overview of these investments' environmental impact, only equity investors could actively oppose the investment decisions.

4.2.2 Hypotheses

Our first hypothesis, and the starting point of our analysis, is that negative news about climate change lead to profitable investment opportunities. Investors typically attach a premium to stocks or bonds from carbon-intensive firms (e.g., Bolton and Kacperczyk, 2023; Huynh and Xia, 2021). This premium increases when institutional investors face additional pressure to divest from polluting firms or industries. Examples of divestment pressures due to public awareness include Müller et al. (2024) and Ceccarelli et al. (2024), who examine banks and mutual funds, respectively. Moreover, Choi et al. (2020) show that retail investors sell polluting firms' stocks when they experience abnormally hot weather. Hence, we hypothesize that loans issued by polluting firms experience greater price pressure during periods of heightened media attention to climate change.

Hypothesis 1: *Brown loans trade at a discount in times of heightened media attention to climate change.*

The main conjecture in this Chapter 4 is that CLOs exploit elevated attention to climate change and purchase more brown loans. We argue that CLOs are uniquely positioned to exploit of these pricing effects because they are unregulated, non-bank entities with limited public scrutiny of their investments. As discussed in Section 4.2.1, the fact that investors cannot withdraw at short notice, combined with the opacity of their investments, allows them to invest more heavily in brown industries without facing investor redemptions. This view of CLOs as potential arbitrageurs resonates with Cordell et al. (2023), who show that CLOs often act as regulatory arbitrageurs. Building on Hypothesis 1, our second hypothesis is that CLOs exploit elevated media attention to climate change.

Hypothesis 2: *CLOs exploit price discounts during heightened attention to climate change by increasing their net purchases of loans to brown companies.*

The PRI is a network of financial firms initiated in 2006 by the UN. As of November 2023, the network comprised 5,363 signatories with approximately $121.3 trillion of AUM and is described as the "world's leading proponent of responsible investment" (PRI website). One important aspect of signing the PRI is the commitment to incorporating Environmental, Social and Governance (ESG) criteria into investment decisions. Signatories of the PRI face the additional disclosure responsibility of reporting their progress toward climate goals, which theoretically incentivizes environmentally responsible actions. However, Ehlers et al. (2021) and Giannetti et al. (2023) find that "green" banks do not behave significantly differently from the other banks. We use the date a CLO joins the PRI as an indicator of their commitment to climate goals. Despite the previous evidence, we hypothesize that CLOs change their behavior of exploiting media attention to climate change after signing the PRI.

Hypothesis 3: *CLOs whose managers signed the PRI do not increase their net purchases of loans to brown companies during heightened attention to climate change.*

Our final hypothesis revolves around the role of CLOs as regulatory arbitrageurs. Cordell et al. (2023) suggest that CLOs give banks a capital-efficient way of getting exposure to leveraged loans. To reduce their capital requirements, banks can sell leveraged loans to CLOs and subsequently invest in CLO tranches (which have lower capital requirements). Following this logic, it is plausible that banks can use CLOs to get indirect exposure to loans from carbon-intensive industries. We therefore hypothesize that CLOs of bank-affiliated managers take more advantage of heightened attention to climate change and increase their exposure to high-emission industries more strongly.

Hypothesis 4: *CLOs whose managers are affiliated with a bank increase their net purchases of loans to brown companies during heightened attention to climate change more strongly than other CLOs.*

4.3 Data and Variable Construction

In this section, we first describe the data underlying our study. We then explain our approach to identify brown firms and periods of heightened media attention to climate risks.

4.3.1 CLO and Loan Transaction Data

The data source for our analysis is the Creditflux CLO-i database. Creditflux aggregates information from CLO reports to their trustees and provides three data sets. First, details on the CLO structure, performance, and compliance with constraints. Second, monthly or quarterly CLO portfolio holdings. These holdings data also contain basic information about the underlying asset, such as maturity, issuer industry, and credit rating. Third and finally, loan transactions, executed by the CLO manager. These data include the trade date, face amount traded, trade direction (purchase or sale), transaction price, and loan type. We supplement these trading data with credit ratings, industries, and loan times to maturity from the holdings data.

Following other CLO studies (e.g., Fabozzi et al., 2021), we start our analysis with data from January 2010 because the years before 2010 contain a small amount of loan transactions (see Figure C.2 in the Appendix). We further apply the following four filters. First, we require information on the CLO's debt and equity tranches. We drop CLOs without this basic information, as their reports on holdings and transactions are often incomplete or missing. Second, we only include transactions of dollar-denominated term loans. Third, we remove transaction observations with reporting errors in the transaction size (zero or negative) and a price above \$120 or below \$15. Finally, we restrict our analysis to the CLOs' "reinvestment period". This period starts six months after the CLO issuance (after the CLO has purchased most of its collateral assets) and ends when the CLO enters the amortization period and stops reinvesting the proceeds from maturing loans. During this period, the CLO manager faces the lowest trading constraints. See the life cycle of a CLO in Figure C.1 and the distribution of trades over the life cycle of our CLOs in Figure C.3, both in the Appendix.

We focus our main analysis on the January 2010 to May 2018 period because our main measure of media attention to climate change ends in May 2018 (see Section 4.3.4 below). Over this time period, we observe 1,562 CLOs from 161 managers. Table 4.1 shows descriptive statistics of our data, and we list definitions of all variables used in the order as they appear in the text in Table C.6 in the Appendix. Panel A contains summary statistics for the CLOs in our sample. The average CLO has \$545 million in assets under management (AUM) and is almost three years old, measured after closing date.[2] We also retrieve a list of signatories from the PRI website (accessed August 2, 2023) and find that 14.3% of the CLO-month observations in our sample belong to CLOs from managers who have signed the PRI (Table C.1 in the Appendix contains a list of PRI signatories). The average Herfindahl-Hirschman Index (HHI) of the CLO portfolio is 0.007 and suggests strong diversification across borrowers. Further, we convert the credit rating to numerical values with AAA corresponding to 23, AA to 22, . . . , C to 3, and D to 2. We then compute the weighted average rating for each CLO portfolio. The last row in Panel A shows the average portfolio credit rating is 9.7, which corresponds to a rating between B+ and B. This relatively low rating confirms that CLOs mainly invest in leveraged loans, which

[2]See the life cycle of a CLO in Figure C.1 in the Appendix.

Table 4.1: **Summary statistics.** In Panel A, we report summary statistics on CLO's assets under management (AUM), industry concentration (HHI), portfolio rating, and on whether the CLO manager has signed the PRI. In Panel B, we report summary statistics for the sample underlying equation 4.2, which encompasses daily purchases during investment periods of CLOs over the time period 1 January 2010 - 31 May 2018. Price is reported as % of notional amount. We follow Fabozzi et al. (2021) and exclude observations where price=100. Face amount is notional amount in log $. TTM is time to maturity in months. We encode the rating according to AAA=23, AA+= 22 . . . C=3, D=2. The mean rating is 9.671, (i.e. between B and B+, meaning issuers are, on average, rated below investment grade).

	#Obs	Mean	Std. dev.	Min	Max
Panel A: CLO data					
AUM (million USD)	24,598	545.483	299.235	0.000	12558.750
Age (years)	27,909	2.988	1.714	0.058	10.340
Signed PRI	27,909	0.143	0.350	0.000	1.000
HHI	24,596	0.007	0.012	0.002	1.000
Portfolio Rating	18,044	9.734	0.264	4.577	14.000
Panel B: Transaction data					
Price	317,660	98.680	3.178	60.500	101.630
Face amount (logs)	317,660	13.560	1.217	-27.726	18.419
TTM (months)	317,660	63.282	18.104	0.000	186.000
Rating	317,660	9.671	1.212	2.000	22.000

typically have a credit rating below investment grade.

We observe a total of 759,718 transactions (purchases or sales) during CLOs' reinvestment period over our sample period of which 467,518 or 61.5% are purchases. This higher amount of purchases compared to sales is in line with Kundu (2022) and reflects CLO managers' practice of reinvesting loan proceeds as loans mature. Panel B of Table 4.1 contains summary statistics for the transactions in our sample. We report summary statistics for *purchases* of loans, discarding transactions at a price of exactly $100,[3] and report prices for loans that CLOs purchased during their reinvestment period. Loan prices are expressed as percentages of their face value. The average price of a loan purchased by CLOs between 1 January 2010 and 31 May 2018 is 98.68% of its notional amount, suggesting loans are, on average, purchased at a discount to their face value. The average notional amount and time-to-maturity (expressed in months) are $\$1,337$ million (13.6 in logs) and 63 months. The average loan rating is 9.671 and comparable to the average loan rating of the CLO portfolio.

[3] We following Fabozzi et al. (2021) who exclude these loan transactions as they are likely not genuine transactions but are internal rebalancing across CLO families.

4.3.2 Classifying Brown and Green Borrowers

For our empirical analysis, we need to distinguish brown firms from other firms. Doing so is challenging because the only available borrower information in CLO-i is the issuer name and the Moody's industry classification. We therefore begin our analysis by using the available industry classifications to group borrowers into "brown" and "other" firms. Although it would be ideal to use firm-level emission data or environmental ratings when identifying brown borrowers, a vast majority of loans in our sample are issued by private companies with virtually no available information on their operations. Hence, it is plausible that investors in the leveraged loan market also use the available industry classifications to determine whether a firm is considered brown.

To group the Moody's industry classifications based on their environmental impact, we draw on the work of Bolton and Kacperczyk (2021), who use greenhouse gas emissions from individual companies to estimate an industry's overall environmental impact. Bolton and Kacperczyk (2021) use emissions from production (Scope 1); indirect emissions from consumption of purchased electricity, heat or steam (Scope 2); and other indirect emissions from the production of purchased materials, such as product use and waste disposal (Scope 3) as their industry classifications. The classifications from Bolton and Kacperczyk (2021) are based on Global Industry Classification Standards (GICS), which differ slightly from Moody's industry classifications. Tables C.2 and C.3 in the Appendix give a detailed overview of our approach to linking the GICS to Moody's industry classifications. Table 4.2 provides an overview of the different industries in our sample and our ranking based on environmental impact.

Based on our matching approach, we identify eight brown industries. As shown in Table 4.2 these industries are: (i) Automobile; (ii) Cargo Transport; (iii) Chemicals & Plastics & Rubber; (iv) Containers & Packaging & Glass; (v) Mining & Steel & Iron & Non-Precious Metals; (vi) Oil & Gas; (vii) Personal Transportation; and (viii) Utilities. We further identify four industries with evidently low environmental impact from their operations: (i) Banking; (ii) Broadcasting & Entertainment; (iii) Finance; (iv) Healthcare, Education, Childcare. Additionally, Table 4.2 lists the remaining industries that we categorize as neither brown nor green.

Our industry classification is a pragmatic categorization. Given that most firms in our sample are private, it is likely that investors do not have additional information on these firms' climate change risk and rely on a similar industry classification. Hence, even though this is a broad categorization and could understate our findings, it is likely that CLOs and other loan

Table 4.2: **Industry summaries.** This table contains descriptive statistics of the transactions in our sample across 34 Moody's industries in *Panel A*. $Number\ Issuers$ is the number of issuers per industries. $Fraction\ Public$ is the fraction of firms that are publicly traded and for which we see balance sheet information in Eikon. $Total\ Purchases$ is the sum of all purchases in billion USD. $Net\ Purchase_{i,b,t}$ is defined as in Equation (4.1) in *Panel A*: $Net\ Purchase_{i,r,t} = \frac{(Notional\ Purchases_{i,r,t} - Notional\ Sales_{i,r,t})}{DealBalance_{i,t-1}}$, where i is CLO, r is industry, and t is year-month. In *Panel B*, we use the Refinitiv-Eikon emission data described in Section 4.3.2 and put the firms in our sample into quintiles according to their Eikon industries' emission intensities. We then define net purchases analogous to Panel A, replacing industries with emission quintiles.

				Net Purchases			
Industry	Number Issuers	Fraction Public	Total Purchases	mean	25%	50%	75%
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Industries							
Brown industries							
Automobile	205	0.15	23.39	0.14	-0.06	0.12	0.35
Cargo Transport	102	0.13	7.57	0.11	-0.06	0.09	0.25
Chemicals, Plastics and Rubber	204	0.14	26.90	0.13	-0.07	0.12	0.34
Containers, Packaging and Glass	122	0.12	18.25	0.12	-0.07	0.11	0.31
Mining, Steel, Iron and Non-Precious Metals	73	0.27	8.42	0.07	-0.09	0.08	0.26
Oil and Gas	221	0.15	16.34	0.10	-0.09	0.10	0.29
Personal Transportation	8	0.12	0.10	0.02	-0.13	0.01	0.21
Utilities	147	0.22	17.91	0.11	-0.07	0.11	0.32
Green industries							
Banking	290	0.16	48.34	0.20	-0.05	0.17	0.45
Broadcasting and Entertainment	117	0.15	34.29	0.17	-0.07	0.16	0.43
Finance	101	0.09	2.64	0.08	-0.10	0.09	0.29
Healthcare, Education and Childcare	637	0.14	75.48	0.25	-0.07	0.22	0.57
Other industries							
Aerospace and Defense	187	0.10	17.55	0.12	-0.07	0.11	0.32
Beverage, Food and Tobacco	276	0.11	24.52	0.12	-0.09	0.10	0.31
Buildings and Real Estate	200	0.14	16.38	0.10	-0.06	0.10	0.28
Diversified Natural Resources, Precious Metals and Minerals	10	0.00	-0.31	-0.20	-0.42	-0.19	0.11
Diversified/Conglomerate Manufacturing	248	0.11	22.63	0.11	-0.08	0.09	0.29
Diversified/Conglomerate Service	555	0.10	58.18	0.19	-0.07	0.16	0.46
Ecological	70	0.10	5.20	0.06	-0.07	0.06	0.20
Electronics	408	0.18	64.68	0.23	-0.06	0.19	0.51
Farming and Agriculture	11	0.18	0.28	0.04	-0.11	0.09	0.23
Grocery	7	0.14	0.23	0.07	-0.14	0.10	0.29
Home and Office Furnishings, Housewares and Durable Consumer Products	32	0.00	0.09	-0.02	-0.20	-0.01	0.19
Hotels, Motels, Inns and Gaming	224	0.14	34.48	0.15	-0.09	0.13	0.39
Insurance	15	0.27	0.57	0.13	-0.06	0.16	0.38
Leisure, Amusement and Entertainment	40	0.17	0.80	0.05	-0.15	0.11	0.34
Machinery (Non-Agriculture, Non-Construction and Non-Electronic)	169	0.17	18.92	0.10	-0.07	0.09	0.27
Personal and Non-Durable Consumer Products (Manufacturing Only)	69	0.16	2.37	0.09	-0.08	0.10	0.26
Personal, Food and Miscellaneous Services	56	0.11	1.84	0.11	-0.08	0.12	0.30
Printing and Publishing	138	0.08	8.19	0.03	-0.12	0.05	0.25
Retail Stores	229	0.17	33.43	0.14	-0.11	0.12	0.39
Sovereign and Supranational	13	0.15	0.23	0.08	-0.21	-0.06	0.35
Telecommunications	191	0.18	29.84	0.14	-0.10	0.12	0.38
Textiles and Leather	15	0.27	0.17	0.16	0.01	0.14	0.28
Panel B: Activity rankings							
Q1				0.39	-0.03	0.28	0.70
Q2				0.27	-0.03	0.20	0.50
Q3				0.35	-0.05	0.24	0.64
Q4				0.39	-0.01	0.29	0.71
Q5				0.29	-0.06	0.20	0.53

investors use a similar method to understand private firms' carbon footprint.

In addition to our industry classifications, we match each issuer in our sample to Refinitiv-Eikon. To that end, we first match the loan transactions in our sample to Thomson-Reuters Dealscan and use the Refinitiv identifiers provided in Dealscan for the matched loans. For loans that we cannot match to Dealscan, we manually search for the borrower in Eikon and find their identifier. We can match approximately 85% of the borrowers in our sample to Eikon. Column (3) of Table 4.2 shows the fraction of borrowers in each industry with publicly available information. On average, 14% of the issuers in our sample have either equity prices or report their total assets in Eikon. Hence, more than 85% of the issuers in our sample are private companies with limited information about their operations.

Given this small subsample of public companies, we do not focus on individual companies' carbon emissions. Instead, we rank the companies in our sample based on the carbon emissions of their industries, proceeding in three steps. First, we use the Eikon identifiers to obtain "the Refinitiv Business Classification" (TRBC) on the activity level, which provides the most granular industry classification in Eikon and contains 494 activities for the 5,390 issuers in our sample. Second, every year and for each activity, we compute the average sum of Scope 1, 2, and 3 carbon emissions for all loan borrowers in Dealscan. Third, based on these activity group averages we put the firms in our sample into quintiles with the first quintile containing firms with the lowest carbon emissions and the fifth quintile containing firms with the highest carbon emissions. We repeat this ranking procedure using the average emission intensity, measured as the ratio of carbon emissions to market capitalization, instead of the average emissions.

4.3.3 Measuring Net Purchases

To measure the trading behavior of CLOs during heightened attention to climate change, we construct the variable *Net Purchase* as follows:

$$Net\,Purchase_{i,r,t} = \frac{Notional\,Purchases_{i,r,t} - Notional\,Sales_{i,r,t}}{DealBalance_{i,t-1}}, \tag{4.1}$$

where *Notional Purchases* is the sum of the notional of all loans in industry r, purchased by CLO i in month t. Similarly, *Notional Sales* is the sum of the notional of all loans in industry r, sold by CLO i in month t. Deal Balance measures total holdings of leveraged loans by CLO

i in the previous period, t-1. For most CLOs, this variable is only available quarterly. Hence, we use the most recent available deal balance in Equation (4.1). Throughout this Chapter 4, we winsorize *Net Purchase* at the 1% and 99% percentile to avoid our results being driven by large outliers.

We report summary statistics of *Net Purchases* in Columns (4) to (7) of Table 4.2. The numbers are expressed as percentages of the deal balance. As we can see from Table 4.2, the average net purchases are positive for all but two industries and exhibit large cross-sectional variation ranging from −0.20 for "Natural Resources" to 0.23 for "Healthcare". Within industries, *Net Purchases* also exhibits significant variation, with negative 25% quantiles for all but one industry and positive 75% quantiles for all industries. Similar to the net purchases of different industries, Panel B of Table 4.2 provides an overview of the emission intensity quintiles. As we can see from the table, each quintile shows a large time-series variation with a negative 25% quantile and a 75% quantile above 0.5.

4.3.4 Measuring Attention to Climate Change

To identify episodes of elevated attention to climate change, we use the Crimson Hexagon Negative News Index (CHNeg) index, developed by Engle et al. (2020). This index is derived from the Crimson Hexagon database, which includes a collection of newspaper articles from the WSJ, the NYT, The Washington Post, Reuters, BBC, CNN, and Yahoo News. Engle et al. (2020) search for the keyword "climate change" and let the Crimson Hexagon database apply sentiment analysis to obtain negative news related to climate change. The index is available on a monthly basis from June 2008 to May 2018.

Figure 4.1 shows the monthly time series of the index. While the original index is available from 2008 onward, we restrict Figure 4.1 to our sample period and calculate the 80% quantile of the index based on data between January 2010 and May 2018. The red dots in Figure 4.1 highlight months with "attention spikes" when the index exceeds its 80% quantile. For our main analysis, we define periods of heightened media attention to climate change as periods when the CHNeg index exceeds its 80% quantile.

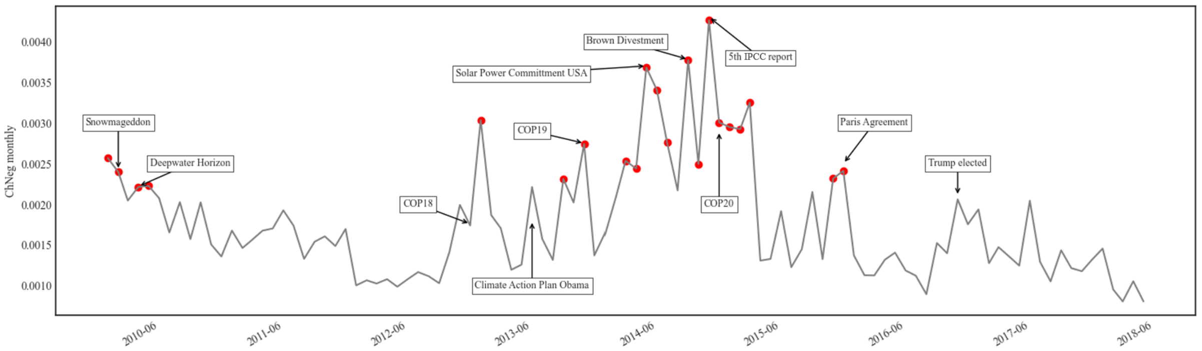

Figure 4.1: **News attention index.** This figure was created using the monthly Crimson Hexagon Negative News Index (CHNeg) of Engle et al. (2020). Red dots mark the 80th percentile for each index. Events are the following: Snowmaggeddon (large snowstorm in Washington, DC); Deepwater Horizon (large oil spill in the Gulf of Mexico); COP 18 in Qatar, which faced criticism as Qatar is the largest greenhouse gas emitter per capita; Obama releases Climate Action Plan to reduce green house gas emissions; COP 19 in Poland shows no crucial outcomes; Obama announces the nation's large solar power commitments; Brown Divestment (over 800 global investors announce fossil fuel divestment); IPCC's fifth assessment report (science report and basis for climate-relevant decisions); COP20 in Peru preparing Paris; Paris Agreement commits world leaders to limit global warming to 1.5 degrees Celsius; Election of President Donald Trump.

4.4 Loan Prices and Media Attention to Climate Change

In this section, we test Hypothesis 1 by investigating how heightened attention to climate change affects the prices of brown loans. Because investors in brown loans might be more eager to reduce their brown loan holdings, we expect CLOs to purchase brown loans at a discount. To test this assertion, we estimate the following regression equation:

$$Price_{l,r,\tau} = \beta\, Brown_r \times Attention_{t-1} + \gamma_x\, \mathbf{X}_t + \alpha_\tau + \alpha_{rat} + \alpha_r + \alpha_i \times \alpha_t + \epsilon_{l,r,\tau}, \tag{4.2}$$

where $Price_{l,r,\tau}$ is the price of loan l from industry r on day τ. The main variable of interest is the interaction between two indicators $Brown_r$, which equals one if industry r is classified as brown, and $Attention_{t-1}$, which equals one for periods where the CHneg index by Engle et al. (2020) is above its 80 % quantile. This interaction captures the prices of brown loans during times of heightened attention to climate change.

In all specifications, we follow Fabozzi et al. (2021) and include the following six controls and fixed effects: (i) the logarithm of the notional amount traded; (ii) the months until the loan matures; (iii) loan rating FE fixed effects (α_{rat}); (iv) issuer industry fixed effects (α_r); (v) CLO fixed effects (α_i); and (vi) year-month fixed effects (α_t). For a clear interpretation of the regression coefficients, we only include the prices of purchases during the CLOs' reinvestment phase in our analysis. Finally, throughout our analysis, we double-cluster the standard errors on

Table 4.3: **Prices**. In this table we report results from estimating the following equation (equation 4.2): $Price_{l,r,\tau} = \beta\ Attention_{t-1} \times Brown_r + \gamma_x\ \mathbf{X}_t + \alpha_\tau + \alpha_{rat} + \alpha_r + \alpha_i \times \alpha_t + \epsilon_{l,r,\tau}$. The sample covers daily purchases during reinvestment periods of CLOs over the period 1 January 2010 - 31 May 2018. Price is reported as % of the notional amount. We follow Fabozzi et al. (2021) and exclude observations where price=100. *Brown* is an indicator which equals 1 for brown loans from six industries, defined as in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is more than 80 % above the whole sample period. We include the following controls ($\mathbf{X}_t$): face amount (logs), time-to-maturity, CLO portfolio concentration measured by the Herfindahl score of industry concentration (HHI) and CLO portfolio rating in columns (2) and (5). We use loan rating FE fixed effects (α_{rat}), issuer industry fixed effects (α_r), CLO fixed effects or CLO-year-month fixed effects ($\alpha_i \times \alpha_t$). We also use date fixed effects (date of trade, (α_t)) in columns (2) and (5). Standard errors are clustered at the CLO level and year-month level and are reported in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	brown versus all			brown versus green		
$Attention_{t-1} \times Brown_r$	-0.915***	-0.912***	-0.915***	-0.868***	-0.864***	-0.932***
	(0.281)	(0.281)	(0.271)	(0.329)	(0.328)	(0.301)
$Face\ amount_l$	0.086***	0.086***	0.078***	0.053***	0.053***	0.051***
	(0.012)	(0.012)	(0.012)	(0.015)	(0.015)	(0.014)
$Months\ to\ maturity_l$	0.014***	0.014***	0.016***	0.008***	0.008***	0.009***
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
$HHI_{i,t-1}$		11.128			3.251	
		(8.741)			(7.575)	
$Portfoliorating_{i,t-1}$		0.245**			0.322**	
		(0.114)			(0.144)	
Number of Observations	317,660	317,660	317,660	131,723	131,723	131,723
Adjusted R2	0.644	0.644	0.695	0.686	0.686	0.752
Mean Dependent	98.68	98.68	98.68	98.679	98.679	98.679
SD Dependent	3.178	3.178	3.178	3.217	3.217	3.217
Rating FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
CLO FE	Yes	Yes	-	Yes	Yes	-
Year-Month FE	Yes	Yes	-	Yes	Yes	-
CLO-Year-Month FE	-	-	Yes	-	-	Yes
Date FE	-	-	Yes	-	-	Yes

the CLO and year-month level.[4]

Column (1) of Table 4.3 shows the results of our baseline analysis. The regression coefficient on the interaction between brown loans *(Brown)* and elevated media attention to climate change *(Attention)* is negative and statistically significant at the 1% level. The average loan transaction

[4]An alternative would be triple-clustering at the CLO, date, and industry level. However, as discussed by Angrist and Pischke (2009), clustering with a small number of groups can result in less reliable standard errors. Moreover, triple-clustering would have a minor effect on our results. For instance, the standard error in Column (1) would increase from 0.249 to 0.283.

price is \$98.680 (i.e., the average loan trades at a discount) and brown loans trade at a discount of \$0.915 during times of elevated media attention, corresponding to a discount of 0.93%.

Next, we add controls for CLO portfolio concentration, measured by the portfolio HHI index, and CLO portfolio rating as controls. The results remain unchanged. Finally, we tighten our regression specification and interact the CLO fixed effects with year-month fixed effects to capture any potential fluctuations in CLO trading behavior over time. We further include date fixed effects to capture potential fluctuations in loan prices within the year-months. As we can see from column (3), these additional fixed effects do not affect the economic significance of media attention to climate change on brown loans. Brown loans trade at a discount of \$0.915, during times of elevated media attention, corresponding to a discount of 0.93%. In columns (4) - (6) we drop the sub sample of industries that are neither defined as brown nor green and directly compare brown loans to green loans. When considering the most conservative model including CLO-year-month fixed effects and date fixed effects in column (6), the discount increases to \$0.932, or 0.94%, respectively.[5]

We conclude this section by examining the relative transaction prices of brown loans as a function of media attention to climate change. Specifically, we focus on the regression specification used in Column (3) of Table 4.3 and estimate the regression coefficient on the interaction between brown loans and an indicator variable that equals one if media attention to climate change is its i-th quintile. As we can see from Figure 4.2, the more intense the attention to climate change, the larger the discount for brown loans in comparison to non-brown loans.

4.5 CLOs Exploit Media Attention to Climate Change

In this section, we test Hypotheses 2-4. Starting with Hypothesis 2, we first investigate if CLOs increase their net purchases of loans from issuers in brown industries when media attention to climate change is elevated. We use the proxy *Net Purchase* defined in Equation (4.3) and control for CLO times year-month fixed effects in all specifications. This addresses the following two issues. First, CLOs face complex incentives and constraints, which might lead them to buy and sell for reasons such as test breaches, par building, or enhancing performance through active

[5] In the Appendix in Table C.4, we present results excluding industry fixed effects, which shows us that in general, there is no price discount between brown loans and non-brown loans (column (1), coefficient for *Brown* is not distinguishable from zero). But if we compare the extremes, i.e. brown versus green loans, we find an average price discount of -0.229 which compares to a sample mean of 98.679, see column (2).

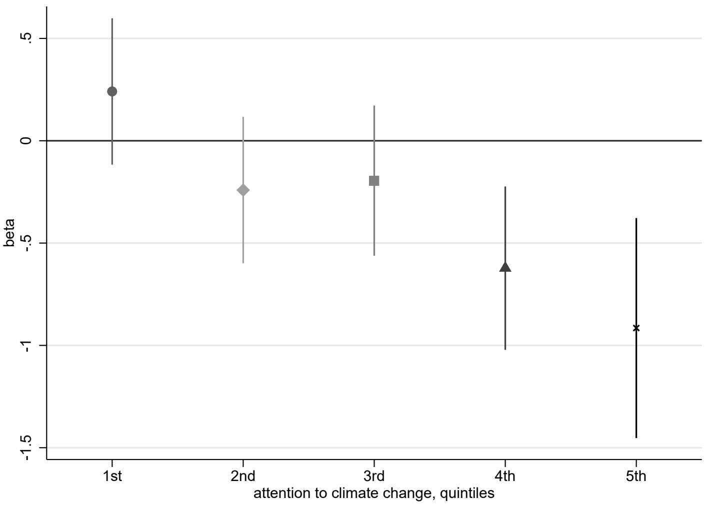

Figure 4.2: **Intensity of attention to climate change and prices.** In this figure we present results from estimating the following equation (Equation 4.2): $Price_{l,r,\tau} = \beta\ Attention_{t-1} \times Brown_r + \gamma_x\ \mathbf{X}_t + \alpha_\tau + \alpha_{rat} + \alpha_r + \alpha_i \times \alpha_t + \epsilon_{l,r,\tau}$. The sample covers daily purchases during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. Price is reported as % of the notional amount. We follow Fabozzi et al. (2021) and exclude observations where price=100. *Brown* is an indicator that equals 1 for brown loans from six industries, defined as in Section 4.3. We define $Attention_{t-1}$, or "attention to climate change", as a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above the first quintile (1st), above second quintile (2nd), and so on, up to the fifth quintile (5th) over the whole sample period. We include log face amount of the loan as well as time-to-maturity as controls ($\mathbf{X}_t$). We use loan rating FE fixed effects (α_{rat}), issuer industry fixed effects (α_r), CLO-year-month fixed effects ($\alpha_i \times \alpha_{ym}$), as well as date fixed effects (date of trade, α_τ). We cluster standard errors on CLO and year-month level. We report confidence intervals at the 5% significance level.

trading. Including the interaction between CLO and time fixed effects controls for this behavior. Second, focusing only on brown purchases could be misleading because heightened attention to climate change might correlate with tighter credit conditions, during which CLOs tend to purchase more. By including CLOs' portfolio tilting in other (non-brown) industries as a control group, we can test whether CLOs increase their net purchases of brown loans relative to their net purchases of other loans.

4.5.1 Analysis Based on Moody's Industry Groups

In this section, we estimate the following regression equation:

$$Net\ Purchases_{i,r,t} = \beta\ Attention_{t-1} \times\ Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}, \tag{4.3}$$

where $Net\ Purchases_{i,r,t}$ are the net purchases, as defined in Equation (4.1), from CLO i in industry r at time t. As with the previous pricing tests, our main focus is on the interaction term between the two indicators $Brown_r$ (equal to one if industry r is classified as brown) and $Attention_{t-1}$ (equal to one if the CHneg index in month $t-1$ is above its 80% quantile). The regression coefficient β for this interaction shows us whether CLOs change their trading behavior of brown industries compared to non-brown industries when media attention to climate change is high. We control for CLO fixed effects interacted with year-month fixed effects ($\alpha_i \times \alpha_t$) and industry fixed effects (α_r). As before, we double-cluster the standard errors on CLO and year-month levels.

Table 4.4: **Net purchases.** In this table we report results from estimating the following equation (equation 4.3): $Net\,Purchase_{i,r,t} = \beta\, Attention_{t-1} \times Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during reinvestment periods of CLOs over the period 1 January 2010 - 31 May 2018. $Net\,Purchase_{i,b,t}$ is defined as in equation (4.1). $Brown$ is an indicator, which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above 80 % over the whole sample period. In columns (1) and (2), we compare net purchases in brown industries to all other industries. In columns (3) and (4), we restrict the sample to brown and green industries, only. We cluster standard errors on the CLO and year-month levels and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	brown versus all		brown versus green	
$Attention_{t-1} \times Brown_r$	0.055***	0.054***	0.082***	0.087***
	(0.015)	(0.016)	(0.022)	(0.023)
Number of Observations	294,265	293,476	122,097	119,751
Adjusted R2	0.087	0.161	0.097	0.167
Mean Dependent	0.212	0.212	0.217	0.216
SD Dependent	0.531	0.53	0.543	0.539
Industry FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
Year-Month FE	Yes	-	Yes	-
CLO-Year-Month FE	-	Yes	-	Yes

We estimate Equation (4.3) to examine how heightened media attention to climate change affects CLOs' net purchases of loans from issuers in brown industries. First, we include net purchases for all available industries and control for industry, CLO, and year-month fixed effects. As shown in Column (1), when media attention to climate change is high, CLOs increase their net purchases of loans to borrowers in brown industries. The effect is statistically significant at a 1% and corresponds to 0.055% of the CLO portfolio. To interpret the economic significance of this estimate, it is important to note that the average monthly net purchase is 0.212%. Hence, this estimate suggests that CLOs allocate 25% more of their net purchases to loans from issuers in brown industries when media attention to climate change is elevated. Column (2) shows that the results remain virtually unchanged when we control for CLO's general trading behavior by interacting CLO fixed effects with year-month fixed effects.

As an additional check, we compare the net purchases of brown loans to net purchases of green loans. As we can see in Columns (3) and (4), the coefficients for this comparison increase substantially. CLOs increase their net purchases of brown loans by between 0.082% and 0.087%. In terms of economic magnitude, this coefficient estimate translates to an increase of between 37.8% and 41.0% compared to the average net purchases.

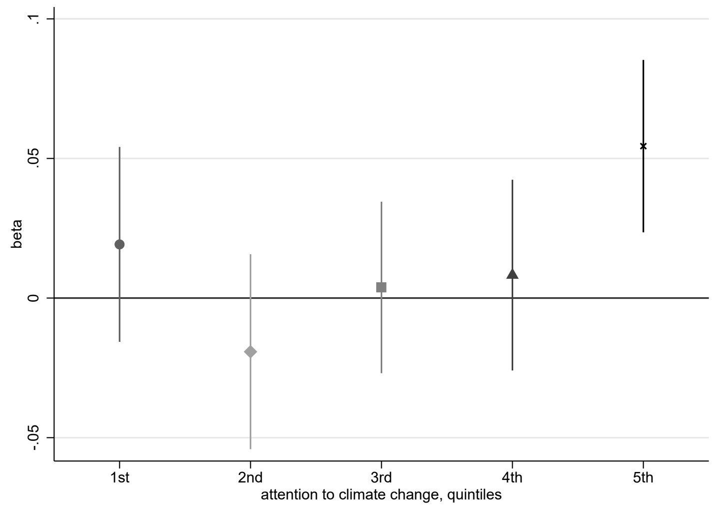

Figure 4.3: **Intensity of attention to climate change and net purchases.** In this figure we present results from estimating the following equation (Equation 4.3): $Net\ Purchases_{i,r,t} = \beta\ Attention_{t-1} \times Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during reinvestment periods of CLOs over the period 1 January 2010 - 31 May 2018. $Net\ Purchases_{i,r,t}$ is defined as in Equation 4.1. *Brown* is an indicator which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above the 1st, 2nd ... 5th quintile over the whole sample period. We cluster standard errors on the CLO and the year-month levels.

To conclude, we estimate Equation (4.3) with *Net Purchases* as dependent variable with varying levels of media attention to climate change. As shown in Figure 4.3, CLOs do not change net purchases of brown loans compared to non-brown loans for low levels of media attention to climate change. The effect on purchases of brown loans only becomes significant during months with elevated attention to climate change (i.e., when the CHneg index is above its 80% quantile).

4.5.2 Analysis Based on Eikon Classifications

In this section, we repeat our analysis using the emission rankings based on carbon emissions data from Refinitiv-Eikon. We modify Equation (4.3) and use *Net Purchases* defined for each emission quintile as dependent variable. Further, we replace $Brown_r$ with *High_Emissions*, which ranges from 0.2 for firms in the lowest emission quintile to 1 for firms in the highest emission quintile.

Table 4.5: **Net purchases (emissions).** In this table we report results from estimating the following equation (equation (4.3) on the basis of emission quintiles q): $Netpurchase_{i,q,t} = \beta\, Attention_{t-1} \times High\_emissions_q + \alpha_i \times \alpha_t + \alpha_q + \epsilon_{i,q,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per emission quintile and CLO i in year-month t during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. $NetPurchase_{i,q,t}$ is defined as in equation (4.1) but, instead of industries, we aggregate on emission quintiles of activities q. *High_emissions* are quintiles in terms of emission intensity of the activity of the issuer as defined in Section 4.3.2. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above 80 % over the whole sample period. In columns (1) and (2), we use emission intensity to rank activities. In columns (3) and (4), we use average emissions to rank activities. We cluster standard errors on the level of the CLO and year-month levels and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	Emission intensity		Average emissions	
$Attention_{t-1} \times High\_emissions_q$	0.184***	0.186***	0.162**	0.161**
	(0.054)	(0.057)	(0.077)	(0.079)
Number of Observations	109,872	108,285	108,181	106,515
Adjusted R2	0.103	0.255	0.099	0.232
Mean Dependent	0.378	0.379	0.382	0.384
SD Dependent	0.771	0.769	0.791	0.789
Rank FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
Year-Month FE	Yes	-	Yes	-
CLO-Year-Month FE	-	Yes	-	Yes

Starting with rankings based on emission intensities, Columns (1) and (2) of Table 4.5 show that CLOs' net purchases of loans with higher emissions are greater during times of elevated media attention to climate change. As before, we first use CLO and year-month fixed effects in Column (1) and then tighten the specification using the interaction between CLO fixed effects and year-month fixed effects. We find a significant link between *Net Purchases* and *High_Emissions* during periods of heightened media attention to climate change. Specifically, a one quintile increase in emission intensities corresponds to an increase of CLO purchases by 3.68% (= 0.184 × 0.2). To put these numbers into perspective with the results in Section 4.5.1, we note that When comparing the lowest quintile to the highest quintile, we move 0.8 units up in terms of $High\_Emissions$ (from 0.2 to 1), and therefore net purchases increase by 0.8 × 0.186 = 0.1488, or by 39% in comparison to the sample mean of 0.379. This is the same magnitude as we find in Section 4.5.1 when comparing brown to green loans.

To conclude, we repeat our analysis using average total emissions instead of emission intensities. Columns (3) and (4) show that the results remain largely unchanged. CLOs tilt their portfolios more toward issuers with higher emission rankings when media attention to climate change is heightened.

4.5.3 Signatories of the PRI Do Not Behave Differently

In theory, and according to Hypothesis 3, CLOs managed by PRI signatories should behave differently from CLOs managed by non-signatories. To test this hypothesis, we modify the analysis from Section 4.5.1 by interacting all variables with another indicator that equals one when the CLO manager joined the PRI (including the previous month). Specifically, we run regressions of the following form:

$$\begin{aligned} Net\ Purchase_{i,r,t} = {} & \beta\ Attention_{t-1} \times Brown_r + \\ & + \gamma\ Attention_{t-1} \times Brown_r \times PRI_{t-1} + \cdots + (\alpha_i \times \alpha_t) + \alpha_r + \epsilon_{i,r,t} \end{aligned} \tag{4.4}$$

According to Hypothesis 3, CLOs whose managers committed to considering sustainability criteria in their investment decisions should not tilt their portfolios toward brown loans in periods of high attention to climate change. Hence, we expect to find a negative effect on the triple interaction $Attention \times Brown \times PRI$ captured by the coefficient γ. We estimate Equation (4.4) with $Net\ Purchase$ as dependent variable and present the results in Table 4.6.

The interaction $Brown \times Attention$ remains positive and significantly different from zero at virtually the same level as in our baseline specification. Contrasting with our hypothesis, the coefficient γ is insignificant and close to zero in all specifications. In addition, CLO managers who joined the PRI the previous quarter do not invest differently across brown or non-brown loans in general. Thus, CLOs managed by PRI signatories do not behave differently than those

Table 4.6: **Net purchases and PRI membership**. In this table we report results from estimating the following equation (equation 4.3): $Net\ Purchase_{i,r,t} = \beta\ Attention_{t-1} \times Brown_r + \gamma\ Attention_{t-1} \times Brown_r \times PRI_{t-1} + \cdots + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. $Net\ Purchase_{i,r,t}$ is defined as in equation 4.1. *Brown* is an indicator which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above 80 % over the whole sample period. PRI_{T-1} is an indicator of whether the CLO manager is a member or has joined the PRI the previous month. In columns (1) and (2), we compare net purchases in brown industries to all other industries. In columns (3) and (4), we restrict the sample to brown and green industries, only. We cluster standard errors on the CLO and year-month levels and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	brown versus all		brown versus green	
$Attention_{t-1} \times Brown_r$	0.053***	0.053***	0.089***	0.087***
	(0.016)	(0.016)	(0.020)	(0.021)
$Attention_{t-1} \times Brown_r \times PRI_{t-1}$	0.009	0.007	-0.009	0.005
	(0.029)	(0.029)	(0.033)	(0.035)
PRI_{t-1}	0.020		0.001	
	(0.018)		(0.022)	
$Attention_{t-1} \times PRI_{t-1}$	-0.005		0.014	
	(0.023)		(0.029)	
$Brown_r \times PRI_{t-1}$	-0.014	-0.013	-0.015	-0.016
	(0.011)	(0.010)	(0.015)	(0.015)
Number of Observations	294,259	293,471	136,448	134,297
Adjusted R2	0.087	0.161	0.096	0.167
Mean Dependent	0.212	0.212	0.210	0.209
SD Dependent	0.531	0.530	0.531	0.527
Industry FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
Year-Month FE	Yes	-	Yes	-
CLO-Year-Month FE	-	Yes	-	Yes

managed by non-signatories.

Our result is in line with three other studies: (i) Gibson Brandon et al. (2022) show US signatories do not meet the PRI standards; (ii) Kim and Yoon (2023) find that PRI signatories do not differ from non-signatories in terms of ESG performance; (iii) Hale et al. (2024) show that banks who signed the PRI do not behave differently than banks that did not sign the PRI.

4.5.4 CLOs Affiliated With Banks React More Strongly

In this section, we test our fourth hypothesis and examine whether CLOs of bank-affiliated managers behave differently in the context of attention to climate change. Similarly to the previous section, we estimate the following triple interaction model:

$$\begin{aligned} Net\,Purchase_{i,r,t} &= \beta\, Attention_{t-1} \times Brown_r + \\ &+ \gamma\, Attention_{t-1} \times Brown_r \times Affiliated_i + \cdots + (\alpha_i \times \alpha_t) + \alpha_r + \epsilon_{i,r,t} \end{aligned} \tag{4.5}$$

where $Affiliated_i$ is an indicator variable that equals one for CLOs whose managers are affiliated with a bank. 28% of the CLOs in our sample belong to a bank-affiliated manager.

Table 4.7 shows the results of our analysis. While $Brown \times Attention$ remains positive and statistically significant, the coefficient β decreased compared to our baseline specification. The coefficient γ that captures the triple interaction we are interested in is positive and statistically significant. Hence, bank-affiliated CLO managers tilt their loan investments more towards brown loans in times of heightened attention to climate change. In addition, $Attention \times Affiliated$ is negative and statistically significantly different from zero in the specifications where we do not control for the interaction between CLO and year-month fixed effects. Hence, CLOs that are affiliated with a bank tend to reduce net purchases of all loans in times of heightened attention to climate change.

4.6 Loan Mutual Funds and the Paris Agreement

In Section 4.2.1, we highlight the competitive advantages of CLOs for investing in polluting firms. Loan mutual funds are the second largest participants in the leveraged loan market. In contrast to CLOs, loan mutual funds are open-ended mutual funds in which investors can withdraw their capital on a daily basis. In addition, loan mutual funds are more visible to regulators and the public because their main investors are retail investors. Hence, investors in loan mutual funds could pressure fund managers to divest from polluting firms or industries. We therefore expect a qualitatively different investment behavior for loan mutual funds.

We obtain data from the Morningstar Database on month-end holding weights of all US mutual funds invested in bank loans for the short sample period from September 2015 to

Table 4.7: **Net purchases and bank affiliation.** In this table we report results from estimating the following equation (equation 4.3): $Net\ Purchase_{i,r,t} = \beta\ Attention_{t-1} \times Brown_r + \gamma\ Attention_{t-1} \times Brown_r \times Affiliated_i + \cdots + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. $Net\ Purchase_{i,r,t}$ is defined as in equation 4.1. $Brown$ is an indicator which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above 80 % over the whole sample period. $Affiliated_i$ is an indicator of whether the CLO manager is affiliated with a bank. In columns (1) and (2), we compare net purchases in brown industries to all other industries. In columns (3) and (4), we restrict the sample to brown and green industries, only. We cluster standard errors on the CLO and year-month levels and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	brown versus all		brown versus green	
$Attention_{t-1} \times Brown_r$	0.045***	0.044**	0.076***	0.072***
	(0.016)	(0.017)	(0.021)	(0.021)
$Attention_{t-1} \times Brown_r \times Affiliated_i$	0.035**	0.040**	0.047**	0.061***
	(0.017)	(0.016)	(0.020)	(0.021)
$Attention_{t-1} \times Affiliated_i$	-0.030**		-0.049***	
	(0.014)		(0.016)	
$Brown_r \times Affiliated_i$	0.013	0.010	-0.012	-0.017
	(0.009)	(0.009)	(0.012)	(0.013)
Number of observations	294,265	293,476	136,451	134,300
Adjusted R2	0.087	0.161	0.096	0.167
Mean dependent	0.212	0.212	0.21	0.209
SD dependent	0.531	0.53	0.531	0.527
Industry FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
CLO-year-month FE	-	Yes	-	Yes
Year-month FE	Yes	-	Yes	-

February 2016. This short sample period has the drawback that we cannot compare the two investor types for our full sample period. However, the advantage of this shorter period is that it centers around the Paris Agreement, which we then use as a natural experiment for a shock to public attention to climate change. This setting allows to mitigate the potential concern that there could be omitted variables (e.g., government policies) that affect both attention to climate change and CLO trading behavior.

Following Degryse et al. (2023); Delis et al. (2024); Ehlers et al. (2021); Müller and Sfrappini (2021); Reghezza et al. (2022), among others, we use the Paris Agreement as an unexpected shock to climate change attention, assessing how trading behavior changes around this event. The Conference of the Parties (COP) 21 occurred from November 30 until December 12, 2015. The resulting Paris Agreement represents a landmark decision in the fight against climate change. We compare CLO transactions of brown loans of both pre- and post-event windows and assess whether net purchase of high emission loans differ in comparison to other loans.

Because we cannot match a large part of the loans held by loan mutual funds to the holdings in CLO-i, Moody's industry classifications are not available for a large part of the loan mutual fund holdings sample. To address this concern, we hand-match all loan issuers held by loan mutual funds to Refinitiv-Eikon and repeat our alternative analysis using emission quintiles. We begin by assessing net purchases of CLOs, and then repeat the analysis with data from Loan mutual funds. We use the specification from Table 4.5, Column (2) and estimate the following difference-in-differences regression around the Paris agreement for CLOs:

$$Net\ Purchase_{i,q,t} = \beta\ Post_t \times High\ Emission_q + \alpha_i \times \alpha_t + \alpha_q + \epsilon_{i,t,q} \tag{4.6}$$

The sample covers monthly net purchases during CLO investment periods over the period October 1, 2015 to February 29, 2016. We include CLO-time fixed effects as well as rank fixed effects, and cluster standard errors on the CLO level. *Post* is an indicator variable that equals one for January 2016, and zero for November 2015. We omit the conference month, December 2015, in our analysis. The coefficient of interest is β, which captures how net purchases of high emission loans have changed after the Paris Agreement compared to net purchases of low-emission loans. As we can see from Column (1) of Table 4.8, CLOs increase net purchases of brown loans one month after the Paris Agreement, compared to the month preceding Paris. In a second step, we repeat the analysis including two months before and after the Paris agreement. However, as shown in Column (2), the coefficient in this specification turns insignificant.

Turning to loan mutual funds, we have a sample of 57 unique US bank loan funds reporting portfolio shares of loans for the respective time period. We exclusively consider loan holdings and therefore discard other asset classes. We sum all position weights (shares) into five emission ranks as described in Section 4.3.2. To approximate *Net purchases*, we calculate changes of the position weights (shares) per fund, rank and month. This leaves us with changes in shares for October 2015 - February 2016. Similarly to the CLO regressions around the Paris Agreement, we estimate Regression Equation (4.6) with Loan mutual fund-year-month fixed effects and cluster standard errors on the Loan mutual fund level.

Columns (3) and (4) in Table 4.8 show the results of our estimation. As expected, Loan mutual funds sell when CLOs are buying. β is negative and statistically significant focusing on one month before and after the Paris agreement. Loan mutual funds show a lower change of loans in their holdings of high emission ranked activities right after the Paris Agreement in comparison to before. Compared to the sample mean of 4.5, the lower change of 1.18 corresponds to a reduction of 40% in comparison to the average change of shares.

Table 4.8: **Paris Agreement with loan mutual funds**. In this table we report results from estimating the following equation (Equation 4.6): $Y_{i,q,t} = \beta\ Post_t \times High\_emissions_q + \alpha_i \times \alpha_t + \alpha_q + \epsilon_{i,t,q}$. Y is $Net\ Purchase_{i,q,t}$ for CLOs and is defined as in Equation (4.1) but, instead of industries, we aggregate on emission quintiles of activities q. Y is defined as the change of the share of the rank in the holdings of the loan mutual fund in columns (3)-(4). $High\_emissions$ are quintiles in terms of emission intensity of the activity of the issuer as defined in Section 4.3.2. $Post_t$ is a binary variable that equals 0 for November 2015 and 1 for January 2016 in column (1) and (3), and extends by +/- one month in columns (2) and (4). We cluster standard errors on the CLO or Loan mutual fund level, respectively, and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	CLOs		Loan Mutual Funds	
	+/- 1 months	+/- 2 months	+/- 1 months	+/- 2 months
$Post_t \times High\ Emissions_q$	0.293***	0.036	-1.182***	0.070
	(0.077)	(0.056)	(0.398)	(0.162)
Number of observations	2,548	5,362	568	1,136
Adj R2	0.410	0.275	0.753	0.754
Mean Dependent	0.270	0.241	4.542	4.516
SD Dependent	0.773	0.667	4.839	4.826
Rank FE	Yes	Yes	Yes	Yes
CLO-Year-Month FE	Yes	Yes	-	-
Fund-Year-Month FE	-	-	Yes	Yes

4.7 Further Results

In this section, we present four additional results that supplement our main analysis. First, we highlight that CLOs do not reverse their increased purchases immediately after heightened media attention to climate change. Second, we highlight that our results are robust to modifying our definition of brown industries. Third, we show that our results are robust to the choice of attention indices. Finally, we show that CLOs with higher leverage do not act differently from CLOs with lower leverage.

4.7.1 CLOs Do Not Reverse Trades

We test whether CLOs reverse the increases of net purchases of brown loans in the periods following peak attention to climate change by selling more brown loans. To that end, we estimate the following regression:

$$Net\ Purchases_{i,r,t} = \sum_{\tau=t-1}^{t-4} \beta_\tau\ Attention_\tau \times \text{Brown}_r + (\alpha_i \times \alpha_t) + \alpha_r + \epsilon_{i,r,t}, \tag{4.7}$$

including up to four lags of our measure of attention to climate change Index. We present results for β_{t-1}, β_{t-2}, β_{t-3} and β_{t-4} in Figure 4.4. If CLOs reversed increases in net purchases of brown loans, we would expect to find negative coefficients for the lags of the attention index. However, we find that CLOs do not reverse their trades within the months following peak media attention to climate change. β_{t-1} is positive and statistically significantly different from zero, but the coefficients for the following lags are zero.

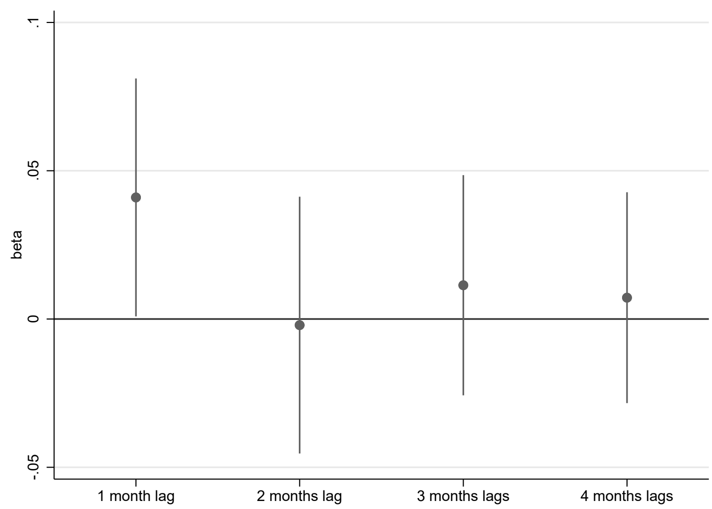

Figure 4.4: **Net purchases and lags of attention to climate change.** In this figure we present results from estimating the following equation (Equation 4.7): $Net\ Purchases_{i,r,t} = \sum_{\tau=t-1}^{t-4} \beta_\tau\ Attention_\tau \times \text{Brown}_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. $Net\ Purchases_{i,r,t}$ is defined as in Equation 4.1. $Brown$ is an indicator that equals 1 for brown loans from six industries, as defined in Section 4.3. $Index_{T-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above the 1st, 2nd ... 5th quintile over the whole sample period. We cluster standard errors on the CLO and year-month levels and report confidence intervals at the 5% level.

4.7.2 Robustness to the Selection of Brown Industries

In this section, we assess how robust our results are toward the selection of brown industries based on the Moody's industry classifier. For this purpose, we re-estimate Equation (4.3) with different definitions of brown industries. In our original definition of brown, as described in Section 4.3, we have eight industries classified as brown and 26 industries classified as non-brown. In the following tests, we define one of the brown industries as non-brown and replace it consecutively with one of the remaining 26 industries. We repeat this exercise 8×26 times. If our classification of brown industries is valuable, we expect that the mass of results using other industries as brown lie to the left of our coefficient. Meanwhile, we expect that the coefficients with varying definitions of brown are not significantly different from our result, indicating that our result does not hinge on one single industry.

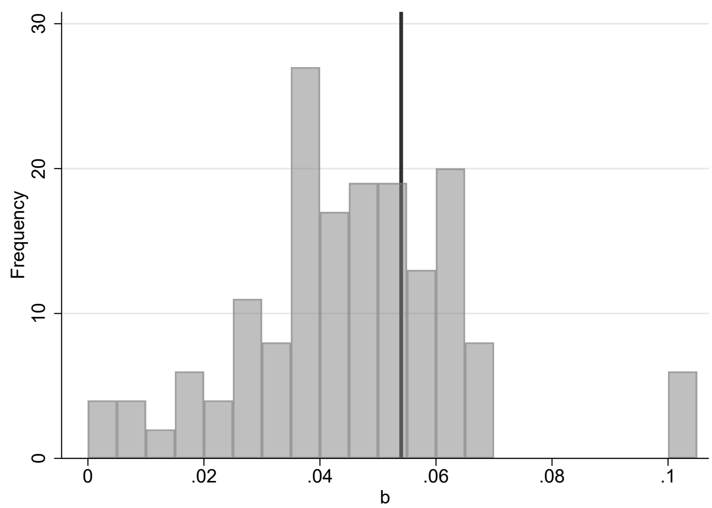

Figure 4.5: **Varying classification of brown industries.** In this histogram we plot the distribution of β by estimating 208 times the following equation with varying definitions for brown industries (Equation 4.3): $Net\,Purchases_{i,r,t} = \beta\,Attention_{t-1} \times Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during reinvestment periods of CLOs over the period 1 January 2010 - 31 May 2018. $Net\,Purchases_{i,r,t}$ is defined as in Equation 4.1. *Brown* is an indicator that equals 1 for brown loans from eight industries, and 0 for the remaining 26 industries. For each regression, we replace one of the brown industries as defined in Section 4.3, with one previously non-brown industry. We repeat the exercise 8×26 times. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above the 5th quintile over the whole sample period. We cluster standard errors on the CLO and year-month level.

We present the distribution of the resulting βs in Figure 4.5. The black line shows the result from our original estimation ($\beta = 0.054$ from Table 4.4, column (2)). We find that, out of 208 regressions, most lie within a similar range of our previous result and the mass lies toward the left of our original result. Out of the 208, the figure shows that four coefficient estimates are close to zero. These outliers are driven by dropping one of the four largest brown industries and substituting it with "Electronics". The electronics industry is the largest in terms of total purchases. It is also arguably a green industry, as electronics (e.g., batteries) are often used to reduce carbon emissions. In fact, Giannetti et al. (2023) define the manufacturing of electronic equipment as a green industry. Hence, we conclude that our definition of brown industries is robust toward alternative classifications.[6]

4.7.3 Robustness to the Selection of Attention Indices

To test the robustness of our results toward the choice of the CHNeg attention index, we consider two more attention indices to identify episodes of elevated attention to climate change. We use the WSJ Index which is based on articles in the WSJ and was also developed by Engle et al. (2020). To construct the index, Engle et al. (2020) use a vocabulary list compiled from more than 70 authoritative texts from, for example, the Intergovernmental Panel on Climate Change (IPCC), as well as other governmental and research organizations and calculate the correlation between the vocabulary list and newspaper articles in the WSJ. The index is available on a monthly basis from January 1984 to June 2017. Moreover, the Media Climate Change Concerns (MCCC) index was developed by Ardia et al. (2023) and is based on ten different newspaper outlets. The authors rely on the publishers' tags of articles linked to climate change and also perform a sentiment analysis of the tagged newspaper articles to obtain negative news on climate change. This index is available on a daily level from January 1, 2003 to August 31, 2022.

To capture information from all attention measures, we aggregate the three indices in two steps. First, we construct monthly "attention spikes" as instances when an index exceeds its 80% quantile, measured during the overlapping time period from January 1, 2010 to June 30, 2017. Second, we sum the quarterly number of spikes across all indices in each quarter. We plot all three attention indices on a monthly level in Figure C.4 in the Appendix. The grey-shaded areas correspond to quarters in which we observe at least three spikes. The red dots in Figure C.4 mark instances in which the respective index exceeds it 80% quantile and the text highlights some of the significant new events in these periods.

[6] 2014-2016 there was a strong decline in oil prices. To rule out that loans from Oil & Gas drive our results, we additionally replicate our main analysis while dropping all loans from Oil & Gas. Our results remain stable and we are confident that the Oil & Gas sector are not the sole drivers of our results.

We define periods of heightened media attention to climate change as periods when we observe at least three or four, respectively, attention spikes within a given quarter. We re-estimate Equation (4.3) on a quarterly level, using our aggregated indices as attention measures and report results in Table 4.9. The results remain largely unchanged. In phases of elevated media attention to climate change, CLOs increase net purchases of brown loans in comparison to non-brown loans. β is positive and significantly different from zero at the 1 % level in all specifications. In terms of economic magnitudes, we observe an increase of brown net purchases compared to non-brown net purchases by 18.7% (0.047/0.251 = 0.187) in Column (2), or 27.5% (0.069/0.251 = 0.275) in Column (4), which corresponds to the magnitudes we found before. We conclude that our results do not hinge on our choice of the CHneg index to measure media attention to climate change.

Table 4.9: **Net purchases with aggregated attention indices.** In this table we report results from estimating the following equation: $NetPurchase_{i,r,t} = \beta Attention\_aggregated_{t-1} \times Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,r,t}$. The sample covers *quarterly* aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-quarter t during reinvestment periods of CLOs over the period 1 January 2010 - 30 September 2017. $Net\,Purchase_{i,r,t}$ is defined as in equation (4.1), where t is year-quarter. *Brown* is an indicator, which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention\_aggregated_{t-1}$ is a binary variable that equals 1 if we observe in total at least three monthly spikes per quarter in the attention indices CHneg and WSJ by Engle et al. (2020) and MCCC by Ardia et al. (2023) in columns (1) and (2), or four spikes in columns (3) and (4) (see Section 4.7.3). We cluster standard errors on the CLO level and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
Attention intensity	Three		Four	
$Attention\_aggregated_{t-1} \times Brown_r$	0.047***	0.047***	0.061***	0.069***
	(0.011)	(0.010)	(0.013)	(0.012)
Number of Observations	145,861	145,861	145,861	145,861
R2	0.212	0.458	0.212	0.458
Mean Dependent	0.251	0.251	0.251	0.251
SD Dependent	1.066	1.066	1.066	1.066
Industry FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
Year-Quarter FE	Yes		Yes	-
CLO-Year-Quarter FE	-	Yes	-	Yes

4.7.4 The Role of CLO Leverage

We test whether CLO leverage can predict CLO investment into high-emission industries when attention to climate change is high. A higher leverage can increase the return of the equity holders and therefore might incentivize for a stronger participation in the increase of net purchases during phases of high attention to climate change. We therefore estimate Equation (4.4) with Leverage of CLO i interacted with $Brown_r \times Attention_{t-1}$. We report results in Table C.5 in the Appendix. We do not find differential effects for highly versus lowly levered CLOs and conclude that CLO leverage is not a driving force in explaining our results.

4.8 Conclusion

We investigate how CLOs adapt their trading behavior regarding loans from firms in high-polluting industries when public attention to climate change is elevated. CLOs are the largest investors in leveraged loan markets and therefore decisive for firms' refinancing opportunities. We show that when climate change gains media attention, CLOs take the opportunity to increase their exposure to brown firms at lower prices. In addition, we provide evidence that CLOs managed by signatories of the PRI do not behave differently, and still tilt their portfolios toward brown loans when prices fall. In contrast, bank-affiliated CLO managers rather increase brown investments more strongly during times of heightened media attention to climate change.

Chapter 5

Summary and Outlook

This dissertation investigates to what degree textual news data can help explain prevailing challenges in asset pricing and increase understanding of markets and investor behavior.

Chapter 2 leverages an extensive dataset of compiled newspaper articles and devises monthly news-derived measures of various economic indicators, many of which have hitherto been accessible solely on an annual or quarterly basis. Our innovative approach elevates annual data to higher frequencies, considering stationary and time-aggregation of data while adjusting for leap years and seasonal fluctuations to enhance economic analysis. We demonstrate the application of our news-based indicators in the area of consumption-based asset pricing. Furthermore, our news-based economic time series pave the way for additional applications, thereby presenting new avenues for economists and researchers in the field. Follow-up research can explore such further applications as well as broaden the usage of our approach by e.g. increasing frequencies to weekly periods.

Chapter 3 explores the economic rational behind the commonly used asset pricing factors HML, SMB, RMW, and CMA. While effective in practice, the fundamental economic origins of these factors remain unclear. With the help of human coders as well as ChatGPT, we create a novel data set which classifies large factor returns into 16 categories using news articles from the following day. We find that news on the macroeconomic outlook, monetary policy, and corporate earnings reports are the most common causes of large absolute returns across all factors. We develop novel topical factors using principal components on news days and discover that HML is associated with macroeconomic updates, whereas CMA is connected to news about commodities. The interpretation of SMB and RMW is more nuanced: SMB is related to exchange rate news and a sentiment-based factor, while RMW is affected by firm-specific information. These findings indicate that both risk-related and behavioral factors are involved, adding to the discussion on whether factor risk premiums arise from rational reasons or market mispricing.

Eventually, Chapter 4 delves into the complex dynamics of CLOs, which are non-bank entities that engage in the securitization of corporate loans. It scrutinizes their trading maneuvers concerning loans from high-emission sectors, particularly amid periods of heightened climate change attention. As dominant actors within the leveraged loan sector CLOs play a crucial role in determining the refinancing prospects for corporations. Our analyses provide evidence that during episodes of intensified media attention to climate change CLOs strategically amplify their engagement with brown firms, thereby capitalizing on the opportunity to acquire such loans at reduced prices. Intriguingly, our findings indicate that even CLOs administered by signatories of the PRI do not deviate from this trend. They as well adjust their portfolios towards brown loans during price declines. Conversely, CLO managers with banking affiliations exhibit an even more pronounced appetite for brown investments. We conclude that CLOs seize brown loans in the wake of divestitures driven by public concern over climate change. We can attribute this behavior to two plausible rationales: CLOs may exhibit an indifference towards the environmental designation of their holdings, whether brown or green, concentrating solely on the evaluative metrics of the firm or loan. This dispassionate stance enables them to absorb brown loans shed by other entities, such as mutual funds, who are motivated by climate-conscious considerations. CLOs may strategically purchase brown loans amid increased climate risk consciousness, predicated on the conviction that such loans are temporarily undervalued, thereby manifesting a contrarian perspective to prevailing public sentiment. For future research, it would be interesting to further investigate the motives behind the behavior of CLO managers. A deeper look at the question of who invests in CLOs and how these investors are constrained could provide a more profound insight into their motives.

Bibliography

Adämmer, P. and R. A. Schüssler (2020): "Forecasting the Equity Premium: Mind the News!" *Review of Finance*, 24, 1313–1355.

Adrian, T., E. Etula, and T. Muir (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *The Journal of Finance*, 69, 2557–2596.

Ai, H. and R. Bansal (2018): "Risk preferences and the macroeconomic announcement premium," *Econometrica*, 86, 1383–1430.

Ait-Sahalia, Y., J. A. Parker, and M. Yogo (2004): "Luxury goods and the equity premium," *The Journal of Finance*, 59, 2959–3004.

Albuquerque, P. H. M., Y. Peng, and J. P. F. da Silva (2022): "Making the whole greater than the sum of its parts: A literature review of ensemble methods for financial time series forecasting," *Journal of Forecasting*, 41, 1701–1724.

Aleti, S. and T. Bollerslev (2024): "News and Asset Pricing: A High-Frequency Anatomy of the SDF," *Review of Financial Studies*, forthcoming.

Angrist, J. D. and J.-S. Pischke (2009): *Mostly harmless econometrics: An empiricist's companion*, Princeton University Press.

Apergis, N., P. Artikis, and J. Sorros (2011): "Asset pricing and foreign exchange risk," *Research in International Business and Finance*, 25, 308–328.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2023): "Climate Change Concerns and the Performance of Green vs. Brown Stocks," *Management Science*, 69, 7607–7632.

Aretz, K., S. M. Bartram, and P. F. Pope (2010): "Macroeconomic risks and characteristic-based factor models," *Journal of Banking and Finance*, 34, 1383–1399.

Baker, S., N. Bloom, S. Davis, and M. Sammon (2021): "What Triggers Stock Market Jumps?" Working Paper, Kellogg School of Management.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, 131, 1593–1636.

Bali, T. G., H. Beckmeyer, and T. Wiedemann (2023): "Expected Mispricing," Working Paper, Georgetown University.

Bansal, R. and A. Yaron (2004): "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 59, 1481–1509.

Belo, F., V. D. Gala, and J. Li (2013): "Government spending, political cycles, and the cross section of stock returns," *Journal of Financial Economics*, 107, 305–324.

Ben-Rephael, A., B. I. Carlin, Z. Da, and R. D. Israelsen (2021): "Information Consumption and Asset Pricing," *The Journal of Finance*, 76, 357–394.

Benmelech, E., J. Dlugosz, and V. Ivashina (2012): "Securitization without adverse selection: The case of CLOs," *Journal of Financial Economics*, 106, 91–113.

Bernanke, B. S. and K. N. Kuttner (2005): "What explains the stock market's reaction to federal reserve policy?" *The Journal of Finance*, 60, 1221–1257.

Bhardwaj, A., S. Ge, and S. Mukherjee (2022): "The Effect of Investor Demand on Loan Securitization and Firm Financing," Working paper, New York University.

Bok, B., D. Caratelli, D. Giannone, A. M. Sbordone, and A. Tambalotti (2018): "Macroeconomic Nowcasting and Forecasting with Big Data," *Annual Review of Economics*, 10, 615–643.

Böll, J., F. Meng, J. Thimme, and M. Uhrig-Homburg (2024): "Following the Footprints: Towards a Taxonomy of the Factor Zoo," Working Paper, Karlsruhe Institute of Technology.

Bolton, P. and M. Kacperczyk (2021): "Do investors care about carbon risk?" *Journal of Financial Economics*, 142, 517–549.

——— (2023): "Global pricing of carbon-transition risk," *The Journal of Finance*, 78, 3677–3754.

Bord, V. M. and J. A. Santos (2015): "Does securitization of corporate loans lead to riskier lending?" *Journal of Money, Credit and Banking*, 47, 415–444.

Bybee, L., B. Kelly, A. Manela, and D. Xiu (2024): "Business News and Business Cycles," *The Journal of Finance*, 79, 3105–3147.

Bybee, L., B. Kelly, and Y. Su (2023): "Narrative Asset Pricing: Interpretable Systematic Risk Factors from News Text," *Review of Financial Studies*, 36, 4759–4787.

Carrieri, F., V. Errunza, and B. Majerbi (2006): "Does emerging market exchange risk affect global equity prices?" *Journal of Financial and Quantitative Analysis*, 41, 511–540.

Ceccarelli, M., S. Ramelli, and A. F. Wagner (2024): "Low carbon mutual funds," *Review of Finance*, 28, 45–74.

Chen, A. Y. and C. Dim (2023): "High-Throughput Asset Pricing," *Working Paper, Federal Reserve Board.*

Chen, N.-F. (1991): "Financial investment opportunities and the macroeconomy," *The Journal of Finance*, 46, 529–554.

Chen, N.-f., R. Roll, and S. A. Ross (1986): "Economic Forces and the Stock Market," *Journal of Business*, 59, 383–403.

Choi, D., Z. Gao, and W. Jiang (2020): "Attention to global warming," *The Review of Financial Studies*, 33, 1112–1145.

Cooper, I., A. Mitrache, and R. Priestley (2022): "A global macroeconomic risk model for value, momentum, and other asset classes," *Journal of Financial and Quantitative Analysis*, 57, 1–30.

Cordell, L., M. R. Roberts, and M. Schwert (2023): "CLO performance," *The Journal of Finance*, 78, 1235–1278.

Daniel, K., D. Hirshleifer, and L. Sun (2020): "Short-and long-horizon behavioral factors," *Review of Financial Studies*, 33, 1673–1736.

De Santis, G. and B. Gerard (1998): "How big is the premium for currency risk?" *Journal of Financial Economics*, 49, 375–412.

Degryse, H., R. Goncharenko, C. Theunisz, and T. Vadasz (2023): "When green meets green," *Journal of Corporate Finance*, 78, 102355.

Delis, M., K. Greiff, K. De, and S. Ongena (2024): "Being stranded with fossil fuel reserves? Climate policy risk and the pricing of bank loans," *Financial Markets, Institutions & Instruments*, 33, 239–265.

Didisheim, A., S. Ke, B. Kelly, and S. Malamud (2024): "APT or "AIPT"? The Surprising Dominance of Large Factor Models," Working Paper, University of Melbourne.

Dittmar, R. F. and C. T. Lundblad (2017): "Firm characteristics, consumption risk, and firm-level risk exposures," *Journal of Financial Economics*, 125, 326–343.

Dittmar, R. F., C. Schlag, and J. Thimme (2024): "Non-substitutable consumption growth risk," *Management Science*, forthcoming.

Dumas, B. and B. Solnik (1995): "The world price of foreign exchange risk," *The Journal of Finance*, 50, 445–479.

Ehlers, T., F. Packer, and K. Greiff (2021): "The pricing of carbon risk in syndicated loans: Which risks are priced and why?" *Journal of Banking and Finance*, 136, 106180.

Elkamhi, R., R. Li, and Y. Nozawa (2023): "A benchmark for collateralized loan obligations," *HKUST Business School Research Paper*.

Elkamhi, R. and Y. Nozawa (2022): "Fire-sale risk in the leveraged loan market," *Journal of Financial Economics*, 146, 1120–1147.

Emin, M., C. M. James, T. Li, and J. Lu (2024): "Institutional Synergies and the Fragility of Loan Funds," Working Paper, University of Alabama.

Engle, R. F., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020): "Hedging Climate Change News," *The Review of Financial Studies*, 33, 1184–1216.

Epstein, L. G. and S. E. Zin (1989): "Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, 57, 937––969.

Estrella, A. and G. A. Hardouvelis (1991): "The term structure as a predictor of real economic activity," *The Journal of Finance*, 46, 555–576.

Fabozzi, F. J., S. Klingler, P. Mølgaard, and M. S. Nielsen (2021): "Active loan trading," *Journal of Financial Intermediation*, 46, 100868.

Fama, E. F. and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33, 3–56.

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116, 1–22.

Fedyk, A. (2024): "Front-Page News: The Effect of News Positioning on Financial Markets," *The Journal of Finance*, 79, 5–33.

Fisher, A., C. Martineau, and J. Sheng (2022): "Macroeconomic Attention and Announcement Risk Premia," *Review of Financial Studies*, 35, 5057–5093.

Fleckenstein, Q. (2022): "Intermediary frictions and the corporate credit cycle: Evidence from CLOs," Working Paper, HEC Paris.

Francis, B. B., I. Hasan, and D. M. Hunter (2008): "Can hedging tell the full story? Reconciling differences in United States aggregate-and industry-level exchange rate risk premium," *Journal of Financial Economics*, 90, 169–196.

Frenkel, J. A. (1981): "Flexible exchange rates, prices, and the role of" news": lessons from the 1970s," *Journal of Political Economy*, 89, 665–705.

Frey, J. (2023): "Which stock return predictors reflect mispricing?" Working Paper, University of Oxford.

Friedman, B. M. and K. N. Kuttner (1998): "Indicator properties of the paper—bill spread: Lessons from recent experience," *Review of Economics and Statistics*, 80, 34–44.

Froot, K. and P. Klemperer (1989): "Exchange rate pass-through when market share matters," *American Economic Review*, 79, 637–654.

Gao, L., S. Hitzemann, I. Shaliastovich, and L. Xu (2022): "Oil volatility risk," *Journal of Financial Economics*, 144, 456–491.

García, D. (2013): "Sentiment during Recessions," *The Journal of Finance*, 68, 1267–1300.

Gentzkow, M., B. Kelly, and M. Taddy (2019): "Text as Data," *Journal of Economic Literature*, 57, 535–574.

Giannetti, M., M. Jasova, M. Loumioti, and C. Mendicino (2023): ""Glossy Green" Banks: The Disconnect Between Environmental Disclosures and Lending Activities," Working Paper, Swedish House of Finance.

Gibson Brandon, R., S. Glossner, P. Krueger, P. Matos, and T. Steffen (2022): "Do responsible investors invest responsibly?" *Review of Finance*, 26, 1389–1432.

Gilchrist, S., V. Yankov, and E. Zakrajšek (2009): "Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets," *Journal of Monetary Economics*, 56, 471–493.

Ginglinger, E. and Q. Moreau (2023): "Climate Risk and Capital Structure," *Management Science*, 69, 7492–7516.

Griffin, J. M. and J. Nickerson (2023): "Are CLO collateral and tranche ratings disconnected?" *The Review of Financial Studies*, 36, 2319–2360.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1, 55–92.

Görgen, M., A. Jacob, M. Nerlinger, R. Riordan, M. Rohleder, and M. Wilkens (2020): "Carbon risk," Working Paper, University of Augsburg.

Hackenberg, K. (2022): "Trading in the shadows: three essays on the secondary market for leveraged loans," Ph.D. thesis, Catholic University Eichstätt-Ingolstadt.

Hackenberg, K., V. Klaus, S. Klingler, and T. Sondershaus (2024): "Exploiting Media Attention to Climate Change: CLO Brown Loan Trading," Working Paper, Catholic University Eichstätt-Ingolstadt.

Hahn, J. and H. Lee (2006): "Yield Spreads as Alternative Risk Factors for Size and Book-to-Market," *Journal of Financial and Quantitative Analysis*, 41, 245–269.

Hale, G., B. Meisenbacher, and F. Nechio (2024): "Industrial Composition of Syndicated Loans and Banks' Climate Commitments," *Federal Reserve Bank of San Francisco Working Paper Series*, 2024, 01–37.

Hall, R. E. (1988): "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 96, 339–57.

Hansen, L. P. (1982): "Large sample properties of generalized method of moments estimators," *Econometrica: Journal of the Econometric Society*, 1029–1054.

Harvey, C. R. (1988): "The real term structure and consumption growth," *Journal of Financial Economics*, 22, 305–333.

He, Z. and A. Krishnamurthy (2013): "Intermediary asset pricing," *American Economic Review*, 103, 732–770.

Hsu, P.-h., K. Li, and C.-y. Tsou (2023): "The pollution premium," *The Journal of Finance*, 78, 1343–1392.

Huynh, T. D. and Y. Xia (2021): "Climate change news risk and corporate bond returns," *Journal of Financial and Quantitative Analysis*, 59, 1985–2009.

——— (2023): "Panic Selling When Disaster Strikes: Evidence in the Bond and Stock Markets," *Management Science*, 69, 7448–7467.

Jeon, Y., T. H. McCurdy, and X. Zhao (2022): "News as sources of jumps in stock returns: Evidence from 21 million news articles for 9000 companies," *Journal of Financial Economics*, 145, 1–17.

Jordà, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019): "The Rate of Return on Everything, 1870–2015," *Quarterly Journal of Economics*, 134, 1225–1298.

Jordà, O., B. Richter, M. Schularick, and A. M. Taylor (2021): "Bank capital redux: solvency, liquidity, and crisis," *Review of Economic Studies*, 88, 260–286.

Jordà, O., M. Schularick, and A. M. Taylor (2017): "Macrofinancial History and the New Business Cycle Facts," *NBER Macroeconomics Annual 2016*, 31.

Jorion, P. (1990): "The exchange-rate exposure of US multinationals," *Journal of Business*, 331–345.

Kacperczyk, M. and J. Peydro (2021): "Carbon emissions and the bank-lending channel," Working Paper, Imperial College London.

Kelly, B., S. Malamud, and K. Zhou (2024): "The Virtue of Complexity in Return Prediction," *The Journal of Finance*, 79, 459–503.

Kelly, B. T., S. Malamud, and K. Zhou (2022): "The virtue of complexity everywhere," *Working Paper, Yale School of Management.*

Kerssenfischer, M. and M. Schmeling (2024): "What moves markets?" *Journal of Monetary Economics*, 145, 103560.

Kim, S. and A. Yoon (2023): "Analyzing Active Fund Managers' Commitment to ESG: Evidence from the United Nations Principles for Responsible Investment," *Management Science*, 69, 741–758.

Klaus, V. and J. Thimme (2024a): "Text-based Macro Data and Asset Prices," Working Paper, Karlsruhe Institute of Technology.

——— (2024b): "Understanding Asset Pricing Factors," Working Paper, Karlsruhe Institute of Technology.

Kogan, L. and D. Papanikolaou (2014): "Growth opportunities, technology shocks, and asset prices," *The Journal of Finance*, 69, 675–718.

Kollmorgen, L. and S. Oh (2022): "Seeing Beyond the Complexity: An Introduction to Collateralized Loan Obligations," *www.pinebridge.com/en/insights/clo-beyond-the-complexity (accessed: 01/23/2023).*

Kozak, S., S. Nagel, and S. Santosh (2018): "Interpreting factor models," *The Journal of Finance*, 73, 1183–1223.

Kroencke, T. A. (2017): "Asset Pricing without Garbage," *The Journal of Finance*, 72, 47–98.

Kroencke, T. A., M. Schmeling, and A. Schrimpf (2021): "The FOMC risk shift," *Journal of Monetary Economics*, 120, 21–39.

Krueger, P., Z. Sautner, and L. T. Starks (2020): "The importance of climate risks for institutional investors," *The Review of Financial Studies*, 33, 1067–1111.

Kundu, S. (2022): "The Anatomy of Corporate Securitizations and Contract Design," *Journal of Corporate Finance*, 81, 102195.

——— (2024): "Financial covenants and fire sales in closed-end funds," *Management Science*, 70, 860–884.

Liebscher, R. and T. Mählmann (2017): "Are professional investment managers skilled? Evidence from syndicated loan portfolios," *Management Science*, 63, 1892–1918.

Liew, J. and M. Vassalou (2000): "Can book-to-market, size and momentum be risk factors that predict economic growth?" *Journal of Financial Economics*, 57, 221–245.

Liu, Y. and B. Matthies (2022): "Long-Run Risk: Is It There?" *The Journal of Finance*, 77, 1587–1633.

Lopez-Lira, A. (2019): "Risk Factors That Matter: Textual Analysis of Risk Disclosures for the Cross-Section of Returns," Working Paper, University of Florida.

Loughran, T. and B. Mcdonald (2016): "Textual Analysis in Accounting and Finance: A Survey," *Journal of Accounting Research*, 54, 1187–1230.

Loumioti, M. and F. P. Vasvari (2019): "Portfolio performance manipulation in collateralized loan obligations," *Journal of Accounting and Economics*, 67, 438–462.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jørgensen (2009): "Long-run stockholder consumption risk and asset returns," *The Journal of Finance*, 64, 2427–2479.

Manela, A. and A. Moreira (2017): "News implied volatility and disaster concerns," *Journal of Financial Economics*, 123, 137–162.

Mehra, R. and E. C. Prescott (1985): "The equity premium: A puzzle," *Journal of Monetary Economics*, 15, 145–161.

Mullainathan, S. and A. Shleifer (2005): "The Market for News," *The American Economic Review*, 95, 1031–1053.

Müller, I., H. Nguyen, and T. Nguyen (2024): "Carbon transition risk and corporate loan securitization," Working Paper, IWH.

Müller, I. and E. Sfrappini (2021): "Climate change-related regulatory risks and bank lending," Working Paper, ECB.

Nguyen, J. and H. Phan (2020): "Carbon risk and corporate capital structure," *Journal of Corporate Finance*, 64, 101713.

Nicolai, F. (2020): "Contagion in the Market for Leveraged Loans," Working Paper, BI Norwegian Business School.

Nimark, K. P. and S. Pitschner (2019): "News media and delegated information choice," *Journal of Economic Theory*, 181, 160–196.

Patton, A. J. and A. Timmermann (2011): "Predictability of output growth and inflation: A multi-horizon survey approach," *Journal of Business & Economic Statistics*, 29, 397–410.

Patton, A. J. and M. Verardo (2012): "Does beta move with news? Firm-specific information flows and learning about profitability," *Review of Financial Studies*, 25, 2789–2839.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021): "Responsible investing: The ESG-efficient frontier," *Journal of Financial Economics*, 142, 572–597.

Peristiani, S. and J. A. Santos (2019): "CLO trading and collateral manager bank affiliation," *Journal of Financial Intermediation*, 39, 47–58.

PETKOVA, R. (2006): "Do the Fama-French Factors Proxy for Innovations in Predictive Variables?" *The Journal of Finance*, 61, 581–612.

PETTENUZZO, D., R. SABBATUCCI, AND A. TIMMERMANN (2020): "Cash Flow News and Stock Price Dynamics," *The Journal of Finance*, 75, 2221–2270.

POPPER, K. R. (1968): *The Logic of Scientific Discovery*, London: Hutchinson & Co.

PUKTHUANTHONG, K., J. SHEN, AND R. WANG (2023): "Asset Pricing with Daily Shopper Spending," Working Paper, University of Missouri.

REGHEZZA, A., Y. ALTUNBAS, D. MARQUES-IBANEZ, C. RODRIGUEZ D'ACRI, AND M. SPAGGIARI (2022): "Do banks fuel climate change?" *Journal of Financial Stability*, 62, 101049.

RIGOBON, R. AND B. SACK (2004): "The impact of monetary policy on asset prices," *Journal of Monetary Economics*, 51, 1553–1575.

ROGOFF, K. (1996): "The purchasing power parity puzzle," *Journal of Economic Literature*, 34, 647–668.

ROSS, S. A. (1976): "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, 13, 341–360.

SAUTNER, Z., L. VAN LENT, G. VILKOV, AND R. ZHANG (2023): "Firm-level climate change exposure," *The Journal of Finance*, 78, 1449–1498.

SAVOR, P. AND M. WILSON (2013): "How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements," *Journal of Financial and Quantitative Analysis*, 48, 343–375.

——— (2014): "Asset pricing: A tale of two days," *Journal of Financial Economics*, 113, 171–201.

SAVOV, A. (2011): "Asset Pricing with Garbage," *The Journal of Finance*, 66, 177–201.

SCHMELING, M. AND C. WAGNER (2016): "Does central bank tone move asset prices?" *Journal of Financial and Quantitative Analysis*, 1–48.

SMOLA, A. J. AND B. SCHÖLKOPF (2004): "A tutorial on support vector regression," *Statistics and Computing*, 14, 199–222.

STARKS, L. T. AND K. D. WEI (2003): "Foreign exchange exposure and short-term cash flow sensitivity," Working Paper, University of Texas at Austin.

STUNDZIENE, A., V. PILINKIENE, J. BRUNECKIENE, A. GRYBAUSKAS, M. LUKAUSKAS, AND I. PEKARSKIENE (2024): "Future directions in nowcasting economic activity: A systematic literature review," *Journal of Economic Surveys*, 38, 1199–1233.

TETLOCK, P. C. (2007): "Giving content to investor sentiment: The role of media in the stock market," *The Journal of Finance*, 62, 1139–1168.

——— (2011): "All the news that's fit to reprint: Do investors react to stale information?" *Review of Financial Studies*, 24, 1481–1512.

TETLOCK, P. C., M. SAAR-TSECHANSKY, AND S. MACSKASSY (2008): "More Than Words: Quantifying Language to Measure Firms' Fundamentals," *The Journal of Finance*, 63, 1437–1467.

THIMME, J. (2017): "Intertemporal substitution in consumption: A literature review," *Journal of Economic Surveys*, 31, 226–257.

VASSALOU, M. (2003): "News related to future GDP growth as a risk factor in equity returns," *Journal of Financial Economics*, 68, 47–73.

VISSING-JØRGENSEN, A. (2002): "Limited asset market participation and the elasticity of intertemporal substitution," *Journal of Political Economy*, 110, 825–853.

WACHTER, J. A. AND Y. ZHU (2022): "A model of two days: Discrete news and asset prices," *Review of Financial Studies*, 35, 2246–2307.

YOGO, M. (2006): "A consumption-based explanation of expected stock returns," *The Journal of Finance*, 61, 539–580.

Appendix A

Additional Information on *Text-based Macro Data and Asset Prices*

A.1 Time-Aggregation Problem

To overcome the time aggregation problem, note that we can use for an exemplary economic time series Consumption $C_{l,h} = C_{l-1,12}\prod_{\tau=1}^{h} G_{l,\tau}$, with $G_l = exp(g_l)$, to reformulate time aggregated annual growth as

$$
\begin{aligned}
g_l &= \log\left(\frac{C_l}{C_{l-1}}\right) = \log\left(\frac{\sum_{h=1}^{12} C_{l,h}}{\sum_{h=1}^{12} C_{l-1,h}}\right) = \log\left(\frac{\sum_{h=1}^{12} C_{l-2,12}\prod_{\tau=1}^{12} G_{l-1,\tau}\prod_{\tau=1}^{h} G_{l,\tau}}{\sum_{h=1}^{12} C_{l-2,12}\prod_{\tau=1}^{h} G_{l-1,\tau}}\right) \\
&= \log\left(\prod_{\tau=1}^{12} G_{l-1,\tau}\frac{\sum_{h=1}^{12}\prod_{\tau=1}^{h} G_{l,\tau}}{\sum_{h=1}^{12}\prod_{\tau=1}^{h} G_{l-1,\tau}}\right) = \sum_{h=1}^{12} g_{l-1,h} + \log(\sum_{h=1}^{12}\prod_{\tau=1}^{h} G_{l,\tau}) - \log(\sum_{h=1}^{12}\prod_{\tau=1}^{h} G_{l-1,\tau})
\end{aligned}
$$

The key idea now is to use a first order Taylor series expansion to approximate the logs of the two sums. It turns out that this approximation is very close as long as the growth rates G are close to one (i.e. somewhere between 0.9 and 1.1). This is typically the case for consumption growth rates. Note that $\log(\sum_{h=1}^{12}\prod_{\tau=1}^{h} G_{l-1,\tau})$ is a function with 12 variables, so we have to do a Taylor expansion with 12 dimensions, which is tedious but straight-forward. When the high frequency is daily, we have 365 dimensions. The general principle should be clear from the following example:

Image a year had only two months (or, alternatively, our high frequency is semi-annual). Then, according to the formula above,

$$
\begin{aligned}
g_l &= g_{l-1,1} + g_{l-1,2} + \log(G_{l,1} + G_{l,1}G_{l,2}) - \log(G_{l-1,1} + G_{l-1,1}G_{l-1,2}) \\
&= g_{l-1,1} + g_{l-1,2} + g_{l,1} + \log(1 + G_{l,2}) - g_{l-1,1} - \log(1 + G_{l-1,2}) \\
&= g_{l-1,2} + g_{l,1} + \log(1 + \exp(g_{l,2})) - \log(1 + \exp(g_{l-1,2}))
\end{aligned}
$$

Now we do a Taylor series expansion of $F(x) = \log(1 + \exp(x))$ around $x_0 = 0$. The first derivatives are given by $F^{(0)}(0) = \log(2)$, $F^{(1)}(0) = \frac{1}{2}$, $F^{(2)}(0) = \frac{1}{4}$, $F^{(3)}(0) = 0$, $F^{(4)}(0) = -\frac{1}{8}$, $F^{(5)}(0) = 0$, $F^{(6)}(0) = \frac{1}{4}$, $F^{(7)}(0) = 0$, ... Everything above second order is absolutely negligible. We performed a simulation study and found that even the second order term is not important in applications. Thus, we can approximate

$$\log(1 + \exp(g)) \approx \log(2) + \frac{1}{2}g,$$

and so

$$
\begin{aligned}
g_l &= g_{l-1,2} + g_{l,1} + \log(2) + \frac{1}{2}g_{l,2} - \log(2) - \frac{1}{2}g_{l-1,2} \\
&= \frac{1}{2}g_{l-1,2} + g_{l,1} + \frac{1}{2}g_{l,2}. \qquad \text{(A.1.1)}
\end{aligned}
$$

When assuming three months per year, we can do the same type of Taylor approximation, which leads to

$$g_l = \frac{1}{3}g_{l-1,2} + \frac{2}{3}g_{l-1,3} + g_{l,1} + \frac{2}{3}g_{l,2} + \frac{1}{3}g_{l,3} \tag{A.1.2}$$

and this pattern generalizes to more months/days. When partitioning the year into J subperiods, the general formula is

$$g_l = \sum_{h=2}^{H} \frac{h-1}{H} g_{l-1,h} + \sum_{h=1}^{H} \frac{H-h+1}{H} g_{l,h}.$$

A.2 Additional Figures and Tables

Table A.1: **Context examples from NYT.** This table shows examples of bigrams which got allocated a large coefficient by our models as presented in Table 2.4. *PI* is Production Index, *CPI* is Consumer Price Index, and *UE* is Unemployment. *Date* is the print date of the NYT. The exemplary context passage is captured after initial data cleaning and therefore missing punctuation.

Bigram	**Time Series**	**Date**	**Text Passage**
advertis,caution	PI	06.08.1989	[. . .] loan environment banks no longer urge mortgagors to pay off their debts in fact they warn against it in their advertisements they caution that a paid off home is an unused asset it is the means they suggest for acquiring a second home [. . .]
		23.02.1939	[. . .] those answering advertisements are cautioned not to enclose original references or other valuable material copies serve the purpose and avoid possible lose of originals [. . .]
avail,everi	PI	09.01.1983	[. . .] designer of the proposal is apple computer corporation whose chairman steve jobs has repeated a vow to make a computer available to every school in america but the company is competitors have not endorsed it saying it would accentuate problems caused by providing [...]
		28.03.2004	[. . .] fact that will be the case with the bloomfield apartments the idea is to make sure there are rental units available at every level of income all across the state said annmarie webbing the preservation corporations director in new jersey [. . .]
call,product	CPI	23.01.1954	[. . .] wiley who is chairman of the senate foreign relations committee hit at the brick amendment in fall its phases he called it a product of hysteria that might undo the delicate balance of power that our founding fathers created in this constitution this nation [...]
		25.08.1970	[. . .] who is fast accurate knowledgeable in all recons we would like to speak to you full my benefits office atmosphere call production manager typist will train good typist interested in learning statistical typing for a from must have good figure [. . .]
end,end	PI	10.01.1938	[. . .] antisemitism movement was spreading to rumanian that unless the forces of liberalism and democracy can rescue the country it will end as it is ending in germany other officers of the federation elected were henry lineman of detroit ira a younker and william rose wald [...]
		02.02.1977	[. . .] since the freeze really hardened dec of say the ice will not break before the season ends at the end march short term unemployment benefits of a to of a a week began [. . .]
girl,go	CPI	22.01.1926	[. . .] belief that women would always wear them for the sake of health and modesty and while the custom of modern girls to go corset less grew and seriously affected his business the aged manufacturer still continued it the latest financial statement of the [...]
		28.03.1953	[. . .] degree when a poor village woman needs an operation and can not pay for it the jefe pays when a girl can not go to her first communion because she occasions including of of he was reduced to captains that he could win his [. . .]
glass,menageri	UE	19.03.2010	[. . .] new drama playwrights horizons peter jay sharp theatre of a west and street clinton of a playwright horizons org the glass menagerie in previews opens on wednesday gordon ede stein directs a revival of this tennessee william drama with judith ideas the [...]
		28.09.1950	[. . .] which gable himself stars will be made for metro release at the radio city music hall the attraction

			is the glass menagerie the adaptation by tennessee williams and peter be is of or williams play featuring gertrude lawrence jane woman kirk douglas [...]
includ,power	CPI	19.02.1930	[. . .] these appointments could be marshalled together they would indeed make an impressive trust service hall of fame the roster would include powerful financiers whose wealth was measured in millions big bankers and directors of giant corporations famous artists and sculptors inventors of [...]
		01.09.1950	[. . .] said under any theory of world government the united states would be required to sacrifice in large part its sovereignty including the power to control and direct its national armed forces and defense to relinquish the integrity of its form of government to [...]
market,organ	CPI	09.02.1928	[. . .] forthright cutting down in the urges farm relief by market board small loans for organization expenses only to these commodity marketing organizations to be repaid with per cent interest bankruptcy proceedings banker predicts uncontrolled wastes in distribution [...]
		06.11.1962	[...] a bright view the head of the national security traders association sees a rosy future for the overt counter securities market the organization includes most traders in overt counter activity charles a bodie s or president told the annual convention of then [...]
never,worn	PI	27.02.1983	[. . .] the world ready for a new go round of structured fashions and fairly elaborate accessories are young women who have never worn these kinds of clothes eager to try the glamour circuit the new designers are certainly adept at reviving the old [. . .]
		25.11.1930	[. . .] coats reasonable genuine fox scarfs of open till up a manner of west with lady sell handsome new fur coat never worn of dealer riverside of of and price was times oil paintings art works cash [. . .]
organ,individu	CPI	01.08.2004	[. . .] federal oversight unions have never gained any footing despite numerous efforts by boxers to organize themselves and by advocates to organize for them and individual stars whose big paydays distort boxing overall economics drive purses and television contracts in short boxing [...]
		26.12.1988	[. . .] support of people of good will may the world see an increase of the generous reaction which has mobilized governments organizations and individuals in a wonderful chain of solidarity do not lose hope moving to another dominant global issue aids john paul said [...]
pool,gym	UE	07.04.2021	[. . .] after vice president dan quayle began soliciting donations for various home improvements including a pool a gym and putting green he is my favourite vice president president biden a former resident and fan of the pool [. . .]
		11.03.1937	[. . .] noise on your doorstep the low rate for you comfortable room include social activities free use of natural salt water pool gym rooms with bath business opportunities weekdays of of a line sundays has three business references required closing time for sunday [...]
present,busi	PI	01.01.1931	[. . .] extreme depression as a familiar recurring phenomenon which must within reasonable time be followed by recovery the idea that the present business crisis is exceptional that it foreshadow permanent decline is rejected everywhere opinion differs chiefly on the question whether recovery will [...]
		28.10.1927	[. . .] in the wage scale of workers to did not believe he intimated that wages should be reduced in view of present business methods mentioning particularly the deferred credit system which has been employed quite largely in purchase by workers or grace said [...]
seventh,fleet	UE	01.01.1972	[. . .] thus with communist china such was the situation or alsop wrote when president nixon ordered elements of the united states seventh fleet to team toward the indian ocean in his column or alsop attributed to or kissinger a comparison between hitler preoccupation [...]
		06.03.1966	[. . .] public record the alternative chosen by or johnson was to utilize the provocation of the tonkin gulf attack on the

			seventh fleet by north vietnamese gunboats to get a generalized expression of support from congress this worked well enough until it was [...]
state,ticket	PI	21.06.1993	[. . .] companies under most such programs it requires of of a in card charges to qualify for a free domestic united states ticket and of of a for an international flight chase manhattan exudes optimism about its new affinity partner we are looking [. . .]
		06.10.2019	[. . .] a a million have been set aside for overseas visitors sold by authorized ticket resellers for those in the united states ticket sales will be handled by sport and went on sale in july as of this writing all available tickets have [. . .]
without,addit	CPI	17.02.1962	[. . .] trial offer send is for the next weekly issues of united reports ind receive without additional charge this new research report on of a long time dividend payers special over open to new readers only [. . .]

Table A.2: **Chosen bigrams per macroeconomic time series.**

Time Series		**Model**	**Top 5 Bigrams**
Banks, capital ratio	bigram (+)	lasso	[time,club, credit,cent, sale,director, share,averag, smith,corona]
		svr	[tax,cost, even,entertain, time,club, short,sell, mrs,collin]
	bigram (-)	lasso	[big,price, oil,gasolin, van,loon, refriger,free, knee,length]
		svr	[offer,hous, eve,ford, servic,mani, big,price, present,music]
Banks, loans-to-deposits ratio	bigram (+)	lasso	[open,white, church,england, offic,handl, plan,stock, offer,unit]
		svr	[tax,cost, open,white, news,flash, unit,larg, even,entertain]
	bigram (-)	lasso	[passion,play, effect,economi, oil,gasolin, condit,real, bank,holiday]
		svr	[perfect,locat, bank,holiday, present,music, leav,everi, del,mont]
Banks, noncore funding ratio	bigram (+)	lasso	[news,flash, open,six, abl,follow, trade,zone, spa,sky]
		svr	[tax,cost, ocean,hill, news,flash, govern,board, figur,skate]
	bigram (-)	lasso	[bank,receiv, axe,rod, passion,play, harrison,said, even,summer]
		svr	[lam,duck, sale,novemb, even,summer, maintain,close, arthur,kill]
Broad Money	bigram (+)	lasso	[american,land, john,chamberlain, servic,local, car,purchas, minist,inform]
		svr	[get,park, car,purchas, enjoy,good, string,orchestra, warn,issu]
	bigram (-)	lasso	[governor,roosevelt, manag,sinc, share,expens, port,princ, hamburg,germani]
		svr	[plus,minus, plus,us, plus,plus, minus,plus, us,plus]
Confidence	bigram (+)	lasso	[work,job, class,graduat, page,eight, record,week, center,unit]
		svr	[page,eight, senat,lafollett, probabl,lineup, work,job, set,march]
	bigram (-)	lasso	[sixth,day, best,equip, open,addit, complet,independ, receipt,shipment]
		svr	[open,addit, west,bryant, fla,novemb, american,light, complet,independ]
Consumer Price Index	bigram (+)	lasso	[price,increas, organ,individu, call,product, treasur,island, market,organ]
		svr	[treasur,island, mover,system, monday,monday, queen,victoria, countrywid,weather]
	bigram (-)	lasso	[girl,go, end,tuesday, team,match, without,addit, includ,power]
		svr	[citi,better, convent,visitor, girl,go, visitor,bureau, elev,frigidair]
Consumption (Barro)	bigram (+)	lasso	[receiv,payment, advisori,commiss, bond,without, cathol,protest, see,job]
		svr	[trade,wind, project,approv, room,music, advisori,commiss, decemb,also]
	bigram (-)	lasso	[secretari,stimson, rise,cost, barrel,crude, ziegfeld,theatr, intern,settlement]
		svr	[even,entertain, chines,troop, store,grand, japanes,chines, minus,minus]
Corporate Debt	bigram (+)	lasso	[home,industri, albani,counti, per,florin, ski,lift, water,tank]
		svr	[anoth,fine, arrow,sale, home,industri, seek,work, made,see]
	bigram (-)	lasso	[start,train, govern,citi, offic,month, alon,new, compani,contract]
		svr	[residu,chief, fuel,shortag, compani,contract, energi,crisi, mile,st]
Dividends	bigram (+)	lasso	[itali,would, drug,depart, entir,home, transfer,power, list,time]
		svr	[general,clark, move,st, howard,dean, govern,council, jan,low]
	bigram (-)	lasso	[short,sell, chines,armi, california,coast, island,main, well,open]
		svr	[chines,armi, japanes,chines, chines,japanes, young,young, justic,black]
Employment	bigram (+)	lasso	[charli,chaplin, present,emerg, also,arrang, lot,share, apart,mani]

		svr	[charli,chaplin, ideal,time, new,wing, jim,brown, globe,democrat]
	bigram (-)	lasso	[rise,unemploy, one,touch, sky,new, sunday,news, near,offic]
		svr	[avail,next, sky,new, weill,hartmann, plymouth,cherbourg, student,american]
Expenditure	bigram (+)	lasso	[wkli,room, carpet,floor, industri,help, central,front, job,provid]
		svr	[wkli,room, case,continu, civil,work, big,price, carpet,floor]
	bigram (-)	lasso	[direct,charl, high,polish, new,regular, buy,build, wife,want]
		svr	[lillian,gish, window,sill, good,care, piec,piec, meet,high]
Exports	bigram (+)	lasso	[psycholog,test, fee,charg, store,restaur, order,larg, us,citizen]
		svr	[psycholog,test, suppli,larg, take,ad, size,larg, order,larg]
	bigram (-)	lasso	[street,mortgag, north,philadelphia, increas,unemploy, point,return, news,radio]
		svr	[show,earth, news,radio, entir,product, morn,melodi, servic,get]
GDP	bigram (+)	lasso	[call,servic, civilian,conserv, possibl,loss, much,cent, goe,effect]
		svr	[control,law, yet,close, order,full, beverag,control, citi,bill]
	bigram (-)	lasso	[mar,via, short,sell, intern,settlement, investor,new, citi,daili]
		svr	[intern,settlement, tax,cost, japanes,chines, even,entertain, avail,list]
Housing capital gain	bigram (+)	lasso	[state,ticket, thoma,day, nation,move, control,go, job,men]
		svr	[state,ticket, candid,repres, vote,governor, candid,state, candid,unit]
	bigram (-)	lasso	[school,establish, street,scene, lot,trade, time,english, major,member]
		svr	[case,continu, civil,work, citi,bill, week,beauti, continu,cold]
House Prices	bigram (+)	lasso	[gotten,man, learn,high, hand,busi, region,offic, time,yes]
		svr	[state,ticket, vote,governor, elect,governor, candid,state, washington,plan]
	bigram (-)	lasso	[short,sell, one,modern, dinett,bath, portug,spain, perman,guest]
		svr	[case,continu, short,sell, intern,settlement, tax,cost, theori,practic]
Housing rental return	bigram (+)	lasso	[ecuador,peru, team,stand, candid,repres, offer,wide, open,fifth]
		svr	[candid,repres, washington,plan, candid,unit, yes,sir, candid,state]
	bigram (-)	lasso	[general,good, freedom,press, place,made, kong,berg, propos,intern]
		svr	[newark,airport, passeng,leav, general,good, first,lesson, ice,snow]
Housing rental yield	bigram (+)	lasso	[current,busi, one,modern, east,wind, short,sell, rate,also]
		svr	[case,continu, tax,cost, build,open, current,busi, compani,less]
	bigram (-)	lasso	[immedi,payment, state,ticket, girl,live, servic,mani, purpl,heart]
		svr	[state,ticket, vote,governor, bank,holiday, elect,governor, friend,make]
Housing total return	bigram (+)	lasso	[state,ticket, thoma,day, nation,move, control,go, campaign,headquart]
		svr	[state,ticket, candid,repres, candid,state, vote,governor, candid,unit]
	bigram (-)	lasso	[school,establish, street,scene, lot,trade, auto,mart, time,english]
		svr	[case,continu, civil,work, citi,bill, ice,snow, continu,cold]
Investment-to-GDP ratio	bigram (+)	lasso	[chancellor,hitler, california,nation, bank,holiday, free,good, manag,food]
		svr	[bank,holiday, theatr,shop, candid,repres, co,born, judith,anderson]
	bigram (-)	lasso	[garden,swim, chines,japanes, employ,advertis, even,entertain, short,sell]
		svr	[tax,cost, even,entertain, japanes,chines, chines,japanes, intern,settlement]
Imports	bigram (+)	lasso	[put,full, use,great, candid,repres, receiv,payment, saint,john]
		svr	[put,full, candid,repres, amount,work, water,year, water,need]

	bigram (-)	lasso	[short,sell, north,texa, famous,beauti, coupl,also, rental,rate]
		svr	[east,open, theori,practic, chines,armi, show,ever, short,sell]
Long Term Interest Rate	bigram (+)	lasso	[continu,first, rail,road, meet,mayor, guarante,rate, box,box]
		svr	[hart,said, gari,hart, affair,would, colonel,north, death,squad]
	bigram (-)	lasso	[american,hospit, practic,problem, economi,move, air,india, polic,sergeant]
		svr	[american,hospit, artifici,heart, air,india, mrs,good, tax,propos]
Mortgage loans to non-financial private sector	bigram (+)	lasso	[rout,east, high,polish, meet,unit, busi,cd, line,train]
		svr	[excel,educ, fact,find, tour,avail, part,price, arizona,new]
	bigram (-)	lasso	[new,social, hand,first, drastic,cut, plan,bank, mani,direct]
		svr	[sky,new, pictur,exhibit, district,line, budget,tax, golan,height]
Narrow Money	bigram (+)	lasso	[buck,hill, bond,insur, cost,food, end,saturday, somewhat,better]
		svr	[new,night, sea,wall, wiki,leak, price,book, great,french]
	bigram (-)	lasso	[governor,landon, st,board, green,hous, new,differ, assembl,one]
		svr	[wood,ford, green,hous, assembl,one, greec,itali, man,field]
PI - Consumer Goods	bigram (+)	lasso	[room,music, put,full, new,wing, chou,eli, ride,free]
		svr	[satin,canton, new,wing, put,full, advisori,commiss, water,year]
	bigram (-)	lasso	[buy,buy, rise,cost, heavi,equip, one,cover, per,gallon]
		svr	[buy,buy, plymouth,cherbourg, st,pierr, need,never, report,area]
PI - Energy	bigram (+)	lasso	[seventh,fleet, connect,larg, famili,sketch, box,box, employ,work]
		svr	[connect,larg, tour,countri, blue,flower, time,hall, posit,best]
	bigram (-)	lasso	[weill,hartmann, today,get, princip,creditor, sale,season, price,crude]
		svr	[mrs,hall, sourc,hold, today,get, web,bush, ohm,alley]
PI - Manufacturing	bigram (+)	lasso	[ahead,year, end,end, advertis,caution, avail,everi, carri,nation]
		svr	[candid,unit, candid,state, state,ticket, candid,repres, vote,governor]
	bigram (-)	lasso	[short,sell, never,worn, present,busi, three,meet, influenc,peopl]
		svr	[island,open, young,young, chines,armi, intern,settlement, short,sell]
Pers. Consumption	bigram (+)	lasso	[california,nation, free,good, regist,owner, seventh,fleet, help,see]
		svr	[california,nation, help,see, opportun,help, see,job, drive,year]
	bigram (-)	lasso	[employ,advertis, save,fuel, build,high, oil,suppli, industri,help]
		svr	[fuel,shortag, spare,tire, station,owner, take,ad, station,oper]
Population	bigram (+)	lasso	[secretari,general, school,financ, trade,public, heat,low, garag,left]
		svr	[take,ad, brown,also, sovereign,state, former,east, firm,long]
	bigram (-)	lasso	[wireless,new, church,perform, transfer,tax, matter,relat, public,date]
		svr	[help,see, opportun,help, matter,relat, offic,continu, see,job]
Producer Price Index	bigram (+)	lasso	[rise,price, size,larg, night,music, bar,restaur, plan,plant]
		svr	[rise,price, size,larg, south,long, price,rise, price,rais]
	bigram (-)	lasso	[fall,price, french,citi, share,averag, hour,fli, law,graduat]
		svr	[young,young, well,heat, live,busi, music,children, yangtz,river]
Production Index	bigram (+)	lasso	[ahead,year, state,ticket, end,end, advertis,caution, avail,everi]
		svr	[candid,unit, state,ticket, candid,state, candid,repres, vote,governor]
	bigram (-)	lasso	[short,sell, never,worn, increas,unemploy, present,busi, influenc,peopl]

		svr	[island,open, young,young, chines,armi, intern,settlement, short,sell]
Public debt-to-GDP ratio	bigram (+)	lasso	[two,japanes, wage,cut, use,full, price,diamond, declin,year]
		svr	[tax,cost, even,entertain, avail,list, mani,plan, product,west]
	bigram (-)	lasso	[king,solomon, owner,plan, averag,earn, store,display, servic,world]
		svr	[main,featur, control,go, ahead,year, defens,order, king,solomon]
Real GDP (barro)	bigram (+)	lasso	[civilian,conserv, note,answer, italian,troop, restaur,also, pro,con]
		svr	[year,attract, continu,north, good,boy, direct,state, even,park]
	bigram (-)	lasso	[ziegfeld,theatr, iowa,year, short,sell, mar,via, posit,go]
		svr	[japanes,chines, tax,cost, intern,settlement, even,entertain, ziegfeld,theatr]
Real GDP (mad)	bigram (+)	lasso	[italian,troop, mckay,inc, restaur,also, pro,con, give,better]
		svr	[year,attract, candid,unit, continu,north, good,boy, big,price]
	bigram (-)	lasso	[posit,go, mar,via, short,sell, iowa,year, ziegfeld,theatr]
		svr	[home,christma, japanes,chines, tax,cost, tour,avail, intern,settlement]
Revenue	bigram (+)	lasso	[person,engag, broadcast,new, hour,also, east,coast, young,st]
		svr	[salari,first, decemb,open, card,game, mile,drive, cost,week]
	bigram (-)	lasso	[wind,near, sport,review, dalli,rate, time,club, girl,go]
		svr	[tax,cost, mani,plan, wind,near, time,club, avail,list]
Sales	bigram (+)	lasso	[star,movi, via,victoria, packard,limousin, refin,surround, store,way]
		svr	[via,victoria, doubl,reason, march,day, glass,menageri, port,de]
	bigram (-)	lasso	[product,declin, result,great, econom,slowdown, report,larg, farm,near]
		svr	[set,plan, call,young, lead,educ, use,regular, give,ad]
Short Term Interest Rate	bigram (+)	lasso	[open,white, best,play, valuabl,time, call,build, point,increas]
		svr	[open,white, give,day, first,busi, order,full, valuabl,time]
	bigram (-)	lasso	[declin,first, passion,play, point,cut, hous,industri, want,retir]
		svr	[made,see, anoth,fine, main,room, naval,confer, right,angl]
Unemployment	bigram (+)	lasso	[pool,gym, offer,low, news,farm, rise,unemploy, product,declin]
		svr	[schuyler,west, news,farm, homelik,hotel, pool,gym, avail,next]
	bigram (-)	lasso	[seventh,fleet, also,arrang, glass,menageri, citi,set, green,park]
		svr	[right,day, old,vic, seventh,fleet, pioneer,new, annul,associ]
Total loans to business	bigram (+)	lasso	[amount,said, provis,said, russian,front, deed,trust, program,product]
		svr	[dealt,unlist, amount,work, put,full, desk,stenograph, annex,salesmen]
	bigram (-)	lasso	[cent,dollar, paid,first, slash,price, must,quick, west,help]
		svr	[plymouth,cherbourg, weather,market, time,weather, postseason,rental, may,full]
Total loans to households	bigram (+)	lasso	[full,employ, pennsylvania,west, diamond,diamond, honor,new, general,new]
		svr	[war,power, hey,man, debtor,second, american,marin, word,insert]
	bigram (-)	lasso	[program,expand, bank,charg, also,build, dead,sea, new,social]
		svr	[sale,even, compani,experienc, district,line, heavi,equip, product,declin]
Total loans to non-financial private sector	bigram (+)	lasso	[high,polish, cement,block, rout,east, atom,bomb, call,fair]
		svr	[put,full, water,year, amount,work, fact,find, gallon,water]
	bigram (-)	lasso	[island,main, hand,first, institut,govern, econom,condit, tompkin,squar]
		svr	[island,main, arthur,schwartz, sky,new, free,storag, compani,opportun]

Wage	bigram (+)	lasso	[control,law, increas,cost, door,room, product,manag, beverag,control]
		svr	[civil,work, beverag,control, control,law, addit,person, open,point]
	bigram (-)	lasso	[wage,cut, sport,review, leagu,council, wind,feet, manag,sinc]
		svr	[avail,list, minus,minus, plus,minus, minus,plus, plus,st]
Women Employment	bigram (+)	lasso	[plan,inc, total,day, new,bedford, bureau,west, via,victoria]
		svr	[via,victoria, larg,use, also,capabl, barrett,corp, globe,democrat]
	bigram (-)	lasso	[rise,unemploy, sale,declin, music,love, rubber,futur, result,great]
		svr	[rubber,futur, twelv,clock, news,wiz, downtown,salesmen, mut,second]

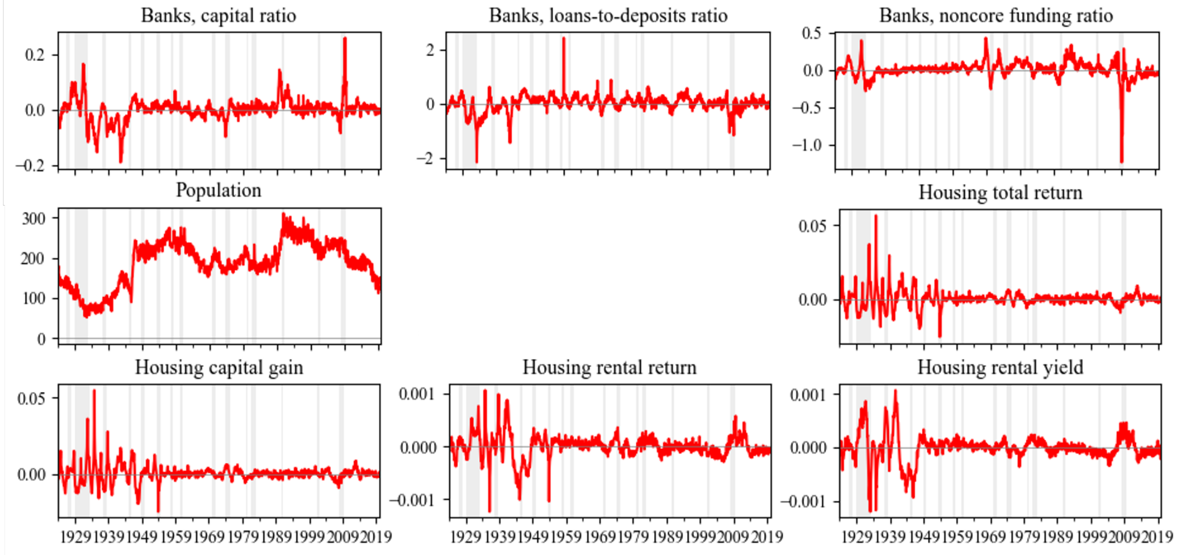

Figure A.1: **Additional text-based time series.** This Figure shows monthly text-based time series. Further details are in Appendix Table 2.5. Grey shaded areas are US recessions according to NBER definition.

Table A.3: **JST Marcohistory Database variables: not used.** This table reports variables not used from Jordà et al. (2017, 2019, 2021). Variables used for our analysis can be found in Table 2.5.

Variable	**Description**
eq_tr	Equity total return, nominal. r[t] = [[p[t] + d[t]] / p[t-1]] - 1
bond_tr	Government bond total return, nominal. r[t] = [[p[t] + coupon[t]] / p[t-1]] - 1
bill_rate	Bill rate, nominal. r[t] = coupon[t] / p[t-1]
rent_ipolated	1 if housing rental yields interpolated e.g. wartime
housing_capgain_ipolated	1 if housing capital gains and total returns interpolated e.g. wartime
eq_capgain	Equity capital gain, nominal. cg[t] = [p[t] / p[t-1]] - 1
eq_dp	Equity dividend yield. dp[t] = dividend[t]/p[t]
eq_capgain_interp	1 if equity cap. gain interpolated to cover exchange closure
eq_tr_interp	1 if equity total return interpolated to cover exchange closure
eq_dp_interp	1 if equity dividend interpolated or assumed zero to cover exchange closure
bond_rate	Gov. bond rate, rate[t] = coupon[t] / p[t-1], or yield to maturity at t
eq_div_rtn	Equity dividend return. dp_rtn[t] = dividend[t]/p[t-1]
capital_tr	Tot. rtn. on wealth, nominal. Wtd. avg. of housing, equity, bonds and bills
risky_tr	Tot. rtn. on risky assets, nominal. Wtd. avg. of housing and equity
safe_tr	Tot. rtn. on safe assets, nominal. Equally wtd. avg. of bonds and bills

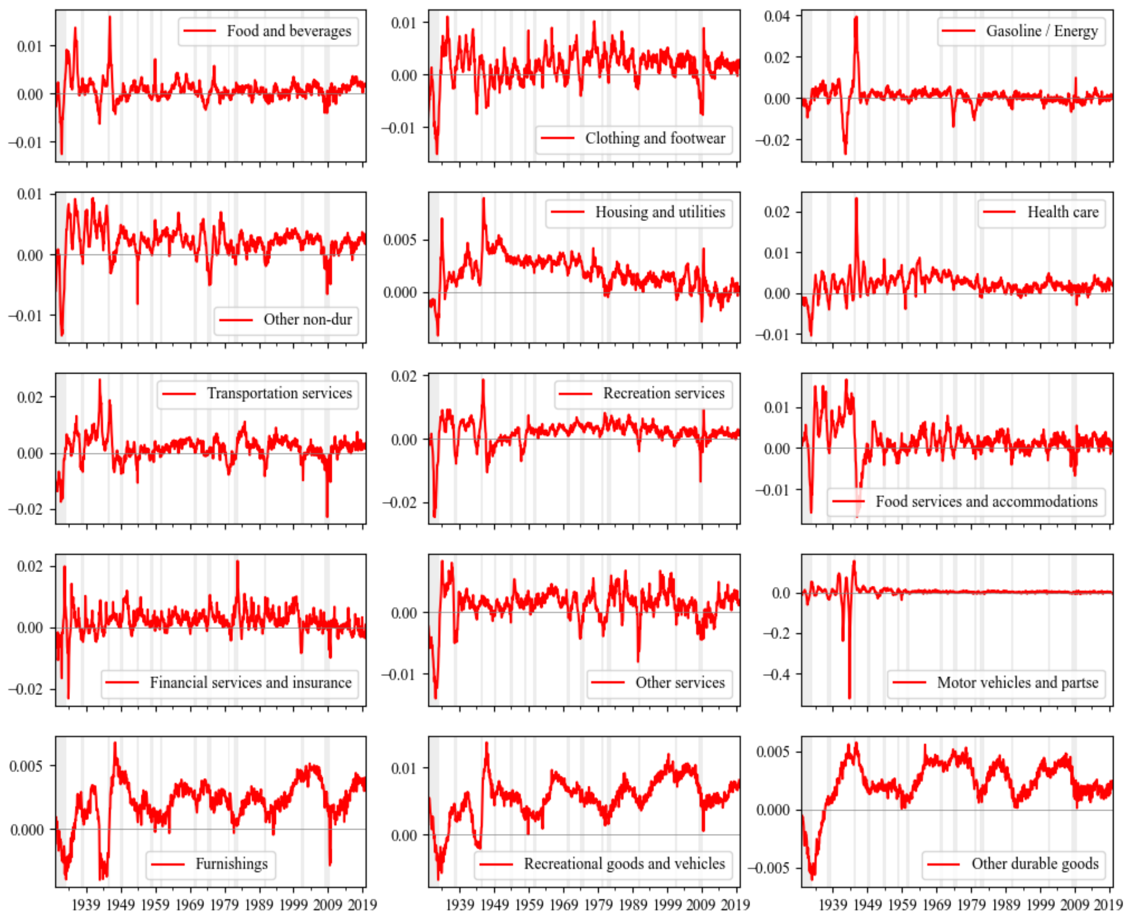

Figure A.2: **Consumption - quantity indices.** This Figure shows monthly text-based consumption quantity indices from NIPA (log growth rates). Grey shaded areas are US recessions according to NBER definition.

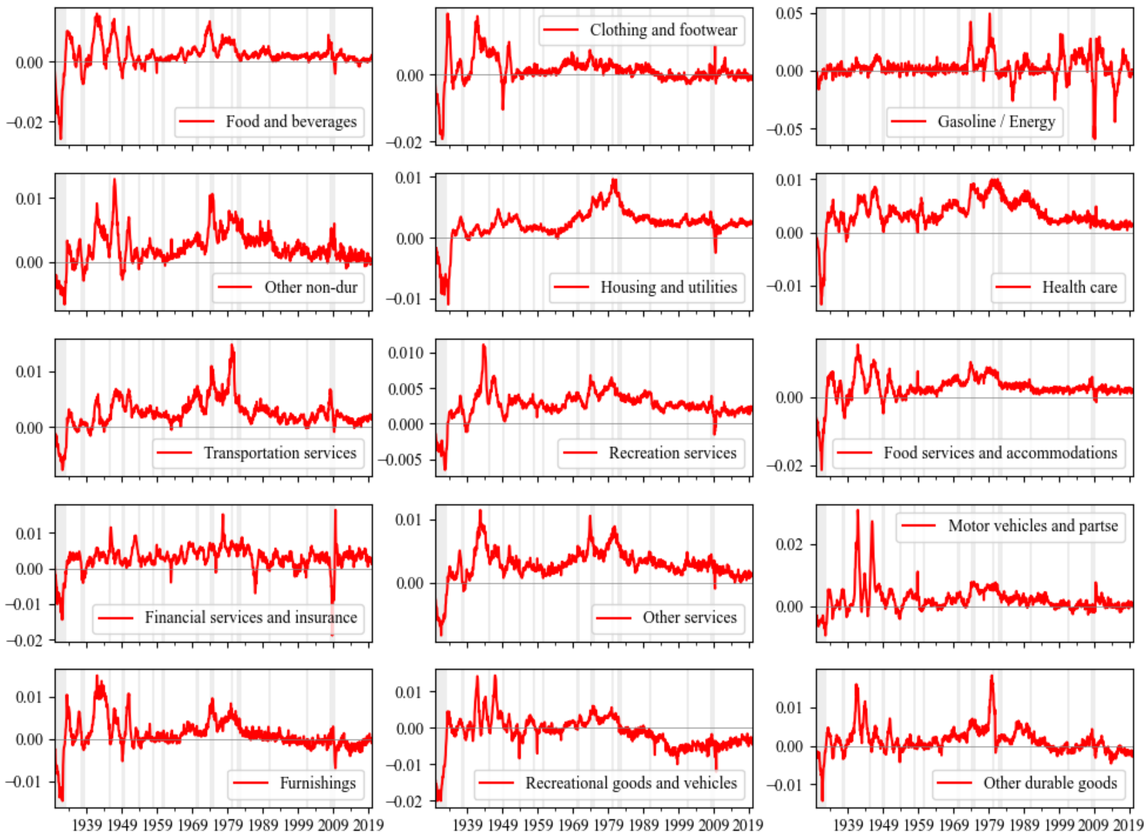

Figure A.3: **Consumption - price indices.** This Figure shows monthly text-based consumption price indices from NIPA (log growth rates). Grey shaded areas are US recessions according to NBER definition.

Table A.4: **Chosen bigrams per consumption index.** Annual data is obtained from the Bureau of Economic Analyses' website `https://apps.bea.gov/`. In specific, we use Table 2.4.3 (Real Personal Consumption Expenditures by Type of Product, Quantity Indexes) of the National Income and Production Accounts (NIPA) as well as Table 2.4.4 (Price Indexes for Personal Consumption Expenditures by Type of Product).

Type	**Consumption Subcategory**		**Model**	**Top 5 Bigrams**
Price Index	Clothing and footwear	bigram (+)	lasso	[servic,bar, price,rais, bar,bar, deputi,administr, broadcast,new]
			svr	[beverag,control, control,law, nation,free, mile,drive, price,rais]
		bigram (-)	lasso	[street,heir, citi,better, offic,window, price,valu, news,flash]
			svr	[tax,cost, news,flash, ideal,one, even,entertain, squar,subway]
	Financial services and insurance	bigram (+)	lasso	[time,paid, compani,tax, strike,would, call,pay, opportun,help]
			svr	[model,two, everi,famili, octob,stock, american,peac, market,collaps]
		bigram (-)	lasso	[short,sell, machin,age, share,averag, fall,day, room,view]
			svr	[arm,shipment, blue,velvet, import,detail, justic,rehnquist, arm,sale]
	Food and beverages	bigram (+)	lasso	[want,grand, control,price, price,rise, south,long, call,servic]
			svr	[south,long, price,rais, rise,price, requir,industri, yet,close]
		bigram (-)	lasso	[sport,review, danc,privat, build,perman, yangtz,river, offic,window]
			svr	[squar,subway, sport,review, deleg,larg, theori,practic, street,heir]
	Food services and accommodations	bigram (+)	lasso	[increas,price, want,grand, current,high, call,servic, product,manag]
			svr	[south,long, control,go, side,unit, us,also, set,exercis]
		bigram (-)	lasso	[sport,review, offic,window, busi,employ, danc,privat, build,perman]
			svr	[squar,subway, deleg,larg, theori,practic, deleg,nation, intern,settlement]
	Furnishings	bigram (+)	lasso	[air,sea, particular,need, increas,price, price,rais, war,front]
			svr	[particular,need, price,rais, ohio,river, mile,drive, rise,price]
		bigram (-)	lasso	[street,heir, girl,go, feet,also, citi,better, bridg,open]
			svr	[greatest,show, tax,cost, squar,subway, news,flash, mrs,collin]
	Gasoline / Energy	bigram (+)	lasso	[station,close, time,central, ad,appear, highest,possibl, site,site]
			svr	[manner,ann, bradley,said, one,air, fuel,shortag, station,close]
		bigram (-)	lasso	[side,show, ship,captain, sourc,hold, fall,price, novemb,american]
			svr	[side,show, sourc,hold, west,africa, sweet,briar, euro,one]
	Health care	bigram (+)	lasso	[report,pm, increas,averag, increas,cost, open,product, report,manag]
			svr	[model,two, thank,thank, hous,direct, author,propos, offic,administr]
		bigram (-)	lasso	[wind,near, offic,window, manag,sinc, easi,sell, easi,money]
			svr	[news,flash, tax,cost, ideal,one, wind,near, even,entertain]
	Housing and utilities	bigram (+)	lasso	[relat,industri, rise,cost, reader,letter, tabl,set, check,account]
			svr	[john,gielgud, school,establish, meet,high, daili,use, conserv,program]
		bigram (-)	lasso	[lowest,rental, titl,guar, disarma,confer, reich,tag, imperi,chemic]
			svr	[lowest,rental, morn,glori, danc,orch, beverag,control, drive,high]
	Motor vehicles and partse	bigram (+)	lasso	[increas,cost, tour,avail, broadcast,new, mani,posit, car,product]
			svr	[defens,bond, buy,buy, york,seek, broadcast,new, foot,forward]
		bigram (-)	lasso	[broadway,station, compani,meet, campaign,speech, line,phone, beau,champ]

			svr	[get,broadway, station,take, get,park, paterson,new, candid,repres]
	Other durable goods	bigram (+)	lasso	[price,rise, applic,may, broadcast,new, product,manag, soviet,troop]
			svr	[broadcast,new, mani,direct, winner,one, break,away, buy,buy]
		bigram (-)	lasso	[news,flash, three,girl, stori,ever, long,reach, human,race]
			svr	[chines,japanes, japanes,chines, news,flash, get,broadway, tax,cost]
	Other non-dur	bigram (+)	lasso	[increas,cost, produc,nation, suppli,would, industri,servic, need,product]
			svr	[queen,victoria, week,regular, spare,tire, co,avenu, web,mac]
		bigram (-)	lasso	[public,serv, busi,employ, name,month, peopl,refer, offic,window]
			svr	[name,day, goe,new, busi,employ, girl,go, citi,better]
	Other services	bigram (+)	lasso	[rise,cost, director,offic, plant,need, regul,permit, sunday,matine]
			svr	[fuel,shortag, ohio,river, energi,crisi, arab,oil, requir,industri]
		bigram (-)	lasso	[wage,cut, public,serv, peopl,refer, young,french, econom,depress]
			svr	[congress,adjourn, danc,privat, present,busi, greatest,show, publish,time]
	Recreation services	bigram (+)	lasso	[book,publish, hous,hour, control,act, manag,plant, price,control]
			svr	[bloomingdal,bloomingdal, releas,statement, ziegfeld,folli, essenti,worker, need,releas]
		bigram (-)	lasso	[public,serv, de,grass, color,scheme, mauric,chevali, wage,cut]
			svr	[calvin,coolidg, better,show, act,mayor, weekday,cent, mile,one]
	Recreational goods and vehicles	bigram (+)	lasso	[increas,cost, lillian,hellman, robert,gordon, obtain,write, new,send]
			svr	[defens,bond, book,john, york,seek, control,legisl, buy,buy]
		bigram (-)	lasso	[land,trust, daili,edit, manag,privat, grand,hotel, set,want]
			svr	[get,broadway, station,take, get,park, paterson,new, enjoy,good]
	Transportation services	bigram (+)	lasso	[rise,cost, account,may, replac,cost, deliveri,date, republ,nation]
			svr	[america,call, bring,friend, consid,nation, design,women, order,public]
		bigram (-)	lasso	[year,train, feet,new, talk,mrs, report,compani, becom,acquaint]
			svr	[cost,free, blue,law, merri,around, deleg,larg, french,belgian]
Quantity Index	Clothing and footwear	bigram (+)	lasso	[exempt,secur, store,may, hey,man, charli,chaplin, famili,church]
			svr	[trade,wind, heart,citi, develop,person, project,approv, hey,man]
		bigram (-)	lasso	[plus,plus, extraordinari,opportun, addit,capit, time,club, apart,fob]
			svr	[plus,plus, plus,minus, york,book, minus,plus, day,secur]
	Financial services and insurance	bigram (+)	lasso	[park,zoo, offer,hous, yorker,magazin, assembl,elect, tuesday,elect]
			svr	[new,congress, tuesday,elect, state,ticket, candid,unit, offer,hous]
		bigram (-)	lasso	[case,continu, organ,known, room,sever, york,night, itali,new]
			svr	[week,top, civil,work, case,continu, globe,democrat, control,law]
	Food and beverages	bigram (+)	lasso	[help,see, glass,menageri, first,lesson, hale,selassi, privat,dine]
			svr	[ref,ser, hale,selassi, servic,pay, project,approv, add,abba]
		bigram (-)	lasso	[news,flash, teach,posit, line,product, increas,rate, train,profession]
			svr	[news,flash, tax,cost, ideal,one, noel,coward, time,club]
	Food services and accommodations	bigram (+)	lasso	[old,red, continu,success, men,white, attack,north, restaur,store]
			svr	[civil,work, addit,person, men,white, citi,bill, control,law]
		bigram (-)	lasso	[way,white, weekday,cent, world,capit, also,type, fact,find]
			svr	[friend,make, calvin,coolidg, weekday,cent, month,car, slash,price]

Furnishings	bigram (+)	lasso	[tonight,tomorrow, sell,busi, use,also, resort,citi, american,white]
		svr	[compani,cd, departur,date, tonight,tomorrow, summer,travel, virginia,beach]
	bigram (-)	lasso	[servic,reduc, liber,educ, bank,examin, govern,guarante, increas,unemploy]
		svr	[servic,reduc, small,bank, unemploy,worker, increas,unemploy, chorus,orchestra]
Gasoline / Energy	bigram (+)	lasso	[california,nation, regist,owner, free,good, suit,go, glass,menageri]
		svr	[california,nation, help,see, opportun,help, see,job, drive,year]
	bigram (-)	lasso	[employ,advertis, mani,direct, build,high, oil,suppli, save,fuel]
		svr	[fuel,shortag, spare,tire, station,owner, take,ad, energi,crisi]
Health care	bigram (+)	lasso	[job,appli, help,see, item,one, california,nation, great,tradit]
		svr	[opportun,help, fact,find, help,see, california,nation, find,board]
	bigram (-)	lasso	[via,victoria, week,less, instrument,use, need,present, suggest,make]
		svr	[pictur,exhibit, spare,tire, oper,file, take,ad, new,congress]
Housing and utilities	bigram (+)	lasso	[domest,import, union,strike, dollar,price, us,post, averag,earn]
		svr	[civil,work, live,good, control,law, beverag,control, week,top]
	bigram (-)	lasso	[citi,larg, citi,far, loan,money, invest,public, leader,busi]
		svr	[citi,larg, calvin,coolidg, west,last, weekday,cent, elect,governor]
Motor vehicles and partse	bigram (+)	lasso	[first,armi, cours,high, present,music, time,athen, aircraft,design]
		svr	[get,broadway, paterson,new, get,park, cours,high, leav,everi]
	bigram (-)	lasso	[bloomingdal,bloomingdal, defens,bond, greatest,show, slide,scale, buy,buy]
		svr	[bloomingdal,bloomingdal, releas,statement, essenti,worker, need,releas, buy,buy]
Other durable goods	bigram (+)	lasso	[news,broadcast, nazi,war, call,war, web,time, book,order]
		svr	[year,cd, system,product, general,pinochet, daimler,chrysler, trite,baum]
	bigram (-)	lasso	[econom,condit, kan,texa, actual,cost, educ,work, product,season]
		svr	[new,district, bank,credit, mutual,benefit, econom,condit, extend,credit]
Other non-dur	bigram (+)	lasso	[well,servic, men,white, charli,chaplin, goe,new, italian,troop]
		svr	[trade,wind, charli,chaplin, loan,year, addit,person, control,law]
	bigram (-)	lasso	[avail,list, week,becom, increas,capit, short,sell, near,offic]
		svr	[even,entertain, tax,cost, news,flash, japanes,chines, avail,list]
Other services	bigram (+)	lasso	[pasadena,calif, room,door, court,opinion, well,man, assembl,elect]
		svr	[one,master, civil,work, addit,person, citi,bill, propos,agreement]
	bigram (-)	lasso	[secur,bank, plus,plus, barrel,crude, governor,roosevelt, cash,busi]
		svr	[farm,bill, iraq,kuwait, young,young, island,open, japanes,chines]
Recreation services	bigram (+)	lasso	[store,may, california,nation, citi,bill, term,high, march,end]
		svr	[drive,year, california,nation, opportun,help, see,job, high,train]
	bigram (-)	lasso	[minus,minus, short,sell, morn,melodi, month,hour, ideal,one]
		svr	[even,entertain, news,flash, tax,cost, ideal,one, time,club]
Recreational goods and vehicles	bigram (+)	lasso	[world,capit, happi,birthday, sell,part, traffic,congest, music,store]
		svr	[york,suburb, call,board, american,clipper, year,enjoy, share,week]
	bigram (-)	lasso	[servic,reduc, make,loan, product,season, paper,money, eth,cd]
		svr	[servic,reduc, make,loan, larg,bank, product,declin, increas,unemploy]
Transportation services	bigram (+)	lasso	[educ,nation, drive,year, orson,well, pc,east, market,advanc]

	svr	[drive,year, present,music, big,price, opportun,help, get,park]
bigram (-)	lasso	[short,sell, hous,sell, state,ask, sit,room, north,pole]
	svr	[iraq,kuwait, spare,tire, pictur,exhibit, rate,person, play,john]

Appendix B

Additional Information on *Understanding Asset Pricing Factors*

B.1 Excerpt from Coders' Guide

- Macroeconomic news & outlook: News relating to macroeconomic conditions, forecasts or reports such as inflation, housing prices, unemployment, employment, personal income, industrial production, manufacturing activity, etc. Also included are the following:
  → News about credit conditions and financial crisis developments that does not fall into another category such as "Monetary Policy & Central Banking".
  → News about trade matters (trade surplus/deficit) and exchange rates NOT due to policy developments.
  → Articles that attribute stock market moves to a shift in sentiment about the macroeconomic environment, even when the article does not point to a specific piece of news about the macroeconomic outlook.
- Monetary policy & central banking: Actions, possible actions, and concerns related to the conduct and policies of the central bank or similar authority. Such actions and policies pertain to interest rate changes and monetary policy announcements, inflation control, liquidity injections by the monetary authority, changes in currency-gold convertibility under a gold standard, changes in reserve requirements or other bank regulations used by the monetary authority to exercise control over monetary conditions, lender-of-last resort actions, and extraordinary actions by the monetary authority in response to bank runs, systemic financial crisis and threats to the payments system.

  Distinguishing Monetary Policy & Central Banking from Macroeconomic News & Outlook: Some news articles that discuss market reactions to macro developments also discuss the Fed's normal response to the macro development. Generally, we code an article as "Macro News & Outlook" if it attributes the market move to news about the macro economy. We code it as "Monetary Policy & Central Banking" if the article attributes the market move to (a) shifts in how the Fed responds to a given macro development or (b) news about unexpected consequences of Fed actions. It is helpful to approach this classification issue from a Taylor Rule perspective. Consider the following cases:
  1. Macro news: The market moves because it anticipates or speculates (or sees) that the Fed will respond in its usual manner to news about the macro economy. That is, the market anticipates or speculates that the Fed will respond to macro developments according to a Taylor Rule or other well-defined, well-understood description of the Fed's interest-rate setting behavior.
  2. Monetary policy: The market moves because of a surprise change in the policy interest rate – i.e., a surprise conditional on the state of the macro economy. From a Taylor Rule perspective, we can think of this change as a new value for the innovation term in the Taylor rule.
  3. Monetary policy: The market moves because of an actual or potential change in the Fed's policy rule. From a Taylor Rule perspective, this event corresponds to an actual or potential change in the form of the Taylor Rule or a change in specific parameter values. A concrete example would be a big market response to proposals to increase the target interest rate.
  4. Monetary policy: The market moves because of news that leads to revised views or concerns about the consequences of the Fed's actual or anticipated actions.

  If an article fits the description for Case 1 above, code its category as Macroeconomic News & Outlook. If an article fits the descriptions in Items 2, 3 or 4, code its category as "Monetary Policy & Central Banking."

B.2 ChatGPT

B.2.1 Prompt

"You are a language model which determines the main reason for stock market movements in each text file provided. You only know the following 16 categories as reasons: Commodities; Corporate earnings & outlook; Elections & Political Transitions; Exchange Rate Policy & Capital Controls; Foreign Stock Markets; Government spending; International trade policy; Macroeconomic news & outlook; Monetary policy & central banking; Other non-policy; Other policy; Regulation; Sovereign military & security; Taxes; Terrorist attacks & large-scale violence by non-state actors; Unknown & No Explanation. You cannot use any other category for your classification.

[The 16 categories along with examples are described in detail here.]

You can only return your result in the following format:

[date;primaryCategory;journalistConfidence;easeOfCoding;keyPassage], where

- date: write the publishing date of the article.
- primaryCategory: write the one category from the 16 above that describes best the reason for stock market movements in the article.
- journalistConfidence: write either "low", "medium" or "high" to describe how confident the article is about the reasons it gives for stock market movements. You cannot use any other values.
- easeOfCoding: write either "easy", "medium" or "hard" to describe the difficulty level of making a category choice. You cannot use any other values.
- keyPassage: quote the key passage of the article that best explains your classification choice. Limit this keyPassage to one sentence.

Only return the provided output format and no additional explanations.

A valid output example is [1965_06_28;Taxes;high;hard;Biden increases taxes on the wealthy, which causes stock prices to fall]"

B.2.2 Examples of ChatGPT Category Choices Compared to Human Coders' Choices

- 1974-09-09: ChatGPT categorizes the jump as **Commodities** based on the following NYT article passage:

  "The decline on the New York Stock Exchange was broadly based, with the biggest losses in the gold mining and silver issues. Most of the gold mining issues fell more than 8 points."

  Evaluation: This passage mentions big losses in gold and silver issues. However, this is not the reason for stock price fluctuations. In an earlier passage it is stated: "Investors' continued concern about the economy and over President Ford's decision to pardon Richard M. Nixon sent the stock market down sharply yesterday", which reflects jump causes to origin from changes in investors' macroeconomic outlook. In the sampled cases the categorization as Commodities is not suitable. We find words related to commodities (e.g. "gold", "fueled") in the cited key passage, yet these words do not indicate reasons for stock movements.
- 1974-09-16: ChatGPT categorizes the factor innovation as **Commodities** based on the following NYT article passage:

  "The stock market ended its five-session losing streak with a late technical recovery, resulting in the Dow rising 12.59 points to close at 639.78, supported by increased buying especially in blue-chip stocks. Turnover on the Big Board surged to 18.37 million shares, the largest in over three months. The market was fueled by hopes of a decline in the prime rate and short-covering by traders, while some analysts expect the recent technical rally to turn into a worthwhile advance for investors."

  Evaluation: This passage does not mention commodity influences. The only word often connected with a commodity (i.e. oil) is "[...] market was *fueled* by hopes" which, however, is only a figurative word choice here.

- 1966-08-26: ChatGPT categorizes the factor innovation as **Foreign** stock markets based on the following NYT article passage:

  "In a new low for the year, the stock market experienced steep declines, with concentrations in the electronics and aviation sectors. The Dow-Jones industrial average hit its lowest point since January 1964. Notable stocks like Motorola, Fairchild Camera, and Texas Instruments all suffered significant losses."

  Evaluation: From the text passage ChatGPT chose it is not clear why foreign stock markets are responsible for stock price movements in the US. The stock market (in the US) is stated as declining with especially strong losses in the electronics and aviation industries. From this passage alone a categorization as **Corporate** would be better fitting.

- 2001-02-07: ChatGPT categorizes the factor innovation as **Foreign** stock markets based on the following NYT article passage:

  "Stock prices fell yesterday, as traders and investors sold shares after a disappointing earnings report from Cisco Systems, the maker of computer-networking equipment. [...] But a rumor that a settlement had been reached in the Microsoft anti-trust case promoted a turnaround, narrowing the day's losses"

  Evaluation: ChatGPT and the human coders baser their categorization on the same text extract but came to different conclusions. While our human coders correctly understood that the key driver for stock movements are news on Cisco's weak earnings, ChatGPT does not make this connection. Instead, the AI finds **Foreign** stock markets as the key driver for price movements, although foreign markets are not mentioned at all.

- 2009-04-16: ChatGPT categorizes the factor innovation as **Foreign** stock markets based on the following NYT article passage:

  "Solid profits at the banking giant JPMorgan Chase gave investors another reason to feel optimistic on Thursday, and they pushed stocks sharply higher, shrugging off another batch of shaky economic data."

  Evaluation: Again, ChatGPT and the human coders baser their categorization on the same text extract but came to different conclusions. While our human coders correctly infer that the key driver for stock movements are news on JPMorgan's strong profits causing optimism in the market, ChatGPT does not make this connection. **Foreign** markets are not mentioned at all.

- 2011-10-17: ChatGPT categorizes the factor innovation as **Foreign** stock market, based on the following NYT article passage:

  "Unlike the previous week's strong gains, stocks on Wall Street retreated on Monday as the outlook for a broad solution to the European debt crisis appeared to wane. The Dow and the Nasdaq indexes were down for the year, and the S&P pressed deeper into negative territory for the year. Stocks vulnerable to economic stress took hits, while bond prices in the United States rose with the Treasury's 10-year note yield falling to 2.16 percent"

  Evaluation: The main reason for stocks retreating is the negative outlook on a solution to the European debt crisis, according to the article. Although the origin of the news innovation stems from Europe, the reason is not stock market movements but political decision making. The human coders chose **Macro** instead which seems plausible as the overall macroeconomic outlook is worsened through the lack of solution finding in Europe. Alternatively, **Other (political)** would have been a suitable category.

- 2011-12-08: ChatGPT categorizes the factor innovation as **Macro**economic news & outlook, based on the following NYT article passage:

  "Stocks fell sharply on Wall Street on Thursday after the European Central Bank appeared to dampen expectations for an expanded bond-buying program and as leaders gathered for a summit meeting in Brussels aimed at resolving the sovereign debt crisis in Europe. The E.C.B cut its benchmark interest rate for the second month in a row and expanded the emergency funding it provides to cash-starved banks."

  Evaluation: The bond-buying program of the European Central Bank is an instrument of monetary policy. Here, **Monetary** would have been the better choice.

- 2022-07-29: ChatGPT categorizes the factor innovation as **Macro**economic news & outlook, based on the following NYT article passage:

  "The S&P 500 had its best month since November 2020. The rebound in stocks is a reflection that the current round of updates from corporate America are not as bad as feared. But there are still reasons for investors to be wary."

  Evaluation: The reason for the rebound in stocks is that corporate earnings postings are better than expected. Here, **Corporate** would have been the better choice.

B.3 Clarity

Following Baker et al. (2021) we decide to measure overall clarity of a jump cause. This is necessary because newspaper articles differ in communicating how explicit the economic reason of a jump is. Early in our coding process we realized that there are articles about the stock market without a clear opinion of the journalist on why stock prices moved. This lead us to introduce an additional measure "own evaluation". So we end up measuring clarity with three different proxies:

1. Own Evaluation was selected by coders if they saw a reason for stocks jumping on a day but cannot find a suitable article or the article does not offer an explanation for stock jumps. Coders then filled out the primary and secondary category based on their own evaluation.
2. Ease of Coding (EoC) defines how easy it is to code the "Primary Category".
3. Journalist Confidence (JC) characterizes the confidence/assurance/certainty with which the article describes an explanation for the defined jump.

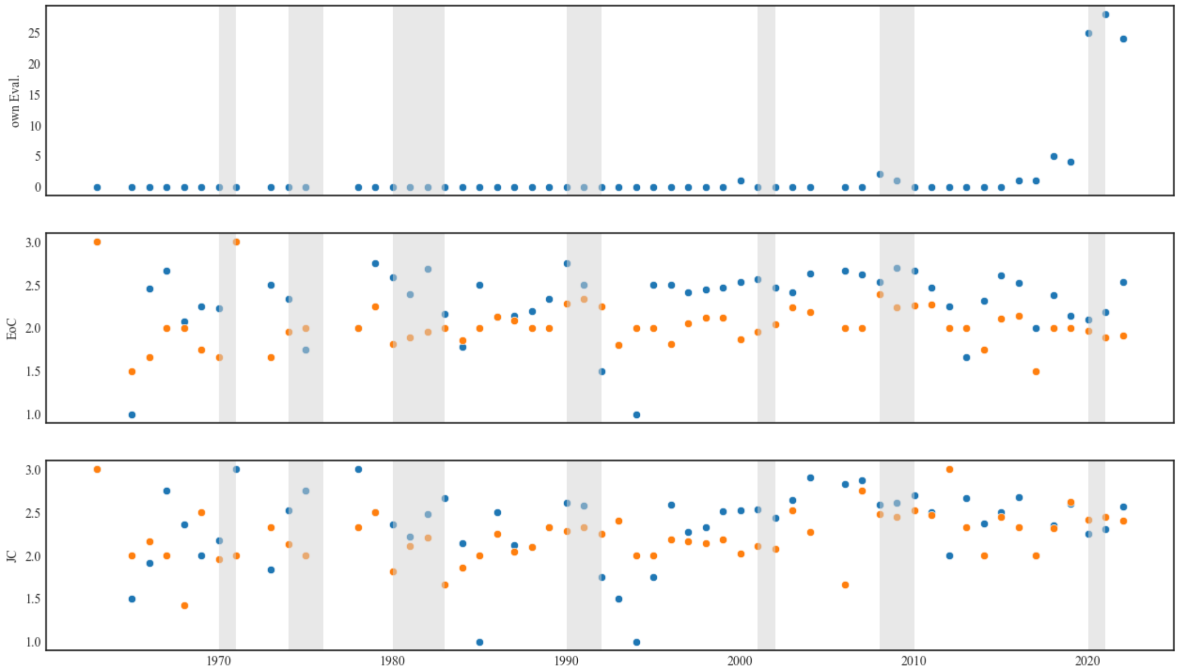

Figure B.1: **Clarity plots.** Blue points represent our data set with only human coders, orange points are with AI coder. Grey shaded areas indicate US recessions according to NBER's definition. EoC and JC are annual averages. Own Eval. is the annual sum. Coders choose "Own Evaluation" if the coder based their category choice on their own evaluation as no journalist explanation is given. EoC is "Ease of Coding" (1 = difficult, 2 = medium, 3 = easy). JC is "Journalist Confidence" and characterizes the confidence/assurance/certainty with which the article advances an explanation for the jump (1 = low, 2 = medium, 3 = high).

B.4 Additional Figures

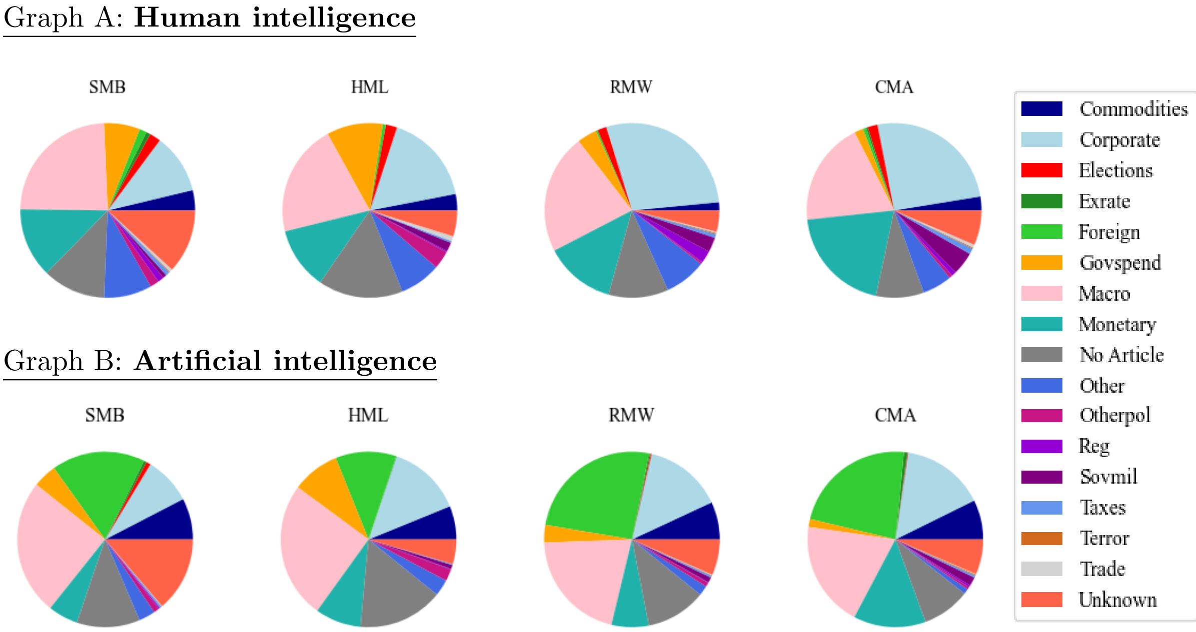

Figure B.2: **Relative shares of quadratic variation.** This figure shows the share each category contributes to the sum of squared factor returns between 1963 and 2022. Share of category i in factor k are calculated as follows $share_{k,i} = \sum_{t\epsilon T_i} r_k^2 / \sum_T r_k^2$.

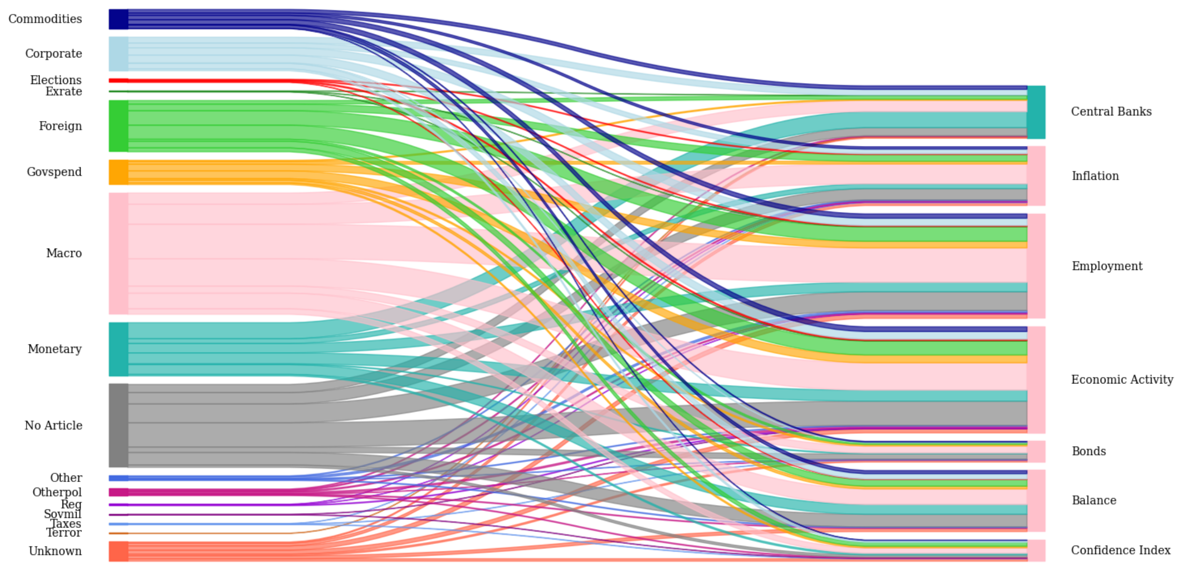

Figure B.3: **Announcements (AI).** This figure shows the scheduled announcement categories (right hand side) as well as our human coded categories on announcement days (left hand side). Colors on the right hand side indicate the major contributor (pink = macro, green = monetary). Announcement dates are retrieved from the economic calendar from `https://www.investing.com/economic-calendar/`.

Appendix C

Additional Information on *CLO Trading of Brown Loans*

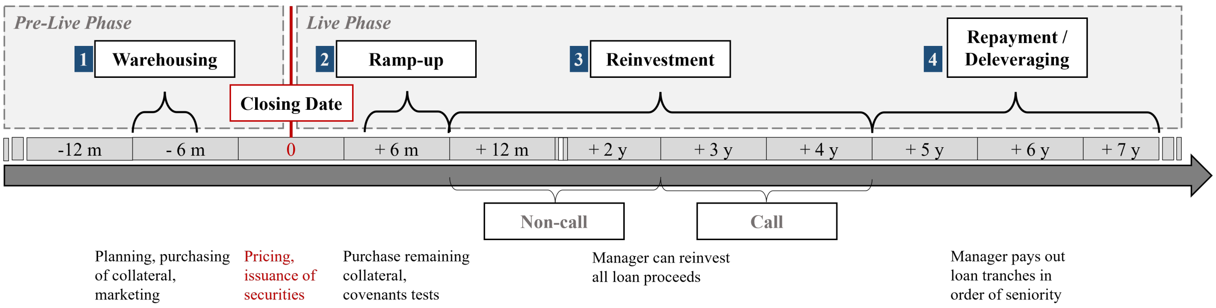

Figure C.1: **CLO life cycle.** This Figure shows a schematic timeline of an exemplary CLO. Age count starts with the closing date. As CLOs can be called any time after the non-call period ends, CLOs can be redeemed before their natural maturity date. Source: Kundu (2022) and Kollmorgen and Oh (2022), authors' own illustration.

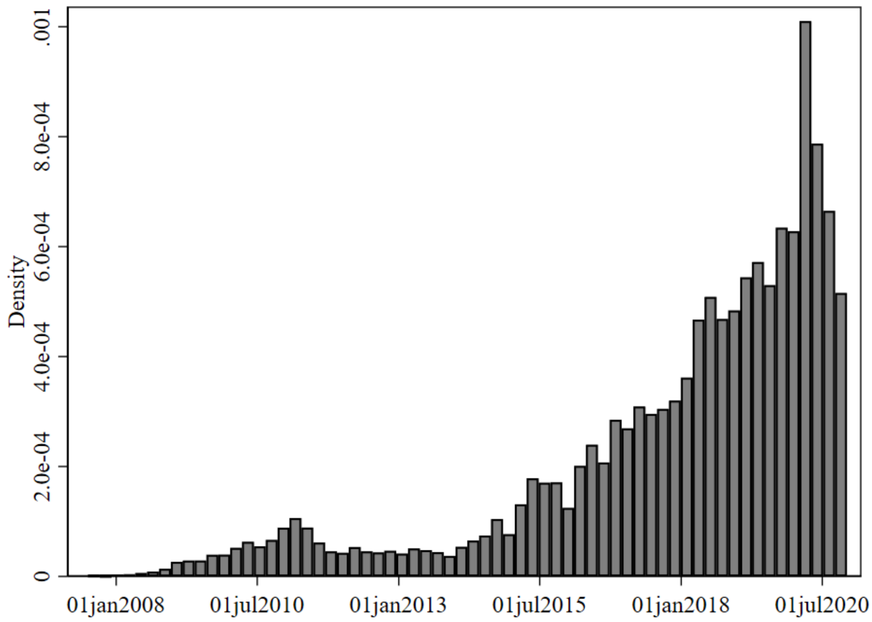

Figure C.2: **Histogram of CLO trades between 2008 and 2020.** Figure C.2 shows the growing importance of CLOs on the secondary market for leveraged loans through plotting the number of CLO trades in our data between 2008 and 2020 within our CLO-i dataset. Note that after the global financial crisis, CLO loan trading began. After stabilizing between 2011 and 2014, CLOs' trading entered a strong growth phase lasting until 2020.

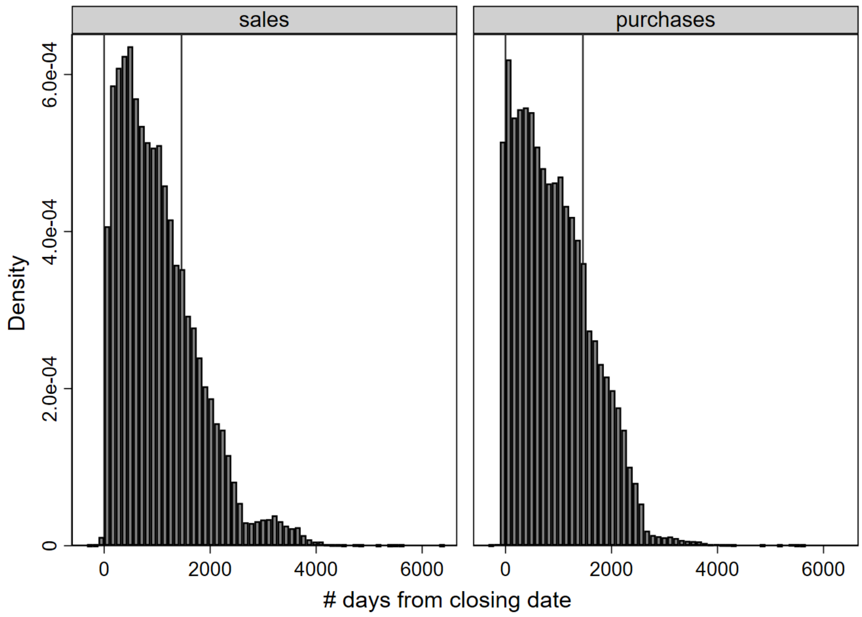

Figure C.3: **Histogram of CLO age and loan transactions.** This figure shows a histogram of CLO transactions. Purchases are on the right side, sales on the left. Vertical lines indicate life cycle phases of CLOs: First, warehousing (<0 days), second, ramping up and reinvestment phase (0 to 1460 days), and third, deleveraging (> 1460 days).

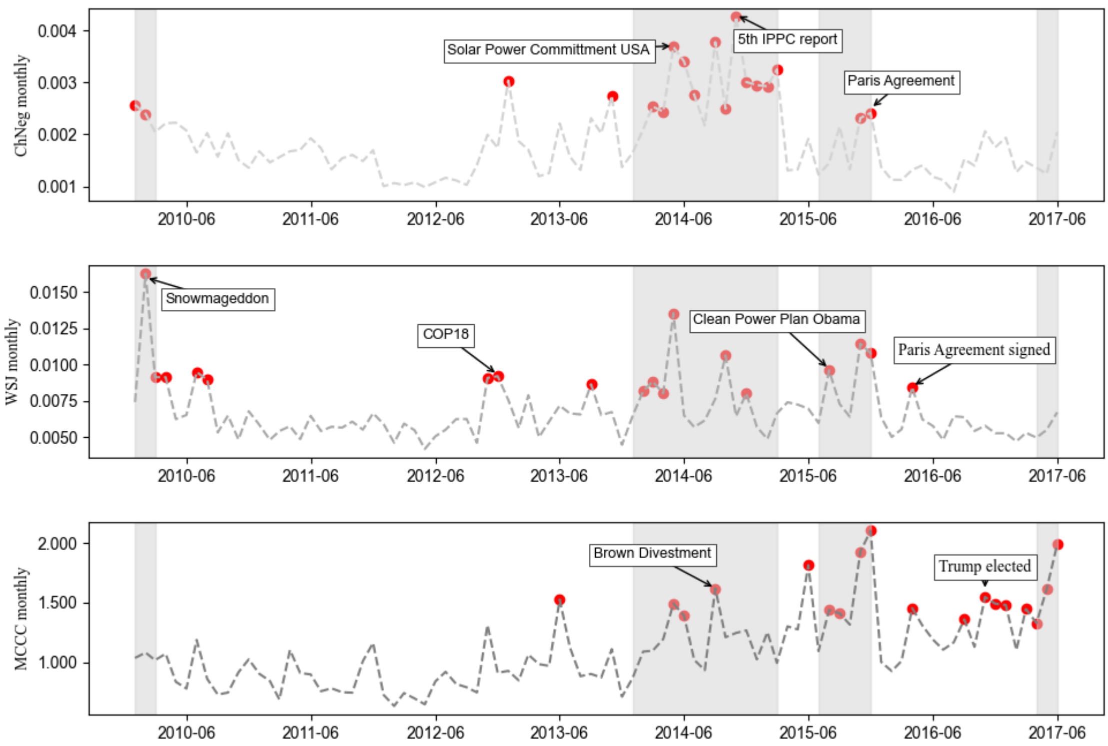

Figure C.4: **Three news attention indices.** This figure was created using monthly climate change attention indices of Engle et al. (2020), i.e. the CHneg index and the WSJ index, and of Ardia et al. (2023), i.e. the MCCC index. Shaded areas are quarters in which indices spike overall at least three times. Red dots mark the 80th percentile for each index.

Table C.1: **CLO asset managers and PRI joining date.** Information on asset managers and their joining date is obtained from the website of the PRI.

Asset Managers	**joined PRI on**	**Asset Managers**	**joined PRI on**
40/86 Advisors	-	Insight Investment	27.04.2006
Acis Capital Management	-	Intermediate Capital Group plc	29.04.2013
Aegon N.V.	03.11.2022	Invesco Ltd	12.07.2013
AGL Credit Management	20.12.2021	Investcorp Holdings B.S.C.(c)	09.02.2021
AIG Asset Management	-	Jefferies Finance LLC	02.12.2022
Airlie Capital Management	-	JMP Group	-
Alcentra	20.06.2018	Kayne Anderson Capital Advisors	05.09.2019
AllianceBernstein	-	KCAP Financial	-
Allianz Global Investors	23.04.2007	Kingsland Capital Management	-
Allstate Investment Management	-	KKR	-
Alvarez & Marsal	-	KVK Credit Strategies	-
American Capital Strategies	-	LCM Asset Management	28.11.2022
American Money Management	-	Loomis, Sayles & Company, L.P.	22.05.2015
Anchorage Capital Group, L.L.C.	07.10.2020	Lufkin Advisors	-
Angelo Gordon	25.10.2021	M&G (Prudential Assurance Company)	16.06.2021
Antares	25.04.2023	Madison Investment Holdings, Inc.	18.10.2021
Apex Group Ltd	20.09.2019	Man Group	04.08.2017
Apollo Global Management, Inc.	14.10.2020	Maranon Capital, LP	19.07.2021
Ares Management Corporation	21.05.2020	Marathon Asset Management, LP	04.08.2017
ArrowMark Partners	-	Marble Point Capital Management	-
Avenue Capital Group	12.06.2020	MatlinPatterson Global Advisers	-
AXA Investment Managers	29.05.2007	Medalist Partners	-
Bain Capital	17.05.2022	Mercer Advisors, Inc	18.09.2019
Ballyrock Investment Advisors	-	MFS Investment Management	01.02.2010
Bardin Hill Investment Partners LP	06.12.2021	MidCap	-
Barings LLC	02.01.2014	MidOcean Partners	09.01.2020
Benefit Street Partners	-	MJX Asset Management LLC	24.08.2021
Birch Grove Capital LP	22.07.2020	Monroe Capital LLC	01.04.2021
Black Diamond Capital Management	28.01.2020	Morgan Stanley Investment Mgmt.	30.10.2013
BlackRock	07.10.2008	Napier Park Global Capital	10.02.2020
BlueMountain Capital Management	-	Nassau Financial Group, L.P.	07.02.2023
BlueMountain Fuji CLO Management	-	Neuberger Berman Group LLC	29.06.2012
BMO Global Asset Management	20.12.2019	New York Life Investment Mgmt. LLC	04.11.2019
BNP Paribas Asset Management	27.04.2006	Newfleet Asset Management	-
Bradford and Marzec	-	NewMark Capital Management	-
Brigade Capital Management, LP	26.03.2020	NewStar Capital	-
Canaras Capital Management	-	NewStar Financial	-
Canyon Partners, LLC	26.01.2021	Nicholas-Applegate Capital Mgmt.	-
Carlson Capital	-	NXT Capital	-
Carlyle	22.04.2022	Oak Hill Advisors	11.03.2019
CarVal CLO Management	-	Oaktree Capital Management	27.09.2019
CBAM	-	Och Ziff	-
Cerberus Capital Management, L.P.	02.03.2023	Octagon Credit Investors, LLC	04.02.2020
Chicago Fundamental Investment Partners	-	Onex	29.10.2021
Churchill Investment Management	07.09.2021	Orchard Capital Management	02.07.2020
CIFC LLC	16.01.2020	Owl Rock Capital Advisors	-
CIT Asset Management	-	Pacific Life	-
Columbia Management	-	Palmer Square Capital Management	23.01.2020
Covenant Capital Group, LLC	18.09.2018	Par Four Investment Management	-
Credit Suisse Group AG	06.01.2014	Park Avenue Institutional Advisers	-
Credit Value Partners	-	Partners Capital	15.06.2020
Crescent Capital Group	21.02.2018	PGIM Fixed Income	10.02.2015
Crestline Denali Capital	-	PIMCO	13.09.2011
CVC Credit Partners	29.04.2021	PineBridge Investments	22.06.2015
DFG	-	Post Advisory Group	-
Diameter Capital Partners LP	27.02.2023	PPM America, Inc.	02.10.2018
Doral Money Management	-	Pretium	01.09.2022
DoubleLine	27.01.2023	Princeton Advisory Group	-
DWS Group	29.02.2008	Redding Ridge	-
Eaton Vance	-	Rockford Tower Capital Management	-
Ellington Management Group	28.10.2021	Saratoga Investment Management	-
Elmwood Asset Management LLC	05.02.2021	Sculptor Capital Management, Inc.	11.01.2021
Fidelity Investments	23.02.2017	Seix Investment Advisors LLC	16.07.2021
First Eagle Investments	23.09.2020	SHENKMAN CAPITAL MGMT., INC.	25.08.2017
Five Arrows	14.09.2012	Silvermine Capital Management	-
Fore Research and Management	-	SilverPoint Capital	-
Fort Washington Investment Advisors, Inc	17.11.2016	Sound Harbor Partners	-
Fortis Investments	-	Sound Point Capital Management, LP	10.05.2021
Fortress Investment Group	-	Steele Creek	-
Four Corners Capital Management	-	Symphony Asset Management	-
Franklin Templeton Investments	01.05.2013	Tall Tree Investment Management	-
FS KKR Capital Corp	-	Telos Asset Management	-
Gallatin Loan Management	-	Tennenbaum Capital Partners	-
Garrison Investment Group	-	The TCW Group, Inc.	07.02.2019
Global Leveraged Capital	-	THL Credit Advisors	-
GoldenTree Asset Management	30.04.2021	TIAA	-
Goldman Sachs Asset Management (GSAM)	15.12.2011	TICC Management	-
Golub Capital	21.01.2022	TPG Capital Advisors, LLC	17.06.2013
Grandview Capital Management	-	Tricadia Capital Management	-
Greywolf Capital Management LP	12.06.2023	Trimaran Advisors	-
GSO Capital Partners	-	Trinitas Capital Management, LLC	09.11.2022
Guggenheim Partners Investm. Mgmt., LLC	27.02.2020	Voya Investment Management	21.12.2017
HalseyPoint Asset Management	-	Wellfleet Credit Partners	-
HarbourView Asset Management	-	Western Asset Management Company	02.02.2016
Hayfin	12.06.2018	White Star Capital	04.10.2019
HIG WhiteHorse	-	WR Huff Asset Management	-
Highland Capital Management	-	York Capital Management	-
HPS Investment Partners	29.01.2021	Z Capital Group, LLC	12.05.2023
ICE Canyon	-	ZAIS Group, LLC	03.07.2019
Indosuez Capital	-		

Table C.2: **Industry match GICS to Moody's classification.** We match Moody's industry classifications, found in our CLO-i data set, to Global Industry Classification Standards (GICS 6) from Bolton and Kacperczyk (2021).

GICS Industry	Moody's Industry
Air Freight & Logistics	Cargo Transport
Airlines	Personal Transportation
Auto Components	Automobile
Automobiles	Automobile
Banks	Banking
Beverages	Beverage, Food and Tobacco
Biotechnology	Healthcare, Education and Childcare
Capital Markets	Finance
Chemicals	Chemicals, Plastics and Rubber
Construction Materials	Diversified/Conglomerate Manufacturing
Consumer Finance	Finance
Diversified Consumer Services	Healthcare, Education and Childcare
Diversified Telecommunication Services	Telecommunications
Electric Utilities	Utilities
Equity Real Estate Investment Trusts (REITs)	Buildings and Real Estate
Food & Staples Retailing	Grocery
Food Products	Beverage, Food and Tobacco
Gas Utilities	Utilities
Health Care Technology	Healthcare, Education and Childcare
Household Products	Home and Office Furnishings, Houseware
Independent Power and Renewable Electricity Producers	Utilities
Industrial Conglomerates	Diversified/Conglomerate Manufacturing
Interactive Media & Services	Broadcasting and Entertainment
Internet Software & Services	Diversified/Conglomerate Service
Life Sciences Tools & Services	Healthcare, Education and Childcare
Marine	Aerospace and Defense
Metals & Mining	Mining, Steel, Iron, and Non-Precious Metals
Mortgage Real Estate Investment Trusts (REITs)	Buildings and Real Estate
Multiline Retail	Retail Stores
Multi-Utilities	Utilities
Oil, Gas & Consumable Fuels	Oil and Gas
Paper & Forest Products	Diversified Natural Resources
Professional Services	Diversified/Conglomerate Service
Real Estate Management & Development	Buildings and Real Estate
Software	Diversified/Conglomerate Service
Thrifts & Mortgage Finance	Finance
Tobacco	Beverage, Food and Tobacco
Water Utilities	Utilities

Table C.3: **Industry match GICS to Moody's classification: Emissions.** We use Bolton and Kacperczyk (2021), who report the top and bottom 10 of GIC 6 industries in terms of average emission production in scope 1, scope 2, and scope 3 (sample period 2005–2017). The emissions are expressed in tons of CO2. See C.2 for the matching between GICS 6 and Moody's.

Moody's Industry	**Category**	**Scope 1**	**Scope 2**	**Scope 3**	**Total (in T)**
Aerospace and Defense			21,798		21.80
Automobile	brown		2,094,174	21,985,134	24,079.31
Banking	green	6,965	45,627	116,073	168.67
Beverage, Food and Tobacco				24,753,485	24,753.49
Broadcasting and Entertainment	green	7,649			7.65
Buildings and Real Estate		8,770	36,013	214,133	258.92
Cargo Transport	brown	4,316,221			4,316.22
Chemicals, Plastics and Rubber	brown	3,280,770	1,475,783		4,756.55
Containers, Packaging and Glass	brown				
Diversified Natural Resources, Precious Metals and Minerals		3,286,922	1,375,637		4,662.56
Diversified/Conglomerate Manufacturing		3,827,648	1,014,037	6,575,213	11,416.90
Diversified/Conglomerate Service		35,925	22,653	114,132	172.71
Ecological					
Electronics					
Farming and Agriculture					
Finance	green	14,837	37,160	15,193	67.19
Grocery			2,163,081	5,882,429	8,045.51
Healthcare, Education and Childcare	green	11,657	80,512	361,180	453.35
Home and Office Furnishings Housewares and Durable Consumer Products			994,783	4,313,762	5,308.55
Hotels, Motels, Inns and Gaming					
Insurance					
Leisure, Amusement and Entertainment					
Machinery (Non-Agriculture Non-Construction Non-Electronic)					
Mining, Steel Iron and Non-Precious Metals	brown	6,343,545	1,749,360	3,580,245	11,673.15
Oil and Gas	brown	6,302,663	820,777	6,049,237	13,172.68
Personal Transportation	brown	17,600,000			17,600.00
Personal and Non-Durable Consumer Products (Manufacturing Only)					
Personal Food and Miscellaneous Services					
Printing and Publishing					
Retail Stores			825,501		825.50
Sovereign and Supranational					
Telecommunications			1,219,956		1,219.96
Textiles and Leather					
Utilities	brown	81,200,000	39,177	41,182	81,280.36

Table C.4: **Prices without Industry FE.** In this table we report results from estimating the following equation (equation 4.2): $Price_{l,r,\tau} = \beta\, Attention_{t-1} \times Brown_r + \gamma_x \mathbf{X}_t + \alpha_\tau + \alpha_{rat} + \alpha_i \times \alpha_t + \epsilon_{l,r,\tau}$. The sample covers daily purchases during reinvestment periods of CLOs over the period 1 January 2010 - 31 May 2018. Price is reported as % of the notional amount. We follow Fabozzi et al. (2021) and exclude observations where price=100. *Brown* is an indicator which equals 1 for brown loans from six industries, defined as in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is more than 80 % above the whole sample period. We include the following controls ($\mathbf{X}_t$): face amount (logs), time-to-maturity, CLO portfolio concentration measured by the Herfindahl score of industry concentration (HHI) and CLO portfolio rating in columns (2) and (5). We use loan rating FE fixed effects (α_{rat}), CLO fixed effects (α_i) and year-month fixed effects (α_t). Standard errors are clustered at the CLO level and year-month level and are reported in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1) brown versus all	(2) brown versus green
$Brown_r$	-0.064	-0.229**
	(0.089)	(0.107)
$Attention_{t-1} \times Brown_r$	-0.899***	-0.877**
	(0.281)	(0.339)
$Face\ amount_l$	0.086***	0.056***
	(0.012)	(0.015)
$Months\ to\ maturity_l$	0.014***	0.009***
	(0.002)	(0.003)
Number of observations	317,660	131,723
Adjusted R2	0.637	0.680
Mean dependent	98.680	98.679
SD dependent	3.178	3.217
Rating FE	Yes	Yes
Industry FE	-	-
CLO FE	Yes	Yes
YM FE	Yes	Yes

Table C.5: **Net purchases and CLO leverage.** In this table we report results from estimating the following equation (equation 4.3): $Net\,Purchase_{i,r,t} = \beta\, Attention_{t-1} \times Brown_r \times Leverage_i + \cdots + \gamma\, Brown_r + \alpha_i \times \alpha_t + \alpha_r + \epsilon_{i,b,t}$. The sample covers monthly aggregates of net purchases (purchases - sale) of loans per industry and CLO i in year-month t during CLO reinvestment periods over the period 1 January 2010 - 31 May 2018. $Net\,Purchase_{i,r,t}$ is defined as in equation 4.1. *Brown* is an indicator which equals 1 for brown loans from six industries, as defined in Section 4.3. $Attention_{t-1}$ is a binary variable that equals 1 if the climate change attention index CHneg by Engle et al. (2020) is above 80 % over the whole sample period. $Leverage_i$ is the leverage ratio of CLO i. In columns (1) and (2), we compare net purchases in brown industries to all other industries. In columns (3) and (4), we restrict the sample to brown and green industries, only. We cluster standard errors on the CLO and year-month levels and report them in parentheses. *, **, *** indicate significant coefficients at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	brown versus all		brown versus green	
$Attention_{t-1} \times Brown_r$	0.051***	0.048***	0.087***	0.087***
	(0.016)	(0.017)	(0.020)	(0.022)
$Attention_{t-1} \times Brown_r \times Leverage_i$	0.005	0.013	0.001	0.000
	(0.019)	(0.019)	(0.019)	(0.020)
$Attention_{t-1} \times Leverage_i$	-0.017		-0.020	
	(0.014)		(0.018)	
$Brown_r \times Leverage_i$	-0.012	-0.014	-0.003	-0.003
	(0.009)	(0.009)	(0.011)	(0.012)
Number of Observations	293,896	293,109	136,270	134,127
Adjusted R2	0.087	0.161	0.096	0.167
Mean Dependent	0.211	0.211	0.21	0.208
SD Dependent	0.531	0.53	0.531	0.527
Industry FE	Yes	Yes	Yes	Yes
CLO FE	Yes	-	Yes	-
CLO-year-month FE	-	Yes	-	Yes
Year-month FE	Yes	-	Yes	-

Table C.6: **Variable definitions.** In this table we present variable definitions in the order they appear in the text.

Variable	Unit	Definition
AUM	Mio USD	Assets under management in million USD.
Age	years	CLO age in years.
PRI	0/1	Binary that equals 1 for CLO managers that joined the Principles of Responsible Investments, and 0 otherwise.
HHI	number	Herfindahl-Hirschman Index (HHI) of the CLO portfolio measures industry concentration.
Portfolio Rating	number	Mean of loan rating (see below) of CLO portfolio.
Price	%	Price of loan, expressed as % of notional amount.
Face amount	log USD	Notional amount of loan in logs of USD.
TTM	months	Time-to-maturity of loans in months.
Rating	(2 (1) 23)	Encoded rating according to AAA=23, AA+= 22 ... C=3, D=2
Brown	0/1	Binary that equals 1 for loans from the following industries and 0 otherwise: (i) Automobile; (ii) Cargo Transport; (iii) Chemicals & Plastics & Rubber; (iv) Containers & Packaging & Glass; (v) Mining & Steel & Iron & Non-Precious Metals; (vi) Oil & Gas; (vii) Personal Transportation; and (viii) Utilities.
Green	0/1	Binary that equals 1 for loans from the following industries and 0 otherwise: (i) Banking; (ii) Broadcasting & Entertainment; (iii) Finance; (iv) Healthcare, Education, Childcare.
High_emissions	(0.2 (0.2) 1)	Ranking based on emissions per Eikon activity. We use Refinitiv Business Classification (TRBC) on the activity level, which provides the most granular industry classification in Eikon and contains 494 activities for 85% of the borrowers in our total sample. As we only find emission data for 14% of our issuers, we aggregate per TRBC all available Scope 1, 2 and 2 emissions available in Dealscan. We rank industries into quitiles with the first quintile containing firms with the lowest carbon emissions and the fifth quintile containing firms with the highest carbon emissions. We repeat the ranking procedure using the average emission intensity, measured as the ratio of carbon emissions to market capitalization.
Net purchases_i,r,t	%	We define the % change in Net purchases of loans in industry r, month t by CLO i according to $Net\,Purchase_{i,r,t} = \frac{Notional\ Purchases_{i,r,t} - Notional\ Sales_{i,r,t}}{DealBalance_{i,t-1}}$. We multiply by 100 such that net purchases is to read in % of Deal Balance.
Notional Purchases	USD	Sum of the notional of all loans in industry r, purchased by CLO i in month t.
Notional Sales	USD	Sum of notional sales in industry r, sold by CLO i in month t.
Deal Balance	USD	Total holdings of leveraged loans by CLO i in month t. When we use the lag (t-1), we use the latest available Deal Balance lagged by one month or more.
Attention	0/1	Binary that equals 1 if the monthly Crimson Hexagon Negative News Index (CHNeg) of Engle et al. (2020) is at or above its 80th percentile, and 0 otherwise.
Net purchases_i,q,t	0/1	Similarly defined as Net purchases_i,r,t, but we calculate net purchases per CLO i, emission rank q and month t, instead of by Moody's industry r.
Affiliated	0/1	Binary that equals 1 for CLOs with a manager affiliated to a bank and 0 otherwise.
Net purchases, LMF	Δ Share	We have position weights (shares) of all loan holdings in Loan mutual funds. We sum position weights by their Eikon activity into five emission ranks q as described in Section 4.3.2. To approximate best our dependent variable *Net purchases* in the CLO analyses, we use the changes of the aggregated position weights (shares) per fund i, month t and rank q.
Attention_aggregated	0/1	Binary variable that equals 1 if we observe in total at least three monthly spikes per quarter in the attention indices CHNeg and WSJ by Engle et al. (2020) and MCCC by Ardia et al. (2023), and 0 otherwise. In a second version, it equals 1 if there are four spikes, and 0 otherwise.
Leverage	%	Leverage of CLO i in month t, or latest available month. Defined as $\frac{1}{Equity\ ratio}$.