Fear, Anger, and the Price of Risk

Andrea Guerrieri D'Amati[1]

[1]University of Warwick

This version: September 24, 2026
PRELIMINARY DRAFT: DO NOT SHARE

Abstract

This paper asks whether emotions in economic news affect the market price of risk. I combine 9.83 million UK newspaper articles with FTSE 100 returns from 2003 to 2022. I measure daily fear and anger using dictionary word shares and LLM-classified article shares, and proxy the price of risk using a model-motivated mean–variance ratio. I instrument economic-news emotions with matching emotions in non-economic and sports news, and impose a timing structure between news and returns. A one-standard-deviation increase in instrument-predicted dictionary anger raises the price-of-risk proxy by 28.3 per cent of its mean with non-economic-news instruments and by 27.9 per cent with sports-news instruments, implying lower risk appetite.

Keywords: Asset Pricing, Behavioural Finance, Sentiment Analysis, Textual Analysis, Machine Learning.
JEL Classification: G12, G41, C55, E71

1 Introduction

The price of risk is the compensation investors require for bearing more risk. As such, it can vary over time. Asset-pricing models give competing explanations for this variation, such as external habit (Campbell and Cochrane, 1999), long-run growth and volatility risks (Bansal and Yaron, 2004), rare disasters (Barro, 2009), and learning about uncertain mean returns (Brennan, 1998). However, identifying the underlying state variables driving these changes remains empirically difficult. This paper asks whether an observable part of the information environment, that is, the emotional content of economic news, predicts short-run variation in the risk-return trade-off.

I answer this question using a corpus of 9.83 million UK newspaper articles published between 2003 and 2022 and FTSE 100 total returns. I identify economic articles through dictionaries and large language models (LLMs) and construct daily indices for six emotions: fear, anger, sadness, disgust, happiness, and surprise. The analysis focuses on fear and anger; I use the full six-emotion system as an extension. To address news endogeneity with respect to market conditions, I instrument emotions in economic news with the same emotions measured in non-economic news and in sports news. The theoretical model supplies the empirical price of risk. In a partial-equilibrium portfolio-choice framework, a time-varying price of risk links the conditional mean of excess returns to their conditional variance. I estimate this object using an AR(1)–GARCH(1,1) mean–variance ratio.

The idea that markets respond to emotions is not new in the literature. Keynes (1936) describes investment as shaped by "animal spirits", or a "spontaneous urge to action rather than inaction". Graham (1949) describes a similar idea in the figure of Mr. Market, whose manic-depressive characteristics influence his willingness to invest in the stock market. Similar waves of moods from optimism to pessimism moving asset prices away from fundamentals can also be found in Shiller (2000). More recently, Farmer (2012, 2018) and Farmer and Platonov (2019) provide formal general-equilibrium foundations for these metaphors by looking at overlapping-generations models in which non-fundamental shocks to beliefs generate persistent movements in asset prices and risk premia even when fundamentals are constant and financial markets are complete.

Investors receive information about the economy through the news. However, this information may differ significantly from the original. News providers may simplify the information to make it more relatable or might choose to highlight certain messages over others. In doing so, they may change the type and quantity of information that is received by the final consumers, which can then affect their investment decisions. For example, institutions may move markets just by changing emotional, vocal, or even facial cues in their communications (Gorodnichenko et al., 2023; Kanelis and Siklos, 2026; Barry et al., 2025). Natural language processing and sentiment analysis make this interpretation measurable by classifying language as positive, negative, or neutral. However, this approach reduces emotional content to a single dimension, often referred to as "polarity". Indeed, a polarity index can tell us whether language is favourable or unfavourable, but it cannot tell us the underlying emotion. This distinction matters because emotions sharing the same polarity can produce different choices (Raghunathan and Pham, 1999; Pertl et al., 2024). Fear is associated with pessimistic risk estimates and risk-averse choices, whereas anger is associated with more optimistic estimates and risk-seeking choices through appraisals of certainty and control (Lerner and Keltner, 2001). The effect of anger on risk taking is nevertheless context-dependent (Baumann and DeSteno, 2012), and panel evidence identifies happiness, anger, and fear as correlates of within-person changes in risk attitudes (Meier, 2022). These psychological mechanisms motivate the paper's core question: do different emotions in economic news predict different movements in the compensation investors require for risk?

The answer is yes: a one-standard-deviation increase in a dictionary-measured daily index of anger-related language increases the price of risk the following day by 28.3% of its mean under non-economic-news instruments and by 27.9% of its mean under sports-news instruments. Conditional on instrument relevance, independence, exclusion, and timing, the IV coefficients identify causal effects on the price-of-risk state at $t + 1$. Because a higher (lower) mean–variance ratio is a higher (lower) price of risk, this movement implies lower (higher) risk appetite. Dictionary fear is negative, whereas LLM fear is positive under both IV strategies. Fear therefore changes sign across measurement systems.

This paper adds to three literatures. First, it brings the emotional content of public information into the empirical study of time-varying risk prices. Although existing models allow the price of risk to vary over time, the state variables driving this variation are often latent. This paper tests whether emotions in economic news provide an observable signal of these shifts. This empirical approach complements general-equilibrium models in which changes in beliefs affect asset prices and economic activity (Farmer, 2012, 2018; Farmer and Platonov, 2019). Second, it contributes to research on emotions and financial-market behaviour (Ackert et al., 2003; Breaban and Noussair, 2018; Consoli et al., 2021) by studying whether emotions embedded in broad newspaper coverage affect the market price of risk over a long daily sample. Third, it connects sentiment-based text analysis in economic and financial news (Shapiro et al., 2022) to research on discrete emotions in news content (Soroka et al., 2015) by separating fear and anger from broad sentiment and comparing dictionary counts with contextual article-level measures.

These findings suggest that financial stability authorities could monitor how their communications are reframed in the news. Emotional measures derived from news could complement market stress indicators. Authorities could also track emotional responses to announcements to review how their messages were interpreted and adjust their approach if necessary.

The paper is structured as follows: Section 2 derives the price-of-risk relation used in the empirical work; Section 3 describes the newspaper corpus, emotion measures, and financial data; Section 4 explains the construction of the risk-price proxy and the empirical specifications; Section 5 reports the results of the estimations and provides robustness checks; Section 6 concludes.

2 Model Setup

This section develops the pricing measure used in the empirical analysis. The model considers an investor who chooses consumption and exposure to a risky asset in a partial-equilibrium portfolio choice setting. The portfolio first-order condition links expected excess returns to risk through a time-varying price-of-risk measure.

2.1 Preferences

At date t, a price-taking investor chooses consumption c_t and the share π_t of investable wealth a_t to hold in a risky asset. The risky asset has gross return R_{t+1}; the remaining wealth is invested at the safe gross return R_f. Preferences are CRRA with time-varying curvature,

$$u(c_t; \rho_t) = \frac{c_t^{1-\rho_t}}{1-\rho_t}, \tag{1}$$

where ρ_t is the coefficient of relative risk aversion and follows an AR(1) process.

2.2 Return Process

Let excess returns be

$$x_{t+1} \equiv R_{t+1} - R_f. \tag{2}$$

Let $\mathcal{F}_t$ denote the information available when the investor chooses c_t and π_t. This information includes the history of realised returns and the current wealth, preference, and return-process states. The conditional mean and variance are

$$\mu_{t+1} \equiv \mathbb{E}[x_{t+1} \mid \mathcal{F}_t], \qquad h_{t+1} \equiv \mathrm{E}[x_{t+1}^2 \mid \mathcal{F}_t].$$

For brevity, $\mathbb{E}_t[\cdot]$ below denotes expectation conditional on $\mathcal{F}_t$. The investor takes these conditional moments as given when choosing the portfolio share; wealth and risk aversion affect that choice but do not determine the distribution of returns, as preference shocks are independent of return shocks. The gross portfolio return is

$$R_{p,t+1} \equiv R_f + \pi_t x_{t+1}, \tag{3}$$

so next-period wealth is

$$w_{t+1} = a_t\big(R_f + \pi_t x_{t+1}\big). \tag{4}$$

2.3 Price of Risk

The portfolio first-order condition says that the investor chooses π_t so that the expected marginal-utility-weighted excess return is zero. Appendix A derives this condition and applies a local approximation around the steady state. The resulting pricing relation is

$$\mathbb{E}_t[x_{t+1}] \approx \lambda_t \, \mathbb{E}_t[x_{t+1}^2] \tag{5}$$

where λ_t is the price of risk. This object should not be considered equivalent to the risk aversion parameter ρ_t in the utility function. Indeed, λ_t combines the curvature of continuation value, the investor's risky-asset position, and hedging terms from changes in risk aversion and return moments. A higher λ_t means that the investor requires a higher expected excess return for a given amount of conditional risk. The price of risk can therefore move because fundamentals change, because willingness to bear risk changes, or because both change. Equation (5) implies

$$\lambda_t \approx \frac{\mu_{t+1}}{h_{t+1}}. \tag{6}$$

The empirical proxy uses this mean–variance interpretation: an AR(1) specification conditions the estimated mean on the lagged excess return, while a GARCH(1,1) specification updates the conditional variance using the previous squared return innovation and the previous conditional variance.[1] The observed return history therefore provides an empirical approximation to the return-relevant information in $\mathcal{F}_t$.

3 Data

The data combine UK newspaper articles published between 2003 and 2022 with daily FTSE 100 total-return, risk-free-yield, and trading-volume series. The unit of observation in the text data is an article; the unit of observation in the regressions is a trading day. I use the article data to construct daily measures of emotions in economic, non-economic, and sports news, and I merge those measures with financial data used to measure the market price of risk.

3.1 Financial Data

Financial-market data come from the London Stock Exchange Group (LSEG) and span 1 July 2003 to 31 December 2022. The risky asset is the FTSE 100 Total Return Index, which measures the performance of the 100 most highly capitalised blue-chip companies listed on the London Stock Exchange and includes capital gains and reinvested dividends. I calculate the excess returns of the stock market as the return on the risky asset minus the two-year UK gilt yield, which proxies for the risk-free rate. I also construct the ten-year minus two-year gilt spread as a control for the slope of the yield curve.

3.2 Text Data

The newspaper corpus comes from the online editions of four major UK newspapers: The Daily Mail, The Guardian, The Independent, and Metro. These outlets cover a substantial

[1] Because h_{t+1} is the conditional variance, GARCH models it as $\mathbb{E}_t[x_{t+1}^2] = h_{t+1} + \mu_{t+1}^2$. At a daily level, μ_{t+1}^2 can be dropped as a second-order term, so $\mathbb{E}_t[x_{t+1}^2] \approx h_{t+1}$.

segment of UK print and online newspaper readership (Ofcom, 2025)[2]. The raw corpus contains 9,832,307 articles between 2003 and 2022, spanning the Global Financial Crisis, the European sovereign debt crisis, the Brexit referendum, the COVID-19 pandemic, and the subsequent inflationary episode. Table 1 reports the raw number of articles by newspaper and year. The Daily Mail contributes the largest share of articles, followed by The Guardian, The Independent, and Metro. I then retain articles that contain at least 100 words and were published during the trading week, and align the resulting sample period with that of the financial data. The filtered corpus contains 7,709,678 articles.

Table 1: Raw UK Newspaper Articles by Outlet and Year, 2003–2022

Year	Daily Mail	Guardian	Independent	Metro	Total
2003	23,142	70,334	38,312	0	131,788
2004	29,372	74,361	41,771	0	145,504
2005	25,799	74,013	43,283	2,160	145,255
2006	38,077	79,165	42,481	14,376	174,099
2007	68,394	91,029	23,089	45,036	227,548
2008	92,777	102,048	45,468	58,397	298,690
2009	89,081	91,883	61,596	56,193	298,753
2010	103,065	83,049	78,807	43,189	308,110
2011	130,081	89,246	76,891	34,785	331,003
2012	167,026	87,417	80,700	36,862	372,005
2013	181,050	89,575	78,130	48,272	397,027
2014	334,974	99,427	73,775	44,458	552,634
2015	468,409	115,887	83,010	55,881	723,187
2016	555,385	107,980	72,670	70,562	806,597
2017	549,941	83,480	71,714	79,753	784,888
2018	606,307	78,095	64,312	87,359	836,073
2019	571,000	76,905	65,533	92,184	805,622
2020	545,428	77,082	78,029	87,295	787,834
2021	597,590	75,149	144,331	79,779	896,849
2022	468,995	77,124	181,962	80,760	808,841
Total	**5,645,893**	**1,723,249**	**1,445,864**	**1,017,301**	**9,832,307**

Notes: Cells count articles collected from the online editions of the Daily Mail, The Guardian, The Independent, and Metro in each calendar year. These are raw corpus counts before any filter is applied to the corpus. The Metro series begins in 2005. The final row sums articles across 2003–2022.

3.2.1 Dictionary-based measures

The dictionary-based emotion measures use the NRC Word–Emotion Association Lexicon (EmoLex) developed by Mohammad and Turney (2010) and Mohammad and Turney (2013). I construct article-level shares for fear, anger, sadness, disgust, happiness, and surprise, defined as the fraction of words associated with each category relative to the total number

[2]Ofcom's 2025 News Consumption Survey reports that, among adults using print or online newsbrands, the Daily Mail/Mail on Sunday and Guardian/Observer were the two most widely used newsbrands, reaching 33 and 31 per cent, respectively; Metro reached 17 per cent and The Independent 7 per cent.

of words in an article. I also use the Harvard-IV4 (HIV4) sentiment lexicon as a benchmark for broad sentiment. HIV4 is a word list that classifies terms into positive, negative, and related categories. For each article, I construct a measure of polarity as $(P - N)/(P + N)$, where P and N are the article's counts of HIV4 positive and negative words. The measure ranges from -1 to 1: higher (lower) values indicate more (less) positive sentiment.

I classify articles as economic or sports news using the dictionaries in Table 2. An article is classified as economic if it contains at least one economy term. The sports rule is narrower: an article is classified as sports if it contains a competition phrase, or if it contains one of the listed venue names together with at least one sports-context term within 15 tokens. Using these labels, I average the article-level emotion shares by trading day to construct separate dictionary emotion indices for economic news, non-economic news, and sports news. The economic-news indices are the regressors of interest. The non-economic and sports-news indices supply the excluded instruments in the dictionary IV specifications.

Table 2: Dictionary Rules Used to Classify Economic and Sports News

Dictionary component	Terms and phrases
Economy terms	GDP; inflation; interest rate; interest rates; monetary policy; central bank; BoE; bank of england; Federal Reserve; Fed; ECB.
Sports competitions	Premier League; Champions League; FA Cup; League Cup; World Cup; European Championship; Euros; Test match; Ashes; Wimbledon; Grand Prix; Formula 1; Six Nations; Premiership.
Sports venues	Old Trafford; Anfield; Wembley.
Sports context	football; cricket; rugby; tennis; match; game; fixture; final; semi-final; semifinal; goal; goals; score; scored; manager; coach; club; team; league; season; stadium; fans; captain; striker; midfielder; defender; tournament; cup; championship.

Notes: The table defines the dictionary rules used to classify newspaper articles before constructing the dictionary emotion measures. Matching is case-insensitive. An article is classified as economic news when it contains any economy term. It is classified as sports news when it contains any sports-competition phrase, or when a sports-venue phrase appears within 15 tokens of a sports-context term.

Table 3 compares dictionary emotion shares in economic and non-economic news and, separately, in sports and non-sports news. Economic articles contain more fear, anger, and sadness than non-economic articles, but less happiness and surprise. Sports articles contain less fear, anger, sadness, and disgust than non-sports articles; their happiness share is only slightly lower, while surprise is the same to three decimal places. The greater concentration of negative emotional language in economic news is consistent with evidence that negative news elicits stronger and more sustained reactions (Soroka and McAdams, 2015) and that negative headline words can increase online news consumption (Robertson et al., 2023).

Table 3: Mean Dictionary Emotion Word Shares by News Type

	Economic news	Non-economic news	Sports news	Non-sports news
Fear	0.025 (0.017)	0.022 (0.017)	0.017 (0.013)	0.023 (0.018)
Anger	0.018 (0.013)	0.016 (0.013)	0.013 (0.011)	0.016 (0.014)
Happiness	0.015 (0.010)	0.018 (0.012)	0.017 (0.010)	0.018 (0.012)
Sadness	0.018 (0.011)	0.016 (0.012)	0.014 (0.009)	0.017 (0.012)
Disgust	0.008 (0.007)	0.008 (0.008)	0.006 (0.005)	0.008 (0.008)
Surprise	0.010 (0.006)	0.011 (0.007)	0.011 (0.007)	0.011 (0.007)
N	1,075,394	6,634,284	1,180,899	6,528,779

Notes: Cells report the mean article-level dictionary emotion word share, with its standard deviation in parentheses. Economic and non-economic news are articles for which the dictionary economy indicator equals one and zero, respectively. Sports and non-sports news are articles for which the sports indicator equals one and zero, respectively. The sports classification is selected separately and may overlap the economic or non-economic classification. A share is the number of words associated with an NRC emotion category divided by the article's word count. The sample keeps newspaper articles with at least 100 words, published Monday–Friday from 1 July 2003 to 31 December 2022. N is the number of articles in each column.

Figure 1 compares the six dictionary emotions with bins constructed from the HIV4 polarity measure. The figure highlights the need to use granular emotions rather than a single sentiment index. Fear, anger, and sadness are more common in negative-polarity bins, but the mapping is not one-for-one, and emotions on the same side of the polarity scale differ sharply in frequency.

Figure 2 shows the monthly evolution of dictionary emotion shares. Emotional content varies substantially over time. Fear and sadness rise around the Global Financial Crisis, the Brexit referendum, the COVID-19 pandemic, and the subsequent inflationary episode. Happiness is more subdued in those periods, while surprise rises around some sharp policy or macroeconomic announcements.

3.2.2 LLM-based measures

As a complementary measurement exercise, I develop measures based on large language models (LLMs). LLM metrics provide a contextual comparison with dictionary-based counts. A dictionary classifies a word as emotional or economic whenever it appears in a predefined list. The LLM approach, on the other hand, treats the article itself as the object to be classified. Contextual methods are capable of capturing semantic distinctions that word lists miss. In central bank communications, Baumgärtner and Zahner (2025) demonstrate that domain-specific embeddings outperform dictionaries and generic embeddings in predicting monetary policy shocks. Their corpus and model differ from those used in this study, but their results justify treating lexical counts and contextual measures as distinct representations of communication.

The construction proceeds in two stages. First, a GPT model labels a random subsample

Figure 1: Composition of Dictionary Emotion Words by HIV4 Polarity Bin

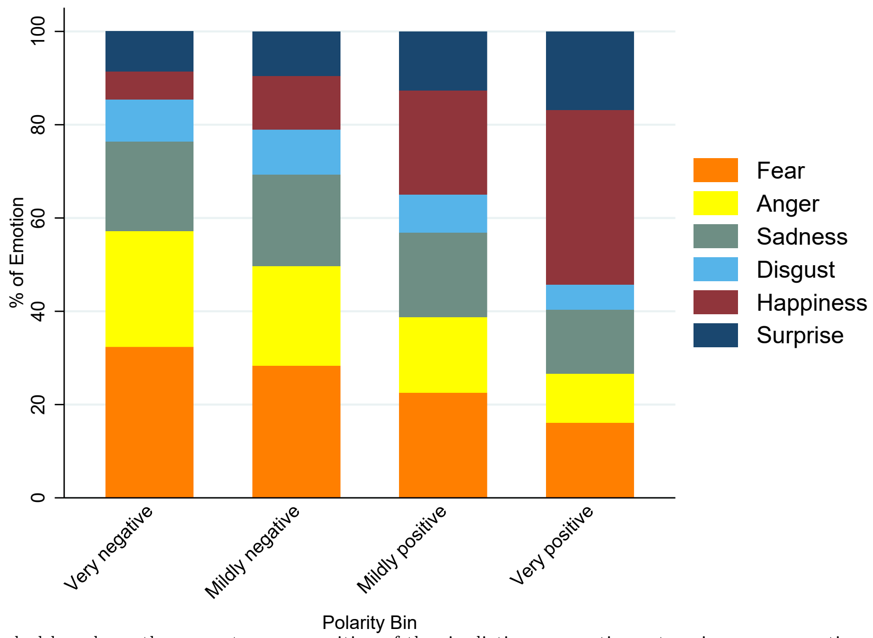

Notes: Each stacked bar shows the percentage composition of the six dictionary emotion categories among emotion-associated words in UK newspaper articles assigned to the indicated HIV4 polarity bin; each bar sums to 100. HIV4 polarity is $(P - N)/(P + N)$, where P and N are the article's counts of HIV4 positive and negative words; bins to the right indicate more favourable language. Dictionary emotions are identified with the NRC Word–Emotion Association Lexicon.

Figure 2: Monthly Average Dictionary Emotion Word Shares in UK News, 2003–2022

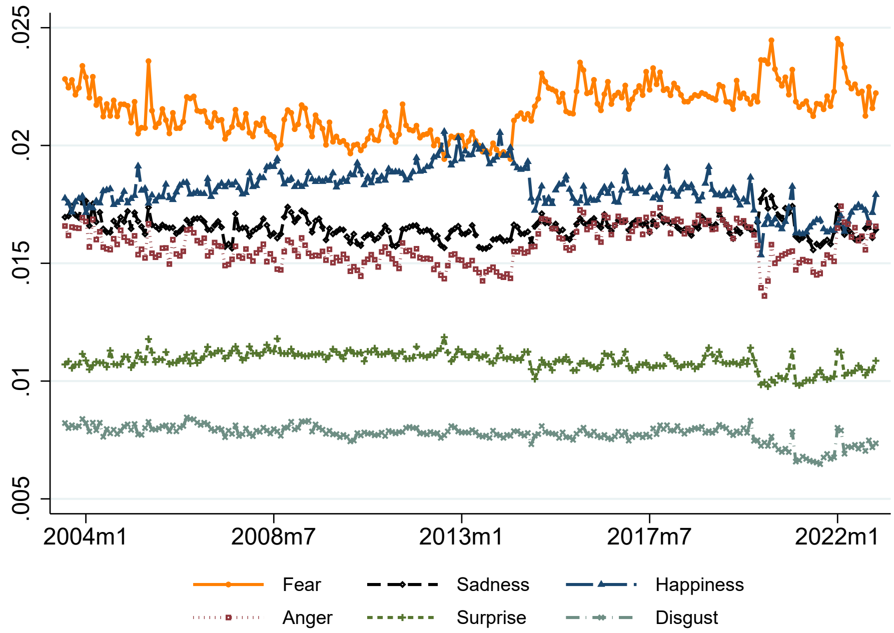

Notes: Each line is the monthly mean of an article-level dictionary emotion share in the UK newspaper corpus. For each article and emotion, the share equals the number of words associated with that NRC emotion category divided by the article's total word count. Values are proportions, so 0.02 denotes two per cent of words.

of articles. For each article, it records whether the article concerns the economy and assigns percentage shares to the six emotion categories, plus a "neutral" category to describe residual/non-emotional content. The model is given the following prompt using definitions of emotions from the basic-emotion tradition of Ekman (1971) and the appraisal-based characterisation from Lerner and Keltner (2001):

> *"Given the news article below, answer the following questions: 1. Is the article about economic conditions in the UK? Reply with 'Yes' or 'No'. 2. What percentage share of the article's content is associated with each emotion? The percentage for each emotion should reflect its share in the article and be between 0 and 100; the percentages cannot sum to more than 100. Use these definitions for emotions: - Anger: An emotional response to perceived injustice or goal obstruction, characterized by a sense of certainty and individual control. - Disgust: A feeling of revulsion toward something offensive, contaminating, or violating moral/social norms, often leading to rejection or avoidance. - Fear: A reaction to perceived threat or danger, with a sense of uncertainty and low control, motivating avoidance of harm. - Happiness: A state of joy, satisfaction, or positive affect arising from goal attainment or favorable outcomes, enhancing openness and social approach. - Sadness: An emotional response to loss, disappointment, or adverse outcomes, often leading to withdrawal and reflection to realign with goals. - Surprise: A brief emotional response to unexpected events, leading to heightened attention and reevaluation of the situation. - Neutral: None of the above emotions. Output Format: - Always output exactly two lines: Line 1: Yes or No Line 2: The percentage share of each emotion in the article. Do not include reasoning or explanations. Here is the article: "*

Second, I train ModernBERT models on those GPT-labelled examples. ModernBERT is a text prediction model: it converts an article into numerical representations of its words and context, then learns how those representations correspond to GPT's labels. After training, I apply the ModernBERT models to articles that GPT has not labelled directly. This step is necessary because the full corpus contains millions of articles, so directly labelling each article with GPT would be impractical and costly.

I train two separate ModernBERT models because the tasks differ: the economy classifier predicts a yes/no outcome, while the emotion model predicts continuous category shares. The GPT-labelled sample is split 75/25 into training and held-out validation sets. The training set is shown to the model while it learns, whereas the held-out set is not used to fit the model. Performance on the latter set therefore indicates how well the fitted model reproduces GPT labels for previously unseen articles.

I select the economy classifier using its F1 score on the held-out set. F1 is the harmonic mean of precision, P, and recall, R: $F1 = 2PR/(P + R)$. It ranges from zero to one and is high only when both precision and recall are high. Precision is the share of articles predicted to be economic that GPT also labelled as economic. Recall is the share of GPT-labelled economic articles that the classifier successfully identifies. On the held-out sample, the economy classifier has an F1 score of 0.904, precision of 0.908, and recall of 0.899. Put differently, about 91 per cent of articles that the classifier identifies as economic also have a GPT economic label, and it recovers about 90 per cent of all articles that GPT labels as economic. The F1 score summarises the balance between these two properties rather than the overall percentage of articles classified correctly.

I select the emotion model using held-out mean squared error, which measures how close its predicted emotion shares are to the GPT-labelled shares. Across the seven emotion shares, the held-out mean squared error is 0.0129 and the mean absolute error is 0.0795. The mean absolute error corresponds to an average absolute difference of approximately eight percentage points per emotion category.

The final daily LLM series are simple averages of article-level predictions. For a given trading day, I average the predicted emotion share across articles published that day. These daily averages are then standardised before entering the regressions. Figure 3 compares the LLM-predicted emotion shares with polarity bins constructed from the dictionary-based HIV4 sentiment measure. Negative LLM emotions remain prominent as dictionary polarity becomes more positive, and predicted fear is highest in the most positive polarity bin. This does not imply that positive language causes fear. It shows that word-count polarity and contextual article-level classification can assign different meanings to the same text: an article may use positive words while still discussing threats, uncertainty, or fearful reactions.

Figure 3: Composition of LLM-Predicted Article Emotions by HIV4 Polarity Bin

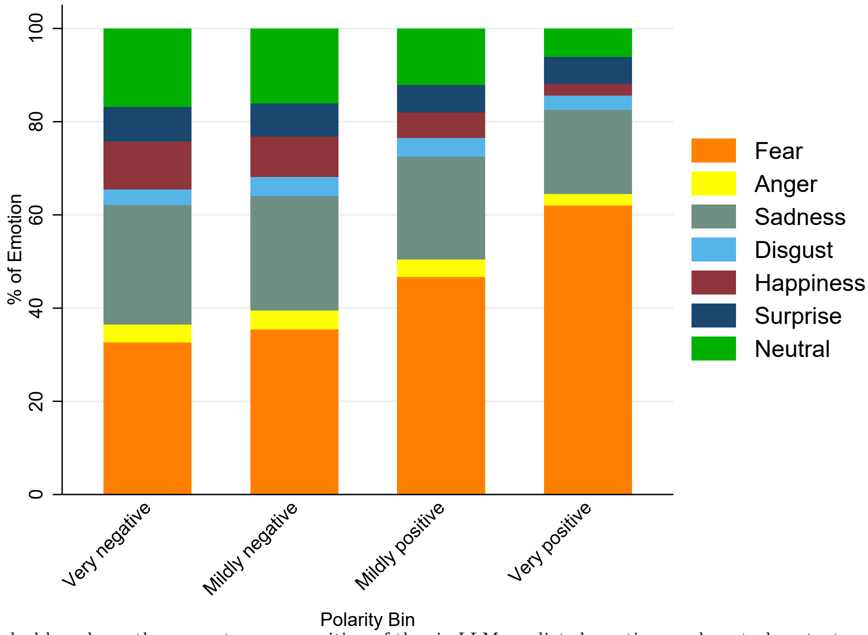

Notes: Each stacked bar shows the percentage composition of the six LLM-predicted emotions and neutral content among UK newspaper articles assigned to the indicated HIV4 polarity bin; each bar sums to 100. HIV4 polarity is $(P - N)/(P + N)$, where P and N are the article's counts of HIV4 positive and negative words; bins to the right indicate more favourable language. The emotion shares are article-level ModernBERT predictions from models trained on GPT-labelled articles.

Figure 4 shows the monthly evolution of the LLM-based emotion shares. Fear and sadness account for the largest shares throughout the sample, while anger, happiness, surprise, and disgust are lower and move in a narrower range. Fear falls around the Global Financial Crisis period, rises again in parts of the post-crisis sample, and becomes more volatile around the COVID-19 and inflationary periods. Sadness is also persistent and elevated in several later episodes.

Table 4 reports the corresponding descriptive statistics for the LLM measures. Because the LLM assigns contextual category shares, their levels are not directly comparable to dictionary word-count shares. Fear and sadness have the largest average shares. Economic articles have less predicted fear but more sadness and neutral content than non-economic articles. Sports articles have more predicted fear than non-sports articles, while anger, happiness, sadness, disgust, surprise, and neutral content are all lower. For this comparison, the emotion shares are generated by the LLM, while the sports and non-sports labels come from the dictionary-based sports indicator.

Figure 4: Monthly Average LLM-Predicted Emotion Shares in UK News, 2003–2022

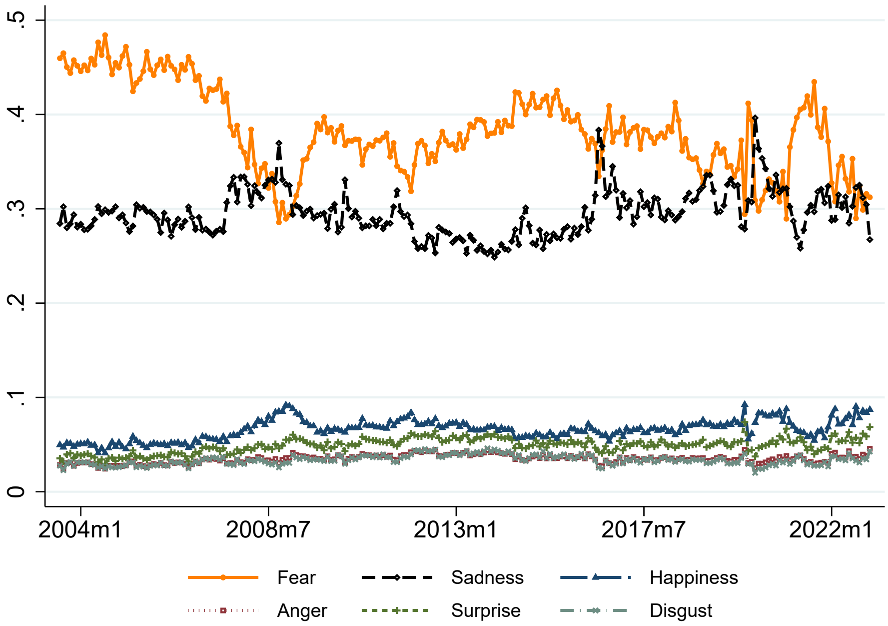

Notes: Each line is the monthly mean of an article-level LLM-predicted emotion share in the UK newspaper corpus. ModernBERT predicts the share of each article associated with an emotion after training on GPT-labelled articles. Values are proportions, so 0.40 denotes 40 per cent of predicted article content. Neutral content is omitted from the figure.

Table 4: Mean LLM-Predicted Emotion Shares by News Type

	Economic news	Non-economic news	Sports news	Non-sports news
Fear	0.384 (0.235)	0.459 (0.293)	0.531 (0.285)	0.440 (0.288)
Anger	0.033 (0.038)	0.036 (0.041)	0.032 (0.042)	0.037 (0.040)
Happiness	0.064 (0.066)	0.059 (0.066)	0.043 (0.058)	0.062 (0.067)
Sadness	0.303 (0.174)	0.220 (0.147)	0.206 (0.131)	0.230 (0.154)
Disgust	0.032 (0.045)	0.040 (0.055)	0.037 (0.055)	0.039 (0.054)
Surprise	0.047 (0.062)	0.062 (0.097)	0.051 (0.091)	0.063 (0.096)
Neutral	0.136 (0.127)	0.123 (0.121)	0.100 (0.119)	0.129 (0.122)
N	540,113	7,169,565	1,180,899	6,528,779

Notes: Cells report the mean article-level LLM-predicted emotion share, with standard deviation in parentheses. Economic and non-economic news are articles for which the ModernBERT economy classifier predicts yes and no, respectively. The sports and non-sports labels come from the dictionary-based sports indicator. The sports classification is selected separately and may overlap the economic or non-economic classification. Shares range from zero to one. The sample keeps newspaper articles with at least 100 words, published Monday–Friday from 1 July 2003 to 31 December 2022. N is the number of articles in each column.

Because LLM measures are jointly estimated as emotional scores, the non-neutral emotion scores are highly collinear when included together in the same regression. This makes it difficult to attribute variations in the price of risk to a single LLM emotion. I therefore use the LLM measures primarily as a contextual comparison between two emotions and consider the full LLM specification as an extension.

3.2.3 Comparing the measures

Although the dictionary measure is replicable, it can be affected by topic composition, negation, institutional vocabulary and quoted speech. The LLM measure uses article context but relies on GPT-generated training labels and the quality of ModernBERT predictions.

If the dictionary measure is treated as a benchmark, the LLM measure should be understood as a model-dependent approximation to latent emotion. Indeed, its predictions depend on the prompt used, the examples supplied to the model, the version of the model used, the training procedure for the ModernBERT architecture, and fine-tuning choices. Changes in these components could alter the estimated emotion shares even when the underlying article does not change. Differences in the emotion measures may thus reflect different dimensions of emotional content, measurement error, or a combination of both. This sensitivity has implications for reproducibility and interpretation, and illustrates a broader issue in using LLM-generated data in empirical work.

Table 5 shows the standard Pearson correlations between daily emotion measures and the HIV4 sentiment benchmarks. The first column compares the measures derived from the dictionary with those obtained via LLM for the same emotion. Fear shows a negative correlation between the two systems, while anger is close to zero. Sadness and disgust show the strongest agreement, but their correlations are only about 0.37. The two methods therefore measure different features of emotional language.

The remaining columns compare each set of emotions with the negativity and polarity of HIV4. The dictionary metrics follow a polarity trend for fear, anger, sadness, and disgust, while the LLM correlations are more heterogeneous, particularly for fear and happiness. The differences between the emotional categories may reflect the fact that dictionary counts, LLM contextual labels, and general sentiment scores measure related but different characteristics of news language.

Appendix Table B.1 compares the average percentage of each emotion between economic and non-economic articles, as well as between sports and non-sports articles. All estimated differences are statistically significant at the 1 percent level, which is not surprising given the large sample size. Their direction is more significant. According to the dictionary-based measure, economic articles contain more fear, anger, and sadness than non-economic articles. According to the LLM measure, articles classified as economic have lower predicted percentages of fear and anger but a higher predicted percentage of sadness. Sports articles contain less fear and anger, according to the dictionary-based measure, than non-sports articles. The LLM measure also assigns sports articles lower anger scores and lower percentages for most other emotions, but higher fear scores.

Table 5: Correlations Between Dictionary, LLM, and HIV4 Text Measures

		Dictionary–HIV4		LLM–HIV4	
Emotion	Dict.–LLM	Negative	Polarity	Negative	Polarity
Fear	-0.417	0.234	-0.605	-0.475	0.538
Anger	-0.059	0.102	-0.649	0.349	-0.275
Sadness	0.370	0.232	-0.510	0.022	-0.146
Happiness	0.083	-0.045	0.378	0.464	-0.499
Surprise	0.291	0.034	0.104	0.440	-0.295
Disgust	0.374	0.117	-0.428	0.246	-0.170

Notes: Pearson correlations use daily observations from 1 July 2003 to 31 December 2022. Each row identifies an emotion. The Dict.–LLM column correlates the daily standardised NRC dictionary word-share measure with the daily standardised LLM-predicted article-share measure for that emotion. The remaining columns correlate each emotion measure with the HIV4 negative-word share and HIV4 polarity. Polarity is $(P-N)/(P+N)$, where P and N are the article's counts of HIV4 positive and negative words. The omitted same-series correlation equals one by definition.

4 Empirical Setup

The empirical analysis asks whether emotions in economic news are related to short-run movements in the market price of risk. The text analysis produces six daily dictionary-based emotion indices from economic news: fear, anger, sadness, disgust, happiness, and surprise. The two-emotion model focuses on fear and anger because they are both negative emotions but are conceptually different categories and keep the regressions parsimonious. The six-emotion model is an extension that adds the other individual emotions.

The main specification is organised around three threats to inference and identification. First, the outcome variable and emotion regressors are both constructed rather than directly observed. Conventional standard errors would not account for the uncertainty in the generated variables nor the serial dependence in returns and news. I therefore use a moving-block bootstrap that preserves short-run dependence and, within every draw, reconstructs the proxy, returns, and controls before re-estimating the OLS and IV models. The daily emotion series are resampled in the same blocks, although the dictionary and LLM measurement systems are not re-estimated, so inference remains conditional on those measurement procedures. Second, publication and trading overlap within the day, creating temporal ambiguity and the possibility of same-day feedback from markets to reporting. The chosen specification therefore relates emotions, excluded instruments, the return control, and the remaining controls at t to the price-of-risk state at $t+1$, represented by the conditional mean–variance ratio for the return at $t+2$. This ordering separates the news measure from the contemporaneous price-of-risk state. Third, economic-news emotions may be endogenous because common macroeconomic or financial shocks can move both reporting and the price of risk. I address this concern with two IV strategies: IV–NE uses matching emotions in non-economic news, and IV–Sports uses matching emotions in sports news.

Let x_t denote the daily FTSE 100 excess return over the two-year gilt yield. Equation (5) implies that the price-of-risk state at $t+1$ is approximated by the conditional mean–variance ratio for the return at $t+2$:

$$\hat{\lambda}_{t+1} = \frac{\hat{\mu}_{t+2|t+1}}{\hat{h}_{t+2|t+1}}.$$

This distinction between the state date and the conditional-return date is important. The AR(1)–GARCH(1,1) fitted moments on the return date $t+2$ condition on the information available at $t+1$; their ratio therefore represents λ_{t+1}, not λ_{t+2}. A higher (lower) value indicates a higher (lower) price of risk and a lower (higher) risk appetite, holding investment opportunities fixed.

Let $m \in \{D, L\}$ index dictionary and LLM emotion measures. The common regression specification is

$$\hat{\lambda}_{t+1} = \alpha_m + \boldsymbol{E}'_{m,t}\boldsymbol{\beta}_m + \delta_t + X'_{m,t}\boldsymbol{\gamma}_m + u_{m,t}, \tag{7}$$

where $\boldsymbol{E}_{m,t}$ contains fear and anger in the main model and all six emotions in the robustness specification. The term δ_t captures day-of-week fixed effects. The vector $X_{m,t}$ includes the excess return at t, log article count, log trading volume, inflation- and monetary-policy-announcement dummies, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty from Baker et al. (2016), and HIV4 polarity. Table 6 describes the variables used in the regressions.

Table 6: Description of variables

Variable	Role	Construction
$\hat{\lambda}_{t+1}$	GARCH price-of-risk outcome	Price-of-risk state at $t+1$, represented by the AR(1)–GARCH(1,1) conditional mean–variance ratio $\hat{\mu}_{t+2\|t+1}/\hat{h}_{t+2\|t+1}$ for the return at $t+2$.
$\boldsymbol{E}_{D,t}$	Dictionary emotions	Daily dictionary-based emotion indices in economic news. The main model includes fear and anger; the robustness specification adds sadness, disgust, happiness, and surprise.
$\boldsymbol{E}_{L,t}$	LLM emotions	Daily ModernBERT-predicted emotion shares in economic news, trained on GPT-labelled articles.
$\boldsymbol{Z}_{m,t}^{NE}$	Non-economic instruments	Matching emotions measured in non-economic news, using the same measurement family m as the endogenous economic-news emotions.
$\boldsymbol{Z}_{m,t}^{SP}$	Sports-news instruments	Matching sports-news dictionary emotions. In LLM specifications, the sports instrument set uses sports-news dictionary emotions.
$X_{m,t}$	Controls	Excess return, log article count, log trading volume, Bank of England and CPI announcement dummies, ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, and HIV4 polarity.
δ_t	Calendar controls	Day-of-week fixed effects.

Notes: The table defines the dependent variable, endogenous emotion measures, excluded instruments, controls, and fixed effects in the daily regressions in Equation (7). The sample runs from 1 July 2003 to 31 December 2022. HIV4 polarity is $(P-N)/(P+N)$, where P and N are the article's counts of HIV4 positive and negative words. "Family-matched" means that dictionary regressions use polarity for dictionary-classified economic articles, while LLM regressions use polarity for LLM-classified economic articles. Newspaper variables come from the four-outlet article corpus; dictionary variables use the NRC emotion and HIV4 sentiment lexicons; LLM variables use ModernBERT predictions trained on GPT-labelled articles. FTSE 100 total-return prices, trading volume, and gilt yields come from LSEG. Bank of England and CPI announcement indicators come from the event calendar merged to the trading-day data. The UK Economic Policy Uncertainty index comes from Baker et al. (2016). All regression emotion variables are standardised.

These controls absorb observable variation in market conditions, as well as the information environment, the yield curve, and broad sentiment. Conditional on them, variation in $\hat{\lambda}_{t+1}$ coming from emotions may reflect the effect of the preference component ρ from the model, although the emotion measures may still covary with unobserved changes in the investment opportunity set, so λ cannot be considered a direct measure of ρ. However, changes in λ may serve as an indirect proxy for changes in ρ since greater risk aversion raises the compensation investors require for risk.

The emotional variables have been standardised, meaning that the coefficients measure the change in risk pricing associated with an increase of one standard deviation in the relevant emotional index. In the IV specifications, the interpretation applies to the part of economic-news emotions predicted by the excluded instruments. A positive (negative) coefficient implies a higher (lower) price of risk and lower (higher) risk appetite. I refer to coefficients as a percentage of the estimation-sample mean of the outcome to make them more interpretable.

The six-emotion robustness specification expands $\boldsymbol{E}_{m,t}$ to include sadness, disgust, happiness, and surprise. It assesses how the estimated emotion coefficients change under a richer decomposition of emotional content.

5 Results

Table 7 reports the daily estimates, with dictionary measures in Panel A and LLM measures in Panel B. Dictionary anger rises from 0.51 in OLS (6.9 per cent of the mean outcome) to 2.08 under IV–NE (28.3 per cent) and 2.05 under IV–Sports (27.9 per cent). The similarity of the two IV magnitudes shows that the anger result is stable across instrument sets. Dictionary fear is negative: the IV–NE coefficient is -1.55 (-21.0 per cent of the mean), while the imprecise IV–Sports estimate is -0.83 (-11.3 per cent). In Panel B, LLM fear rises from 1.08 in OLS (14.6 per cent) to 3.20 under IV–NE (43.5 per cent) and 4.01 under IV–Sports (54.5 per cent); only the IV–NE estimate excludes zero at the 95 per cent level. LLM anger is 1.61 under IV–NE (21.8 per cent) and 2.38 under IV–Sports (32.4 per cent); again, only the IV–NE estimate is statistically precise. Thus, positive dictionary anger is stable across instrument sets, LLM IV–NE also supports a positive anger effect, and the sign of fear remains measurement-dependent. This divergence may reflect differences in the emotional content captured by the two measures, as well as measurement error in the LLM-generated shares. Using AI/ML-generated regressors can introduce coefficient bias and invalidate inference (Battaglia et al., 2025). The bootstrap treats the fitted measurement models as fixed and does not correct for this potential measurement-error bias, which could explain the sign divergence.

The regression table also reports the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions for each IV specification. In the main two-emotion model, the minimum ranges from 93.3 to 372.0 for the dictionary designs and equals 245.8 for LLM IV–NE. The value for LLM IV–Sports is 12.5, indicating a comparatively weaker first stage.

Table 7: Price-of-Risk Regressions, Two Emotions

	(1) OLS	(2) IV–NE	(3) IV–Sports
Panel A: Dictionary emotions			
Fear	-0.716***	-1.547***	-0.833
	[-1.244, -0.300]	[-2.681, -0.625]	[-2.076, 0.264]
Anger	0.510***	2.080***	2.054**
	[0.141, 0.997]	[0.592, 4.114]	[0.340, 4.353]
Panel B: LLM emotions			
Fear	1.076***	3.197***	4.005*
	[0.450, 1.784]	[1.400, 5.566]	[-0.168, 10.10]
Anger	0.498***	1.606***	2.383
	[0.196, 0.858]	[0.563, 2.989]	[-2.714, 8.500]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	372.03	93.34
Min. SW F, LLM	–	245.79	12.47
Mean of proxy	7.354	7.354	7.354
Observations	4,927	4,927	4,927
Bootstrap replications	5,000	5,000	5,000

Notes: The table reports daily OLS and IV regressions for the estimation sample, 2 July 2003–30 December 2022. The dependent variable is an AR(1)–GARCH(1,1) conditional-mean-to-variance ratio. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. The dependent price-of-risk state is dated $t+1$ and is represented by the conditional mean–variance ratio for the return at $t+2$; emotions, excluded instruments, and the other controls are dated t. Emotion indices are standardised. Coefficients are estimated using OLS and IV. A positive coefficient indicates a higher price of risk and lower risk appetite. NE denotes non-economic news. IV–NE instruments economic-news emotions with matching non-economic-news emotions. IV–Sports instruments economic-news dictionary emotions with sports-news dictionary emotions; the LLM panel uses sports-news dictionary instruments. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, family-matched HIV4 polarity, and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. HIV4 polarity is $(P-N)/(P+N)$, where P and N are the article's counts of HIV4 positive and negative words. Each moving-block draw reconstructs the generated proxy and lagged controls. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, **, and * denote exclusion of zero from the 99, 95, and 90 per cent percentile intervals, respectively.

5.1 Robustness

I test the results from Section 5 against a series of alternative data, measurements, and specifications. First, I look at a broader decomposition of emotional content using all six emotions. The results are qualitatively similar for dictionary fear and anger, whereas no six-emotion LLM coefficient is statistically precise. Their very large bootstrap intervals might be due to high collinearity among the six predicted emotion shares. Appendix Table B.3 reports the full robustness specification.

Second, as the dictionary and LLM classifiers select different sets of economic articles, I reconstruct both emotion measures on a common article dataset. The union includes every article classified as economic by either system. On this common universe, dictionary anger remains positive under both instrument sets, while LLM anger remains positive under IV–NE. Dictionary fear remains negative under IV–NE, and LLM fear is positive under both instrument sets. Appendix Table B.4 reports the full two-emotion estimates.

Third, the price-of-risk proxy may be persistent and thus affect estimation. Therefore, I estimate a separate specification that adds a lag of the proxy. Relative to the main estimates, controlling for persistence reduces dictionary anger by 66–69 per cent and LLM fear by 54–64 per cent. Dictionary anger and LLM fear retain their signs and significance under IV–NE, while their IV–Sports intervals include zero. Appendix Table B.5 reports the full estimates.

Fourth, the instruments might respond to prior market conditions. To test for this, I estimate eight feedback regressions using fear and anger in non-economic and sports news at $t + 1$ on returns at t, conditional on the same controls as the main specifications. Seven of the eight coefficients are not significant, with the exception of LLM non-economic anger. However, the latter does not survive a multiple-testing correction, thus proving the instruments' temporal independence from past returns. Appendix Table B.6 provides the results of the estimation.

Fifth, the full-sample estimation of the GARCH parameter may introduce look-ahead bias by allowing later returns to influence the proxy at earlier dates. I therefore reconstruct the proxy using an expanding window, and each monving-block bootstrap repeates this construction within each draw[3]. Dictionary anger remains positive under IV–NE, while the IV–Sports estimate mantain the sign but is imprecise. LLM fear remains positive under

[3] I estimate the model using a 1,008-day initial window, which is approximately 4 trading years. The boostrap is resampled a lower amount of times to keep the estimation feasible.

IV–NE, and dictionary fear remains negative under the same instrument set. The LLM IV–Sports estimates are imprecise. These estimates retain the qualitative patterns of the main specification, although the training window reduces the sample to 3,921 observations, so the comparison also reflects a change in sample coverage. Appendix Table B.7 reports the full estimates.

6 Conclusions

This paper examines whether emotions in economic news predict the market price of risk. It combines 9.83 million newspaper articles with daily FTSE 100 returns from 2003 to 2022 and measures fear and anger using dictionary and LLM methods. The main regressions relate emotions and controls at t to the price of risk at $t+1$. I instrument economic-news emotions with matching emotions in non-economic and sports news.

The most stable result is for dictionary anger. A one-standard-deviation increase in instrument-predicted anger raises the price-of-risk proxy by 28.3% of its mean under non-economic-news instruments and by 27.9% under sports-news instruments, implying lower risk appetite. Fear is measurement-dependent: dictionary fear is negative, whereas LLM fear is positive. This sign reversal, together with the limited correlation between the two measurements, shows that dictionary counts and contextual predictions capture different features of emotional language.

The robustness checks narrow the scope of these findings. The anger and fear results persist on an alternative economic news classifier. Dictionary anger also remains positive under both instrument sets in the six-emotion model, with marginal precision under IV–Sports. Controlling for the price-of-risk state at t attenuates the emotion coefficients: IV–NE continues to support positive dictionary anger and LLM fear, while the IV–Sports emotion estimates become imprecise. The evidence therefore shows that emotional content in public economic news predicts the next day's inferred price of risk, but the strength and interpretation of that relationship depend on the emotion measure, instrument set, and identifying assumptions.

Finally, these findings suggest that news-based emotion measures could play a role in monitoring financial stability. Authorities could use these measures alongside market stress indicators to track the emotional tone of coverage following a major announcement, allowing them to monitor how their messages are interpreted.

References

Ackert, L. F., Church, B. K. and Deaves, R. (2003), 'Emotion and financial markets', *Federal Reserve Bank of Atlanta Economic Review* **88**(2).

Baker, S. R., Bloom, N. and Davis, S. J. (2016), 'Measuring economic policy uncertainty', *The Quarterly Journal of Economics* **131**(4), 1593–1636.

Bansal, R. and Yaron, A. (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *The Journal of Finance* **59**(4), 1481–1509.

Barro, R. J. (2009), 'Rare disasters, asset prices, and welfare costs', *American Economic Review* **99**(1), 243–264.

Barry, M.-L., Bruns, B. J., Kandemir, S., Klose, J., Smirnov, V. and Tillmann, P. (2025), The emotions of monetary policy, Working paper, SSRN.

Battaglia, L., Christensen, T., Hansen, S. and Sacher, S. (2025), 'Inference for regression with variables generated by AI or machine learning', arXiv preprint arXiv:2402.15585. Revised April 2025.
URL: *https://arxiv.org/abs/2402.15585v5*

Baumann, J. and DeSteno, D. (2012), 'Context explains divergent effects of anger on risk taking.', *Emotion* **12**(6), 1196.

Baumgärtner, M. and Zahner, J. (2025), 'Whatever it takes to understand a central banker—embedding their words using neural networks', *Journal of International Economics* **157**, 104101.

Breaban, A. and Noussair, C. N. (2018), 'Emotional state and market behavior', *Review of Finance* **22**(1), 279–309.

Brennan, M. J. (1998), 'The role of learning in dynamic portfolio decisions', *Review of Finance* **1**(3), 295–306.

Campbell, J. Y. and Cochrane, J. H. (1999), 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**(2), 205–251.

Consoli, S., Pezzoli, L. T. and Tosetti, E. (2021), 'Emotions in macroeconomic news and their impact on the european bond market', *Journal of International Money and Finance* **118**, 102472.

Ekman, P. (1971), Universals and cultural differences in facial expressions of emotion., *in* 'Nebraska symposium on motivation', University of Nebraska Press.

Farmer, R. E. (2012), 'Confidence, crashes and animal spirits', *The Economic Journal* **122**(559), 155–172.

Farmer, R. E. (2018), 'Pricing assets in a perpetual youth model', *Review of Economic Dynamics* **30**, 106–124.

Farmer, R. E. and Platonov, K. (2019), 'Animal spirits in a monetary model', *European Economic Review* **115**, 60–77.

Gorodnichenko, Y., Pham, T. and Talavera, O. (2023), 'The voice of monetary policy', *American Economic Review* **113**(2), 548–584.

Graham, B. (1949), *The Intelligent Investor: A Book of Practical Counsel*, Harper & Brothers, New York.

Kanelis, D. and Siklos, P. L. (2026), 'Emotion in euro area monetary policy communication and bond yields: The draghi era', *Journal of Economic Behavior & Organization* **245**, 107525.

Keynes, J. M. (1936), *The General Theory of Employment, Interest and Money*, Macmillan, London.

Lerner, J. S. and Keltner, D. (2001), 'Fear, anger, and risk.', *Journal of Personality and Social Psychology* **81**(1), 146.

Meier, A. N. (2022), 'Emotions and risk attitudes', *American Economic Journal: Applied Economics* **14**(3), 527–558.

Mohammad, S. M. and Turney, P. D. (2013), 'Crowdsourcing a word–emotion association lexicon', *Computational Intelligence* **29**(3), 436–465.

Mohammad, S. and Turney, P. (2010), Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon, *in* 'Proceedings of the NAACL HLT 2010 workshop on computational approaches to analysis and generation of emotion in text', pp. 26–34.

Ofcom (2025), News consumption in the UK: 2025 research findings, Report, Ofcom. Published 21 July 2025.

Pertl, S. M., Srirangarajan, T. and Urminsky, O. (2024), 'A multinational analysis of how emotions relate to economic decisions regarding time or risk', *Nature Human Behaviour* **8**(11), 2139–2155.

Raghunathan, R. and Pham, M. T. (1999), 'All negative moods are not equal: Motivational influences of anxiety and sadness on decision making', *Organizational Behavior and Human Decision Processes* **79**(1), 56–77.

Robertson, C. E., Pröllochs, N., Schwarzenegger, K., Pärnamets, P., Van Bavel, J. J. and Feuerriegel, S. (2023), 'Negativity drives online news consumption', *Nature Human Behaviour* **7**(5), 812–822.

Shapiro, A. H., Sudhof, M. and Wilson, D. J. (2022), 'Measuring news sentiment', *Journal of Econometrics* **228**(2), 221–243.

Shiller, R. J. (2000), *Irrational Exuberance*, Princeton University Press, Princeton, NJ.

Soroka, S. and McAdams, S. (2015), 'News, politics, and negativity', *Political Communication* **32**(1), 1–22.

Soroka, S., Young, L. and Balmas, M. (2015), 'Bad news or mad news? sentiment scoring of negativity, fear, and anger in news content', *The ANNALS of the American Academy of Political and Social Science* **659**(1), 108–121.

Appendix

A Model Derivation

This appendix derives the Bellman equation, the portfolio first-order condition, and the local approximation used in the main text.

A.1 Bellman equation

Let the state vector be

$$s_t \equiv (w_t, \rho_t, \mu_t, h_t).$$

The representative investor solves

$$V_t(s_t) = \max_{c_t,\pi_t} \left\{ u(c_t;\rho_t) + \beta \mathbb{E}_t \left[V_{t+1}(s_{t+1}) \right] \right\}, \tag{A.1}$$

subject to

$$w_{t+1} = a_t\big(R_f + \pi_t x_{t+1}\big). \tag{A.2}$$

The associated first-order conditions are

$$c_t^{-\rho_t} = \beta \mathbb{E}_t \left[V_{w,t+1}(s_{t+1})\, R_{p,t+1} \right], \tag{A.3}$$

and

$$\mathbb{E}_t \left[V_{w,t+1}(s_{t+1})\, x_{t+1} \right] = 0. \tag{A.4}$$

A.2 Local approximation

Let $\bar{s} \equiv (\bar{w}, \bar{\rho}, \bar{\mu}, \bar{h})$ denote the steady state, and define

$$\Delta w_{t+1} \equiv w_{t+1} - \bar{w}, \qquad \Delta\rho_{t+1} \equiv \rho_{t+1} - \bar{\rho}, \qquad \Delta\mu_{t+1} \equiv \mu_{t+1} - \bar{\mu}, \qquad \Delta h_{t+1} \equiv h_{t+1} - \bar{h}.$$

Define also

$$G_{t+1} \equiv x_{t+1} V_{w,t+1}(s_{t+1}),$$

so that (A.4) can be written as $\mathbb{E}_t[G_{t+1}] = 0$.

A second-order Taylor expansion of G_{t+1} around the steady state, centred at $\bar{x} = 0$, yields

$$\begin{aligned} 0 \approx &(A_t + B_t m_t)\mathbb{E}_t[x_{t+1}] + a_t \pi_t B_t\, \mathbb{E}_t[x_{t+1}^2] + C_t\, \mathbb{E}_t[\Delta\rho_{t+1} x_{t+1}] + \\ &D_t\, \mathbb{E}_t[\Delta\mu_{t+1} x_{t+1}] + H_t\, \mathbb{E}_t[\Delta h_{t+1} x_{t+1}], \end{aligned} \tag{A.5}$$

where $A_t \equiv V_{w,t+1}(\bar{s})$, $B_t \equiv V_{ww,t+1}(\bar{s})$, $C_t \equiv V_{w\rho,t+1}(\bar{s})$, $D_t \equiv V_{w\mu,t+1}(\bar{s})$, $H_t \equiv V_{wh,t+1}(\bar{s})$, and $m_t \equiv a_t R_f - \bar{w}$.

Define the hedging terms

$$\kappa_{\rho,t} \equiv \frac{C_t}{A_t + B_t m_t} \frac{\mathbb{E}_t[\Delta\rho_{t+1} x_{t+1}]}{\mathbb{E}_t[x_{t+1}]}, \tag{A.6}$$

$$\kappa_{\mu,t} \equiv \frac{D_t}{A_t + B_t m_t} \frac{\mathbb{E}_t[\Delta\mu_{t+1} x_{t+1}]}{\mathbb{E}_t[x_{t+1}]}, \tag{A.7}$$

$$\kappa_{h,t} \equiv \frac{H_t}{A_t + B_t m_t} \frac{\mathbb{E}_t[\Delta h_{t+1} x_{t+1}]}{\mathbb{E}_t[x_{t+1}]} \tag{A.8}$$

Let

$$\kappa_{H,t} \equiv 1 + \kappa_{\rho,t} + \kappa_{\mu,t} + \kappa_{h,t}.$$

Then the local portfolio condition becomes

$$0 \approx (A_t + B_t m_t)\mathbb{E}_t[x_{t+1}]\,\kappa_{H,t} + a_t \pi_t B_t\,\mathbb{E}_t[x_{t+1}^2]. \tag{A.9}$$

Rearranging yields

$$\mathbb{E}_t[x_{t+1}] \approx \lambda_t\,\mathbb{E}_t[x_{t+1}^2], \qquad \lambda_t \equiv -\frac{a_t \pi_t B_t}{(A_t + B_t m_t)\kappa_{H,t}}. \tag{A.10}$$

This is the reduced-form object used in the main text. For a myopic investor, the hedging channel is shut down by setting $\kappa_{H,t} = 1$.

B Additional Tables

B.1 Mean Differences Across News Types

Table B.1: Emotion Differences Across News Types

Emotion	Econ. − NE	Sports − NS
Panel A: Dictionary emotion shares		
Fear	0.327***	-0.579***
	(0.003)	(0.002)
Anger	0.174***	-0.319***
	(0.002)	(0.002)
Sadness	0.125***	-0.289***
	(0.002)	(0.002)
Happiness	-0.283***	-0.116***
	(0.002)	(0.002)
Surprise	-0.083***	0.039***
	(0.001)	(0.001)
Disgust	0.020***	-0.232***
	(0.001)	(0.001)
Panel B: LLM emotion shares		
Fear	-6.111***	8.690***
	(0.077)	(0.066)
Anger	-0.490***	-0.514***
	(0.009)	(0.010)
Sadness	8.404***	-1.785***
	(0.054)	(0.028)
Happiness	0.189***	-1.866***
	(0.020)	(0.012)
Surprise	-1.547***	-1.283***
	(0.016)	(0.024)
Disgust	-0.880***	-0.362***
	(0.011)	(0.013)
N	7,709,678	

Notes: Cells report percentage-point coefficients, with publication-date-clustered standard errors in parentheses, from article-level regressions of the emotion share on economic-news and sports-news indicators entered jointly. "NE" stands for non-economic, "NS" stands for non-sports. All models include publication-date and newspaper-outlet fixed effects. The economic coefficient compares economic with non-economic news conditional on sports classification; the sports coefficient compares sports with non-sports news conditional on economic classification. Dictionary economic news uses the rule-based economy indicator; LLM economic news uses the ModernBERT economy classifier. The sports classification is rule-based in both panels. Dictionary outcomes are NRC emotion-word shares; LLM outcomes are ModernBERT-predicted contextual emotion shares. The sample contains articles with at least 100 words published Monday–Friday from 1 July 2003 to 31 December 2022. ***, **, and * denote significance at the 1, 5, and 10 per cent levels.

B.2 Robustness: Excluding Daily Mail

Table B.2: Price-of-Risk Regressions, Daily Mail Excluded, λ_{t+1}

	OLS	IV–NE	IV–Sports
Panel A: Dictionary emotions			
Fear	-0.728***	-1.982***	-0.954*
	[-1.284, -0.280]	[-3.577, -0.746]	[-2.178, 0.164]
Anger	0.514***	2.080**	1.243
	[0.132, 1.000]	[0.409, 4.327]	[-0.485, 3.348]
Panel B: LLM emotions			
Fear	0.805***	1.730**	1.382
	[0.275, 1.455]	[0.235, 3.377]	[-5.871, 9.000]
Anger	0.584***	2.844***	3.535
	[0.222, 1.018]	[0.841, 5.732]	[-13.40, 28.47]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	325.42	119.96
Min. SW F, LLM	–	79.43	6.20
Mean of proxy	7.354	7.354	7.354
Observations	4,927	4,927	4,927

Notes: The table reports daily OLS and IV regressions after removing Daily Mail articles before reconstructing economic-, non-economic-, and sports-news dictionary measures and the corresponding LLM measures. The estimation sample runs from 2 July 2003 to 30 December 2022. The dependent price-of-risk state is dated $t + 1$ and is represented by the conditional mean–variance ratio for the return at $t + 2$; emotions, excluded instruments, and the other controls are dated t. The dependent variable is an AR(1)–GARCH(1,1) conditional-mean-to-variance ratio. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. Each panel jointly includes fear and anger. Emotion indices are standardised. NE denotes non-economic news. IV–NE instruments economic-news emotions with matching non-economic-news emotions. IV–Sports instruments economic-news dictionary emotions with sports-news dictionary emotions; the LLM panel uses sports-news dictionary instruments. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, family-matched HIV4 polarity, log UK Economic Policy Uncertainty, and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. HIV4 polarity is $(P - N)/(P + N)$, where P and N are the article's counts of HIV4 positive and negative words. Each moving-block draw reconstructs the generated proxy and constructed controls. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, **, and * denote exclusion of zero from the 99, 95, and 90 per cent percentile intervals, respectively.

B.3 Robustness: Six-Emotion Model

Table B.3: Price-of-Risk Regressions, Six Emotions

	(1) OLS	(2) IV–NE	(3) IV–Sports
Panel A: Dictionary emotions			
Fear	-0.448***	-1.412**	-1.848*
	[-0.893, -0.104]	[-3.206, -0.011]	[-4.706, 0.329]
Anger	0.567***	1.888**	1.723*
	[0.141, 1.124]	[0.238, 4.094]	[-0.165, 4.195]
Sadness	-0.442***	-0.355	-0.171
	[-0.818, -0.141]	[-1.392, 0.644]	[-1.576, 1.274]
Happiness	0.035	0.516	-1.187
	[-0.241, 0.338]	[-1.037, 2.329]	[-3.977, 1.234]
Surprise	-0.151	-1.623***	-0.468
	[-0.388, 0.046]	[-3.007, -0.573]	[-1.786, 0.755]
Disgust	-0.016	-0.192	0.656
	[-0.255, 0.213]	[-1.390, 0.913]	[-0.783, 2.367]
Panel B: LLM emotions			
Fear	2.768***	-6.155	-9.462
	[0.710, 5.340]	[-1498.8, 1793.2]	[-516.1, 698.3]
Anger	2.784***	13.57	-17.30
	[0.979, 4.985]	[-841.5, 925.0]	[-557.6, 768.9]
Sadness	1.240*	-2.810	-7.059
	[-0.095, 2.877]	[-993.6, 1202.2]	[-414.2, 554.1]
Happiness	0.365	-9.989	-3.758
	[-0.894, 1.781]	[-847.3, 992.7]	[-246.5, 320.8]
Surprise	0.366	-3.453	-1.551
	[-0.071, 0.847]	[-134.0, 157.7]	[-98.78, 100.7]
Disgust	-1.268***	-9.260	11.14
	[-2.381, -0.351]	[-279.1, 248.0]	[-347.2, 288.0]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	168.42	54.18
Min. SW F, LLM	–	0.36	0.52
Mean of proxy	7.354	7.354	7.354
Observations	4,927	4,927	4,927
Bootstrap replications	5,000	5,000	5,000

Notes: The table reports daily OLS and IV regressions for the estimation sample, 2 July 2003–30 December 2022. The dependent variable is an AR(1)–GARCH(1,1) conditional-mean-to-variance ratio. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. The dependent price-of-risk state is dated $t+1$ and is represented by the conditional mean–variance ratio for the return at $t+2$; emotions, excluded instruments, and the other controls are dated t. Emotion indices are standardised. Coefficients are estimated using OLS and IV. A positive coefficient indicates a higher price of risk and lower risk appetite. NE denotes non-economic news. IV–NE instruments economic-news emotions with matching non-economic-news emotions. IV–Sports instruments economic-news dictionary emotions with sports-news dictionary emotions; the LLM panel uses sports-news dictionary instruments. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, family-matched HIV4 polarity, and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. HIV4 polarity is $(P-N)/(P+N)$, where P and N are the article's counts of HIV4 positive and negative words. Each moving-block draw reconstructs the generated proxy and lagged controls. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, **, and * denote exclusion of zero from the 99, 95, and 90 per cent percentile intervals, respectively.

B.4 Robustness: Articles Classified as Economic by Either Method

Table B.4: Price-of-Risk Regressions, Union Article Universe

	(1) OLS	(2) IV–NE	(3) IV–Sports
Panel A: Dictionary emotions			
Fear	-0.974***	-1.550***	-1.075*
	[-1.633, -0.434]	[-2.759, -0.587]	[-2.415, 0.090]
Anger	0.732***	2.293***	2.294**
	[0.274, 1.342]	[0.733, 4.347]	[0.476, 4.736]
Panel B: LLM emotions			
Fear	1.489***	1.923***	2.076***
	[0.678, 2.342]	[0.702, 3.561]	[0.562, 3.908]
Anger	0.779***	0.701**	0.058
	[0.361, 1.263]	[0.082, 1.495]	[-2.177, 2.171]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	462.89	112.48
Min. SW F, LLM	–	265.97	43.36
Mean of proxy	7.354	7.354	7.354
Observations	4,927	4,927	4,927
Bootstrap replications	5,000	5,000	5,000

Notes: The table reports daily OLS and IV regressions for the estimation sample, 2 July 2003–30 December 2022. The dependent variable is an AR(1)–GARCH(1,1) conditional-mean-to-variance ratio. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. The dependent price-of-risk state is dated $t + 1$ and is represented by the conditional mean–variance ratio for the return at $t + 2$; emotions, excluded instruments, and the other controls are dated t. Both panels define economic news using the union of articles classified as economic by the rule-based dictionary classifier or ModernBERT. IV–NE instruments use articles classified as non-economic by both systems. The HIV4 polarity control is averaged over the same union-economic articles. IV–Sports retains the rule-based sports-news universe. Emotion indices are standardised. Coefficients are estimated using OLS and IV. A positive coefficient indicates a higher price of risk and lower risk appetite. NE denotes non-economic news. IV–NE instruments economic-news emotions with matching non-economic-news emotions. IV–Sports instruments economic-news dictionary emotions with sports-news dictionary emotions; the LLM panel uses sports-news dictionary instruments. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, family-matched HIV4 polarity, and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. HIV4 polarity is $(P - N)/(P + N)$, where P and N are the article's counts of HIV4 positive and negative words. Each moving-block draw reconstructs the generated proxy and lagged controls. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, **, and * denote exclusion of zero from the 99, 95, and 90 per cent percentile intervals, respectively.

B.5 Robustness: Including a Lagged Price-of-Risk Control

Table B.5: Price-of-Risk Regressions, Lagged Price-of-Risk Control

	(1) OLS	(2) IV–NE	(3) IV–Sports
Panel A: Dictionary emotions			
Fear	-0.339***	-0.659***	-0.330
	[-0.724, -0.072]	[-1.439, -0.105]	[-1.034, 0.276]
Anger	0.207**	0.646**	0.694
	[0.020, 0.527]	[0.046, 1.674]	[-0.190, 2.109]
Panel B: LLM emotions			
Fear	0.376***	1.476**	1.461
	[0.113, 0.795]	[0.178, 3.490]	[-1.119, 6.615]
Anger	0.190***	0.873***	0.601
	[0.045, 0.428]	[0.163, 2.053]	[-1.927, 4.972]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	373.03	94.33
Min. SW F, LLM	–	183.06	14.19
Mean of proxy	7.354	7.354	7.354
Observations	4,927	4,927	4,927
Bootstrap replications	5,000	5,000	5,000

Notes: The table reports daily OLS and IV regressions for the estimation sample, 2 July 2003–30 December 2022. The dependent variable is an AR(1)–GARCH(1,1) conditional-mean-to-variance ratio. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. The dependent price-of-risk state is dated $t+1$ and is represented by the conditional mean–variance ratio for the return at $t+2$; emotions, excluded instruments, and the other controls are dated t. Each specification also includes the GARCH price-of-risk state at t, the first lag of the dependent state. Emotion indices are standardised. Coefficients are estimated using OLS and IV. A positive coefficient indicates a higher price of risk and lower risk appetite. NE denotes non-economic news. IV–NE instruments economic-news emotions with matching non-economic-news emotions. IV–Sports instruments economic-news dictionary emotions with sports-news dictionary emotions; the LLM panel uses sports-news dictionary instruments. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, family-matched HIV4 polarity, and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. HIV4 polarity is $(P-N)/(P+N)$, where P and N are the article's counts of HIV4 positive and negative words. Each moving-block draw reconstructs the generated proxy and lagged controls. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, **, and * denote exclusion of zero from the 99, 95, and 90 per cent percentile intervals, respectively.

B.6 Robustness: Feedback from Market Returns to Instruments

Table B.6: Feedback from Market Returns to Instrument Emotions

	Dictionary specification		LLM specification	
	Fear	Anger	Fear	Anger
Panel A: Non-economic-news instrument emotions				
Return at t	0.117	1.839	-0.754	2.970
	[-1.798, 1.941]	[-0.385, 4.182]	[-3.261, 2.042]	[0.546, 5.704]
Holm-adjusted p-value	1.000	0.606	1.000	0.211
Panel B: Sports-news instrument emotions				
Return at t	0.363	1.227	-0.212	0.582
	[-2.193, 3.104]	[-1.080, 3.529]	[-2.786, 2.557]	[-1.784, 2.888]
Holm-adjusted p-value	1.000	1.000	1.000	1.000
Main controls	Yes	Yes	Yes	Yes
EPU	Yes	Yes	Yes	Yes
Observations	4,928	4,928	4,928	4,928
Bootstrap replications	5,000	5,000	5,000	5,000

Notes: Each column reports a separate daily OLS regression of an instrument-emotion measure at $t+1$ on the FTSE 100 excess return at t. Panel A uses emotions in non-economic news. Panel B uses dictionary emotions in sports news; the dictionary and LLM specifications differ in their family-matched polarity control. All regressions control for log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, family-matched HIV4 polarity, log EPU, and day-of-week effects. Emotion indices are standardised. Point estimates are bootstrap means; brackets contain 95 per cent moving-block-bootstrap percentile intervals. Holm-adjusted p-values correct jointly across the eight feedback regressions. The block length is 22 trading observations.

B.7 Robustness: Expanding-Window GARCH

Table B.7: Price-of-Risk Regressions, Expanding-Window GARCH

	(1) OLS	(2) IV–NE	(3) IV–Sports
Panel A: Dictionary emotions			
Fear	-0.725***	-1.579***	-0.843
	[-1.423, -0.236]	[-3.122, -0.354]	[-2.358, 0.509]
Anger	0.512**	2.103**	2.015*
	[0.051, 1.118]	[0.087, 4.782]	[-0.110, 4.956]
Panel B: LLM emotions			
Fear	1.095***	3.254***	3.549
	[0.367, 1.991]	[0.944, 6.493]	[-2.215, 12.41]
Anger	0.499***	1.633**	2.637
	[0.145, 0.957]	[0.278, 3.626]	[-4.847, 11.80]
Controls	Yes	Yes	Yes
Min. SW F, dictionary	–	499.99	102.34
Min. SW F, LLM	–	172.88	0.45
Mean of proxy	7.689	7.689	7.689
Observations	3,921	3,921	3,921
Bootstrap replications	1,000	1,000	1,000

Notes: The table reports daily OLS and IV regressions for 25 June 2007–28 December 2022. Panel A uses NRC dictionary word shares; Panel B uses LLM-predicted article emotion shares. Emotion indices are standardised. The dependent variable is the price-of-risk state at $t+1$, represented by the one-step conditional-mean-to-variance forecast for the return at $t+2$. An AR(1)–GARCH(1,1) model is fitted to an initial 1,008-return window and re-estimated at every step. IV–NE uses matching non-economic-news emotions; IV–Sports uses sports-news dictionary emotions in both panels. Controls are the excess return, log article count, log trading volume, Bank of England and CPI announcement indicators, the ten-year minus two-year gilt spread, log UK Economic Policy Uncertainty, family-matched HIV4 polarity and day-of-week effects. The minimum SW F is the minimum Sanderson–Windmeijer conditional first-stage statistic across the endogenous emotions in the indicated panel; it is undefined for OLS. Point estimates are bootstrap means; brackets are 95 per cent percentile intervals. ***, ** and * denote exclusion of zero from the 99, 95 and 90 per cent percentile intervals.