



Facial expressions and the business cycle[☆]

J. Daniel Aromi^a, Adam Clements^{b,*}

^a FCE, Universidad de Buenos Aires, CONICET-Universidad de Buenos Aires Instituto Interdisciplinario de Economía Política de Buenos Aires (IIEP-BAIRES),

FCE Universidad Católica Argentina, Argentina

^b School of Economics and Finance, Queensland University of Technology, Australia



ARTICLE INFO

JEL classification:

D91

E30

E70

Keywords:

Emotions

Facial expressions

Economic growth

Business cycle

ABSTRACT

This paper is the first to consider the link between information in the facial expressions of economic actors and economic activity. While much research has focused on text based sentiment, little is understood about the possible information conveyed by facial expressions. A collection of media photographs corresponding to the US economy for the period 1996–2018 is used to construct indices of emotions communicated by facial expressions. The indices are correlated with business cycle conditions, and also contain information about the future state of the business cycle, information that is incremental to that contained in commonly used measures of economic conditions. Social cognition means that emotions transmitted via facial expressions can influence expectations regarding economic prospects, decision making and hence future economic outcomes. The results here open up the possibility of using information from facial expressions to augment existing methods for using text-based sentiment for macroeconomic or financial forecasting.

1. Introduction

There is a well established literature that considers the information contained in text communicated through mass media outlets, social networks or even macroeconomic announcements (Tetlock, 2007; Bloom et al., 2016; Smales and Apergis, 2017; Ardia et al., 2019; Audrino et al., 2020; Liang et al., 2021; Mazumder, 2021), among many others. The information extracted from text data, often in the form of sentiment scores, has been used in various contexts. The applications range from measuring uncertainty in economic conditions through to forecasting equity returns, volatility or macroeconomic indicators. Gentzkow et al. (2019) and Algaba et al. (2020) provide wide ranging overviews of the literature in this growing area.

Moving beyond textual analysis, this paper takes a novel direction. It considers if the emotions shown in the facial expressions of different economic actors (taken from media photographs of policy maker, traders, corporate managers etc.) convey important information relating to economic conditions. This issue can now be addressed exploiting enhanced access to images in digital format and advances in machine learning algorithms. In this study, computer vision tools are used to rate

photographs taken by the media (related to a range of economic and financial topics) in terms of emotions communicated by facial expressions. The ratings of emotions from individual images are then used to construct indices of emotional states. Examining the informational content of these indices in relation to economic conditions is the main focus of this paper. In contrast to textual analysis, image analysis in this area is a more recent strand of research. While there may be a well established role for emotions in explaining social judgements in different contexts such as leadership and voting outcomes (Little et al., 2007; Re et al., 2013), decision-making in competitive games (Rojas et al., 2011) and trustworthiness (Todorov and Oosterhof, 2011), large scale image processing to measure emotions is a developing area. Boxell (2018) and Peng (2018) examine the emotions conveyed by politicians in media photographs, along with bias across different media outlets. Blankespoor et al. (2017), Graham et al. (2017), Huang et al. (2018) and Hu and Ma (2020) are recent examples of using information extracted from facial expressions to explain a range of corporate and financial outcomes, IPO performance, entrepreneurial funding and perceived CEO competence and firm performance. Joo and Steinert-Threlkeld (2018)

[☆] We thank the editor, Sushanta Mallick, and two anonymous referees for comments and suggestions that greatly improved the paper. J. Daniel would also like to thank Daniel Heymann, Eduardo Olaberria and Marcelo Resico for helpful comments and discussions. Adam Clements would also like to thank Benno Torgler and Stephen Whyte for helpful insights on preliminary work. All remaining errors are our own.

* Corresponding author.

E-mail addresses: aromi.daniel@gmail.com (J.D. Aromi), a.clements@qut.edu.au (A. Clements).

<https://doi.org/10.1016/j.econmod.2021.105563>

Received 28 September 2020; Received in revised form 5 May 2021; Accepted 31 May 2021

Available online 4 June 2021

0264-9993/© 2021 Elsevier B.V. All rights reserved.

provide an in-depth overview of automated methods for visual content analysis relevant to empirical work in this area.

The role for emotions in this context is based on the concept of social cognition, how individuals process and act on information about other people and social situations (Fiske and Taylor, 1991; Kunda, 1999), specifically emotions in facial expressions here. The idea that emotions are contagious is an old one (Smith, 1759; Hatfield et al., 1993) and implies that emotions can be transmitted from the facial expressions of economic actors to individuals. Emotions transmitted via facial expressions can influence beliefs and expectations regarding economic prospects (Johnson and Tversky, 1983; Arkes et al., 1988; Damasio, 1994; Lerner et al., 2015), and current decision making and hence future economic outcomes (Fenton-O’Creevy et al., 2012; Bénabou and Tirole, 2016; Hirshleifer, 2020). Abstracting from facial expressions and emotions, Blinder and Krueger (2004), Barbaglia et al. (2020) and Ellingsen et al. (2020) suggest that the sentiment contained in text data affects economic outcomes through broadly similar channels. Media is an important source of information about the state of the economy with news influencing expectations, hence consumption and investment decisions and finally economic activity.

To address the central issue here, indices of emotions are constructed from media photographs by employing facial recognition technology. To examine the informational content of these indices, the relationship between indices of emotions and economic activity is considered from a number of different perspectives. Emotions are found to be correlated with changes in both current economic conditions and future activity. Positive emotions (happiness) increase and negative emotions (anger, sadness and fear) decrease with improved economic conditions. Granger causality results show that the movements in the emotion indices help explain future activity indicating that the indices may contain useful information in relation to business cycle conditions. Statistically significant relationships between indices of emotion and contemporaneous and future business cycle conditions are found, which are robust to a number of commonly used measures of macroeconomic conditions. Gains in forecast accuracy are also observed from using emotion indices, particularly during periods of crisis. These results, show that the emotions conveyed by individuals in the public spotlight convey important information regarding business conditions, and that beyond their spoken words, their facial expressions also convey information.

Section 2 presents the methodology and underlying data sources. Section 3 provides a broad preliminary analysis of the behaviour of the emotion indices constructed. Section 4 reports the empirical results for the more formal analysis of the information contained in the indices, including an analysis of both the contemporaneous and future links with economic activity. A number of robustness checks are also provided along with a small forecasting exercise. Conclusions and a number of potentially interesting avenues for future research are presented in Section 5.

2. Data and methodology

This study constructs indicators of emotions from the facial expressions inferred from media photographs of a range of economic actors. The information content of these indicators is then evaluated in terms of their association with business cycle conditions, taking into account a number of macroeconomic control variables. The indices of emotion are evaluated based on their association with the quarterly seasonally adjusted real GDP growth. This indicator of economic activity is procured from the data portal maintained by Federal Reserve Bank of Saint Louis.¹ A detailed description of how the indices of emotion are constructed is now presented. A brief discussion of the macroeconomic control variables then follows.

2.1. Indices of emotion from facial expressions

The pictures used to construct the indices of emotion correspond to previews available from a public platform maintained by a news agency.² Importantly, each image displayed in the platform is accompanied by information on date of creation, date of submission and a brief text that describes the content of the picture. This information is exploited to identify the relevant sets of pictures used for different specifications of the indicators. The construction of the indicators of facial expressions can be described as a three step process: selection of photographs, processing of images and aggregation of information inferred from individual faces. The first stage requires specifying a set of economic topics and establishing a strategy to select the relevant images. In the second step, a set of computer vision tools are used to extract faces from pictures and rate them in terms of emotional content. In the final stage, the flow of information provided by these classified images is combined to obtain a quarterly indicator of facial expressions. Below, each step is described in more detail.

2.1.1. Selection of images

The indicators of facial expressions are built targeting different aspects of the US economy. The three topics or areas are: Policy Making, Stock Market and Corporations. To identify relevant photographs in each of these cases, the text describing each photograph is evaluated. In the first step of the process, we use a set of economic keywords to identify a subset of pictures using the search tool of the platform (AP Images: www.apimages.com). The list of terms used is: ‘business’, ‘economy’, ‘finance’, ‘economic’, ‘financial’, ‘bank’, ‘stock market’, ‘industrial’, ‘industry’ and ‘interest rate’. Next, we inspect the text describing each picture in this subset to select those pictures related to specific economic topics. More precisely, we evaluate if the corresponding text includes keywords that are associated with the three economic aspects: Policy Making, Stock Market and Corporations, see Table 1 for a list of the keywords.³ For example, in the case of policy making ‘Federal Reserve’ and ‘U.S. Treasury’ are among the proposed keywords. If the description of a picture includes any of these terms or any other term in the list of keywords, the picture is selected for the computation of the index associated with policy making. One particular concern is that prominent companies are mentioned in very different contexts. For example, sport arenas and performing arts centers are often sponsored by, and hence named after corporations. To account for this, in the case of the Corporations index, selected pictures were required to satisfy an additional condition. We kept the 66% of pictures whose text is most closely related to economic issues. More specifically, we used BERT (Devlin et al., 2018), a standard pre-trained language model, to estimate the extent to which the text describing the photo is related to economic topics versus sport events or entertainment. In the appendix we provide a detailed description of how we implemented this additional filter, along with a sample of text ranked from least to most related to economic concepts (Table 12). A number of recent works suggest using a similar approach where a final computational step is taken to classify text data, see Bao et al. (2020), Wu et al. (2020), Wan et al. (2020) and Khan et al. (2020). A small scale evaluation of the classification strategy (100 pictures and associated text) results in an estimated precision of 0.78 and estimated recall of 0.80.

We have excluded pictures that are submitted well after their creation date to try and minimise the effect of the submission of stock type photographs relating to a particular topic or individual that had been taken much earlier, thus reducing the effect of editorial bias. In such an instance, the emotions conveyed by these photographs may not

² The platform is AP Images: www.apimages.com.

³ Alternative list of keywords have been evaluated with similar results. For example, lists with a larger group of government agencies or a larger set of market venues have been considered. The results are available upon request.

¹ <https://fred.stlouisfed.org/>.

Table 1

Keywords by economic topic. Given an economic topic, a picture is selected if the text describing the picture includes at least one of the corresponding keywords.

Policy Making:	Federal Reserve, US Treasury, U.S. Treasury, Secretary of the Treasury, Treasury Department, Department of the Treasury, Treasury Secretary.
Stock Market:	New York Stock Exchange.
Corporations:	Top 300 companies in the Forbes 2000 list of US corporations. (To avoid false positives, names with less than 5 characters are excluded)

be related to conditions at the time at which they are submitted. To achieve this, we construct a filter exploiting the image creation and submission dates provided by the repository, by selecting pictures only if the month of the creation date and submission date match.

2.1.2. Processing images

The faces in pictures are detected and classified using Python package *fer*.⁴ According to this package, face detection is carried out using a multitask cascade convolutional neural network as proposed in Zhang et al. (2016). The task of emotion classification is performed by the model implemented in Arriaga et al. (2017). This implementation involves model training using FER-2013 dataset. This is a collection of 35,887 grayscale images. Each image is classified as: ‘angry’, ‘disgust’, ‘fear’, ‘happy’, ‘sad’, ‘surprise’ or ‘neutral’ (in the paper we simply use the noun forms of each of the emotions, e.g. ‘angry’ and ‘happy’ in ‘fer’ is denoted as Anger and Happiness). More specifically, the trained model assigns a probability to each category. Using the category with the highest probability as a prediction, the trained model achieves an accuracy of 66%, which is approximately human level accuracy as documented in Goodfellow et al. (2013). This degree of accuracy is based on the FER-2013 dataset used for training. If the images contained in the sample from AP Images differ systematically from those in the training dataset, the degree of accuracy will be compromised. Measuring accuracy on the basis of images that have undergone human validation would give a clearer picture of the performance of the classification method in the current sample. A more detailed description of the classification process is contained in the Appendix.

2.1.3. Indicators of facial expressions

The previous steps result in a collection of faces, with corresponding creation dates and facial expression scores. An indicator of facial expression is built summarizing this information at a quarterly frequency. Initially, quarterly indicators corresponding to each facial expression are computed using the average value of the score for the corresponding pictures. Then, the information corresponding to these seven indicators is summarized computing the first principal component. Finally, to account for different intensities of information arrival, the first principal component is multiplied by the logarithm of the number of faces corresponding to each quarter.

More formally, let n_t be the number of faces collected for quarter t and $\{r_{it}\}_{i=1}^{n_t}$ represent the seven-dimensional scores for each emotion linked to the collection of n_t faces. Then, the average scores for pictures created on that quarter is given by the seven-dimensional vector

$$\bar{r}_t = \sum_{i=1}^{n_t} r_{it}/n_t. \tag{1}$$

After standardization of the seven series, we compute the first principal component, through PCA, to get a one-dimensional series r_t^* . In the final step, this indicator is multiplied by $\log(n_t)$. In summary, the index of facial expressions I_t is given by

$$I_t = \log(n_t)r_t^*. \tag{2}$$

The indices of emotion reflect both the emotions in facial expressions, and the volume of the photographs in a similar way that the sentiment indices constructed by Barbaglia et al. (2020) capture both the sentiment and volume of text based news.⁵ To ensure that the subsequent results are not simply driven by the volume of news captured by $\log(n_t)$, a small number of results will be presented based solely on the emotions, $I_t = r_t^*$.

It is worth noting that the indices are the combined outcome of facial expressions expressed by economic actors (policy makers, traders), and the views of the photographers and editors of the news agency in terms of the images they view as worth distributing. This could be affected by variations in the type of events covered and the selection of pictures by members of the news organization. For example, during expansions the frequency of IPOs increase (Pástor and Veronesi, 2005), and as a result the index could vary because more pictures cover this type of event. Also, reporters might select pictures to distribute according to their views regarding the extent to which a given image, and its associated emotion, is representative of the event. In other words, the documented regularities involve both perceived emotions and opinions regarding relevant perceived emotions. This distinction needs to be kept in mind for a proper interpretation of the results.

2.2. Macroeconomic control variables

In examining the informational content of the indices of emotions, two sets of macroeconomic control variables are used. The first is the News-based component of the Economic Policy Uncertainty (EPU) index of Bloom et al. (2016) which is based on newspaper coverage of policy-related economic uncertainty.⁶ Bloom et al. (2016) find the full EPU index helps explain future movements in major macroeconomic indicators including output and employment. The news based uncertainty is chosen here as it is based on the sentiment extracted from news articles related to economic conditions. It is the logical choice to compare to the sentiment (in the form of emotions) extracted from facial expressions. McCracken and Ng (2020) developed the FRED-QD macroeconomic database, designed to offer a publicly available source of data. It contains more than 200 quarterly macroeconomic series ranging across GDP, Industrial Production, Prices, Housing, Exchange and Interest Rates and Stock Markets to name a few. To construct a macroeconomic factor, we transform each raw series following the benchmark transformation suggested by McCracken and Ng (2020), standardize each series and compute the first principal component.

3. Preliminary analysis of estimated indicators

Facial expression indices are computed according to the procedure described above. Three indicators correspond to the three specific economic topics discussed above. In addition, a general indicator is computed combining the photographs corresponding to the three topics,

⁵ In addition to this index, we also evaluated the performance of the four indices smoothed using the Kalman filter as a robustness check with respect to the degree of noise in the indices. This approach leads to similar results reported in the main text, and are available upon request.

⁶ Data was obtained from https://www.policyuncertainty.com/us_monthly.html.

⁴ <https://github.com/justinshenk/fer>.

Table 2
Descriptive statistics of emotions by topic. Figures correspond to probabilities assigned to each extracted face by the facial expression classification model.

	Policy Making		Stock Market		Corporations	
	Mean	S.Dev.	Mean	S.Dev.	Mean	S.Dev.
Anger	0.215	0.209	0.206	0.194	0.201	0.208
Disgust	0.001	0.011	0.001	0.009	0.001	0.010
Fear	0.084	0.111	0.089	0.111	0.073	0.102
Happiness	0.180	0.290	0.215	0.303	0.241	0.337
Neutral	0.274	0.242	0.234	0.217	0.264	0.245
Sadness	0.215	0.189	0.228	0.194	0.194	0.187
Surprise	0.028	0.079	0.025	0.066	0.024	0.071
Number of classified faces	11,553		18,583		9892	

Table 3
Mean and volatility of the emotions (in the form of probability) by stage of the business cycle. Difference (Dif.) between expansion (Exp.) and Recession (Rec.). Two recessions are contained in the sample period: March 2001–November 2001 and December 2007–June 2009.

		Policy Making			Stock Market			Corporations		
		Exp.	Rec.	Dif.	Exp.	Rec.	Dif.	Exp.	Rec.	Dif.
		Anger	Mean	0.209	0.241	-0.032	0.204	0.221	-0.017	0.196
	St.Dev.	0.206	0.219	-0.013	0.194	0.197	-0.003	0.207	0.215	-0.008
Disgust	Mean	0.001	0.001	0.000	0.001	0.001	0.000	0.001	0.001	0.000
	St.Dev.	0.011	0.010	0.001	0.009	0.011	-0.002	0.010	0.008	0.002
Fear	Mean	0.084	0.085	-0.001	0.087	0.105	-0.018	0.072	0.081	-0.009
	St.Dev.	0.111	0.114	-0.003	0.108	0.125	-0.017	0.101	0.104	-0.003
Happiness	Mean	0.187	0.152	0.035	0.221	0.169	0.052	0.251	0.190	0.061
	St.Dev.	0.294	0.267	0.027	0.308	0.259	0.049	0.344	0.296	0.048
Neutral	Mean	0.273	0.278	-0.005	0.236	0.217	0.019	0.267	0.249	0.018
	St.Dev.	0.241	0.247	-0.006	0.219	0.200	0.019	0.248	0.224	0.024
Sadness	Mean	0.216	0.211	0.005	0.225	0.255	-0.030	0.187	0.225	-0.038
	St.Dev.	0.190	0.185	0.005	0.192	0.200	-0.008	0.185	0.195	-0.010
Surprise	Mean	0.028	0.030	-0.002	0.025	0.030	-0.005	0.024	0.023	0.001
	St.Dev.	0.078	0.082	-0.004	0.064	0.079	-0.015	0.071	0.073	-0.002

denoted below as the All index. For each version of the indicator, Table 2 reports both the number of detected pictures linked with each topic and the rate of occurrence of each emotion linked to each topic. Irrespective of the subset of pictures, on average, the highest average probabilities are associated with neutral, closely followed by sadness, happiness and anger. There are clear similarities in the average probabilities estimated for the different sets of images. Out of the seven emotions, happiness exhibits the highest volatility in the probabilities followed by neutral.

Table 3 reports the contemporaneous association between communicated emotions and the business cycle. According to NBER’s Business Cycle Dating Committee (<https://www.nber.org/cycles.html>), two recessions are covered by the sample period (start of 1996 to the end of 2018): March 2001–November 2001 and December 2007–June 2009. During recessions, the probability associated with happiness is consistently lower across the different topics. There are smaller increases in neutral expressions associated with the Stock Market and Corporations topics. Changes in the probabilities of the other emotions are generally much smaller in magnitude. This is the first indication of an association between communicated emotions and the state of the business cycle, highlighting in particular the importance of the happiness in facial expressions.

The results in Table 3 provide some very preliminary results showing that emotions, mainly happiness, change according to the state of the business cycle. However, to consider the information conveyed in facial expressions, the links between emotions and the future evolution of the business cycle (in the form of GDP growth as opposed to expansion vs. recession) are of interest here. To provide preliminary evidence of these links, Table 4 reports correlations between the probability of emotions (from all selected pictures) and current, and future

Table 4
Correlation between the probability of each emotion and GDP growth (taken across all selected pictures). *, ** and *** indicate significance at 10%, 5% and 1% respectively.

	Current quarter	Next quarter	Next year
Anger	-0.412***	-0.355***	-0.442***
Disgust	-0.042	0.105	-0.010
Fear	-0.392***	-0.235**	-0.237**
Happiness	0.462***	0.292**	0.390***
Neutral	0.062	0.098	0.071
Sadness	-0.376***	-0.206*	-0.279**
Surprise	-0.100	-0.098	-0.134

GDP growth. There is strong positive correlation between happiness and GDP growth, consistent with the preliminary findings in Table 3. Very interestingly, these results reveal strong negative correlation between GDP growth and the negative emotions of anger, fear, sadness and to a lesser degree surprise.

As described in Section 2, once the images are classified, an index is computed averaging scores observed in each time period, computing the first principal component of the average scores. Table 5 shows, for alternative sets of images, the factor loadings corresponding to each first principal component. The loadings suggest that a structural property is being captured by each first principal component as common patterns can be observed. First, happiness always displays a positive, and the largest loading in absolute terms. The loading on neutral is somewhat mixed. Whereas the loadings on the negative emotions, in particular anger, fear and sadness are all negative. The emotions with the highest factor loadings are also the most strongly correlated with

Table 5
Factor loadings for each index on the individual emotions. PCA is computed after standardization of the respective indices.

	Policy Making	Stock Market	Corporations	All
Anger	-0.416	-0.394	-0.351	-0.450
Disgust	-0.311	-0.042	-0.105	-0.063
Fear	-0.380	-0.512	-0.448	-0.445
Happiness	0.476	0.510	0.597	0.535
Neutral	0.199	0.099	-0.184	0.153
Sadness	-0.564	-0.426	-0.500	-0.508
Surprise	-0.029	-0.360	-0.156	-0.167

GDP (Table 4). Happiness (Anger, Fear and Sadness) exhibits a positive (negative) loading on the factor and is positively (negatively) correlated with GDP growth. The PCA results indicate the first principal components in each case reflect an easily interpretable indices of emotion.

Fig. 1 shows the trajectory of the indices computed using PCA, with the shaded areas representing recessions. Some common patterns linking the indices to macroeconomic and financial market conditions can be observed. Relatively low values are observed for most of the indices around the time of the 2001 recession. This effect is not as pronounced in the Stock Market index. The indices also exhibit negative values around the 2008–2009 recession in the aftermath of the Global Finan-

cial Crisis. Policy drops early in this period, with Stock Market following and Corporations index being consistently negative during this period.

Fig. 1 shows that the different indices share some common behaviour, in particular around recession periods. Table 6 provides further evidence of the relative behaviour of the indices, reporting the volatilities of the individual indices, and the correlation between them. Of the three individual indices, Stock Market is the most volatile and Policy Making the least, with the combined All index a little more volatile than any of the three individual indices. The correlations between the individual topic indices are relatively low (0.402 is the highest between Policy and Stock), indicating the individual indices may contain different sets of information that should be reflected in the combined All index based on all of the collected photographs. The correlations between the All index and the three individual topics are generally higher, which is to be expected as they share common information with the All index that contains the same photographs.

Unit root test (Dickey-Fuller) results (not reported here but available upon request) show that the null hypothesis of a unit root can be rejected meaning the levels of the indices can be used in standard regression models in the subsequent empirical analysis. While the preliminary results so far are based on correlations, Granger causality results in Table 7 provide the first insights into the directional links between the emotion indices and economic activity. The results reveal that while movements in GDP growth do not explain changes in future

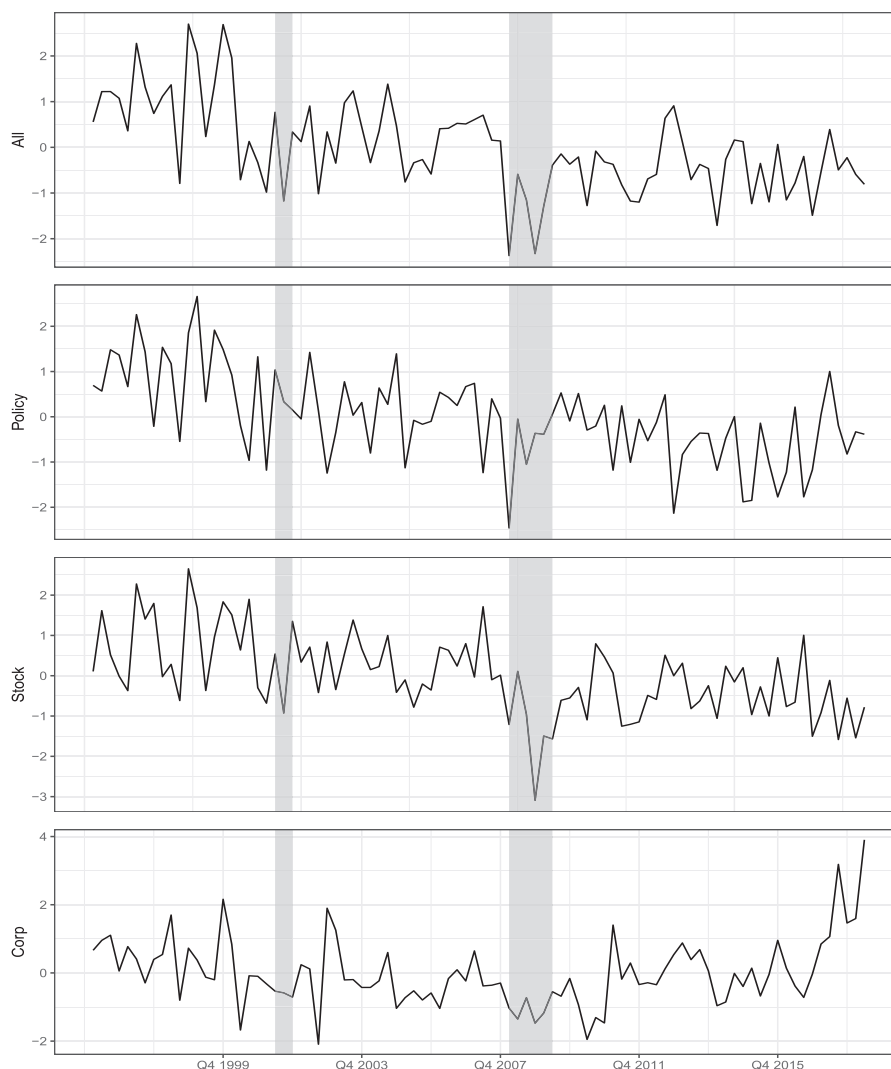


Fig. 1. Indices of communicated emotions. Shaded areas are quarters associated with recessions. Indices are standardized.

Table 6
Standard deviation of, and correlations between the indices of emotions. *, ** and *** indicate significance at 10%, 5% and 1% respectively.

	St. Dev.	Correlations			
		Policy Making	Stock Market	Corporations	All
Policy Making	4.974		0.402***	0.101	0.672***
Stock Market	6.593			0.125	0.827***
Corporations	5.562				0.355***
All	7.982				1

Table 7
Granger causality results. *p*-values are reported. Lags are selected according to AIC with maximum length 10. Estimations implemented using R package ‘vars’.

H0: GDP does not Granger-cause Policy	0.2079
H0: Policy do not Granger-cause GDP	0.0325
H0: GDP does not Granger-cause Stock Market	0.049
H0: Stock Market do not Granger-cause GDP	0.721
H0: GDP does not Granger-cause Corporations	0.182
H0: Corporations do not Granger-cause GDP	0.030
H0: GDP does not Granger-cause All	0.0709
H0: All do not Granger-cause GDP	0.0664

values for the emotions related to Policy or Corporations, these emotion indices influence future economic activity. The Stock Market index does not influence GDP but changes in economic activity do influence future emotions associated with the Stock Market. The broader All index show bi-directional linkages with economic activity at just over a 5% level of significance. Overall the results show that there are potential causal relationships worthy of further investigation.

The preliminary analysis presented in this section shows how both positive and negative emotions, and the ensuing indices of emotion are related to broad changes in business cycle conditions. When combined with the Granger causality results, these results motivate the more

formal analysis of the informational content of the indices presented below.

4. Results

In this Section, the links between the indices of emotion and economic activity are examined more formally. The information content of the indices is first examined through an autoregressive model for GDP growth that is extended using the indices as explanatory variables. First, the case of contemporaneous associations is considered. Next, the information provided by lagged values of the indices is estimated. This is followed by a number of robustness exercises taking into account the additional measures of economic conditions. The results of a small forecasting exercise are then presented.

4.1. Contemporaneous associations

A simple autoregressive model for GDP growth is extended incorporating the contemporaneous value of an indicator of facial expressions. Formally, the model is given by:

$$g_t = \alpha + \beta_{-1}g_{t-1} + \beta_1 I_t + u_t \tag{3}$$

where g_t is the quarterly growth rate of GDP for quarter t , I_t is an index of facial expressions and u_t is an error term.

Table 8 shows the estimation results for a range of models based on Equation (3). Panel A contains results for models containing either

Table 8
Estimation results for Equation (3). Significance levels: “*” 0.10, “***” 0.05 and “****” 0.01. Standard errors are estimated following Newey and West 1987, Newey and West 1994. Parameter estimates are standardized; absolute t-statistics in brackets.

	[1]	[2]	[3]	[4]	[5]
A.					
$\hat{\beta}_{-1}$	0.390** [2.3]	-	-	-	-
$\hat{\beta}_{Pol.Making}$	-	0.414*** [4.9]	-	-	-
$\hat{\beta}_{StockMarket}$	-	-	0.536*** [2.9]	-	-
$\hat{\beta}_{Corporations}$	-	-	-	0.178 [1.1]	-
$\hat{\beta}_{All}$	-	-	-	-	0.549*** [3.3]
Adj. R^2	0.142	0.162	0.280	0.021	0.293
B.					
$\hat{\beta}_{-1}$	0.390** [2.3]	0.316* [1.6]	0.205 [1.2]	0.369** [2.3]	0.206 [1.4]
$\hat{\beta}_{Pol.Making}$	-	0.346*** [4.8]	-	-	-
$\hat{\beta}_{StockMarket}$	-	-	0.453*** [3.5]	-	-
$\hat{\beta}_{Corporations}$	-	-	-	0.097 [1.1]	-
$\hat{\beta}_{All}$	-	-	-	-	0.468*** [4.3]
Adj. R^2	0.142	0.249	0.307	0.142	0.322

g_{t-1} or the indices of emotion individually. Panel B reports results for the models with both the indices of emotion and lagged growth. The information content of the indices is highlighted by the respective estimated coefficients on the emotion indices along with the variation in the adjusted R^2 .

With the exception of the Corporations index, Panel A shows that the estimated coefficients associated with the other indices are positive and significant and offer improvements in R^2 relative to the model that only contains lagged growth. The strongest link corresponds to the All index, that is, the index which results from combining the pictures related to Policy Making, the Stock Market and Corporations. In this case, the estimated coefficient indicates that a one standard deviation increase in the index is associated with an increase of 0.55 standard deviations in GDP growth. As discussed earlier in Section 2.1.3, emotion indices are also computed by excluding the volume of images. Defining the ALL index based only on $I_t = r_t^*$ leads to a very similar result to that in the last row of Panel A: $\hat{\beta}_{All} = 0.514$, a t-statistic of 3.5 and an adjusted $R^2 = 0.256$ are obtained which are very similar to the result reported in Panel A. The results in Panel B show that even in the presence of lagged growth, the estimated coefficients on the indices are positive and significant with improvements in R^2 relative to the model containing only lagged growth. The importance of lagged growth is also reduced. Again, excluding volume from the ALL index leads to a very similar outcome to that reported in Panel B, $\hat{\beta}_{All} = 0.428$, a t-statistic of 4.5 and an adjusted $R^2 = 0.291$ are obtained. These two extra results based on $I_t = r_t^*$ show that the information in the emotion indices is not simply driven by the volume of images.

The estimated models also point to differences in the information content of alternative specifications of the index. The weak association in the case of the index linked to Corporations, may reflect the types of announcements made when photographs are taken. It is possible that many announcements made by companies are related to the future outlook of investment, profits or earnings etc., and hence may contain less information about current economic conditions.

4.2. Associations with future activity

It is possible that the model in Equation (3) suffers from endogeneity given that both emotions captured in I_t and economic activity, g_t may be jointly determined. Before formally considering forecasting style models, we examine the information conveyed in the indices of emotion again while mitigating the impact of possible endogeneity, a framework similar in spirit to that of Barbaglia et al. (2020) is used. This framework breaks the contemporaneous link between emotions and growth and examines how the explanatory power of emotions changes with differences in timing. The setting here differs from Barbaglia et al. (2020) in that there is no real-time forecasting. While Barbaglia et al. (2020) mix frequencies, e.g. quarterly and monthly, preliminary experiments to construct monthly indices of emotion here led to noisier indices as the number of pictures per month was lower. Here we consider the information in the emotion indices from three month rolling windows that are updated more frequently than lagged growth. This is based on a forecast exercise that takes place k months before the end of the target quarter. In the baseline model, the information set is the latest realized quarterly GDP growth. In the extended model, the most recent emotions observed at that point in time are used. For example, consider the case of forecasting growth for the 1st quarter of 2000, and $k = 4$, meaning the simulated forecast takes place in November 1999. This means the latest realized quarterly GDP growth rate is for the 3rd quarter of 1999. The baseline model for this exercise is:

$$g_t = \alpha + \beta_{AR}g(t, k) + u_t \tag{4}$$

where again g_t is the quarterly GDP growth for quarter t and $g(t, k)$ is the most recent realization of quarterly GDP growth observed as of k months before the end of quarter t . In the extended model, this information set is expanded using facial expression information that is updated




Fig. 2. R^2 from the baseline model in Equation (4) (solid line) and extended model in Equation (5) (dashed line) across a range of values for lags k .

at a higher frequency. More specifically, the index corresponds to a 3-month window that covers the current month and the two immediately preceding months. For example, a forecast generated in November of 1999 would be conditioned on facial expression information corresponding to September, October and November of 1999. Formally, the extended model is given by:

$$g_t = \alpha + \beta_{AR}g(t, k) + \beta_I I(t, k) + u_t \tag{5}$$

where $I(t, k)$ is the value of the facial expression index using information corresponding to $k, k + 1$ and $k + 2$ months before the end of quarter t .

Fig. 2 shows the R^2 for both the baseline (from Equation (4) - solid line) and extended model (from Equation (5) - dashed line) containing the indices of emotion for different lags k . Overall, the results show that the more frequently updated indices of emotion provide additional explanatory power over the baseline model containing the most recently observed growth rates as of k months ago. This additional explanatory power is consistently observed up to a lag of $k = 8$ months. This result extends those based on Equation (3) and shows that moving beyond a contemporaneous relationship (minimising the impact of endogeneity), the indices of emotion contain incremental information relative to the most recently observed measure of growth.

While contemporaneous associations between economic conditions and emotions are interesting, the link between current emotions and future activity is of central importance in relation to the informational content of the emotions. Economic decisions influenced by the emotions conveyed by economic actors will likely have a lasting effect on economic activity. Therefore examining the link between the indices of emotion and future economic activity is crucially important in addressing the informational content of communicated emotions. This is achieved using a slightly modified version of the model in Equation (3):

$$g_{[t,t+h]} = \alpha + \beta_0 g_t + \beta_I I_t + u_{[t,t+h]} \tag{6}$$

where $g_{[t,t+h]}$ is the cumulative growth rate of GDP from quarter t through quarter $t + h$, g_t is the growth rate for quarter t , I_t is an index of facial expressions and $u_{[t,t+h]}$ is an error term. Four forecast horizons are considered: $h \in \{1, 2, 4, 8\}$. Table 9 shows the estimation results based on Equation (6). In the case of the All index, and indices linked to Policy Making and Corporations, positive and statistically significant coefficients are observed in all cases with the exception of $h = 1$ and

Table 9
 Estimation results for Equation (6). Significance levels: “*” 0.10, “**” 0.05 and “***” 0.01. Standard errors are estimated following Newey and West 1987, Newey and West 1994. Parameter estimates are standardized; absolute t-statistics in brackets.

		h = 1	h = 2	h = 4	h = 8
Baseline adj. R ²		0.143	0.186	0.149	0.081
Policy Making	$\hat{\beta}_{PolMaking}$	0.037	0.137**	0.227**	0.185**
	t-stat.	[0.4]	[2.3]	[2.3]	[2.0]
	Adj. R ²	0.134	0.193	0.183	0.100
Stock Market	$\hat{\beta}_{StockMarket}$	0.061	0.000	0.032	0.004
	t-stat.	[0.8]	[0.0]	[0.4]	[0.0]
	Adj. R ²	0.136	0.177	0.140	0.070
Corporations	$\hat{\beta}_{Corporations}$	0.211**	0.262**	0.278*	0.359***
	t-stat.	[2.5]	[2.2]	[1.9]	[2.8]
	Adj. R ²	0.177	0.245	0.216	0.157
All	$\hat{\beta}_{All}$	0.213**	0.251***	0.309***	0.202
	t-stat.	[2.1]	[3.9]	[3.2]	[1.0]
	Adj. R ²	0.165	0.222	0.208	0.099
All: $I_t = r_t^*$	$\hat{\beta}_{All}$	0.192*	0.232***	0.304***	0.228
	t-stat.	[1.9]	[3.6]	[3.6]	[1.4]
	Adj. R ²	0.160	0.217	0.208	0.109

Table 10
 Evaluation of information content for alternative indicators of the business cycle. Significance levels: “*” 0.10, “**” 0.05 and “***” 0.01. Standard errors are estimated following Newey and West 1987, Newey and West 1994. Parameter estimates are standardized; absolute t-statistics in brackets.

	[1]	[2]	[3]	[4]	[5]
A. Current quarter					
$\hat{\beta}_{-1}$	0.389**	-	-	-	0.126
	[2.3]				[1.1]
$\hat{\beta}_{Fac.Expr.}$	-	0.549***	-	-	0.340***
		[3.3]			[4.0]
$\hat{\beta}_{EPU}$	-	-	0.456***	-	0.147**
			[2.6]		[1.9]
$\hat{\beta}_{FRED-QD}$	-	-	-	0.463*	0.254*
				[2.0]	[1.8]
Adj. R ²	0.142	0.293	0.199	0.205	0.393
B. Next quarter					
$\hat{\beta}_{-1}$	0.390**	-	-	-	0.229
	[2.3]				[1.4]
$\hat{\beta}_{Fac.Expr.}$	-	0.363***	-	-	0.190*
		[2.6]			[1.8]
$\hat{\beta}_{EPU}$	-	-	0.261*	-	0.032
			[1.9]		[0.3]
$\hat{\beta}_{FRED-QD}$	-	-	-	0.275	0.090
				[1.4]	[0.9]
Adj. R ²	0.143	0.122	0.057	0.065	0.154
C. Cum. growth next 4 quarters					
$\hat{\beta}_{-1}$	0.399***	-	-	-	0.265**
	[3.0]				[2.0]
$\hat{\beta}_{Fac.Expr.}$	-	0.435***	-	-	0.344**
		[2.9]			[2.3]
$\hat{\beta}_{EPU}$	-	-	0.176	-	0.102
			[1.6]		[0.7]
$\hat{\beta}_{FRED-QD}$	-	-	-	0.186*	-0.018
				[1.7]	[0.2]
Adj. R ²	0.149	0.179	0.020	0.024	0.198

the Policy Making index, and $h = 8$ with the All index. Beyond the estimated coefficients, it is worth noting, that in all cases, the inclusion of the index results in consistent increases in adjusted R². The final panel presents the results based on defining the All index only on $I_t = r_t^*$. Again these are consistent with the results in the panel immediately above, again confirming that the information in the indices is robust to excluding the volume of images.

The result in relation to the information in the Corporations index is consistent with the preliminary correlation analysis in Section 3 and may reflect corporate announcements that are generally forward looking. If corporate announcements are generally related to forward looking issues (investment, growth in profits or earnings etc.) then the emotions conveyed during these announcements would likely contain information about future economic conditions.

4.3. Robustness to macroeconomic variables

This section presents results to examine the robustness of the information in emotions to the inclusion of commonly used measures of economic conditions.

Results checking the robustness of the results to changes in the empirical model (Equations (3) and (6)), and to smoothing the indices using the Kalman Filter are available upon request, and show that the earlier results are robust to these changes.

Table 10 presents results relating to the robustness of the original results with respect to the inclusion of alternative macroeconomic indicators, the EPU index and FREQ-QD factor here. These results are based on the combined All index. The results in Panel A of Table 10 extend those of Section 4.1 by extending Equation (3). The results in Panels B and C extend those in Section 4.2 by extending Equation (6). The results in columns 1–4 of Panel A show that the emotions and EPU indices, and the FRED-QD factors are contemporaneously related to growth (along with lagged growth). When all indices are included in the full extended version of Equation (3), the results show that $\hat{\beta}_{Fac.Expr.}$ continues to be significant even in the presence of EPU and FRED-QD. When considering the links with future economic activity, Panels B and C show broadly similar results. In most cases, the results in columns 1–4 show that the emotions and EPU indices along with the FRED-QD are individually significant. Importantly, the results in column 5 show that $\hat{\beta}_{Fac.Expr.}$ continues to be significant in the presence of EPU and FRED-QD in an extended version of Equation (6). Overall, these results show that the index of emotion contains information beyond that contained in commonly used measures of economic conditions, namely the EPU index and the factor extracted from the FREQ-QD database.

4.4. Out-of-sample forecasting

We implement an out-of-sample forecasting exercise to provide further insights into the information content of the facial expression indices. The sample is divided into two subsets. The in-sample period is 1996–2006 and the out-of-sample period corresponds to 2007–2018. An expanding forecasting scheme is used. The in-sample period of 1996–2006 is used as the initial estimation window, which then expands through the out-of-sample period. Following the full-sample analysis described above, autoregressive models are estimated to generate baseline forecasts. These forecasts are then compared to forecasts generated by models that extend the autoregressive model incorporating an indicator of facial expressions. Three exercises are considered: nowcasts, quarter ahead forecasts and one year ahead forecasts.⁷

It must be noted that the sample is relatively short for this type of demanding forecasting exercise. As a result, the results described below can be considered to be informative, but should not be interpreted as a definitive test of the information content of the indices. Development of novel forecasting methods to harness the indices of emotion would be an interesting avenue for future research, one that is beyond the scope of this paper. Table 11 reports RMSE results for the baseline model containing only lagged growth and models containing emotion indices (All index) based on the current quarter, and average of last 2 or 4 quarters to examine if any benefits stem from more smoothing. In all cases, forecasts that use the facial expression indices are associated with lower RMSEs than those corresponding to the baseline autoregressive models. While the differences in accuracy fail to be consistently significant, the reductions in RMSE (in the form of the ratios) are large in magnitude, often around 10%. While RMSE is easily interpreted, and understood as a loss function, Patton and Sheppard (2007) provide simulation based evidence showing (R)MSE provides only low power in statistical tests

⁷ In the case of nowcasts, the index of facial expression for quarter t is added as a regressor to an autoregressive model in which the level of activity on quarter t is a function of its one quarter lagged value.

Table 11
Out of sample forecast results: RMSE. Significance levels: “*” 0.10, “**” 0.05 and “***” 0.01. p-values are computed using bootstrap following Faust et al. (2013).

	Nowcast	One quarter ahead	One year ahead
Baseline			
RMSE (A)	0.0069	0.0068	0.0203
Quarterly index			
RMSE(B)	0.0058	0.0065	0.0190
Ratio (B/A)	0.849	0.962	0.938
p-value	0.01	0.34	0.42
2-quarter index			
RMSE(C)	0.0060	0.0062	0.0181
Ratio (C/A)	0.879	0.922	0.893
p-value	0.02	0.12	0.19
4-quarter index			
RMSE(D)	0.0061	0.0062	0.0185
Ratio (D/A)	0.892	0.918	0.912
p-value	0.10	0.27	0.37




Fig. 3. Difference in cumulative sum of squared forecast error (CSSE) between forecasts from the baseline model (Equation (6) with lagged growth only) and the model (Equation (6) containing the ALL emotion index). Positive values in the plots indicate that the extended model containing emotions is providing more accurate forecasts. The top panel: Nowcasts. Middle panel: 1 quarter ahead. Bottom panel: 4 quarters ahead. Shaded areas are quarters associated with recessions.

comparing forecast performance. This observation, coupled with the small sample sizes here, limit any power in identifying statistically significant differences in forecast accuracy for the models containing the emotion indices.

To gain a deeper understanding of the information in the emotions index in the context of forecast performance, Fig. 3 plots the difference in cumulative sum of squared forecast error (CSSE) between forecasts from the baseline model (with lagged growth only) and the extended

model (including the ALL emotion index). Positive values in the plots indicate that the extended model containing emotions is providing more accurate forecasts. The shaded area represents the 2008–2009 recession in the aftermath of the GFC. These results show the largest improvements in forecast performance are due to using the All emotions index during the recession period. Smaller yet consistent gains are observed during the remaining part of the forecast period. Overall these results show that while the index of emotions offers information to improve forecast accuracy, the information is most valuable during periods of crisis.

5. Conclusions

This study examined the information content of the facial expressions of various economic actors in relation to the state of the business cycle. To address this issue, indices of emotions were constructed from media photographs by employing facial recognition technology. It was shown that emotions are correlated with changes in both current economic conditions and future activity. Positive emotions (happiness) increase and negative emotions (anger, sadness and fear) decrease as economic conditions improve. Statistically significant relationships, both with contemporaneous and future business cycle conditions were found. These links were also found to be robust to a number of commonly used measures of macroeconomic conditions. These results, show that the emotions conveyed by individuals in the public spotlight convey important information regarding business conditions. The role of emotions here is consistent with the idea of social cognition in that emotions transmitted via facial expressions influence beliefs and expectations regarding economic prospects. An implication of these results, is that beyond their spoken words, such individuals should be careful

Appendix

Picture selection for corporations

Given a collection of images available on the platform,⁸ we first use a set of economic keywords to identify a collection of pictures using the search tool of the platform (the list of terms is: ‘business’, ‘economy’, ‘finance’, ‘economic’, ‘financial’, ‘bank’, ‘stock market’, ‘industrial’, ‘industry’ and ‘interest rate’). Next, we inspect the text describing each picture in this subset to select those pictures related to specific economic topics. More precisely, we evaluated if the corresponding text includes keywords that are associated with three economic aspects: policy making, stock market and corporations (see Table 1).

In the case of corporations, the reference to economic topics is less evident since the name of corporations could be used in different contexts. For example sports and performing arts venues are named after corporations. Given this concern, in this case we implement an additional filter. The text accompanying the image is required to be related to economic issues as inferred from a language model.

To implement this task, we use DistilBERT, a lighter version of the well-established BERT model proposed in Devlin et al. (2018). More specifically, we use the pre-trained model ‘distilbert-base-uncased’ from python package ‘transformers’ (<https://pypi.org/project/transformers/>). This type of model can be used to assess associations between different texts by inspecting the similarity of the respective numerical representation produced by the model. In our case, we evaluated whether the representation of the text describing an image (text embedding) was similar to the representations of two sequences of words. The first sequence of words, labeled ‘econ’, is given by: ‘economic, economics, finance and financial’. The second sequence of words, labeled ‘sport&ent’, is simply given by: ‘sport and entertainment’. A robustness check was undertaken by using the.

Formally, given a text describing an image, the model produces a vector \mathbf{v}_{image} that summarizes the information in that text. This output is known as text embedding and, under this model configuration, it has 768 dimensions. Similarly given the two sequences of thematic words, we obtain corresponding word embeddings: \mathbf{v}_{econ} and $\mathbf{v}_{sport\&ent}$. Given a pair of vectors \mathbf{v} and \mathbf{w} , their association can be assessed computing their cosine similarity: $\cos(\mathbf{v}, \mathbf{w}) = \frac{\mathbf{v}\mathbf{w}}{\|\mathbf{v}\|\|\mathbf{w}\|}$. Then, we approximate the association of an image with economic issues by computing the difference: $\cos(\mathbf{v}_{image}, \mathbf{v}_{econ}) - \cos(\mathbf{v}_{image}, \mathbf{v}_{sport\&ent})$. We only kept the 66% of images that, according to this metric, are more closely associated to economic events. A small scale evaluation of the classification strategy (100 pictures and associated text) results in an estimated precision of 0.78 and estimated recall of 0.80. In Table 12 we provide sample texts that have been ranked as most(least) related to economic issues. A subsequent check of the classification was undertaken by excluding ‘and’ from the word embeddings. This had a relatively small impact, only changing the classification outcomes in 4% of cases. It would be possible to explore this issue further by generating a new training sample and considering a more specific classification model.

regarding the emotions they convey as they contain information.

These findings could be understood as a lower bound on the informational content of emotions in facial expressions. Interesting directions for future research would include examining more comprehensive collections of pictures, possibly collecting images in real-time directly after events, and or improved computer vision techniques that may lead to more useful indicators of emotion. An experiment designed to gain a deeper understanding of the impact of images on economic expectations would be interesting to consider. The results would shed more light on the way individuals respond to observed facial expressions, and how they alter their expectations. A number of interesting avenues for empirical research stem from these results. Given that bond and equity returns are linked to economic conditions, these results have potentially interesting implications for investment strategies. Considering how to harness these indices to improve portfolio allocation would be an interesting direction to take, which may require more sophisticated modeling and forecasting approaches than standard multivariate volatility models. Considering the impact of emotions conveyed during economic or corporate announcements may also be a potentially fruitful direction for future research.

Declaration of competing interest

The authors whose names are listed immediately below certify that they have NO affiliations with or involvement in any organization or entity with any financial interest (such as honoraria; educational grants; participation in speakers’ bureaus; membership, employment, consultancies, stock ownership, or other equity interest; and expert testimony or patent-licensing arrangements), or non-financial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

⁸ AP Images: www.apimages.com.

Table 12
Sample classified text according to the implemented Language Model.

<p>Least associated to economic topics: ‘Ed Whitacre, chairman and CEO of SBC Communications Inc., jokes with the San Antonio Spurs mascot at the SBC Center ...’ ‘Country music recording artist Kellie Pickler is seen at Second Harvest Food Bank of Middle Tennessee during the ConAgra Foods Foundation Hunger-Free Summer Kick Off ...’ ‘Ryan Scott, Market & Rye executive chef and former Top Chef contestant, speaks to California Culinary Academy students at a McDonald’s culinary event ...’ ‘Mary West, right, part-owner of Room Service, hugs jockey Shaun Bridgmohan after he rode Room Service to a dead heat with Rosalind...’ ‘Fans are shown at Road America in Elkhart Lake, Wisc., during the NASCAR Nationwide series ...’</p>
<p>Most associated to economic topics: ‘Standing left to right, Director of Global Health Initiative USA Christopher Murray, Norway’s Prime Minister Jens Stoltenberg, Chairman of Microsoft Corporation Bill Gates.....at the World Economic Forum in Davos, Switzerland’ ‘Witnesses representing the Citigroup Subprime-Related Structured Products and Risk Management operation testify as the Financial Crisis Inquiry Commission (FCIC)...’ ‘UNITED STATES - SEPTEMBER 20: From left, Sens. Elizabeth Warren, D-Mass.,at Wells Fargo featuring testimony by CEO John Stumpf...’ ‘German Labor Minister Olaf Scholz arrives for a meeting about the future of Opel at the Chancellery in Berlin, on Wednesday, May 27, 2009. General Motors Corp....’ ‘Russ Smyth from the U.S., left, President of McDonald’s Europe,during the presentation of the new McDonald’s marketing campaign...’</p>

Image processing

Package “fer”⁹ is used to carry out two tasks: face extraction and facial expression classification. The default option for face extraction is a ‘haar cascade model’ We selected the more accurate multitask cascade convolutional neural network (MTCNN). The task of emotion classification is performed by a different convolutional neural network. More specifically, this task is performed by the model implemented in Arriaga et al. (2017). This implementation involves model training using FER-2013 dataset. This is a collection of 35,887 grayscale images. Each image is classified as: ‘angry’, ‘disgust’, ‘fear’, ‘happy’, ‘sad’, ‘surprise’ or ‘neutral’. The trained model achieves an accuracy of 66% which is, approximately, human level accuracy as documented in Goodfellow et al. (2013). We apply the command ‘detect_emotion’ to each selected image to retrieve two elements for each identified face. The first element indicates the location and the size (in pixels) of each identified face. The second element is a vector with 7 elements indicating the probability of each facial expression according to the trained model: $r = \{r_{anger}, \dots, r_{surprise}\}$.

References

- Albaga, A., Ardia, D., Bluteau, K., Borms, S., Boudt, K., 2020. Econometrics meets sentiment: an overview of methodology and applications. *J. Econ. Surv.* 34 (3), 512–547.
- Ardia, D., Bluteau, K., Boudt, K., 2019. Questioning the news about economic growth: sparse forecasting using thousands of news-based sentiment values. *Int. J. Forecast.* 35 (4), 1370–1386.
- Arkes, H., Herren, L., Isen, A., 1988. The role of potential loss in the influence of affect on risk taking behavior. *Organ. Behav. Hum. Decis. Process.* 42, 181–193.
- Arriaga, O., Valdenegro-Toro, M., Plöger, P., 2017. Real-time Convolutional Neural Networks for Emotion and Gender Classification. arXiv preprint arXiv:1710.07557.
- Audrino, F., Sigrist, F., Ballinari, D., 2020. The impact of sentiment and attention measures on stock market volatility. *Int. J. Forecast.* 36 (2), 334–357.
- Bao, W., Che, H., Zhang, J., 2020. Willgo at Semeval-2020 Task 3: an Accurate Model for Predicting the (Graded) Effect of Context in Wordsimilarity Based on Bert. <https://arxiv.org/pdf/2005.01006.pdf>.
- Barbaglia, L., Consoli, S., Manzan, S., 2020. Forecasting with Economic News.
- Bénabou, R., Tirole, J., 2016. Mindful economics: the production, consumption, and value of beliefs. *J. Econ. Perspect.* 30 (3), 141–164.
- Blankespoor, E., Hendricks, B., Miller, G., 2017. Perceptions and price: evidence from CEO presentations at IPO roadshows. *J. Account. Res.* 55 (2), 275–327.
- Blinder, A., Krueger, A., 2004. What Does the Public Know about Economic Policy, and How Does it Know it? Technical Report. National Bureau of Economic Research.
- Bloom, N., Baker, S., Davis, S., 2016. Measuring economic policy uncertainty. *Q. J. Econ.* 131 (4), 1593–1636.
- Boxell, L., 2018. Slanted Images: Measuring Nonverbal Media Bias. University Library of Munich, Germany. MPRA Paper 89047 <https://ideas.repec.org/p/pru/mprapa/89047.html>.
- Damasio, A., 1994. *Descartes’ Error: Emotion, Reason and the Human Brain*. Putnam, New York.
- Devlin, J., Chang, M.-W., Lee, K., Toutanova, K., 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv preprint arXiv:1810.04805.
- Ellingsen, J., Larsen, V., Thorsrud, L., 2020. News Media vs. FRED-MD for Macroeconomic Forecasting. CESifo Working Paper 8639 <https://ssrn.com/abstract=3720396>.
- Faust, J., Gilchrist, S., Wright, J., Zakrajšek, E., 2013. Credit spreads as predictors of real-time economic activity: a Bayesian model-averaging approach. *Rev. Econ. Stat.* 95 (5), 1501–1519.
- Fenton-O’Creevy, M., Lins, J.T., Vohra, S., Richards, D.W., Davies, G., Schaaff, K., 2012. Emotion regulation and trader expertise: heart rate variability on the trading floor. *Journal of Neuroscience, Psychology, and Economics* 5 (4), 227–237.
- Fiske, S., Taylor, S., 1991. *Social Cognition*. McGraw-Hill Company.
- Gentzkow, M., Kelly, B., Taddy, M., 2019. Text as data. *J. Econ. Lit.* 57 (3), 535–574.
- Goodfellow, I.J., Erhan, D., Carrier, P.L., Courville, A., Mirza, M., Hamner, B., Cukierski, W., Tang, Y., Thaler, D., Lee, D.-H., et al., 2013. Challenges in representation learning: a report on three machine learning contests. In: *International Conference on Neural Information Processing*. Springer, pp. 117–124.
- Graham, J., Harvey, C.R., Puri, M., 2017. A corporate beauty contest. *Manag. Sci.* 63 (9), 3044–3056.
- Hatfield, E., Cacioppo, J., Rapson, R., 1993. Emotional contagion. *Curr. Dir. Psychol. Sci.* 2, 96–100.
- Hirshleifer, D., 2020. Presidential Address: social transmission bias in economics and finance. *J. Finance* 75.
- Hu, A., Ma, S., 2020. Human Interactions and Financial Investment: A Video-Based Approach.
- Huang, X., Ivkovic, Z., Jiang, J., Wang, I., 2018. Swimming with the sharks: entrepreneurial investing decisions and first impression. https://neeley.tcu.edu/uploadedFiles/LKCM_Center/Conference/SwimmingWithTheSharks.pdf.
- Johnson, E., Tversky, A., 1983. Affect, generalization, and the perception of risk. *J. Pers. Soc. Psychol.* 45 (1), 20–31.
- Joo, J., Steinert-Threlkeld, Z., 2018. Image as data: automated visual content analysis for political science. <https://arxiv.org/abs/1810.01544v1>.
- Khan, A., Shah, Q., Uddin, M., Ullah, F., Alharbi, A., Alyami, H., Gul, M., 2020. Sentence Embedding Based Semantic Clustering Approach for Discussion Thread Summarization. *Complexity*. Article ID 4750871.
- Kunda, Z., 1999. *Making Sense of People*. MIT Press.

⁹ <https://pypi.org/project/fer/>.

- Lerner, J., Li, Y., Valdesolo, P., Kassam, K., 2015. Emotion and decision making. *Annu. Rev. Psychol.* 66, 799–823.
- Liang, Q., Sun, Q., Li, W., Yu, F., 2021. Media effects matter: macroeconomic announcements in the gold futures market. *Econ. Modell.* 96, 1–12. <https://www.sciencedirect.com/science/article/pii/S0264999320312839>.
- Little, A., Burriss, R., Jones, B., Roberts, S., 2007. Facial appearance affects voting decisions. *Evol. Hum. Behav.* 28 (1), 18–27.
- Mazumder, S., 2021. The reaction of inflation forecasts to news about the fed. *Econ. Modell.* 94, 256–264. <https://www.sciencedirect.com/science/article/pii/S0264999320312104>.
- McCracken, M., Ng, S., 2020. Fred-QD: A Quarterly Database for Macroeconomic Research.
- Pástor, L., Veronesi, P., 2005. Rational IPO waves. *J. Finance* 4, 1713–1757.
- Patton, A., Sheppard, K., 2007. Evaluating volatility forecasts. In: Andersen, T.G., Davis, R.A., Kreiss, J.P., Mikosch, T. (Eds.), *Handbook of Financial Time Series*. Springer-Verlag, Berlin.
- Peng, Y., 2018. Same candidates, different faces: uncovering media bias in visual portrayals of presidential candidates with computer vision. *J. Commun.* 68 (5), 920–941, <https://doi.org/10.1093/joc/jqy041>.
- Re, D., Hunter, D., Coetzee, V., Tiddeman, B., Xiao, D., DeBruine, L., Jones, B., Perrett, D., 2013. Looking like a leader—facial shape predicts perceived height and leadership ability. *PloS One* 12 (8), 1–10.
- Rojas, Q., Masip, D., Todorov, A., Vitria, J., 2011. Automatic prediction of facial trait judgments: appearance vs. structural models. *PloS One* 6 (8), 373–380.
- Smales, L., Apergis, N., 2017. Understanding the impact of monetary policy announcements: the importance of language and surprises. *J. Bank. Finance* 80, 33–50.
- Smith, A., 1759. *The Theory of Moral Sentiment*. T. Longman, London.
- Tetlock, P., 2007. Giving content to investor sentiment: the role of media in the stock market. *J. Finance* 62, 1139–1168.
- Todorov, A., Oosterhof, N., 2011. Modeling social perception of faces. *Soc. Sci.* 104 (3), 535–559.
- Wan, C., Jiang, S., Wang, C., Yuan, Y., Wang, C., 2020. A novel sentence embedding based topic detection method for microblogs. *IEEE Access* 8, 202980–202992.
- Wu, Y., Zhao, S., Li, W., 2020. Phrase2vec: phrase embedding based on parsing. *Inf. Sci.* 517, 100–127.
- Zhang, K., Zhang, Z., Li, Z., Qiao, Y., 2016. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.* 23 (10), 1499–1503.