

Sentiment fuels fear: on the transmission of social mood to market volatility



Stefano Borgioli | European Central Bank (ECB)

Giampiero M. Gallo | The Italian Court of Audits & New York University in Florence

Chiara Ongari | European Central Bank (ECB)

Keywords: VAR, Granger Causality, sentiment analysis, financial market, forecasting

JEL codes: C1, C32, C53, G4

Abstract

In 1936, John Maynard Keynes proposed that emotions and instincts are pivotal in decision-making, particularly for investors. Both positive and negative moods can influence judgments and decisions, extending to economic and financial choices. Measuring mood or sentiment is challenging, but surveys and new data collection methods offer some solutions. Recently, the availability of web data, including search engine queries and social media activity, has provided high-frequency sentiment measures. For example, the Italian National Statistical Institute's Social Mood on Economy Index (SMEI) uses X (formerly Twitter) data to assess economic sentiment in Italy. The relationship between SMEI and financial market activity, specifically the FTSE MIB index and its volatility, is examined using a trivariate Vector Autoregressive model, taking into account the impact of the COVID-19 pandemic.

Disclaimer: This policy brief is based on [ECB Working Paper series, No 2999](#) "Financial returns, sentiment and market volatility. A dynamic assessment". The views expressed are those of the authors and do not necessarily reflect those of the ECB, nor of the Italian Court of Audits.

Introduction

As early as 1936, John Maynard Keynes introduced the concept that emotions and instincts ("*animal spirits*") play a pivotal role in decision-making, especially among investors.

Moods—whether positive or negative—profoundly influence judgment and decision-making, often triggered by unrelated events. This dynamic extends to economic and financial decisions as well, as Nobel laureate Daniel Kahneman described, trying to understand how intuition, emotional states, and biases shape our judgments, thoughts, behaviours, and choices. Information alone does not dictate outcomes, and rationality is not always the primary driver of our actions.

While the theoretical rationale is well-established, measuring collective mood or sentiment presents notable challenges. To isolate an aggregate measure, various approaches have been proposed in the literature. One option is to conduct surveys and collect data about how economic agents judge the evolution of the general economic conditions or their own.

More recently, the widespread availability of online data has inspired new sentiment indicators derived from search engine activity. This approach allows researchers to examine both sentiment levels (akin to confidence) and sentiment dynamics (mood swings), with the significant advantage of providing data at a daily frequency—a critical feature for analysing financial markets.

X (formerly Twitter) seems to be the natural outlet for this expression of sentiments, helped by a large number of single messages, the possibility of replying to one another, and of classifying the content using “tags”.

Financial market activity can be viewed as a reflection of (collective) beliefs and sentiments, shaping equilibrium prices and returns through trading; by the same token, market volatility (i.e. the variability of returns) can be seen as inversely related to the consensus on how information reaching the market points to the evolution of the market itself. Typically, a downturn in the market is a reaction to bad news and is characterized by high volatility. Since the early 1990s, a market-based measure of volatility extracted from the implied volatilities of put and call options (at the money- 30 days to expiration) on a market index came to be known as the “fear and greed index” (the VIX is derived from the S&P500, Whaley, 1993, but other option-based volatility indices are available).

The data


The “Social Mood on Economy Index” (SMEI) is an experimental index introduced by ISTAT (Italian National Institute of Statistics) in October 2018, with daily data available as from February 10, 2016. It provides a daily measure of Italian public sentiment on the economy, derived from real-time analysis of approximately 26,000 tweets per day. These tweets are selected based on a predefined set of keywords that were designed by subject-matter experts. The SMEI measure of daily sentiment swings rather wildly, while a smoother pattern is achieved by a monthly moving average. Oscillations in both directions alternate until the absolute minimum of the indicator coinciding with the first COVID-19 lockdown occurred in March 2020. Other minima are then found in October 2020, when the second lockdown was announced, and in February 2022, following the Russian aggression to Ukraine and the ensuing banking and economic sanctions. Starting from the raw daily data of the SMEI, a trend is extracted upon removing two seasonal components. Despite the limitations of this Twitter-based indicator (e.g. SMEI does not produce a representative sample, neither of the whole Italian population nor of the FTSE MIB investors), we follow similar studies in using it as a proxy to measure sentiment swings.

The primary benchmark Index for the Italian equity markets is the FTSE MIB, measuring the performance of 40 Italian stocks (approximately 80% of the domestic market capitalization). The FTSE Implied Volatility Index (FTSE IVI) is a series of end-of-day mean volatility, derived from the at-the-money put and call implied volatilities on the FTSE MIB index options (indices for 30, 60, 90 and 180 day implied volatility estimates). Looking closely to the FTSE (30 days)

IVI, two remarkable volatility peaks happen on 18 March 2020 and 07 March 2022, in conjunction with the outbreak of the COVID-19 pandemic in Western Europe and with the Russian aggression against Ukraine respectively.

Next to the FTSE (30 days) IVI, two alternative volatility measures can be calculated (Parkinson volatility estimator, VP and Garman-Klass volatility estimator, GK), using four commonly available intra daily prices (Open, High, Low, and Close) for the index. Both are end-of-day measures of the volatility, and they share much information with the FTSE (30-day) IVI. Their features are graphically summarized in Figure 1.

Figure 1. Bivariate scatter plots of the three volatility measures of the FTSE MIB (45-degree)




Two key remarks emerge. First, both indicators closely resemble the pattern of the FTSE (30-day) IVI, with the GK measure showing a particularly strong linear relationship (high positive correlation, 0.65, with the FTSE (30-day) IVI, compared to 0.62 for the Parkinson measure). This strong correlation allows these indicators to serve as proxies for the FTSE (30-day) IVI. Second, the Garman-Klass and Parkinson volatility indices are highly correlated, with a coefficient of 0.93 over the sample period. Given their similarity, the models are applied to the Garman-Klass volatility, expressed in annualized percentage terms.

Some empirical evidence on the relationship between the SMEI and the FTSE-MIB

Figure 2 presents the daily observations of GK volatility alongside the corresponding values of the SMEI and its trend. It is evident that negative peaks in the SMEI almost always align with positive peaks in GK volatility, suggesting a potential relationship between these two components. This interplay indicates that they may influence each other, possibly reflecting a lower sentiment among the Italian population during periods of heightened uncertainty in the Italian equity markets, as represented by the MIB.

The figure, supported by the data (for more information see ECB WP 2999), reveals also a negative correlation between the FTSE-MIB closing prices and both the daily values and trend of the SMEI for the overall period. To shed light on the dynamic relationship, we then perform a Granger causality test to understand possible predictive relationships between the SMEI series and the return and the volatility of the MIB.

Figure 2. GK and SMEI, sample from 10/02/2016 to 30/09/2023



Granger causality in-sample

The basis of the analysis is a Vector Auto-regressive Model (VAR) with the SMEI, the MIB return and the GK estimator for the MIB volatility, using five lags and a sample of daily observations from 10 February 2016 to 8 March 2020. The Granger causality test applies to the null hypothesis that lagged values of X have no predictive value for Y and viceversa (unidirectional or bidirectional). Our results clearly show how such an hypothesis is strongly rejected when the dependent variable used is the Garman-Klass volatility of the FTSE-MIB. In particular, the lagged return of the MIB (as expected) alone, and both the lagged return of the MIB and SMEI when put together, have a strong influence on the current Garman-Klass volatility.

When looking at the other two dependent variables, the null hypothesis of no Granger causality cannot be rejected and we could say that the SMEI receives a certain influence by factors as the return of the MIB and its volatility but there are probably also other major external factors contributing to its behaviour.

We also performed the Granger causality test on the same variables and models, but restricted to the period between 9 March 2020 and 30 September, 2023, i.e. the COVID-19 and post-COVID-19 period. Please refer to [ECB WP 2999](#) for the detailed regression tables.


In this period, characterised by a unique emergency with major impact on the economic landscape and on the life of the Italian population, we see results that differ from the ones analysed before. The null hypothesis of no Granger causality can be rejected in most of the cases, and it is interesting to observe how the joint impact of the lagged SMEI and GK volatility “Granger causes” the return of the MIB. The main result we focus on is the explainability and predictability of the volatility of the MIB using the lagged SMEI and the return of the MIB.

Out of sample- Diebold-Mariano forecast comparison test

To complete our analysis, we performed an extension of the Granger causality test to ascertain whether extra information is valuable in an out-of-sample framework. The question is then whether the forecasts for a variable produced by a simple AR with five lags, that is, using its past can be significantly outperformed by the forecasts obtained using the corresponding equation of a VAR model that includes additional variables (with the same number of lags). The comparison is done by the Diebold-Mariano (DM) test, where the null hypothesis is one of equal performance of the two sets of forecasts according to a simple loss function, say, the Mean Square Error (MSE) or the Mean Absolute Error (MAE). The forecasts are generated recursively, by fixing the initial parameter estimation period between Feb. 10, 2016, and Mar. 8, 2020 (corresponding to the same 1087 observations used in-sample before), and producing one-step ahead results for 66 periods (approximately, three months) using the historical values for the lagged variables in the models. The fixed window is then moved forward by 66 periods, keeping 1087 observations for the new estimation (until Jun. 27, 2023), and 66 for projection (until Sep. 29, 2023). Refer to [ECB WP 2999](#) for the detailed Diebold Mariano (DM) test statistics.

The evidence somewhat complements the outcome of the in-sample analysis: out of the three variables, only GK volatility benefits from the extended information set; in our setup, the two models can be considered equivalent for the other two variables, with the interpretation that only their own past should be considered as relevant.

Figure 3. GK volatility versus its one-step ahead VAR forecasts generated on windows of 66 days before re-estimating



Source: Eurostat, own calculations.

Projecting the behaviour of volatility from the VAR can be appreciated graphically between the end of May 2020 and the end of September 2023, as in Figure 3. Except for the burst of volatility corresponding to the Russian aggression in Ukraine in February 2022, the profile of the one-step ahead forecast follows the actual values rather closely.

Concluding remarks

The widespread participation of users in social forums like X (formerly Twitter) presents a significant challenge in validating the informational accuracy of individual messages. To address this, various studies have focused on identifying relevant keywords and consolidating them into indices, designed to track prevailing sentiments over time, offering insights into the economic environment. The daily Social Mood on Economy Index (SMEI) produced by the Italian National Institute of Statistics (ISTAT) is such an example.

In this work, we investigated some properties of the SMEI in its relationship with the performance of the Milan Stock Exchange, as represented by the daily FTSE-MIB index, both for returns and volatility. The research question focuses on examining the relationship between these two time series and the SMEI within a stationary VAR model, utilizing an in-sample Granger causality test to uncover any potential causal links. To account for the significant impact of the COVID-19 pandemic, we divided the full sample period (February 10, 2016, to September 30, 2023) into two sub-samples. The pandemic, which emerged in March 2020, triggered unprecedented emergency measures that profoundly disrupted regular economic and social activities, reshaping perceptions of uncertainty and risk over an extended period.

The results show that for the pre COVID-19 period, the only variable significantly being affected by lagged values of other variables is the volatility singularly for the returns, presumably due to the so-called leverage effect by which past negative returns increase market volatility. By contrast, when the second sub-sample is considered, we notice that, for single variable tests, lagged volatility Granger-causes both SMEI and returns, lagged returns affect volatility, and lagged SMEI this time affects volatility. The picture given, therefore, is one in which the pandemic turns out to significantly change the market activity as represented by dynamic relationships among the variables considered.

The question is also addressed dynamically in an out-of-sample context, whereby we resort to the well-known Diebold-Mariano test of superior predictive ability, holding a univariate autoregression as the benchmark. The framework we built is one in which we use rolling regressions, holding an estimation sample to a window of 1087 observations, producing 66 one-step ahead forecasts with both the univariate and the VAR models. In this case, the output shows that the only variable for which the VAR is predictively superior to the AR model is the range-based volatility, indicating that both lagged SMEI and returns possess valuable information for forecasting market activity turbulence.

About the author(s)

Stefano Borgioli is a Senior Team Lead in the Banking Supervision Data Division of the Directorate General Statistics of the European Central Bank (ECB). He joined the ECB in April 1999, after a 3-year working experience as an economist in the Macro-Prudential Department of Banca d'Italia. He had temporary stints in both the Directorate General Macroeconomic Policy & Financial Stability and the Directorate General Economics of the ECB. He has worked extensively on banking statistics, statistics for macroprudential analysis and financial integration. He published in these areas and he frequently gives lectures at Universities (Politecnico di Milano, University of Florence, University of Frankfurt).

Giampiero M. Gallo (Ph.D. in Economics - UPenn) Judge in the Italian Court of Audits, Part-time Professor of Economics at NYU in Florence, Fellow of the Society for Financial Econometrics, former Professor of Econometrics at the Università di Firenze and former President of the Italian Econometric Association (SIDE-IEA). Research interests include financial volatility, financial crises, macroeconomic forecasting, and public debt sustainability, with publications in top journals. He is married with two children, an avid urban biker, and a collector of Le Petit Prince editions.

Chiara Ongari is a Supervision Analyst at the European Central Bank (ECB), with an academic background in mathematics, stochastic processes, and econometrics. She joined the ECB nearly four years ago, initially gaining experience in the areas of statistics and financial stability before transitioning to the Supervisory Technology Team. Her work focuses on cutting-edge projects, specializing in textual analysis and large language models (LLMs).

SUERF Policy Notes and Briefs disseminate SUERF Members' economic research, policy-oriented analyses, and views. They analyze relevant developments, address challenges and propose solutions to current monetary, financial and macroeconomic themes. The style is analytical yet non-technical, facilitating interaction and the exchange of ideas between researchers, policy makers and financial practitioners.

SUERF Policy Notes and Briefs are accessible to the public free of charge at <https://www.suerf.org/publications/suerf-policy-notes-and-briefs/>.

The views expressed are those of the authors and not necessarily those of the institutions the authors are affiliated with.

© SUERF – The European Money and Finance Forum. Reproduction or translation for educational and non-commercial purposes is permitted provided that the source is acknowledged.

Editorial Board: Ernest Gnan, David T. Llewellyn, Donato Masciandaro, Natacha Valla

Designed by the Information Management and Services Division of the Oesterreichische Nationalbank (OeNB)

SUERF Secretariat
c/o OeNB, Otto-Wagner-Platz 3A-1090 Vienna, Austria
Phone: +43 1 40 420 7206
E-Mail: suerf@oenb.at
Website: <https://www.suerf.org/>