



Learning about unprecedented events: Agent-based modelling and the stock market impact of COVID-19

Daive Bazzana^{a,b}, Michele Colturato^c, Roberto Savona^{d,*}

^a Department of Economics and Management, University of Brescia, via San Faustino 74/b, 25122 Brescia, Italy

^b Fondazione Eni Enrico Mattei, Corso Magenta, 63, 20123 Milan, Italy

^c Department of Mathematics, University of Pavia, Via Ferrata, 5, 27100 Pavia, Italy

^d Department of Economics and Management, University of Brescia, C/da S. Chiara 50, 25122 Brescia, Italy

ARTICLE INFO

JEL Classification:

G11
G12
G14
C63

Keywords:

Agent-based model
Representativeness
Unprecedented events
Asset pricing model
Heterogeneous expectations

ABSTRACT

We model the learning process of market traders during the unprecedented COVID-19 event. We introduce a behavioural heterogeneous agents' model with bounded rationality by including a correction mechanism through representativeness (Gennaioli et al., 2015). To inspect the market crash induced by the pandemic, we calibrate the STOXX Europe 600 Index, when stock markets suffered from the greatest single-day percentage drop ever. Once the extreme event materializes, agents tend to be more sensitive to all positive and negative news, subsequently moving on to close-to-rational. We find that the deflation mechanism of less representative news seems to disappear after the extreme event.

1. Introduction

The COVID-19 pandemic has had such an unprecedented large-scale impact on stock markets as to represent a natural experiment to explore how economic agents react to unknown events without any historical episode providing useful insights. Recent papers inspected stock market dynamics and panics along the pandemic evolution and government restriction measures implemented to counteract COVID-19 (e.g., Alfaro et al., 2020; Baker et al., 2020; Ramelli and Wagner, 2020; Aggarwal et al., 2021; Scherf et al., 2022; Yu and Xiao, 2023). We complement this literature by focusing on the description of what may have happened inside the mind of market traders when an unprecedented and unknown event, such as the COVID-19, materializes.

To do this, we introduce an agent-based model (ABM) for the learning and decision-making process of stock market traders when catastrophic and unprecedented events occur within a behavioural heterogeneous agents' context with bounded rationality (Brock and Hommes, 1997; Anufriev and Hommes, 2012). Agents process news according to the Bayesian posterior correction mechanism through the representativeness introduced in Gennaioli et al. (2015). This is a key assumption in our model. Indeed, when mental shifts take place in agents because the probability of the less representative state suddenly becomes more representative, they overreact by reassessing their price expectation proportionally to the variation of the corrected posterior probability of the most representative state. Speed of change in agents' confidence depends on: (1) how fast good and bad news accumulates over time; (2) how far agents look back in time when processing information before taking their investment decisions (memory); (3) how much importance agents

* Corresponding author.

E-mail address: roberto.savona@unibs.it (R. Savona).

attach to news (weight). By inspecting the stock market crash induced by the pandemic in February-March 2020, we show that our model matches the pattern of the STOXX Europe 600 Index exhibited before and after the big shock.

The proposed framework has two main objectives. First, it is intended to replicate empirically-observed stylized facts in financial markets both considering potentially tranquil and crisis periods, then including unprecedented tail events as a special case: the price expectation mechanism is indeed based on the disproportionate weight attached to the event arrival and the resulting jump size, which reflects the catastrophic potential perceived and evaluated by agents. In this perspective, our analysis of the stock market impact induced by the COVID-19 is a natural laboratory we explore to understand how economic agents react to unknown events without any historical episode providing useful insights.

Second, our model intends to artificially derive the decision-making mechanism followed by the traders while preserving a reasonable fit to the data. In a sense, we create a tension between an *abstract* modelling (Gilbert, 2008), to demonstrate the micro-founded (unobservable) learning and decision-making process assumed by stock market traders, and a *facsimile* modelling (Gilbert, 2008), to validate the simulated model by minimizing the mismatch between observed and simulated data. In this perspective, our paper also innovates on the calibration and validation procedure, as we provide time-varying parameter's estimation, which allows to dynamically inform the ABM with the changing market conditions.

2. Market price dynamics

The model setup follows market equilibrium dynamics with heterogeneous beliefs as in Brock and Hommes (1997). A population of $\sum_j = J$ heterogeneous and bounded rational traders invest their wealth in one risk-free asset and in one risky asset. They are myopic mean variance maximizer trading in a discrete-time market solving the following problem:

$$\max_s E_{h,t}(W_{j,t+1}) - \frac{a}{2} V_{h,t}(W_{j,t+1}), \tag{1}$$

where $s_{j,t}$ is the demand/supply for risky asset share, a denotes the risk aversion parameter assumed as being equal for all agents; $E_{h,t}(W_{j,t+1})$ and $V_{h,t}$ denote the conditional expectation and conditional variance of tomorrow's wealth based upon the informational set available at time t of the agent j following the trading rule h (Brock and Hommes, 1998). The wealth dynamics W is as follows:

$$W_{j,t+1} = RW_{j,t} + s_{j,t} W_{j,t} [E_{h,t}(D_{t+1}) + E_{h,t}(p_{t+1}) - Rp_t], \tag{2}$$

where R is the gross return paid by the risk-free asset with $R = 1 + r$ and r is the constant risk-free rate of return; p_t is the ex-dividend price of the risky asset at time t . The terms within the square brackets denote the risk premium, where $E_{h,t}(D_{t+1})$ is the expectation at time t for traders following the h -th trading rule of tomorrow's dividend D_{t+1} which, in turn, is assumed to be exogenous and deterministic (Bottazzi et al., 2015); $E_{h,t}(p_{t+1})$ is tomorrow's expected price p_{t+1} of the risky asset predicted by all investors following the h -th trading rule.

Investors form their expectations on future price dynamics by extrapolating past prices using the following h heuristics (Anufriev and Hommes 2012; Bao et al., 2017):

$$E_{NAI,t}(p_{t+1}) = p_t; \tag{3}$$

$$E_{BIA,t}(p_{t+1}) = p_t(1 \pm b); \tag{4}$$

$$E_{WCH,t}(p_{t+1}) = p_t \pm \alpha(p_t - p_{t-1}); \tag{5}$$

$$E_{SCH,t}(p_{t+1}) = p_t \pm \beta(p_t - p_{t-1}), \tag{6}$$

where b represents a bias, and α and β are trend extrapolation coefficients with $|\beta| > \alpha$.

Equation (3) describes the naive heuristic (NAI) in which expectations are formed using the last observed price. Equation (4) represents biased (BIA) expectations as in Brock and Hommes (1998): agents are optimistic (pessimistic) about future prices which follow an in(de)creasing path (positive/negative b). Equations (5) and (6) describe chartist strategies where the extrapolation coefficients α and β measure the strength of the adjustment (weak or strong, WCH or SCH respectively): positive coefficients are with weak and strong trend followers; contrarians are with negative coefficients; trend followers (contrarians) make buy (sell) orders when the price trend is positive and vice versa (Schmitt and Westerhoff, 2021).

Once traders have formed their expectations about the future price, they submit their orders. The market price for the risky asset is set by market clearing, i.e., when demand equals supply:

$$\sum_{j=1}^J s_{j,t} = 0. \tag{7}$$

The price p_t at time t depends on h -based price average predictions $E_{h,t-1}(p_t)$ and the fundamental price $p_{f,t}$ as follows:

$$p_t = p_{f,t} + \frac{1}{R} (E_{h,t-1}(p_t) - p_{f,t}) + \varepsilon_t \tag{8}$$

where $p_{f,t} = \frac{D_t}{r}$ and ε_t is IID noise term. As Hommes et al. (2021) point out, the price equation for p_t has rational bubble solutions, with $(E_{h,t-1}(p_t) - p_{f,t})$ growing at the risk-free rate r , although these rational bubbles are often excluded by imposing transversality conditions.

At the end of the period, agents update their price expectations – also deciding to maintain or switch their heuristics – based on the following two steps:

First, they assess the prediction ability of the selected heuristic and decide to select the forecasting strategies that have performed better in the recent past according to the following discrete choice model (Hommes et al., 2005):

$$q_{h,t+1} = (1 - v) \frac{e^{\vartheta U_{h,t}}}{\sum_{h=1}^H e^{\vartheta U_{h,t}}} + v q_{h,t}, \tag{9}$$

where q_h is the share of investors following the h -th trading rule, $\vartheta \geq 0$ represents the choice intensity, $U_{h,t} = [p_t - E_{h,t-1}(p_t)]^2$ is the forecasting error of the h -th trading rule, and v is a persistence parameter introducing a status quo stickiness (Kahneman et al., 1991).

Second, they incorporate new information signals and update price expectations based on their expected posterior probability ($\pi_{x,t}$) about the state of the market:

$$\pi_{x,t} = \frac{\pi_{x,t-1} + n_{x,t}}{1 + n_{x,t} + n_{-x,t}}, \tag{10}$$

where x denotes the state of the market which is good g or bad b ; $n_{x,t}$ is the number of the x -th type news at time t with $n_{x,t} = \sum_{c=0}^m \varphi_i n_{x,t-c}$ where $\varphi_i > 0$ is the weight assigned to each type of news and m is the memory of past news that agents accumulate up to time $t - 1$ when updating their expected posterior probability about the state of the market. To make the intuition behind the mechanism as clear as possible, we focus on the simplest case where we have only two types of news, market activity-based (mkt) and public ($publ$), therefore $i = mkt, publ$. We assume that some information is perceived as more important than others, thus reflecting a higher weight. Market activity-based information measures are, (1) the stock price dynamics and, (2) the stock price volatility (standard deviation computed on the past m trading days). News is classified as good ($n_{g,t}$)/bad ($n_{b,t}$) whenever, (a) the change in the price occurred in the last period is positive/negative and, (b) the stock price volatility in the last period is decreasing/increasing. Public information is exogenous, timetabled and randomly good or bad; with this type of pseudo-random string of news, we rely on possibly scheduled announcements having powerful source of price movements, such as macroeconomic and political announcements (Birz and Lott, 2011; Baker et al., 2019) and central bank communications (Andrade and Ferroni, 2020).

Given new information at time t , every agent revises the posterior probability of the market state x (good/bad) by inflating the likelihood of the most representative state and deflating the less representative one. Contextualized within the financial market dynamics, representativeness (Kahneman and Tversky, 1972) leads investors to overweight the probability of events that have become more likely in light of recent news (Gennaioli et al., 2015; Bordalo et al., 2018). Therefore, after a period of good news, investors tend to judge positive future outcomes in an overly optimistic way while neglecting the bad ones; and vice versa.

In every period t , agents first compute representativeness R of state x as:

$$R_{x,t} = \frac{\pi_{x,t}}{\pi_{x,t-1}}, \tag{11}$$

and next they revise the posterior probability $\pi_{x,t}^p$ for market states based on the following rule:

$$\text{if } R_{x,t} > R_{-x,t} \rightarrow \pi_{x,t}^p = \frac{\pi_{x,t-1} + n_{x,t}}{(\pi_{x,t-1} + n_{x,t}) + \delta(\pi_{-x,t-1} + n_{-x,t})}; \pi_{-x,t}^p = \frac{\delta(\pi_{-x,t-1} + n_{-x,t})}{(\pi_{x,t-1} + n_{x,t}) + \delta(\pi_{-x,t-1} + n_{-x,t})}, \tag{12}$$

where δ is the discount factor that modulates the severity of the probability deflation with $0 \leq \delta \leq 1$.

For every time step t we have:

$$R_{x,t} > R_{-x,t} \Leftrightarrow \pi_{x,t} > \pi_{x,t-1} \Leftrightarrow \pi_{x,t} > \frac{n_{x,t}}{n_{x,t} + n_{-x,t}}, \tag{13}$$

with $\pi_{x,t}^p + \pi_{-x,t}^p = 1$.

Probabilities are dynamically updated with arriving strings of good and bad news. Representativeness causes changes in agents' beliefs when the less representative state suddenly becomes the more representative one. Therefore, the amount of good and bad news and the weights attached to it induces changes from one mood to another producing jumps in price expectations as follows:

$$\begin{aligned} E_{h,t}^p(p_{t+1}) &= [E_{h,t}(p_{t+1}) - \gamma_t] \text{ if } \pi_{g,t-1}^p > \pi_{b,t-1}^p \cup \pi_{g,t}^p < \pi_{b,t}^p \cup \frac{n_{b,t}}{n_{g,t}} > \frac{\pi_{b,t-1}^p + n_{b,t-1}}{\pi_{g,t-1}^p + n_{g,t-1}}; \\ &= [E_{h,t}(p_{t+1}) + \gamma_t] \text{ if } \pi_{b,t-1}^p > \pi_{g,t-1}^p \cup \pi_{b,t}^p < \pi_{g,t}^p \cup \frac{n_{g,t}}{n_{b,t}} > \frac{\pi_{g,t-1}^p + n_{g,t-1}}{\pi_{b,t-1}^p + n_{b,t-1}}; = E_{h,t}(p_{t+1}) \text{ otherwise,} \end{aligned} \tag{14}$$

Table 1
Parameters' value and initial conditions: Calibration.

Parameter	Label	Value
Number of periods	T	250
Number of traders	J	200
Risk-free rate	r	0.5%
Weak chartist coefficient	α	0.7
Strong chartist coefficient	β	1.3
Bias	b	2.5%
Risk aversion	a	20
Discount factor	δ	calibrated
Memory	m	calibrated
Number of market activity-based news per trade	n_{mkt}	2
Weight of market activity-based news	φ_{mkt}	1
Number of public news per trade	n_{publ}	2
Weight of public news	φ_{publ}	2
Weight of extreme event	φ_{ee}	calibrated
Jump amplitude	η_x	calibrated

Note: see Appendix B.

where $\gamma_t = \eta_x(\pi_{x,t}^p - \pi_{x,t-1}^p)$ with $\eta_x > 0$ measuring the perceived boom/bust, also allowing for possible asymmetric impacts with $\eta_b > < \eta_g$. Our framework includes unprecedented tail events as a special case, being modulated based on the magnitude of the jump size: $\eta_x > 0$ reflects the catastrophic potential perceived and evaluated by the agents. Note that the jump is proportional to change in the adjusted posterior probability of the most representative state. The stronger the surprise of the unexpected event, the greater the probability correction and, as a result, the price overreaction. Such a price overreaction modelling is consistent with the bimodal risk perception logic (McClelland et al., 1993), as the event that suddenly comes "on screen" can produce changes in behaviour and price, only if the associated risk probability crosses a certain threshold, in our case equal to 0.5.

To assess the consistency of the market dynamics, in Appendix A we run a numerical simulation showing how our model is able to reproduce many market stylized facts, such as bubbles and crashes, excess volatility, fat-tailed return distributions, uncorrelated price changes and volatility clustering.

3. Calibration

Having assessed the model's ability in replicating stylized facts in financial markets, we now explore the stock market impact induced by the COVID-19, as it represents a natural laboratory to explore how market traders react to unprecedented events. We do this by deriving their decision-making mechanism while minimizing a loss function that measures the distance between simulated and historical stock market index time series.

We calibrate the model using daily values of the STOXX Europe 600 Index over the period from 10 July 2019 to 30 June 2020, using $T = 250$ observations. The values assigned to the model's parameters are in Table 1. Since COVID-19 produces a structural break in the data, we split the time series into two sub-periods by fixing the time threshold Z , $1 < Z < T$, on 21 February 2020 ($Z = 160$), when international media brought the COVID-19 news "on screen" focusing on the Italian epidemic outbreak.¹ The extreme event arrival Z_{ee} is set on 12 March 2020 ($Z_{ee} = 173$), when stock markets experienced one of the highest losses ever.

We run 100 simulations of the model over the 250 trades, then calibrating: (a) the discount factor δ and the jump amplitude η_x on a daily basis; (b) m and τ (eq. 14) as constant estimates for pre- Z ($\hat{m}_A, \hat{\tau}_A$) and post- Z , ($\hat{m}_B, \hat{\tau}_B$); (c) the weight attached to the extreme event φ_{ee} (12 March 2020; $Z_{ee} = 173$).

Computationally, we follow the interior-point method (Pólik and Terlaky, 2010) to minimize the Euclidean distance between actual and predicted price (eq. 8); see Appendix B for technical details. As discussed in the introduction, the proposed calibration procedure is time-varying, then allowing to inspect how the model's parameters changed over time. As a result, our modelling goes through the use of data to initialize, build and update the model, and not only for validation purposes, which is similar in spirit to Hassan et al. (2008).

4. Results

Fig. 1 reports estimates for the price pattern p (Fig. 1a), the posterior probability of bad state π_b^p (Fig. 1b), δ (Fig. 1c) and η (Fig. 1d), whereas the value for m and τ ($p_{f,t}$ follows the process: $p_{f,t} = \alpha + \tau p_{f,t-1} + \varepsilon_t$; see Appendix B) the two sub-periods pre- Z (A) and post- Z (B) are:

$$\hat{m}_A = 9, \hat{m}_B = 7, \hat{\tau}_A = 0.0106, \hat{\tau}_B = 0.0983.$$

m shrinks from pre- to post-COVID outbreak, moving from 9 to 7 past days: agents cover a very short-term past window, which shortens even more when the extreme event materializes reflecting on high price volatility.

¹ <https://www.reuters.com/article/us-china-health-italy/coronavirus-outbreak-grows-in-northern-italy-16-cases-reported-in-one-day-idUSKBN20FOUI>.




Fig. 1. Calibration results

Note: The Figure reports results from calibration experiment for the STOXX Europe 600 index over the period from 10 July 2019 to 30 June 2020. We report average values computed on 100 simulations: grey area in Fig.1a denotes the min-max range. Vertical black line is the time threshold Z (21 February 2020), while vertical red line is the extreme event arrival Z_{ee} (12 March 2020).

$\hat{\tau}_A$ and $\hat{\tau}_B$ coefficients reflect the substantial flat fundamental price path before and after the pandemic outbreak with a coefficient close to zero before COVID outbreak, afterwards denoting a weak, but still low in level, coefficient.

The time-varying δ denotes 4 regimes. The first is from 10 July 2019 to 13 August 2019 with values around 0.75 on average: since

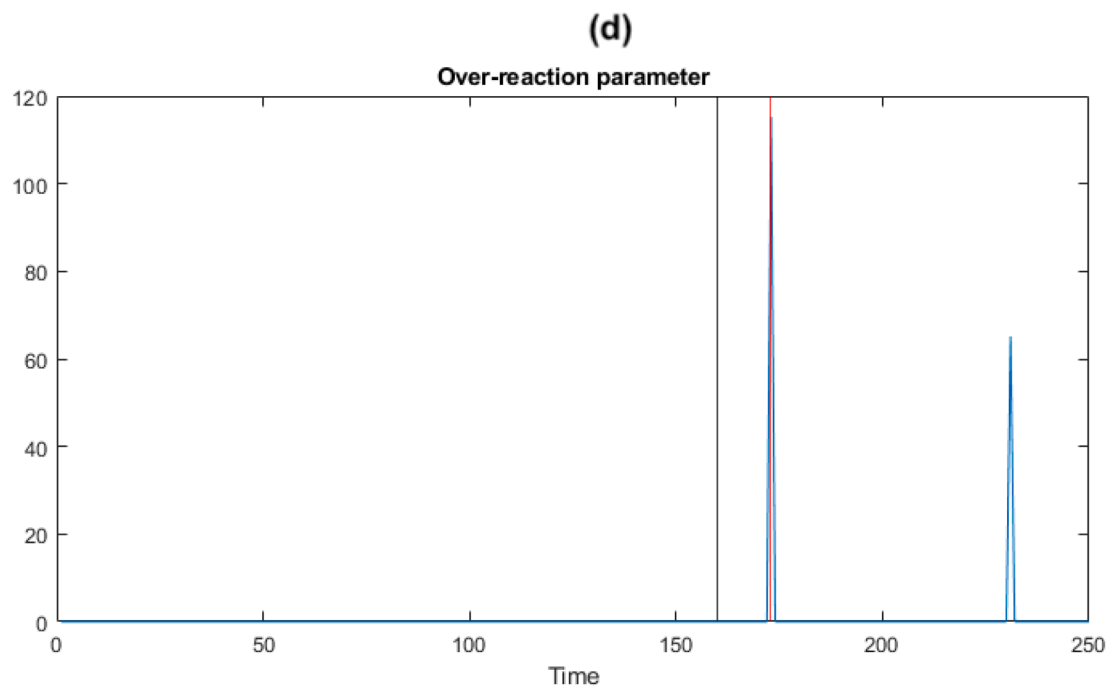
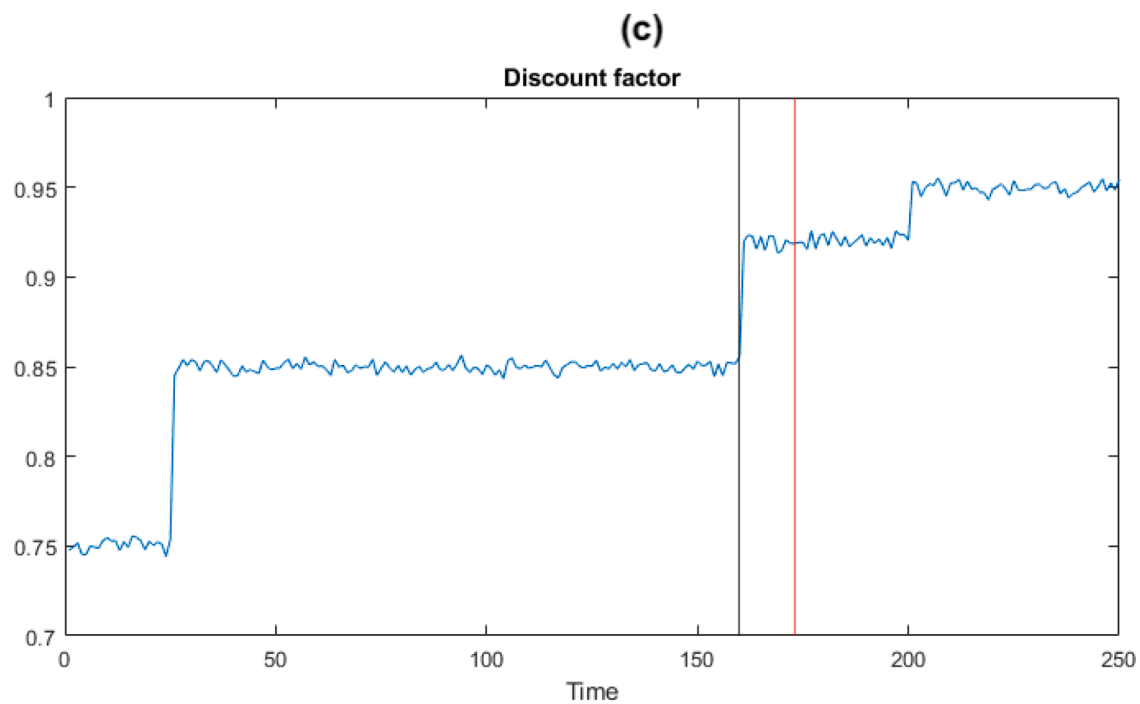


Fig. 1. (continued).

we are in a good state, agents tend to deflate bad news substantially. The second regime is from 14 August 2019 to 21 February 2020, when discount factors jump to 0.85 on average, thus deflating bad news. The third regime is from 22 February to 21 April 2020 and includes the extreme event, when bad state suddenly becomes the most representative, the price jump magnitude η_b is close to 120 and the weight of extreme event φ_{ee} is 950, which is equivalent to more than half a year of continuing daily bad news.² The discount factor moves around 0.92 on average, documenting how agents move towards a rational-based framework (good and bad news have the same weight in forming market beliefs). This tendency is reinforced in the fourth regime, from 22 April 2020 to the end of the period, when stock markets recovered substantially; the market shifts to a good state on 3 June 2020 when we have a good price jump around 60, which documents the asymmetric impacts with $\eta_b > \eta_g$. Afterwards, posterior probability remains lower than the threshold $(\pi_{b,t}^p < \frac{1}{2})$. Over this final regime, investors tend to hold close-to-rational expectations since δ is around 0.95 on average.

As a whole, our results, while confirming that COVID-19 impact on stock markets occurred about 3 weeks after the newspapers reported the pandemic arrival (Baker et al., 2020), indicate that agents maintained their beliefs even after the news was already circulated worldwide, since the posterior probability of a bad state was contained under 15 percent (see Fig. 1), while they started to process information by deflating the bad news less (discount factor is getting closer to 1). This phase seems characterized by an *unknown risk* perception (Peters and Slovic, 1996), for which the hazard of the extreme event is, in some sense, unobservable and unknown with harmful impacts judged as potentially substantial, but nevertheless delayed.

The extreme negative impact of COVID-19 seems to reflect a *dread risk* perception: the perceived lack of control and the catastrophic perspectives reflect a jump of the bad state posterior probability close to 100 percent.

As Fig. 1.a shows, the simulated prices are not tracking exactly to the empirical prices. As discussed in the introduction, our aim is not to get the best fit to the data *per se*, but rather to infer unobservable parameters describing the learning and decision process of single agents then aggregating into market dynamics. In this perspective, the quantitative calibration is intended to provide a formal (data-driven) validation of the conjectures proposed in our ABM by trading off between *abstract* and *facsimile* modelling (Gilbert, 2008) objecting functions.

5. Conclusion

This paper presents an agent-based model for the learning and decision-making process of stock market traders when catastrophic and unprecedented events materialize within a behavioural heterogeneous agent context with bounded rationality. Agents are assumed to shift from excessive optimism to excessive pessimism and vice versa, reflecting on large swings in agents' confidence and price expectations, through representative diagnostics.

The calibration experiment we ran on the STOXX Europe 600 index over the period from 10 July 2019 to 30 June 2020 to explore the anatomy of the COVID-19 impact offers two key insights. First, the extreme event impact on stock price is delayed and needs confirmation of bad news to align single risk perceptions. Second, once the extreme event occurred, agents are moving towards a rational-based framework.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.frl.2023.104085](https://doi.org/10.1016/j.frl.2023.104085).

References

- Alfaro, L., Chari, A., Greenland, A.N., Schott, P.K., 2020. Aggregate and firm-level stock returns during pandemics, in real time (No. w26950). Natl Bureau Econ. Res. <https://doi.org/10.3386/w26950>.
- Andrade, P., Ferroni, F., 2020. Delphic and odyssean monetary policy shocks: evidence from the euro area. J. Monet. Econ. <https://doi.org/10.1016/j.jmoneco.2020.06.002>.
- Anufriev, M., Hommes, C., 2012. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. Am. Econ. J. Microecon. 4, 35–64. <https://doi.org/10.1257/mic.4.4.35>.
- Baker, S.R., Bloom, N., Davis, S.J., Kost, K.J., 2019. Policy news and stock market volatility (No. w25720). Natl Bureau Econ. Res. <https://doi.org/10.3386/w25720>.
- Baker, S.R., Bloom, N., Davis, S.J., Kost, K.J., Sammon, M.C., Viratyosin, T., 2020. The unprecedented stock market impact of COVID-19 (No. w26945). Natl Bureau Econ. Res. <https://doi.org/10.3386/w26945>.

² Since the model generates 2 daily market activity-based information with weight of $\varphi_{mkt} = 1$, and 2 public information, each one with a weight of $\varphi_p = 2$, by assuming that all news in one day is bad, we have a total weight attached to bad news equal to 6 (see Table 1). Since the weight attached to the COVID-based extreme event is 950 (calibrated as explained in Appendix B), we can calculate that this extreme weight is equivalent to $\frac{950}{6} \cong 158$ days of continuing bad news.

- Bao, T., Hommes, C., Makarewicz, T., 2017. Bubble Formation and (In)Efficient Markets in Learning-to-forecast and optimise Experiments. *Econ. J.* 127, F581–F609. <https://doi.org/10.1111/ecoj.12341>.
- Birz, G., Lott, J.R., 2011. The effect of macroeconomic news on stock returns: new evidence from newspaper coverage. *J. Bank. Financ.* 35, 2791–2800. <https://doi.org/10.1016/j.jbankfin.2011.03.006>.
- Bordalo, P., Gennaioli, N., Shleifer, A., 2018. Diagnostic expectations and credit cycles. *J. Finance* 73, 199–227. <https://doi.org/10.1111/jofi.12586>.
- Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dynam. Contr.* 22, 1235–1274. [https://doi.org/10.1016/S0165-1889\(98\)00011-6](https://doi.org/10.1016/S0165-1889(98)00011-6).
- Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. *Econometrica* 65, 1059–1095. <https://doi.org/10.2307/2171879>.
- Gennaioli, N., Shleifer, A., Vishny, R., 2015. Neglected risks: the psychology of financial crises. *Am. Econ. Rev.* 105, 310–314. <https://doi.org/10.1257/aer.p20151091>.
- Gilbert, N. 2008. Agent-based models. Sage Publications Inc. URL <https://us.sagepub.com/en-us/nam/agent-based-models/book251732>.
- Hassan, S., Antunes, L., Pavon, J., Gilbert, N., 2008. Stepping on earth: a roadmap for data-driven agent-based modelling. In: *Proceedings of the 5th Conference of the European Social Simulation Association 2008*.
- Hommes, C., Kopányi-Peuker, A., Sonnemans, J., 2021. Bubbles, crashes and information contagion in large-group asset market experiments. *Exp. Econ.* 24, 414–433. <https://doi.org/10.1007/s10683-020-09664-w>.
- Kahneman, D., Knetsch, J.L., Thaler, R.H., 1991. Anomalies: the endowment effect, loss aversion, and status quo bias. *J. Econ. Perspect.* 5, 193–206. <https://doi.org/10.1257/jep.5.1.193>.
- Kahneman, D., Tversky, A., 1972. Subjective probability: a judgment of representativeness. *Cogn Psychol* 3, 430–454. [https://doi.org/10.1016/0010-0285\(72\)90016-3](https://doi.org/10.1016/0010-0285(72)90016-3).
- McClelland, G.H., Schulze, W.D., Coursey, D.L., 1993. Insurance for low-probability hazards: a bimodal response to unlikely events. *J. Risk Uncertainty* 7, 95–116. <https://doi.org/10.1007/BF01065317>.
- Peters, E., Slovic, P., 1996. The role of affect and worldviews as orienting dispositions in the perception and acceptance of nuclear power1. *J. Appl. Soc. Psychol.* 26, 1427–1453. <https://doi.org/10.1111/j.1559-1816.1996.tb00079.x>.
- Pólik, I., Terlaky, T., 2010. Interior point methods for nonlinear optimization. In: Bomze, I.M., Demyanov, V.F., Fletcher, R., Terlaky, T., Di Pillo, G., Schoen, F. (Eds.), *Nonlinear Optimization: Lectures given At the C.I.M.E. Summer School Held in Cetraro, Italy, July 1-7, 2007. Lecture Notes in Mathematics*. Springer, Berlin, Heidelberg, pp. 215–276. https://doi.org/10.1007/978-3-642-11339-0_4.
- Ramelli, S., Wagner, A.F., 2020. Feverish Stock Price Reactions to Covid-19 (SSRN Scholarly Paper No. ID 3560319). *Social Science Research Network, Rochester, NY*.
- Scherf, M., Matschke, X., Rieger, M.O., 2022. Stock market reactions to COVID-19 lockdown: a global analysis. *Finance Res. Lett.* 45, 102245 <https://doi.org/10.1016/j.frl.2021.102245>.
- Schmitt, N., Westerhoff, F., 2021. Pricking asset market bubbles. *Finance Res. Lett.* 38, 101441 <https://doi.org/10.1016/j.frl.2020.101441>.
- Yu, X., Xiao, K., 2023. COVID-19 Government restriction policy, COVID-19 vaccination and stock markets: evidence from a global perspective. *Finance Res. Lett.* 53, 103669 <https://doi.org/10.1016/j.frl.2023.103669>.