Learning about Unprecedented Events: Agent-Based Modelling and the Stock Market Impact of COVID-19

Davide Bazzana[*]

Michele Colturato[†]

Roberto Savona[‡]

This draft: September 14, 2021

Abstract

We model the learning process of market traders during the unprecedented COVID-19 event. We introduce a behavioral 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.

Keywords: Agent-Based Model, Representativeness, Unprecedented Events.

JEL Classification: G11, G12, G14, C63

[*] Department of Economics and Management, University of Brescia, via San Faustino 74/b, 25122 Brescia (Italy). E-mail: davide.bazzana@unibs.it.

[†] Department of Mathematics, University of Pavia, Via Ferrata, 5, 27100 Pavia (Italy). E-mail: michele.colturato@unipv.it.

[‡] *Corresponding author*: Roberto Savona, Department of Economics and Management, University of Brescia, C/da S. Chiara 50, 25122 Brescia (Italy). E-mail: roberto.savona@unibs.it.

"*If you want to know what will bring down this bull market, look at factors beyond old age. … Think too about what economists like to call "exogenous shocks". Coronavirus may fit into that category. It is easy to overreact, and markets tend to, but they can also be remarkably resilient and more often than not bounce back strongly.*"
(Simon Edelsten, February 26, 2020, Financial Times)

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 show how stock market dynamics followed pandemic evolution and government restriction measures were implemented to counteract COVID-19. Ramelli and Wagner (2020) identify phases of investor behavior along "incubation" (early January 2020 after cases of pneumonia detected in Wuhan, China, reported to the World Health Organization [WHO]) with no substantial stock price moves, "outbreak" (after January 20, when the WHO issued the first situation report) when stocks began suffering strongly, and "fever" (after February 24, when Italy implemented a strict lockdown) in which panic selling materialized, moving stock markets to one of the most dramatic crashes in history. Kumer Dey et al. (2020) find that COVID-19 cases and deaths made US investors panic and increased their anxiety, as reflected in Google search intensity for COVID-19, and Alfaro et al. (2020) show that unexpected changes in the trajectory of COVID-19 infections predict US stock returns. Baker et al. (2020) provide evidence that the unprecedented stock market reaction is much explained by the role of government restrictions imposed on social mobility and commercial activities.

The literature on the economic understanding of tail events enriched by research on the probability weighting function suggests that people tend to overestimate the likelihood of these events, also tending to overweight their decision-making process. As Barberis (2013) clarifies, such a decision-making process comprises a first step in which an individual assesses the probability of a tail event, and a second step, where a decision is made conditional on the probability the individual assigns to that event.

Tversky and Kahneman (1973) posit that a person evaluates probability by *availability*, i.e. by the ease with which relevant instances come to mind. The challenging question COVID-19 has brought to the fore is how people react to unprecedented and unknown tail events. Availability heuristic judgments should suggest a first underestimation of the likelihood of the event, until the point in which the worst-case scenario becomes vivid and easy to visualize, to then reflect on the decision-making process. This is consistent with experimental works indicating that people act in a bimodal setting (McClelland et al., 1993), paying little or no attention to events when the risk probability is below a certain threshold, while risks that suddenly come "on screen" produce significant changes in behavior. This is a judgmental decision-making process that also relates to *conservatism*, a psychological phenomenon defined as the slow updating of models in the face of new evidence (Edwards, 1968).

Recent anecdotal evidence on belief dynamics and trading activity on the stock market during the COVID-19 market crash seems to confirm such a psychological scheme in which behavioral reactions to risky situations are driven by emotions experienced at the moment of decision-making, together with cognitive evaluations which rely on probabilities and expected outcomes. Giglio et al. (2020) surveyed retail investors during February-March 2020, documenting a substantial increase during the crash in the dispersion of beliefs across investors, who turned more pessimistic and increased their perceived probabilities of catastrophic events in terms of real economic outcomes and further stock market declines. Interestingly, the authors document the largest decline in the expectations of the most optimistic investors in February, who then went on to sell the most equity in March; instead, respondents who were the most pessimistic in February largely left their portfolios unchanged during and after the crash.

From a psychological perspective, the time-varying risk perception of COVID-19 as measured by stock market dynamics seems consistent with the characterization of the dual risk perception proposed by Peters and Slovic (1996) along *unknown risk*, referring to a hazard which is judged to be unobservable, unknown, new, and delayed in producing perceptions of harmful impacts, and *dread risk*, as defined by the extent of perceived lack of control, feelings of dread, and perceived catastrophic potential. Indeed,

the risk perception of COVID-19 showed up, at first, as an unknown risk with uncertain economic impacts[4] qualifying, in a next phase, as *dread risk*, with extreme perceived catastrophic perspectives that were reflected in a substantial stock market decline.

In this paper, we introduce an agent-based model (ABM) for the learning and decision-making process of stock market traders when catastrophic and unprecedented events materialize within a behavioral heterogeneous agents' context with bounded rationality. We model a market for one risky asset paying dividends and one risk-free asset with a fixed interest rate (Chiarella et al., 2008; Kaizoji et al. 2015, Westphal and Sornette, 2020), in which a population of traders, grouped within four styles of investing and adopting different price expectation rules (naïve, biased, weak and strong chartists), receive market-based and public information news every period and subsequently reallocate their wealth between risk-free and risky assets, as in Chiarella and He (2001, 2003) and Hommes (2006). Agents process news according to the Bayesian posterior correction mechanism through the representativeness introduced in Gennaioli et al. (2015). Bordalo et al. (2018) describe how representativeness impacts on making judgments about the probability of an event under uncertainty: "agents overweight those future states whose likelihood increases the most in light of current news relative to what they know already. Thus, just as doctors overestimate the probability of sickness after a positive test result, agents overestimate the probability of a good future state when the current news is good" (p. 200). As the authors note, the approach has significant implications in terms of excessive optimism/pessimism based on the path of good/bad news processed by the agent: (1) excessive optimism is a path of good news leading an agent to focus on positive future outcomes while neglecting the bad ones, while (2) excessive pessimism results in a path of bad news leading the agent to focus on negative future outcomes and neglecting the good ones.

In our setting, agents rectify Bayesian posterior probabilities based on the amount of good and bad news they receive over time and the weight they attach to it, as we assume that some news is more important

[4] As we emphasize in the opening quotation to the paper (Simon Edelsten, February 26, 2020, Financial Times).

than others and receives a higher weight. When good news exceeds bad news in value, the likelihood of excessive optimism (representative state) is inflated, while the likelihood of bad news (non-representative state) is deflated. Symmetrically, when excessive pessimism is the representative state, bad news has more weight than good news. The severity of the probability deflation is modulated by a parameter δ, ranging from 0, corresponding to a complete neglect of risk (investors only process the most representative state), to 1, corresponding to investors holding rational expectations; intermediate values of δ relate to investors overestimating the likelihood of representative states (see Gennaioli et al., 2015).

Changes from one mood to another (from excessive optimism to excessive pessimism and *vice-versa*) occur through representativeness, causing large swings in agents' confidence and price expectations. 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 then overreact by reassessing their price expectation proportionally to the variation of the corrected posterior probability of the most representative state. Our modelling of overreactions and changes in behavior is consistent with the bimodal risk perception setting (McClelland et al., 1993), with which an event suddenly comes "on screen", producing significant changes in behavior whenever the associated risk probability crosses a certain threshold.

Speed of change in agents' confidence depends on three factors: 1) how fast news accumulates over time (amount of good and bad news); 2) how far agents look back in time when processing information before taking their investment decisions (memory); 3) how much importance agents attach to news (weight). Overreactions in price expectations are made-up of price jumps whose amplitude is modulated by a parameter $\eta \geq 0$ which reflects the perceived catastrophic potential. Therefore, our framework is so general as to include unprecedented tail events as a special case, being classified as such by their posterior probability and the disproportionate weight attached to bad news, first, and by the jump size in price expectations, second.

A higher number of changes in beliefs according to the representativeness mechanism reflects on more frequent (positive and negative) price jumps, thus contributing to price volatility. The effect is stronger with high neglect of risk and short memory windows upon which traders collect and process their information in forming price expectations.

By inspecting the stock market crash induced by the pandemic in February-March 2020 and the subsequent partial stock market recovery, we show that our model matches the pattern of the STOXX Europe 600 Index exhibited before and after the big shock.

Settling on 21 February 2020, the pandemic-based time threshold (international media put the COVID-19 news "on screen" focusing on the Italian epidemic epicenter), we find that the memory window shrinks from pre- to post-COVID outbreak. We estimate a posterior probability of a bad state under 30 percent even after the news was already circulated worldwide; at the same time, the discount factor is getting closer to 1, indicating that agents started to process information by deflating the bad news less. We classify this first phase as unknown risk perception: the extreme event impacts are unobservable, unknown and judged as potentially substantial but nevertheless delayed.

The second phase is when the extreme negative impact of COVID-19 materializes. The impact we estimate in one day (12 March 2020) is comparable to more than half a year (in business days) of continuing bad news on a daily basis. We classify this second phase as dread risk perception, where catastrophic perspectives reflect a jump close to 100 percent in the bad state posterior probability, continuing for nearly 3 months. The third phase is when investors are making close-to-rational expectations (δ is around 0.95) by denoting higher sensitivity to flows of news and judging positive or negative market outcomes as almost equally likely (calibrated posterior probabilities are both around 0.5 on average).

The paper proceeds as follows: Section 2 presents the model setup and shows how representativeness impacts on agents' confidence and price expectations in a dynamic setting; Section 3 focuses on price

overreaction; in Section 4, we simulate the full model, while Section 5 presents and comments on the results from our calibration experiment; Section 6 concludes.

2 Market price dynamics

The model setup follows market equilibrium dynamics with heterogeneous beliefs as in Brock and Hommes (1997). The market is composed of $\sum j = J$ heterogeneous and bounded rational traders who invest their wealth in a risk-free asset and in a risky asset. Assuming that each *j-th* investor is a myopic mean variance maximizer, demand for the risky asset share s solves the following problem:

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

where a denotes the risk aversion parameter which we assume 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. Tomorrow's wealth for the *j-th* investor is computed as:

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

where $s_{j,t}$[5] is the share of wealth invested in the risky asset at time t, 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, while $E_{h,t}(p_{t+1})$ is tomorrow's expected price p_{t+1} of the risky asset predicted by all investors following the *h-th* trading rule.

Investors are assumed to form their expectations on future price dynamics by extrapolating past prices following four different *h* trading rules:

[5] As in Brock and Hommes (1998), the solution for the share of wealth invested in the risky asset at time t is: $s_{j,t} = \frac{E_{h,t}(D_{t+1}) + E_{h,t}(p_{t+1}) - Rp_t}{aV_{h,t}(W_{j,t+1})}$.

(1) $$E_{h,t}(p_{t+1}) = p_t;$$

(2) $$E_{h,t}(p_{t+1}) = p_t(1 + b);$$

(3) $$E_{h,t}(p_{t+1}) = p_t + \alpha(p_t - p_{t-1});$$

(4) $$E_{h,t}(p_{t+1}) = p_t + \beta(p_t - p_{t-1}),$$

where $b > 0$ represents a positive bias and α and β are positive trend extrapolation coefficients with $\beta > \alpha$.

Equation (1) is the naive heuristic. According to this rule, agents form their expectations using the last observed price. This is the simplest case tracing back to Ezekiel (1938) for which past prices are assumed to prevail in the future. Equation (2) is the biased rule and represents optimistic expectations as in Brock and Hommes (1998). Agents are optimistic about future prices which are assumed to follow an increasing path by adding a constant b to past prices. Equations (3) and (4) describe chartist strategies where the extrapolation coefficients α and β measure the strength of the adjustment. These strategies have been deeply analyzed both in laboratory experiments (e.g. Hommes, 2011; Anufriev and Hommes, 2012) and in empirical studies on financial market dynamics (e.g. Frankel and Froot, 1990).

Once the *j-th* trader has formed her/his expectations about a future price, she/he submits a limit order ($l_{j,t}$) to the central order book (Staccioli and Napoletano, 2020). A limit order is characterized by the *desired* quantity and market order (i.e. to sell or buy):

$$l_{j,t} = (\hat{Q}_{j,t}, o_{j,t});$$

where the first element in the brackets denotes the quantity that the trader would like to share, $\hat{Q}_{j,t} = s_{j,t}W_{j,t}$, and $o_{j,t}$ represents the order type. We assume that all traders can short-sell but, in line with Raberto et al. (2001), they do not have access to external financing, thereby avoiding budget constraints.

Once all orders have been collected, the stock market volume at time t is computed as the minimum between buy and sell order quantities:

(5) $$V_t = min\left[\sum_{j=1}^{Y}\left(\hat{Q}_{j,t}|o_{j,t} = buy\right); \sum_{j=1}^{J-Y}\left(\hat{Q}_{j,t}|o_{j,t} = sell\right)\right].$$

Equation (5) implies that some orders may not be executed for a given market price. Namely, for some investors the *desired* quantity they want to exchange is higher than the actual traded quantity, i.e. $\hat{Q}_{j,t} > Q_{j,t}$. In this case, we randomly select traders who have decided to exchange in order to get $\hat{Q}_{j,t} = Q_{j,t}$.

As in Hommes et al. (2020), the realized price p_t at time t depends on the *h*-based price average predictions $E_{h,t-1}(\bar{p}_t)$ and the rational expectation equilibrium fundamental price $p_{f,t}$ as follows:

(6) $$p_t = p_{f,t} + \frac{1}{R}\left(E_{h,t-1}(\bar{p}_t) - p_{f,t}\right) + \varepsilon_t,$$

where $p_{f,t} = \frac{D_t}{r}$ and ε_t is an IID noise term. As Hommes et al. (2020) point out, the price equation for p_t has rational bubble solutions, with $\left(E_{h,t-1}(\bar{p}_t) - p_{f,t}\right)$ growing at the risk-free rate r, although these rational bubbles are often excluded by imposing transversality conditions.

Agents are allowed to maintain or switch their trading rule in every period based on the prediction ability of the forecasting heuristic they use. The logic comes from Brock and Hommes (1997), where agents tend to switch towards forecasting strategies that have performed better in the recent past, and are consistent with the diagnostic expectations (DE) introduced in Bordalo et al. (2018) and Bordalo et al. (2020a). Within the DE framework, investors receive noisy private information every period, to then update their beliefs by putting more weight on the states of the world whose objective likelihood has increased the most in light of recent news. In our model, traders are forward-looking in forming their beliefs and after formulating their (noisy) price expectation, they are assumed to check whether the forecasting power of the trading rule is good enough to maintain the strategy or switch to another. Living in a noisy environment, as traders receive good signals (with low noise), they grow more confident about the value of the asset they expect, thereby maintaining the same trading rule in forming price expectations. As we discuss in the next section, another important ingredient in our model which makes it more consistent with diagnostic expectations is the assumption that investors incorporate their signals more aggressively into their beliefs according to the representativeness heuristic (Kahneman and Tversky,

1972). This refers to the notion that, in forming probabilistic assessments, individuals put too much weight on outcomes that are more likely in relative rather than in absolute terms.

Agents assess their prediction ability according to the following forecasting diagnostic (Hommes, 2011):

$$U_{h,t} = \left[p_t - E_{h,t-1}(p_t)\right]^2 \tag{7}$$

where p_t is the actual price at time t and $E_{h,t-1}(p_t)$ is the expected price for period t formed in $t-1$ with the *h-th* trading rule. The switching mechanism is based on a relative performance evaluation process that every agent executes in each period relative to the average forecasting diagnostic exhibited by all agents, $\overline{U}_t = J^{-1} \cdot \sum_{h=1}^{H}(U_{h,t} \cdot J_h)$. The relative performance evaluation between $U_{h,t}$ and $\overline{U}_t$ causes agents to switch to forecasting strategies according to the following Proposition 1.

Proposition 1: In every period t*, every agent has three possible actions they can follow based on* $Q_{j,t}$*,* $U_{h,t}$ *and* $\overline{U}_t$*:*

1. *If she/he has traded in* t *and the performance measure* $U_{h,t}$ *is lower than the average forecasting error* $\overline{U}_t$*, she/he maintains both the same h-th heuristic and the same order type in the next period* $t+1$.
2. *If she/he has traded in* t *and the performance measure* $U_{h,t}$ *is higher than the average forecasting error* $\overline{U}_t$*, she/he switches the strategy randomly towards another h-th heuristic while maintaining the same order type in the next period* $t+1$.
3. *If in period* t *he/she has not traded, he/she randomly changes both the h-th heuristic and the order type in the next period* $t+1$.

Mathematically:

$$\left(h_{j,t+1}; o_{j,t+1}\right) = \begin{cases} \left(h_{j,t}; o_{j,t}\right) & if \;\; Q_{j,t} > 0 \; and \; U_{h,t} < \overline{U}_t \\ \left(\rho_{j,t+1}; o_{j,t}\right) & if \;\; Q_{j,t} > 0 \; and \; U_{h,t} \geq \overline{U}_t \\ \left(\rho_{j,t+1}; \varrho_{j,t+1}\right) & if \;\; Q_{j,t} = 0 \end{cases} \tag{8}$$

where $\rho_{j,t}$ is a random integer parameter uniformly distributed on the support (1, H), and $\varrho_{j,t+1}$ is a binary categorial variable randomly uniformly distributed on the domain (buy, sell).

The switching mechanism does not work for all agents, since we assume that a sample of investors maintain their forecasting strategy regardless of the prediction ability of the rule. This is consistent with the status quo effect (Kahneman et al., 1991), in which individuals have a strong tendency to remain at the status quo, because the disadvantages of leaving it are greater than the corresponding advantages.

3 Representativeness and Price Overreaction

In our model setup, once the exchange takes place, the informational set available to traders is enriched by the information on pure price and trading movements (market activity-based information), together with general economic and political news, in addition to firm-specific information (public information). Agents use both strings of information to update their price expectation, an assumption which is consistent with empirical literature confirming the robust relationship between public information and market activity measures (e.g. Mitchell and Mulherin, 1994). Market activity-based information measures are (1) stock price dynamics; (2) trading volume; (3) the number of traders who actually executed buy or sell orders. The news is classified as *good* (n_g) whenever the change that occurred in each of the 3 measures in the last period is ≥ 0; instead, the news is classified as *bad* (n_b) when the change is negative.

Public information could instead rely on scheduled and unscheduled announcements. The logic is consistent with the empirical findings on public information and the role it plays as a powerful source of price movements, such as macroeconomic and political announcements (Aït-Sahalia et al., 2012; Birz and Lott, 2011; Baker et al., 2019), central bank communications (Andrade and Ferroni, 2020) and firm-specific announcements, such as earnings announcements (Bernard and Thomas, 1990).

3.1 News, memory, and weight of news

Agents receive strings of good and bad news in each period for both types of information (market activity-based and public) and accumulate a memory of past news (*m*) over time up to time $t-1$, used to update their expected posterior probability ($\pi_{x,t}$) concerning the state of the market with the new piece of information of period t:

(9) $$\pi_{x,t} = \frac{\pi_{x,t-1}+n_{x,t}}{1+n_{g,t}+n_{b,t}};$$

where $n_{x,t} = \sum_{i=1}^{m} \varphi_i n_{i=x,t}$, with x denoting the state of the market (good, g, or bad, b), $n_{x,t}$ is the total news of *x-th* type which arrived at time t. $\varphi_i > 0$ is the weight assigned to each news since we assume that some information is perceived as more important than other information, thus reflecting a higher weight. For every time step t we have:

$$\pi_{g,t} + \pi_{b,t} = 1.$$

See Proposition A.1 in the Appendix.

It is at this point that representativeness comes into play. Kahneman and Tversky (1972) define the heuristic as follows: "an attribute is representative of a class if it is very diagnostic; that is, the relative frequency of this attribute is much higher in that class than in a relevant reference class". Within a stock market setup like the current one, representativeness 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; in the same way, when past news is bad, investors are excessively pessimistic when future outcomes are negative and neglect good news because it is less representative.

Proposition 2: 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.

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

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

and next they revise the posterior probability for good, $\pi^p_{g,t}$, and bad states, $\pi^p_{b,t}$, based on the following rule:

$$if\ R_{g,t} > R_{b,t} \begin{cases} \pi_{g,t}^{p} = \dfrac{\pi_{g,t-1} + n_{g,t}}{\left(\pi_{g,t-1} + n_{g,t}\right) + \delta\left(\pi_{b,t-1} + n_{b,t}\right)} \\ \pi_{b,t}^{p} = \dfrac{\delta\left(\pi_{b,t-1} + n_{b,t}\right)}{\left(\pi_{g,t-1} + n_{g,t}\right) + \delta\left(\pi_{b,t-1} + n_{b,t}\right)} \end{cases}$$

(10)

$$if\ R_{b,t} > R_{g,t} \begin{cases} \pi_{b,t}^{p} = \dfrac{\pi_{b,t-1} + n_{b,t}}{\left(\pi_{b,t-1} + n_{b,t}\right) + \delta\left(\pi_{g,t-1} + n_{g,t}\right)} \\ \pi_{g,t}^{p} = \dfrac{\delta\left(\pi_{g,t-1} + n_{g,t}\right)}{\left(\pi_{b,t-1} + n_{b,t}\right) + \delta\left(\pi_{g,t-1} + n_{g,t}\right)} \end{cases}$$

where δ *is the discount factor that modulates the severity of the probability deflation with* $0 \leq \delta \leq 1$. *We have three scenarios: (1) when* $\delta = 0$, *investors only process the most representative state; (2) when* $\delta = 1$, *investors hold rational expectations; (3) with intermediate values of* δ, *investors overestimate the likelihood of representative states* (Gennaioli et al., 2015).

For every time step t we have:

$$R_{g,t} > R_{b,t} \ \Leftrightarrow\ \pi_{g,t} > \pi_{g,t-1} \Leftrightarrow \pi_{g,t} > \frac{n_{g,t}}{N}, \text{ with } N = n_{g,t} + n_{b,t},$$

see Proposition A.2 in the Appendix; and:

$$\pi_{g,t}^{p} + \pi_{b,t}^{p} = 1,$$

see Proposition A.3 in the Appendix.

Suppose investors observe a longer string of good news, obtaining $\pi_{g,t}^{p} > \pi_{b,t}^{p}$. Proposition 2 states that, while traders are surfing the wave of optimism by neglecting the probability of bad events, the *actual* posterior probability of the negative market state is not changed, even though bad news has been less frequent. Traders are therefore overestimating the probability of positive outcomes endogenously, thereby reinforcing the market trend until new information is strong enough to change the probability of the less representative state suddenly becoming the more representative one. In short, the change in an agent's belief depends on three factors:

1) The amount of good and bad news: strings of market activity-based and public information accumulate over time, providing positive and negative signals to the agent who assesses probability estimates for good and bad market states;
2) Memory: agents look back in time when processing information in a way which is consistent with Bordalo et al. (2020b), who postulate that experiences from the past are first consolidated within a norm, then individuals adjust valuations in the direction of any discrepancy between the estimated and observed attributes. In our setup, memory relates to the length of the time interval over which past news is recorded and evaluated through representativeness leading to shape and reshape beliefs and price expectations.
3) Importance attached to the news (weight): news is not all the same but differs in its contribution to agents' beliefs based on the importance individuals attach to it. Changes in agents' beliefs can therefore occur through flows of news on the same type (good or bad) or because of the high impact the news exerts on risk/return perception.

3.2 Agents' confidence and Price Overreaction

First of all, representativeness causes changes in agents' beliefs, and this occurs when the less representative state suddenly becomes the more representative one. Probabilities are dynamically updated with arriving strings of good and bad news, some of which might be so "extreme" as to overturn the previous probability state. Therefore, the amount of good and bad news and the weights attached to it induces changes from one mood to another. These sudden changes in beliefs produce jumps in price expectations, whose amplitude is modulated by a parameter $\eta \geq 0$ which reflects the perceived boom/bust potential.

Proposition 3: When the probability of the less representative state suddenly becomes the highest, the price expectation exhibits a jump which is proportional to change in the adjusted posterior probability of the most representative state.

$$(11)\quad E^p_{h,t}(p_{t+1}) = \begin{cases} E_{h,t}(p_{t+1}) - \gamma_t \quad if\ \pi^p_{g,t-1} > \pi^p_{b,t-1} \cup\ \pi^p_{g,t} < \pi^p_{b,t} \cup \frac{n_{b,t}}{n_{g,t}} > \frac{\pi^p_{b,t-1}+n_{b,t-1}}{\pi^p_{g,t-1}+n_{g,t-1}} \\ E_{h,t}(p_{t+1}) + \gamma_t \quad if\ \pi^p_{b,t-1} > \pi^p_{g,t-1} \cup\ \pi^p_{b,t} < \pi^p_{g,t} \cup \frac{n_{g,t}}{n_{b,t}} > \frac{\pi^p_{g,t-1}+n_{g,t-1}}{\pi^p_{b,t-1}+n_{b,t-1}}; \\ E_{h,t}(p_{t+1}) \quad otherwise; \end{cases}$$

where $\gamma_t = \eta_x(\pi^p_{x,t} - \pi^p_{x,t-1})$ *with* $\eta_x > 0$ *measuring the perceived boom/bust, also allowing for possible asymmetric impacts with* $\eta_b \lessgtr \eta_g$.

Eq. (11) can be reformulated as:

$$E^p_{h,t}(p_{t+1}) = \begin{cases} E_{h,t}(p_{t+1}) - \gamma_t \quad if\ \pi^p_{g,t-1} > \frac{1}{2} \wedge\ \pi^p_{g,t} < \frac{1}{2} \wedge \frac{n_{b,t}}{n_{g,t}} > \frac{\pi^p_{b,t-1} + n_{b,t-1}}{\pi^p_{g,t-1} + n_{g,t-1}} \\ E_{h,t}(p_{t+1}) + \gamma_t \quad if\ \pi^p_{g,t-1} < \frac{1}{2} \wedge\ \pi^p_{g,t} > \frac{1}{2} \wedge \frac{n_{g,t}}{n_{b,t}} > \frac{\pi^p_{g,t-1} + n_{g,t-1}}{\pi^p_{b,t-1} + n_{b,t-1}}; \\ E_{h,t}(p_{t+1}) \quad otherwise; \end{cases}$$

See Proposition A.4 in the Appendix.

To simplify further, we can finally re-manipulate the equation as:

$$(12)\quad E^p_{h,t}(p_{t+1}) = E_{h,t}(p_{t+1}) + \mathrm{P}_t\gamma_t, \text{ with}$$

$$\mathrm{P}_t = \begin{cases} -1 \quad if\ \pi^p_{g,t-1} > \frac{1}{2} \wedge\ \pi^p_{g,t} < \frac{1}{2} \wedge \frac{n_{b,t}}{n_{g,t}} > \frac{\pi^p_{b,t-1} + n_{b,t-1}}{\pi^p_{g,t-1} + n_{g,t-1}} \\ +1 \quad if\ \pi^p_{g,t-1} < \frac{1}{2} \wedge\ \pi^p_{g,t} > \frac{1}{2} \wedge \frac{n_{g,t}}{n_{b,t}} > \frac{\pi^p_{g,t-1} + n_{g,t-1}}{\pi^p_{b,t-1} + n_{b,t-1}}; \\ 0 \quad otherwise. \end{cases}$$

Equation (12) can be interpreted as a mean-reverting doubly stochastic process with Poisson jumps. Specifically, if we consider the event:

$$A_t = \left\{ \left(\pi^p_{g,t-1} > \frac{1}{2} \wedge\ \pi^p_{g,t} < \frac{1}{2} \wedge \frac{n_{b,t}}{n_{g,t}} > \frac{\pi^p_{b,t-1} + n_{b,t-1}}{\pi^p_{g,t-1} + n_{g,t-1}} \right) \right.$$
$$\left. \vee \left(\pi^p_{g,t-1} < \frac{1}{2} \wedge\ \pi^p_{g,t} > \frac{1}{2} \wedge \frac{n_{g,t}}{n_{b,t}} > \frac{\pi^p_{g,t-1} + n_{g,t-1}}{\pi^p_{b,t-1} + n_{b,t-1}} \right) \right\},$$

and its complementary event $B_t = A_t^C$, then λ_t is the probability of A_t, as in (Duffie, 2005) and (Duffie and Singleton, 2003).

Note that our framework includes unprecedented tail events as a special case. Suppose the tail event is negative, this is what materializes in the market. The news first arrives on screen and agents attach a disproportionate weight; second, the bad state posterior probability is adjusted upward thus exceeding the threshold $\frac{1}{2}$; third, agents change their beliefs and, conditional on the new adjusted posterior probabilities, reassess their price expectations along a jump-based process formalized in eq. (12), in which the jump size $\eta_x > 0$ reflects the catastrophic potential perceived and evaluated by agents. This is consistent with the probability weighting function suggesting that people tend to overestimate the likelihood of tail events in their decision-making process (Barberis, 2013). Moreover, our price overreaction is also consistent with the bimodal risk perception logic (McClelland et al., 1993), as the event that suddenly comes "on screen" can produce changes in behavior and price, only if the associated risk probability crosses a certain threshold, in our case set at $\frac{1}{2}$ (see eq. (12) and Proposition A.4).

4 Model Simulation

4.1 Simulation setup and market dynamics

Since the model includes heterogenous, bounded rational agents and cannot be solved analytically, in this section we run a numerical simulation to explore the entire market dynamics. This exercise shows how our model is able to produce many market stylized facts, such as bubbles and crashes, excess volatility, fat-tailed return distributions, uncorrelated price changes and volatility clustering.

The values assigned to the model's parameters are in Table 1. We start by assuming a market with 200 traders over 250 trading periods. At time 0, agents are equally distributed among the four *h* trading rules, then switching (or maintaining) their trading rule in every period based on the prediction ability of the forecasting heuristic they use, as discussed in Section 2. Following Anufriev and Hommes (2012), we set the trend coefficients for the weak (α) and strong (β) chartist (eq. 3 and eq. 4) at 0.4 and 1.3, respectively,

and the risk-free rate r at 5 percent (Hommes et al., 2020). As in Bao et al. (2017) we take the fundamental price as fixed over time, also resulting in a constant dividend. Specifically, we assume a fundamental price $p_{f,t} \equiv \overline{p_f} = 100$, implying a dividend $D_t \equiv \overline{D} = r \cdot p_{f,t} = 5$. The conditional variance $V_{h,t}(W_{j,t+1})$ is computed in each time t, based upon prices p_t (eq. 6) up to time $t-1$. The bias parameter b in equation (2) for optimistic expectations is 2.5% (De Grauwe and Rovira Kaltwasser, 2012) and loss aversion $a = 20$ (Tversky and Kahneman, 1992; Aït-Sahalia and Brandt, 2001; Pruna et al., 2020). The discount factor that modulates the severity of the probability deflation within the representativeness framework is $\delta = 0.8$, therefore overestimating the likelihood of representative states. Memory is set at $m = 6$ periods, over which agents evaluate bad and good news through representativeness. After each trade, investors receive a total amount of news $n = n_{mkt} + n_p = 5$ of which: (i) $n_{mkt} = 3$ is market activity-based information (price trend, volume exchanged and active traders in the past run) with a weight attached to each stream of news $\varphi_{mkt} = 1$, and (b) $n_p = 2$ is public information (e.g. public macroeconomic/political announcements) with a weight attached to each stream of news $\varphi_p = 2$, since we assume that this information has a higher impact.

Good and bad public information news is randomly simulated and is therefore exogenous, while market activity-based news is endogenously formed by the simulation exercise. Accordingly, overreactions are both endogenous and exogenous. Together with public and market activity-based information, we also include a number of exogenous extreme events $n_{ee,t}$ which are randomly generated and can be both positive or negative with $0 \leq n_{ee,t} \leq 10$. In our exercise, a number of $n_{ee,t} = 5$ extreme events were simulated, which are depicted and numbered in Figure 1 by means of vertical dotted lines (blue is for positive extreme events, red for negative ones). We arbitrarily attach to each extreme event a weight $\varphi_{ee} = 36$, which persists over the next 6 trades with a constant decrease and zeroing at the sixth trade ($\varphi_{ee,t} = 36 - 6t$ with $t = 1, \dots, 6$). The logic we follow in simulating this news is that, once the extreme event occurs, the impact is definitely greater than common news and remains vivid, although disappearing piecemeal over the short run. Boom burst patterns are driven by

the flow of good and bad news arriving over time to agents, who tend to neglect the probability of the less representative state because the corresponding news registered over their moving memory windows are less frequent and/or have lower weights. However, when new strings of information of the less probable state are high enough to change the agent's beliefs, the less representative state suddenly becomes the more representative one, traders overreact and the market price exhibit jumps with an amplitude that reveals the perceived boom/bust. The amplitude of jumps is set at $\eta_b = 60$ and $\eta_g = 30$ for bad- and good-based jumps, respectively, then assuming a sensitivity to losses greater than that of gains (Thaler et al., 1997).

It is interesting to note that not all the 5 extreme events affect market dynamics in the same way. The first two vertical dotted red lines in Figure 1 are two negative events having the same impact, but the effect triggered by each on price dynamics differs substantially. In the first extreme event, the bad news does not activate an overreaction in price expectations and, therefore, we do not observe a negative spike in the market price. This is because traders have already recorded higher amounts of bad news, thereby still being in a bad state. For this reason, the new event, while extreme, is not strong enough to reflect into $P_t = -1$ (eq. 12); see Figure 1. As a result, investors continue to deflate good news because this is less representative. Instead, when the second extreme event happens, the market is in a good state and traders are underestimating the probability of bad news. Once the negative event materializes, agents correct their posteriors and, since the bad state is no longer the less representative state, they start deflating good news. As a whole, extreme event-driven overreactions appear when the vertical dotted lines intersect with the spikes in the dashed lines, where we have $P_t = \pm 1$. Instead, the other overreactions are driven by the strings of market activity-based and public news.

Market price dynamics depend on the trading rule switching mechanism also followed by the agents (Proposition 1). In Figure 2, we note that the optimistic (biased) rule is chosen by the lower share of agents, while strong trend follower and naïve rules are the strategies mostly selected by investors, together accounting for around 40-50 percent of the market over the 250 trading periods. This is due to the relative performance evaluation process (eq. 6), leading agents to switch towards forecasting strategies that have

performed better (according to Proposition 1). Strong trend follower and naïve are likely to be the rules that better capture jump-based price dynamics. Moreover, once overreactions take place, both heuristics tend to move with the market price, to then reflect in a higher share of traders following such rules since they perform better in terms of forecasting error.

The market price dynamics we simulate denote many of the market stylized facts documented in the empirical literature on bubbles and crashes, excess volatility, fat-tailed return distributions, uncorrelated price changes and volatility clustering. The top left- to right-hand panel of Figure 3 reports log price dynamics relative to its fundamental value (blue line) to emphasize the bubbles and anti-bubbles generated by the model through representativeness-based overreactions. Also, the top-right side panel exhibits the return behavior, which clearly denotes high volatility. The corresponding return distribution and autocorrelation function are in the bottom left- to right-hand panel. Note the non-normal distribution with fat tails, negative asymmetry and excess kurtosis, all characterizing an equity return distribution far from normality when extreme events occur. The autocorrelation function (ACF) of raw returns for the first 100 lags computed together with their 95 percent confidence bands denote no statistically significant return autocorrelation. This is interesting since no statistical artefacts arise even when including possible extreme events. Indirectly, this is consistent with the view on price latency[6] as a primary source of autocorrelations in equity returns (e.g. Atchison et al., 1987), which are indeed not included in our simulation.

4.2 Sensitivity Analysis

In order to inspect how market dynamics are sensitive to changes in the key model parameters within the same simulation design (Table 1), we run 50 Monte Carlo simulations by changing the memory (m) and the discount factor (δ), which assume a major role in the way agents collect and process information in forming their price expectations. Specifically, we generated a total of 6 extreme events, 4 positive and

[6] Price latency, also referred to as nonsynchronous trading, relies on delays in transaction price adjustments due to market friction, such as liquidity. Other market frictions, such as price stabilization mechanisms (designed to control price volatility), may also contribute to the observed autocorrelation; see, Harris (1989).

2 negative, over the following time periods: period # 24, 35, 49 and 175 for good extreme events, and period # 157 and 193 for bad extreme events. We first ran the simulations using $m = 1;\ 6; 12$, thereby inspecting how the short- to medium-term memory windows impact on market dynamics. Next, we changed the representative-based deflation mechanism for news using $\delta = 0.6;\ 0.8; 1$, and then compared the price formation, when the more representative state was overestimated ($\delta = 0.6;\ 0.8$), with the rational expectation hypothesis ($\delta = 1$) in which more and less representative states are processed in the same way with no deflation mechanism.

Figure 4 reports market price distribution for different values of memory. Note that as memory increases, price dispersion reduces substantially. When processing more pieces of information collected over higher time intervals, agents formulate their price expectations by containing their forecasting error, which governs the switching mechanism of trading strategies leading the market price on a higher level with a low dispersion. Interestingly, note how extreme events impact differently on the three scenarios. When traders use the informational set of the last trade only ($m = 1$), the market price denotes a higher volatility because of the higher number of changes in beliefs according to the representativeness mechanism, subsequently reflecting on more frequent (positive and negative) price jumps. As a result, the price level tends to be lower on average with strong dispersion. Indeed, having less information upon which to compute representativeness, the impact of unexpected news, i.e. with lower representativeness, is likewise high enough to change agents' beliefs. To put the point into perspective, we calculate that under the $m = 1$ scenario, overreactions were almost one-third of total price simulations (3,732 out of 12,500[7]), of which only 3.91% ascribed to extreme events. On the other hand, when traders move on longer memory windows, the informational set enriches and good and bad news tend to mix together reflecting on more stable representativeness, and then, on less changes in beliefs. Overreactions are reduced to about one-fifth (2,480 out of 12,500) and overreactions due to extreme events are 5%. Market prices denote low dispersion with higher persistence around the trend.

[7] Since $T = 250$, and 50 were the Monte Carlo simulations, we have $250 \times 50 = 12{,}500$ price simulations as a whole.

Consider now the role played by the discount parameter δ. Figure 5 reports price distributions for the three scenarios $\delta = 0.6;\ 0.8; 1$. Note that when the severity of the distortion is stronger (i.e. $\delta = 0.6$), market price dynamics are more dispersed, also impacting on the lower average price value. This result can be explained by the nature of the representativeness mechanism which leads traders to process good and bad news asymmetrically, in that less representative news are deflated while posterior probabilities of the more representative news are over-weighted. Instead, when traders become more "rational" (i.e., $\delta = 1$), they do not deflate any news, resulting in lower impacts on price dynamics. As a result, the market price level is higher on average and less volatile.

5 Calibrating the Stock Market Impact of COVID-19

The COVID-19 pandemic is a natural laboratory to explore how the learning process of market traders evolves when an unprecedented and unknown event occurs. The pandemic was perceived as an unprecedented global crisis, with hundreds of countries having implemented varying degrees of restrictions on population movement to slow the spread of the severe acute respiratory syndrome (Han et al., 2020). As pointed out by the World Bank[8], COVID-19 caused a global recession whose depth over the past century and a half was surpassed only by the two World Wars and the Great Depression.

How did the stock market react when the world grasped the devastating health and economic crisis caused by the pandemic? In retrospect, stock market dynamics during the COVID-19 pandemic denote three main phases. In the first phase, markets simply ignored the potential impacts of the pandemic (until 21 February 2020). In the second phase, from 23 February to 20 March 2020, the fear of COVID caused stock markets to plummet, reaching a global crash on 12 March 2020; stock markets suffered from the greatest single-day percentage fall since the 1987 crash. Lack of investor confidence was also exacerbated by the European Central Bank, since the Governing Council decided to keep key ECB interest rates

[8] See https://www.worldbank.org/en/publication/global-economic-prospects

unchanged[9] despite market expectations.[10] In the third phase, from 23 March to 20 April 2020, stock market prices rebounded worldwide following massive interventions by central banks.

In order to examine the pandemic's impact on stock markets, we focus on the STOXX Europe 600[11] over the period from 10 July 2019 to 30 June 2020. This period includes 250 daily closing values of the index we used to calibrate our agent-based model. The following sections present the calibration set-up and discuss main results.

5.1 Calibration set-up

We used the $T = 250$ observations (from 10 July 2019 to 30 June 2020) of daily values of the STOXX Europe 600 Index $\boldsymbol{d} = (d_1, d_2, \dots, d_T)$ to calibrate the model, assuming each value as a single trade. We calibrated the discount factor, the jump amplitude and the fundamental price trend, whereas the remaining model parameters maintained the same settings as before (see Section 4.1, Table 1), except for the timing of the extreme event and the corresponding weight, which we discuss later.

Denoting the discrete time interval upon which the model is calibrated as $I \coloneqq [1, T] \cap \mathbb{N}$, for every time step $t \in I$ we assumed a time-varying discount factor δ and a jump amplitude η, having $\delta_t \in [0,1]$ and $\eta_t \in (0, \infty]$. The fundamental price $p_{f,t}$ is assumed as a first-order autoregressive process following

$$p_{f,t} = \alpha + \tau_t p_{f,t-1} + \varepsilon_t \tag{13}$$

with $\tau_t \in [0,1]$, thus imposing stationarity. We set $p_{f,1} = 350$, corresponding to the average computed over the period from May to June 2020. Memory m_t is also time-varying, and we set $1 \leq m_t \leq K$ for every $t \in I$, thereby assuming that traders update their expectations looking at a time window only ranging from the last trade to K past days, which we set at 20 to enable the modelling of the essential aspect of learning from experience, namely the gradual loss of memory (Nagel and Xu, 2019).[12] Finally,

[9] See https://www.ecb.europa.eu/press/pressconf/2020/html/ecb.is200312~f857a21b6c.en.html
[10] S&P 500 futures dropped off more than 200 points in less than an hour.
[11] The STOXX Europe 600 Index includes a fixed number of 600 stocks and represents large, mid and small capitalization companies across 17 countries in the European region: Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, Norway, Poland, Portugal, Spain, Sweden, Switzerland and the United Kingdom.
[12] As pointed out by Nagel and Xu (2019), the dynamics of subjective beliefs and asset prices crucially depend on agents' memory, as past data moves into beliefs, decisions, and, ultimately, prices. While full memory of all past observations is a

$p_t \in \mathbb{R}_+$ is the theoretical market price obtained as a solution of Equation (6), and the corresponding vector over the time interval I is

$$\boldsymbol{p} = (p_1, p_2, \dots, p_T), \qquad \boldsymbol{p} \in \mathbb{R}_+^T .$$

We then have:

$$\boldsymbol{\delta} = (\delta_1, \delta_2, \dots, \delta_T), \qquad \boldsymbol{\delta} \in [0,1]^{\boldsymbol{T}},$$

$$\boldsymbol{\eta} = (\eta_1, \eta_2, \dots, \eta_T), \qquad \boldsymbol{\eta} \in (0, \infty]^{\boldsymbol{T}},$$

$$\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_T), \qquad \boldsymbol{\tau} \in [0,1]^{\boldsymbol{T}},$$

$$\boldsymbol{m} = (m_1, m_2, \dots, m_T), \qquad \boldsymbol{m} \in [1, K]^{\boldsymbol{T}} \cap \mathbb{N}_+^T.$$

Since COVID-19 produces a structural break in the data, we split the time series into two sub-periods by fixing a time threshold Z, $1 < Z < \mathrm{T}$. As a result, the discrete time interval I is split into two subsets $I_A \coloneqq [1, Z] \cap I$ and $I_B \coloneqq [Z+1,\ T] \cap I$, respectively. The time threshold is set on 21 February 2020 (corresponding to $Z = 160$), when international media brought the COVID-19 news "on screen" focusing on the Italian epidemic outbreak[13]. The extreme event arrival Z_{ee} is set on 12 March 2020 (corresponding to $Z_{ee} = 174$), when stock markets experienced one of the highest losses ever. The weight attached to such an event is set at $\varphi_{ee} = 950$, also assuming weight decay over 50 trades[14]. The value for φ_{ee} was chosen for minimization purposes after running many simulations $(p_{174} - d_{174})^2$.

To take into account the change in regime that occurred in the parameters' dynamics, we assume m_A , $m_B \in [1, K] \cap \mathbb{N}$ and $\tau_A, \tau_B \in [0,1]$, such that:

---

standard assumption for econometricians, memory decay seems to be a more appropriate conjecture. This is also consistent with standard Bayesian parameter learning models. The authors observe that as a consequence of memory decay, learning is perpetual and there is a persistent time-varying wedge between agents' subjective beliefs and the objective beliefs implied by the true parameters of the process generating asset payoffs.

[13] See Reuters: https://www.reuters.com/article/us-china-health-italy/coronavirus-outbreak-grows-in-northern-italy-16-cases-reported-in-one-day-idUSKBN20F0UI

[14] We chose the time decay of 50 days by mirroring pandemic lockdown durations in main European countries (e.g. France, 55; Italy, 70; Germany, 28; Ireland, 67; Spain, 56; UK, 103). The rationale for the time decay assumed for φ_{ee} is consistent with the literature on extreme returns and informational releases. For e.g., for a subset of stocks from 1990 to 1992 that had extreme returns, (Pritamani and Singal, 2001) collect news from the Wall Street Journal and Dow Jones News Wire and find both the positive and negative abnormal return drift for up to 20 days after a news story.

$$m_t = m_A \quad \text{and} \quad \tau_t = \tau_A \qquad \text{for every} \quad t \in I_A,$$

$$m_t = m_B \quad \text{and} \quad \tau_t = \tau_B \qquad \text{for every} \quad t \in I_B.$$

The whole parameters' space is thus specified as:

$$X = \{\boldsymbol{x} \mid \boldsymbol{x} = (\boldsymbol{\delta}, \boldsymbol{\eta}, m_A, m_B, \tau_A, \tau_B) \in [0,1]^{\boldsymbol{T}} \times \mathbb{R}_+^T \times \mathbb{N}_+ \times \mathbb{N}_+ \times [0,1] \times [0,1]\}.$$

At this point, to estimate $\boldsymbol{X}$ we start from the market price vector $\boldsymbol{d} \in \mathbb{R}_+^T$ and look for the parameters' vector $\widehat{\boldsymbol{x}} \in X$ which minimizes the Euclidean distance between $\boldsymbol{d}$ and $\boldsymbol{p}$. To do this, we define the map:

$$\Psi\colon X \to \mathbb{R}_+^T, \qquad \Psi(\mathbf{x}) = \boldsymbol{p_x}$$

which associates the corresponding solution $\boldsymbol{p_x}$ of Equation (6) to the parameters' vector $\boldsymbol{x} \in \boldsymbol{X}$. Computationally, we consider the following convex cost functional:

$$\Gamma\colon X \to \mathbb{R}_+, \qquad \Gamma(\mathbf{x}) = \frac{1}{2} \| \Psi(\mathbf{x}) - \boldsymbol{d} \|^2_{\mathbb{R}^T} = \frac{1}{2} \sum_{t=1}^{T} |\Psi(\mathbf{x})_t - \boldsymbol{d}_t|^2,$$

next solving the following control problem (Quarteroni, 2014):

(13) $$\widehat{\boldsymbol{x}} := \min_{\mathbf{x} \in \mathbf{X}} \Gamma(\mathbf{x}).$$

We compute eq. (14) following the interior-point method (Pólik and Terlaky, 2010).

5.2 Results

The solution for the minimization problem (eq. [14]), $\widehat{\boldsymbol{x}} = \left(\widehat{\boldsymbol{\delta}}, \widehat{\boldsymbol{\eta}}, \widehat{m}_A, \widehat{m}_B, \hat{\tau}_A, \hat{\tau}_B, \right) \in X$ led to optimal vectors $\widehat{\boldsymbol{\delta}}$ and $\widehat{\boldsymbol{\eta}}$ depicted in Figure 4, whereas the value for memory and autoregressive coefficients $\boldsymbol{\tau}$ (eq. [13]) for the two sub-periods $\boldsymbol{I_A}$ and $\boldsymbol{I_B}$ are:

$$\widehat{m}_A = 8, \qquad \widehat{m}_B = 6, \qquad \hat{\tau}_A = 6.1 \cdot 10^{-3}, \qquad \hat{\tau}_B = 0.875.$$

The resulting estimated price $\boldsymbol{p}$ is reported in Figure 5, where we compare the observed (actual) versus the simulated index values.

The major outcomes of the calibration experiment are four. Foremost, consider first the memory and the fundamental price trend. The memory window shrinks from pre- to post-COVID outbreak, moving

from 8 to 6 past days. The gradual loss of memory (Nagel and Xu, 2019) of agents seems to cover a very short-term past window, which shortens even more when the extreme event materializes. This is also due to the weight attached to the news of 12 March 2020, which is so strong that it is equivalent to more than half a year (in business days) of continuing daily bad news. Indeed, since the model generates 3 market activity-based information daily, each one having a 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 7 (see Section 4.1). Hence, since the weight attached to the COVID-based extreme event is 950, we can calculate that this extreme weight is equivalent to $\frac{950}{7} \cong 136$ days of continuing bad news. Predominant concern on very recent streams of news reflects on high price volatility, which increases just after the extreme event when the memory window shrinks. Moreover, as we discuss in Section 4.2, with less information upon which to compute representativeness, the impact of news with lower representativeness could be high enough to change agents' beliefs more frequently. This is what happened on 22 April 2020, when we move from a bad to a good state, as we comment later on price calibration results. Consider now $\hat{\tau}_A$ and $\hat{\tau}_B$. The two autoregressive coefficients reflect the substantial flat fundamental price path before the pandemic outbreak with a coefficient close to zero, and a strong mean reversion towards pre-COVID values after 12 March 2020 with a coefficient near 0.9. This value documents the substantial rebound of stock markets when central banks' massive interventions were announced all over the world.

The second key finding is on discount factor dynamics. The time-varying δ in Figure 6 denotes 4 regimes. The first approximately covers the period from 10 July 2019 to 13 August 2019 with values around 0.75 on average: since 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, the time threshold Z (COVID-19 outbreak), when discount factors jump to 0.85 on average, thus deflating bad news to a lesser extent. The third regime is from 22 February to 21 April 2020 and includes the extreme event exogenously fixed on 12 March 2020 (Z_{ee}; see Section 5.1) when we have the extreme negative overreaction, as a bad state suddenly becomes the most representative. Over this third regime, the discount factor moves around 0.92 on average, documenting

how agents tend to process information without asymmetrical deflations and then approaches a rational-based framework in which good and bad news have the same weight in forming market beliefs. This tendency is reinforced in the fourth and final regime, from 22 April 2020 to the end of the period when stock markets recovered substantially; the market shifts to a good state on 2 June 2020 to then bounce to bad-good-bad (12-15-16 June 2020) while maintaining a corrected bad posterior probability slightly higher than the threshold $\left(\pi_{b,t}^{p} > \frac{1}{2}\right)$ which causes the market to remain in a bad state until the end of the period. Over this final regime, investors tend to hold close-to-rational expectations since δ is around 0.95 on average.

Our third key outcome refers to the jump amplitude η. The parameter governs price overreaction when a change in beliefs occurs, giving measure to the perceived price impact of extreme events, as well as to endogenous and smooth changes in representativeness. Indeed, as discussed in Section 3.2, the number and the weight of good and bad news induces changes from one state to another, producing price jumps. Figure 6 denotes the 5 overreactions. The first is on 12 March 2020 corresponding to the exogenous COVID-based extreme event, the second is on 2 June 2020 which is endogenously driven by the flows of good news received after the negative peak which led to change from a bad to a good state when the COVID-based extreme event had gradually lost its weight, whereas good news accumulates to the point of overturning representativeness. Finally, it arrives at the rebound of 12-15-16 June 2020, when we move on to a bad-good-bad state. Calibrated values for the corresponding jump amplitude parameter are:

$$\hat{\eta}_{3/12/2020} = 103, \quad \hat{\eta}_{6/2/2020} = 55, \quad \hat{\eta}_{6/12/2020} = 50, \quad \hat{\eta}_{6/15/2020} = 52, \quad \hat{\eta}_{6/16/2020} = 64.$$

As expected, in Z_{ee} the parameter is almost double the others, which relate to endogenous overreactions after the extreme and unprecedented event materialized; agents now tend to be more sensitive to all news and they are "living on the hedge", being ready to change their beliefs, incidentally, by processing information on a shortened time window (as already stated, memory shortens from the pre- to post-COVID outbreak).

After having discussed the key calibrated parameters, consider now the resulting simulated price obtained as the solution of Equation (6) reported in Figure 7, compared with the actual values of the STOXX Europe 600 index. The calibrated values of the index show a good fit especially in replicating price dynamics after the COVID-19 outbreak, when we have the extreme negative return followed by a substantial market recovery. As a whole, the entire price calibration experiment makes two points.

First, the unprecedented COVID-19 impact on stock markets occurred about 3 weeks after the arrival of the pandemic was reported on the first page of newspapers around the world (Baker et al., 2020). Agents maintained their beliefs even after the news was already circulated worldwide, since the posterior probability of a bad state was contained under 30 percent (see Figure 7), 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.

Second, the extreme negative impact of COVID-19, which was so strong that in one day it assumed the same impact as more than half a year (in business days) of continuing bad news on a daily basis, 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 continuing for nearly 3 months (Figure 7). In this time interval, the stock market is recovering from the extreme negative return and is approaching a new normal, with investors making close-to-rational expectations (δ is around 0.95) and judging positive or negative market outcomes as almost equally likely (posterior probabilities are both around 0.5 on average).

6 Conclusion

In this paper, we introduce an agent-based model for the learning and decision-making process of stock market traders when catastrophic and unprecedented events materialize within a behavioral 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. This is our key assumption to explain when overreactions occur and the extent to which they impact on price dynamics. The price jump amplitude reflects the perceived catastrophic potential and, as such, unprecedented tail events are special cases in our modelling, since they are classified as extreme because of their posterior probability, the disproportionate weight attached to the news, and the jump size in price expectation.

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 about the learning process of market traders when an unprecedented and extreme event occurs.

First, the extreme event impact on stock price is delayed and needs confirmation of bad news to align single risk perceptions. The impact we estimate on 12 March 2020, which is equivalent to more than half a year (in business days) of continuing bad news on a daily basis, materialized after flows of bad news mounted over time and a price jump occurred when the lack of investor confidence was exacerbated by the decision to keep key ECB interest rates unchanged[15].

Second, once the extreme event occurred, agents tended to be more sensitive to all positive and negative news by changing their beliefs more frequently as they processed information on shortened past time windows. Investors are successively moving on close-to-rational expectations, assessing positive/negative market outcomes as almost equally likely. Hence, representativeness seems to be time

[15] Referring to calls for the ECB to go further and cut interest rates to ease borrowing costs for highly indebted eurozone countries, President Christine Lagarde said: "We are not here to close spreads, there are other tools and other actors to deal with these issues".

dependent and conditional on states of the economy. Our evidence on an extreme negative event proves that the deflation mechanism of less representative news seems to disappear after the shock materializes.

APPENDIX

Proposition A.1

For every time step $t = 1, \dots, T$ we have:

$$\pi_{g,t} + \pi_{b,t} = 1.$$

Proof. By definition (eq. [8]), we have:

$$\pi_{g,t} = \frac{\pi_{g,t-1} + n_{g,t}}{1 + n_{b,t} + n_{g,t}},$$

$$\pi_{b,t} = \frac{\pi_{b,t-1} + n_{b,t}}{1 + n_{b,t} + n_{g,t}}.$$

Since in the first time step $\pi_{g,1} + \pi_{b,1} = 1$, by applying a recursive argument, we infer that:

$$\pi_{g,t} + \pi_{b,t} = \frac{\pi_{g,t-1} + n_{g,t} + \pi_{b,t-1} + n_{b,t}}{1 + n_{b,t} + n_{g,t}} = \frac{1 + n_{g,t} + n_{b,t}}{1 + n_{b,t} + n_{g,t}} = 1.$$

Proposition A.2

For every time step t:

$$R_{g,t} > R_{b,t} \;\Leftrightarrow\; \pi_{g,t} > \pi_{g,t-1} \Leftrightarrow \pi_{g,t} > \frac{n_{g,t}}{N},$$

where $N = n_{g,t} + n_{b,t}$.

Proof. From eq. (8) we obtain:

$$\pi_{g,t} = \frac{\pi_{g,t-1} + n_{g,t}}{1 + N},$$

$$\pi_{g,t-1} = \frac{\pi_{g,t-2} + n_{g,t-1}}{1 + N}.$$

Combining the result of Prop. 1 with the definition of $R_{x,t} = \frac{\pi_{x,t}}{\pi_{x,t-1}}$ we infer that:

$$R_{g,t} = \frac{\pi_{g,t}}{\pi_{g,t-1}} = \frac{\pi_{g,t-1} + n_{g,t}}{1 + N} \cdot \frac{1 + N}{\pi_{g,t-2} + n_{g,t-1}} = \frac{\pi_{g,t-1} + n_{g,t}}{\pi_{g,t-2} + n_{g,t-1}},$$

$$R_{b,t} = \frac{\pi_{b,t}}{\pi_{b,t-1}} = \frac{1 - \pi_{g,t}}{1 - \pi_{g,t-1}} = \left(1 - \frac{\pi_{g,t-1} + n_{g,t}}{1 + N}\right) : \left(1 - \frac{\pi_{g,t-2} + n_{g,t-1}}{1 + N}\right).$$

By comparison, the inequality

$$R_{g,t} > R_{b,t}$$

can be rewritten as:

$$\frac{\pi_{g,t}}{\pi_{g,t-1}} > \frac{1 - \pi_{g,t}}{1 - \pi_{g,t-1}}.$$

Exploiting calculations, we have:

$$\pi_{g,t}(1 - \pi_{g,t-1}) > \pi_{g,t-1}(1 - \pi_{g,t}),$$

Hence, it follows that:

$$\pi_{g,t} - \pi_{g,t-1}\pi_{g,t} > \pi_{g,t-1} - \pi_{g,t-1}\pi_{g,t}.$$

Subsequently, after suitable cancellation, we obtain:

$$\pi_{g,t} > \pi_{g,t-1}$$

and the first chain of implications is proved. Moreover, recalling that

$$\pi_{g,t} = \frac{\pi_{g,t-1} + n_{g,t}}{1 + N}$$

and combining this with the previous equation, we conclude that:

$$\pi_{g,t} > \pi_{g,t-1} = (1 + N)\,\pi_{g,t} - n_{g,t} = \pi_{g,t} + N\,\pi_{g,t} - n_{g,t},$$

We can then infer the second chain of implications, namely:

$$\pi_{g,t} < \frac{n_{g,t}}{N}.$$

Proposition A.3

For every time step t:

$$\pi_{g,t}^{p} + \pi_{b,t}^{p} = 1.$$

Proof. According to eq. (9) we have:

$$\pi_{g,t}^{p} = \frac{\pi_{g,t-1} + n_{g,t}}{(\pi_{g,t-1} + n_{g,t}) + \delta(\pi_{b,t-1} + n_{b,t})},$$
$$\pi_{b,t}^{p} = \frac{\delta(\pi_{b,t-1} + n_{b,t})}{(\pi_{g,t-1} + n_{g,t}) + \delta(\pi_{b,t-1} + n_{b,t})}.$$

Adding up the previous equations, it follows that:

$$\pi_{g,t}^{p} + \pi_{b,t}^{p} = \frac{\pi_{g,t-1}+n_{g,t}}{(\pi_{g,t-1}+n_{g,t})+\delta(\pi_{b,t-1}+n_{b,t})} + \frac{\delta(\pi_{b,t-1}+n_{b,t})}{(\pi_{g,t-1}+n_{g,t})+\delta(\pi_{b,t-1}+n_{b,t})} = \frac{\pi_{g,t-1}+n_{g,t}+\delta(\pi_{b,t-1}+n_{b,t})}{(\pi_{g,t-1}+n_{g,t})+\delta(\pi_{b,t-1}+n_{b,t})} = 1.$$

Proposition A.4

Eq. (10) can be written as:

$$E^p_{h,t}(p_{t+1}) = \begin{cases} E_{h,t}(p_{t+1}) - \gamma_t \quad if\ \pi^p_{g,t-1} > \frac{1}{2} \wedge \pi^p_{g,t} < \frac{1}{2} \wedge \pi^p_{g,t-1} > n_{g,t}(N+1) - n_{g,t-1}, \\ E_{h,t}(p_{t+1}) + \gamma_t \quad if\ \pi^p_{g,t-1} < \frac{1}{2} \wedge \pi^p_{g,t} > \frac{1}{2} \wedge \pi^p_{g,t-1} < \frac{n_{g,t}}{N}\left(1 - n_{g,t-1} + n_{g,t}\right), \\ E_{h,t}(p_{t+1}) \quad if\ otherwise, \end{cases}$$

where $N = n_{g,t} + n_{b,t}$.

Proof. By mean of Prop. A.2, it easily follows that:

$$\pi^p_{g,t-1} > \pi^p_{b,t-1} \Leftrightarrow \pi^p_{g,t-1} > 1 - \pi^p_{g,t-1} \Leftrightarrow \pi^p_{g,t-1} > \frac{1}{2},$$

$$\pi^p_{g,t} < \pi^p_{b,t} \Leftrightarrow \pi^p_{g,t} < 1 - \pi^p_{g,t} \Leftrightarrow \pi^p_{g,t} < \frac{1}{2},$$

$$\pi^p_{b,t-1} > \pi^p_{g,t-1} \Leftrightarrow 1 - \pi^p_{g,t-1} > \pi^p_{g,t-1} \Leftrightarrow \pi^p_{g,t-1} < \frac{1}{2},$$

$$\pi^p_{b,t} < \pi^p_{g,t} \Leftrightarrow 1 - \pi^p_{g,t} < \pi^p_{g,t} \Leftrightarrow \pi^p_{g,t} > \frac{1}{2}.$$

Moreover, by denoting $N = n_{g,t} + n_{b,t}$ and applying the result of Prop. A.2, we have:

$$\frac{n_{b,t}}{n_{g,t}} = \frac{1 - n_{g,t}}{n_{g,t}} > \frac{\pi^p_{b,t-1} + n_{b,t-1}}{\pi^p_{g,t-1} + n_{g,t-1}} = \frac{1 - \pi^p_{g,t-1} + N - n_{g,t-1}}{\pi^p_{g,t-1} + n_{g,t-1}}.$$

Then, we obtain:

$$\left(1 - n_{g,t}\right) \cdot \left(\pi^p_{g,t-1} + n_{g,t-1}\right) > \left(1 - \pi^p_{g,t-1} + N - n_{g,t-1}\right) \cdot n_{g,t}\ ,$$

Hence, exploiting cancellation, we conclude that:

$$\pi^p_{g,t-1} > n_{g,t}(N+1) - n_{g,t-1}.$$

Finally, applying a completely analogous argument, we infer that:

$$\frac{n_{g,t}}{n_{b,t}} > \frac{\pi^p_{g,t-1} + n_{g,t}}{\pi^p_{b,t-1} + n_{b,t}} \quad \Leftrightarrow \quad \pi^p_{g,t-1} < \frac{n_{g,t}}{N}\left(1 - n_{g,t-1} + n_{g,t}\right).$$

References

Aït-Sahali, Y., Brandt, M.W., 2001. Variable Selection for Portfolio Choice. The Journal of Finance 56, 1297–1351. https://doi.org/10.1111/0022-1082.00369

Aït-Sahalia, Y., Andritzky, J., Jobst, A., Nowak, S., Tamirisa, N., 2012. Market response to policy initiatives during the global financial crisis. Journal of International Economics, Symposium on the Global Dimensions of the Financial Crisis 87, 162–177. https://doi.org/10.1016/j.jinteco.2011.12.001

Alfaro, L., Chari, A., Greenland, A.N., Schott, P.K., 2020. Aggregate and Firm-Level Stock Returns During Pandemics, in Real Time (No. w26950). National Bureau of Economic Research. https://doi.org/10.3386/w26950

Andrade, P., Ferroni, F., 2020. Delphic and odyssean monetary policy shocks: Evidence from the euro area. Journal of Monetary Economics. https://doi.org/10.1016/j.jmoneco.2020.06.002

Anufriev, M., Hommes, C., 2012a. Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments. American Economic Journal: Microeconomics 4, 35–64.

Anufriev, M., Hommes, C., 2012b. Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments. American Economic Journal: Microeconomics 4, 35–64. https://doi.org/10.1257/mic.4.4.35

Atchison, M.D., Butler, K.C., Simonds, R.R., 1987. Nonsynchronous Security Trading and Market Index Autocorrelation. The Journal of Finance 42, 111–118. https://doi.org/10.1111/j.1540-6261.1987.tb02553.x

Baker, S.R., Bloom, N., Davis, S.J., Kost, K.J., 2019. Policy News and Stock Market Volatility (No. w25720). National Bureau of Economic Research. 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). National Bureau of Economic Research. https://doi.org/10.3386/w26945

Bao, T., Hommes, C., Makarewicz, T., 2017. Bubble Formation and (In)Efficient Markets in Learning-to-forecast and optimise Experiments. The Economic Journal 127, F581–F609. https://doi.org/10.1111/ecoj.12341

Barberis, N., 2013. The Psychology of Tail Events: Progress and Challenges. American Economic Review 103, 611–616. https://doi.org/10.1257/aer.103.3.611

Bernard, V.L., Thomas, J.K., 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. Journal of Accounting and Economics 13, 305–340. https://doi.org/10.1016/0165-4101(90)90008-R

Birz, G., Lott, J.R., 2011. The effect of macroeconomic news on stock returns: New evidence from newspaper coverage. Journal of Banking & Finance 35, 2791–2800. https://doi.org/10.1016/j.jbankfin.2011.03.006

Bordalo, P., Gennaioli, N., Kwon, S.Y., Shleifer, A., 2020a. Diagnostic bubbles. Journal of Financial Economics. https://doi.org/10.1016/j.jfineco.2020.06.019

Bordalo, P., Gennaioli, N., Shleifer, A., 2020b. Memory, Attention, and Choice*. The Quarterly Journal of Economics 135, 1399–1442. https://doi.org/10.1093/qje/qjaa007

Bordalo, P., Gennaioli, N., Shleifer, A., 2018. Diagnostic Expectations and Credit Cycles. The Journal of 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. Journal of Economic Dynamics and Control 22, 1235–1274. 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

De Grauwe, P., Rovira Kaltwasser, P., 2012. Animal spirits in the foreign exchange market. Journal of Economic Dynamics and Control, Quantifying and Understanding Dysfunctions in Financial Markets 36, 1176–1192. https://doi.org/10.1016/j.jedc.2012.03.008

Duffie, D., 2005. Credit risk modeling with affine processes. Journal of Banking & Finance, Thirty Years of Continuous-Time Finance 29, 2751–2802. https://doi.org/10.1016/j.jbankfin.2005.02.006

Duffie, D., Singleton, K.J., 2003. Credit Risk, Princeton University Press.

Ezekiel, M., 1938. The Cobweb Theorem. Q J Econ 52, 255–280. https://doi.org/10.2307/1881734

Frankel, J.A., Froot, K.A., 1990. Chartists, Fundamentalists, and Trading in the Foreign Exchange Market. The American Economic Review 80, 181–185.

Gennaioli, N., Shleifer, A., Vishny, R., 2015. Neglected Risks: The Psychology of Financial Crises. American Economic Review 105, 310–314. https://doi.org/10.1257/aer.p20151091

Giglio, S., Maggiori, M., Stroebel, J., Utkus, S., 2020. Inside the Mind of a Stock Market Crash (No. w27272). National Bureau of Economic Research. https://doi.org/10.3386/w27272

Han, E., Tan, M.M.J., Turk, E., Sridhar, D., Leung, G.M., Shibuya, K., Asgari, N., Oh, J., García-Basteiro, A.L., Hanefeld, J., Cook, A.R., Hsu, L.Y., Teo, Y.Y., Heymann, D., Clark, H., McKee, M., Legido-Quigley, H., 2020. Lessons learnt from easing COVID-19 restrictions: an analysis of countries and regions in Asia Pacific and Europe. The Lancet 396, 1525–1534. https://doi.org/10.1016/S0140-6736(20)32007-9

Harris, L., 1989. The October 1987 S&P 500 Stock-Futures Basis. The Journal of Finance 44, 77–99. https://doi.org/10.1111/j.1540-6261.1989.tb02405.x

Hommes, C., 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. Journal of Economic Dynamics and Control 35, 1–24. https://doi.org/10.1016/j.jedc.2010.10.003

Hommes, C., Kopányi-Peuker, A., Sonnemans, J., 2020. Bubbles, crashes and information contagion in large-group asset market experiments. Exp Econ. 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. Journal of Economic Perspectives 5, 193–206. https://doi.org/10.1257/jep.5.1.193

Kahneman, D., Tversky, A., 1972. Subjective probability: A judgment of representativeness. Cognitive Psychology 3, 430–454. https://doi.org/10.1016/0010-0285(72)90016-3

Kumer Dey, A., Haq, T., Das, K., Panovska, I., 2020. Quantifying the impact of COVID-19 on the US stock market: An analysis from multi-source information. arXiv e-prints 2008, arXiv:2008.10885.

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

Mitchell, M.L., Mulherin, J.H., 1994. The Impact of Public Information on the Stock Market. The Journal of Finance 49, 923–950. https://doi.org/10.1111/j.1540-6261.1994.tb00083.x

Nagel, S., Xu, Z., 2019. Asset Pricing with Fading Memory (No. w26255). National Bureau of Economic Research. https://doi.org/10.3386/w26255

Peters, E., Slovic, P., 1996. The Role of Affect and Worldviews as Orienting Dispositions in the Perception and Acceptance of Nuclear Power1. Journal of Applied Social Psychology 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

Pritamani, M., Singal, V., 2001. Return predictability following large price changes and information releases. Journal of Banking & Finance 25, 631–656. https://doi.org/10.1016/S0378-4266(00)00091-1

Pruna, R.T., Polukarov, M., Jennings, N.R., 2020. Loss aversion in an agent-based asset pricing model. Quantitative Finance 20, 275–290. https://doi.org/10.1080/14697688.2019.1655784

Quarteroni, A., 2014. Numerical Models for Differential Problems, 2nd ed, MS&A. Springer-Verlag, Mailand. https://doi.org/10.1007/978-88-470-5522-3

Raberto, M., Cincotti, S., Focardi, S.M., Marchesi, M., 2001. Agent-based simulation of a ÿnancial market. Physica A 9.

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.

Staccioli, J., Napoletano, M., 2020. An agent-based model of intra-day financial markets dynamics. Journal of Economic Behavior & Organization. https://doi.org/10.1016/j.jebo.2020.05.018

Thaler, R.H., Tversky, A., Kahneman, D., Schwartz, A., 1997. The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test*. The Quarterly Journal of Economics 112, 647–661. https://doi.org/10.1162/003355397555226

Tversky, A., Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. Cognitive Psychology 5, 207–232. https://doi.org/10.1016/0010-0285(73)90033-9

Tables

Parameter	Label	Value
Number of periods	T	250
Number of traders	J	200
Initial population for each *h-th* heuristic	J_h	50
Risk-free rate	r	5%
Weak chartist coefficient	α	0.4
Strong chartist coefficient	β	1.3
Bias	b	+2.5%
Risk aversion	a	20
Discount factor	δ	0.8
Memory	m	6
Number of market activity-based news per trade	n_{mkt}	3
Weight of market activity-based news	φ_{mkt}	1
Number of public news per trade	n_p	2
Weight of public news	φ_p	2
Number of extreme events over T	$n_{ee,t}$	[0,10]
Weight of extreme events	φ_{ee}	36
Bad-based Jump amplitude	η_b	60
Good-based Jump amplitude	η_g	30

Table 1. Parameters' value and initial conditions.

Figures

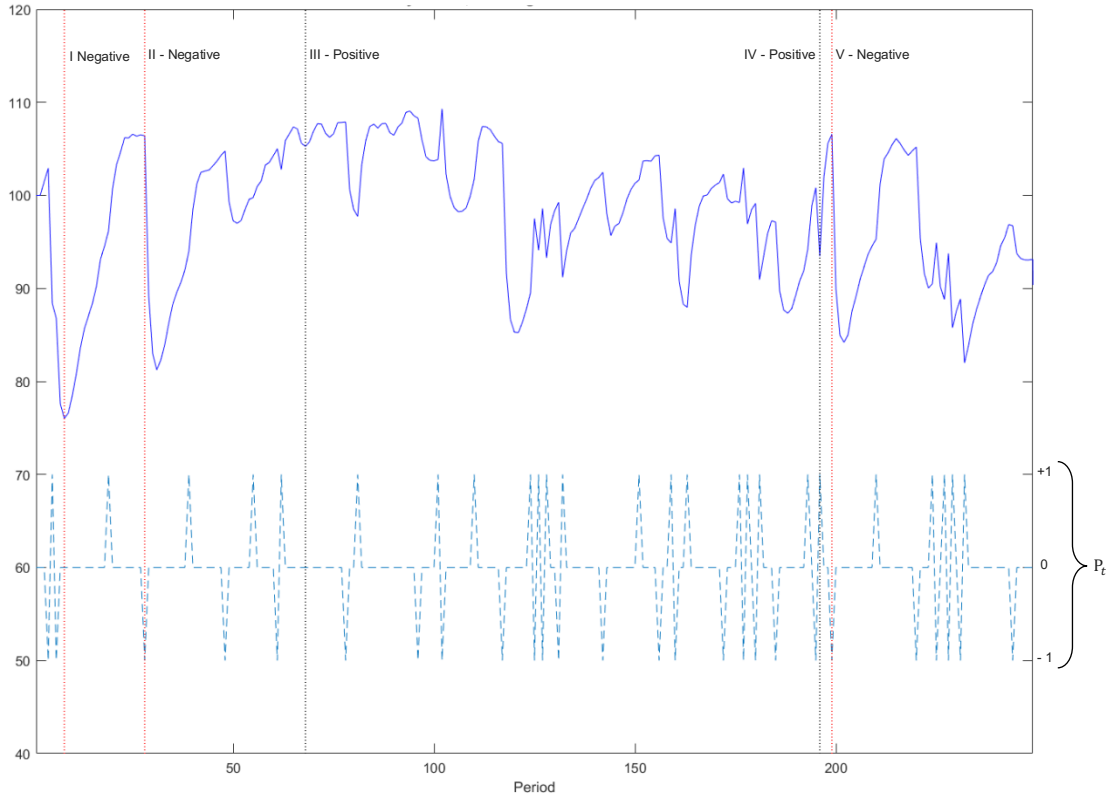

Figure 1. Price dynamics (solid line), overreactions in the price expectations mechanism P_t (dashed line) and positive/negative extreme events (dotted line).

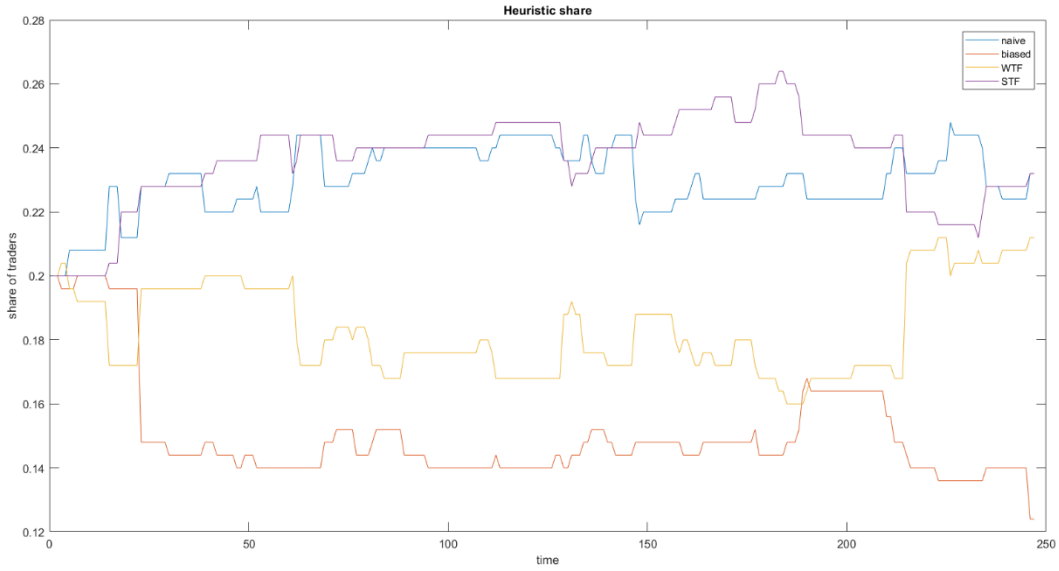

Figure 2. Heuristic share in the population: the blue line represents naïve traders, the red line depicts biased traders, the yellow line is the weak trend heuristic and the purple line shows the strong trend chartist.

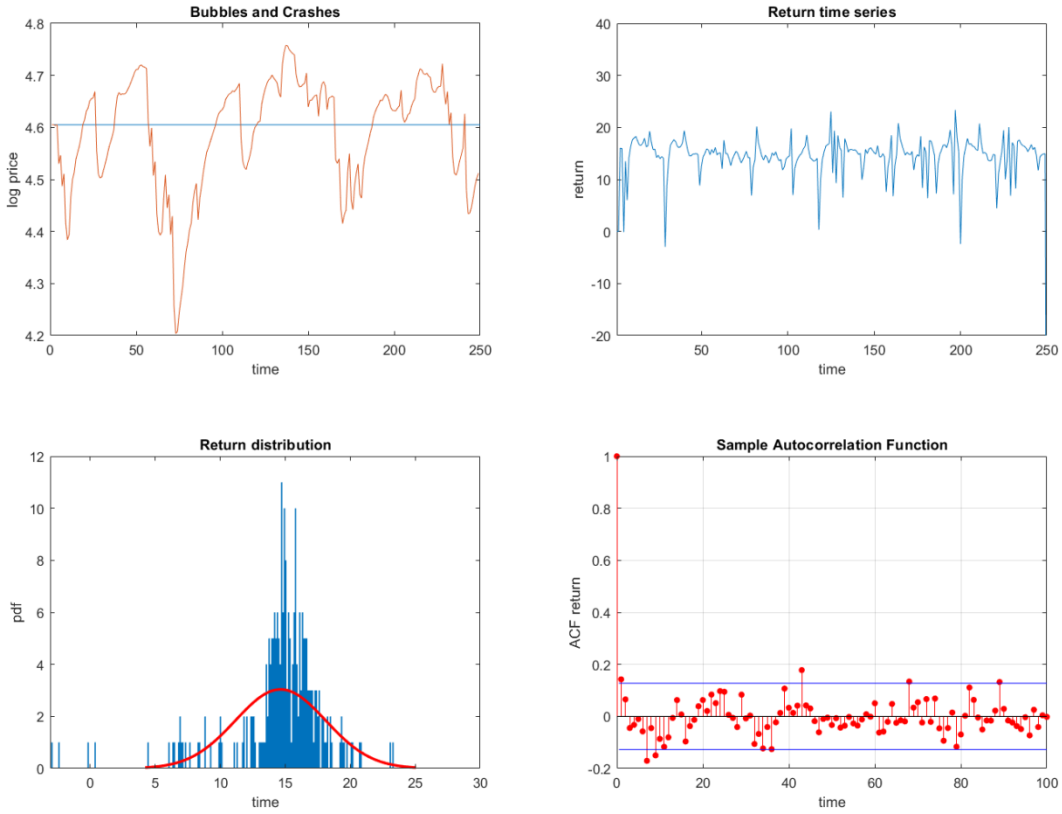

Figure 3. Matching the stylized facts. Top-left panel: log price dynamics. Top-right panel: asset returns. Bottom-left panel: distribution of normalized returns. Bottom-right panel: autocorrelation function of returns.

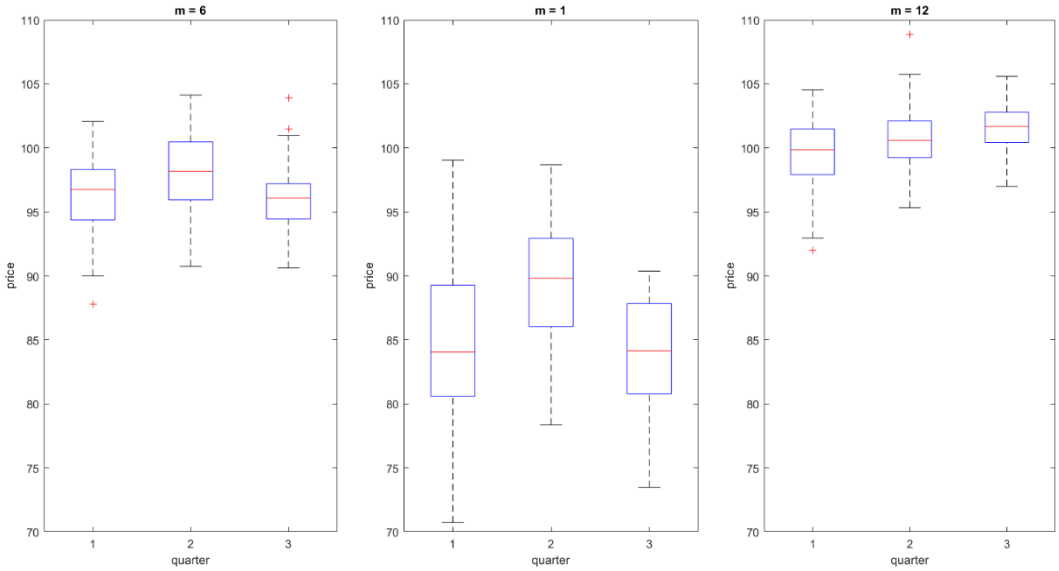

Figure 4: Quarterly price distribution of 12,500 observations with different traders' memory, $m = 6, 1, 12$.

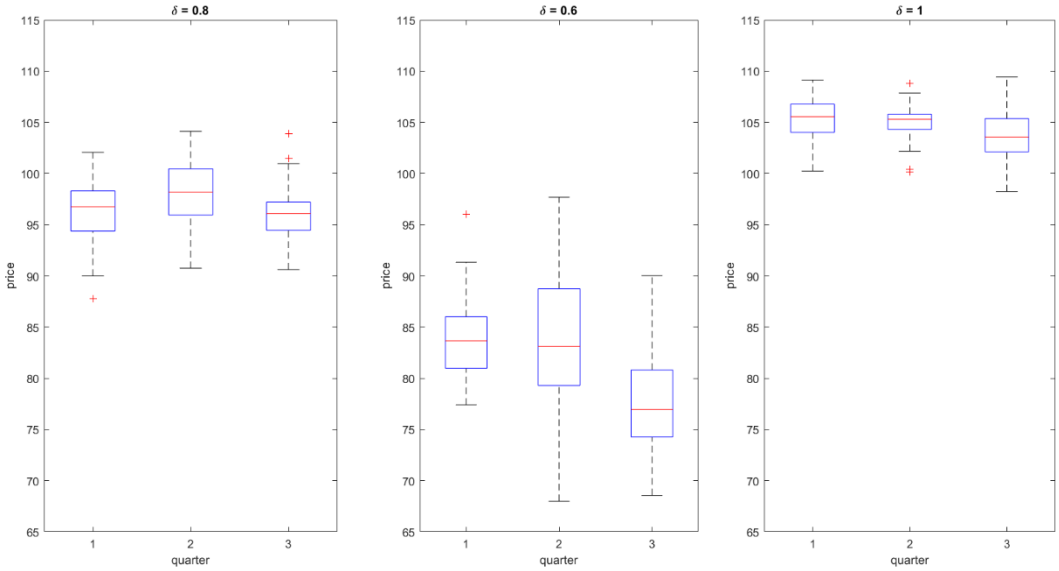

Figure 5: Quarterly price distribution of 12,500 observations with different discount parameter values, $\delta = 0.8, 0.6, 1$.

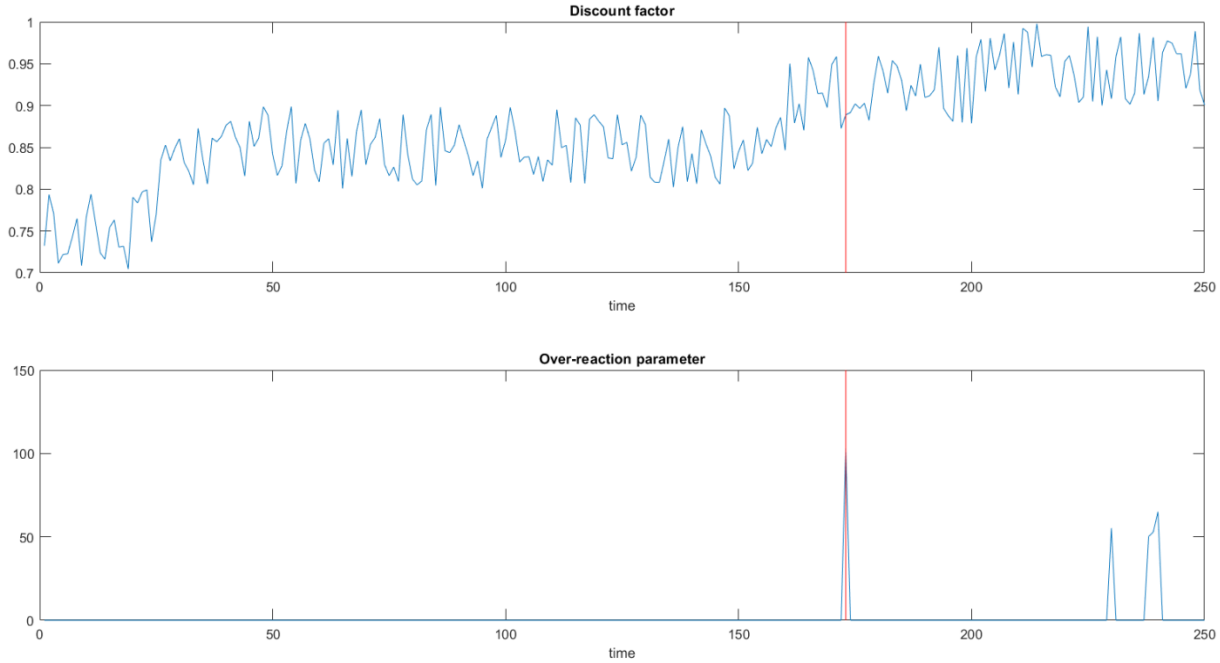

Figure 6. Time-varying discount factor and jump amplitude parameters.

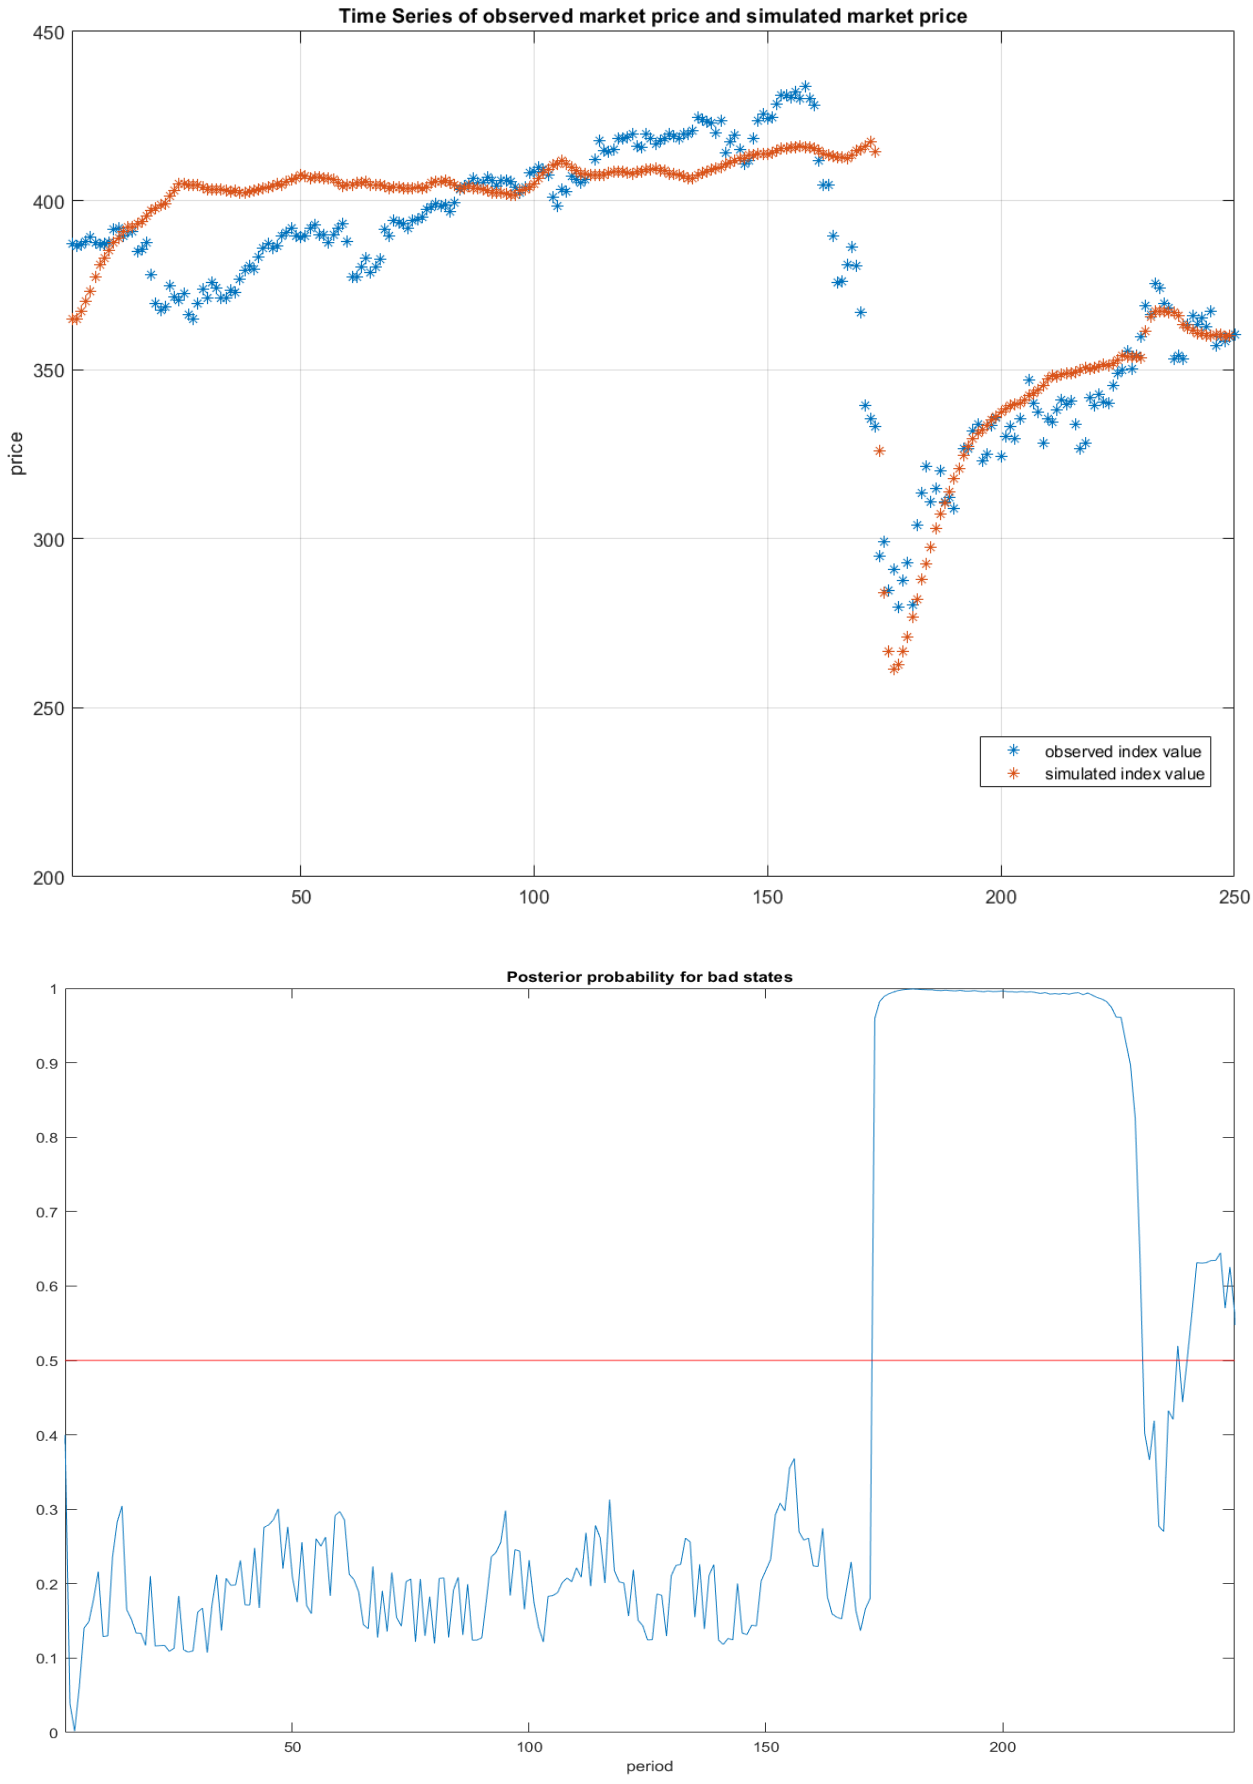

Figure 7. Observed vs. calibrated STOXX Europe 600 and corrected posterior probability of bad states.