Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Hidden Markov Modelling of Earthquake-Related Market Regimes and State-Dependent Forecasts: Evidence from the Turkish Stock Market

Gamze Özel

gamzeozl@hacettepe.edu.tr

Hacettepe University

Research Article

Keywords:

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11318417/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Hidden Markov Modelling of Earthquake-Related Market Regimes and State-Dependent Forecasts: Evidence from the Turkish Stock Market

Abstract

This study investigates whether major earthquakes are systematically associated with changes in the latent return and volatility regimes of the Borsa Istanbul 100 (BIST 100) index. The empirical sample contains 4,334 daily index observations from 5 January 2009 to 3 April 2026. A three-state Gaussian hidden Markov model is estimated for daily logarithmic returns, and the resulting states are interpreted ex post as crisis, uncertainty and normal regimes. The earthquake catalogue covers the common period from 19 June 2009 to 27 October 2025. Ninety-six events with magnitude $M \geq 5.5$ are reduced to 74 representative events through a 48-hour main-event filter and are aligned with the contemporaneous or next trading day. The crisis regime is characterized by a mean daily return of -0.79%, volatility of 3.45% and an expected duration of 5.3 days. Earthquake-linked trading days are distributed across the crisis, uncertainty and normal regimes in counts of 7, 35 and 32, respectively. Their distribution does not differ significantly from the unconditional state composition, and the reported magnitude-group analysis does not establish an association with next-day regime changes. State-conditional forecasts nevertheless reveal substantial differences in short-horizon downside risk. The pooled results do not establish a systematic earthquake–regime association, whereas the fitted model yields state-dependent predictive risk distributions. A time-varying-transition Markov-switching extension is proposed for future causal and predictive analysis.

Keywords: Hidden Markov model, regime switching, state-dependent forecasting, earthquake, BIST 100, volatility

JEL Classification: C22 , G14 , G17 , Q54

1 Introduction

Large earthquakes can affect financial markets through several channels that operate at different speeds. Physical destruction may interrupt production, transportation, communication, and energy networks, while insurance liabilities and reconstruction expenditures may alter firms' expected cash flows. Disasters may also increase fiscal financing requirements, sovereign risk, liquidity pressures, and investor uncertainty.

Consequently, the financial impact of an earthquake may differ across firms, sectors, and periods rather than producing a uniform average return response. Previous evidence shows that the magnitude of market reactions depends on factors such as geographical exposure, sectoral activity, firm size, pre-event volatility, investor composition, and post-disaster reconstruction expectations (Worthington & Valadkhani, 2004; Hood et al., 2013; Aksoy & Yılmaz, 2024; Pandey et al., 2024).

The response of a financial market may also evolve over time. A geographically distant or economically limited earthquake may leave the market in a relatively stable state, whereas a major disaster occurring during a period of macro-financial fragility may shift the market into a high-volatility state. Between these two cases, the market may enter an intermediate regime characterized by elevated uncertainty but without a persistent collapse in average returns. Evidence from the 2011 Tohoku earthquake and the 2023 Türkiye–Syria earthquakes indicates that investor behaviour, trading activity, market proximity, and sectoral exposure can generate heterogeneous and time-dependent reactions (Hood et al., 2013; Sakariyahu et al., 2023; Pandey et al., 2024). Most empirical studies examining the financial consequences of earthquakes employ event-study methods, abnormal-return models, intervention analysis, generalized autoregressive conditional heteroskedasticity models, causality tests, or counterfactual time-series approaches. These methods provide important evidence on short-run abnormal returns, volatility changes, and the causal impact of individual disasters (Worthington & Valadkhani, 2004; Cılek & Ergün, 2023; Aksoy & Yılmaz, 2024; Khan et al., 2024; Pandey et al., 2024). However, event windows generally measure deviations from a single underlying return-generating process. They do not explicitly allow the distribution of returns, market volatility, and transition dynamics to change across unobserved market conditions. This limitation is important because financial returns commonly display volatility clustering, fat tails, and periods of persistently high or low risk. The behaviour of financial returns may therefore be better represented by models that allow parameters to change according to an unobserved state (Schwert, 1989). Hamilton's (1989) Markov-switching model represents structural changes through a latent state governed by a first-order Markov process, while hidden Markov models provide a probabilistic mechanism for estimating latent-state sequences from observed data (Rabiner, 1989).

In finance, regime-switching models have been used to distinguish crash, recovery, low-growth, and bull-market conditions and to demonstrate that expected returns, risks, and portfolio allocations can vary considerably across regimes (Turner et al., 1989; Ang & Bekaert, 2002; Guidolin & Timmermann, 2007; Kim & Nelson, 1999). A regime-based approach is particularly appropriate for analysing earthquakes because an earthquake may influence the probability of moving into a stress regime without generating the same average return response on every event day. The latent state should not be interpreted as a directly observed economic or disaster variable. Instead, it summarizes recurring statistical configurations of return, volatility, and persistence. A major earthquake may therefore increase the probability of entering a high-risk regime, while trading restrictions, policy responses, prior market conditions, or reconstruction expectations may moderate the observed return within that regime. Time-varying transition probability models provide a natural extension by

allowing observable variables to affect the probabilities of moving between latent states (Filardo, 1994).

The BIST 100 provides an informative setting for such an analysis because Türkiye is highly exposed to seismic risk and the index represents the principal benchmark of the domestic equity market. The February 2023 Kahramanmaraş earthquakes produced substantial negative abnormal returns in many Borsa İstanbul sectors, although the magnitude and direction of the effects varied across industries and firms (Aksoy & Yılmaz, 2024; Pandey et al., 2024). Related research also reports significant changes in investor sentiment, BIST 100 dynamics, and banking-sector returns following the earthquakes (Cılek & Ergün, 2023; Sakariyahu et al., 2023). The 2023 episode further illustrates why market and disaster measurements cannot always be treated independently. Trading on Borsa İstanbul was suspended after substantial index declines and subsequently resumed under a set of market-supporting and regulatory measures. As a result, the observed return sequence around the disaster reflects not only the earthquake shock but also market closure, trading rules, liquidity conditions, and policy interventions. A narrow event window may therefore combine several related but analytically distinct mechanisms (Aksoy & Yılmaz, 2024; Khan et al., 2024). Against this background, the present study addresses three related research questions. First, can BIST 100 daily returns be represented by a small number of economically interpretable latent market regimes? Second, are trading days associated with major earthquakes disproportionately classified into the highest-risk regime, and do earthquake characteristics predict an immediate change in market state? Third, how does the most recently estimated regime affect short-horizon forecasts of future state probabilities and cumulative market returns?

The study contributes to the literature in four respects. First, it uses a long BIST 100 sample containing multiple seismic events instead of treating a single major earthquake as the only source of information. Second, it separates the estimation of latent market regimes from the subsequent examination of earthquake-related trading days. This design reduces the risk of defining market states solely according to known disaster dates. Third, the study reports transition probabilities, regime persistence, expected state durations, and regime-conditional forecasts, thereby extending the conventional focus on average abnormal returns. Fourth, it proposes a time-varying Markov-switching specification in which earthquake magnitude, depth, distance, and post-event indicators may enter either the conditional return equation or the transition probabilities. This distinction makes it possible to examine whether earthquakes affect returns within an existing regime or instead alter the probability of moving into a different regime.

The remainder of the article is organized as follows: Section 2 presents the theoretical background and reviews the related empirical literature. Section 3 describes the BIST 100 and earthquake data, the hidden Markov and Markov-switching specifications, the event-regime association tests, and the forecasting procedure. Section 4 reports the empirical findings. Section 5 discusses the economic implications and interpretation of the findings. Section 6 concludes.

2 Theoretical background and previous studies

2.1 Why disasters may generate state-dependent market responses

The effect of an earthquake on equity values can be written conceptually as a change in expected future cash flows, discount rates and risk premia. Damage to plants and infrastructure lowers expected operating cash flows for exposed firms, whereas construction, materials or selected service companies may benefit from reconstruction demand. Insurance losses, public expenditure and changes in sovereign risk can affect discount rates. Simultaneously, information uncertainty and investor attention may increase the risk premium required to hold equities. The aggregate index response is therefore a weighted combination of heterogeneous firm and sector effects.

These mechanisms imply two sources of nonlinearity. The first is event heterogeneity: magnitude, depth, location, fatalities, economic exposure and market closure differ across earthquakes. The second is market-state dependence: the same earthquake may have a larger financial effect when liquidity is already weak, the exchange rate is under pressure or global risk aversion is elevated. In this setting, a constant-parameter model averages over distinct conditions and can understate tail behavior. A latent-regime model does not by itself identify a causal earthquake effect. Its role is different: it provides a probabilistic description of market conditions and allows the researcher to ask whether earthquake-linked observations are unusually frequent in a stress state. Causal interpretation requires additional assumptions or a counterfactual design. Accordingly, the present study uses conditional and associational language throughout.

2.2 Markov-switching and hidden Markov foundations

Hamilton (1989) formalized regime switching for economic time series by allowing model parameters to depend on an unobserved finite-state Markov chain. Hidden Markov models use the same latent-state logic but describe observations through state-dependent emission distributions (Rabiner, 1989; Zucchini et al., 2016). In financial applications, these models are useful because returns may alternate between low-volatility, uncertainty and crisis conditions without the regime being observed directly.

Let $S_t \in \{1, \ldots, \mathrm{K}\}$ denote the unobserved market regime on trading day t. The regime sequence is assumed to follow a first-order Markov chain:

$$\Pr(S_t = j \mid S_{t-1} = i, S_{t-2}, \ldots) = \Pr(S_t = j \mid S_{t-1} = i) = p_{ij} \tag{1}$$

where p_{ij} is the probability of moving from regime i to regime j. For each initial regime, the transition probabilities satisfy

$$\sum_{j=1}^{K} p_{ij} = 1 \tag{2}$$

In the basic hidden Markov specification, the daily BIST 100 return r_t follows a state-dependent distribution:

$$r_t \mid S_t = k \sim \mathcal{N}(\mu_k, \sigma_k^2) \tag{3}$$

The parameters μ_k and σ_k^2 represent the conditional mean and variance of regime k. A Student-t emission distribution can also be considered when the return series displays pronounced excess kurtosis. The regimes are inferred from recurring combinations of return level, volatility and persistence; they are not defined in advance according to earthquake dates. A Hamilton-type Markov-switching autoregressive model extends the basic HMM by allowing the conditional mean to depend on lagged returns:

$$r_t = \mu(S_t) + \phi(S_t) r_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}\big(0, \sigma^2(S_t)\big) \tag{4}$$

This specification separates short-run return dependence from regime-specific market conditions. It also allows mean return, persistence and volatility to differ across stable, uncertainty and crisis regimes. Model estimation yields filtered and smoothed state probabilities. The filtered probability uses information available up to day t, whereas the smoothed probability uses the complete sample:

$$\Pr(S_t = k \mid r_1, \ldots, r_t) \tag{5}$$

$$\Pr(S_t = k \mid r_1, \ldots, r_T) \tag{6}$$

Filtered probabilities are appropriate for real-time monitoring and forecasting. Smoothed probabilities are more suitable for the retrospective identification of historical market regimes. The most likely complete state sequence can additionally be obtained through the Viterbi algorithm. The diagonal transition probabilities measure regime persistence. The expected duration of regime k is

$$D_k = \frac{1}{1 - p_{kk}} \tag{7}$$

where p_{kk} is the probability of remaining in the same regime. Multi-step regime forecasts follow directly from the transition matrix P:

$$\boldsymbol{\pi}_{t+h|t} = \boldsymbol{\pi}_{t|t} \mathbf{P}^h \tag{8}$$

where $\boldsymbol{\pi}_{t|t}$ is the vector of current filtered probabilities. This structure provides short-horizon estimates of the probability that the BIST 100 will remain in, or move toward, a high-volatility regime.

In this study, the theoretical quantities serve three purposes. State-dependent means and variances characterize the economic content of the latent regimes. Transition probabilities and expected durations measure the persistence of market stress. Finally, filtered probabilities generate short-horizon forecasts of future regimes and cumulative returns. Earthquake dates are matched with the estimated state sequence only after the regimes have been identified, reducing the risk of mechanically imposing disaster-related classifications on the return series.

2.3 Earthquake-dependent transition probabilities

A constant-transition-probability model assumes that the mechanism governing movements between latent market states is unchanged throughout the sample. This restriction is convenient, but it may be too strong when the arrival and severity of external shocks vary over time. Earthquakes differ in magnitude, focal depth, location, economic exposure and timing relative to the trading session. These characteristics may therefore alter the probability that the BIST 100 moves from a stable regime to an uncertainty or crisis regime. Following Filardo (1994), this possibility is represented by allowing the transition probabilities to depend on a vector of observable earthquake-related variables.

$$\Pr(S_t = j \mid S_{t-1} = i, \boldsymbol{z}_t) = \frac{\exp\left(\alpha_{ij} + \boldsymbol{\beta}'_{ij}\boldsymbol{z}_t\right)}{\sum_{l=1}^{K} \exp(\alpha_{il} + \boldsymbol{\beta}'_{il}\boldsymbol{z}_t)} \tag{9}$$

Here, S_t denotes the unobserved market regime at trading day t, while $\boldsymbol{z}_t$ is the vector of time-varying earthquake covariates. For a given origin state i, the coefficients describe how the covariates change the log-odds of moving to destination state j relative to a reference destination state. One destination category is normalized for identification, and the resulting probabilities are positive and sum to one across all possible destination states. Because the model is nonlinear, the sign of β_{ij} indicates the direction of the association in log-odds terms, whereas the economic magnitude should be evaluated using predicted probabilities or average marginal effects.

In this study, $\boldsymbol{z}_t$ may include an event-day indicator, standardized earthquake magnitude, focal depth, distance from major economic centres and indicators for the first trading days following an event. Earthquakes occurring after the market closes are aligned with the next trading day so that the information set is consistent with actual trading opportunities. Continuous earthquake characteristics are centred and scaled before estimation. Event-specific terms are set to zero on non-event days and are used jointly with an event indicator, which prevents zero values from being interpreted as physically meaningful magnitudes or distances. Earthquake variables may influence the market through two distinct channels. The first is a within-regime return effect, in which seismic information changes the conditional return while the latent state remains unchanged. A simple representation is

$$r_t = \mu_{S_t} + \phi_{S_t} r_{t-1} + \boldsymbol{\delta}'_{S_t}\boldsymbol{z}_t + \varepsilon_t, \qquad \varepsilon_t \mid S_t = k \sim \mathcal{N}(0, \sigma_k^2) \tag{10}$$

The second is a transition effect, in which the earthquake changes the probability of moving to another latent state through Equation (9). This second channel motivates the proposed extension. A severe earthquake may increase the likelihood of entering a high-risk regime even when the event-day return is moderated by a temporary market closure, regulatory intervention, liquidity support or offsetting macroeconomic news. The two channels should not be combined mechanically: the return-equation coefficients measure state-conditional changes in returns, whereas the transition coefficients measure changes in the probability of regime movement.

In a future empirical implementation, the constant-transition and earthquake-dependent specifications could be compared using the log-likelihood, Akaike information criterion and Bayesian information criterion. The joint contribution of the

earthquake covariates could additionally be assessed with likelihood-ratio or Wald tests, subject to the regularity conditions of the fitted models. Interpretation would require predicted transition probabilities for meaningful earthquake scenarios rather than raw logit coefficients alone. These comparisons and covariate-dependent estimates are not reported in the present analysis. This makes it possible to quantify, for example, how the predicted probability of moving from a stable regime to a crisis regime changes with earthquake magnitude while the previous market state and the remaining covariates are held fixed. The earthquake-dependent specification is not, by itself, a causal design. The estimated relationships may also reflect contemporaneous policy measures, market closures, global risk conditions or other macro-financial shocks. Accordingly, the coefficients are interpreted as conditional associations after controlling for the preceding market state. A causal interpretation would require an additional counterfactual strategy or stronger identifying assumptions. Within these limits, fitting the extension would provide a test of whether seismic information contributes to changes in BIST 100 regime dynamics.

2.4 Earthquakes and stock markets

International evidence does not support a universal market response to earthquakes. Ferreira and Karali (2015), using aggregate indices from 35 markets, found broad resilience and emphasized that market development, trade exposure and earthquake characteristics can mediate effects in market-specific ways. Scholtens and Voorhorst (2013) similarly documented heterogeneity in domestic market reactions. These findings caution against treating magnitude alone as a sufficient predictor of financial losses.

Research on the 2023 Türkiye–Syria earthquakes provides stronger evidence of short-run disruption but again reveals substantial heterogeneity. Pandey et al. (2024) reported negative abnormal returns for most sectors and showed that larger and less volatile firms were more resilient. Aksoy and Yilmaz (2024) found that corporate sustainability performance could moderate downside risk among BIST firms. Sakariyahu et al. (2023) documented negative reactions in the stock markets of Türkiye's major trading partners, especially in geographically or economically close countries. Khan, Cifuentes-Faura and Shahbaz (2024), using Bayesian structural time series, estimated a negative causal effect on the Borsa İstanbul index relative to a counterfactual path. Cılek and Ergün (2023) and Büyükoğlu, Özpolat and Nakıpoğlu Özsoy (2024) examined BIST 100 dynamics through causality and cross-quantile approaches. The empirical literature therefore establishes that disaster effects can be negative, persistent for several days and heterogeneous across firms, sectors and countries. What remains less developed is the link between earthquake observations and latent market-state dynamics. Event studies identify abnormal returns around a prespecified date, but they do not directly estimate the probability that the market has shifted into a recurring stress regime or the expected duration of that regime.

2.5 Regime models in Borsa İstanbul research

Regime-switching methods have been used in Borsa İstanbul for questions unrelated to earthquakes. Günay (2019) applied a Markov regime-switching VAR to BIST 100 volatility and credit-risk indicators. Burhan and Acar (2021) used an HMM to investigate return predictability under the adaptive market hypothesis. Samırkaş (2021) identified bull and bear regimes with a two-state Markov-switching model and examined time-varying transition probabilities linked to exchange-rate changes. Kaya and Yarbaşı (2021) estimated a three-regime MS-GARCH model for BIST 100 volatility. These studies establish the empirical relevance of nonlinear state changes in the Turkish market, but they do not connect seismic events to the probability of occupying or entering a high-risk state.

Table 1: Selected previous studies and the research gap

Study	Market / event	Method	Main finding
Ferreira & Karali (2015)	35 national stock indices; major earthquakes	Panel/event-based analysis	Mostly resilient aggregate markets; mediating factors are market-specific.
Günay (2019)	BIST 100 and Turkish credit-risk indicators	Markov regime-switching VAR	Credit-risk effects differ between low- and high-volatility regimes.
Burhan & Acar (2021)	BIST 100, 1988–2017	Hidden Markov model	Return predictability changes across latent market conditions.
Samırkaş (2021)	BIST 100, 2010–2020	MS(2)-AR(0), TVTP	Bull/bear regimes and exchange-rate-linked transition probabilities.
Sakariyahu et al. (2023)	Türkiye's major trading partners	Event study	Negative short-run reactions, stronger for closer markets.
Pandey et al. (2024)	382 Turkish listed firms	Firm-level event study	Negative responses for most sectors; resilience varies by firm risk.
Aksoy & Yilmaz (2024)	BIST All and BIST 100 firms	Event study and sustainability comparison	Sustainability performance is associated with downside resilience.
Khan et al. (2024)	Borsa İstanbul index	Bayesian structural time series	Negative index effect relative to a counterfactual path.

The reviewed studies leave a specific gap. Earthquake research usually estimates abnormal returns or a counterfactual price path, while BIST regime studies usually condition on macro-financial variables rather than seismic events. The present design joins these literatures by first estimating market regimes from BIST 100 returns and then examining earthquake-linked days within the estimated state process. This sequence avoids defining "crisis" directly from the earthquake indicator and reduces the risk of circular classification. Our research hypotheses are given below:

H1: BIST 100 daily returns are generated by multiple latent regimes with distinct means and volatilities,

H2: The highest-risk regime is less persistent than the lower-risk regimes but displays substantially greater return dispersion,

H3: Earthquake-linked trading days are more likely than ordinary trading days to occur in the crisis regime,

H4: Larger earthquakes are more likely to be followed by an immediate market-regime change,

H5: State-conditional forecasts differ materially, even when the same transition matrix governs all future paths.

3 Data and methodology

3.1 BIST 100 data and return construction

The financial dataset contains 4,334 daily BIST 100 observations from 5 January 2009 to 3 April 2026. The analysis uses actual trading observations; weekends and exchange holidays are not interpolated. One anomalous record dated 1 January 2013 was removed during data quality control. Because the first price observation is used to construct the initial difference and additional alignment is required, the decoded HMM sequence contains 4,332 return observations.

$$r_t = 100\,[\ln(P_t) - \ln(P_{t-1})] \tag{11}$$

where P_t is the BIST 100 closing level on trading day t and r_t is the daily percentage logarithmic return. The use of log returns makes multi-day outcomes additive before conversion back to cumulative simple returns. The July 2020 removal of two zeros from BIST index values is checked before return calculation so that a mechanical change in the quoted index level is not interpreted as an economic return.

3.2 Earthquake catalogue, filtering and trading-day alignment

Table 2 summarizes the financial and earthquake data used in the empirical analysis. The earthquake catalogue and the financial series overlap from 19 June 2009 to 27 October 2025. The primary threshold is $M \geq 5.5$. Ninety-six events meet this threshold. Because aftershock sequences can produce several events within a short interval, treating every observation as an independent market shock would over-weight clustered seismic episodes. A 48-hour main-event filter is therefore applied, leaving 74 representative events for the event-regime analysis. Events occurring during trading hours are assigned to the contemporaneous trading day. Events occurring after the market close, at weekends or on exchange holidays are mapped to the next available trading day. If several retained events map to the same trading day, the day is counted once in the regime-distribution analysis and event characteristics are summarized by the largest magnitude. The magnitude analysis distinguishes M 5.5–6.49 from $M \geq 6.5$ events. The latter group contains only eight observations, so exact or Monte Carlo tests are preferred to asymptotic tests when expected cell frequencies are small.

Table 2: Financial and earthquake data used in the empirical analysis

Component	N	Period	Analytical role
BIST 100 price observations	4,334	5 Jan 2009–3 Apr 2026	Daily trading observations
Decoded return observations	4,332	5 Jan 2009–3 Apr 2026	After return construction and cleaning
Earthquake–market overlap	—	19 Jun 2009–27 Oct 2025	Common observation period
Earthquakes with $M \geq 5.5$	96	Common period	Before clustering control
Filtered representative events	74	Common period	After 48-hour main-event filter
M 5.5–6.49 events	66	Common period	Magnitude subgroup
$M \geq 6.5$ events	8	Common period	High-magnitude subgroup

3.3 Gaussian hidden Markov model

Let $S_t \in \{1,2,3\}$ denote the unobserved market state at trading day t. The state follows a first-order, time-homogeneous Markov chain. Conditional on S_t=k, the daily return follows a state-specific Gaussian distribution:

$$r_t \mid (S_t = k) \sim \mathcal{N}(\mu_k, \sigma_k^2), \qquad k = 1, 2, 3 \tag{12}$$

The transition probability from state i to state j is

$$p_{ij} = \Pr(S_t = j \mid S_{t-1} = i), \qquad p_{ij} \geq 0, \qquad \sum_j p_{ij} = 1 \tag{13}$$

The model parameters consist of the initial state probabilities, the 3×3 transition matrix P, and the state-specific means and variances. They are estimated by maximum likelihood through the expectation–maximization algorithm (Dempster et al., 1977). Multiple starting values are used to reduce sensitivity to local optima. The forward recursion evaluates the likelihood efficiently:

$$\alpha_t(j) = f_j(r_t) \sum_i \alpha_{t-1}(i) p_{ij} \tag{14}$$

where $f_j(r_t)$ is the Gaussian density in state j.

The backward recursion combines future information with the filtered probabilities to obtain the smoothed state probability $\gamma_t(k) = \Pr(S_t = k \mid r_1, \ldots, r_T)$. The decoded state is the state with the highest smoothed posterior probability, but all forecasts use probabilities rather than treating decoded labels as error-free observations. State labels are assigned only after estimation. The state with the largest volatility and a negative mean is labelled crisis; the medium-volatility state is labelled uncertainty; and the lowest-volatility state is labelled normal. This ex post ordering prevents the numerical state labels generated by the optimizer from determining the economic interpretation. Persistence is summarized by the diagonal elements p_{kk}. Under a geometric duration distribution, the expected number of consecutive trading days in state k is

$$\mathrm{E}(D_k) = \frac{1}{1 - p_{kk}} \tag{15}$$

3.4 State-number selection and proposed Markov-switching extension

Two-, three- and four-state HMMs were considered. Preliminary information criteria favored additional flexibility in the four-state solution, but the additional state divided the broad lower-risk portion of the sample and weakened economic interpretability. The three-state specification is therefore used as the primary parsimonious model because it preserves a clear crisis–uncertainty–normal ordering and supports transparent event and forecast interpretation. The four-state result is treated as a sensitivity indication rather than as the primary classification.

The HMM is the empirical core of this article. A Hamilton-type Markov-switching regression provides a natural extension when lagged returns or external variables are introduced. A general state-dependent return equation is

$$r_t = \mu(S_t) + \phi(S_t) r_{t-1} + \boldsymbol{\beta}(S_t)' \boldsymbol{x}_t + \varepsilon_t, \qquad \varepsilon_t \mid S_t \sim \mathcal{N}\big(0, \sigma^2(S_t)\big) \tag{16}$$

where $\boldsymbol{x}_t$ can include an earthquake-day indicator, magnitude, depth, distance, exchange-rate changes, market volatility or global risk controls. Equation (16) tests whether a covariate changes returns within a regime. A distinct question is whether the covariate changes the probability of moving between regimes. This can be represented through time-varying transition probabilities:

$$p_{ij,t} = \frac{\exp\big(\boldsymbol{\eta}_{ij}' \boldsymbol{z}_t\big)}{\sum_l \exp(\boldsymbol{\eta}_{il}' \boldsymbol{z}_t)} \tag{17}$$

where $\boldsymbol{z}_t$ contains earthquake and macro-financial information. This distinction is central to the proposed extension: an earthquake may have a small direct coefficient in the return equation but still increase the probability of entering a high-volatility state. The time-varying transition probability (TVTP) and autoregressive extensions are not estimated in the present article; they specify the next stage of the research design.

3.5 Event–regime association tests

The first event analysis compares the observed distribution of the 74 filtered earthquake-linked trading days across the three HMM regimes with the unconditional regime shares in the decoded sample. A Pearson chi-square goodness-of-fit statistic is used because the expected count in each regime exceeds five. Failure to reject equality indicates insufficient evidence of a difference under this test; it does not establish equality of the distributions. The second analysis cross-classifies event days by magnitude group and regime. Because the $M \geq 6.5$ category is small, the result is interpreted descriptively. The third analysis records whether the decoded regime changes on the next trading day. Fisher's exact test evaluates whether this immediate change indicator is associated with the two magnitude categories. These tests do not establish

causality; they evaluate whether simple event characteristics align with the estimated state sequence.

3.6 State-conditional forecasting

Let $\boldsymbol{\pi}_t$ denote the row vector of filtered state probabilities at time t. The h-step regime probability forecast is obtained from the transition matrix:

$$\boldsymbol{\pi}_{t+h|t} = \boldsymbol{\pi}_t \mathbf{P}^h \tag{18}$$

Because the final filtered state is not imposed as a single real-time condition in this article, forecasts are reported as three transparent scenarios: the current state is crisis, uncertainty or normal. This avoids presenting one endpoint classification as certain and makes the implications of each starting condition directly comparable. The one-step predictive return density is a mixture of state-specific Gaussian densities weighted by the forecast state probabilities. Thirty-trading-day cumulative return scenarios are generated by Monte Carlo simulation. At each step, the next state is sampled from P and the return is sampled from the corresponding Gaussian emission. The cumulative simple return is $\exp(\sum_h r_{t+h}/100) - 1$. One hundred thousand paths are simulated with a fixed random seed. The intervals reflect transition and emission uncertainty conditional on the estimated parameters; they do not include parameter-estimation uncertainty and should not be interpreted as investment advice.

4 Empirical results

4.1 Results of HMM

We report the characteristics of the three BIST 100 regimes in Table 3. The selected three-state HMM produces a sharply ordered volatility structure. Table 3 shows that the crisis regime contains 315 decoded return observations, or 7.27% of the 4,332-day state sequence. Its mean daily log return is -0.79% and its standard deviation is 3.45%. The uncertainty regime contains 2,238 days (51.66%), with a positive mean of 0.20% and volatility of 1.16%. The normal regime contains 1,779 days (41.07%), with a mean of 0.19% and volatility of 0.50%. The near equality of the uncertainty and normal mean returns shows that the main distinction between these two states is risk rather than average return.

Table 3: Estimated characteristics of the three BIST 100 regimes

Regime	Days	Share (%)	Mean return (%)	Std. dev. (%)	Interpretation
Crisis	315	7.27	-0.79	3.45	High dispersion and negative mean
Uncer-tainty	2,238	51.66	0.20	1.16	Intermediate volatility
Normal	1,779	41.07	0.19	0.50	Low-volatility market state

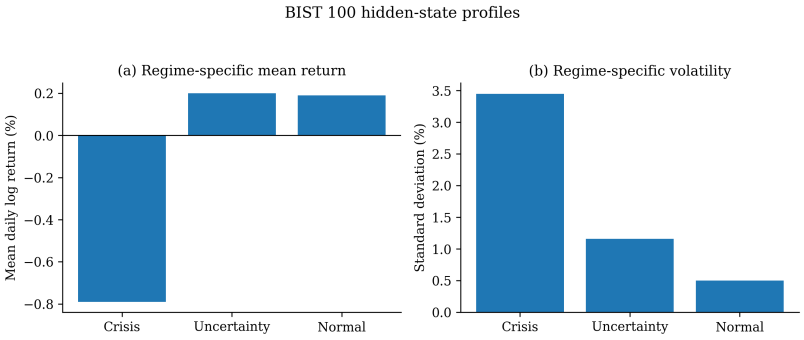


Fig. 1 Regime-specific mean returns and volatility

BIST 100 hidden state profiles are shown in Figure 1. The crisis-state standard deviation is almost seven times the normal-state standard deviation. This difference is consistent with H1 and the dispersion component of H2 within the fitted model. The state labels are not judgments about market quality. They are concise descriptions of the fitted return distributions. Figure 1 presents the estimated mean return and volatility profiles of the three latent BIST 100 regimes. The crisis regime is clearly distinguished by both a negative average daily log return of approximately -0.8% and a standard deviation of about 3.45%. Its volatility is nearly three times that of the uncertainty regime and almost seven times that of the normal regime, indicating that crisis periods combine pronounced losses with exceptionally large return fluctuations. By contrast, the uncertainty and normal regimes have similar positive mean returns, approximately 0.20% and 0.19%, respectively. Their main distinction therefore arises from volatility: the standard deviation of the uncertainty regime is about 1.15%, more than twice the 0.50% observed in the normal regime. These fitted profiles are consistent with the distinct latent conditions described in H1; they do not constitute a formal test of the number of states. They also confirm the substantially greater return dispersion of the highest-risk regime stated in H2. However, the persistence component of H2 should be evaluated using the estimated transition probabilities and expected regime durations rather than Figure 1 alone. Estimated transition matrix and expected regime durations (days) by HMM are presented in Table 4.

4.2 Transition probabilities and persistence

Table 4: Estimated transition matrix and expected regime durations

Origin regime	To crisis	To uncertainty	To normal	Expected duration (days)
Crisis	0.8110	0.1890	0.0000	5.29
Uncertainty	0.0306	0.9176	0.0518	12.14
Normal	0.0059	0.0577	0.9364	15.72

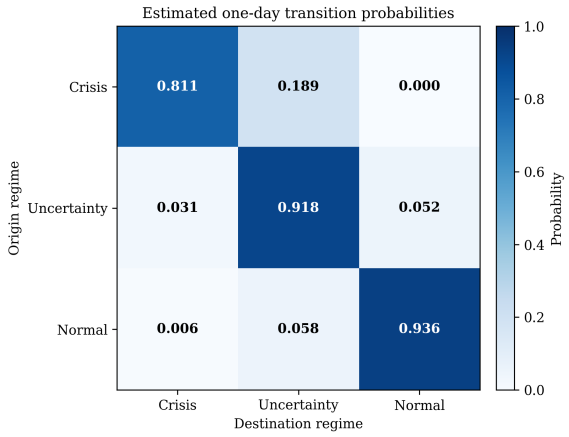


Fig. 2 One-day transition probability matrix

One-day transition probability matrix is presented in Figure 2. All three regimes are persistent, but their persistence differs. Conditional on being in crisis, the probability of remaining there on the next trading day is 0.811, corresponding to an expected duration of 5.29 days. The estimated direct crisis-to-normal transition probability rounds to zero; exits from crisis therefore occur predominantly through uncertainty in this fitted model. This ordered path is economically plausible because markets generally pass through a period of elevated uncertainty before returning to low volatility. The uncertainty regime has a self-transition probability of 0.918 and an expected duration of 12.14 days. Its probability of moving to normal is 0.052, while the probability of moving to crisis is 0.031. The normal regime is the most persistent, with $\Pr(N \to N)$=0.936 and an expected duration of 15.72 days. The stationary distribution implied by the estimated transition matrix is 9.36% crisis, 49.96% uncertainty and 40.69% normal, which is broadly consistent with the decoded sample shares.

4.3 Earthquake-linked trading days and regime membership

This subsection examines whether earthquake occurrence and earthquake magnitude are systematically associated with the latent market regimes identified by the HMM. Each retained earthquake is matched to the relevant BIST 100 trading day, taking account of weekends, public holidays and events occurring after market closure. Two complementary questions are considered. First, the distribution of earthquake-linked trading days across the crisis, uncertainty and normal regimes is compared with the unconditional regime distribution of the full return sample. This comparison provides a direct test of H3. Second, the analysis evaluates whether larger earthquakes are more frequently followed by an immediate change in the decoded market state, thereby testing H4. Because some regime–magnitude combinations contain small numbers of observations, both chi-square and exact testing procedures are used where appropriate. These analyses are interpreted as conditional association tests rather than causal

estimates, since earthquake characteristics, market conditions and policy responses may jointly affect the observed regime classification.

Earthquake-linked trading days across HMM regimes are presented in Table 5. Among the 74 filtered $M \geq 5.5$ events, 7 earthquake-linked trading days are classified as crisis, 35 as uncertainty and 32 as normal. These correspond to event-day shares of 9.46%, 47.30% and 43.24%. The unconditional decoded shares are 7.27%, 51.66% and 41.07%. A chi-square goodness-of-fit test gives χ^2=0.845 with p=0.655. Thus, the earthquake-day distribution does not differ significantly from the overall state composition, and H3 is not supported for the pooled multi-event sample.

Table 5: Earthquake-linked trading days across HMM regimes

Event group	Total	Crisis	Uncertainty	Normal	Crisis share (%)
All filtered $M \geq 5.5$ events	74	7	35	32	9.46
M 5.5–6.49	66	6	31	29	9.09
$M \geq 6.5$	8	1	4	3	12.50

Distribution of filtered earthquake events across BIST 100 regimes is presented in Figure 3.

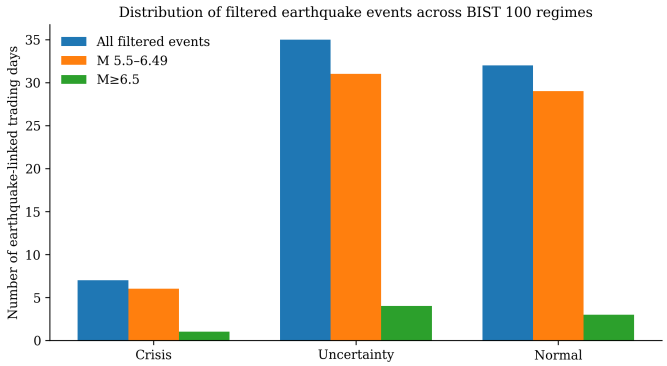


Fig. 3 Earthquake-event counts by magnitude group and market regime

As seen in Figure 3, the magnitude split yields a similar pattern. Of the 66 events in the M 5.5–6.49 group, 6 are aligned with crisis, 31 with uncertainty and 29 with normal. Of the eight $M \geq 6.5$ events, one is aligned with crisis, four with uncertainty and three with normal. Only five of the 74 event days are followed by a change in the decoded state. Fisher's exact test comparing magnitude group and next-day state change yields p=0.952. H4 is therefore not supported. This null result should not be interpreted as evidence that severe earthquakes have no financial impact. The test asks a narrower question: across all retained $M \geq 5.5$ events, does magnitude alone predict an immediate transition in the aggregate index regime? The reported

analysis does not provide evidence of such an association. Exceptional events may still generate economically large losses, as demonstrated by firm-level and counterfactual studies of the 2023 earthquake. The pooled result indicates that earthquake magnitude is insufficient without information about exposure, location, market conditions and policy response.

4.4 State-conditional regime forecasts

Table 6 reports state-conditional forecasts of future regime probabilities obtained from successive powers of the estimated HMM transition matrix. Each forecast assumes that the market is located with certainty in the indicated initial regime at the forecast origin. When the forecast starts from the crisis regime, the probability of remaining in crisis is 81.1% after one trading day, but declines to 38.7% after five days and 20.1% after ten days. Most of this probability shifts toward the uncertainty regime rather than directly toward the normal regime, suggesting that market stress tends to dissipate through an intermediate adjustment phase. The uncertainty regime is also highly persistent, with a one-day continuation probability of 91.8%, while the normal regime has the highest one-day persistence at 93.6%. Nevertheless, the influence of the initial state gradually weakens as the forecast horizon increases, and the conditional distributions move toward the long-run regime probabilities. These results support H5 by showing that short-horizon market-state forecasts differ materially according to the regime prevailing at the forecast origin.

Table 6: State-conditional HMM probability forecasts

Initial regime	Horizon	Pr(crisis)	Pr(uncert.)	Pr(norm.)
Crisis	1	0.811	0.189	0.000
Crisis	5	0.387	0.544	0.069
Crisis	10	0.201	0.615	0.184
Uncertainty	1	0.031	0.918	0.052
Uncertainty	5	0.090	0.715	0.195
Uncertainty	10	0.105	0.604	0.291
Normal	1	0.006	0.058	0.936
Normal	5	0.030	0.225	0.745
Normal	10	0.055	0.345	0.600

State conditional HMM regime forecasts are presented in Figure 4.

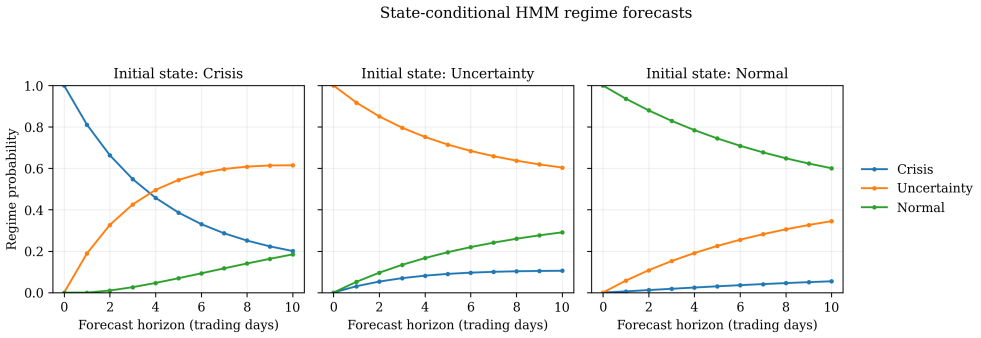


Fig. 4 Ten-trading-day regime probability forecasts under alternative starting states

The forecast paths in Figure 4 illustrate how persistence and transition ordering interact. Starting from crisis, the probability of remaining in crisis is 0.811 after one day and 0.201 after ten days. The uncertainty state becomes dominant by the fifth day, while the normal-state probability increases gradually. Starting from uncertainty, the probability of that state remains 0.604 at day ten, and the normal-state probability rises to 0.291. Starting from normal, the normal-state probability remains 0.600 at day ten, while the crisis probability is only 0.055. The forecasts converge toward the stationary distribution as the horizon increases, but convergence is not immediate. Consequently, the current regime contains useful short-horizon information even though all scenarios share the same long-run transition matrix. This result supports H5.

Table 7 reports the one-day predictive moments obtained by combining the estimated transition probabilities with the state-specific return distributions. The reported means and standard deviations are therefore predictive mixture moments rather than the parameters of the current regime alone. Starting from the crisis state produces a negative one-day expected return of -0.603% and a predictive standard deviation of 3.171%, reflecting both the high probability of remaining in crisis and the substantial dispersion associated with that regime. By contrast, the uncertainty and normal states generate positive expected returns and considerably lower predictive volatility. The normal state provides the most favourable one-day risk profile, with an expected return of 0.185%, a standard deviation of 0.623% and a 93.6% probability of remaining in the normal regime.

Table 7: One-day predictive moments and regime probabilities

Current state	**One-day mean (%)**	**One-day SD (%)**	**Pr(crisis)**	**Pr(uncert.)**	**Pr(norm.)**
Crisis	-0.603	3.171	0.811	0.189	0.000
Uncertainty	0.169	1.281	0.031	0.918	0.052
Normal	0.185	0.623	0.006	0.058	0.936

To examine longer-horizon forecast uncertainty, 30-trading-day cumulative returns are simulated using 100,000 Monte Carlo paths for each initial regime. At each step, the latent state is first advanced according to the estimated transition matrix, after which the daily return is drawn from the Gaussian distribution associated with the simulated state. The results in Table 8 show that the forecast distribution depends strongly on the regime prevailing at the forecast origin. Starting from crisis yields an average cumulative return of only 0.03%, a 46.0% probability of a negative return and an 89.9% probability of visiting the crisis regime during the forecast horizon. Starting from uncertainty or normal produces higher expected returns, narrower prediction intervals and lower downside probabilities. In particular, the normal-state forecast has the highest mean return and the lowest probability of a negative 30-day outcome. These findings support H5 by demonstrating that initial market conditions materially affect both expected performance and forecast dispersion, even though the same estimated transition mechanism governs all simulated paths.

Table 8: Thirty-trading-day Monte Carlo forecast summaries. Loss probability denotes a negative cumulative return; crisis visitation refers to the future 30 trading days, excluding the initial state

Initial regime	**Mean (%)**	**Median (%)**	**2.5%**	**97.5%**	**Loss prob. (%)**	**Crisis visit (%)**
Crisis	0.03	1.06	-24.44	19.97	46.0	89.9
Uncertainty	3.50	4.40	-17.52	18.89	28.1	47.8
Normal	4.28	5.09	-14.01	16.88	20.5	35.0

Figure 5 presents the state-conditional 30-trading-day return forecasts. The crisis-start scenario has an expected 30-day cumulative return close to zero, a median of 1.06% and a 95% simulation interval from -24.44% to 19.97%. The probability of a negative cumulative return is 46.0%. The near-zero mean does not imply a stable path: it reflects the combination of large crisis-state dispersion and a high probability of moving from crisis into uncertainty during the forecast horizon. The uncertainty-start scenario has a mean cumulative return of 3.50%, a 95% interval from -17.52% to 18.89% and a 28.1% probability of a negative outcome. The normal-start scenario has the narrowest downside interval and the lowest negative-return probability, although the possibility of visiting crisis during 30 trading days is still 35.0%.

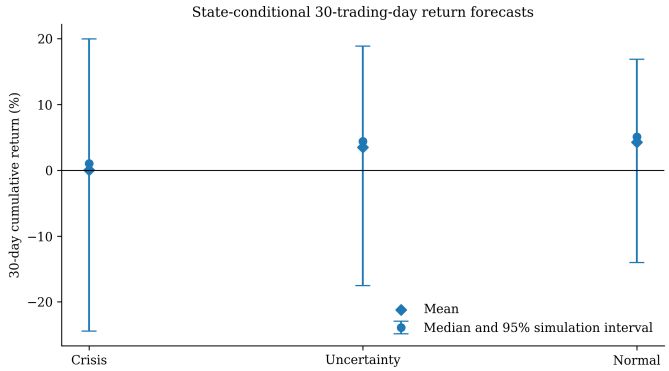


Fig. 5 State-conditional cumulative-return forecasts over 30 trading days

These simulations show why a point forecast alone is inadequate. Regime persistence and the probability of state transitions materially affect the distribution of future returns. These simulations are intended to illustrate regime-conditioned forecast dispersion and should not be interpreted as real-time investment forecasts, since they do not incorporate parameter uncertainty, transaction costs, macroeconomic covariates or model-estimation risk.

5 Discussion

This study develops a regime-based analysis of the relationship between earthquakes and BIST 100 daily returns over the 2009–2026 period. The empirical results reveal a pronounced three-regime market structure consisting of crisis, uncertainty, and normal conditions. These regimes differ substantially in their conditional means, volatility levels, persistence, and expected durations, providing a model-based description of heterogeneity in the return distribution. The findings are therefore consistent with the broader regime-switching literature, which emphasizes that financial-market behaviour changes across tranquil and turbulent periods. They also complement previous BIST 100 studies that identified bull–bear or low–high-volatility regimes using hidden Markov, Markov-switching, and MS-GARCH models. The crisis regime is relatively rare, negative on average, and characterized by exceptionally high return dispersion. Its estimated mean daily return is -0.79%, its volatility is 3.45%, and its expected duration is approximately 5.3 trading days. By contrast, the uncertainty and normal regimes account for most observations and exhibit positive average returns with substantially lower volatility. The uncertainty regime is the most frequently observed intermediate state, whereas the normal regime is the most persistent low-volatility condition. These findings support the view that market risk evolves through distinct latent configurations rather than through gradual changes around one unconditional mean and variance.

The estimated transition matrix provides additional information beyond the regime-specific return parameters. Although the crisis regime is not absorbing, it is sufficiently persistent to produce clusters of elevated market stress. More importantly,

the estimated transition structure indicates that the market does not generally move directly from crisis to normal conditions. Instead, crisis periods tend to be followed by the uncertainty regime, which functions as the principal adjustment channel between high- and low-risk market conditions. The uncertainty regime should therefore not be interpreted as a residual or weakly defined category. It is both the dominant market state and the main bridge through which stress accumulates or dissipates. From a risk-monitoring perspective, a fall in the probability of crisis should not automatically be interpreted as a complete return to normal market conditions. The earthquake-related findings are more cautious than the conclusions commonly obtained from studies focusing on a single catastrophic event. Among the 74 representative earthquakes with magnitude $M \geq 5.5$, earthquake-linked trading days are not statistically overrepresented in the crisis regime. Their regime distribution is statistically consistent with the unconditional composition of the BIST 100 market states. Similarly, the division of earthquakes into M5.5–6.49 and $M \geq 6.5$ groups does not significantly explain whether the decoded market regime changes on the following trading day. Accordingly, the pooled evidence does not support the hypotheses that earthquake-linked trading days are systematically more likely to occur in crisis or that earthquake magnitude alone predicts an immediate regime transition.

These results should not be interpreted as evidence that severe earthquakes have no financial consequences. Rather, they show that the occurrence and magnitude of an earthquake are insufficient, on their own, to explain market-wide regime changes across a long multi-event sample. This conclusion does not contradict previous event studies reporting negative abnormal returns following the 2023 Kahramanmaraş earthquakes. The two approaches address different questions. Event studies generally estimate the effect of one exceptional disaster over a narrow event window and may use firm-level or sector-level exposure measures. The present analysis pools numerous earthquakes over more than sixteen years and asks whether earthquake days, considered as a general class, systematically coincide with a latent crisis state in the aggregate BIST 100 index. A highly destructive event may generate large financial losses, while many other earthquakes of similar magnitude may have limited implications for listed-company cash flows, investor expectations, or market liquidity. The pooled null results also demonstrate the importance of the proposed earthquake-dependent Markov-switching extension. Magnitude is an incomplete measure of financial exposure because earthquakes with similar magnitudes may differ greatly in depth, distance from production centres, affected population, infrastructure damage, timing relative to the trading session, and institutional response. A time-varying transition probability specification can incorporate these characteristics together with pre-event volatility, exchange-rate movements, global risk conditions, market closure, and regulatory intervention. Such a model can evaluate whether seismic and macro-financial variables alter the probability of moving from normal or uncertainty into crisis after conditioning on the preceding market state. It can also allow the same earthquake characteristic to have different implications depending on the origin regime. An event may have a limited effect during normal conditions but become destabilizing when the market is already in an uncertainty state.

The forecasting results further demonstrate the usefulness of regime identification even though the pooled earthquake association is weak. State-conditional forecasts show that the current market regime materially affects short-horizon risk distributions. When the forecast begins in crisis, the probability of remaining in crisis is high over the next trading day, but uncertainty becomes the dominant state as the forecast horizon increases. This pattern is consistent with a gradual recovery process in which market stress passes through an intermediate adjustment state before normal conditions are restored. When the forecast begins in the normal regime, short-run crisis probabilities remain low, although the cumulative probability of visiting crisis becomes non-negligible over a thirty-trading-day horizon. The Monte Carlo simulations lead to a similar conclusion. Starting from the crisis regime produces the lowest expected thirty-day return, the widest prediction interval, and the highest probability of a negative cumulative return. Starting from uncertainty generates a more favourable but still relatively dispersed forecast distribution, whereas the normal regime produces the highest expected return and the lowest downside probability. Thus, even when all simulated paths are governed by the same estimated transition matrix, their predicted returns and risks differ substantially according to the regime prevailing at the forecast origin. This result supports the use of latent-state models as probabilistic risk-monitoring tools.

The forecast analysis converts the estimated transition structure into quantities that are directly interpretable for risk monitoring. Starting from crisis, the expected path is not an immediate return to normal; uncertainty becomes dominant over the next several days. Starting from normal, crisis remains unlikely at short horizons, but its cumulative visitation probability becomes non-negligible over a month. These are scenario forecasts, not directional trading signals. Their value lies in quantifying the consequences of current-state uncertainty for future risk distributions.

6 Conclusions

The fitted three-state Gaussian hidden Markov model distinguishes crisis, uncertainty and normal regimes in BIST 100 returns. The regimes differ primarily in volatility and persistence, and the estimated transition structure associates exits from crisis predominantly with the intermediate uncertainty state. The pooled event analysis does not establish a systematic association between earthquake-linked trading days and crisis-state membership. These findings concern aggregate regime classifications and do not rule out substantial effects of individual disasters on particular firms or sectors.

State-conditional forecasts translate the estimated transition structure into distributions of future regimes and cumulative returns. Their differences describe predictive risk conditional on the fitted model and initial state; they do not demonstrate out-of-sample forecasting superiority. Earthquake-dependent transition and return-equation specifications remain proposed extensions requiring estimation and validation. The framework is therefore an applied probabilistic description of market regimes and conditional risk, with cautious interpretation of the earthquake associations.

References

Aksoy, M., & Yilmaz, M. K. (2024). Stock market response to natural disasters: Does corporate sustainability performance make difference?. *Journal of Cleaner Production, 454*, 142294. https://doi.org/10.1016/j.jclepro.2024.142294

Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies, 15*(4), 1137–1187. https://doi.org/10.1093/rfs/15.4.1137

Burhan, H. A., & Acar, E. (2021). Adaptive market hypothesis and return predictability: A hidden Markov model application in Borsa Istanbul. *Sosyoekonomi, 29*(48), 31–58. https://doi.org/10.17233/sosyoekonomi.2021.02.02

Büyükoğlu, B., Özpolat, A., & Nakıpoğlu Özsoy, F. (2024). The effects of natural disaster on financial markets: The empirical analysis of 2023 earthquakes in Türkiye. *Akademik Araştırmalar ve Çalışmalar Dergisi, 16*(30), 186–197. https://doi.org/10.20990/kilisiibfakademik.1439250

Çılek, A., & Ergün, M. (2023). The impact of the 2023 Kahramanmaraş earthquake on BIST100 and BIST Bank Index: Evidence from Toda–Yamamoto causality test. *PressAcademia Procedia, 17*(1), 92–100. https://doi.org/10.17261/Pressacademia.2023.1760

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B, 39*(1), 1–38. https://doi.org/10.1111/j.2517-6161.1977.tb01600.x

Ferreira, S., & Karali, B. (2015). Do earthquakes shake stock markets?. *PLOS ONE, 10*(7), e0133319. https://doi.org/10.1371/journal.pone.0133319

Filardo, A. J. (1994). Business-cycle phases and their transitional dynamics. *Journal of Business & Economic Statistics, 12*(3), 299–308. https://doi.org/10.1080/07350015.1994.10524545

Guidolin, M., & Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control, 31*(11), 3503–3544. https://doi.org/10.1016/j.jedc.2006.12.004

Günay, S. (2019). An analysis through credit default swap, asset swap and zero-volatility spreads: Coup attempt and BIST 100 volatility. *Borsa Istanbul Review, 19*(2), 158–170. https://doi.org/10.1016/j.bir.2018.11.001

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica, 57*(2), 357–384. https://doi.org/10.2307/1912559

Hood, M., Kamesaka, A., Nofsinger, J. R., & Tamura, T. (2013). Investor response to a natural disaster: Evidence from Japan's 2011 earthquake. *Pacific-Basin Finance Journal, 25*, 240–252. https://doi.org/10.1016/j.pacfin.2013.09.006

Kaya, A., & Yarbaşı, İ. Y. (2021). Forecasting of volatility in stock exchange markets by MS-GARCH approach: An application of Borsa Istanbul. *Journal of Research in Economics, Politics and Finance, 6*(1), 16–35. https://doi.org/10.30784/epfad.740815

Khan, K., Cifuentes-Faura, J., & Shahbaz, M. (2024). Do earthquakes shake the stock market? Causal inferences from Turkey's earthquake. *Financial Innovation, 10*, 145. https://doi.org/10.1186/s40854-024-00664-w

Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime switching*. MIT Press.

Pandey, D. K., Kumar, S., Kumari, V., & Alahdal, W. M. (2024). Quaking the stock market: Event study evidence on the Turkey–Syria earthquake. *International Review of Economics & Finance, 93*, 182–194. https://doi.org/10.1016/j.iref.2024.03.023

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE, 77*(2), 257–286. https://doi.org/10.1109/5.18626

Sakariyahu, R., Lawal, R., Oyekola, O., Dosumu, O. E., & Adigun, R. (2023). Natural disasters, investor sentiments and stock market reactions: Evidence from Turkey–Syria earthquakes. *Economics Letters, 228*, 111153. https://doi.org/10.1016/j.econlet.2023.111153

Samırkaş, M. C. (2021). Ayı piyasası mı, boğa piyasası mı? Markov rejim değişim modeli ile Borsa İstanbul için bir analiz. *Insan ve Toplum Bilimleri Arastirmalari Dergisi, 10*(3), 2227–2249. https://doi.org/10.15869/itobiad.880539

Scholtens, B., & Voorhorst, Y. (2013). The impact of earthquakes on the domestic stock market. *Earthquake Spectra, 29*(1), 325–337. https://doi.org/10.1193/1.4000103

Schwert, G. W. (1989). Why does stock market volatility change over time?. *The Journal of Finance, 44*(5), 1115–1153. https://doi.org/10.1111/j.1540-6261.1989.tb02647.x

Turner, C. M., Startz, R., & Nelson, C. R. (1989). A Markov model of heteroskedasticity, risk, and learning in the stock market. *Journal of Financial Economics, 25*(1), 3–22. https://doi.org/10.1016/0304-405X(89)90094-9

Worthington, A. C., & Valadkhani, A. (2004). Measuring the impact of natural disasters on capital markets: An empirical application using intervention analysis. *Applied Economics, 36*(19), 2177–2186. https://doi.org/10.1080/0003684042000282489

Zucchini, W., MacDonald, I. L., & Langrock, R. (2016). *Hidden markov models for time series: An introduction using r* (2nd ed.). Chapman & Hall/CRC.