National Bank of Moldova

Working Paper 1/2026

CONSTRUCTING A FINANCIAL CYCLE INDEX FOR THE REPUBLIC OF MOLDOVA

Alexandru Fala and Vlad Tampiza

2026

Constructing a Financial Cycle Index for the Republic of Moldova

Alexandru Fala and Vlad Tampiza

This publication is available on the website of the National Bank of Moldova: www.bnm.md

The views expressed in this paper are those of the authors and do not necessarily reflect the official position of the National Bank of Moldova.

Acknowledgments. The authors would like to thank Daniela Mugrabi Otero Farah for the careful peer review of an earlier version of the paper and for her valuable comments and suggestions, which helped improve the manuscript. The final text remains the sole responsibility of the authors.

Reproduction of this publication, in whole or in part, is permitted provided that the customary bibliographic citation is included, with appropriate attribution to the authors and the title of the paper.

Abstract

This study constructs a composite financial cycle index for the Republic of Moldova to support the assessment of cyclical systemic risk and the calibration of macroprudential policy instruments. Using quarterly data for 2007Q2–2025Q4, the cyclical components of selected macro-financial variables are extracted through Complete Ensemble Empirical Mode Decomposition with Adaptive Noise. The extracted cycles are broadly robust to alternative specifications and stable under quasi-real-time estimation. The comparison of candidate indices with different compositions and aggregation schemes supports the selection of an equally weighted index comprising real credit, real residential property prices, and the sign-adjusted current account balance, whereas time-varying weights provide no clear advantage and the real interest rate may generate misleading signals during periods of high inflation. The resulting index provides a reasonable and transparent framework that can be used to support macroprudential policy decisions.

Keywords: financial cycle; cyclical systemic risk; composite index; Complete Ensemble Empirical Mode Decomposition with Adaptive Noise; macroprudential policy

JEL Classification: C43, E32, E44, G28

Contents

Introduction **4**

I. Literature review **6**

II. Methodology **14**

III. Results **27**

IV. Conclusions **51**

Bibliography **53**

Appendix 1. EMD and CEEMDAN methodology **57**

Appendix 2. Decomposition results **62**

Introduction

The global financial crisis demonstrated that vulnerabilities can accumulate within the financial system even during periods of apparently favourable macroeconomic conditions. Rapid credit expansion, increasing asset prices, and growing leverage may reinforce each other during the upward phase of the financial cycle and subsequently amplify economic downturns when accumulated imbalances begin to unwind. These developments strengthened the role of macroprudential policy in identifying and mitigating systemic risks.

One of the main instruments introduced under the Basel III framework is the countercyclical capital buffer (CCyB), which requires banks to accumulate additional capital when cyclical systemic risk is increasing. The credit-to-GDP gap is commonly used as a reference indicator for CCyB calibration. However, its signalling performance differs across countries, and the European Systemic Risk Board recommends that designated authorities complement it with country-specific indicators and expert judgement. Credit developments, property prices, external imbalances, and other macro-financial variables may therefore provide relevant additional information on the position of the financial cycle.

The financial cycle is not directly observable, and its empirical measurement remains subject to considerable uncertainty. Different cycle-extraction methods may produce different estimates of its duration, amplitude, and turning points. Frequency-based filtering approaches generally require the prior specification of a frequency range, while conventional spectral methods typically rely on stationary or previously detrended series, although time-varying spectral and wavelet approaches can accommodate certain forms of non-stationarity. Such requirements may be restrictive when the underlying macro-financial variables are persistent, nonlinear, or non-stationary. Moreover, the construction of a composite indicator requires additional decisions concerning the selection, normalization, and aggregation of its components.

These challenges are especially relevant for the Republic of Moldova. The available macro-financial time series are relatively short, the domestic financial system is predominantly bank-based, and sufficiently long and consistent series for a broader range of financial asset prices are not available. These characteristics affect data availability and the selection of indicators used to capture the domestic financial cycle. Structural changes and exposure to large external shocks may also alter the amplitude and persistence of cyclical fluctuations over time and therefore complicate their empirical identification. These features provide important context for the analysis but do not, by themselves, determine the choice of cycle-extraction method, which is based on the methodological trade-offs discussed in Section II. In addition, the absence of a sufficiently detailed official domestic crisis chronology limits the application of validation methods that require a predefined classification of crisis and non-crisis periods.

Against this background, the objective of the study is to construct and evaluate a composite financial cycle index for the Republic of Moldova that can support the assessment of the build-up and release of cyclical systemic risk. The analysis uses quarterly data for 2007Q2–2025Q4 and applies Complete Ensemble Empirical Mode Decomposition with Adaptive Noise to extract the cyclical components of the selected macro-financial variables. The robustness checks indicate that the resulting cyclical estimates are broadly stable under alternative specifications and over time.

Several versions of the composite indicator are examined by combining credit and residential property prices with the current account balance and the real interest rate and by applying alternative aggregation schemes. The resulting indices are assessed based on their cyclical characteristics, including

duration, historical timing relative to the GDP gap, and the economic interpretation of their component contributions. The comparison supports the selection of an equally weighted three-component index comprising real credit, real residential property prices, and the sign-adjusted current account balance.

Overall, the findings suggest that this three-component specification provides a reasonable and transparent framework for assessing cyclical systemic risk in the Republic of Moldova. The resulting index can be used to support the calibration of macroprudential policy instruments. Its level, direction, and component contributions should also be assessed in conjunction with the credit-to-GDP gap, broader macro-financial evidence, and indicators of banking-sector resilience.

The remainder of the paper is organized as follows. Section I reviews the conceptual and empirical literature on financial-cycle measurement, variable selection, cycle extraction, and aggregation. Section II presents the data and variable selection and describes the methodological framework used to extract, adjust, normalize, aggregate, and evaluate the cyclical components. Section III presents the descriptive analysis, empirical results, robustness assessments, and selection of the preferred composite indicator. Section IV concludes and outlines the implications for macroprudential surveillance and future applications.

I. Literature review

Conceptual framework of the financial cycle

The global financial crisis highlighted the importance of the interactions between the financial system and the real economy. The accumulation of vulnerabilities during periods of rapid credit expansion and rising asset prices can amplify subsequent economic downturns through losses in the banking sector, tighter credit supply, and adverse feedback effects on real activity. This experience strengthened the case for macroprudential policies aimed at identifying and mitigating systemic risks before they materialize.

One of the main macroprudential instruments introduced within the Basel III framework is the countercyclical capital buffer (CCyB), developed by the Basel Committee on Banking Supervision. The CCyB enables designated authorities to require banks to accumulate additional capital during periods of excessive credit growth and elevated cyclical systemic risk. The buffer can subsequently be released when risks materialize or financial conditions deteriorate, thereby supporting the capacity of banks to absorb losses and continue lending to the real economy.

In 2014, the European Systemic Risk Board (ESRB) issued recommendations on the operationalization of the countercyclical capital buffer, primarily addressed to the designated authorities of EU Member States (ESRB, 2014). The document synthesizes empirical evidence into an operational framework that helps national authorities understand the advantages and limitations of the CCyB and other instruments. The recommendation emphasizes the following principles:

- Empirical evidence identifies the credit-to-GDP gap as a useful common reference indicator for the build-up of cyclical systemic risk across EU Member States;
- The signalling performance of the credit-to-GDP gap differs across countries, which supports the use of additional country-specific indicators and expert judgement;
- Complementary indicators, including property prices, credit developments, measures of external imbalances, and other relevant macro-financial variables, may improve the assessment of cyclical systemic risk.

The policy debate that followed the 2008–2009 global financial crisis encouraged designated authorities to complement the credit-to-GDP gap with broader sets of country-specific indicators capable of capturing different dimensions of cyclical systemic risk. In parallel, the concept of the financial cycle became increasingly prominent in both the academic literature and macroprudential policy frameworks. However, the empirical boundary between the financial cycle and cyclical systemic risk remains relatively blurred. The two concepts are frequently operationalized using closely related sets of macro-financial indicators. For instance, ESRB (2014), in its *Recommendation on Guidance for Setting Countercyclical Buffer Rates*, refers to both cyclical systemic risk and financial-cycle developments within the same policy framework.

This overlap becomes particularly evident when comparing the composition of composite indicators proposed in the literature and used by macroprudential authorities. As shown in Table 2, indicators labelled as measures of the financial cycle and those explicitly designed to capture cyclical systemic risk often rely on broadly similar variables. Moreover, the distinction is not determined simply by

whether the underlying variables are narrowly financial. Broader macro-financial indicators, such as the current account balance, are included in measures explicitly targeting cyclical systemic risk, as in Lang et al. (2019), Vella (2024) and Banka Slovenije (2026), but are also incorporated into indicators described as measures of the financial cycle, as in Kupkovič and Šuster (2020) and Czech National Bank (2023).

Given that the term *financial cycle* remains more prevalent in the empirical literature reviewed in this study, we retain it as the main organizing concept throughout the paper. In this study, the financial cycle is treated as a key observable manifestation of the time-varying accumulation and unwinding of systemic vulnerabilities and, therefore, as a central component of cyclical systemic risk. Under this interpretation, a composite financial-cycle indicator can provide policy-relevant information for assessing the build-up of cyclical systemic vulnerabilities and, ultimately, for informing the calibration of macroprudential instruments such as the countercyclical capital buffer.

The concept of the financial cycle builds on earlier analyses of recurrent interactions between credit expansion, asset-price increases, leverage, and subsequent financial contraction. More recent studies have sought to formalize these dynamics empirically and to develop indicators suitable for financial-stability surveillance. Nevertheless, Hiebert et al. (2014) and Russo (2022) emphasize that no universally accepted definition or measurement framework has emerged.

A major development in the literature has been the transformation of the financial cycle from a primarily descriptive concept into a policy-relevant framework. Post-crisis research emphasizes that financial imbalances can accumulate even when inflation and institution-level indicators appear stable, which strengthens the case for monitoring system-wide cyclical vulnerabilities. Accordingly, the financial cycle has become an important input for macroprudential surveillance and for the calibration of instruments such as the CCyB. As noted by Gadanecz and Jayaram (2015), an effective macroprudential framework requires the identification of system-wide vulnerabilities in relation to the position and amplitude of the financial cycle.

The empirical literature can be organized into three closely related strands: the analysis of individual macro-financial variables, the construction of composite financial-cycle indicators, and the assessment of their predictive or policy relevance. These strands are complementary, since the analysis of individual variables generally informs indicator construction, while predictive and historical assessments are used to evaluate the resulting measures (Table 1).

Table 1: Main strands of the empirical financial cycle literature

Authors	Analysis of separate variables	Construction of a composite cyclical indicator	Predictive assessment
Borio (2014)	+		
Che and Shinagawa (2014)	+		
S. Chen and Svirydzenka (2021)	+	+	
Drehmann et al. (2012)	+	+	
Harun et al. (2014)		+	
Hiebert et al. (2014)		+	
Kupkovič and Šuster (2020)		+	
Lang et al. (2019)		+	+
Lee-Poy (2018)		+	
Mugrabi and Rünstler (2025)	+		+
Ng (2011)			+
Russo (2022)		+	
Schüler et al. (2015)	+	+	+
Stremmel (2015)		+	+
Vella (2024)		+	+

Source: authors' compilation

Several empirical regularities are commonly used to characterize and evaluate financial-cycle indicators. Based particularly on Drehmann et al. (2012) and Borio (2014), two features are especially relevant:

- the joint behavior of credit and residential property prices provides the most widely used empirical representation of the financial cycle;
- systemic banking crises frequently occur close to the peak or during the downturn phase of the financial cycle.

These regularities provide useful criteria for evaluating whether an estimated indicator offers an economically meaningful representation of cyclical systemic risk. Additional evidence suggests that the duration and amplitude of financial cycles depend on the prevailing monetary, financial, and institutional regimes. Drehmann et al. (2012) also document an increase in the length and amplitude of financial cycles since the mid-1980s.

The relative duration of financial and business cycles is, however, less clear-cut in the empirical literature. Earlier studies, including Drehmann et al. (2012), Borio (2014), and Schüler et al. (2015), generally characterised financial cycles as longer and lower-frequency than conventional business cycles. However, this empirical regularity is not uniform across indicators, countries, or methodologies. Even Drehmann et al. (2012), while concluding that the composite financial cycle is substantially longer than the business cycle, report turning-point estimates for individual series in which the credit cycle is markedly longer than the GDP cycle, whereas the credit-to-GDP and house-price cycles are of broadly similar duration to the GDP cycle, while equity and aggregate asset-price cycles are somewhat shorter than the GDP cycle. More recent evidence is more cautious. S. Chen and Svirydzenka (2021),

using a broad sample of advanced and emerging economies, find financial cycles to be on par with, or shorter than, business cycles, with substantial heterogeneity across indicators and country groups. Škare et al. (2025) also identify financial cycles that are generally longer than business cycles in their sample, but at the same time emphasise that estimates of financial-cycle dynamics and duration are sensitive to the choice of financial variables, their aggregation, and the method used to identify the cyclical component. Taking these differences into account, the relative duration of the financial and business cycles may be more appropriately viewed as an empirical characteristic that depends on measurement choices rather than as a universal defining property of the financial cycle.

Financial stress represents a related but conceptually distinct dimension of systemic risk. Financial stress indices measure contemporaneous tensions and disruptions in financial markets and institutions. For example, Holló et al. (2012) develop the Composite Indicator of Systemic Stress (CISS) for the euro area to quantify the intensity and systemic character of financial stress. By contrast, financial-cycle indicators focus primarily on the gradual accumulation and subsequent correction of cyclical vulnerabilities. Nevertheless, the boundary between financial-cycle indicators, early-warning indicators, and systemic-risk measures is not always clear in the empirical literature, as emphasized by Kupkovič and Šuster (2020).

The distinction is clearer from a methodological and policy perspective. Financial stress indices generally use monthly or higher-frequency data and combine variables that capture current market and institutional tensions. Financial-cycle indicators are usually based on quarterly data and seek to identify lower-frequency movements through filtering, decomposition, or turning-point methods. From a policy perspective, stress indicators are more closely associated with crisis management and short-term intervention, whereas financial-cycle indicators are primarily used for macroprudential surveillance and the calibration of countercyclical instruments. The two perspectives remain complementary, since the materialization of previously accumulated cyclical vulnerabilities is often accompanied by an increase in financial stress.

Variables used to capture the financial cycle

The construction of a financial-cycle indicator requires the selection of variables that capture the main channels through which cyclical vulnerabilities accumulate. The literature does not provide a universally accepted theoretical or empirical selection framework. Nevertheless, credit developments and residential property prices constitute the most widely used core variables. Existing specifications range from parsimonious two-variable indicators to broader measures that incorporate asset prices, interest rates, leverage, external imbalances, and banking-sector indicators (Table 2).

Table 2: Variables used in the construction of composite indicators

Authors	Name of composite indicator	Variables
Banka Slovenije (2026)	Composite indicator of cyclical systemic risk	Bank credit to the domestic private non-financial sector relative to GDP; overall real growth in credit; the ratio of residential real estate prices to income; the debt-service-to-income ratio; the ratio of the current account balance to GDP
S. Chen and Svirydzenka (2021)	Overheating index	Advanced economies: equity prices and output gap; emerging economies: equity prices, property prices, and credit
Czech National Bank (2023)	Financial cycle indicator	Evolution of new loans; changes in property price index; debt sustainability; lending conditions; PX stock index; adjusted current account-to-GDP ratio
Drehmann et al. (2012)	Financial cycle	Total credit; credit-to-GDP ratio; residential property prices; equity prices; aggregate asset price index
Kupkovič and Šuster (2020)	Financial cycle indicator	Change in the stock of loans to households; change in the stock of loans to non-financial corporations; house price inflation; household debt service ratio; non-financial corporations debt service ratio; credit standards on loans to households; credit standards on loans to non-financial corporations; spread between interest rate on loans to households and three-month EURIBOR inter-bank rate; spread between interest rate on loans to non-financial corporations and three-month EURIBOR inter-bank rate; current account deficit-to-GDP ratio
Harun et al. (2014)	Financial cycle	Broad credit; broad credit-to-GDP ratio; Jakarta Composite Index; Jakarta Property Index
Hiebert et al. (2014)	Financial cycle	Total credit; residential property prices; equity prices; interest rate
Lang et al. (2019)	Domestic cyclical systemic risk indicator	The two-year change in the bank credit-to-GDP ratio; the two-year growth rate of real total credit; the two-year change in the debt-service ratio; the three-year change in the residential real estate price-to-income ratio; the three-year growth rate of real equity prices; the current account-to-GDP ratio
Lee-Poy (2018)	Synthetic financial cycle measure (dynamically weighted composite)	Same set of variables as in Schüler et al. (2015)
Morana (2024)*	Financial cyclical condition	Private credit gap; house price gap; house price-to-income ratio; house price-to-rent ratio; real gold price; European Fama-French market factor return relative to three-month Treasury bill rate
Russo (2022)	Financial cycle indicator	Mortgages; lending to non-financial corporations; sovereign yields; consumer credit; stock market index in France
Schüler et al. (2015)	Composite index of the financial cycle	Total credit; residential property prices; equity prices; benchmark bond yields
Stremmel (2015)	Financial cycle measure	Total credit; credit-to-GDP ratio; residential property prices; short-term funding to total assets; net income to total assets; proportion of loans to total assets
Vella (2024)	Cyclical systemic risk indicator	Real bank credit, 2-year growth rate; debt service to income ratio, 1-year change; house price to income ratio; real total debt, 2-year growth rate

Source: authors' compilation.
*Note: * – Morana (2024) constructs a macro-financial conditions index for the euro area; however, the index includes a financial cycle conditions sub-index with its own set of variables related to the financial cycle.*

The reviewed studies confirm that credit developments and residential property prices constitute the most common empirical core of financial-cycle measurement. Equity prices, bond yields, and

interest-rate indicators are frequently added to capture broader financial-market conditions. Other variables, including debt-service burdens, external imbalances, funding structures, and banking-sector balance-sheet indicators, are used less systematically and generally reflect country-specific transmission channels. The appropriate variable set therefore depends on both the structure of the domestic financial system and the availability and quality of the underlying data.

Empirical approaches to the construction of composite financial cycle indicators

The construction of a composite financial-cycle indicator involves two main methodological stages: the extraction of cyclical components from the selected variables and their aggregation into a single measure. The literature proposes several approaches for both stages, with important implications for the estimated duration, amplitude, and timing of the resulting financial cycle.

A central contribution is provided by Drehmann et al. (2012), who use band-pass filters and turning-point analysis to characterize financial cycles. These approaches offer complementary perspectives: filtering isolates oscillations within predefined frequency ranges, while turning-point methods identify expansions and contractions directly from the observed series.

Using the band-pass filter developed by Christiano and Fitzgerald (2003), Drehmann et al. (2012) separate short-term movements of approximately 1–8 years from medium-term movements of approximately 8–30 years. Their results indicate that medium-term components account for a substantial share of the variation in credit and property prices and are more closely associated with systemic banking crises.

A common limitation of such filtering approaches is their reliance on parameters specified before the cyclical component is estimated. Band-pass filters require the specification of a frequency range. Although parameter values used in the literature provide useful benchmarks, their suitability may differ across countries and financial structures. S. Chen and Svirydzenka (2021) show that a common extraction rule does not necessarily generate comparable results across a heterogeneous country sample, particularly when financial systems differ in depth, maturity, and structural stability.

An alternative model-based approach extracts cyclical components within an unobserved-components or state-space framework. In such models, observed macro-financial variables are decomposed into stochastic trends and latent cyclical components, while parameters governing cycle persistence, frequency, and phase shifts can be estimated within a coherent econometric framework. Lee-Poy (2018), for example, applies a stochastic-cycle model to identify common business and financial cycles, while Mugrabi and Rünstler (2025) estimate credit, housing, and business cycles jointly within a multivariate unobserved-components model. These approaches have the important advantage of exploiting cross-variable information and estimating key characteristics of the cycle rather than relying exclusively on mechanically imposed univariate filters.

Nevertheless, their flexibility remains partly constrained by the functional form used to represent cyclical dynamics. The latent stochastic cycles are constructed through sine–cosine rotations, so that, even when their frequency, persistence, and phase shifts are estimated from the data, the extracted cycle is still encouraged to conform to a particular oscillatory structure. In some applications, additional restrictions are imposed on the admissible cycle duration. This structure is economically convenient and allows a clear interpretation of cycle length and phase relationships, but it may be restrictive if actual financial-cycle dynamics are asymmetric, irregular, or change their shape over time. In such cases, the underlying cycle may not necessarily follow a sufficiently regular harmonic pattern.

Spectral methods offer a data-informed approach to identifying relevant frequencies. Hiebert et al. (2014) use cohesion measures to identify frequency bands characterized by strong common comovement across financial variables, while Schüler et al. (2015) apply power cohesion to real and financial indicators. However, conventional spectral methods generally require stationary series. Transformations such as differencing or preliminary detrending may modify the low- and medium-frequency

properties of the original series and may complicate their economic interpretation, as discussed by S. Chen and Svirydzenka (2021).

An alternative is provided by Empirical Mode Decomposition (EMD), introduced by Huang et al. (1998). EMD is an adaptive, data-driven method that decomposes a signal into a finite set of Intrinsic Mode Functions (IMFs) and a residual component. Unlike conventional band-pass filters, it does not require a pre-specified frequency interval or a fixed functional form for the trend and cycle. Huang et al. (2003a) apply EMD to non-stationary financial time series and show that the method provides a localized representation of time-varying amplitude and frequency, thereby improving the identification of changes in cyclical dynamics over time. The extracted IMFs represent oscillatory modes identified from the local properties of the data, with the first IMF generally capturing the highest-frequency movements and subsequent IMFs representing progressively longer time scales. In this respect, EMD may be interpreted as an adaptive bank of nonlinear filters whose effective frequency bands are determined by the empirical structure of the signal rather than specified in advance. The residual captures the lowest-frequency movement and may be interpreted as the long-term trend. Consequently, the cyclical component can be reconstructed by combining the relevant IMFs while excluding the residual, without imposing an ex ante frequency cutoff. However, the economic interpretation and relevance of the individual modes must be assessed separately.

To address mode mixing and the sensitivity of EMD to intermittent extrema, ensemble-based extensions introduce noise-assisted decomposition. Ensemble Empirical Mode Decomposition (EEMD) reduces mode mixing by averaging the IMFs obtained from multiple noise-perturbed realizations, but the resulting finite-ensemble decomposition does not generally provide exact reconstruction of the original signal. Torres et al. (2011) proposed Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN), which introduces adaptive noise at successive decomposition stages while preserving complete signal reconstruction. CEEMDAN therefore retains the adaptive character of EMD while improving the stability and completeness of the decomposition. The present study applies CEEMDAN to extract the cyclical components of the selected macro-financial variables. A detailed description of EMD, EEMD, and CEEMDAN is provided in Appendix 1.

The final step in constructing the financial cycle indicator is to aggregate the cyclical components of the selected variables into a single measure. A simple and transparent aggregation procedure assigns equal weights to all normalized cyclical components. This approach is used, among others, by Drehmann et al. (2012), Stremmel (2015), and Kupkovič and Šuster (2020). When the underlying variables are expressed in different units, prior normalization is required to place their cyclical components on a comparable scale. Equal weighting avoids estimating additional parameters but implicitly assumes that all included variables are equally informative for measuring the financial cycle.

Time-varying weighting schemes allow the contribution of individual variables to change with their degree of co-movement. One approach is the exponentially weighted moving average (EWMA), which recursively updates covariance and correlation matrices by assigning greater importance to more recent observations. The resulting dependence structure can be used to derive weights from the strength of positive co-movement between each component and the remaining components. EWMA-based aggregation is widely used in financial stress indices, including Holló et al. (2012), Chadwick and Öztürk (2018), and Filatov (2021), and has also been applied to financial-cycle indicators by Schüler et al. (2015) and Lee-Poy (2018).

Principal Component Analysis (PCA) provides another data-driven aggregation method. Hiebert et al. (2014) rescale the selected variables to a common variance and represent the financial cycle using the first principal component. PCA is useful for identifying the linear combination that explains the largest share of common variation. However, maximizing explained variance does not necessarily identify the component that is most relevant for cyclical systemic risk. Moreover, the resulting loadings may be sensitive to the normalization and composition of the dataset and may be less straightforward to interpret from a policy perspective.

The evaluation of financial-cycle indicators remains challenging because the financial cycle is not

directly observable and sufficiently detailed crisis chronologies are unavailable for many countries, including the Republic of Moldova. Existing studies therefore use several complementary validation approaches. These include comparing the duration and persistence of the estimated financial cycle with those of the business cycle, examining the timing of financial-cycle peaks relative to historically relevant downturns or banking crises, assessing predictive performance when a reliable crisis classification is available, and evaluating real-time or quasi-real-time stability. Predictive assessments frequently rely on measures such as the area under the receiver operating characteristic (AUROC) curve, whereas real-time exercises examine the extent to which estimated signals are revised as new observations become available. In the absence of a formal crisis chronology, validation generally requires a combination of statistical evidence, historical interpretation, robustness analysis, and comparison with established macroprudential benchmarks such as the credit-to-GDP gap.

Despite the extensive literature, the measurement of financial cycles in small open economies characterized by short data samples, structural changes, shallow financial markets, and a dominant role of bank-based intermediation remains challenging. Moreover, the use of adaptive decomposition methods in the construction of composite financial-cycle indicators remains comparatively limited. This highlights the need for empirical frameworks that can identify country-specific cyclical dynamics without imposing a fixed frequency structure and that remain sufficiently transparent for macroprudential interpretation.

II. Methodology

General approach

The construction of an indicator intended to provide an analytical basis for assessing the accumulation of cyclical systemic risk involves several methodological challenges. These include the extraction of cyclical components, the adjustment of short-term high-frequency fluctuations, the selection of the most informative variables, and their aggregation into a composite indicator. These challenges are amplified in the case of the Republic of Moldova by the absence of an official chronology specifying the exact starting and ending dates of historical financial crises. Although periods of heightened financial instability can generally be identified by economists and policymakers, the lack of a formally dated crisis benchmark precludes the direct application of validation methods that require a predefined classification of crisis and non-crisis observations.

Several alternative approaches to cycle extraction were considered but were not adopted as the primary methodology. Band-pass filters require the ex ante specification of the frequency range associated with the financial cycle, which may be restrictive when the duration of domestic financial cycles is uncertain and potentially time-varying. Spectral methods provide a more data-informed way of identifying dominant frequencies and common cyclical movements, but their implementation generally requires stationary series or prior transformations to achieve stationarity; such transformations may alter the low- and medium-frequency properties of the original data and complicate their economic interpretation. State-space approaches based on stochastic-cycle specifications offer a more flexible model-based alternative, since parameters governing cycle frequency, persistence, and phase shifts can be estimated jointly within the model. However, the latent cyclical component is still represented through a stochastic sine–cosine structure, which constrains cyclical dynamics to a particular oscillatory functional form and may be restrictive if the underlying financial cycle is asymmetric, irregular, or changes shape over time.

These considerations motivate the use of an adaptive decomposition method that does not require the prior specification of a fixed frequency band or a deterministic trend function. The cyclical components are therefore extracted using Complete Ensemble Empirical Mode Decomposition with Adaptive Noise, an extension of the Empirical Mode Decomposition framework. CEEMDAN is suitable for the present application because it can accommodate nonlinear and non-stationary dynamics and allows oscillatory modes to be identified from the local properties of each series. Nevertheless, the resulting decomposition remains dependent on the selected sifting rule, noise intensity, ensemble size, and treatment of sample boundaries. These elements are therefore explicitly calibrated and evaluated in the subsequent stages of the methodology.

The stability of the CEEMDAN extraction is evaluated in a quasi-real-time exercise. This test is conducted at the level of the individual cyclical components and is intended to assess whether the cycles remain stable as the available sample expands. It is therefore conceptually distinct from the subsequent comparison of the composite indicators.

After the CEEMDAN decomposition, two alternative representations of the cyclical component are retained. The raw cycle is obtained by summing all extracted intrinsic mode functions while excluding the residual component. An adjusted cycle is subsequently constructed by removing IMF1, which generally captures the highest-frequency oscillations. This makes it possible to assess whether the information contained in IMF1 improves the identification of cyclical systemic risk or mainly reflects

short-term fluctuations with limited relevance for medium- and longer-term dynamics. Retaining both the raw and adjusted versions allows the informational contribution of these high-frequency oscillations to be evaluated empirically through the properties of the resulting indicators.

The individual cyclical components are subsequently normalized and aggregated within a comparative framework. Two aggregation schemes are considered: equal weighting and a time-varying weighting approach based on the Exponentially Weighted Moving Average method. Different variable combinations are used to construct candidate composite indicators from both the raw and adjusted cyclical components.

Given the absence of a sufficiently detailed and reliable chronology of financial crises in the Republic of Moldova, the preferred composite indicator is not selected through a formal crisis-classification exercise. Instead, the candidate composite indicators are evaluated along three complementary dimensions: their cyclical duration relative to the real economic cycle, the historical timing of positive financial-cycle peaks before selected macroeconomic downturns, and the economic plausibility of their component contributions. Cycle duration is treated as a descriptive and supporting characteristic rather than as a standalone selection rule. The credit-to-GDP gap is included in the same comparative framework as a conventional macroprudential reference indicator, but it is not treated as a candidate composite specification. The preferred composite indicator is selected on the basis of the overall consistency of the evidence.

Data and variable selection

The selection of variables was guided by two complementary considerations: the evidence provided by the financial-cycle literature and the structural characteristics of the Moldovan economy. The analysis therefore sought to identify indicators that are both theoretically relevant to the accumulation of cyclical systemic risk and empirically meaningful in the context of a small open economy with a relatively shallow financial system.

Credit and residential property prices were treated as the core components of the analysis. Their inclusion reflects the broad consensus in the literature that the interaction between credit expansion and property-price dynamics constitutes one of the central mechanisms underlying financial-cycle fluctuations.

An interest-rate indicator was also included among the candidate variables. This choice was motivated by the frequent use in the literature of lending rates, market yields, and interest-rate spreads as measures of financing conditions and risk-taking incentives. The interpretation of interest rates is nevertheless less straightforward than that of credit or property prices. Low interest rates may contribute to the build-up of vulnerabilities by reducing borrowing costs, stimulating credit demand, and encouraging greater risk-taking.

The current account position was included to capture the external dimension of financial-cycle dynamics. This consideration is particularly relevant for small open economies, in which domestic credit expansion and increases in asset prices may be accompanied by a deterioration in the current account balance. The inclusion of the current account is also consistent with empirical applications developed for small open economies. In particular, Kupkovič and Šuster (2020) include the current account deficit-to-GDP ratio in their financial-cycle indicator for Slovakia. Given Moldova's high degree of external openness, the current account balance was therefore regarded as a potentially important indicator of the accumulation of macro-financial imbalances.

Although the current account is frequently expressed relative to GDP in the financial-cycle literature, the present analysis uses the current account balance in real terms. This choice reflects the objective of extracting the cyclical dynamics of the external balance itself rather than constructing a relative measure of external imbalances. Scaling the current account by GDP would introduce movements originating from both the numerator and the denominator, so changes in the ratio could reflect developments in either the external balance or domestic output. In the present framework, the level-

based measure is preferred because it provides a more direct signal of the accumulation of external vulnerabilities. A widening current account deficit may itself indicate an increase in macro-financial imbalances, whereas the same deterioration may be partly obscured in a current-account-to-GDP ratio if it coincides with strong GDP growth. In this sense, the real current account balance is better suited to the objective of obtaining a cleaner measure of the build-up of cyclical systemic vulnerabilities, while avoiding additional information generated by fluctuations in the denominator.

The final set of candidate variables was also shaped by the intended interpretation of the composite index. The index is designed primarily to assess the build-up of cyclical systemic risk rather than the contemporaneous resilience of the banking sector or the ex post materialization of losses. From this perspective, rapid credit growth, sustained increases in residential property prices, and a widening current account deficit provide relatively direct signals of expanding financial imbalances. Persistently low interest rates may reinforce this process by lowering borrowing and debt-servicing costs and facilitating faster credit expansion.

By contrast, indicators related to the strength of bank balance sheets or the private-sector debt burden primarily describe the financial system's capacity to absorb shocks or the vulnerability of borrowers to adverse developments. Although such indicators are highly relevant for financial-stability surveillance, they do not necessarily provide independent information on the initial build-up of cyclical risk. Moreover, some of them may partly duplicate information already contained in credit dynamics. For instance, an increase in leverage or in the aggregate debt burden may itself result from rapid credit expansion. Similarly, non-performing loans are more naturally interpreted as a manifestation of previously accumulated vulnerabilities and adverse macro-financial developments than as an early indicator of the build-up of financial imbalances.

Their exclusion from the composite financial-cycle index should not, however, be interpreted as implying that these indicators are irrelevant for macroprudential policy. In its "Recommendation on guidance for setting countercyclical buffer rates", the ESRB (2014) identifies, among other variables, measures of bank balance-sheet strength and private-sector debt burden as complementary indicators that designated authorities should monitor when assessing cyclical systemic risk and determining the appropriate policy response. Accordingly, balance-sheet resilience and debt-burden indicators should be considered alongside the composite index when assessing the overall risk environment and, in particular, when calibrating macroprudential instruments. The present variable selection therefore reflects the narrower purpose of the index—capturing the build-up of financial imbalances—rather than the complete information set required for macroprudential decision-making.

Broader financial asset-price indicators could also provide relevant information on financial-cycle conditions. However, their inclusion was constrained by data availability. For Moldova, sufficiently long, continuous, and methodologically consistent time series for financial asset prices were not available over the full estimation period.

The analysis is based on quarterly data covering the period from 2007Q2 to 2025Q4. The candidate variables considered in the construction of the index, together with their definitions, units of measurement, and data sources, are summarized in Table 3. Credit is measured as the sum of claims denominated in domestic and foreign currency on privately owned non-financial corporations and other resident sectors, as reported in the monetary survey. Credit and the nominal RPPI are deflated using the consumer price index, with 2015 as the base year, and the resulting real RPPI is expressed as an index with 2015 equal to 100. The current account balance is first converted from US dollars into Moldovan lei using the MDL/USD exchange rate and is subsequently deflated using the GDP deflator, with 2015 as the base year. The ex-post real lending rate is calculated as the difference between the nominal average interest rate on newly granted bank loans and the year-on-year inflation rate. The quarterly series entering the financial-cycle analysis are not seasonally adjusted. No seasonal adjustment was applied because visual inspection of the series did not indicate a clear and persistent seasonal pattern over the sample period.

The credit series, nominal RPPI, nominal average lending rate, current account balance, and

MDL/USD exchange rate are obtained from the National Bank of Moldova, while the consumer price index and GDP deflator are obtained from the National Bureau of Statistics. To ensure a consistent interpretation across all components, credit and residential property prices retain their original signs, whereas the real interest rate and current account balance are multiplied by -1 during the initial data-transformation stage. Consequently, higher values of all transformed variables consistently indicate a stronger accumulation of cyclical systemic vulnerabilities.

Table 3: Candidate variables considered in the construction of the Financial Cycle Index

Variables	Description	Unit of measure	Source of data
CREDIT	The sum of claims denominated in MDL and foreign currency on privately owned non-financial corporations and other resident sectors, as reported in the monetary survey; expressed in real terms, with 2015 as the base year.	MDL	NBM, NBS
Residential Property Price Index (RPPI)	Residential Property Price Index expressed in real terms, with 2015 as the base year.	Index	NBM, NBS
Real interest rate* (RIR)	The difference between the average interest rate on newly granted bank loans and the year-on-year inflation rate.	%	NBM, NBS
Current account balance* (BCA)	Current account balance expressed in real terms, with 2015 as the base year.	MDL	NBM, NBS

Source: authors' compilation

*Note: * – RIR and BCA are multiplied by −1 during the initial data-transformation stage, so that higher values of the transformed series consistently correspond to a stronger accumulation of cyclical systemic vulnerabilities.*

Mirror-padding selection

Before applying CEEMDAN, symmetric mirror padding is introduced at both ends of the sample in order to mitigate boundary distortions. The selection of the CEEMDAN hyperparameters and the selection of the mirror-padding length are treated as two separate but interrelated exercises. In the implementation used in this work, the Index of Orthogonality is used to select the stopping criterion and the noise intensity, conditional on a given padding specification. It is not used to select the padding length itself because changing the padding modifies the number and structure of the artificially reflected observations entering the decomposition. Consequently, IO values obtained under different padding lengths are not directly comparable as measures of boundary performance and may mechanically favour a particular extension without indicating that the corresponding cycle is more stable over the observed sample. At the same time, an appropriate mirror-padding length is important for the reliability of the subsequent hyperparameter optimization, since inadequately controlled boundary effects may distort the extracted IMFs and lead the IO criterion to select parameters that fit artefacts generated near the sample endpoints. The padding length is therefore evaluated separately through the stability of the extracted cyclical component over the original, non-extended sample, while the IO-based hyperparameter selection is interpreted for each padding specification.

The padding length should be sufficiently large to reduce end effects, while remaining limited enough to avoid an excessive proportion of artificially generated observations and an unnecessary increase in computational burden. As discussed by Stallone et al. (2020), the choice of the signal-extension length involves a trade-off between mitigating boundary effects and limiting the length of the extended signal.

In the present implementation, the padding length refers to the number of reflected observations added separately at each endpoint of the original series. Three specifications are considered: $0.5T$, $0.75T$, and $1T$ at each endpoint, where T denotes the length of the original sample. Thus, under the baseline $0.5T$ specification, $0.5T$ observations are added before the beginning and after the end of the observed sample. The extension is constructed by simple symmetric reflection of the observed series, and no additional windowing or tapering procedure is applied. No additional custom mirroring of local extrema is imposed beyond this sample-level extension; the subsequent envelope construction is handled internally by the Rlibeemd implementation. Accordingly, Stallone et al. (2020) are used to motivate the general trade-off involved in selecting the extension length rather than to indicate that their complete signal-extension procedure is replicated.

Padding lengths exceeding $1T$ are not considered appropriate for the present application. Long signal extensions are more frequently employed in engineering settings in which the underlying oscillatory structure is comparatively stable and can reasonably be continued beyond the observed sample. By contrast, macro-financial series may be affected by structural changes, policy interventions, crises, and shifts in economic regimes and therefore do not generally exhibit stable repetitive patterns. Excessively long mirror extensions may consequently introduce artificial oscillatory structures that are not representative of the observed process and may influence the selection of the CEEMDAN hyperparameters. Limiting the maximum extension to $1T$ prevents an excessive increase in the share of artificial observations and in the computational burden associated with repeated hyperparameter optimization.

The adequacy of the alternative mirror-padding specifications is assessed through a visual comparison of the amplitude, persistence, and overall shape of the cycles over the original sample. The comparison is also supported by three complementary quantitative criteria: Pearson correlation, sign concordance, and turning-point concordance. This set of criteria is used because it captures the main dimensions through which the mirror-padding choice may affect the extracted cycle: the overall cyclical dynamics, the position of the cycle relative to zero, and the timing of the main peaks and troughs. The latter criterion is evaluated through turning-point analysis.

The turning-point analysis originates in the business-cycle dating literature of Burns and Mitchell (1946) and Bry and Boschan (1971) and was adapted to quarterly data by Harding and Pagan (2002). This study follows the dating approach proposed by Harding and Pagan (2002). The method identifies local peaks and troughs subject to window-size and censoring rules. A peak is defined as an observation that is higher than the K observations preceding and following it, while a trough is defined as an observation that is lower than the K observations on both sides. For quarterly data, $K = 2$ is used, implying that each candidate turning point must represent a local extremum relative to the two quarters on each side.

The candidate turning points are subsequently ordered chronologically. If two consecutive peaks are identified, only the higher peak is retained, whereas if two consecutive troughs are detected, only the lower trough is preserved. This procedure imposes strict alternation, so that the resulting sequence follows either a peak–trough–peak–trough pattern or a trough–peak–trough–peak pattern.

Censoring rules are applied by requiring a minimum phase duration of two quarters and a minimum complete-cycle duration of five quarters. The alternation rule is reapplied after censoring. The overall average cycle duration is computed by pooling the valid peak-to-peak and trough-to-trough durations. For the mirror-padding validation, no restriction is imposed on the sign of peaks or troughs, since the purpose of the exercise is to evaluate the sensitivity of the extracted cyclical structure to alternative mirror-padding lengths rather than to identify economically relevant financial cycles.

The $0.5T$ specification is used as the reference, while the turning points obtained under the $0.75T$ and $1T$ alternatives are compared with the reference dates. For each reference turning point, the alternative padding specification is checked for a turning point of the same type within a one-quarter window. Turning-point concordance is then calculated as the share of reference peaks and troughs that are matched by the alternative specification.

Sign concordance is calculated by comparing, quarter by quarter, the sign of the cycles obtained under two alternative padding specifications. A quarter is treated as concordant if both cycles are positive or both cycles are negative. The final measure is given by the share of concordant quarters in the total number of observations over the original, non-extended sample.

The preferred mirror-padding specification is selected on the basis of the overall consistency of Pearson correlation, sign concordance, turning-point concordance, and visual evidence, rather than through a single mechanical criterion.

Cycle extraction

The CEEMDAN decomposition is applied to a set of macro-financial variables selected on the basis of their relevance in the financial-cycle literature, the structure of the Moldovan financial system, and data availability. The variable set includes credit, residential property prices, the current account balance, and the real interest rate. Prior to decomposition, the variables are transformed and oriented so that higher values consistently correspond to a greater accumulation of cyclical systemic risk. The CEEMDAN decomposition is implemented in R version 4.4.3 using the ceemdan() function from the Rlibeemd package, version 1.4.3. Mirror padding, hyperparameter selection, stopping rules, and the subsequent reconstruction of the cyclical components are implemented through custom R routines.

The CEEMDAN decomposition is applied separately to each candidate variable. Each observed series is decomposed into a finite number of intrinsic mode functions and a residual component. The final representation of the series can be written as:

$$x(t) = \sum_{j=1}^{N_{\mathrm{IMF}}} \zeta_j(t) + r_{N_{\mathrm{IMF}}}(t) \tag{1}$$

where $x(t)$ denotes the original series, $\zeta_j(t)$ represents the j-th intrinsic mode function, and $r_{N_{\mathrm{IMF}}}(t)$ is the residual component. In this framework, the residual is treated as an estimate of the long-term trend, while the sum of all extracted IMFs is treated as the preliminary cyclical component. This preliminary component is subsequently adjusted for short-term high-frequency oscillations before normalization and aggregation.

The EMD algorithm starts with the extraction of each IMF through an iterative sifting procedure. For each mode, the local maxima and minima of the signal are first identified and used to construct the upper and lower envelopes. The mean of these two envelopes is then computed and subtracted from the current signal. This procedure is repeated until convergence is reached. In our case, we apply the S-Stoppage Criterion, under which convergence is achieved after S consecutive iterations in which the number of zero-crossings and extrema remains unchanged or differs by at most one.

After an IMF is extracted, it is removed from the current signal, and the same procedure is applied to the remaining residual. The decomposition continues until the residual becomes sufficiently smooth, for example monotonic, characterized by at most one extremum, or no longer containing a meaningful oscillatory structure. In CEEMDAN, the same general EMD logic is preserved, but the decomposition is performed within a noise-assisted framework in order to reduce mode mixing and improve the stability of the extracted components.

The CEEMDAN decomposition requires the specification of the ensemble size (N_{E}), the S-number governing the termination of the sifting process, and the noise intensity (ε). To select the S-number and the noise intensity, a grid-search procedure is applied. The grid ranges are defined ex ante on the basis of parameter values commonly used in the EMD and ensemble-EMD literature: $S \in \{4, 5, 6, 7, 8\}$ and $\varepsilon \in \{0.20, 0.25, 0.30, 0.35, 0.40\}$. The optimization therefore identifies the preferred parameter configuration within a literature-based parameter space rather than over an unrestricted set of values. Further details are provided in Appendix 1.

For each combination of S and ε, the CEEMDAN decomposition is performed using an ensemble size of $N_{\mathrm{E}} = 1{,}000$ realizations. This conservative setting is consistent with the convergence properties of ensemble averaging discussed by Wu and Huang (2009) and is supported by subsequent CEEMDAN and improved-CEEMDAN applications that either employ 1,000 realizations or report stable performance over ensemble sizes ranging from 500 to 2,000 (Yan et al., 2024; He & Tsang, 2021; Li et al., 2019).

The decomposition continues until the residual no longer contains a meaningful oscillatory structure. Operationally, the extraction of additional IMFs terminates when the residual contains fewer than two extrema, indicating that no further intrinsic mode function can be identified, or when an upper computational bound on the number of IMFs is reached. This bound is defined as round$[\log_2(T_e)]$, where T_e denotes the length of the extended series after mirror padding. For the full sample, this rule implies a maximum of seven IMFs.

To select the optimal hyperparameter combination, the quality of the decomposition is evaluated using the Index of Orthogonality. This metric measures the degree of leakage between the components of the decomposition, including both the extracted IMFs and the residual. The index is calculated as:

$$IO = \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^{N_{\mathrm{IMF}}+1} \frac{\left| \sum_t c_i(t) c_j(t) \right|}{\sum_t x^2(t)} \tag{2}$$

where $c_i(t)$ denotes the i-th component of the complete decomposition, including both the extracted IMFs and the residual component. The IO is used as a technical measure of separation between the components of a given decomposition. It is not interpreted as a standalone measure of the economic validity of the extracted cycle.

For each variable the combination of S, ε, and the endogenously determined number of IMFs that minimizes the Index of Orthogonality is selected. The decomposition is subsequently re-estimated using the selected parameters and an ensemble size of 1,000. The residual component is retained as the trend, while the raw cyclical component is obtained as the sum of all extracted IMFs, excluding the residual, as follows:

$$\mathrm{Cycle}_{i,t}^{\mathrm{raw}} = \sum_{j=1}^{N_{\mathrm{IMF},i}} \zeta_{i,j}(t). \tag{3}$$

Cycle adjustment

The extracted cycle may still contain high-frequency movements that are not necessarily relevant for the analysis of medium-term financial dynamics. The adjustment is based on the structure of the CEEMDAN decomposition itself. Under this approach, the adjusted cycle is obtained by excluding the first IMF, since the first IMF usually captures the highest-frequency oscillations, and summing the remaining IMFs:

$$\mathrm{Cycle}_{i,t}^{\mathrm{adjusted}} = \sum_{j=2}^{N_{\mathrm{IMF},i}} \zeta_{i,j}(t) \tag{4}$$

The adjustment is applied separately to each variable. The adjusted cycle is obtained by excluding IMF1, since this component generally contains the highest-frequency movements. Both the raw and adjusted cycles are therefore retained for the subsequent normalization step. Their comparative assessment makes it possible to determine whether IMF1 contributes useful cyclical information or mainly introduces short-term volatility.

Quasi-real time validation

The stability of the CEEMDAN-based cycle extraction is evaluated through a quasi-real-time recursive exercise conducted before normalization and aggregation. The purpose of this exercise is to assess whether the extraction procedure remains stable when applied recursively as new observations become available, which is particularly relevant for an indicator intended to support macroprudential monitoring and policy calibration.

The initial estimation window covers 2007Q2–2019Q4, while the recursive evaluation period extends from 2020Q1 to 2025Q4. This starting point provides a sufficiently long initial sample and allows the performance of the decomposition to be assessed during a period that includes several major macro-financial shocks.

Beginning with 2020Q1, the available sample is expanded by one quarter at each iteration. For every variable and vintage, the complete extraction procedure is repeated using only the observations available up to that quarter. This includes the re-selection of the CEEMDAN parameters and the re-estimation of the decomposition. Only the endpoint estimate is retained. The recursive endpoint estimates for 2020Q1–2025Q4 are then compared with the corresponding full-sample estimates obtained from the decomposition of the full sample.

This exercise follows the logic of real-time reliability assessments used in the financial-cycle literature, where indicators estimated using the full sample are compared with one-sided or recursively estimated counterparts (Lee-Poy, 2018; Kupkovič & Šuster, 2020). In particular, the validation metrics are taken from the approach used by Lee-Poy (2018), who evaluates real-time reliability using Pearson correlation, sign concordance, and a signal-to-noise measure. This set of indicators is used because it captures the main dimensions through which recursive endpoint estimates may differ from the corresponding full-sample estimates: the overall co-movement of the cyclical signal, the classification of observations into positive or negative cyclical phases, and the size of revisions relative to the benchmark cyclical signal.

Unlike the mirror-padding exercise, the quasi-real-time validation focuses on revisions between recursive endpoint estimates and full-sample estimates, rather than on the sensitivity of turning-point dates to alternative technical specifications. Following Lee-Poy (2018), the signal-to-noise SNR measure is applied to assess the magnitude of revisions in the recursive CEEMDAN estimates. It is computed as

$$\mathrm{SNR} = 1 - \frac{\mathrm{sd}(r_t)}{\mathrm{sd}(\hat{c}_t^{\mathrm{full}})} \tag{5}$$

Here, r_t denotes the revision in quarter t, calculated as the difference between the recursive estimate and the corresponding full-sample estimate. The term $\hat{c}_t^{\mathrm{full}}$ denotes the cyclical estimate obtained using the full sample covering 2007Q2–2025Q4. For the calculation of the SNR, however, both r_t and $\hat{c}_t^{\mathrm{full}}$ are restricted to the evaluation period 2020Q1–2025Q4. Consequently, both standard deviations are calculated using the same 24 quarterly observations.

The exercise is performed separately for the raw and adjusted cycles of each variable and is interpreted as a quasi-real-time robustness test of the CEEMDAN extraction procedure.

Normalization

After the cycle adjustment step, the extracted cyclical components (raw cycle and adjusted cycle) are normalized using z-score standardization in order to make them comparable across variables. This step is necessary because cycles may be expressed in different units of measurement. Therefore, before aggregation, they need to be transformed into a common scale. The normalization procedure is applied separately to each variable included in the candidate set.

$$z_{i,t}^{(m)} = \frac{\text{Cycle}_{i,t}^{(m)} - \mu_i^{(m)}}{\sigma_i^{(m)}}, \qquad m \in \{\text{raw}, \text{adjusted}\} \tag{6}$$

where i denotes the variable, t denotes time, and $m \in \{\text{raw}, \text{adjusted}\}$ identifies the cycle specification. The terms $\mu_i^{(m)}$ and $\sigma_i^{(m)}$ denote, respectively, the sample mean and standard deviation of the corresponding cyclical component.

Aggregation

After normalization, the individual cyclical components are aggregated into alternative composite indicators of the financial cycle. The aggregation is performed using several candidate combinations of variables in order to compare alternative representations of cyclical systemic risk. The four-variable specification includes credit, residential property prices, the current account balance, and the real interest rate. The three-variable specification excludes the real interest rate, while the two-variable specification includes only credit and residential property prices.

Two aggregation schemes are considered. The first is a simple equal-weighting approach, where each variable included in a given specification receives the same weight. For a composite financial cycle index (FCI) constructed from N normalized cyclical components, the equal-weighted index is computed as:

$$FCI_t^{(m)} = \frac{1}{N} \sum_{i=1}^{N} z_{i,t}^{(m)}, \qquad m \in \{\text{raw}, \text{adjusted}\} \tag{7}$$

where $z_{i,t}^{(m)}$ denotes the normalized cyclical component of variable i under cycle specification m, and N is the number of variables included in the corresponding composite indicator. This approach provides a transparent benchmark and avoids imposing additional assumptions regarding the relative importance of individual variables.

The second aggregation scheme is based on time-varying weights derived from an EWMA framework. The EWMA procedure is used to estimate time-varying covariance and correlation matrices between the normalized cyclical components. For notational simplicity, $x_{i,t}$ denotes the normalized cyclical component used in the corresponding candidate specification. For each pair of variables i and j, the covariance is updated as follows:

$$\sigma_{ij,t} = \lambda \sigma_{ij,t-1} + (1 - \lambda) \left(x_{i,t} - \bar{x}_{i,t}^{(8)} \right) \left(x_{j,t} - \bar{x}_{j,t}^{(8)} \right). \tag{8}$$

where $\sigma_{ij,t}$ denotes the time-varying covariance between the normalized cyclical components of variables i and j, λ is the decay factor, and $\bar{x}_{i,t}^{(8)}$ and $\bar{x}_{j,t}^{(8)}$ denote the corresponding means calculated over the most recent eight-quarter window, including period t. The decay factor, λ, is set to 0.89, following Schüler et al. (2015). The authors note that this value is slightly lower than the decay factor of 0.93 adopted by Holló et al. (2012). They justify the lower value by the quarterly frequency of their data, in contrast to the daily observations used by Holló et al. (2012), as the lower sampling frequency implies a smaller number of observations. A decay factor of 0.89 is also employed by Lee-Poy (2018), based on the approach proposed by Schüler et al. (2015). The initial covariance matrix is estimated using the first eight observations after centering each normalized cyclical component by its mean over the initialization window. The covariance and correlation matrices obtained from the initialization window are kept constant for the first eight quarters, after which the recursive updating procedure begins. This choice also follows the procedure adopted by Schüler et al. (2015).

The time-varying correlation coefficient is then obtained as:

$$\rho_{ij,t} = \frac{\sigma_{ij,t}}{\sqrt{\sigma_{ii,t}\sigma_{jj,t}}}. \tag{9}$$

The resulting correlations are collected in a time-varying correlation matrix. Since the variables are transformed and oriented before the decomposition stage so that higher values consistently indicate higher cyclical systemic risk, only positive co-movements are retained for the purpose of determining the aggregation weights. Accordingly, and in line with Schüler et al. (2015), non-positive correlations are set to zero. This specification assigns greater weight to components that move jointly in the direction of increasing systemic risk and prevents positive and negative correlations from offsetting each other, which could weaken the ability of the composite indicator to capture the build-up and correction of cyclical financial imbalances. The restricted correlation matrix is defined as follows:

$$\rho^{+}_{ij,t} = \begin{cases} \rho_{ij,t}, & \text{if } \rho_{ij,t} > 0 \text{ and } i \neq j, \\ 0, & \text{if } \rho_{ij,t} \leq 0 \text{ and } i \neq j, \end{cases}$$

$$C_t = \begin{pmatrix} 1 & \rho^{+}_{12,t} & \rho^{+}_{13,t} & \cdots & \rho^{+}_{1N,t} \\ \rho^{+}_{21,t} & 1 & \rho^{+}_{23,t} & \cdots & \rho^{+}_{2N,t} \\ \rho^{+}_{31,t} & \rho^{+}_{32,t} & 1 & \cdots & \rho^{+}_{3N,t} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{+}_{N1,t} & \rho^{+}_{N2,t} & \rho^{+}_{N3,t} & \cdots & 1 \end{pmatrix}. \tag{10}$$

The vector of time-varying weights is then calculated as:

$$w_t = \frac{C_t \mathbf{1}}{\mathbf{1}^T C_t \mathbf{1}}. \tag{11}$$

where w_t is the column vector of time-varying weights, C_t is the restricted time-varying correlation matrix, and $\mathbf{1}$ is a column vector of ones. By construction, the resulting weights are normalized such that $\sum_{i=1}^{N} w_{i,t} = 1$. The EWMA-based composite index is obtained as:

$$FCI_t = w_t^T X_t. \tag{12}$$

where X_t is the column vector of normalized cyclical variables. This aggregation method allows the contribution of each variable to vary over time depending on its co-movement with the other components of the index. For the two-variable specification, the correlation-based weighting formula yields equal weights because the two rows of the symmetric correlation matrix have identical sums. Consequently, the EWMA-weighted and equal-weighted two-variable indicators coincide, and only one version is reported.

The selection of the final composite index

After constructing the alternative composite indicators, a set of assessment dimensions is used to evaluate the candidate specifications and identify the preferred composite indicator. The selection procedure does not rely on a single mechanical rule. Instead, it combines insights from the financial-cycle literature with country-specific economic evidence related to periods of financial and macroeconomic turbulence in the Republic of Moldova.

Since an official and sufficiently detailed chronology of financial crises is not available for Moldova, the validation of the candidate indicators cannot rely on standard crisis-classification methods. The selection procedure therefore considers three complementary dimensions of assessment. First, the average duration of the candidate financial-cycle indicators is compared with the duration of the real

economic cycle, proxied by the GDP gap. This comparison is used to assess whether financial cycles in Moldova tend to be longer than the real economic cycle and whether differences in cyclical duration provide useful information for distinguishing among the candidate indices. Second, the historical timing of financial-cycle peaks is assessed relative to selected episodes of economic downturn. Third, the economic relevance of each candidate is evaluated by comparing the contributions of its underlying components and assessing whether the aggregate signal reflects economically plausible sources of cyclical systemic risk.

The historical timing dimension focuses on the downturns of 2009 and 2015. Both episodes were preceded by the accumulation of identifiable macro-financial imbalances, although the nature of these imbalances differed. The period preceding the 2009 downturn was characterised by sizeable current account deficits and rapidly rising residential property prices, whereas the period preceding the 2015 downturn featured strong credit expansion under comparatively weak prudential conditions. In both cases, the accumulated domestic vulnerabilities increased the sensitivity of the economy to adverse external developments, including the contraction of external demand and the decline in exports and remittance inflows. These episodes are therefore suitable for examining whether the candidate indicators capture the build-up of cyclical systemic risk before a subsequent deterioration in real economic activity.

The downturn associated with the outbreak of the war in Ukraine in 2022 is not included as a target episode in the historical peak-timing dimension. Although credit activity was recovering after the COVID-19 pandemic, the deterioration in macroeconomic conditions in 2022 was primarily triggered by an abrupt and largely unanticipated geopolitical shock. Its transmission to the Moldovan economy occurred particularly through the sharp increase in imported energy and food prices, heightened uncertainty and disruptions to regional trade. Consequently, this episode was not preceded by the same type of sustained and broad-based accumulation of domestic macro-financial imbalances observed before the 2009 and 2015 downturns. Requiring a financial-cycle indicator to display a peak before 2022 would therefore conflate its ability to identify the gradual build-up of cyclical vulnerabilities with its ability to anticipate the timing of an exogenous geopolitical shock. The exclusion of 2022 from the peak-timing dimension does not imply that the episode is economically irrelevant. It remains informative for assessing whether the contributions of individual components preserve a meaningful economic interpretation under exceptionally large external and inflationary shocks.

The third dimension examines the internal composition of the candidate indicators. The contribution of each component is compared across specifications and over time to determine whether changes in the aggregate index are supported by economically plausible developments in the underlying variables. This assessment considers whether the signal is broadly consistent with the expected channels of cyclical risk accumulation, whether it is excessively dominated by a single component and whether alternative weighting schemes provide additional economically meaningful information. The contribution analysis thus complements the duration and historical timing assessments by evaluating the interpretability of the signals generated by each candidate specification.

The three dimensions of assessment are discussed separately and are not combined into a numerical score. The preferred specification is selected on the basis of the overall consistency of the evidence. This approach avoids imposing an arbitrary compensatory rule across dimensions that capture distinct properties of the candidate indicators.

To operationalize these dimensions of assessment, the GDP gap is first calculated and the turning points of both the GDP gap and the candidate financial-cycle indicators are identified. As a macroeconomic benchmark, the output gap is computed from seasonally adjusted real GDP using the Hodrick–Prescott (HP) filter with the conventional smoothing parameter for quarterly data, $\lambda = 1600$. The GDP gap is defined as the deviation of seasonally adjusted real GDP from its HP trend. The HP filter is applied to the GDP series starting in 2006Q1, while the resulting gap is retained from 2007Q2 onward to match the sample available for the financial-cycle indicators. Because the HP-based GDP gap is estimated using the full sample, it is employed only for retrospective historical comparison and

not as a formal crisis chronology.

The turning points used in the selection exercise are identified according to the Harding and Pagan (2002) dating procedure described in the mirror-padding assessment. For this application, a minimum phase length of two quarters is imposed to avoid treating very short-lived fluctuations as separate phases, while a valid cycle is required to have a minimum peak-to-peak length of five quarters. The standard turning-point procedure identifies local maxima and minima subject to these duration and alternation constraints, without imposing restrictions on the sign of the identified extrema.

Given that the composite indicators are standardised around zero, an additional sign restriction is also considered for economic interpretation. Under this specification, only positive peaks and negative troughs are retained, and a valid financial cycle is defined as a positive peak–negative trough–positive peak sequence. Positive values of the index are interpreted as periods in which cyclical systemic risk is above its sample-average level, while negative values indicate below-average cyclical conditions. The sign restriction is therefore intended to focus the dating exercise on broader movements between these two cyclical positions rather than on local reversals occurring on the same side of the zero reference level.

Since the sign restriction is not part of the original Harding–Pagan procedure and may affect the number and timing of identified turning points, cycle duration is assessed under both specifications. Average peak-to-peak durations are therefore reported using, first, the unrestricted local turning points identified by the standard procedure and, second, the additional positive-peak/negative-trough restriction. Comparing the two specifications allows the sensitivity of the estimated cycle duration to the sign convention to be evaluated and helps determine whether duration provides useful information for distinguishing among the candidate indices.

The same turning-point and cycle-duration statistics are also computed for the credit-to-GDP gap under both the unrestricted and sign-restricted specifications. The credit-to-GDP ratio is calculated as the stock of bank credit at the end of each quarter relative to nominal GDP accumulated over the preceding four quarters. The definition of credit is kept consistent with the credit aggregate used in the composite indicators. The credit-to-GDP gap is calculated as the deviation of this ratio from its one-sided HP-filtered trend, using the Basel smoothing parameter of $\lambda = 400{,}000$ for quarterly data. This makes it possible to compare the candidate composite indicators both with the GDP gap, which is used as a proxy for the real economic cycle, and with the conventional credit-gap benchmark for cyclical systemic risk. For each candidate indicator, the average cycle length is computed as the mean duration of all identified peak-to-peak cycles under the corresponding turning-point specification.

The second dimension evaluates the historical timing of positive financial-cycle peaks relative to the subsequent deterioration in real economic activity. For each selected episode, the first positive-to-negative transition of the GDP gap is used as a macroeconomic reference point, indicating the quarter in which real GDP falls below its estimated trend. For the 2009 and 2015 episodes, a candidate indicator is expected to display a peak before this GDP-gap crossing. The corresponding lead is therefore measured as the number of quarters between the financial-cycle peak and the positive-to-negative transition of the GDP gap.

The comparison is completed by examining the internal structure of each candidate indicator. For every specification, the signed contributions of the underlying components are analysed over time and compared across the historically relevant episodes. The assessment examines whether increases and decreases in the composite indicator can be attributed to economically plausible developments in the constituent variables or are instead generated by counterintuitive or disproportionately dominant movements in a particular component. It also evaluates whether the use of time-varying weights adds economically relevant information relative to equal weighting or merely changes the relative magnitude of component contributions without materially improving the aggregate signal.

Particular attention is paid to periods dominated by large exogenous shocks, including 2022. Although such episodes are not used to assess the ability of the indicators to anticipate the gradual accumulation

of vulnerabilities, they provide an informative test of whether individual components can generate signals that are disconnected from prevailing financial conditions. The economic relevance dimension therefore helps distinguish between specifications that produce statistically similar cycle-duration and peak-timing results but differ in the interpretability and robustness of their internal composition. The final specification is selected through a transparent qualitative synthesis of the available evidence, with cycle duration treated as a descriptive and supporting characteristic, alongside the historical timing relative to the GDP gap and the economic plausibility of component contributions.

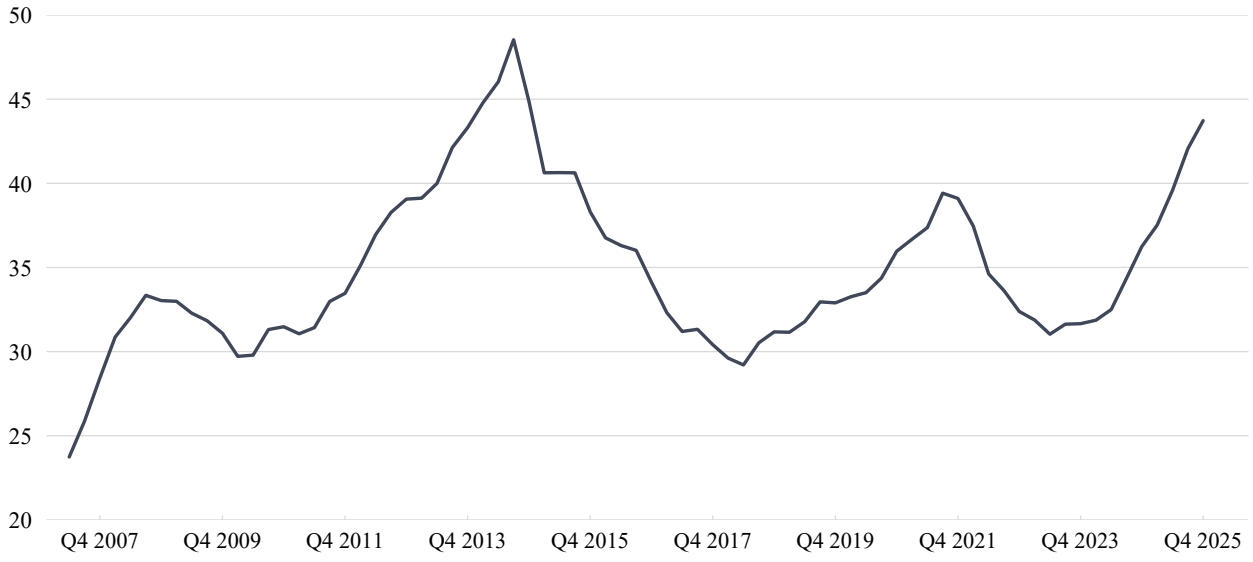
Following its selection, the composite index is subjected to a quasi-real-time validation using the same recursive approach applied to the extracted component cycles. The recursive estimates are compared with the corresponding full-sample results through visual inspection, the Pearson correlation coefficient, sign concordance and a signal-to-noise measure. This exercise assesses whether the subsequent normalization and aggregation of the extracted cycles materially affect the stability of the final index over time.

III. Results

Descriptive analysis and preliminary diagnostics

Before extracting the cyclical components, we first examine the main empirical characteristics of the candidate variables. As shown in Figure 1, following rapid growth in 2007–2008, real credit declined during 2009–2010 amid the repercussions of the global financial crisis. A new expansion subsequently emerged and culminated in 2014, amid non-compliance with prudential requirements in parts of the banking sector (National Bank of Moldova, 2015). Real credit then entered a prolonged contraction that continued until 2018. A subsequent recovery was temporarily interrupted during 2021–2023, against the backdrop of worsening regional economic conditions following Russia's invasion of Ukraine. Toward the end of the sample, credit growth accelerated markedly.

Figure 1: Credit, billion MDL at 2015 prices



Source: NBM, NBS, authors' calculations.

As shown in Figure 2, the real residential property price index declined in 2009 as a result of the economic downturn caused by the global financial crisis, and the decrease continued until 2012. The index subsequently fluctuated through alternating periods of decline and recovery. In 2024–2025, however, it increased sharply.

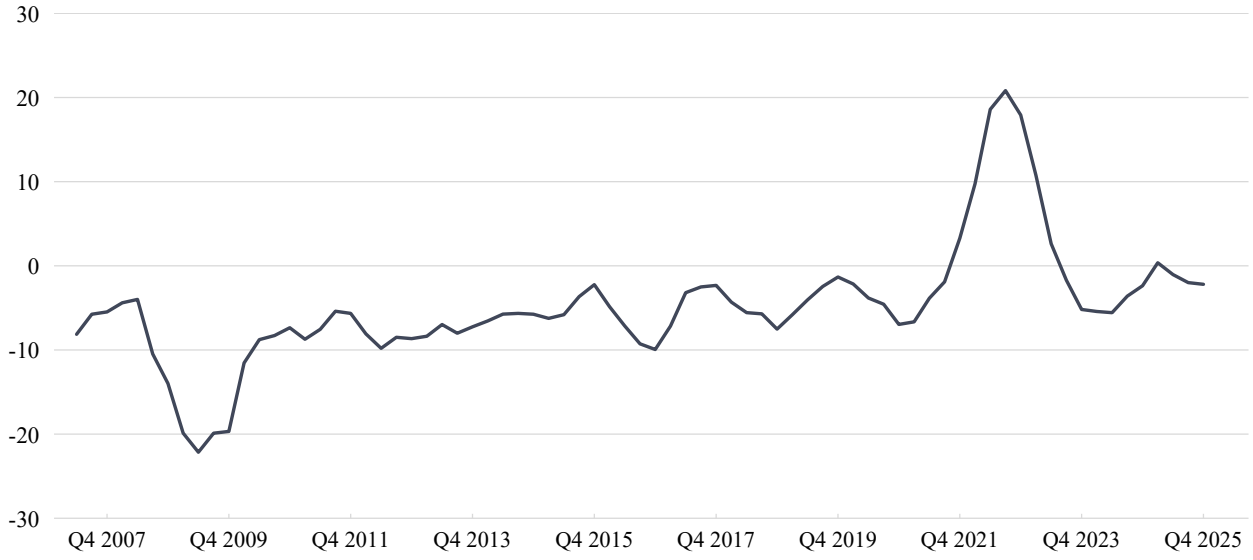
Figure 2: Residential Property Price Index, 2015 = 100



Source: NBM, NBS, authors' calculations.

As shown in Figure 3, the original real interest rate remained above zero during most of the sample period, corresponding to predominantly negative values of the sign-adjusted series, although its magnitude varied considerably. The most pronounced decline in the original real interest rate occurred in 2021–2022, corresponding to a sharp increase in the sign-adjusted series. In 2022, the sharp increase in international energy prices contributed to a substantial acceleration of inflation in Moldova. According to the IMF Primary Commodity Prices database, the energy price index increased by around 66% in 2022; this index covers crude oil, natural gas, coal, and propane prices (International Monetary Fund, 2026). Annual inflation in Moldova exceeded 30% during the second half of the year and was more than twice as high as the nominal interest rate on newly granted loans. Consequently, the original real interest rate became strongly negative, while the sign-adjusted rate moved sharply into positive territory.

Figure 3: Real interest rate, %

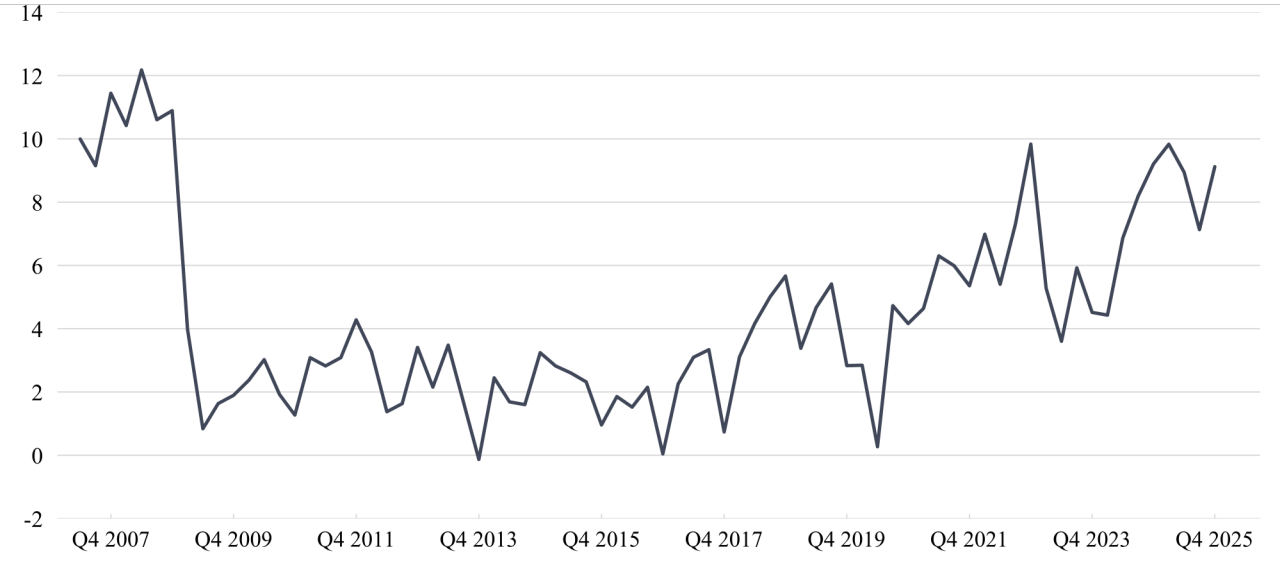


Source: NBM, NBS, authors' calculations.

Note: The series is multiplied by -1 *during the initial data-transformation stage.*

As shown in Figure 4, the current account deficit deteriorated rapidly before the 2009 crisis. In 2009, however, the sharp contraction in imports during the economic downturn led to a substantial narrowing of the deficit, which corresponded to a pronounced decline in the sign-adjusted current account balance. After 2009, the series remained volatile, but an upward tendency gradually became distinguishable, consistent with a widening of the underlying current account deficit and an accumulation of external imbalances. This tendency strengthened toward the end of the sample, particularly in 2024–2025, when the current account deficit increased sharply and the sign-adjusted series recorded a corresponding rapid rise.

Figure 4: Current account balance, billion MDL at 2015 prices



Source: NBM, NBS, authors' calculations.

Note: The series is multiplied by −1 during the initial data-transformation stage.

A visual inspection of the series suggests that they are more likely to exhibit non-stationary properties. Credit and the real interest rate display fluctuations with changing amplitude and persistence, while the residential property price index follows a clearly distinguishable trend. A similar tendency can also be identified in the current account balance. To formally assess the stationarity properties of the series, two complementary tests were applied: the Elliott–Rothenberg–Stock Dickey–Fuller generalized least squares unit-root test (ERS DF–GLS), as described by Elliott et al. (1996), and the Kwiatkowski–Phillips–Schmidt–Shin stationarity test (KPSS), as described by Kwiatkowski et al. (1992). Given the quarterly frequency of the data, the ERS DF–GLS test was implemented with four lagged differences to account for short-run serial dependence over a one-year horizon while preserving a parsimonious specification for the relatively short sample. For the KPSS tests, the long bandwidth option was used. Given $T = 75$, the corresponding rule, $L = \lfloor 12(T/100)^{1/4} \rfloor$, yields 11 truncation lags. The results of the stationarity tests are reported in Table 4.

Table 4: Stationarity tests

Variable	ERS DF–GLS		KPSS	
	Constant	Trend	Level	Trend
RPPI	-0.818	-1.196	0.574**	0.186**
CREDIT	-1.122	-2.721	0.115	0.097
BCA	-1.316	-1.702	0.215	0.186**
RIR	-1.587	-2.681	0.515**	0.065

Source: authors' calculations.
*Note: The table reports test statistics. The ERS DF–GLS test has the null hypothesis of a unit root, while the KPSS test has the null hypothesis of stationarity. ** indicates rejection of the corresponding null hypothesis at the 5% significance level.*

The results reported in Table 4 point to considerable persistence in the variables. The ERS DF–GLS test does not reject the null hypothesis of a unit root for any of the four series, either under the constant specification or under the trend specification. This result suggests that, from the unit-root testing perspective, the variables cannot be confidently treated as stationary processes.

The KPSS results provide additional evidence in the same direction, although the conclusions are not uniform across all variables. For the residential property price index, both level stationarity and trend stationarity are rejected, indicating a clearly non-stationary behaviour. For the real interest rate, level stationarity is rejected, although trend stationarity is not rejected, suggesting that the series is highly persistent and may be better characterized as stationary only around a deterministic trend. For the current account balance, trend stationarity is rejected, while level stationarity is not rejected. Credit is the only variable for which the KPSS test does not reject either level or trend stationarity; however, this result contrasts with the ERS DF–GLS evidence, which does not reject the unit-root null.

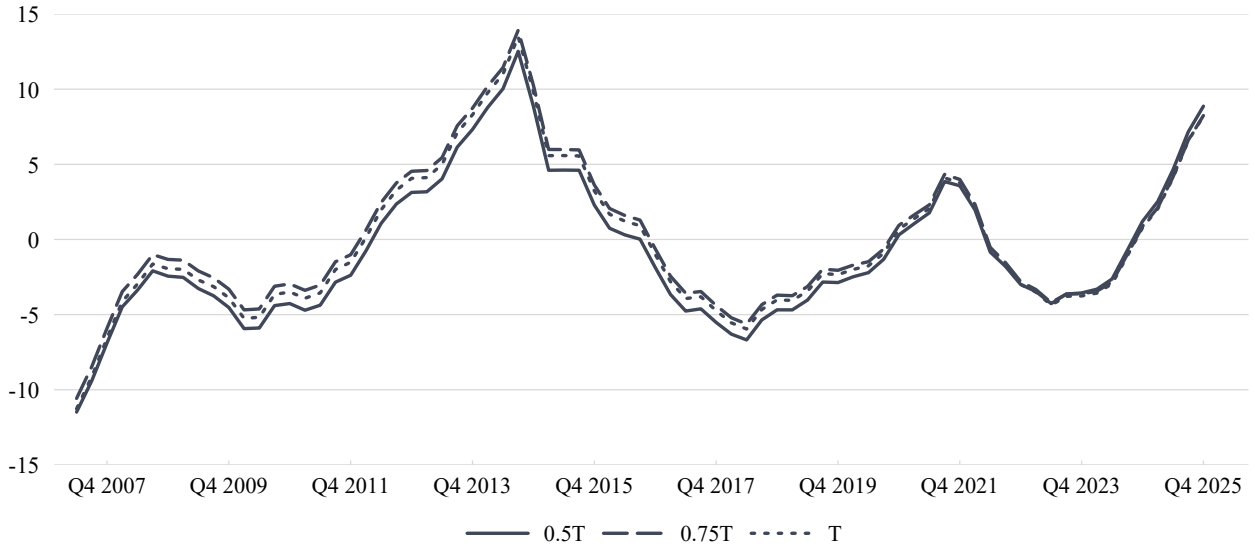
Overall, the testing evidence does not support a strong assumption of stationarity for the variables included in the index. Although some KPSS results are mixed, the combination of visual inspection, the ERS DF–GLS results, and the rejection of stationarity for several specifications indicates that the series are better treated as highly persistent and potentially non-stationary. This supports the use of a decomposition approach that is less dependent on stationarity assumptions. In this respect, CEEMDAN is suitable for the present analysis, as it allows cyclical components to be extracted from non-linear and non-stationary time series without requiring prior detrending or transformation to stationarity.

Mirror-padding selection

The CEEMDAN decomposition, like other data-driven filtering methods, may be sensitive to endpoint conditions, since the behaviour of the signal near the beginning and the end of the sample is less precisely determined than in the interior of the series. To assess the sensitivity of the extracted cycles to endpoint treatment, the decomposition was repeated using three alternative mirror-padding lengths: $0.5T$, $0.75T$, and T, where T denotes the length of the original sample. The purpose of this exercise is to verify whether the selected padding length materially affects the phase, sign, or turning-point structure of the cyclical components.

For credit, the three mirror-padding specifications generate almost identical cyclical dynamics. The main expansionary and contractionary phases are preserved across all padding lengths, while the differences between the extracted cycles are limited mainly to very small amplitude deviations.

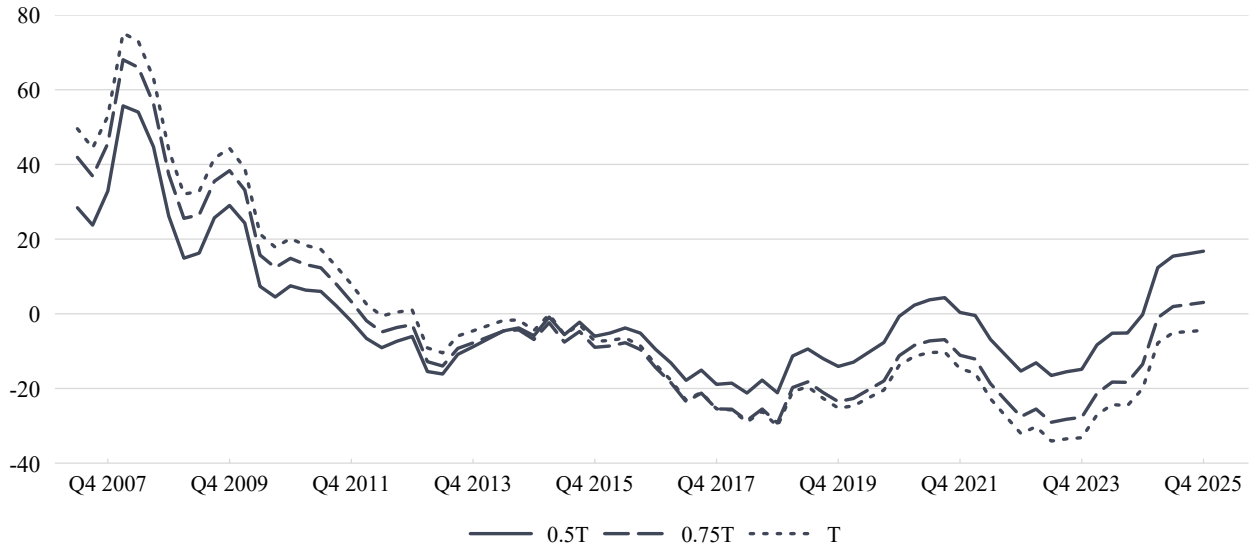
Figure 5: Credit cycle under alternative mirror-padding rules



Source: authors' calculations.

For the residential property price index, the alternative padding rules lead to more visible amplitude differences than in the case of credit. Nevertheless, the broad cyclical pattern remains comparable across specifications, and the main turning points are largely preserved, suggesting that the interpretation of the cyclical phases is not materially altered.

Figure 6: Residential property price cycle under alternative mirror-padding rules



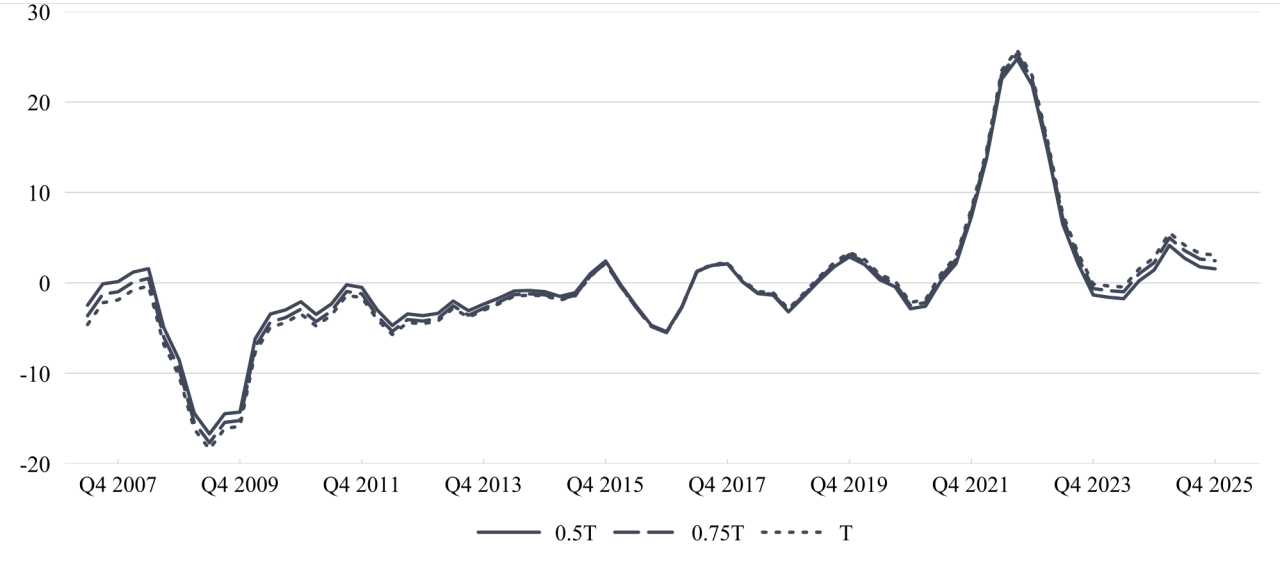
Source: authors' calculations.

The prolonged movement observed in the RPPI cycle across all mirror-padding specifications is largely associated with the lowest-frequency IMF. As shown in Appendix 2, Figure 2.2, IMF6 captures a slow decline followed by a gradual recovery and does not complete a full oscillation within the available sample. Its inclusion in the cyclical component follows from the sequential decomposition procedure adopted in this study. After each additional IMF is extracted, the remaining component is evaluated to determine whether further oscillatory structure can still be identified. The decomposition terminates when the residual contains fewer than two extrema. For RPPI, after IMF6 is extracted, the remaining

component satisfies this stopping condition and is therefore retained as the long-term residual, while IMF6 remains the lowest-frequency oscillatory component identified by the decomposition. Given the relatively short sample, only part of this low-frequency oscillation may be observed; therefore, the absence of a complete oscillation within the sample does not by itself imply that IMF6 should be classified as part of the long-term trend.

For the real interest rate, the extracted cycles are highly stable across the three padding specifications. The cyclical movements are almost fully aligned, including during the large fluctuation observed in 2021–2022, indicating that this episode is not an artefact of the selected endpoint-treatment rule.

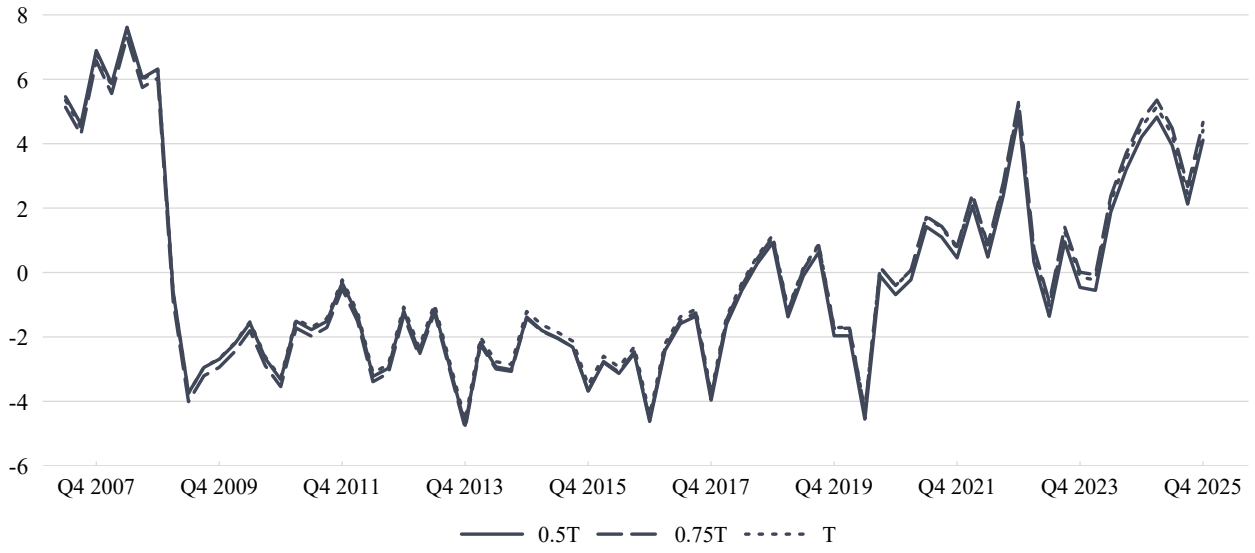
Figure 7: Real interest rate cycle under alternative mirror-padding rules



Source: authors' calculations.

For the current account balance, the cycles obtained under the three padding rules are also very similar. The differences are limited to small amplitude variations, while the phase and turning-point structure remain practically unchanged.

Figure 8: Current account balance cycle under alternative mirror-padding rules



Source: authors' calculations.

The visual comparison is complemented by three quantitative validation measures. Pearson correlation captures the overall co-movement between the cycles obtained under alternative padding rules. Sign concordance measures the share of periods in which two cycles have the same sign, thereby assessing whether the padding choice affects the classification of observations into positive or negative cyclical phases. Turning-point concordance is calculated as the share of turning points identified under the reference specification for which a turning point of the same type is found under the alternative specification, using one-to-one matching and a tolerance of ± 1 quarter. This tolerance accommodates limited dating uncertainty in quarterly data while still requiring the turning points to be identified almost contemporaneously. The $0.5T$ padding rule is used as the reference specification, while $0.75T$ and T represent the alternative specifications.

Table 5: Mirror-padding validation results

Variable	Comparison	Pearson correlation	Sign concordance	Turning-point concordance
CREDIT	$0.5T$ vs. $0.75T$	0.992	0.987	1.000
CREDIT	$0.5T$ vs. T	0.996	0.987	1.000
RPPI	$0.5T$ vs. $0.75T$	0.954	0.920	0.750
RPPI	$0.5T$ vs. T	0.912	0.840	0.750
RIR	$0.5T$ vs. $0.75T$	0.997	0.987	1.000
RIR	$0.5T$ vs. T	0.993	0.947	1.000
BCA	$0.5T$ vs. $0.75T$	0.996	0.947	1.000
BCA	$0.5T$ vs. T	0.999	0.960	1.000

Source: authors' calculations.

The results reported in Table 5 broadly confirm the visual evidence. Across all variables, the cycles obtained under alternative padding lengths are strongly correlated with the $0.5T$ reference cycle. Pearson correlations range from 0.912 to 0.999, indicating that the overall cyclical dynamics are highly stable across endpoint-treatment specifications. Sign concordance ranges from 0.840 to 0.987 and is particularly high for CREDIT, BCA, and RIR. The somewhat lower values obtained for RPPI

indicate greater sensitivity in the classification of individual observations into positive and negative cyclical phases.

Turning-point concordance equals one for CREDIT, BCA, and RIR in all comparisons. This means that every peak and trough identified under the $0.5T$ reference specification has a corresponding turning point of the same type under the alternative padding specification within the permitted one-quarter window. For RPPI, turning-point concordance equals 0.750 under both alternative specifications: five of the six reference peaks and four of the six reference troughs are matched. All matched RPPI turning points occur in the same quarter as their reference counterparts. Therefore, the lower concordance reflects unmatched turning points rather than shifts in the dating of the matched peaks and troughs.

Overall, the selected mirror-padding length does not materially affect the general dynamics of the extracted cycles, although RPPI exhibits some sensitivity in phase classification and in the number of identified turning points. The $0.5T$ rule is therefore retained as the baseline specification. This choice is parsimonious, limits the extent of artificial endpoint extension, and produces cyclical signals that are broadly comparable to those obtained under the longer padding rules.

Cycle extraction and quasi-real-time validation

After selecting the mirror-padding rule, the cyclical components were extracted using CEEMDAN with a padding length equal to $0.5T$. For each variable, the CEEMDAN hyperparameters were selected through a grid-search procedure. The search focused on two parameters: the S-number, which controls the stopping criterion of the sifting process, and the noise intensity, which determines the standard deviation of the added noise relative to the standard deviation of the signal or residual entering the corresponding decomposition stage. The parameter values selected for the full sample analysed in this paper are reported in Table 6.

Table 6: Selected CEEMDAN hyperparameters

Variable	S-number	Noise intensity
CREDIT	7	0.40
RPPI	6	0.40
RIR	5	0.20
BCA	8	0.35

Source: authors' calculations.

The selected hyperparameters differ across variables, reflecting differences in the persistence, volatility, and frequency composition of the underlying series over the sample analysed. These values should therefore be interpreted as the outcome of the grid-search procedure for the present dataset, rather than as fixed parameters intended to remain unchanged as new observations become available. In an operational setting, the same selection procedure can be repeated when the sample is updated.

The cycles extracted using these parameter values represent the full-sample benchmark estimates. However, full-sample decompositions may incorporate information that was not available at each historical point in time, particularly when estimating the cycle near the sample endpoints. Therefore, the next step evaluates the reliability of the extracted cycles in a quasi-real-time setting. The exercise is conducted separately for the raw cycle, defined as the sum of all extracted IMFs excluding the residual component, and the adjusted cycle, obtained by excluding IMF1 and summing the remaining IMFs. Each cycle specification is first estimated over the full sample and subsequently re-estimated at each recursive endpoint as the available sample is expanded one quarter at a time. The resulting recursive endpoint estimates are then compared with the corresponding full-sample benchmarks.

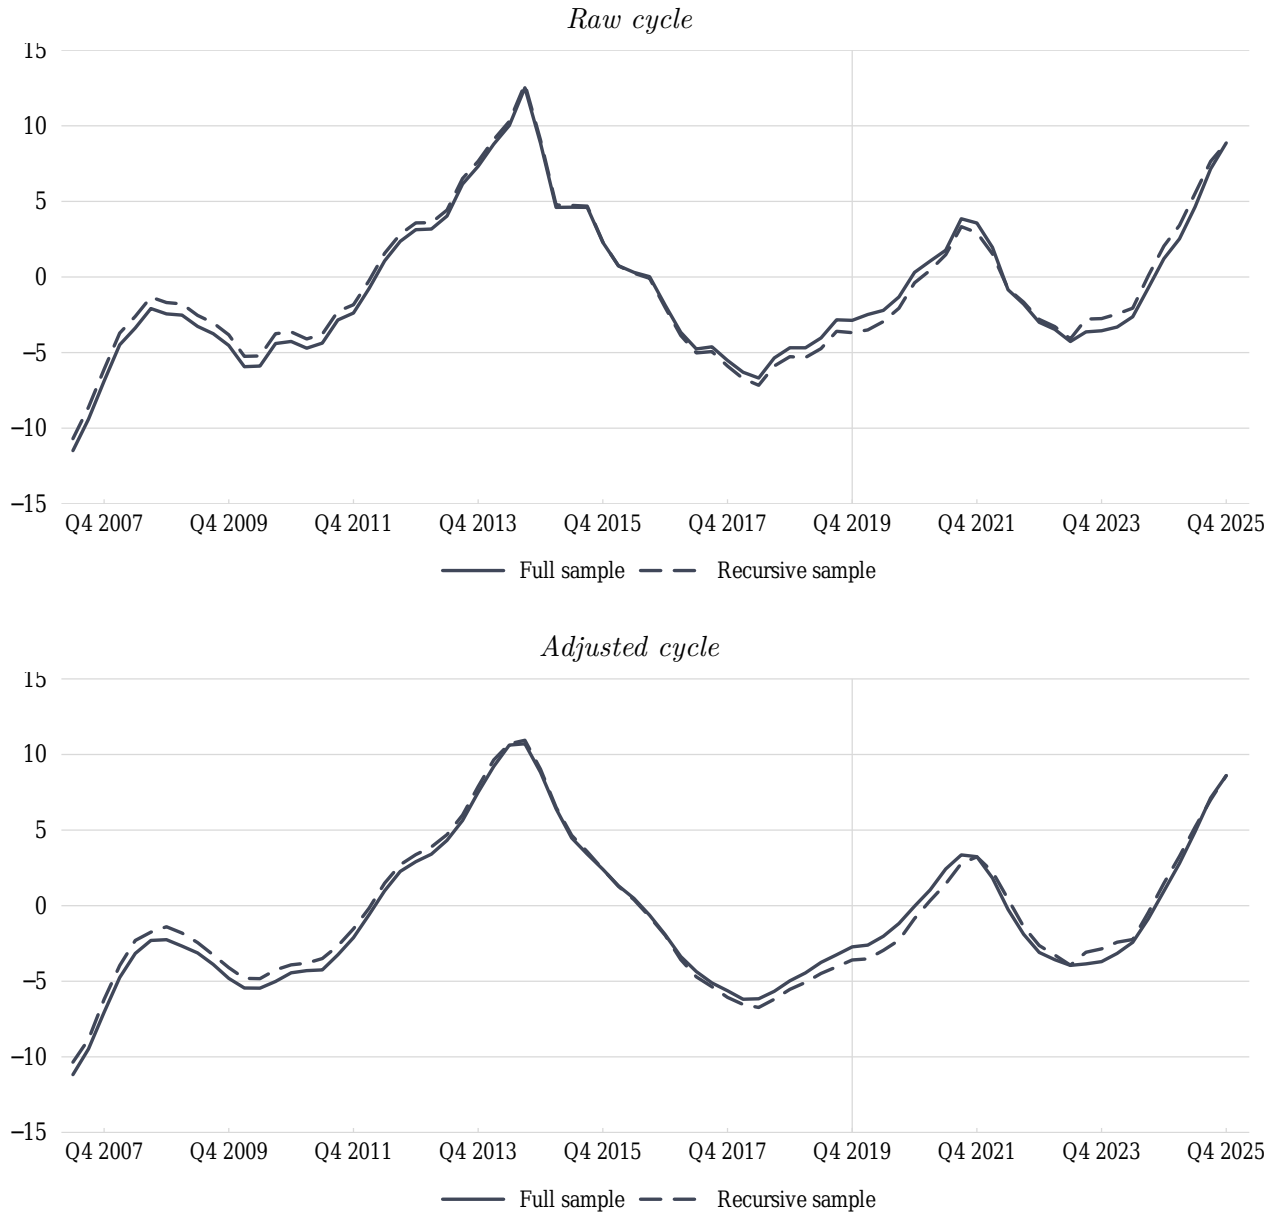
The initial estimation window covers 2007Q2–2019Q4, after which the sample is expanded one quar-

ter at a time through 2025Q4. At each recursive endpoint, the mirror-padding length is updated to $\lfloor 0.5T \rfloor$, and the CEEMDAN hyperparameters are reselected using the same grid-search procedure. The resulting recursive endpoint estimates are compared with the corresponding full-sample benchmarks.

The recursive evaluation period covers an exceptionally turbulent macro-financial environment. In 2020, the Moldovan economy was affected by the contraction associated with the socio-economic containment measures adopted to limit the spread of COVID-19 (National Bank of Moldova, 2021). The subsequent recovery coincided with a sharp increase in international prices (National Bank of Moldova, 2022). In 2022, Russia's invasion of Ukraine generated another major external shock, transmitted through several channels, including higher global price pressures, particularly in energy markets, and disruptions to trade with Russia and Ukraine (National Bank of Moldova, 2023). Calculations based on official foreign-trade statistics indicate that, in 2021, before the invasion, Russia and Ukraine jointly accounted for 11.7 percent of Moldova's total goods exports, while the unit value index of imported mineral products, a category that includes energy products, increased by 84.9 percent in 2022 relative to 2021 (National Bureau of Statistics of the Republic of Moldova, 2026c). Against this background, real GDP contracted by 4.6 percent in 2022, while the annual average inflation rate reached 28.7 percent (National Bureau of Statistics of the Republic of Moldova, 2026b; National Bureau of Statistics of the Republic of Moldova, 2026a). The banking sector also experienced temporary tensions at the beginning of the war, reflected in significant deposit withdrawals (National Bank of Moldova, 2023). Monetary policy was tightened substantially: the NBM base rate rose from 6.5 percent at the beginning of 2022 to 21.5 percent in August, was maintained at that level at the November monetary policy meeting, and was reduced to 20 percent in December (National Bank of Moldova, 2026). Taken together, these developments make the period from 2020Q1 to 2025Q4 a demanding setting for evaluating the temporal stability of the CEEMDAN-based extraction procedure in a quasi-real-time framework.

For credit, the full-sample and recursive-sample estimates are closely aligned over the entire period. Before 2019Q4, the cycle extracted from the initial estimation window reproduces very closely the full-sample benchmark, both in terms of direction and amplitude. During the recursive evaluation period, from 2020Q1 to 2025Q4, the recursive estimates continue to follow the same cyclical pattern as the full-sample cycle, including the expansion observed around 2021, the subsequent decline, and the renewed increase towards the end of the sample. This result is observed for both the raw and adjusted cycles, suggesting that the credit cycle is highly stable under recursive estimation.

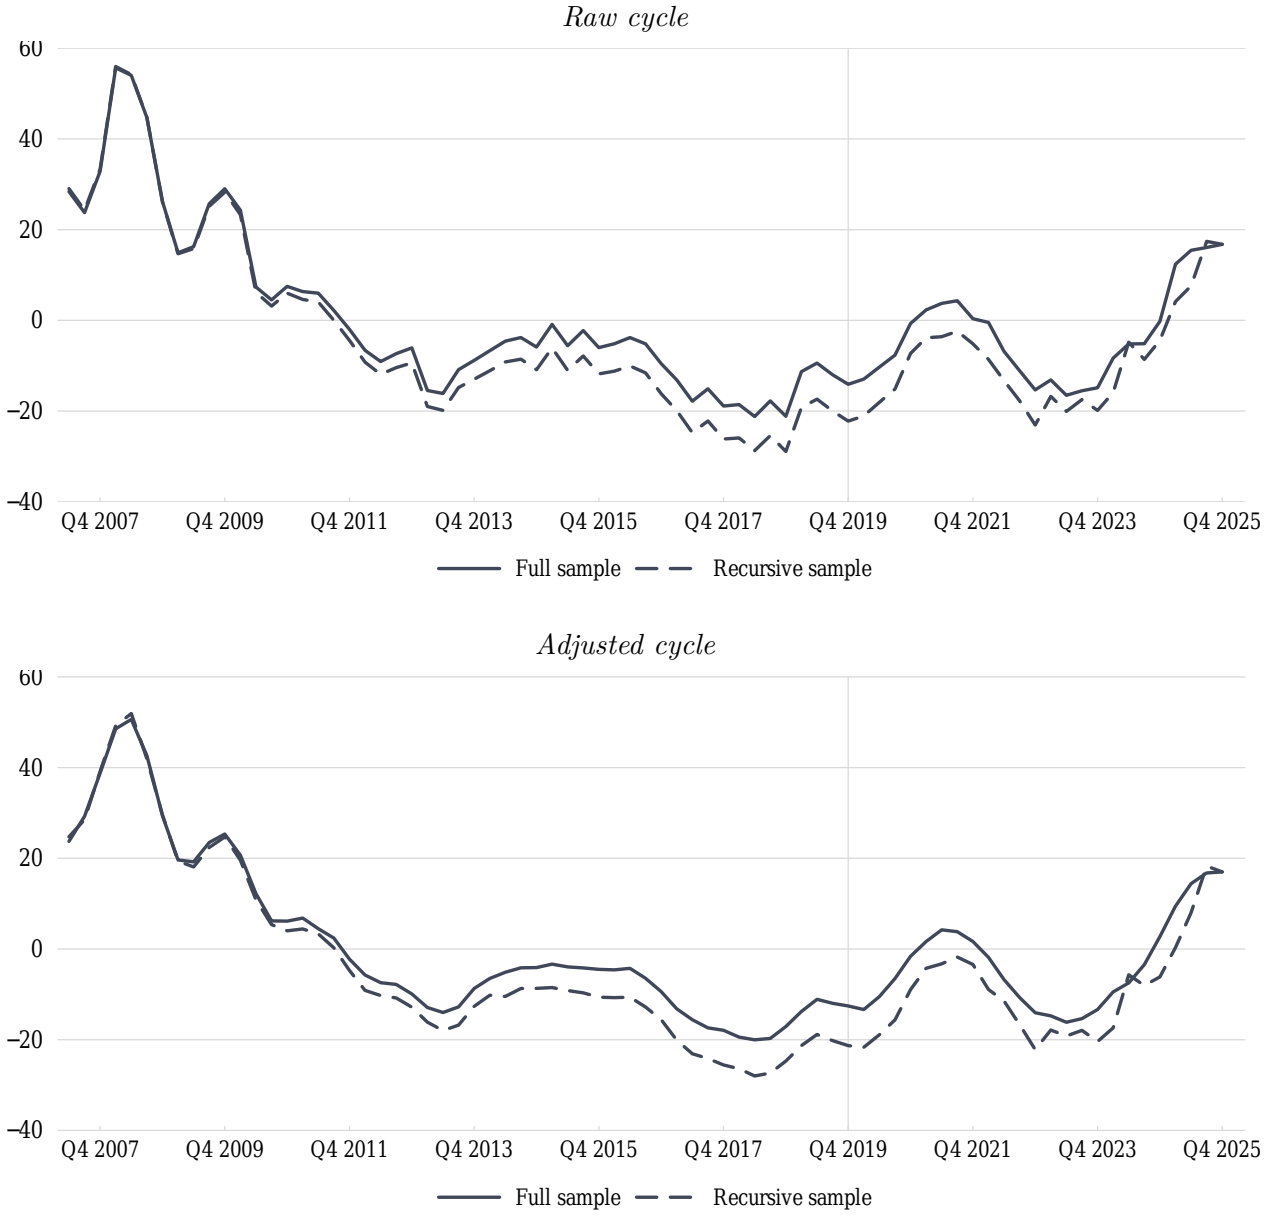
Figure 9: Full-sample vs recursive-sample cycles: Credit, billion MDL at 2015 prices



Source: authors' calculations.
Note: The recursive-sample line combines the cycle estimated over the initial 2007Q2–2019Q4 window with successive endpoint estimates for 2020Q1–2025Q4. Each endpoint estimate is obtained using only the observations available through the corresponding quarter.

For the residential property price index, the recursive-sample estimates display somewhat larger amplitude differences relative to the full-sample benchmark. Nevertheless, the two series remain broadly parallel and preserve the same medium-term direction both before and after 2019Q4. The recursive estimates capture the prolonged negative cyclical position observed after 2011, the temporary improvement around 2020–2021, and the renewed increase towards the end of the sample. The adjusted cycle shows a smoother pattern, but the main conclusion remains unchanged: the recursive estimates preserve the main cyclical phases of the full-sample benchmark.

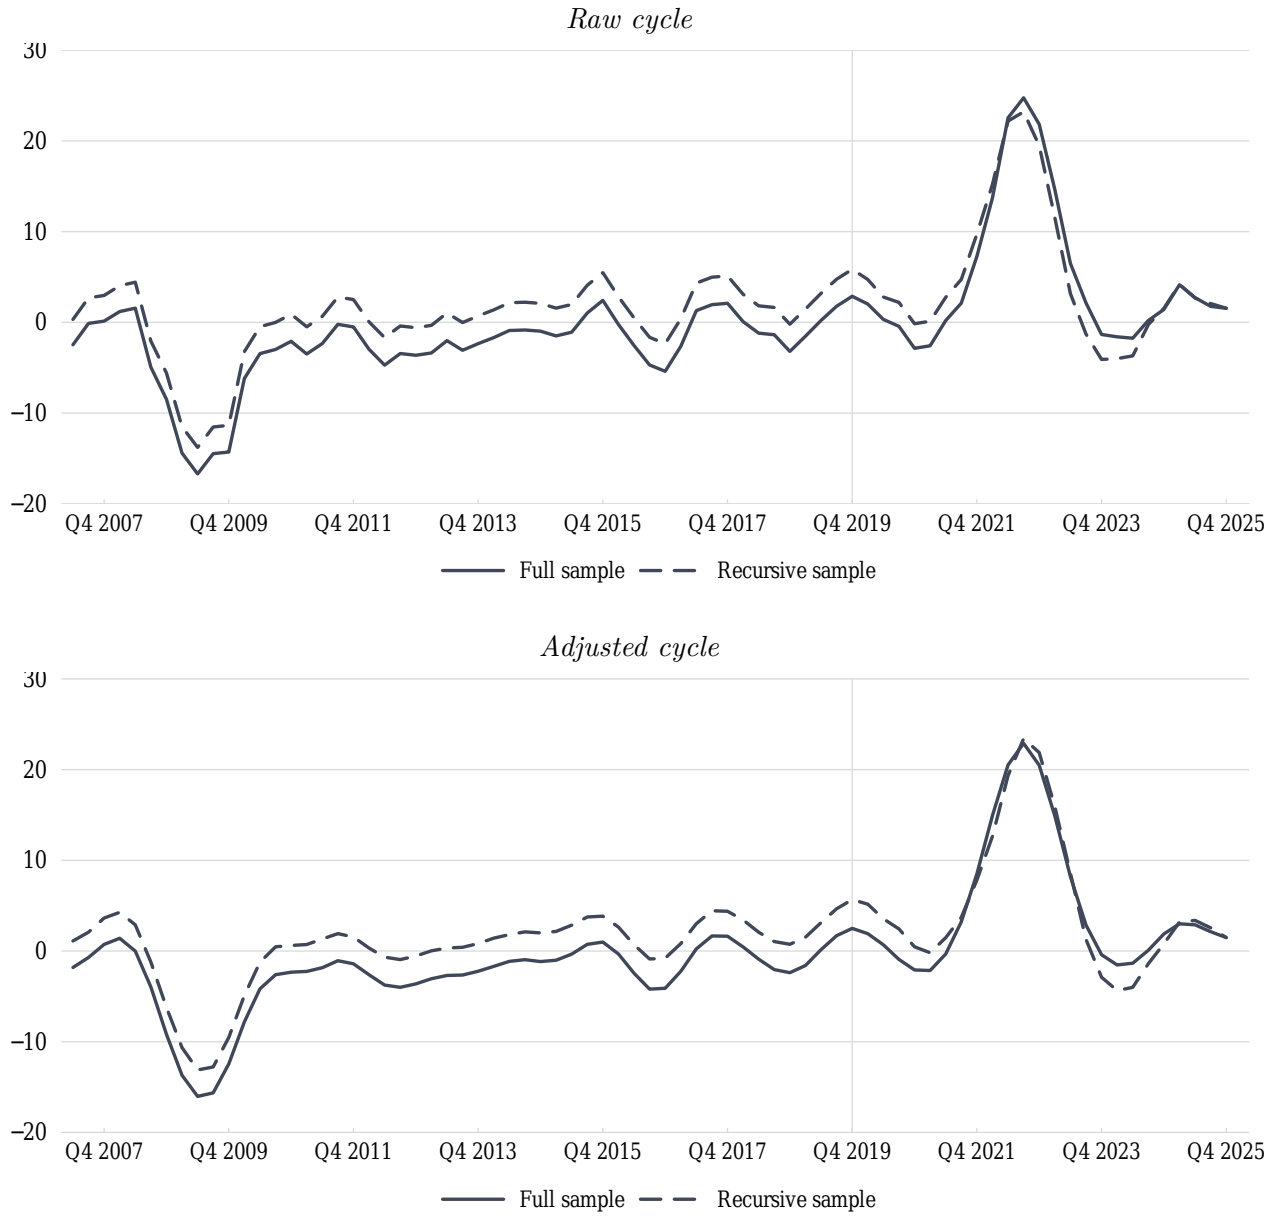
Figure 10: Full-sample vs recursive-sample cycles: Residential property price index, 2015 = 100



Source: authors' calculations.
Note: The recursive-sample line combines the cycle estimated over the initial 2007Q2–2019Q4 window with successive endpoint estimates for 2020Q1–2025Q4. Each endpoint estimate is obtained using only the observations available through the corresponding quarter.

For the real interest rate, the recursive-sample and full-sample cycles are also broadly consistent. Some differences are visible in the amplitude of short-term movements, especially before 2020, but the two series generally move in the same direction. Importantly, the recursive estimates capture the large cyclical movement observed in 2021–2022, which is the most prominent episode in the series. This indicates that the main signal is not driven by full-sample information alone and remains identifiable when the decomposition is updated recursively.

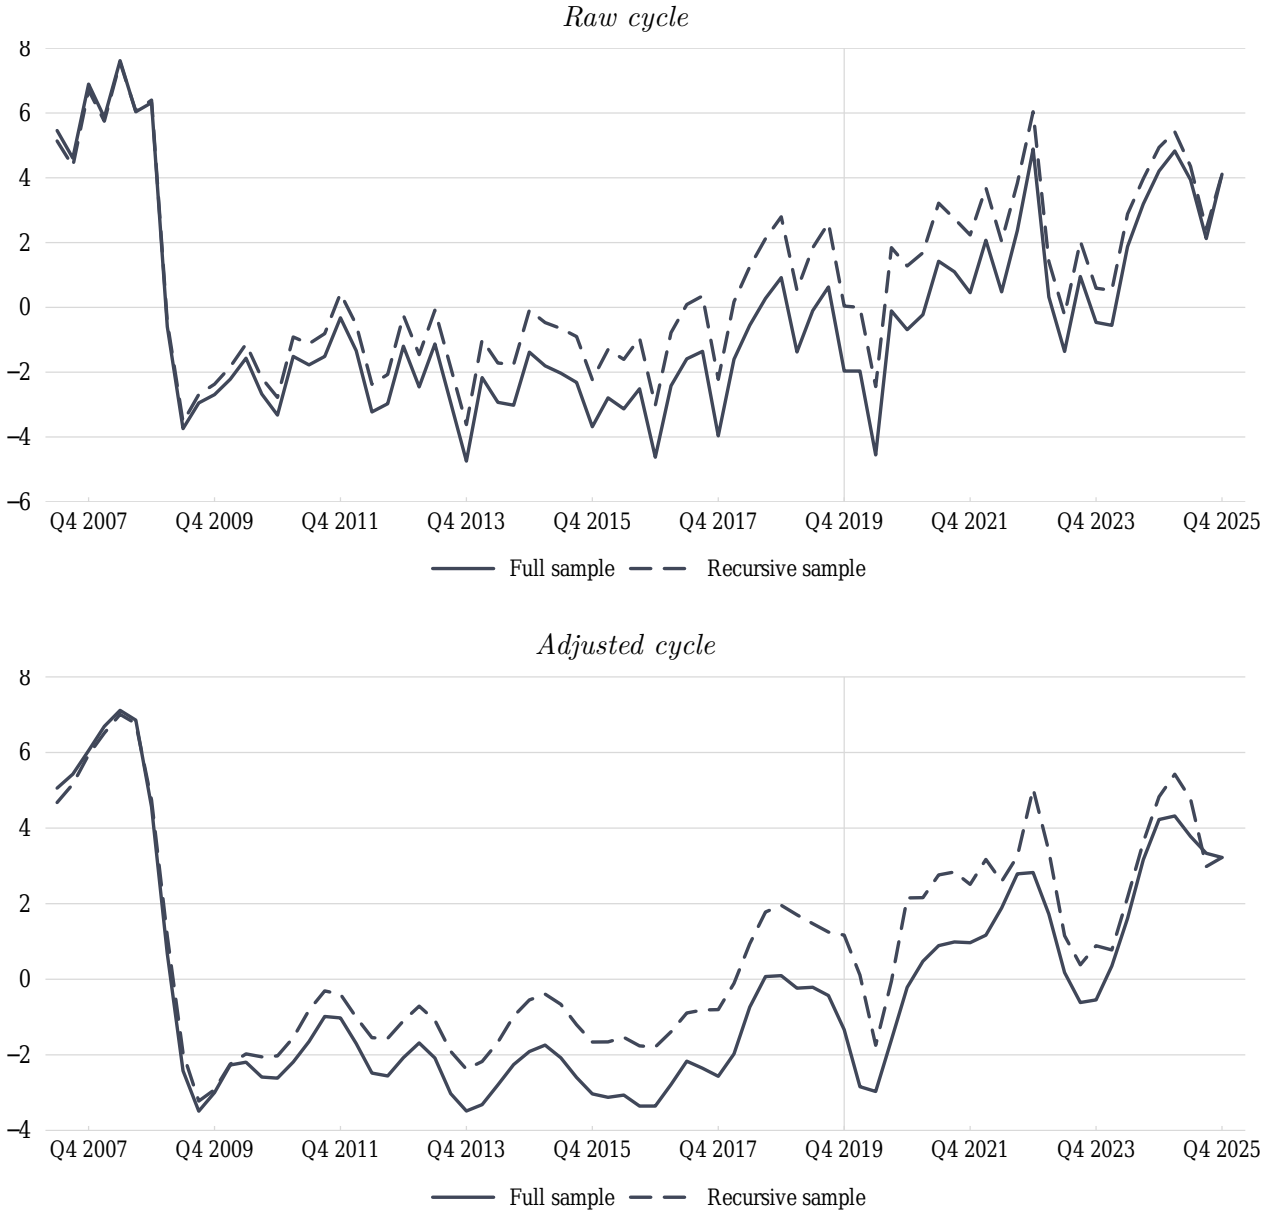
Figure 11: Full-sample vs recursive-sample cycles: Real interest rate, percent



Source: authors' calculations.
Note: The recursive-sample line combines the cycle estimated over the initial 2007Q2–2019Q4 window with successive endpoint estimates for 2020Q1–2025Q4. Each endpoint estimate is obtained using only the observations available through the corresponding quarter.

For the current account balance, the recursive and full-sample estimates are less tightly aligned than in the case of credit, reflecting the more volatile nature of this variable. The raw cycle displays visible short-term fluctuations, while the adjusted cycle provides a smoother representation of the same underlying dynamics. Despite these differences, the recursive estimates broadly preserve the direction of movement and the main cyclical phases, including the improvement after 2020 and the fluctuations observed during the 2022–2025 period. This suggests that, even for the most volatile component, recursive estimation does not generate a fundamentally different cyclical signal.

Figure 12: Full-sample vs recursive-sample cycles: Current account balance, billion MDL at 2015 prices



Source: authors' calculations.
Note: The recursive-sample line combines the cycle estimated over the initial 2007Q2–2019Q4 window with successive endpoint estimates for 2020Q1–2025Q4. Each endpoint estimate is obtained using only the observations available through the corresponding quarter.

The visual evidence indicates that the recursive-sample estimates remain broadly aligned with the full-sample cycles across the four variables. The degree of alignment is strongest for credit and somewhat weaker for the current account balance, but in all cases the recursive estimates preserve the main direction and cyclical phases of the full-sample benchmark. To complement the visual assessment and provide a more precise evaluation of quasi-real-time reliability, three quantitative indicators are used: Pearson correlation, sign concordance, and the signal-to-noise measure. These indicators assess, respectively, the similarity of the overall cyclical dynamics, the consistency of the cyclical position relative to zero, and the size of revisions relative to the full-sample cyclical signal.

The Pearson correlations reported in Table 7 are high for all variables and cycle specifications. They range from 0.924 to 0.984, indicating a strong degree of co-movement between the recursive endpoint

estimates and the corresponding full-sample cycles. The highest values are observed for credit, where both the raw and adjusted cycles record a correlation of 0.984. The lowest value is observed for the adjusted current account balance cycle, at 0.924, but even this result indicates a strong relationship between the recursive and full-sample estimates. Overall, the correlation results confirm that the recursive procedure preserves the main dynamics of the extracted cycles.

Sign concordance also provides supportive evidence. The recursive and full-sample estimates generally place the cycles on the same side of zero, with values ranging from 0.792 to 0.958. Credit again shows the strongest performance, especially for the adjusted cycle, where sign concordance reaches 0.958. For RPPI, RIR, and BCA, the values are somewhat lower, but still indicate that the recursive estimates usually imply the same cyclical phase as the full-sample benchmark. This is important from a policy perspective, since the sign of the cycle is directly related to the interpretation of whether a variable is in a positive or negative cyclical position.

The signal-to-noise measure further supports the quasi-real-time reliability of the extraction procedure. All values are positive and range from 0.614 to 0.821, implying that revisions are smaller than the variability of the corresponding full-sample cyclical signal. The strongest results are again observed for credit, while the lowest value is recorded for the adjusted BCA cycle. This confirms that revisions exist, especially for more volatile variables, but they do not dominate the underlying cyclical signal. Therefore, the recursive estimates can be interpreted as sufficiently informative for analytical monitoring, even though they should not be treated as revision-free estimates.

Table 7: Quasi-real-time validation results

Variable	Cycle type	Pearson correlation	Sign concordance	Signal-to-noise
CREDIT	Raw	0.984	0.917	0.820
CREDIT	Adjusted	0.984	0.958	0.821
RPPI	Raw	0.968	0.833	0.724
RPPI	Adjusted	0.961	0.792	0.689
RIR	Raw	0.963	0.833	0.730
RIR	Adjusted	0.971	0.833	0.759
BCA	Raw	0.976	0.792	0.746
BCA	Adjusted	0.924	0.833	0.614

Source: authors' calculations.

Overall, the quasi-real-time validation provides evidence of temporal stability over the recursive evaluation window. Re-estimating the decomposition recursively and reselecting the hyperparameters as the sample expands generally preserves the sign of the full-sample estimates, although the magnitude of revisions varies across variables. The recursive estimates also remain broadly aligned with the full-sample benchmarks in terms of overall dynamics and revision size. These results support the potential use of the procedure as an analytical input for future updates and macroprudential monitoring.

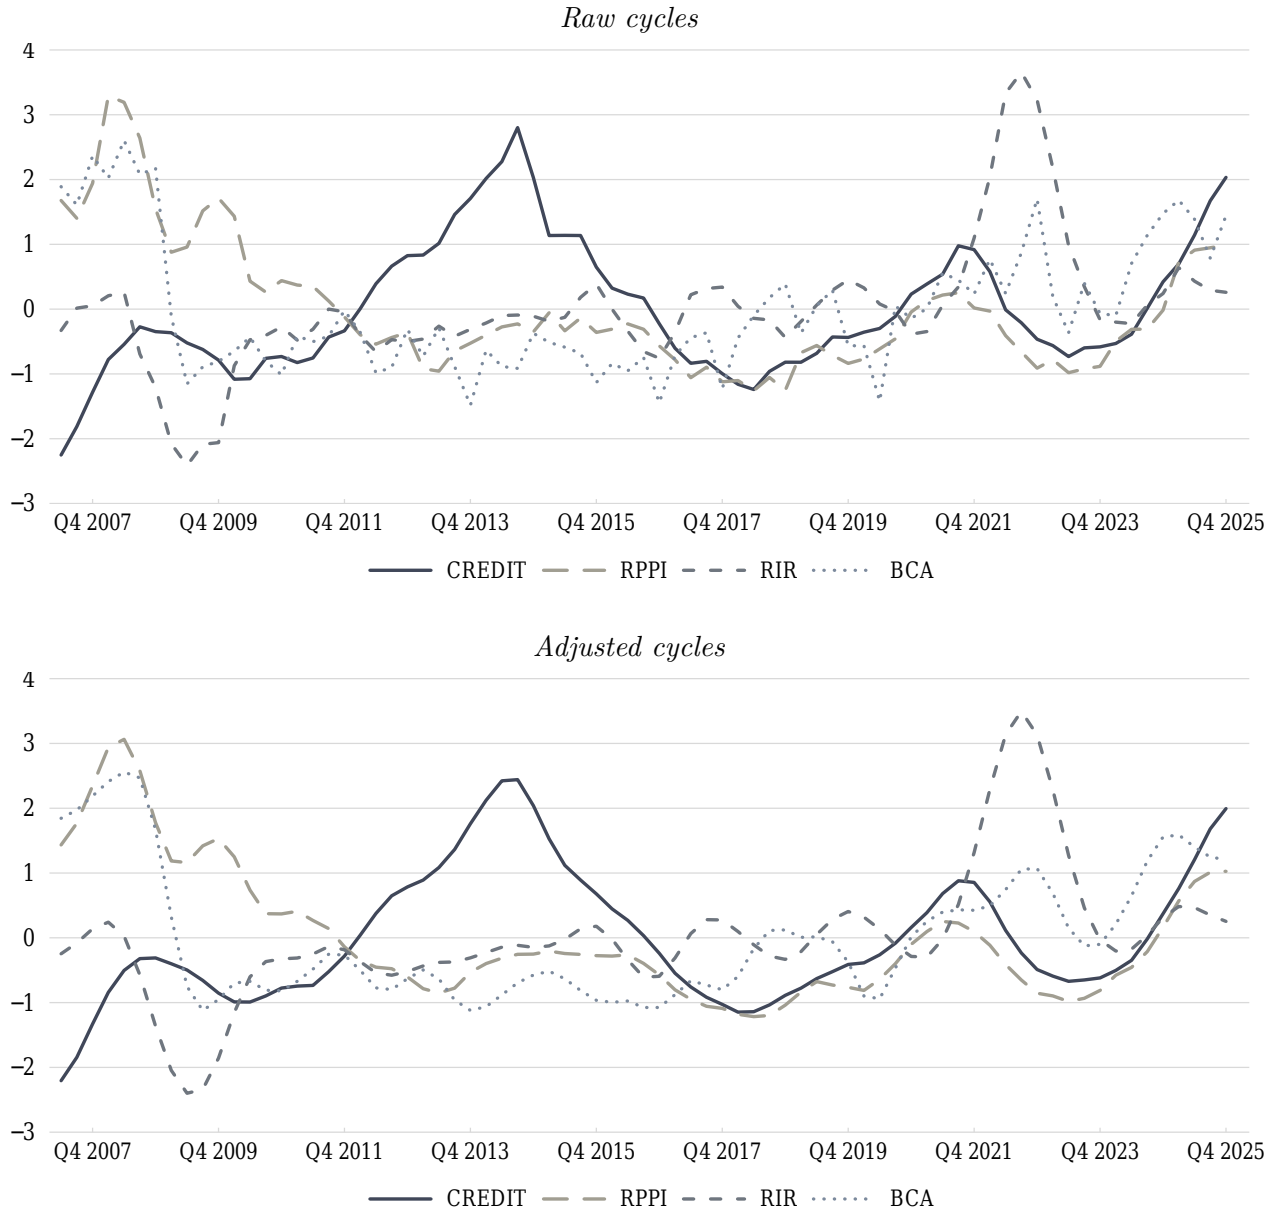
Normalization

After extracting the cyclical components, the raw and adjusted cycles are normalized using the full-sample z-score defined in the methodology, with the mean and standard deviation calculated separately for each variable and cycle specification over 2007Q2–2025Q4. This step is necessary because the original variables are measured in different units and have different ranges of variation. Normalization expresses each cyclical component in standard deviation units, so that no variable dominates the composite indicator mechanically because of its original scale.

To assess the evolution of the normalized cyclical components over time, the raw and adjusted spec-

ifications are examined jointly. Figure 13 illustrates their medium-term dynamics and the extent of short-term fluctuations across the four variables.

Figure 13: Normalized cyclical components, standard deviation units



Source: authors' calculations.

Overall, the two specifications display broadly similar medium-term dynamics. The main expansionary and contractionary phases are preserved after excluding IMF1, and the relative position of the variables remains largely unchanged. This suggests that the raw and adjusted cycles contain a common cyclical signal. At the same time, the adjusted cycles display a visibly smoother profile, especially for the more volatile components. This evidence supports the view that continuing the analysis with both raw and adjusted cycles would provide limited additional information, while increasing the complexity of the aggregation exercise.

The preference for adjusted cycles is also consistent with the economic interpretation of the financial cycle. The objective of the composite indicator is to capture the gradual accumulation of cyclical systemic risk, rather than short-lived fluctuations. In the CEEMDAN decomposition, IMF1 represents the highest-frequency component and may therefore contain movements that are less relevant for medium-term macroprudential analysis. This is particularly important because the financial-cycle

literature generally associates the financial cycle with medium-term fluctuations in credit and asset prices, rather than with very short-term oscillations (Drehmann et al., 2012; Borio, 2014). Consequently, excluding IMF1 helps align the extracted signal more closely with the concept of a financial cycle used for systemic-risk monitoring.

The smoother behaviour of the adjusted cycles is confirmed by the median absolute deviation of second differences. For a normalized cyclical component $x(t)$, the second difference is calculated as $\Delta^2 x(t) = [x(t) - x(t-1)] - [x(t-1) - x(t-2)]$. The median absolute deviation of the second differences is then computed as $\text{MAD}(\Delta^2 x) = \text{median}_t \left| \Delta^2 x(t) - \text{median}_t \left(\Delta^2 x(t) \right) \right|$. This measure captures the dispersion of changes in the slope of the cycle, with lower values indicating fewer abrupt short-term movements and a smoother cyclical trajectory. For all variables, the measure declines after excluding IMF1. For CREDIT, it decreases from 0.148 to 0.071, while for RPPI it declines from 0.261 to 0.098. A reduction is also observed for RIR, from 0.233 to 0.141, and is particularly pronounced for BCA, where the measure falls from 0.666 to 0.159. These results confirm that excluding IMF1 substantially reduces abrupt short-term variations. Together with the visual evidence presented in Figure 13, they also indicate that the adjustment preserves the broader cyclical pattern observed in the raw series.

The quasi-real-time validation provides an additional robustness check for the use of the adjusted cycles in the aggregation stage. As shown in Table 7, both raw and adjusted cycles remain broadly aligned with their corresponding full-sample benchmarks. Although excluding IMF1 does not uniformly improve every quasi-real-time metric, the adjusted cycles continue to display high correlations, generally high sign concordance, and positive signal-to-noise measures when the recursive estimates are compared with their full-sample counterparts. The smoother representation obtained through the adjustment is therefore not associated with a qualitatively different conclusion regarding the temporal stability of the extracted cycles.

Overall, the adjusted normalized cycles are retained for the aggregation stage. The raw cycles remain useful as diagnostic benchmarks, but the adjusted cycles are more appropriate for constructing a composite indicator of cyclical systemic risk, since they better reflect the medium-term nature of financial-cycle dynamics and reduce high-frequency fluctuations that may be less relevant for macroprudential policy analysis.

Although the preceding analysis supports the exclusion of IMF1 relative to the raw-cycle specification, it does not by itself establish whether IMF2 should also be removed. The cycle excluding IMF1 is treated as the reference specification, and turning points are matched using the same one-quarter tolerance applied in the mirror-padding assessment. A targeted sensitivity exercise is therefore conducted by comparing the normalized adjusted cycle excluding IMF1 with a more restrictive reconstruction excluding both IMF1 and IMF2. The comparison uses Pearson correlation, sign concordance and turning-point concordance. The same criteria employed in the mirror-padding assessment are applied because both exercises evaluate whether an alternative technical specification materially changes three distinct properties of the extracted cycle: its overall dynamics, its position relative to the zero reference level and the timing of its principal peaks and troughs. Their repeated use also ensures a consistent and scale-independent assessment across the two robustness exercises.

The sensitivity analysis is limited to the additional exclusion of IMF2 because it is the component immediately adjacent to IMF1 in the frequency ordering and therefore represents the most relevant marginal alternative to the selected reconstruction. Extending the exclusion to IMF3 or higher-order modes would progressively remove more persistent components and would consequently modify the economic object being measured rather than test the robustness of the boundary between short-lived fluctuations and the retained cyclical signal. In the absence of prior evidence that these higher-order modes represent noise, such increasingly restrictive reconstructions are not considered. The exercise is therefore a local sensitivity assessment around the selected specification rather than a search for the smoothest possible cycle. Table 8 reports the resulting comparison.

Table 8: Comparison of normalized cycles under alternative IMF exclusions

Variable	Pearson correlation	Sign concordance	Turning-point concordance
CREDIT	1.000	1.000	1.000
BCA	0.999	0.987	0.765
RPPI	0.999	1.000	0.900
RIR	1.000	0.987	1.000

Source: authors' calculations.

The results indicate that the additional exclusion of IMF2 produces virtually no change in the overall dynamics or sign classification of the normalized cycles. Pearson correlations range from 0.999 to 1.000, while sign concordance ranges from 0.987 to 1.000. Turning-point concordance equals 1.000 for credit and RIR and remains high for RPPI at 0.900. The lower value obtained for BCA, at 0.765, indicates that removing IMF2 alters part of its turning-point structure and may therefore discard information relevant to the timing of the cycle. The more restrictive reconstruction does not materially alter the broad cyclical dynamics or sign classification, while for BCA it reduces the preservation of turning-point timing. The results therefore provide no evidence that the additional filtering restriction improves the economic representation of the extracted cycles. The normalized adjusted cycles excluding IMF1 only are therefore retained for the subsequent aggregation stage.

Aggregation and selection of the composite financial cycle index

Following the normalization of the adjusted cyclical components, alternative aggregation schemes are applied to construct the candidate composite indicators. The specifications differ in both their component composition and weighting scheme, allowing the contribution of additional variables and time-varying weights to be assessed.

Five alternative composite indicators are considered: EW2 combines RPPI and credit using an equal-weighted aggregation scheme; EW3 combines RPPI, credit, and BCA using equal weights; EW4 combines RPPI, credit, BCA, and RIR using equal weights; EWMA3 combines RPPI, credit, and BCA using the time-varying EWMA aggregation scheme; and EWMA4 combines RPPI, credit, BCA, and RIR using the time-varying EWMA aggregation scheme. Because BCA and RIR were multiplied by -1 during the initial data-transformation stage, their normalized cyclical components enter the aggregation with a risk-consistent orientation: a larger current account deficit and a lower real interest rate contribute positively to the accumulation of cyclical systemic risk.

Table 9 reports the identified peaks and troughs and the average peak-to-peak cycle duration under both the sign-restricted and unrestricted turning-point specifications. The comparison indicates that the estimated duration of the financial cycle is sensitive to the definition of the turning points. This finding is consistent with the broader empirical literature, which shows that estimates of financial-cycle duration can vary with the choice of indicators and the method used to identify cyclical fluctuations (Škare et al., 2025).

Under the sign-restricted specification, all candidate composite indicators display substantially longer average cycles than the GDP gap. The average duration is 26.5 quarters for EW2, EW3, and EWMA3 and 28.5 quarters for EW4 and EWMA4, compared with 12.0 quarters for the GDP gap. The sign-restricted specification offers a more direct economic interpretation for the present application because it focuses on broader movements between positive and negative cyclical positions, rather than on local reversals occurring on the same side of the zero reference level.

The results are more heterogeneous when the sign restriction is removed. EW2 remains considerably longer than the GDP gap, with an average duration of 26.5 quarters, while EW3 and EW4 have average

durations of 13.2 and 14.2 quarters, respectively. By contrast, EWMA3 and EWMA4 have shorter average cycles than the GDP gap, at 10.6 and 9.5 quarters, respectively. The unrestricted credit-to-GDP gap has an average duration of 28.5 quarters, whereas under the sign-restricted specification only one positive peak is identified and an average peak-to-peak duration cannot be calculated.

Table 9: Turning points and average cycle duration

	Indicator	Peaks	Troughs	Average peak-to-peak duration, quarters
Sign-restricted turning points	EW2	2008Q2; 2014Q3; 2021Q3	2010Q3; 2018Q2; 2023Q2	26.5
	EW3	2008Q2; 2014Q3; 2021Q3	2010Q3; 2018Q1; 2023Q3	26.5
	EWMA3	2008Q2; 2014Q3; 2021Q3	2010Q3; 2018Q1; 2023Q3	26.5
	EW4	2008Q2; 2014Q3; 2022Q3	2009Q3; 2018Q1; 2023Q4	28.5
	EWMA4	2008Q2; 2014Q3; 2022Q3	2009Q3; 2018Q1; 2023Q4	28.5
	Credit-to-GDP gap	2008Q1	2018Q1	n.a.
	GDP gap	2008Q3; 2011Q3; 2014Q3; 2019Q2; 2021Q1; 2023Q3	2009Q1; 2012Q4; 2015Q4; 2020Q2; 2022Q3; 2024Q3	12.0
No sign-restricted turning points	EW2	2008Q2; 2014Q3; 2021Q3	2010Q3; 2018Q2; 2023Q2	26.5
	EW3	2008Q2; 2011Q3; 2014Q3; 2019Q2; 2021Q3	2010Q3; 2012Q2; 2018Q1; 2020Q1; 2023Q3	13.2
	EWMA3	2008Q2; 2011Q3; 2013Q1; 2014Q3; 2019Q2; 2021Q3	2010Q3; 2012Q2; 2013Q3; 2018Q1; 2020Q1; 2023Q3	10.6
	EW4	2008Q2; 2011Q3; 2014Q3; 2019Q3; 2022Q3	2009Q3; 2012Q2; 2018Q1; 2020Q1; 2023Q4	14.2
	EWMA4	2008Q2; 2011Q3; 2013Q1; 2014Q3; 2017Q2; 2019Q3; 2022Q3	2009Q3; 2012Q2; 2013Q3; 2016Q4; 2018Q1; 2020Q1; 2023Q4	9.5
	Credit-to-GDP gap	2008Q1; 2014Q3; 2022Q2	2011Q1; 2018Q1; 2023Q1	28.5
	GDP gap	2008Q3; 2011Q3; 2014Q3; 2019Q2; 2021Q1; 2023Q3	2009Q1; 2012Q4; 2015Q4; 2020Q2; 2022Q3; 2024Q3	12.0

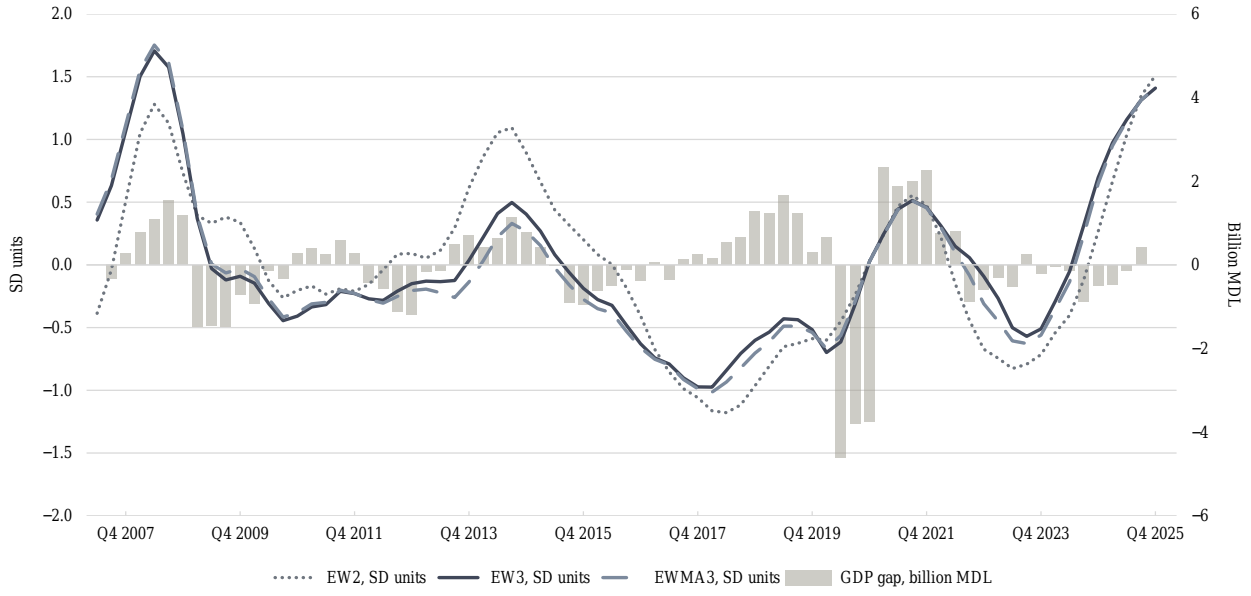
Source: authors' calculations.

Note: The sign-restricted specification retains positive peaks and negative troughs, whereas the specification without sign restrictions identifies local turning points without imposing restrictions on the sign of peaks and troughs.

The differences reported in Table 9 show that the conclusion as to whether a candidate financial-cycle indicator is longer than the real economic cycle depends materially on the turning-point specification. Cycle duration therefore provides useful descriptive information on the cyclical properties of the candidate indicators, but does not offer a sufficiently robust standalone basis for selecting among them. It is consequently treated as supporting evidence alongside the historical timing of financial-cycle peaks and the economic interpretation of the component contributions.

To complement the turning-point evidence reported in Table 9, the temporal relationship between the candidate composite indicators and the real economic cycle is examined graphically. Figure 14 compares EW2, EW3, and EWMA3 with the GDP gap. EW3 and EWMA3 follow almost identical trajectories, indicating that replacing equal weights with time-varying EWMA weights has only a limited effect on the aggregate dynamics. EW2 exhibits some differences in amplitude but preserves the same principal turning points. For both the 2009 and 2015 downturn episodes, all three indicators reach a positive peak approximately one to two quarters before the GDP gap crosses from positive to negative territory.

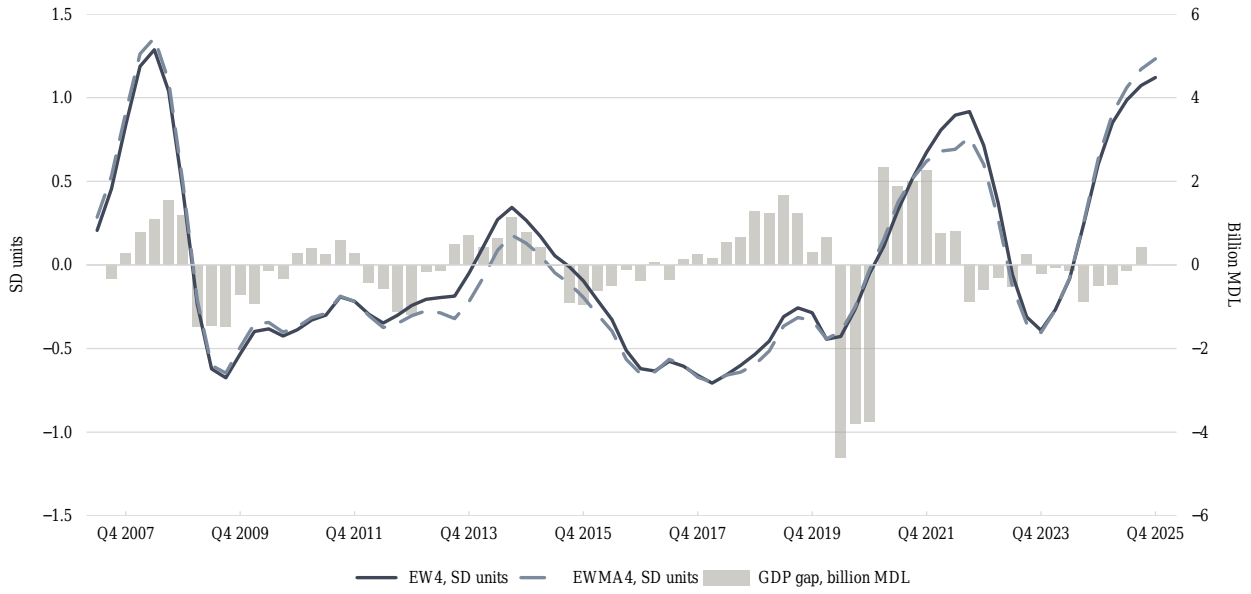
Figure 14: Two- and three-variable specifications and GDP gap



Source: authors' calculations.

Figure 15 presents the corresponding comparison for the four-variable specifications. EW4 and EWMA4 display closely aligned dynamics over almost the entire sample, with only limited differences in amplitude. Their peaks and troughs also occur in the same quarters, suggesting that the time-varying EWMA aggregation scheme does not materially alter the timing of the financial cycle relative to equal weighting. As in the two- and three-variable specifications, both indicators reach positive peaks approximately one to two quarters before the GDP gap crosses from positive to negative territory during the 2009 and 2015 episodes.

Figure 15: Four-variable specifications and GDP gap

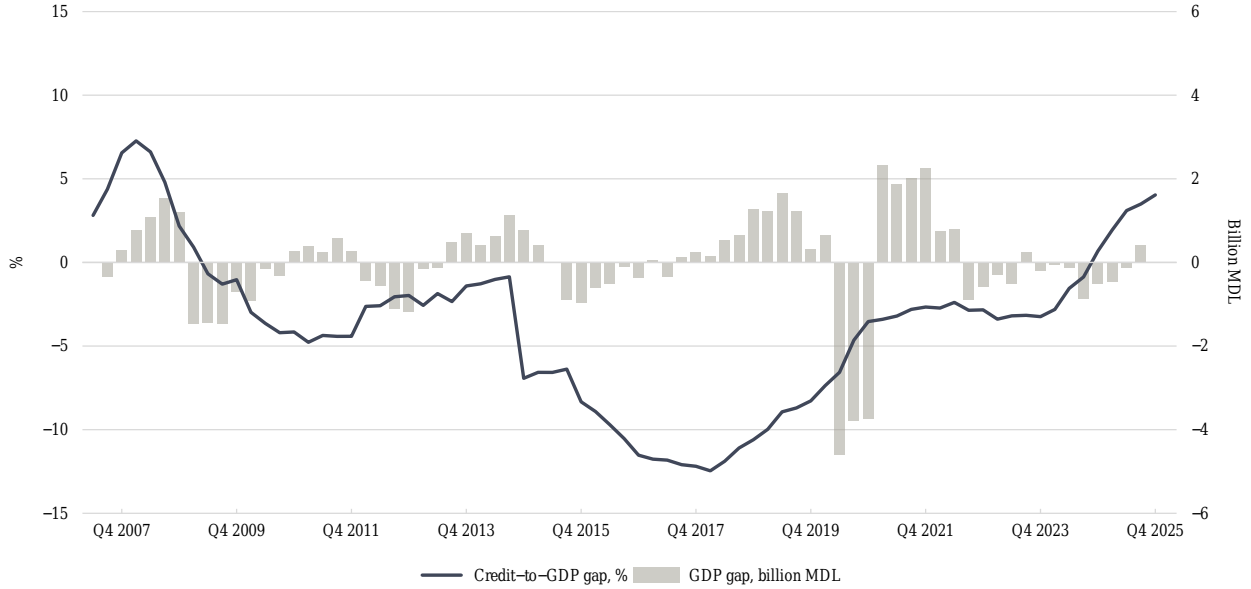


Source: authors' calculations.

Figure 16 compares the conventional credit-to-GDP gap with the GDP gap. The credit-to-GDP gap reaches a positive peak before the deterioration in real economic activity associated with the 2009 episode and subsequently declines as the GDP gap moves into negative territory. However, it

does not identify a comparable positive peak before the 2015 downturn and remains negative for a prolonged period thereafter. The credit-to-GDP gap therefore provides a relevant advance signal for the first episode but exhibits less consistent historical alignment across the two downturns used in the validation exercise.

Figure 16: Credit-to-GDP gap and GDP gap

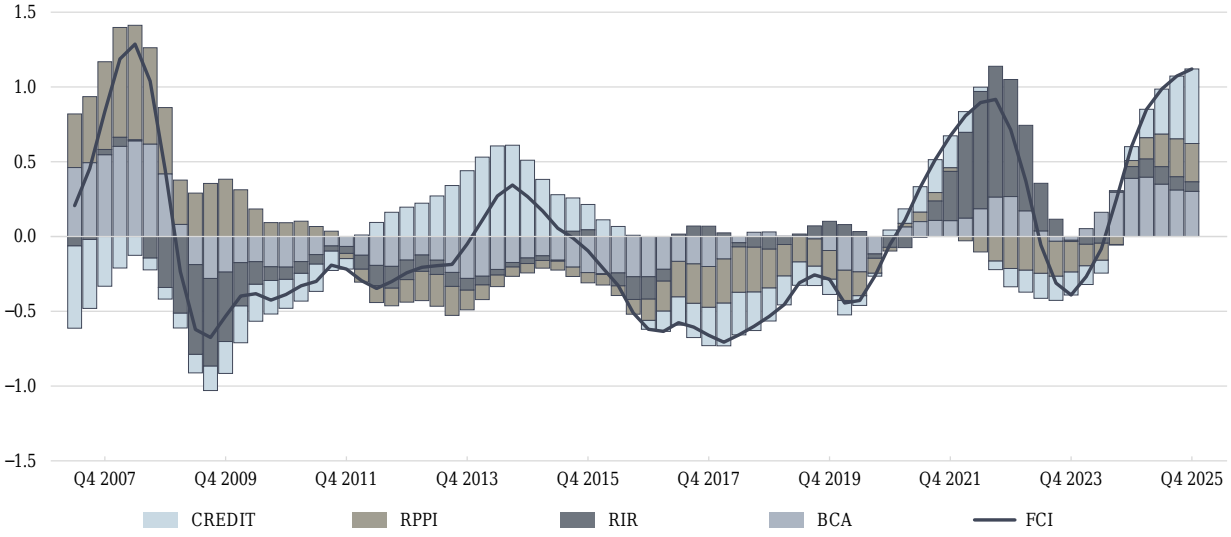


Source: authors' calculations.

Overall, the graphical evidence indicates that, in both selected downturn episodes, all candidate composite indicators reach positive peaks approximately one to two quarters before the GDP gap first moves from positive to negative territory. The equal-weighted and EWMA-weighted specifications also exhibit very similar dynamics and turning-point timing. Consequently, neither average cycle duration nor the historical lead relative to the GDP gap provides a sufficiently clear basis for selecting the preferred candidate. The analysis therefore proceeds to the third selection criterion, which evaluates the economic relevance of the alternative specifications by examining the contributions of their underlying components.

As shown in Figure 17, the decomposition of the EW4 index indicates that the initial upswing was driven mainly by positive contributions from RPPI, RIR, and BCA, which more than offset the negative contribution of credit in 2007–2008. The negative credit cycle during this period does not imply a decline in actual credit, which continued to grow, but indicates that credit remained below its estimated trend. The increase in 2021–2022 included a clearly positive contribution from RIR in 2022, while the recovery toward the end of the sample was supported by all four components.

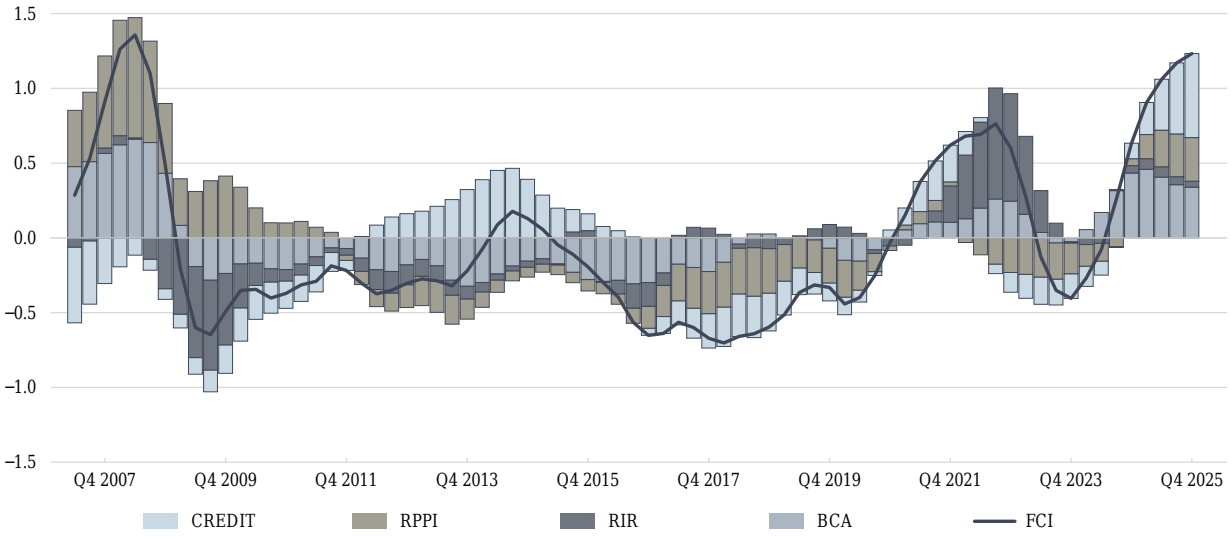
Figure 17: Component contributions to the EW4 index



Source: authors' calculations.

Figure 18 shows that the EWMA4 specification preserves the same signs and broad contribution pattern as EW4, although the time-varying weights modify the relative magnitude of the individual contributions. In particular, credit contributes negatively in 2007–2008 despite its actual expansion, because it remained below trend, while RIR makes a positive contribution in 2022. The final upswing is again supported by positive contributions from all four components.

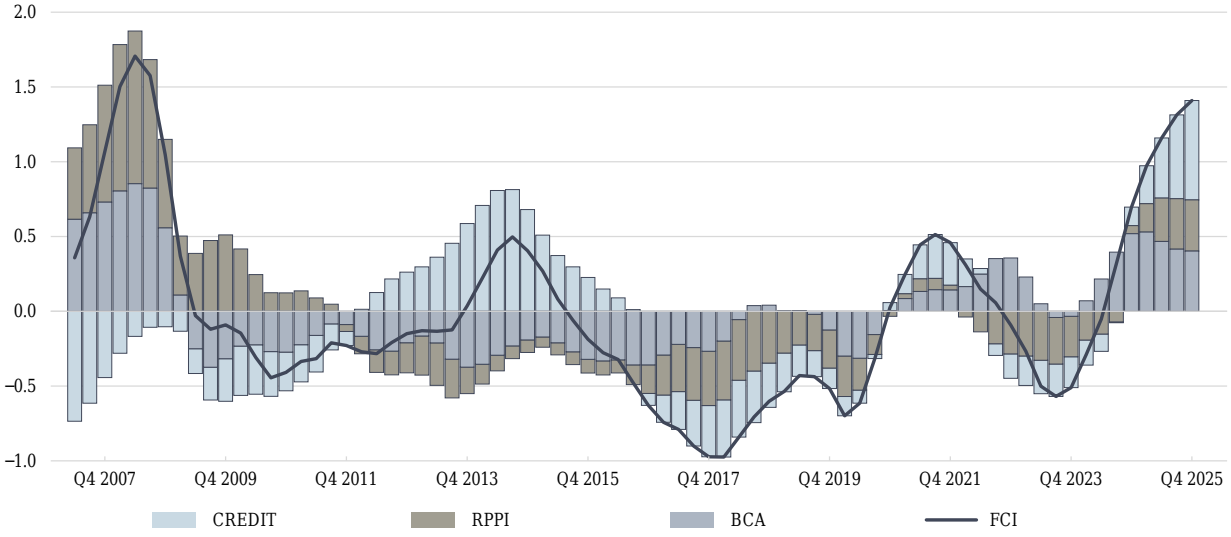
Figure 18: Component contributions to the EWMA4 index



Source: authors' calculations.

As illustrated in Figure 19, the EW3 decomposition indicates that the initial peak was driven by RPPI and BCA, which offset the negative credit contribution in 2007–2008. Although actual credit continued to expand in 2007–2008, its cyclical component was negative because credit remained below its estimated trend. Credit subsequently became the main driver of the expansion during 2012–2015, whereas the downturn around 2017 reflected negative contributions from all three components. The final increase is broad-based, with credit, RPPI, and BCA all contributing positively.

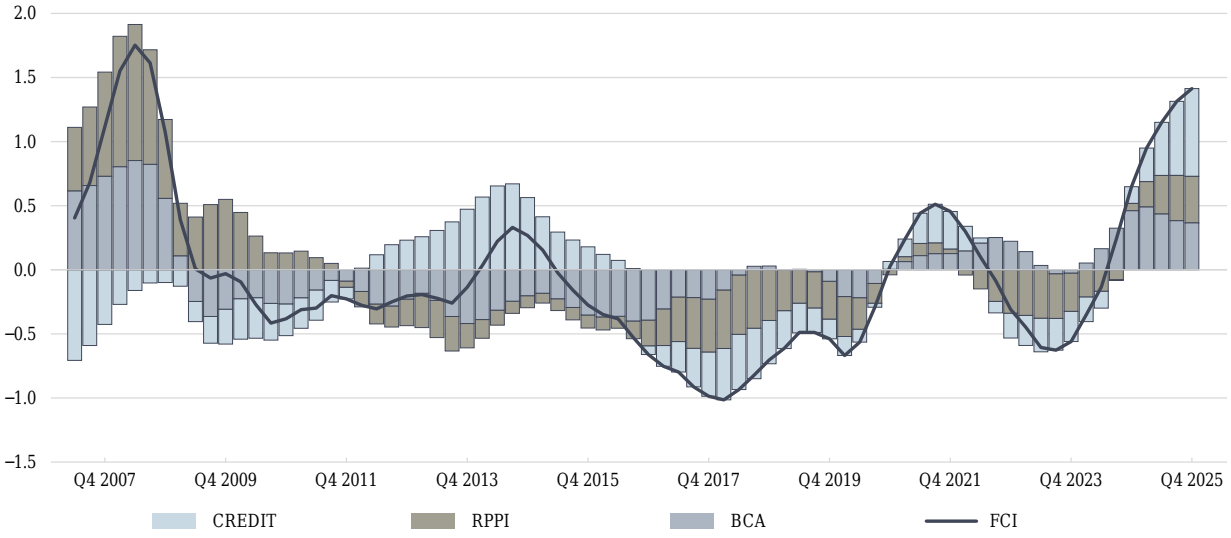
Figure 19: Component contributions to the EW3 index



Source: authors' calculations.

Figure 20 shows that the EWMA3 specification exhibits the same qualitative contribution pattern as EW3, with the differences being confined mainly to the magnitude of the individual contributions. RPPI and BCA dominate the initial expansion, credit plays the central role in the 2012–2015 increase, and all three components support the strong upswing at the end of the sample.

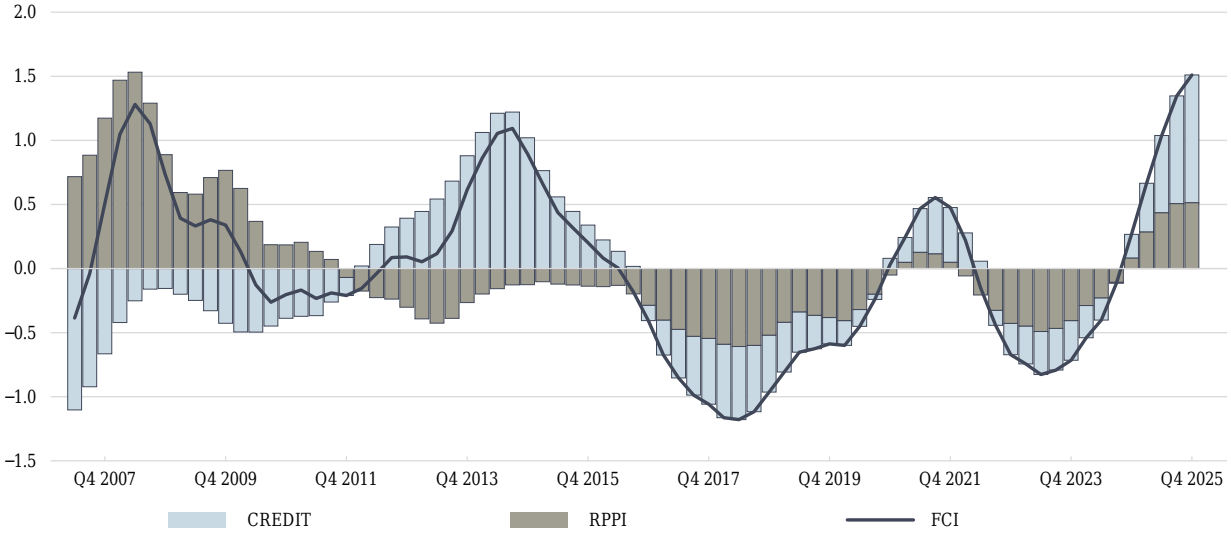
Figure 20: Component contributions to the EWMA3 index



Source: authors' calculations.

As shown in Figure 21, in the EW2 specification, the positive contribution from RPPI more than offsets the negative credit contribution during the initial expansion. The latter reflects credit remaining below its estimated trend in 2007–2008, even though the actual volume of credit continued to increase. Credit subsequently becomes the principal driver of the 2012–2015 upswing, while both components contribute to the downturn around 2017 and to the renewed increase toward the end of the sample.

Figure 21: Component contributions to the EW2 index



Source: authors' calculations.

As reported in Table 9, neither average cycle duration nor the historical lead relative to the GDP gap provides evidence in favour of the more complex EWMA aggregation scheme. The component decompositions also show that EW and EWMA specifications produce essentially the same economic interpretation. Moreover, the structure of the Moldovan economy does not provide an objective basis for assigning differentiated time-varying weights. Consequently, equal weighting offers a sufficiently informative, more parsimonious, and more transparent aggregation scheme.

The contribution of RIR in 2022 nevertheless reveals an important interpretative risk. During this episode, very high inflation exceeded the nominal lending rate, producing a strongly negative ex-post real interest rate which, after sign inversion, generated a positive contribution to the financial cycle. However, this outcome did not imply that effective financial conditions, particularly credit conditions, had become more favourable. Such a distortion is particularly consequential when the composite indicator is used as a reference for calibrating macroprudential instruments. It could overstate the accumulation of cyclical systemic risk and thereby support an untimely or excessively restrictive policy response, including an inappropriate calibration of the CCyB. Therefore, especially during periods of exceptionally high inflation, RIR may generate a misleading cyclical signal, making specifications that include this variable comparatively risky.

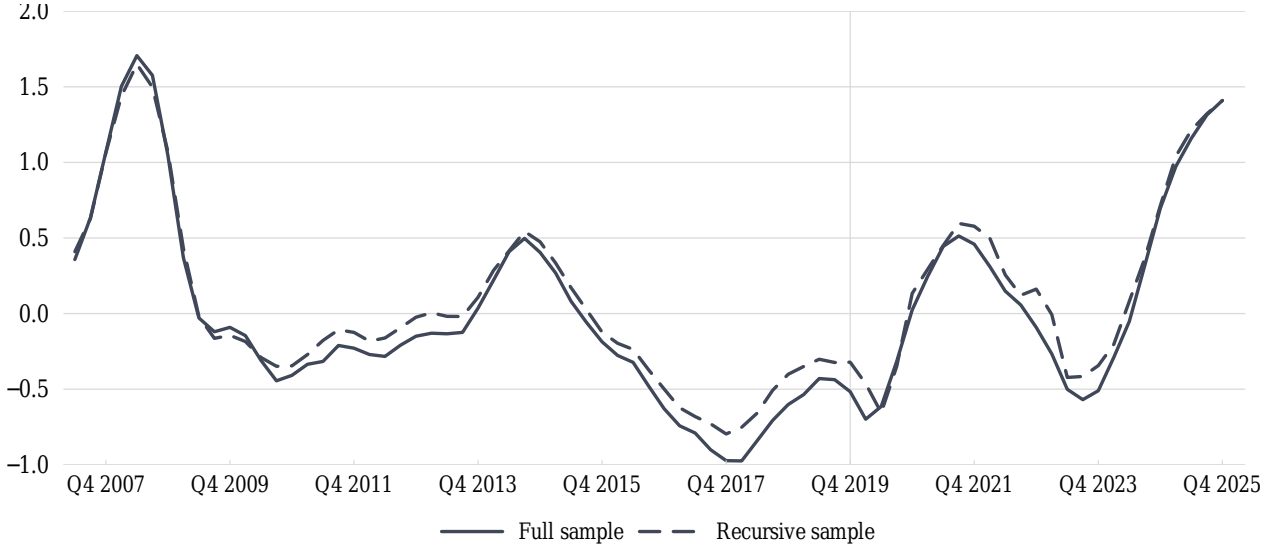
The exclusion of RIR leaves the choice between the EW2 and EW3 specifications. Although the quantitative evidence does not clearly distinguish between them, Moldova is a small open economy in which external imbalances can represent an important dimension of cyclical systemic risk. Including BCA is therefore preferable to restricting the index to credit and property prices alone. On this basis, EW3, comprising equally weighted contributions from credit, RPPI, and BCA, is selected as the preferred composite index.

Although EW3 was selected as the preferred specification, assessing its quasi-real-time performance remains important. The preceding analysis showed that the extraction of the individual component cycles is relatively stable across different sample vintages. Nevertheless, this does not automatically ensure the stability of the final composite index, since the extracted cycles are subsequently normalized and aggregated, and these transformations may propagate or modify the revisions arising from cycle extraction. Therefore, an additional quasi-real-time validation was conducted directly for EW3.

The recursive procedure followed the same approach as the validation of the individual cycles. The initial vintage covered 2007Q2–2019Q4, after which the sample was extended iteratively by one quarter, from 2020Q1 to 2025Q4. For each vintage, the entire estimation procedure was repeated, includ-

ing cycle extraction, normalization, and equal-weight aggregation. The resulting recursive estimates were compared with those obtained using the full sample. Their similarity over the evaluation period 2020Q1–2025Q4 was assessed using the Pearson correlation coefficient, sign concordance, and the signal-to-noise measure.

Figure 22: Quasi-real-time validation of the EW3 composite index



Source: authors' calculations.
Note: The recursive-sample line combines the EW3 index estimated over the initial 2007Q2–2019Q4 window with successive endpoint estimates for 2020Q1–2025Q4. Each endpoint estimate is obtained using only the observations available through the corresponding quarter.

As shown in Figure 22, the recursively estimated EW3 trajectory closely follows its full-sample counterpart, with only limited differences in magnitude. This visual evidence is supported by a Pearson correlation coefficient of 0.991. The sign concordance equals 0.917, indicating that the two estimates identify the same phase of the cycle in 22 of the 24 quarters considered. Furthermore, the signal-to-noise measure reaches 0.862, implying that the standard deviation of the revisions represents only 13.8 percent of the variability of the full-sample index. These results show that the EW3 estimation procedure is relatively stable over time and that extending the information set does not materially alter the cyclical signal provided by the index.

Overall, the findings suggest that a parsimonious composite indicator combining the CEEMDAN-extracted cycles of credit, RPPI, and BCA provides a reasonable framework for assessing cyclical systemic risk in Moldova. The selected EW3 index can therefore serve as a supporting indicator for macroprudential surveillance and the calibration of instruments such as the CCyB. Nevertheless, it should not be interpreted as a mechanical policy rule, and its signals should be assessed alongside broader macro-financial evidence and expert judgement.

IV. Conclusions

This study develops a composite financial cycle index for the Republic of Moldova intended to support the assessment of the build-up and release of cyclical systemic risk. Measuring the financial cycle is particularly challenging in Moldova because the available time series are relatively short and persistent, the domestic financial market is comparatively shallow, and no sufficiently detailed official chronology of financial crises is available. The analysis therefore combines an adaptive cycle-extraction method with statistical robustness checks, historical evidence, and an assessment of the economic interpretation of the resulting indicators.

The cyclical components of the candidate variables are extracted using CEEMDAN, which does not require the prior specification of a fixed frequency band and can accommodate nonlinear and non-stationary dynamics. The empirical results indicate that the extracted cycles are broadly robust to alternative mirror-padding lengths. The quasi-real-time exercise also provides evidence of temporal stability: despite differences in the magnitude of revisions across variables, the recursive endpoint estimates generally preserve the direction, sign, and main dynamics of the corresponding full-sample cycles. These results support the potential use of the extraction procedure in an operational setting in which the indicator is updated as new observations become available.

The comparison between the raw and adjusted cycles supports the exclusion of the first intrinsic mode function, which primarily captures the highest-frequency oscillations. Removing this component produces smoother cyclical signals while preserving their main medium-term dynamics and maintaining satisfactory quasi-real-time properties. At the same time, the additional exclusion of the second intrinsic mode function does not provide a material improvement and, in the case of the current account balance, reduces the preservation of turning-point timing. The adjusted cycles excluding the first intrinsic mode function only are therefore retained as the most appropriate basis for aggregation.

The candidate composite indicators were assessed along three complementary dimensions: their cyclical duration relative to the real economic cycle, the timing of their peaks before the downturns of 2009 and 2015, and the economic plausibility of their component contributions. The duration analysis shows that the relationship between the financial and real economic cycles is sensitive to the turning-point specification. Under the sign-restricted specification, all candidate indicators display substantially longer average cycles than the GDP gap, whereas the unrestricted specification produces more heterogeneous results, with some candidate cycles longer and others shorter than the GDP gap. Cycle duration is therefore treated as a descriptive and supporting characteristic rather than as a standalone basis for selecting the preferred index. In contrast, the historical timing evidence is more consistent across candidate specifications: the composite indicators reach positive peaks approximately one to two quarters before the GDP gap moves into negative territory during the 2009 and 2015 episodes. The credit-to-GDP gap provides a relevant advance signal before the 2009 downturn but does not identify a comparable positive peak before the 2015 episode. This finding supports the use of a broader composite measure as a complement to the conventional credit-gap benchmark.

The comparison of the aggregation schemes does not reveal a clear advantage from using time-varying EWMA weights. Equal-weighted and EWMA-weighted specifications produce very similar dynamics, turning points, and economic interpretations, while the structure of the Moldovan economy does not provide an objective basis for assigning differentiated time-varying weights. Equal weighting is therefore preferred because it offers a more parsimonious and transparent representation of the

financial cycle.

The contribution analysis also indicates that including the real interest rate may create an important interpretative risk. In 2022, exceptionally high inflation produced a strongly negative ex-post real interest rate which, after sign inversion, generated a positive contribution to the financial cycle. This signal did not necessarily correspond to an easing of overall financing and credit conditions. When a composite indicator is used to inform the calibration of macroprudential instruments, such a distortion could overstate cyclical systemic risk and contribute to an inappropriate CCyB response. The real interest rate is therefore excluded from the preferred specification.

The final choice consequently lies between the two-variable index based on credit and residential property prices and the three-variable specification that additionally includes the current account balance. Although the quantitative evidence does not clearly distinguish between them, external imbalances represent a relevant dimension of cyclical systemic risk in a small open economy such as Moldova. On this basis, the preferred index is EW3, which combines equally weighted normalized cycles of real credit, real residential property prices, and the sign-adjusted current account balance. Its historical decomposition provides an economically coherent interpretation: RPPI and BCA were the main sources of the initial expansion preceding the 2009 economic crisis, credit played the central role during the 2012–2015 upswing, and the contraction around 2017 was supported by negative contributions from all three components. Toward the end of the sample, the index records a renewed and broad-based increase, with all three components contributing positively.

The quasi-real-time validation provides further support for the selected composite index. Its recursive estimates closely track the full-sample results, while the observed revisions remain limited and do not materially alter the identified cyclical signal. These findings confirm the stability of the index over time and support its use for the ongoing assessment of cyclical systemic risk.

Overall, the findings suggest that a parsimonious indicator combining credit, residential property prices, and the current account balance provides a reasonable framework for assessing cyclical systemic risk in the Republic of Moldova. The selected index can serve as a supporting indicator for macroprudential surveillance and for the calibration of instruments such as the CCyB. Its level, direction, and component contributions should be assessed alongside the credit-to-GDP gap, indicators of banking-sector resilience and broader macro-financial evidence.

The results should also be interpreted in light of the absence of a formal domestic crisis chronology, the limited number of historical downturn episodes, and data constraints that prevent the inclusion of a broader range of financial asset prices. Future updates will make it possible to assess the performance of the index over a longer sample, examine its signals during new cyclical episodes, and reconsider its composition as additional reliable data become available.

Bibliography

Banka Slovenije (2026). *Countercyclical Capital Buffer: 2nd Quarter of 2026*. Banka Slovenije. (Visited on 09/28/2026).

Bedrosian, Edward (1962). *A Product Theorem for Hilbert Transforms*. Research Memorandum RM-3439-PR. Santa Monica, CA: RAND Corporation. URL: https://www.rand.org/pubs/research_memoranda/RM3439.html.

Boashash, Boualem (1992). "Estimating and Interpreting the Instantaneous Frequency of a Signal. I. Fundamentals". In: *Proceedings of the IEEE* 80.4, pp. 520–538. DOI: 10.1109/5.135376.

Borio, Claudio (2014). "The Financial Cycle and Macroeconomics: What Have We Learnt?" In: *Journal of Banking & Finance* 45, pp. 182–198. DOI: 10.1016/j.jbankfin.2013.07.031.

Bry, Gerhard & Charlotte Boschan (1971). *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*. New York: National Bureau of Economic Research.

Burns, Arthur F. & Wesley C. Mitchell (1946). *Measuring Business Cycles*. New York: National Bureau of Economic Research.

Chadwick, Meltem Gülenay & Hüseyin Öztürk (Oct. 2018). *Measuring Financial Systemic Stress for Turkey: A Search for the Best Composite Indicator*. Working Paper 18/16. Central Bank of the Republic of Turkey.

Che, Natasha Xingyuan & Yoko Shinagawa (Jan. 2014). *Financial Soundness Indicators and the Characteristics of Financial Cycles*. IMF Working Paper 2014/014. International Monetary Fund. DOI: 10.5089/9781484386880.001.

Chen, Qiuhui, Norden E. Huang, Sherman Riemenschneider & Yuesheng Xu (2006). "A B-Spline Approach for Empirical Mode Decompositions". In: *Advances in Computational Mathematics* 24.1–4, pp. 171–195. DOI: 10.1007/s10444-004-7614-3.

Chen, Sally & Katsiaryna Svirydzenka (Apr. 2021). *Financial Cycles—Early Warning Indicators of Banking Crises?* IMF Working Paper 2021/116. International Monetary Fund. DOI: 10.5089/9781513582306.001.

Christiano, Lawrence J. & Terry J. Fitzgerald (2003). "The Band Pass Filter". In: *International Economic Review* 44.2, pp. 435–465. DOI: 10.1111/1468-2354.t01-1-00076.

Cohen, Leon (1995). *Time-Frequency Analysis*. Englewood Cliffs, NJ: Prentice Hall.

Czech National Bank (Sept. 15, 2023). *The CNB's Approach to Setting the Countercyclical Capital Buffer*. Methodological document. Czech National Bank. (Visited on 09/28/2026).

Drehmann, Mathias, Claudio E. V. Borio & Kostas Tsatsaronis (June 2012). *Characterising the Financial Cycle: Don't Lose Sight of the Medium Term!* BIS Working Paper 380. Bank for International Settlements. URL: https://www.bis.org/publ/work380.htm.

Elliott, Graham, Thomas J. Rothenberg & James H. Stock (1996). "Efficient Tests for an Autoregressive Unit Root". In: *Econometrica* 64.4, pp. 813–836. DOI: 10.2307/2171846.

European Systemic Risk Board (June 18, 2014). *Recommendation of the European Systemic Risk Board of 18 June 2014 on Guidance for Setting Countercyclical Buffer Rates (ESRB/2014/1).* Recommendation 2014/C 293/01. European Systemic Risk Board. URL: https://www.esrb.europa.eu/pub/pdf/recommendations/140630_ESRB_Recommendation.en.pdf.

Filatov, Vladyslav (2021). "A New Financial Stress Index for Ukraine". In: *Visnyk of the National Bank of Ukraine* 251, pp. 37–54. DOI: 10.26531/vnbu2021.251.03.

Gabor, Dennis (1946). "Theory of Communication. Part 1: The Analysis of Information". In: *Journal of the Institution of Electrical Engineers—Part III: Radio and Communication Engineering* 93.26, pp. 429–441. DOI: 10.1049/ji-3-2.1946.0074.

Gadanecz, Blaise & Kaushik Jayaram (2015). "Macroprudential Policy Frameworks, Instruments and Indicators: A Review". In: IFC Workshop on Combining Micro and Macro Statistical Data for Financial Stability Analysis: Experiences, Opportunities and Challenges (Warsaw, Dec. 14–15, 2015). Bank for International Settlements. URL: https://www.bis.org/ifc/publ/ifcb41c_rh.pdf.

Harding, Don & Adrian Pagan (2002). "Dissecting the Cycle: A Methodological Investigation". In: *Journal of Monetary Economics* 49.2, pp. 365–381. DOI: 10.1016/S0304-3932(01)00108-8.

Harun, Cicilia A., Aditya Anta Taruna, R. Renanda Nattan & Ndari Suryaningsih (Dec. 2014). *Financial Cycle of Indonesia—Potential Forward Looking Analysis.* Working Paper WP/9/2014. Bank Indonesia.

He, Yaqing & Kim Fung Tsang (2021). "Universities Power Energy Management: A Novel Hybrid Model Based on ICEEMDAN and Bayesian Optimized LSTM". In: *Energy Reports* 7, pp. 6473–6488. DOI: 10.1016/j.egyr.2021.09.115.

Hiebert, Paul, Benjamin Klaus, Tuomas A. Peltonen, Yves S. Schüler & Peter Welz (2014). "Capturing the Financial Cycle in Euro Area Countries". In: *Financial Stability Review* 2, pp. 109–117. URL: https://www.ecb.europa.eu/pub/pdf/fsr/art/ecb.fsrart201411_02.en.pdf.

Holló, Dániel, Manfred Kremer & Marco Lo Duca (Mar. 2012). *CISS—A Composite Indicator of Systemic Stress in the Financial System.* Working Paper Series 1426. Frankfurt am Main: European Central Bank. URL: https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1426.pdf.

Huang, Norden E. (2005). "Introduction to the Hilbert–Huang Transform and Its Related Mathematical Problems". In: *Hilbert–Huang Transform and Its Applications.* Ed. by Norden E. Huang & Samuel S. P. Shen. World Scientific, pp. 1–26. DOI: 10.1142/9789812703347_0001.

Huang, Norden E., Zheng Shen, Steven R. Long, Man-Li C. Wu, Hsing H. Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung & Henry H. Liu (1998). "The Empirical Mode Decomposition and the Hilbert Spectrum for Nonlinear and Non-Stationary Time Series Analysis". In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454.1971, pp. 903–995. DOI: 10.1098/rspa.1998.0193.

Huang, Norden E., Man-Li Wu, Wendong Qu, Steven R. Long & Samuel S. P. Shen (2003a). "Applications of Hilbert–Huang Transform to Non-Stationary Financial Time Series Analysis". In: *Applied Stochastic Models in Business and Industry* 19.3, pp. 245–268. DOI: 10.1002/asmb.501.

Huang, Norden E., Man-Li C. Wu, Steven R. Long, Samuel S. P. Shen, Wendong Qu, Per Gloersen & Kuang L. Fan (2003b). "A Confidence Limit for the Empirical Mode Decomposition and Hilbert Spectral Analysis". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 459.2037, pp. 2317–2345. DOI: 10.1098/rspa.2003.1123.

International Monetary Fund (2026). *Primary Commodity Prices*. URL: https://www.imf.org/en/research/commodity-prices (visited on 07/07/2026).

Kupkovič, Patrik & Martin Šuster (2020). *Identifying the Financial Cycle in Slovakia*. NBS Working Paper 2/2020. Bratislava: National Bank of Slovakia. URL: https://nbs.sk/_img/documents/publik/wp_2_2020_kupkovic_suster_identifying_the_financial_cycle_in_slovakia_en.pdf.

Kwiatkowski, Denis, Peter C. B. Phillips, Peter Schmidt & Yongcheol Shin (1992). "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root: How Sure Are We That Economic Time Series Have a Unit Root?" In: *Journal of Econometrics* 54.1–3, pp. 159–178. DOI: 10.1016/0304-4076(92)90104-Y.

Lang, Jan Hannes, Cosimo Izzo, Stephan Fahr & Josef Ruzicka (Feb. 2019). *Anticipating the Bust: A New Cyclical Systemic Risk Indicator to Assess the Likelihood and Severity of Financial Crises*. Tech. rep. 219. European Central Bank. DOI: 10.2866/468656.

Lee-Poy, Andrew (Oct. 2018). *Characterizing the Canadian Financial Cycle with Frequency Filtering Approaches*. Staff Analytical Note 2018-34. Bank of Canada. URL: https://www.bankofcanada.ca/2018/10/staff-analytical-note-2018-34/.

Li, Taiyong, Yingrui Zhou, Xinsheng Li, Jiang Wu & Ting He (2019). "Forecasting Daily Crude Oil Prices Using Improved CEEMDAN and Ridge Regression-Based Predictors". In: *Energies* 12.19, p. 3603. DOI: 10.3390/en12193603.

Meeson Reginald N., Jr. (2014). "HHT Sifting and Filtering". In: *Hilbert–Huang Transform and Its Applications*. Ed. by Norden E. Huang & Samuel S. P. Shen. 2nd ed. World Scientific, pp. 117–147. DOI: 10.1142/9789814508247_0006.

Morana, Claudio (2024). "A New Macro-Financial Condition Index for the Euro Area". In: *Econometrics and Statistics* 29, pp. 64–87. DOI: 10.1016/j.ecosta.2021.09.005.

Mugrabi, Farah & Gerhard Rünstler (2025). *Housing and Credit Cycles in Ireland*. Tech. rep. 16. Central Bank of Ireland.

National Bank of Moldova (2015). *Raport anual 2014 [Annual Report 2014]*. Chișinău: National Bank of Moldova. URL: https://www.bnm.md/files/RA_main_ro.pdf (visited on 07/21/2026).

— (2021). *Financial Stability Report for 2020*. Chișinău: National Bank of Moldova. URL: https://www.bnm.md/files/RSF-2020_EN.pdf (visited on 07/21/2026).

— (2022). *Financial Stability Report for 2021*. Chișinău: National Bank of Moldova. URL: https://bnm.md/files/RSF-2021_ENG.pdf (visited on 07/21/2026).

National Bank of Moldova (2023). *Financial Stability Report for 2022*. Chișinău: National Bank of Moldova. URL: https://bnm.md/files/RSF_2022_EN_v1.pdf (visited on 07/21/2026).

— (2026). *National Bank of Moldova Base Rate*. URL: https://bnm.md/bdi/pages/reports/dpmc/DPMC2.xhtml (visited on 07/21/2026).

National Bureau of Statistics of the Republic of Moldova (2026a). *Consumer Price Indices*. URL: https://statistica.gov.md/ro/statistic_indicator_details/10 (visited on 07/21/2026).

— (2026b). *National Accounts*. URL: https://statistica.gov.md/ro/statistic_indicator_details/12 (visited on 07/21/2026).

— (2026c). *Unit Value Indices*. URL: https://statistica.gov.md/ro/statistic_indicator_details/19 (visited on 07/21/2026).

Ng, Tim (June 2011). "The Predictive Content of Financial Cycle Measures for Output Fluctuations". In: *BIS Quarterly Review*, pp. 53–65. URL: https://www.bis.org/publ/qtrpdf/r_qt1106g.htm.

Russo, Riccardo (Nov. 2022). *An Indicator of the Financial Cycle in France since the Beginning of the 2000s*. Trésor-Economics 317. Direction générale du Trésor. URL: https://www.tresor.economie.gouv.fr/Articles/2022/11/29/an-indicator-of-the-financial-cycle-in-france-since-the-beginning-of-the-2000s.

Schüler, Yves S., Paul Hiebert & Tuomas A. Peltonen (Sept. 2015). *Characterising the Financial Cycle: A Multivariate and Time-Varying Approach*. Working Paper Series 1846. Frankfurt am Main: European Central Bank. URL: https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1846.en.pdf.

Škare, Marinko, Luis A. Gil-Alana & Małgorzata Porada-Rochoń (2025). "Isolating financial cycles using the fractional cyclical model in selected economies: 1970–2019". In: *Structural Change and Economic Dynamics* 72, pp. 67–77. DOI: 10.1016/j.strueco.2024.10.001.

Stremmel, Hanno (June 2015). *Capturing the Financial Cycle in Europe*. Working Paper Series 1811. Frankfurt am Main: European Central Bank. URL: https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1811.en.pdf.

Torres, María E., Marcelo A. Colominas, Gastón Schlotthauer & Patrick Flandrin (2011). "A Complete Ensemble Empirical Mode Decomposition with Adaptive Noise". In: *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 4144–4147. DOI: 10.1109/ICASSP.2011.5947265.

Vella, Sarah (2024). *Constructing a Cyclical Systemic Risk Indicator (cSRI) for Malta*. Tech. rep. WP/01/2024. Central Bank of Malta. URL: https://www.centralbankmalta.org/site/Publications/Economic%20Research/2024/WP-01-2024.pdf (visited on 09/29/2026).

Wu, Zhaohua & Norden E. Huang (2009). "Ensemble Empirical Mode Decomposition: A Noise-Assisted Data Analysis Method". In: *Advances in Adaptive Data Analysis* 1.1, pp. 1–41. DOI: 10.1142/S1793536909000047.

Yan, Shu-Ping, Li Ji, Ping Zhang, Si-Ming Liu, Lei Lu & Min Long (2024). "A Promising Method for Breaking the Logjam of Time-Frequency Analysis in Astronomy". In: *RAS Techniques and Instruments* 3.1, pp. 56–65. DOI: 10.1093/rasti/rzae001.

Appendix 1. EMD and CEEMDAN methodology

Empirical Mode Decomposition

Empirical Mode Decomposition is an adaptive data-analysis method designed to decompose nonlinear and non-stationary signals into a finite set of Intrinsic Mode Functions and a residual component. The extracted IMFs form an approximately orthogonal, data-driven representation of the signal. Each IMF represents an oscillatory mode identified from the local properties of the data, although its substantive interpretation must be evaluated separately. Together with the residual, the extracted IMFs reconstruct the original signal up to numerical precision.

Unlike conventional Fourier-based representations, which express a signal through globally defined sinusoidal components with fixed frequencies, Empirical Mode Decomposition identifies oscillatory modes directly from the local extrema of the observed series. This distinction is important for non-linear and non-stationary data, for which the amplitude and frequency of oscillations may vary over time and a fixed harmonic representation may generate misleading distortions (Huang et al., 1998). EMD therefore does not require a predetermined functional form or frequency interval and allows the characteristic time scale and amplitude of each extracted component to evolve throughout the sample.

The extracted components are referred to as Intrinsic Mode Functions. Their definition is motivated by the need to obtain locally well-behaved oscillatory components for which instantaneous amplitude and frequency can be meaningfully interpreted through the Hilbert-transform framework (Gabor, 1946; Bedrosian, 1962; Cohen, 1995; Boashash, 1992; Huang et al., 1998). Huang et al. (1998) define an IMF through the following two conditions:

(1) in the whole data set, the number of extrema and the number of zero crossings must either equal or differ at most by one

(2) at any point, the mean value of the envelope defined by the local maxima and the envelope defined by the local minima should be zero (∗)

Intuitively, an IMF can be seen as a counterpart to the simple harmonic function – a simple oscillatory movement respecting a set of predefined conditions. In contrast to simple harmonic functions, IMFs are derived directly from the local properties of the data and represent data-driven oscillatory modes that may exhibit non-stationary behavior. Any complicated signal can be explained by a superposition of Intrinsic Mode Functions and a Residual. For a more detailed understanding of the conditions and how they relate to the definition of a well-defined instantaneous frequency, the reader is advised to read Huang et al. (1998), Huang (2005) and Meeson (2014).

The decomposition of data into IMFs happens via the Sifting Algorithm. Huang et al. (1998) propose the following steps to detect the first IMF of the data:

(1) Identify the local minima and maxima

(2) Connect the local maxima by a cubic spline to create the upper envelope

(3) Repeat the procedure for the local minima

(4) Find the envelope mean (μ_1)

(5) Remove the envelope mean from the data to find the Proto-IMF : $\zeta_1 = x(t) - \mu_1$

After removing the mean μ_1, the resulting proto-IMF should ideally satisfy the IMF conditions described above. In practice, this is rarely achieved after a single iteration, and the sifting procedure must therefore be repeated until a convergence criterion is met. Different convergence criteria have been proposed in the literature. One such criterion is the S-stoppage criterion discussed by Huang et al. (2003b), under which a fixed integer S is selected in advance. The sifting process stops when the numbers of zero crossings and extrema differ by at most one, while both numbers remain unchanged for S consecutive sifting iterations. Huang et al. (2003b) identify values of S between 4 and 8 as an appropriate range. To retain a data-driven selection within this range, the Index of Orthogonality is subsequently used to assess the quality of the resulting decomposition.

By construction, the first IMF generally captures the highest-frequency oscillations present in the signal. If the signal contains longer-period oscillations then the sifting process will push these longer-period oscillations into the residual such that:

$$r_1 = x(t) - \zeta_1 \tag{1}$$

After the first IMF is removed, the resulting residual (r_1) becomes the input signal from which the next IMF is extracted. By applying the sifting algorithm on the residual r_1 we can find the oscillatory modes that were omitted in ζ_1. The second IMF (ζ_2) will generally capture lower-frequency oscillations than ζ_1 since they were already removed from the original data. The extraction stops when the residual contains fewer than two extrema or when the maximum admissible number of modes is reached. The final residual is interpreted as the trend component of the data, while the extracted IMFs represent oscillatory components at progressively longer time scales. Once the algorithm has stopped extracting and sifting IMFs, the final set of functions converges to the form:

$$x(t) = \sum_{i=1}^{N_{\text{IMF}}} \zeta_i + r_{N_{\text{IMF}}} \tag{2}$$

The way in which the extracted IMFs are combined into the final cyclical component is described in the methodology section, where the rules for trend removal and denoising are introduced.

An important limitation of Empirical Mode Decomposition is mode mixing. Mode mixing occurs when a single IMF contains oscillations of clearly different periods or when oscillations of the same periods are split across multiple IMFs. The decomposition thus fails to clearly separate components, and this is precisely the issue that Ensemble EMD is designed to solve.

Complete Ensemble Empirical Mode Decomposition with Adaptive Noise — CEEMDAN

The Ensemble Empirical Mode Decomposition (EEMD) was introduced by Wu and Huang (2009) to address the instability and mode–mixing present in the original EMD. One reason for mode mixing in EMD is that some local extrema are unstable; that is, after a small perturbation, an extremum may no longer remain an extremum. The idea behind EEMD is to stabilize the decomposition by averaging a set of IMFs extracted from multiple noise-perturbed realizations of the data. The EEMD algorithm has the following steps:

(1) Add white-noise of a certain relative intensity (ε) to the original signal:

$$w^{(j)}(t) \sim \mathcal{N}(0,1), \qquad x^{(j)}(t) = x(t) + \beta_0 w^{(j)}(t), \qquad \beta_0 = \varepsilon\,\sigma(x(t))$$

(2) Perform Empirical Mode Decomposition on the new signal $x^{(j)}(t)$ and decompose the series into a complete set of IMFs

(3) Repeat steps (1) — (2) with a newly generated white noise series until ensemble size N_{E} is reached

The final set of IMFs is obtained by averaging each obtained IMF over the ensemble size:

$$\zeta_i = \frac{1}{N_{\mathrm{E}}} \sum_{j=1}^{N_{\mathrm{E}}} \zeta_i^{(j)} \tag{3}$$

The Ensemble Empirical Mode Decomposition improves upon EMD by adding different variations of white noise of a certain intensity to the signal and averaging IMFs over many realizations. The addition of differently generated white-noise realizations helps populate the time-frequency space more uniformly and reduces the sensitivity of the decomposition to intermittent extrema. Averaging across the ensemble is designed to attenuate the added noise and reduce mode mixing. The main drawback of EEMD is that by averaging the IMF components, the final set is not complete and cannot reconstruct the original signal exactly. To fix this issue, Torres et al. (2011) proposed a different version of EEMD, where adaptive noise is selected at each decomposition step so that the final set reconstructs the original signal.

For compactness, define the operator $\mathcal{D}_j\{\cdot\}$ which returns the j–th mode obtained by EMD from given input signal. The Complete Ensemble Empirical Mode Decomposition with Adaptive Noise algorithm can then be described by the following steps:

(1) Add white-noise of a certain relative intensity to the original signal:

$$w^{(j)}(t) \sim \mathcal{N}(0,1), \qquad x^{(j)}(t) = x(t) + \beta_0 w^{(j)}(t), \qquad \beta_0 = \varepsilon\,\sigma(x(t))$$

(2) Decompose and average via EEMD to find the first mode:

$$\zeta_1 = \frac{1}{N_{\mathrm{E}}} \sum_{j=1}^{N_{\mathrm{E}}} \zeta_1^{(j)}$$

(3) Calculate the first residual as:

$$r_1 = x(t) - \zeta_1$$

(4) Generate new ensemble data of the form:

$$r_1 + \beta_1 \mathcal{D}_1\left\{w^{(j)}(t)\right\}, \qquad \beta_1 = \varepsilon\,\sigma(r_1(t)).$$

Note that the adaptive noise is defined as the first IMF after decomposing standard white noise via EMD.

(5) Find the second IMF mode as:

$$\zeta_2 = \frac{1}{N_{\mathrm{E}}} \sum_{j=1}^{N_{\mathrm{E}}} \mathcal{D}_1\left\{r_1 + \beta_1 \mathcal{D}_1\left\{w^{(j)}(t)\right\}\right\}$$

Note that the second IMF is actually the first IMF of the data generated in step (4).

(6) Calculate remaining k residuals as:

$$r_k = r_{k-1} - \zeta_k$$

(7) Generate new ensemble data of the form:

$$r_k + \beta_k \mathcal{D}_k \left\{ w^{(j)}(t) \right\}, \qquad \beta_k = \varepsilon \, \sigma(r_k(t)),$$

and calculate the $(k+1)$-th IMF as:

$$\zeta_{k+1} = \frac{1}{N_{\mathrm{E}}} \sum_{j=1}^{N_{\mathrm{E}}} \mathcal{D}_1 \left\{ r_k + \beta_k \mathcal{D}_k \left\{ w^{(j)}(t) \right\} \right\}$$

(8) Repeat until the residual signal r_{k+1} is monotonic or has at most one extremum.

This noise-adaptive construction is designed so that the perturbation introduced at each decomposition level is related to the corresponding scale of the noise decomposition. By linking the added noise across successive stages, CEEMDAN is intended to reduce mode mixing and improve the completeness and stability of the resulting decomposition (Torres et al., 2011). The main parameters specific to the noise-assisted ensemble procedure are the noise intensity and the ensemble size, while the EMD sifting step additionally depends on the selected S-number.

Hyperparameter tuning

The parameters considered in this study have different roles in the decomposition procedure. The S-number controls the stopping rule of the EMD sifting process and is therefore relevant for EMD-based decompositions, including EEMD and CEEMDAN. The noise intensity ε and the ensemble size N_E are specific to the noise-assisted ensemble procedures. In this study, the ensemble size is fixed, while the S-number and the noise intensity are selected through a grid search method.

The ensemble size (N_{E}), defined as the number of noise-assisted realizations included in the decomposition, was fixed at 1,000. Wu and Huang (2009) show that, provided that the added noise has a moderate amplitude and the ensemble contains a sufficiently large number of realizations, further increases in the noise amplitude and ensemble size produce only limited changes in the resulting decomposition. They also indicate that a larger ensemble is required when the amplitude of the added noise increases.

Evidence from subsequent empirical applications provides additional support for selecting a relatively large ensemble. In an ICEEMDAN-based application to daily crude oil prices, Li et al. (2019) examine ensemble sizes ranging from 20 to 2,000. Their results improve as the number of realizations increases from 20 to approximately 500 and remain broadly stable over the range from 500 to 2,000. In a direct application of CEEMDAN to complex non-stationary signals, Yan et al. (2024) employ an ensemble size of 1,000 and additionally assess decomposition accuracy over a grid comprising 50, 100, 200, 500, 1,000 and 2,000 realizations. The reported relative errors display only limited variation at the larger ensemble sizes. Similarly, He and Tsang (2021) use 1,000 realizations when comparing EEMD, CEEMDAN and improved CEEMDAN within a hybrid model for short-term electricity-consumption forecasting.

Against this background, an ensemble size of 1,000 was adopted as a conservative value located within the stability range documented in previous applications. The relatively large ensemble reduces the dependence of the extracted components on individual noise realizations and supports a stable comparison of alternative hyperparameter combinations through the Index of Orthogonality, while remaining computationally feasible for the quarterly series analysed. However, because no separate sensitivity exercise comparing alternative ensemble sizes is performed for the present dataset, $N_{\mathrm{E}} = 1{,}000$ is interpreted as a literature-informed practical calibration rather than as an empirically optimized value.

A fixed random seed and a single computational thread are used throughout the analysis to ensure the reproducibility of the extracted IMFs.

The two parameters selected through the grid search are the S-Number (S) and the Noise Intensity (ε). The noise intensity ε is specified as a relative scaling parameter. At the first CEEMDAN stage, the noise-scaling coefficient is defined as $\beta_0 = \varepsilon\,\sigma(x(t))$, where $\sigma(x(t))$ denotes the standard deviation of the input signal. At each subsequent stage, the corresponding coefficient is defined relative to the current residual as $\beta_k = \varepsilon\,\sigma(r_k(t))$. Thus, if $\varepsilon = 0.20$, the scaling coefficient applied to the noise component at each stage corresponds to 20% of the standard deviation of the signal or residual entering the respective decomposition stage. The parameter ε therefore represents a relative noise-intensity parameter rather than an absolute noise amplitude. The grid ranges are centered on values commonly employed in the EMD and EEMD literature. Following Huang et al. (2003b) and Wu and Huang (2009), the candidate values considered are $S \in \{4, 5, 6, 7, 8\}$ and $\varepsilon \in \{0.20, 0.25, 0.30, 0.35, 0.40\}$.

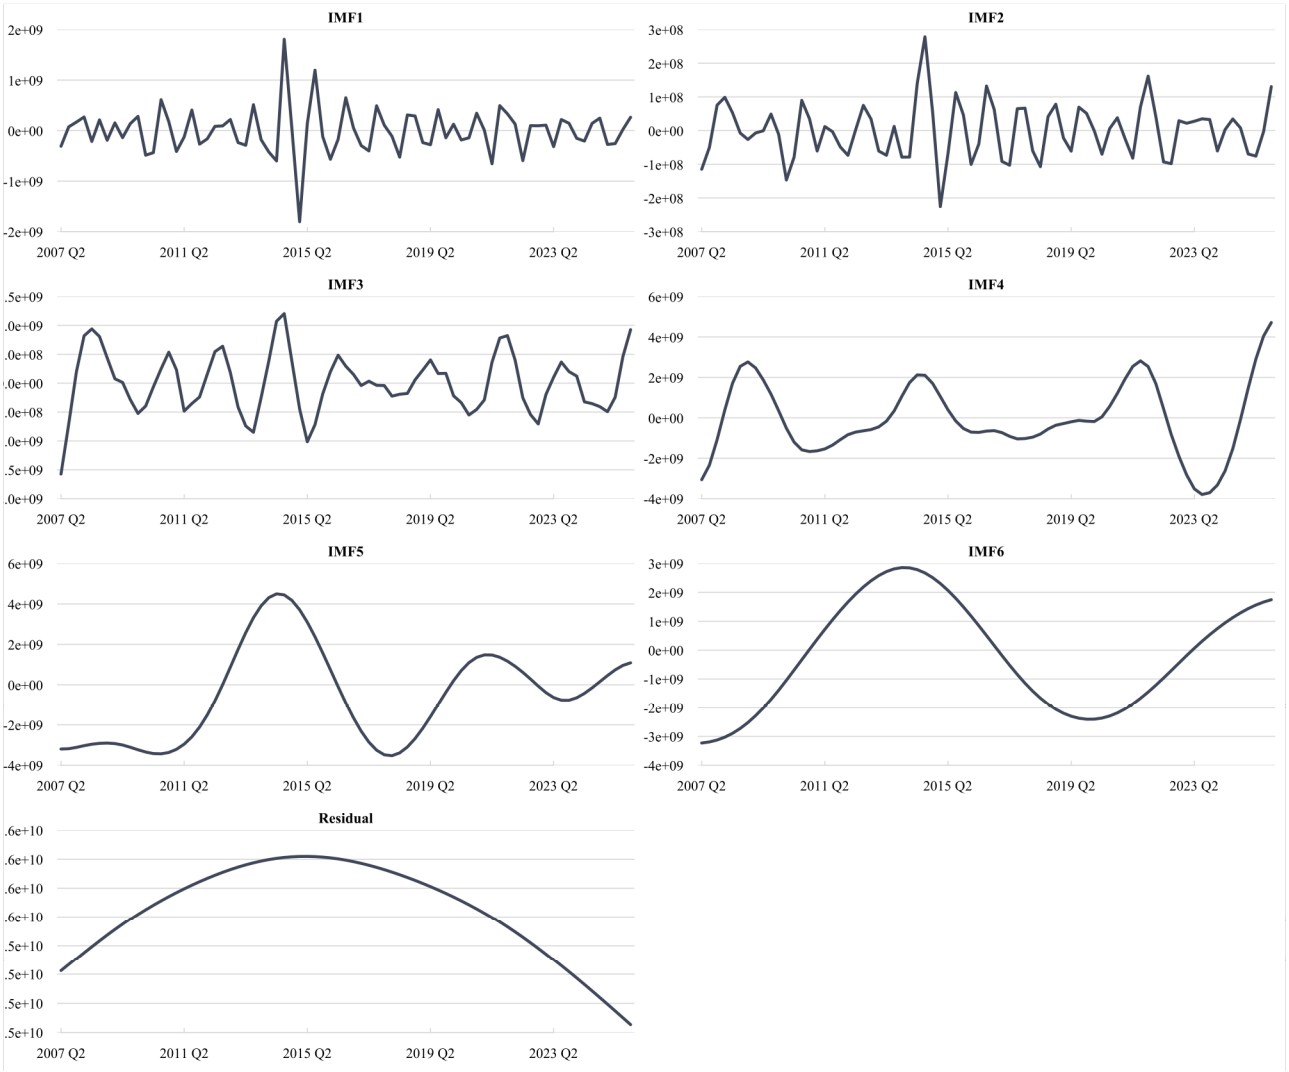
A data-driven choice of the parameters is obtained by defining a grid over the S-number and the noise intensity, and selecting the combination that minimizes the Index of Orthogonality. The performance measure used for this purpose is the Index of Orthogonality (IO) as described in Q. Chen et al. (2006):

$$\mathrm{IO} = \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^{N_{\mathrm{IMF}}+1} \frac{\left|\sum_t c_i(t) c_j(t)\right|}{\sum_t x^2(t)} \tag{4}$$

The Index of Orthogonality measures the degree of leakage between the components of the complete decomposition, including both the extracted IMFs and the residual. Ideally, if the decomposition components are mutually orthogonal, the inner products between distinct pairs of components in Equation (4) should be close to zero. In practice, some leakage may occur, but it is generally expected to remain limited (Huang et al., 1998). The Index of Orthogonality therefore provides a data-driven criterion for selecting the CEEMDAN hyperparameters separately for each variable, conditional on a given padding specification. It is not used to compare or select alternative padding lengths, since each padding specification changes the artificial extension entering the decomposition and, consequently, the basis on which the index is calculated. The mirror-padding length is instead selected through a separate robustness assessment based on the stability of the extracted cycles over the original, non-extended sample.

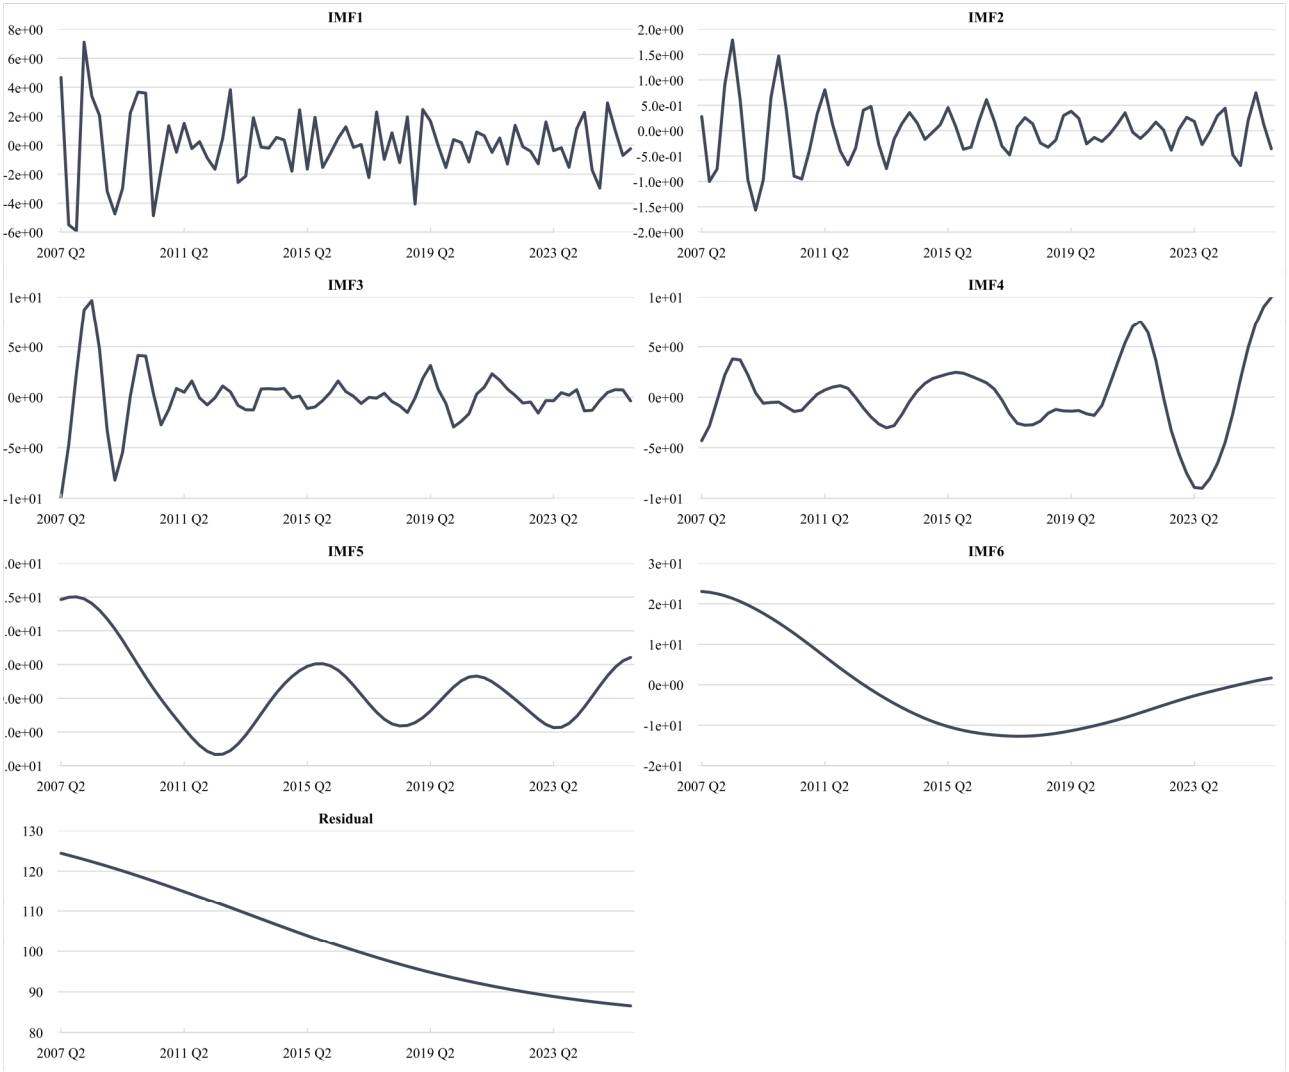
Appendix 2. Decomposition results

Figure 2.1: Credit, MDL at 2015 prices



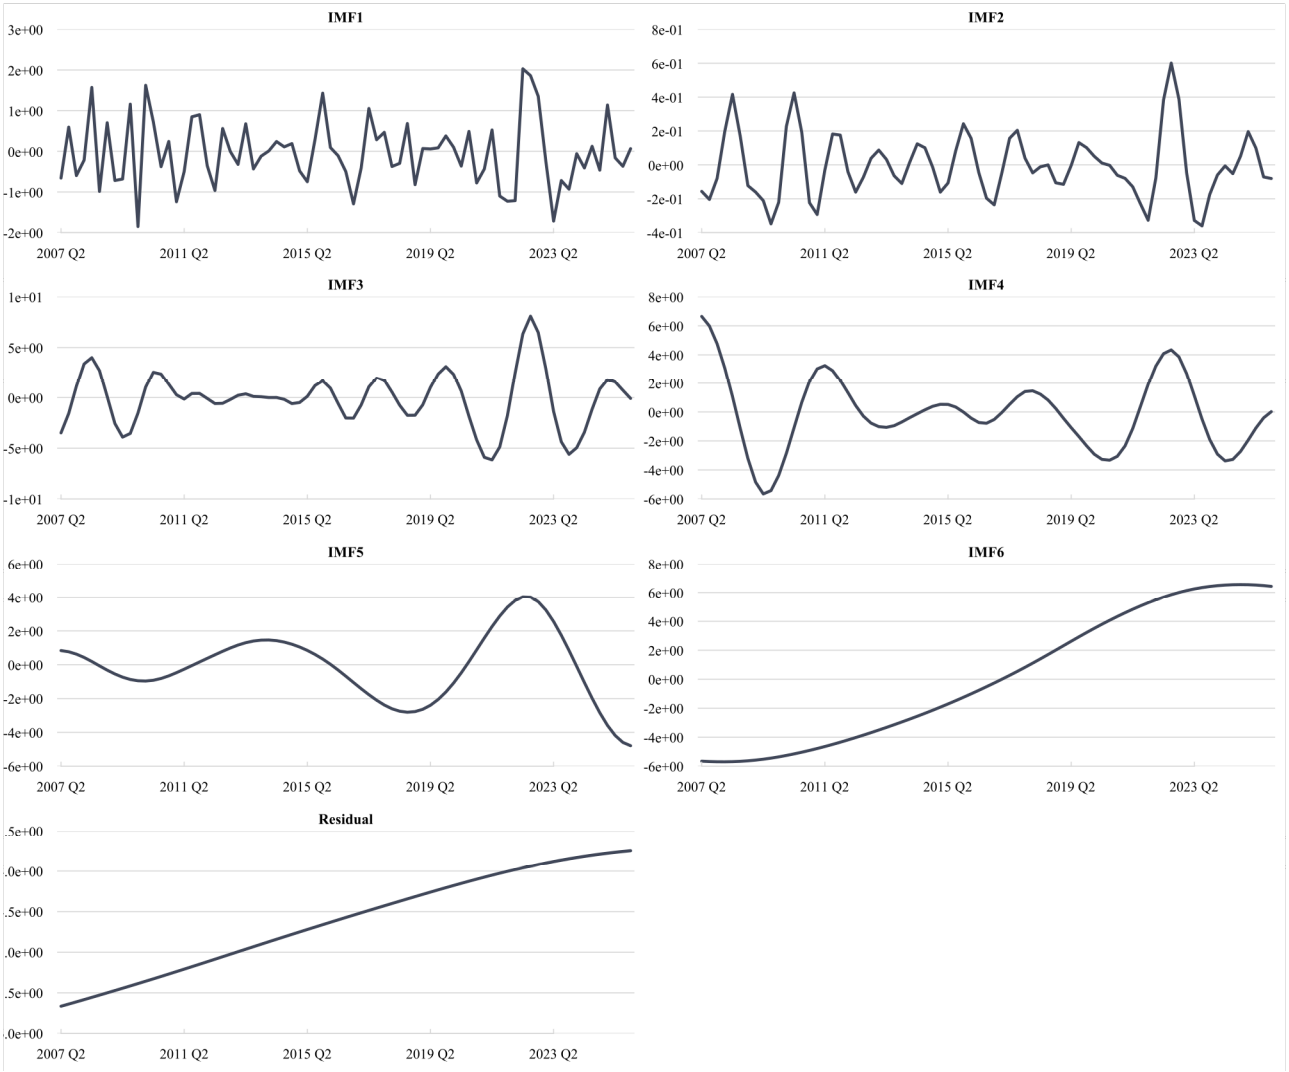
Source: authors' calculations.

Figure 2.2: RPPI, 2015=100



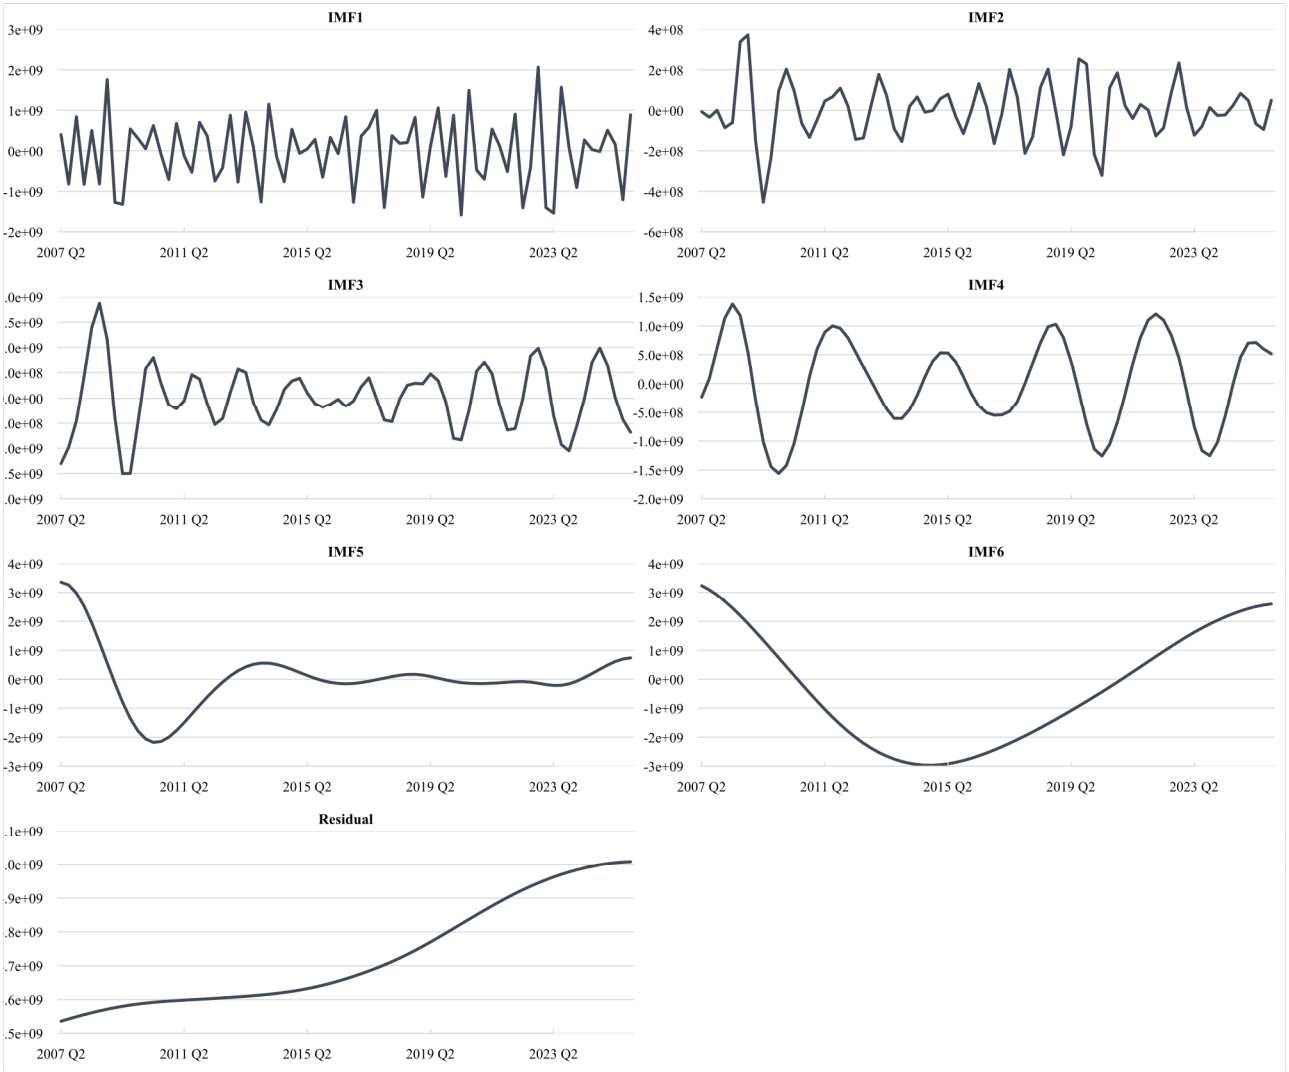
Source: authors' calculations.

Figure 2.3: Real interest rate, percent



Source: authors' calculations.

Figure 2.4: Current account balance, MDL at 2015 prices



Source: authors' calculations.