

South African Reserve Bank

Working Paper Series

WP/26/04

The credit gap as a real-time early warning indicator of financial imbalances: a South African perspective

Wongani Msiska, Ayanda Sikhosana and Cobus Vermeulen

Authorised for publication by Konstantin Makrelov

5 February 2026



SOUTH AFRICAN RESERVE BANK

© South African Reserve Bank

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without fully acknowledging the author(s) and this Working Paper as the source.

South African Reserve Bank Working Papers are written by staff members of the South African Reserve Bank and, on occasion, by consultants under the auspices of the South African Reserve Bank. The papers deal with topical issues and describe preliminary research findings and develop new analytical or empirical approaches in their analyses. They are solely intended to elicit comments and stimulate debate.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the South African Reserve Bank or South African Reserve Bank policy. While every precaution is taken to ensure the accuracy of information, the South African Reserve Bank shall not be liable to any person for inaccurate information, omissions or opinions contained herein.

South African Reserve Bank Working Papers are externally refereed.

Information on South African Reserve Bank Working Papers can be found at <https://www.resbank.co.za/en/home/publications/Papers/working-papers>.

Enquiries relating to the Working Paper Series can be addressed to:

Head: Economic Research Department
South African Reserve Bank
P O Box 427
Pretoria 0001

Tel. +27 12 313 3911

The credit gap as a real-time early warning indicator of financial imbalances: a South African perspective

Wongani Msiska,^{*} Ayanda Sikhosana[†] and Cobus Vermeulen[‡]

Abstract

Excessive credit growth can create systemic risks in the financial sector, which can amplify business cycle fluctuations. Financial crises can result from excessive risk-taking during prolonged periods of growth when accompanied by weak regulation. Financial institutions tend to lend generously during economic booms and then restrict lending during downturns, which can worsen economic recessions. To address this, the Basel Committee on Banking Supervision suggests using the credit-to-GDP gap as an early warning indicator to identify rising credit risks and systemic instability. When this gap exceeds certain thresholds, the countercyclical capital buffer (CCyB) can be deployed to compel banks to hold larger capital buffers, thus strengthening the financial system against potential downturns.

This paper explores different ways to identify the South African credit gap in real time. Using a constructed real-time data set, it evaluates a range of filtering techniques to isolate the credit gap from the credit-to-GDP ratio, and identify the real-time credit gap most consistent with the estimated ex post full sample credit gap. The one-sided Hodrick-Prescott filter produces a reliable real-time estimate of the credit-to-GDP gap, accurately predicting 75% of financial cycle turning points since 1980, and aligning well with full-sample estimates of the credit-to-GDP gap. It is therefore suitable to use as a guide for activating the CCyB.

JEL classification

Basel III, financial sector procyclicality, financial cycle, countercyclical capital buffer, credit-to-GDP gap, Hodrick-Prescott filter

Keywords

C22, E32, E37, E44, G01, G28

* Department of Economics, University of Stellenbosch.

† Financial Stability Department, South African Reserve Bank (SARB).

‡ Economic Research Department, SARB. Corresponding author.

Email: cobus.vermeulen@resbank.co.za

1. Introduction

The 2008 global financial crisis (GFC) highlighted the need for stronger financial systems to detect vulnerabilities before they develop into systemic crises. As part of the Basel III framework, the Basel Committee on Banking Supervision (BCBS) introduced the countercyclical capital buffer (CCyB) to protect the banking sector from excessive credit growth and strengthen its resilience during periods of financial distress. The CCyB ensures that banks build additional capital during economic booms, which can be released during downturns to support lending and absorb losses, preventing financial distress from turning into severe crises.

The BCBS recommends using the credit-to-GDP gap, measured as the deviation of the credit-to-GDP ratio from its long-run trend, as an early warning indicator (EWI) of the build-up of systemic risk. When this gap exceeds certain thresholds, the CCyB can be deployed to increase capital buffers, enabling banks to absorb losses and support credit supply during downturns.

This paper aims to identify the best method for real-time identification of the credit-to-GDP gap in South Africa, which can serve as an EWI for periods of financial distress, proxied here as financial cycle downswings. We specifically aim to identify a credit gap metric that gives a timely warning for upper turning points (peaks) in the South African financial cycle. We compare different filtering approaches to determine which one provides the most accurate and reliable identification of the real-time credit gap for the domestic context. This is crucial for informing the South African Reserve Bank's (SARB's) decisions on the CCyB, helping to mitigate systemic risks and promoting a more resilient domestic financial system.

The paper is structured as follows. Section 2 provides a brief background on the BCBS's guidance and recommendations for estimating the credit-to-GDP gap and how it informs the CCyB. Data and methodology are described in section 3. The results are presented and discussed in section 4, where we pay particular attention to evaluating different real-time credit gaps against an ex post benchmark credit gap, and also explore the implications for activating the CCyB based on different threshold calibrations. Section 5 concludes.

2. Design of the CCyB

2.1 The BCBS motivation

The BCBS's motivation for designing a CCyB scheme is twofold: first, it should limit systemic risk by strengthening the banking system's resilience against shocks, and second, it should prevent the banking system from amplifying economic fluctuations (Drehmann et al. 2010).

To identify the credit-to-GDP gap, the BCBS defines credit as the total nominal credit extended to the domestic private non-financial sector. This broad measure captures debt from bank and non-bank institutions, and domestic and foreign sources, and includes instruments such as bonds and cross-border finance (BCBS 2010). Credit gaps based on total credit, rather than narrower measures such as bank credit alone, are argued to be more effective EWIs of banking crises (Drehmann and Tsatsaronis 2014). The total credit measure also ensures that systemic risk build-up in all parts of the economy is accounted for, and it reduces the incentive for banks to divert credit supply to other parts of the financial system as could be the case when bank credit alone is used as the credit measure.

The credit-to-GDP ratio is then derived by expressing nominal credit extended as a fraction of nominal gross domestic product (GDP); that is, $r_t = \text{credit}_t / \text{GDP}_t$. The credit-to-GDP gap is simply the cyclical component of the credit-to-GDP ratio r_t , and is the difference between the long-term trend component of this ratio and its actual value.

To extract the trend, the BCBS recommends using a one-sided Hodrick-Prescott (HP) filter. The use of a one-sided HP filter to identify the credit gap is encouraged because decisions regarding the CCyB must be made in real time based on data available up to the point of the decision. A conventional two-sided HP filter would require future data that are not available at the time of the CCyB decision, making it unsuitable for real-time policymaking.¹

¹ The question of data requirements and real-time data availability – and their implications for the choice of filter – is central to this study and considered in depth in section 3.

Although the credit-to-GDP gap is the BCBS's recommended indicator, the BCBS advises against a purely mechanical application of the gap for determining the activation of the CCyB. Even though it outperforms other measures as an EWI, the credit-to-GDP gap has some limitations² (Drehmann, Borio and Tsatsaronis 2011). Recognising that no single leading indicator offers a perfect signal for crises, the BCBS (2010) emphasises the importance of using guided discretion and additional macroeconomic and bank performance indicators to complement the credit gap measure. Ideally, the macroprudential authority will have a suite of tools and indicators available to guide policy decisions. The focus of this paper, however, remains on the credit gap as but one metric in this holistic suite of tools.

2.2 Implementing the CCyB

Having a core set of indicators allows policymakers to comprehensively assess the risk environment to determine when cyclical risks are low, beginning to rise, and become elevated. The BCBS (2010) suggests that the CCyB should be activated when the credit-to-GDP gap continually exceeds the lower threshold of 2 percentage points. In this case, banks will be required to hold additional capital – the CCyB – ranging from 0% to 2.5% of their risk-weighted assets, depending on the size of the credit gap and assuming a 0% neutral rate. The CCyB is therefore zero when the gap is below this lower threshold ($L = 2$) and reaches its maximum level of 2.5% when the gap exceeds an upper threshold ($H = 10$). Between these thresholds, the CCyB increases linearly.

The choice of thresholds and the adjustment mechanism are important in ensuring the buffer responds timeously and effectively to changing conditions. The credit gap thresholds of $L = 2$ and $H = 10$ are internationally established benchmarks recommended by the BCBS to guide the activation and maximum levels of the CCyB. These thresholds were derived from historical analyses of banking crises, and are intended to balance early crisis detection with minimising false signals (Drehmann et al. 2010). This design aims to ensure that threshold breaches occur at least 2–3 years before a crisis or period of distress, allowing banks sufficient time to increase their capital. At the same time, the lower threshold must be high enough to avoid the

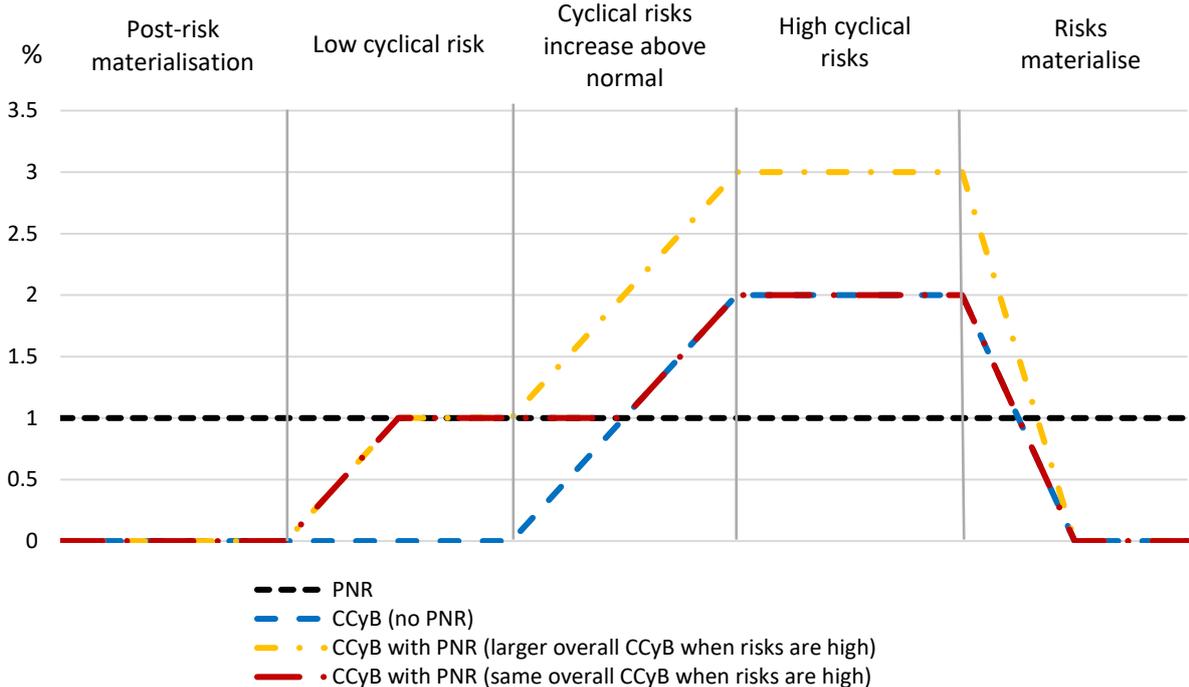
² These limitations – such as its sensitivity to the choice of filter or its potential to generate misleading signals in real time – are explored in more detail below.

imposition of ‘unnecessary’ capital requirements during normal periods, preventing undue strain on financial institutions. The upper threshold ($H = 10$) represents the level at which the CCyB reaches its maximum. Beyond this point, no further capital build-up is required, even if the credit gap continues to grow. The purpose of this upper limit is to ensure maximum readiness for severe crises. The BCBS (2010) shows that $L = 2$ and $H = 10$ achieve a robust trade-off between two key risks: Type I errors, where crises occur without adequate early warning threshold breaches, and Type II errors, where thresholds are breached but no crisis follows (false signals). This balance ensures the thresholds remain effective across diverse jurisdictions.³

Figure 1 compares alternative approaches for implementing the CCyB. The blue dashed line represents jurisdictions with a neutral level of 0% for the CCyB when cyclical risks are assessed to be low. For example, when the credit gap exceeds the lower threshold of 2 percentage points, suggesting that cyclical risks have increased above normal, banks in these jurisdictions are required to hold additional capital. However, for some jurisdictions, the neutral level for the CCyB is set at some level greater than 0% (referred to as the positive cycle-neutral (PCN) CCyB) even when cyclical risks are assessed to be low. Some of these jurisdictions have designed the PCN CCyB to cover risks related to the domestic financial cycle as well as other shocks (i.e. the red line). In this case, the PCN CCyB serves as an early-activation mechanism, triggering the buffer when risks rise above normal levels but not increasing the overall CCyB when risks are already elevated.

³ South African data have been included in international research validating these thresholds, which has found that the performance of these thresholds is not highly sensitive to regional variations (Drehmann et al. 2010).

Figure 1: Evolution of the CCyB with and without a positive neutral rate (PNR)



Source: Adapted from BCBS (2024)

For other jurisdictions, the PCN CCyB has been designed to primarily focus on covering risks unrelated to the domestic financial cycle (i.e. the yellow dash-dotted line). In this case, when cyclical risks begin to rise, the CCyB increases from the level of the PCN CCyB (e.g. 1%) when the credit gap exceeds the lower threshold of 2 percentage points, up to the maximum level of the CCyB (i.e. greater than 2.5%) when the credit gap exceeds the upper threshold of 10 percentage points. Therefore, this results in a larger overall CCyB and more releasable capital when risks materialise for jurisdictions that follow this approach.

3. Data and methodology

3.1 Data

In South Africa, the credit-to-GDP ratio is calculated using credit extended to the domestic private sector by monetary institutions.⁴ While closely following the BCBS’s recommendations, this measure prioritises efficacy and data availability but excludes

⁴ Monetary institutions include the SARB, the former National Finance Corporation, the Corporation for Public Deposits and the so-called ‘pooled funds’ of the former Public Debt Commissioners, the Land Bank, Postbank and private banking institutions (including former banks, discount houses, equity building societies and mutual building societies).

credit extended by non-bank financial institutions, foreign sources and the informal financial sector.

3.1.1 Data sources

Data used in this paper include monthly nominal credit extended (KBP1347M) and quarterly nominal GDP (KBP6006L) over the sample period 1970Q1 to 2024Q3 from the SARB Quarterly Bulletin (SARB 2025).⁵ In each quarter, the final month's observation of credit data (i.e. March, June, September and December) is taken to represent credit extension for that quarter. The credit-to-GDP ratio in quarter t is duly calculated as $r_t = \text{credit}_t / \text{GDP}_t$.

3.1.2 Real-time data

However, published Quarterly Bulletin data represent final revised data and are therefore of limited use in constructing true real-time estimates of the credit gap. This is important, as GDP data, in particular, are prone to updates and revision up to several years after their first publication; therefore, final GDP data can differ significantly from what they appeared to be in 'real time'. Final data can, however, be useful to derive an ex post benchmark measure of the credit gap over the full sample, against which different real-time credit gap estimations can be evaluated.

We therefore complement our full-sample final-data analysis with real-time measures of the credit-to-GDP ratio, using real-time data vintages sourced from the SARB's historical Core and Quarterly Projection Models going back to September 2000.⁶ These vintages represent original observations of GDP and credit extension that policymakers had at hand at the time; combining these vintages therefore yields the closest measure of a real-time credit-to-GDP ratio. Various filters are then tested and applied to this real-time series to identify the calibration that yields a real-time measure of the credit gap that is the most consistent with the full-sample credit gap estimates.

⁵ The credit aggregate is largely in line with the BCBS's recommendations. Bank credit comprises around 97% of nominal credit extended by all monetary institutions.

⁶ We are grateful to M G Ferreira and Rowan Walter for their assistance in sourcing these data vintages.

3.1.3 Real-time vs revised data

The influence of data revisions is significant. Figure 2 illustrates that final revised nominal GDP numbers are generally higher than their real-time counterparts,⁷ while data on credit extension are revised minimally. For this reason, the credit-to-GDP ratio appears significantly higher in real time than ex post, with the real-time estimated credit ratio permanently between 5 and 10 percentage points higher than the final ratio.⁸

It is also interesting to observe that the real-time and ex post ratios converge near the end of the sample. This is likely due to the most recent data not having been fully revised yet, so these observations are likely to still change somewhat over the next few years. However, knowing that credit gap estimates must be derived from data which are bound to change introduces an additional layer of uncertainty in the real-time identification process.

3.2 Methodology: identifying the gap

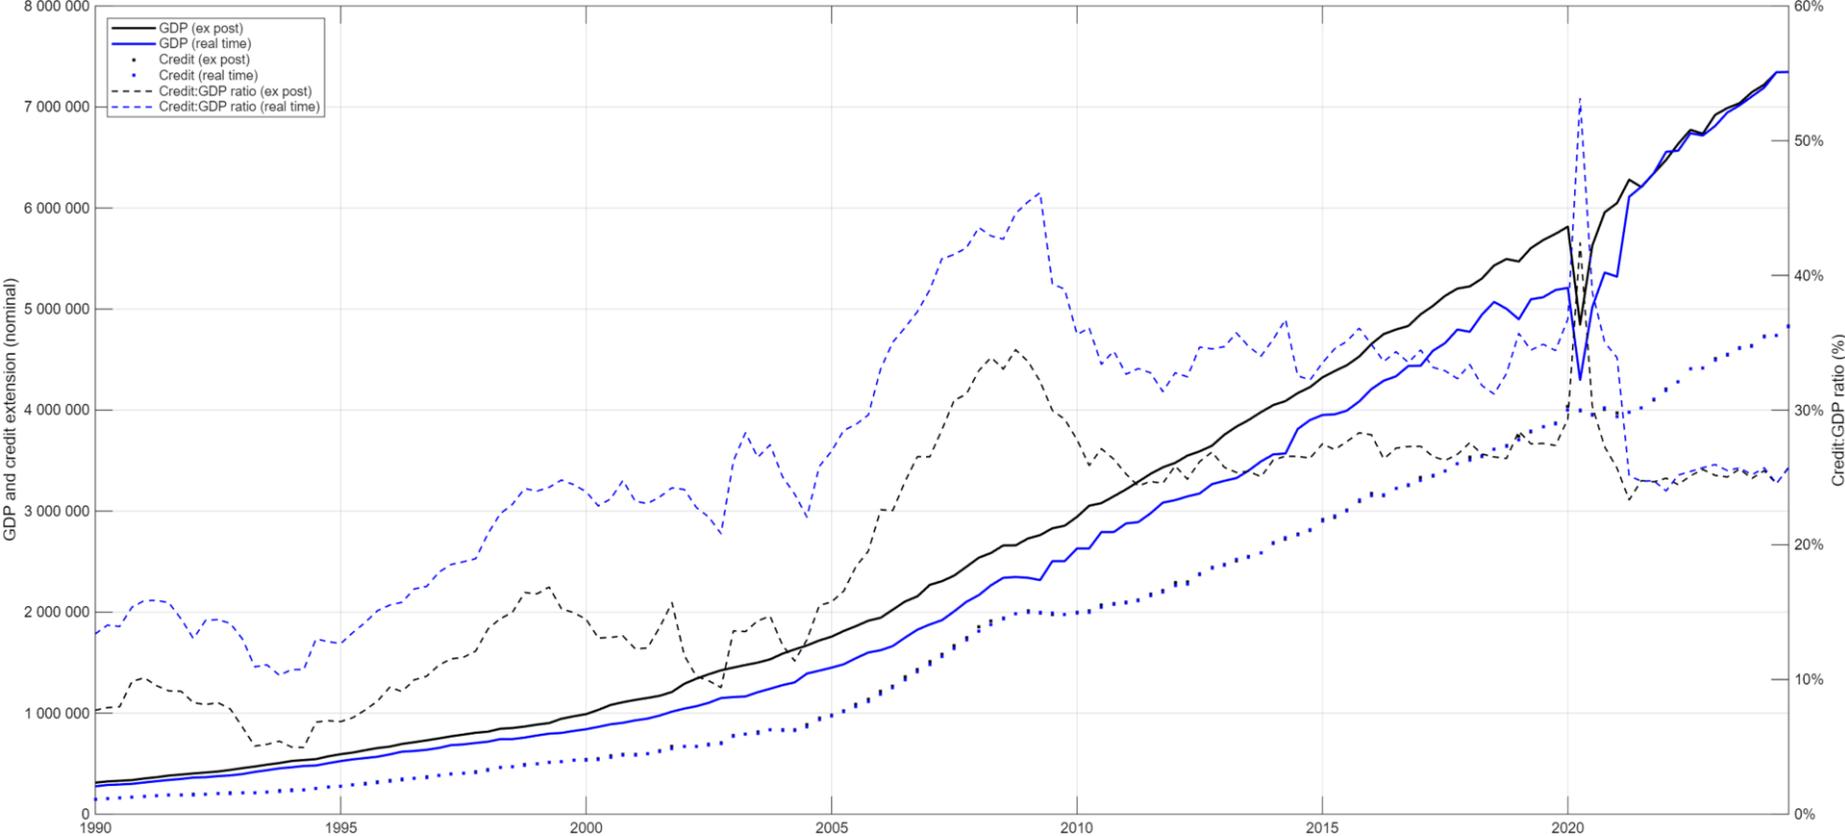
Since any economic time series can be presented as the sum of a trend and cyclical component (Burns and Mitchell 1946; Hodrick and Prescott 1997), the standard approach to extracting the credit-to-GDP gap is to fit a univariate statistical filter on the credit-to-GDP series (r_t) from Figure 2. That is, $r_t = r_t^T + r_t^C$, where r_t^T and r_t^C represent, respectively, the trend and the cyclical components. In this paper's context, the difference between the original credit-to-GDP series and its trend then yields the credit-to-GDP gap, equivalent to the ratio's cyclical component.⁹

⁷ Early prints of GDP numbers may be based on partial or provisional data. Nominal GDP data are also prone to adjustment through the revision of price-level deflators if inflation data are later revised.

⁸ This is the result of the denominator (GDP) being revised upward, resulting in lower final credit ratio estimates relative to real-time estimates.

⁹ The credit gap is therefore calculated as $r_t^C = r_t^T - r_t$.

Figure 2: The credit-to-GDP ratio and data revisions



Sources: Own calculations from SARB (2025) and SARB Core Model

This paper evaluates the well-known HP filter (Hodrick and Prescott 1997) and variations thereof. HP filters are widely used in numerous empirical applications, as they are relatively simple to calculate and easy to interpret. Although HP filters are not the only way to extract trends and cycles¹⁰ from macroeconomic data, they are central to the BCBS's approach to estimating the credit gap, and are therefore a natural starting point for this analysis. Four HP filter variants are used here: the conventional two-sided HP filter, the one-sided HP filter, an adjusted one-sided HP filter and a rolling HP filter. The next subsections will briefly describe the strengths and weaknesses of each filter, with more detailed discussion on each filter provided in Annexure A.

3.2.1 HP filters

Following Hodrick and Prescott (1997), an economic time series y_t can be disaggregated into a long-term trend and short-term cycle by minimising the following loss function:

$$L = \sum_{t=1}^T (y_t - \tau_t)^2 + \lambda \sum_{t=1}^T [(\tau_t - \tau_{t-1}) - (\tau_{t-1} - \tau_{t-2})]^2 \quad (1)$$

where y_t is the original time series, τ_t represents the trend and $y_t - \tau_t$ the cycle. The smoothing parameter λ is a "positive number which penalises variability in the growth component series" (Hodrick and Prescott 1997: 3). The larger the value of λ , the smoother the estimated trend series, while a smaller λ will yield a closer fit of the series.¹¹

The two-sided HP filter

The conventional two-sided HP filter is widely used to extract trends from macroeconomic time series, as it is computationally easy and not subject to rigorous data requirements. However, the HP filter may result in spurious dynamics that are not

¹⁰ One may also consider other univariate filters such as the Beveridge and Nelson (1981), Baxter and King (1999) and Christiano and Fitzgerald (2003) filters, as well as alternative model-based and structural approaches. Model-based approaches, however, introduce additional complexities around model structure and data requirements, which may not necessarily be suitable for real-time estimations.

¹¹ If $\lambda = 0$, the trend will equal the original series, whereas if $\lambda \rightarrow \infty$ the trend converges to a straight line. The choice of the smoothing parameter is discussed in more detail in section 3.3.2.

found in the underlying data (Hamilton 2018), while Orphanides and van Norden (2002: 569) lament the “pervasive unreliability of [its] end-of-sample estimates”. The latter critique is often referred to as the ‘end-point problem’, where the two-sided filter tends to revise its end-of-sample estimates as new data points become available. For this reason, gaps identified by the two-sided HP filter near the end of the sample are generally deemed to be accurate only when they are a few years old. It follows that the standard two-sided HP filter is not the ideal tool for real-time estimations – revisions of past estimates are clearly not desirable – although it can provide a good ex post benchmark against which to compare alternative approaches.

The one-sided HP filter

Given the shortcomings of the conventional two-sided filter, the BCBS recommends using a one-sided HP filter to extract the trend from the credit-to-GDP ratio. The key difference between two is that the one-sided filter uses only past and current observations to estimate the trend at each point in time t ¹² – the one-sided filter is run recursively, with the sample expanding each period – while the two-sided filter uses the entire data sample T to estimate the trend for each point in time t . Since policymakers “can only use the information they have available at each point in time” (Drehmann and Yetman 2018: 3), the BCBS favours the one-sided filter for real-time decision-making in this context.

The one-sided HP filter is less vulnerable to the end-point problem of the two-sided filter. Because the one-sided filter is applied recursively, using only past and current data, end-of-sample distortions are reduced. However, in contrast to the two-sided HP filter, the one-sided HP filter may be vulnerable to the so-called ‘start-point problem’, where the filter’s estimates are sensitive to the chosen start date of the data, particularly if the sample begins near a financial peak or trough. However, given that the primary goal of this analysis is to extract reliable credit gap estimates near the end of the sample, the start-point problem is arguably of less concern.

¹² We denote the full sample by T , while a point in time within the sample is denoted by t . The two-sided filter therefore uses information up to time T to extract the credit gap for period t , while the one-sided filter only uses information up to time t .

Despite this limitation, the one-sided HP filter is praised for its simplicity, ease of use and clear communication of results. Drehmann and Yetman (2018) and Drehmann and Juselius (2014), among others, argue that the one-sided HP filter effectively fulfils its primary purpose of serving as an EWI for financial or banking crises. Policymakers rely on real-time data, not ex post adjustments, and it is argued that the one-sided HP filter remains the best available option compared to alternatives such as the linear projection methodology proposed by Hamilton (2018) or the credit growth rate of Jordà, Schularick and Taylor (2013).

Adjusted one-sided HP filter

Van Vuuren (2012: 316) suggests that well-known problems of the two-sided HP filter can be “alleviated by implementing a single-sided filter”. However, Wolf, Mokinski and Schüler (2024: 1) contend that the standard one-sided HP filter “fails to remove low-frequency fluctuations to the same extent” as the two-sided filter, which results in more volatile trend and cycle estimates.

For this reason, they propose modifying the one-sided HP filter to better align its properties with the two-sided filter. Their approach is closer to a model-based Kalman filter approach than to the univariate optimisation problem of the standard one- and two-sided HP filters. Specifically, they propose a lower value for the smoothing parameter, and a multiplicative rescaling of the extracted cyclical component. This arguably contributes to a more reliable real-time estimate of the credit gap.

The rolling HP filter

The rolling HP filter recursively applies the HP filter over a moving window of fixed length, using only past data at each point in time, instead of applying the HP filter to the entire data set at once. It is more resilient to structural breaks, as old observations eventually drop out of the estimation window. Similar to the one-sided filter, it is also less vulnerable to the two-sided filter’s end-point problems, as it uses only past and

current data to extract the trend and cycle at each point, so it is more suitable for real-time analysis.¹³

Its main shortcoming is that estimates early in the sample may be less reliable as it has fewer data points available in the smaller earlier estimation windows. There is also a wider confidence band in the middle of the sample, where many of the rolling estimation windows overlap.

3.2.2 Calibrating the smoothing parameter (λ)

The second dimension of the credit gap identification approach is to calibrate an appropriate smoothing parameter. Hodrick and Prescott (1997) proposed a smoothing parameter of $\lambda = 1\,600$ for their original minimisation problem (equation 1), based on their observations of post-war US business cycles averaging around 8 years. However, while business cycles typically average 6 to 10 years, financial cycles usually span longer periods.¹⁴ As such, it becomes necessary to adjust the smoothing parameter to correctly capture the underlying cycle: if λ is too big, we risk smoothing away meaningful movements in the data, whereas if λ is too small, we risk mistaking noise for real cycles.

Ravn and Uhlig (2002) explain that smoothing parameters for longer cycles may be calculated through:

$$\lambda = \lambda_0 \left(\frac{f}{f_0} \right)^4 \quad (2)$$

where λ_0 is the benchmark lambda (1 600 for quarterly data), f_0 is the benchmark cycle length, and f is the target cycle length. If, for example, the business cycle averages an eight-year duration ($f_0 = 8$) while the financial cycle averages 24 years ($f = 24$), the ratio f/f_0 would equal 3, and the resultant smoothing parameter would be $1\,600 \times 3^4 = 129\,600$. Similarly, the BCBS endorses a λ of 400 000 based on assumptions that

¹³ We use a one-sided variant of the rolling HP filter. Two-sided variants are also possible, although that introduces the same end-point problems and the requirement of data beyond the current estimation point that plague the standard two-sided HP filter.

¹⁴ Characteristics of the South African financial cycle are explored in Annex B.

financial cycles are typically three to four times longer¹⁵ than business cycles; this aligns with observations that systemic crises tend to occur approximately once every 20–25 years (Drehmann and Tsatsaronis 2014). In the South African context, however, given that the average length of the financial cycle over our sample is around 11 years, (roughly twice as long as the domestic business cycle (see Table B.1)), we propose setting $\lambda = 1\,600 \times 2^4 = 25\,600$ for our credit-gap estimates. Evaluating the relative performance of these alternative smoothing parameters in the South African context could provide more tailored insights into the reliability of the credit-to-GDP gap as an EWI of domestic systemic risk and may better reflect the country’s financial dynamics. Different parameterisations of λ are further discussed in the sensitivity analyses in Annex C.

3.3 Evaluating the predictive performance of the credit-to-GDP gap

In the empirical literature, the predictive performance of the credit-to-GDP gap is commonly assessed using the area under the receiver operating characteristic curve (AUROC), which evaluates the trade-offs between Type I (missed crises) and Type II (false signals) errors, and how well an indicator differentiates between crisis and non-crisis periods (Drehmann and Juselius 2014). A higher AUROC value indicates better predictive performance, though its reliability depends on the availability of sufficient crisis observations. International studies such as Drehmann and Juselius (2014) and Giese et al. (2014) benefit from rich panel data sets across several countries with numerous observations of periods of financial distress.

However, the AUROC approach is less suited to single country analysis, especially given the relatively limited number of crisis periods in South Africa. We instead rely on Farrell and Kemp’s (2020) interpretation of peaks in the financial cycle as “periods where financial conditions are stressed”, and as such, we focus our analyses on the ability of the different credit gaps we construct here to reliably predict these turning points.¹⁶ We therefore formally evaluate various real-time credit gap measures against an ex post credit gap extracted from final revised data through a conventional two-

¹⁵ This calculation for quarterly data yields $1\,600 \times 4^4 = 409\,600$, or approximately 400 000.

¹⁶ Our full sample contains only five financial cycle peaks and downswings (Table B.1), which we interpret as periods of financial distress. However, with the first downswing appearing very early on in the sample, we focus our analyses on the credit gap’s EWI properties from 1980 onwards.

sided HP filter to get a sense of which real-time specification best matches the ex post credit gap.

4. Results and discussion

4.1 Univariate statistical filters

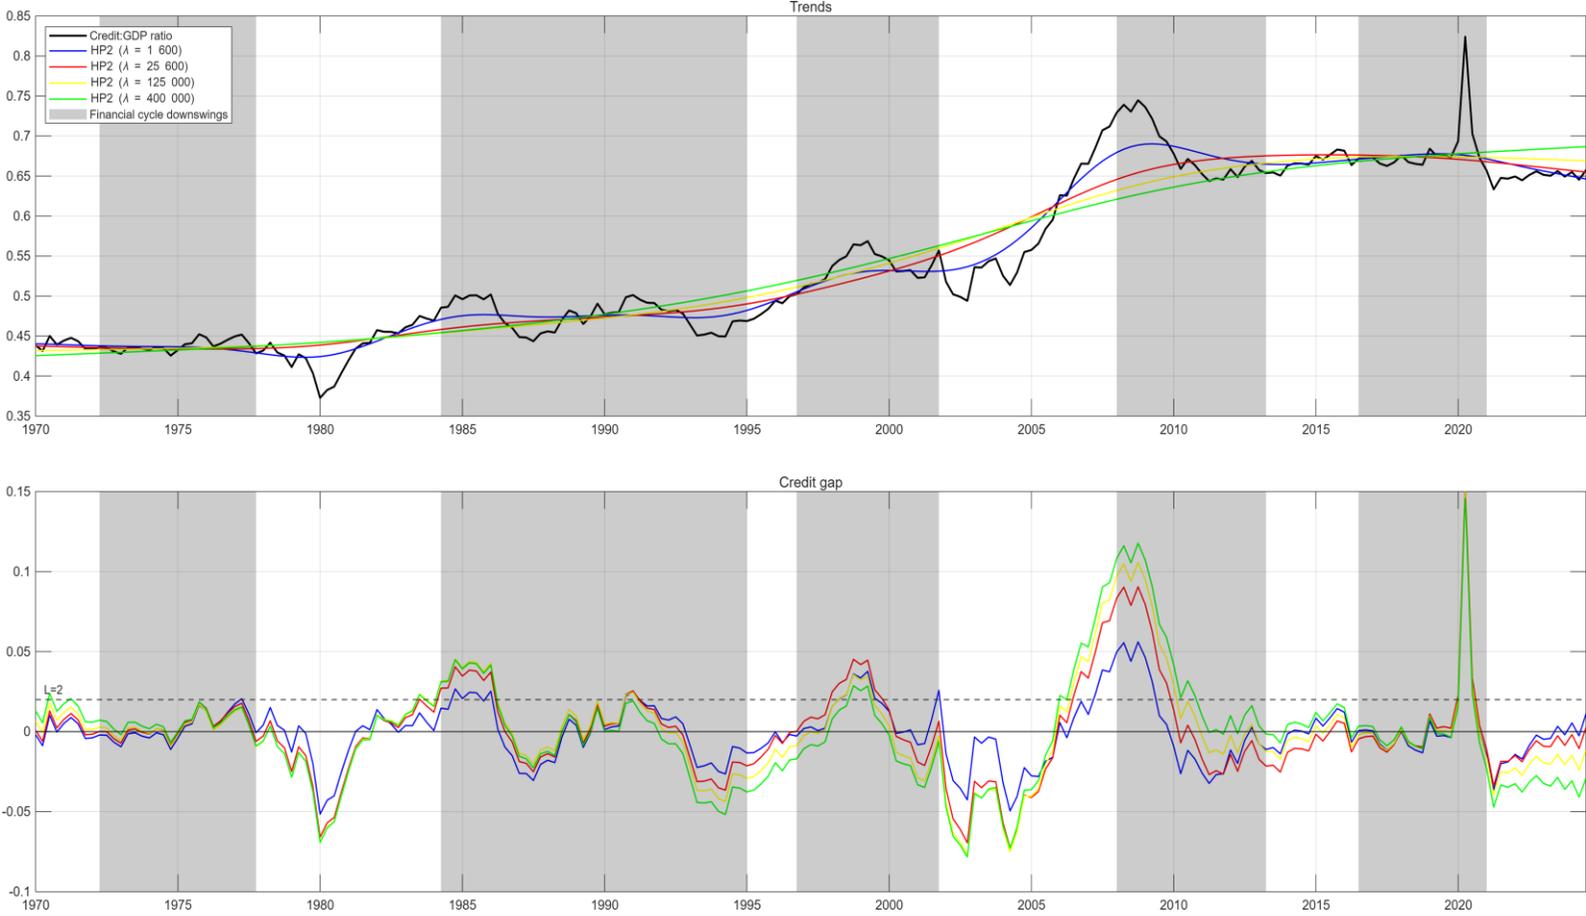
4.1.1 Benchmark two-sided HP filter

As a benchmark we employ the conventional two-sided HP filter (HP2) to extract an ex post measure of the South African credit gap from the final credit-to-GDP ratio (Figure 2).¹⁷ With the benefit of hindsight, it may be argued that the two-sided HP filter yields the ‘best’ final approximation of the credit gap. This can be justified by the fact that the two-sided filter has the benefit of using *all* data points across the full sample, in contrast to the recursive nature of the one-sided filters. Moreover, the majority of data points in the full sample have already been revised, so estimates of earlier trends and cycles should be more stable. As was argued above, the two-sided filter is not ideal for real-time cycle identification and policymaking as it relies on unrealised future observations that are unknown to the policymaker at each point in time; however, the two-sided filter arguably provides a more reliable ex post identification of the cycle against which alternative real-time indicators can be evaluated.

Figure 3 illustrates trends and credit gaps arising from the HP2 filter using final data. It is evident that the larger smoothing parameters yield a much smoother trend, with the trade-off of a somewhat noisier cycle. A smoother trend is consistent with a longer theorised financial cycle; for example, a smoothing parameter of $\lambda = 1\,600$ implies a cycle of 7 years, while the BCBS’s recommended smoothing parameter of $\lambda = 400\,000$ implies a cycle of 20–25 years. We also evaluate $\lambda = 125\,000$, which is consistent with a financial cycle of 17 years (Bosch and Koch 2020; Farrell and Kemp 2020) and is reasonably close to the SARB’s previous parameterisation of $\lambda = 100\,000$ for its benchmark model up to 2024. Finally, we also test a smoothing parameter of $\lambda = 25\,600$, which is more in line with the 11-year average financial cycle duration derived from Table B.1.

¹⁷ Up to 2024, the SARB identified the credit-to-GDP gap by employing a two-sided HP filter with a smoothing parameter of $\lambda = 100\,000$. Forecasts were used to enable the two-sided filter to use data beyond the current period t .

Figure 3: The benchmark two-sided HP filter



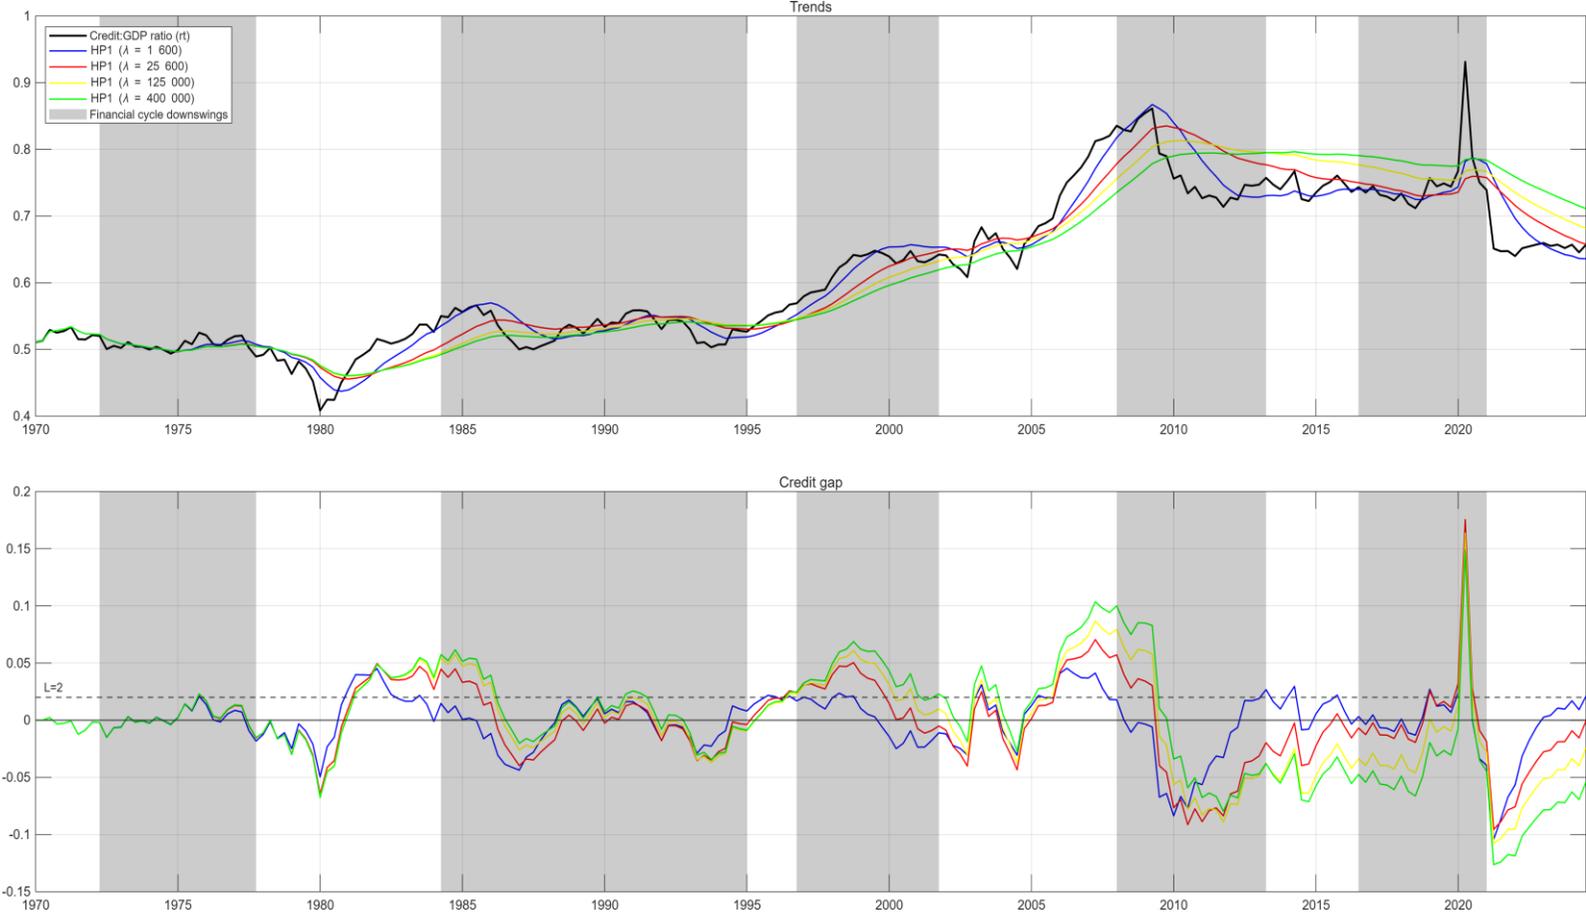
While the different specifications of the two-sided filters yield similar turning points in the credit gap (e.g. 1980, 1984, 1998 and 2008), they are generally 'late' in predicting turning points in the financial cycle. That is, with the exception of the build-up to the GFC, they only breach the CCyB threshold of $L = 2$ *after* the cycle has already turned. Finally, the magnitudes of peaks and troughs in the credit gap are not always consistent, with the larger smoothing parameters generally increasing the credit gap's amplitude.

4.1.2 One-sided HP filter

The BCBS's recommended one-sided HP filter (HP1) is considered next, using the same set of smoothing parameters evaluated above, and is presented in Figure 4. The one-sided filter relies on original unrevised data vintages; it therefore represents the closest to a real-time view that one can achieve. Each financial cycle peak is preceded by the credit gap breaching, or (depending on the parameterisation) getting very close to, the $L = 2$ threshold. This is encouraging as it supports the HP1 filter's real-time EWI properties. However, this does not extend to the very first financial cycle peak in 1972; in fact, the apparent stability of the credit gap during the first eight or so years of the sample may point to the HP1 filter's 'start-point' bias discussed earlier.

Again, when the smoothing parameter is small ($\lambda = 1\ 600$), the credit gap tends to turn earlier and run ahead of other credit gaps. The turning points predicted by our choice of $\lambda = 25\ 600$, relative to $\lambda = 125\ 000$ and the BCBS recommendation of $\lambda = 400\ 000$, are remarkably consistent in terms of timing. However, in terms of magnitude, the larger smoothing parameter generally returns a more pronounced credit gap before the financial cycle turns. This is especially noticeable in the build-up to the GFC, where the BCBS parameterisation suggests the credit gap peaked at 10.4% in 2008, compared with a peak of 7% to 8.7% under our parameterisations. This has implications for the design of the CCyB increments, as it suggests that buffers may not have been at maximum readiness before the biggest financial crisis in living memory.

Figure 4: One-sided HP filter



4.1.3 Adjusted one-sided HP filter

Using Wolf, Mokinski and Schüler's (2024) Matlab routine,¹⁸ and employing their default $\lambda = 1\,600$ against our $\lambda = 25\,600$, as well as $\lambda = 125\,000$ and $\lambda = 400\,000$, yields Figure 8.

The adjusted one-sided HP filter (HP1*) appears to be consistent with the standard HP1 filter (Figure 5), although marginally less volatile. Similar to the standard one-sided HP filter, the HP1* filter consistently breaches the CCyB threshold before the financial cycle turns, with the exception of the 1972 turning point. The HP1* filter is also the only filter which comes close to predicting the 2016 turning point under its smaller calibrations of λ .

4.1.4 Rolling HP filter

Next, we evaluate the rolling HP filter, again using $\lambda = 25\,600$ and real-time data.¹⁹ We choose a window of $W = 60$ quarters, which is slightly longer than the average financial cycle. This allows us to capture longer-term trends without unduly smoothing over structural economic changes.

Figure 6 illustrates that the rolling HP filter produces a significantly noisier credit gap when all the estimation windows overlap,²⁰ making it somewhat challenging for real-time cycle identification. However, it exhibits qualitatively similar results to the one-sided filters with larger smoothing parameters (Figures 4 and 5), in that it identifies comparable peaks and troughs in the credit gap.

¹⁸ The authors' Matlab routine is available at <https://sites.google.com/site/yvesschueler/research>

¹⁹ The differences between various choices of λ are inconsequential, and are therefore not all shown here. Consistent with earlier analyses, the larger smoothing parameters are associated with a generally larger amplitude in the credit gap, but with consistent timing in turning points.

²⁰ The longer the window W , the less noisy the cycle would be; when $W \rightarrow T$ the rolling HP filter would approach the full sample HP1 filter.

Figure 5: Adjusted one-sided HP filter

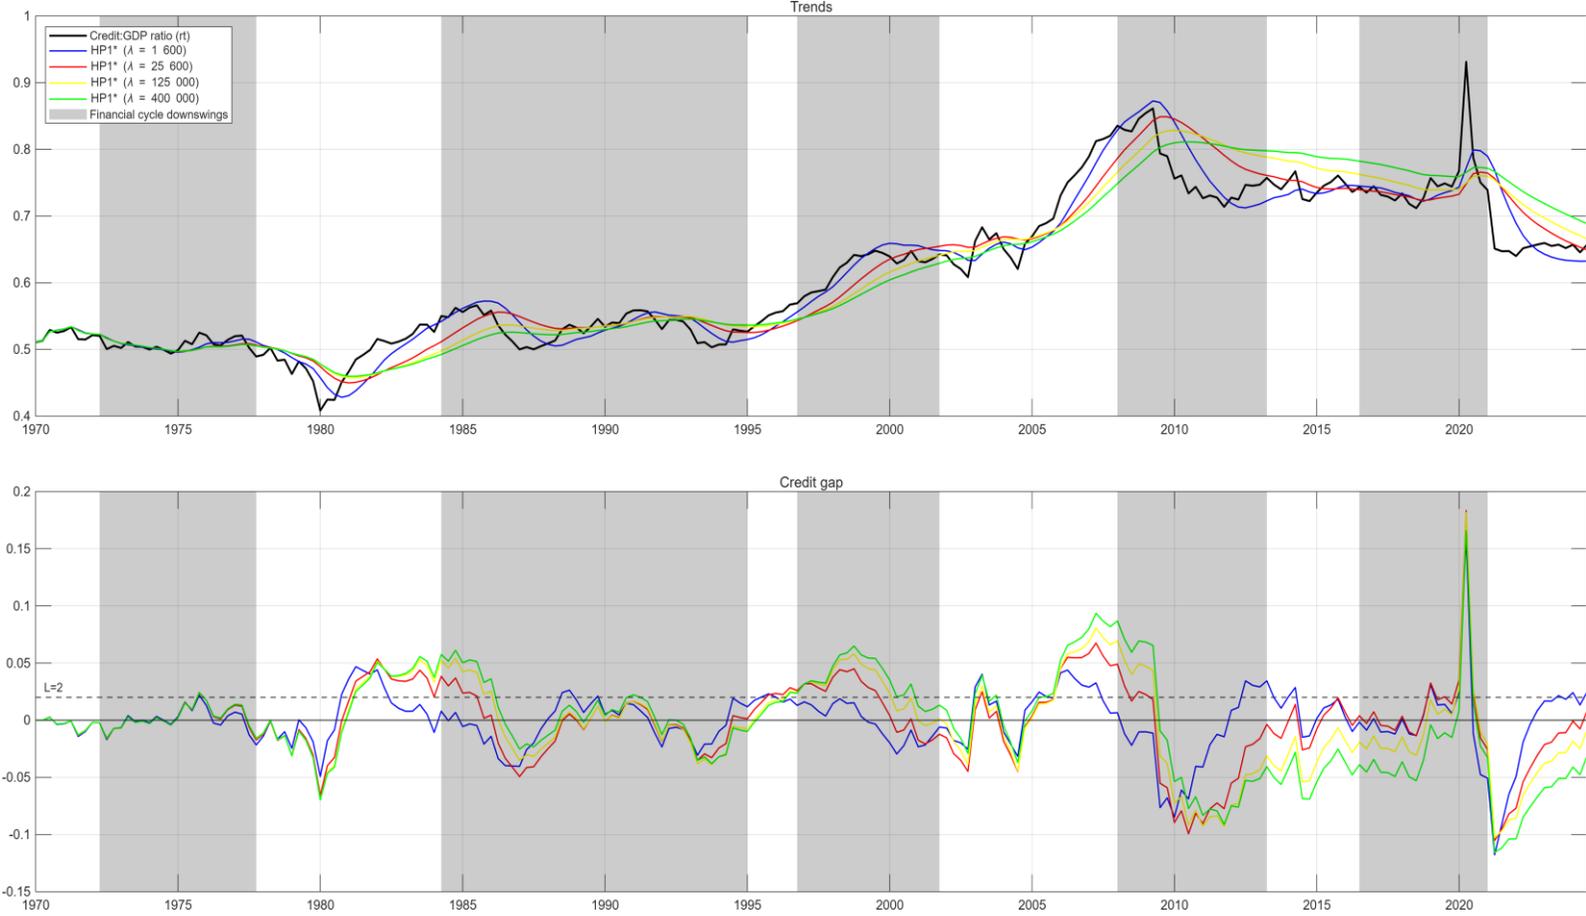
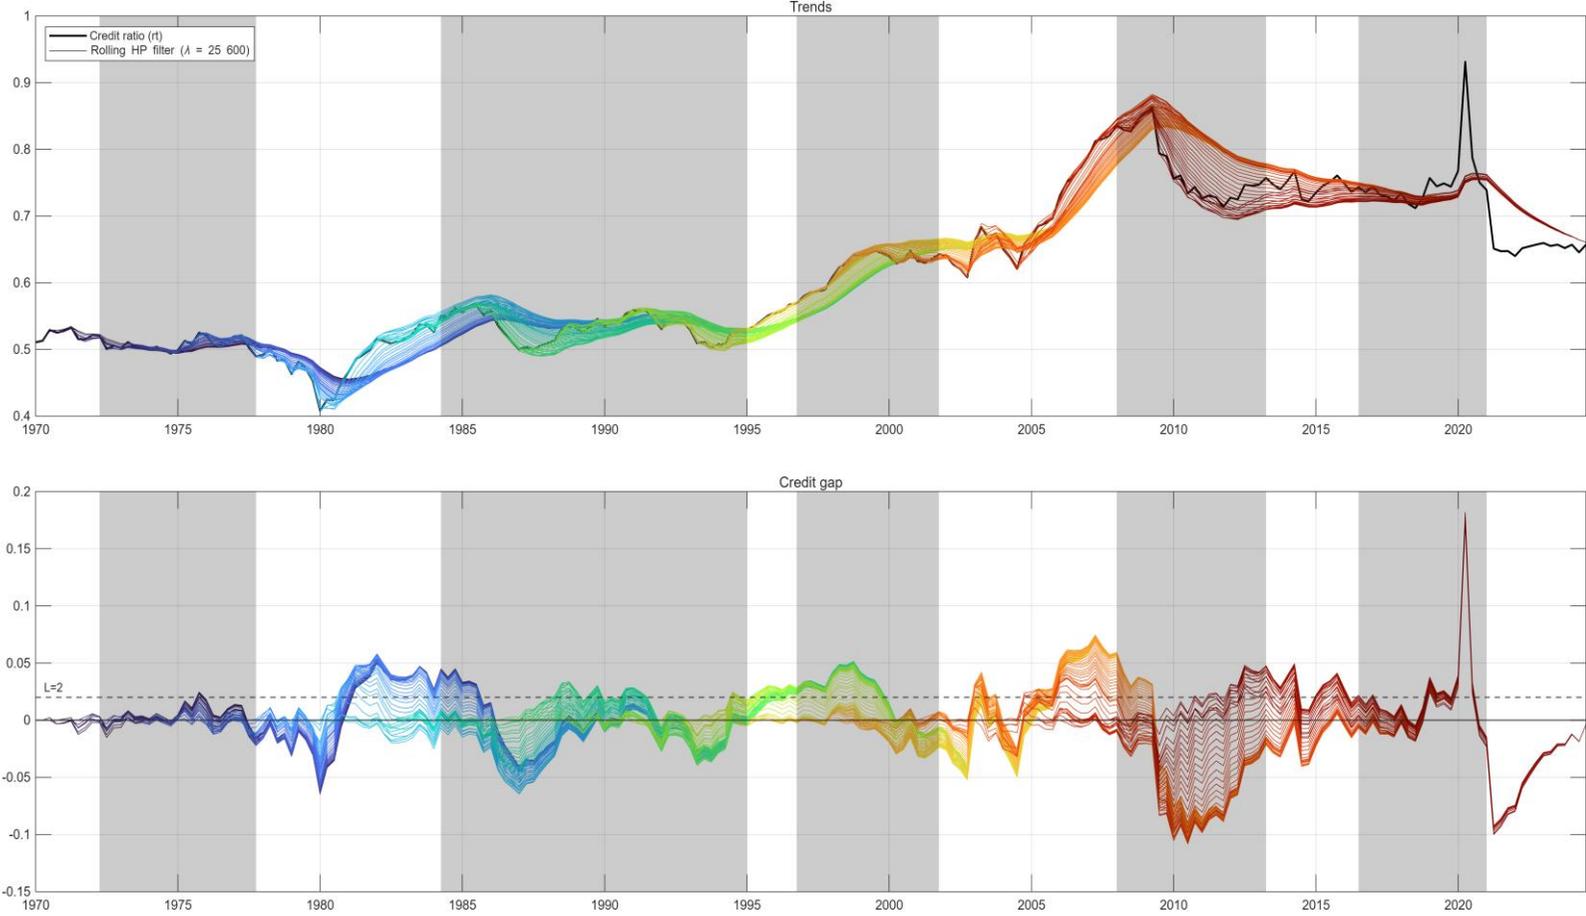


Figure 6: Rolling HP filter (60-quarter rolling window)



4.2 Is there a 'best' real-time filter?

4.2.1 The credit gap and the financial cycle

From Figure 6 it is evident that the rolling HP filter is not particularly suitable for real-time analysis: its noisy overlapping credit gaps make it more difficult to identify periods where the CCyB lower threshold is breached. This may be addressed by averaging out extracted cycles over the different windows; however, this introduces an unnecessary layer of complexity. The two-sided HP filter requires future observations to extract the trend at time t , making it even less suitable for real-time identification of the credit gap. A potential workaround may be to add forecasted observations to the sample (see Annex C); however, accurate identification of the credit gap then becomes sensitive to forecast assumptions and therefore subject to potential errors.

Figures 4 and 5 suggest that the credit gaps identified from one-sided filters quite reliably predict peaks in the financial cycle, although they are less efficient at accurately identifying financial cycle troughs. From 1980, each downswing in the financial cycle was preceded by a build-up in the credit gap, with the credit gap peaking 4–12 quarters before the upper turning points in the financial cycle. Considering that these metrics are concerned with predicting unsustainable credit build-ups *ahead of time*, they seem to serve the intended goal of EWI satisfactorily.

A final consideration is evaluating how long before a potential credit cycle downswing each indicator breaches the CCyB threshold. That is, how suitable is each indicator as an EWI of a period of potential financial distress? As is evident from Figures 4 and 5, from 1980 onwards both real-time HP1 filters breach the $L = 2$ threshold 1–2 years before the start of a financial cycle downswing. An exception is the financial cycle downswing that started in 2016, which was not preceded by a notable credit gap build-up; in fact, the credit gap remained predominantly negative. However, it should be noted that financial cycles are not exclusively driven by credit growth. Other factors, such as developments in equity markets and house prices, also influence the financial

cycle.²¹ This makes it important to not interpret credit gaps mechanically, and for these types of decisions to be made within the context of a broad suite of indicators.

The fact that the credit gap correctly predicts three of the four financial cycle downturns since 1980 – 1984Q2, 1996Q4 and 2008Q1 – suggests that it is indeed suitable as an EWI of potential financial distress. Considering that EWIs “should be evaluated on the basis of their performance relative to the macroprudential policy maker’s decision problem” (Drehmann and Juselius 2014: 759), real-time credit gaps arising from the one-sided HP filter can be seen as a valuable EWI.

Moreover, the one-sided filters hardly create false signals or Type II errors (that is, falsely predicting a turning point that never arrives). The only exception may be the brief 2002/2003 spike (Figures 4 and 5), as house and equity prices started booming shortly after the 2001 rand crisis and the 2002 small banks crisis.

4.2.2 Real-time vs ex post evaluation

To evaluate the real-time approach that is the most consistent with the ex post realised credit gap, we propose the ‘best’ real-time indicator of the credit gap as that measure which is the ‘closest’ to two different benchmark specifications: (i) the SARB’s previous two-sided HP filter, estimated using final and forecast data²² up to 2027 and a smoothing parameter of $\lambda = 100\,000$, and (ii) the one-sided HP filter calibrated to the BCBS default, estimated using final data²³ and a smoothing parameter of $\lambda = 400\,000$.

Figure 7 suggests that the two-sided filter benchmark is often late in predicting financial cycle peaks;²⁴ this is consistent with the results from the standard two-sided filter (Figure 3). The one-sided filters, on the other hand, reliably predict financial cycle peaks. The one-sided filters based on real-time data are potentially better EWIs than

²¹ Indeed, the 2016 financial cycle downswing may have been driven primarily by a downturn in equity prices (SARB 2023), which would not have been observed in considering the credit gap only.

²² These are the SARB’s own forecasts from the Core Model.

²³ Both ex post specifications use final revised, and not real-time, data to identify a *final* credit gap.

²⁴ However, it should be noted that the two-sided filter reliably predicts turning points in the *credit* cycle, which generally lags the financial cycle by 2–3 years. While this falls outside the scope of this analysis and is therefore not shown here, it would be worth evaluating further in future research alongside the interactions between credit and financial cycles.

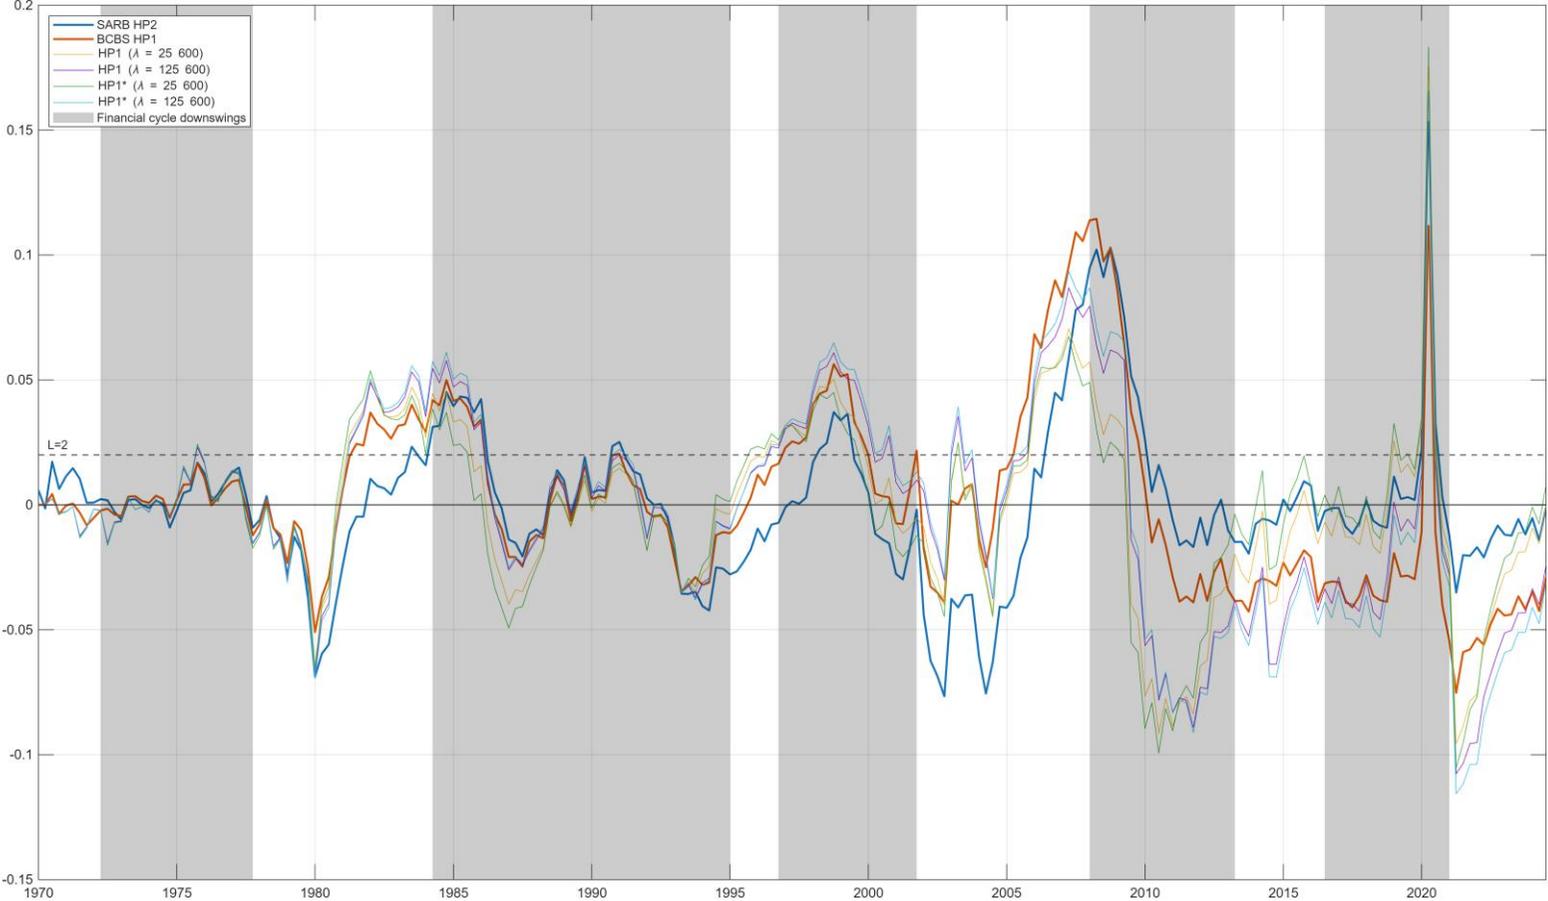
their BCBS final data counterpart, as they give an earlier warning of the 1996 downturn; however, these real-time filters also give a false warning in early 2003. Finally, the one-sided filters with the smaller smoothing parameters, alongside the two-sided filter, suggest a narrower credit gap at the end of the sample, whereas the BCBS benchmark and our filters with $\lambda = 125\,000$ indicate a wider negative credit gap. This may point to lingering effects of the COVID-19 shock that are still mechanically reflected in the various filters, which should normalise over time. However, this recent volatility in the credit gap substantiates the argument that additional indicators need to be considered to ensure a holistic view of credit build-up and financial imbalances across the economy.²⁵

Table 1 provides a comparison of the performance of the one-sided and adjusted one-sided real-time filters with various smoothing parameters presented above against the benchmark final-data filter. A low mean squared error (MSE) and mean absolute error (MAE) indicate a minimal deviation of the real-time credit gaps relative to the benchmark filters based on final data.

Relative to the SARB benchmark, the best-performing real-time filter is the standard one-sided filter with $\lambda = 25\,600$ (filter #3), whereas the one-sided filter with $\lambda = 125\,000$ (filter #2) is the best relative to the BCBS benchmark. This suggests that the 'closest' real-time credit gaps relative to the ex post gaps are indeed the standard one-sided HP filters with a λ of either 125 000 or 25 600. Notably, filters #4 and #8 perform relatively poorly across all metrics. This is perhaps not surprising, as their smoothing parameter of $\lambda = 1\,600$ is not justified by the empirical evidence on the duration of financial cycles (Table B.1), and is far removed from the benchmark filters' larger smoothing parameters.

²⁵ There may also have been a mechanical overshooting of the negative credit gap after the COVID-19 shock. In an effort to ameliorate the massive volatility of the 2020Q2 and Q3 GDP numbers, a smoothing approach is applied to GDP data to extract a smoother credit gap (Annex B.2).

Figure 7: One-sided HP filters against the SARB and BCBS benchmarks



Source: Own calculations from SARB (2023, 2025) and SARB Core Model

Table 1: Evaluating real-time filters relative to ex post final data benchmark

#	Filter	SARB			BCBS		
		MSE	MAE	Corr.	MSE	MAE	Corr.
1	HP1 ($\lambda = 400\ 000$)	0.00145	0.02959	0.58021	0.00037	0.01383	0.92609*
2	HP1 ($\lambda = 125\ 000$)	0.00107	0.02520	0.60904	0.00033*	0.01196*	0.89701
3	HP1 ($\lambda = 25\ 600$)	0.00089*	0.02141*	0.59004	0.00054	0.01535	0.78627
4	HP1 ($\lambda = 1\ 600$)	0.00113	0.02325	0.32522	0.00109	0.02328	0.47730
5	HP1* ($\lambda = 400\ 000$)	0.00121	0.02694	0.60406	0.00035	0.01288	0.90772
6	HP1* ($\lambda = 125\ 000$)	0.00099	0.02345	0.60980*	0.00045	0.01378	0.84251
7	HP1* ($\lambda = 25\ 600$)	0.00101	0.02196	0.53546	0.00075	0.01868	0.70227
8	HP1* ($\lambda = 1\ 600$)	0.00132	0.02563	0.21445	0.00133	0.02634	0.35407

Note: Real-time filters are evaluated relative to a benchmark two-sided HP filter, using forecast data to ameliorate the end-point problem, and a smoothing parameter of $\lambda = 100\ 000$. The BCBS benchmark is a one-sided HP filter with $\lambda = 100\ 000$.

Source: Own calculations

In general, the real-time HP1 filters match the final data full-sample filter well, provided the smoothing parameter is appropriately calibrated. This, coupled with its relative ease of calculation, leads us to favour the standard one-sided HP filter for accurate and reliable real-time credit gap estimation. In addition, setting the smoothing parameter to $\lambda = 25\ 600$ is empirically consistent with the average duration of the South African financial cycle, and we conclude that filter #3 is the most suitable for the South African context.

4.3 Implications for activating the CCyB

4.3.1 Activation thresholds

Drehmann et al. (2010) show that South Africa's credit gap dynamics align with the BCBS framework. However, Farrell (2016) emphasises the need to periodically revisit these thresholds considering local median credit-to-GDP gap figures and the unique characteristics of South Africa's credit cycles.²⁶ Farrell (2016) reports that the spread of credit-to-GDP gap estimates in South Africa historically would have triggered more frequent activation of the CCyB than the BCBS anticipates, suggesting that the default calibration of the buffer thresholds may need to be adjusted for domestic application. This finding underscores the importance of further threshold analysis and highlights

²⁶ Notably, this includes the somewhat shorter South African credit and financial cycles relative to the BCBS assumptions (see Table B.1).

that if the BCBS’s (2010) expectation of activating the buffers only once every 10–20 years is to be met, the default calibration may require modification.

Table 2 presents the implications for how often the CCyB would be activated under various filtering approaches and threshold calibrations. We evaluate the ex post credit gaps arising from the two benchmark filters, as well as the real-time gaps arising from the HP1 and HP1* filters, under two smoothing parameters, $\lambda = 25\,600$ and $\lambda = 125\,000$. The baseline threshold $L = 2$ (tantamount to the credit gap breaching the 2% level) is then compared to alternative thresholds of $L = 3$ and $L = 4$.

Table 2: Calibrating the CCyB lower threshold

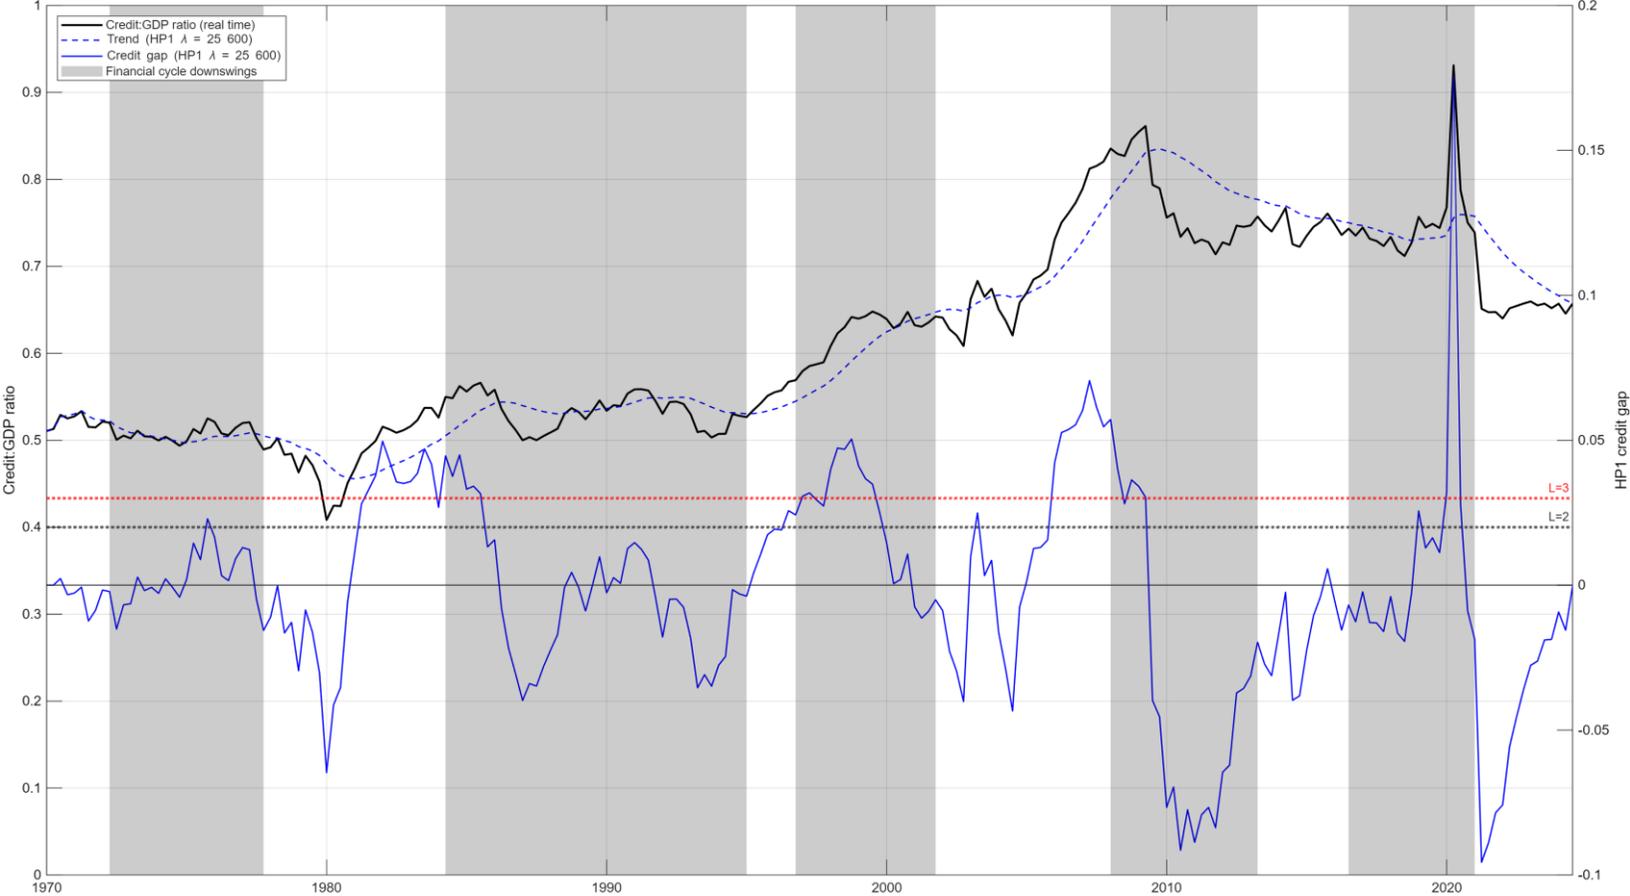
	Thresholds					
	L=2		L=3		L=4	
HP2 SARB ($\lambda = 100\,000$)	34	15.5%	26	11.8%	18	8.2%
HP1 BCBS ($\lambda = 400\,000$)	53	24.2%	40	18.3%	27	12.3%
HP1 ($\lambda = 125\,000$)	55	25.1%	47	21.5%	34	15.5%
HP1 ($\lambda = 25\,600$)	52	23.7%	40	18.3%	21	9.6%
HP1* ($\lambda = 125\,000$)	54	24.7%	45	20.6%	32	14.6%
HP1* ($\lambda = 125\,600$)	53	24.2%	32	14.6%	17	7.8%
Total observations	219					

Source: Own calculations

Our suggested credit gap measure implies that the CCyB will be active between 23% and 25% of the time under the BCBS’s default threshold of $L = 2$. This is somewhat smaller than Farrell’s (2016) result of just under 40%. However, Farrell’s (2016) sample only extends to 2013, whereafter the credit gap has been largely stable between -4% and 1%; these observations naturally bias our result downwards. Naturally, the higher the L , the less frequently the CCyB will be activated. However, if the threshold is set too high, the SARB could be too late to impose effective macroprudential measures; that is, by the time a threshold of, for example, $L = 4$ is crossed, the window for banks to build up credit buffers under a CCyB might be too small.

Figure 8 plots our recommended one-sided HP filter against thresholds of $L = 2$ and $L = 3$. If the threshold is too high, the CCyB might not be deployed in time to ensure sufficient capital buffers for the downswing. For example, if the CCyB is activated at a threshold of $L = 3$, it will be late on the 1996 turning point. On the other hand, if the threshold is too low, it may become burdensome to manage the frequent deployment of the CCyB, or it may lead to false signals.

Figure 8: Recommended one-sided HP filter and CCyB thresholds



Source: Own calculations from SARB (2023, 2025) and SARB Core Model

5. Conclusion

This paper evaluated different methods to estimate the South African credit gap in real time. The credit gap can indicate excessive credit growth relative to economic activity, and can serve as an EWI for financial imbalances. Such metrics can be part of a suite of macroprudential indicators to guide the activation of capital buffer requirements, strengthening the financial system's resilience during distress.

The statistical analyses aimed to show how different models might produce distinct credit gap measures. An effective EWI should accurately and timeously identify financial cycle peaks, especially the start of downturns, to ensure sufficient capital buffers are in place. It should also be simple to calculate without requiring extensive data.

Based on these criteria, the real-time credit gap from a one-sided HP filter with a 'middle-of-the-road' smoothing parameter is deemed to be the best EWI for the local context. Setting the smoothing parameter to 25 600 or 125 000 aligns with the average duration of South Africa's financial cycle (at 11–15 years, roughly 2–3 times the duration of the business cycle), which is shorter than the BCBS's 20–25-year expectation. This credit gap specification correctly predicts three out of four financial cycle peaks since 1980 and matches ex post credit gaps from revised data relatively well. It is also easy to calculate.

We also examined different thresholds for activating the CCyB. Our credit gap measure suggests the CCyB would be active 23–25% of the time at a threshold of 2, but only 8–15% with a threshold of 4. A conservative buffer is preferable, as missing a turning point is arguably more costly than false signals or too frequent activation.

In conclusion, the credit-to-GDP gap is a reliable EWI for potential financial distress. However, it should be part of a broader set of macroprudential tools and not applied too mechanically. Judgement and discretion in capital buffer decisions remain essential.

Annexures

Annex A: Hodrick-Prescott (HP) filter variants

Two-sided HP filter

Following Hodrick and Prescott (1997), an economic time series y_t can be disaggregated into a long-term trend and short-term cycle by minimising the following loss function:

$$L = \sum_{t=1}^T (y_t - \tau_t)^2 + \lambda \sum_{t=1}^T [(\tau_t - \tau_{t-1}) - (\tau_{t-1} - \tau_{t-2})]^2 \quad (\text{A.1})$$

where τ_t represents the trend and $y_t - \tau_t$ the cycle. The smoothing parameter λ is a “positive number which penalises variability in the growth component series” (Hodrick and Prescott 1997: 3). The larger the value of λ , the smoother the estimated trend series, while a smaller λ will yield a better fit of the series. If $\lambda = 0$, the trend will equal the original series. For quarterly data, $\lambda = 1\,600$ is typically chosen to match business cycles with a 7–8-year duration.

Importantly, the minimisation problem is solved over the full sample T . This implies that for any individual time period t within the sample (that is, $1 < t < T$), the two-sided filter uses all data points across the full sample to calculate the trend and cycle at time t . The two-sided filter therefore requires both past and future data observations to extract the trend at every point t – that is, it “revises its inference on all observations in the sample as new observations become available and artificially correlates observations at time t with future observations at $t + h$ ” (Wolf, Mokinski and Schüler 2020: 1). However, this is not a true reflection of real-time trend identification where future data observations are, of course, unknown. Furthermore, if one wishes to extract the trend and cycle near the end of the sample (that is, where t comes close to T), the end-point problem of unreliable estimates near the end of the sample becomes prevalent (Orphanides and van Norden 2002).

One-sided HP filter

The one-sided filter modifies equation A.1 to depend not on the full sample T , but only the series' past and current observations up to time t .

$$L = \sum_{s=1}^t (y_s - \tau_s)^2 + \lambda \sum_{s=1}^t [(\tau_s - \tau_{s-1}) - (\tau_{s-1}^T - \tau_{s-2}^T)]^2 \quad (\text{A.2})$$

Wolf, Mokinski and Schüler (2020, 2024) try to make the one-sided HP filter more reliable in real time by harmonising its properties with those of the two-sided filter. By using a lower value for the smoothing parameter and a multiplicative rescaling of the cyclical component, they are able to minimise the squared distance of its power transfer function (PTF) with the PTF of two-sided filters, improving the one-sided filter's real-time properties.

Rolling HP filter

The rolling HP filter applies the standard one-sided HP filter over a moving window W of fixed length. At each time t , only a subset of past data (e.g. the last 30 or 40 quarters) are used to estimate the trend. Where $W = T$ (i.e. the full sample), the rolling HP filter becomes the standard one-sided filter.

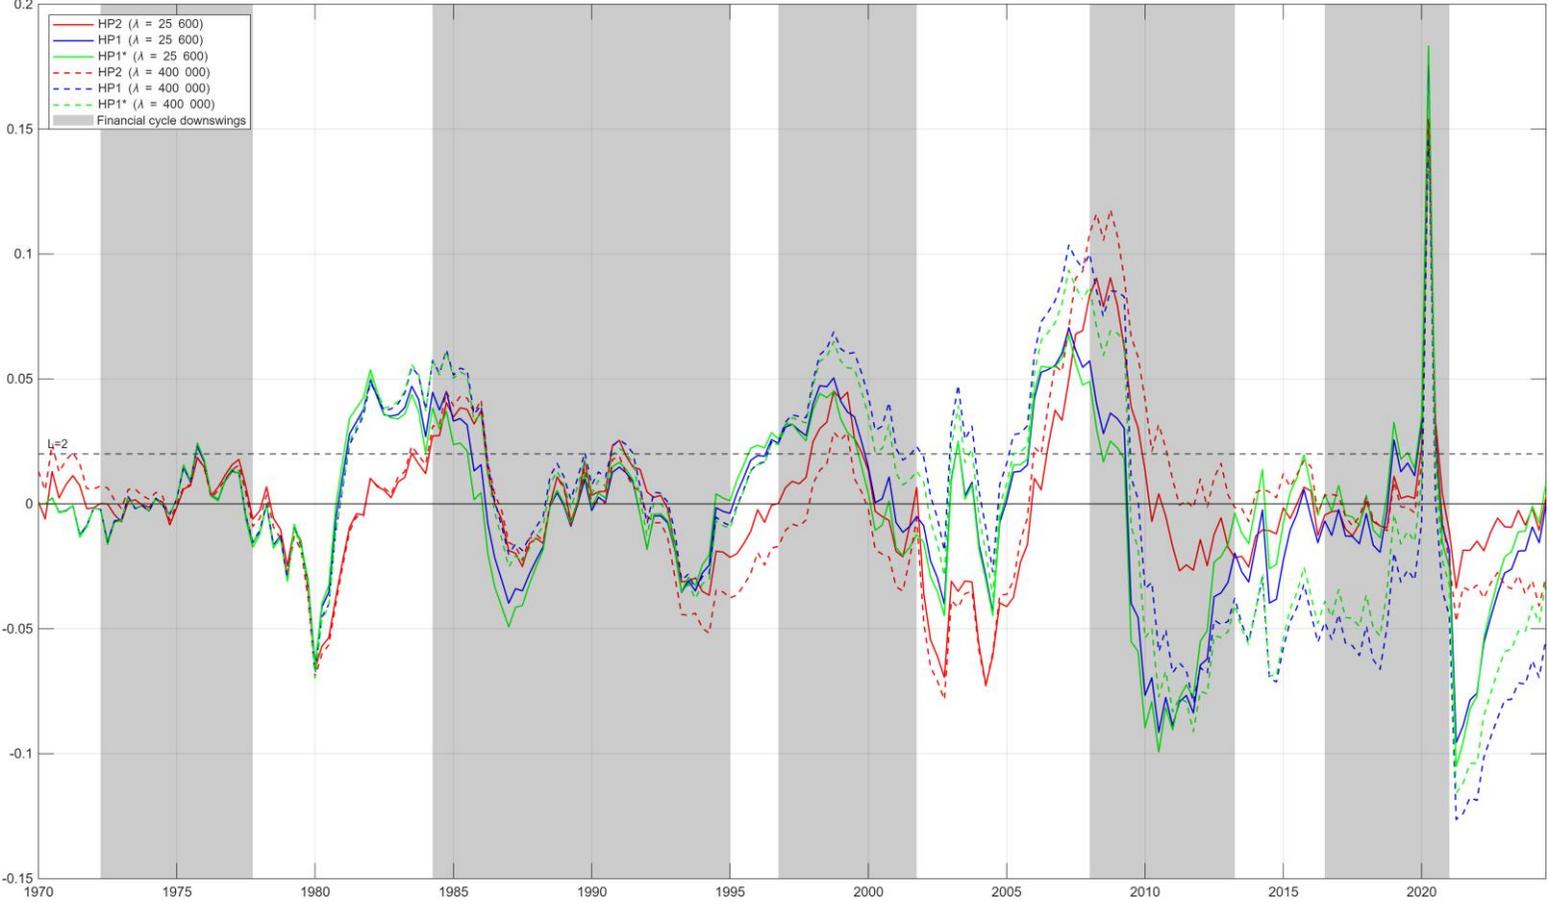
$$L = \sum_{s=t-W+1}^t (y_s - \tau_s)^2 + \lambda \sum_{s=t-W+1}^t [(\tau_s - \tau_{s-1}) - (\tau_{s-1}^T - \tau_{s-2}^T)]^2 \quad (\text{A.3})$$

Comparison of HP filters

The key difference between the one- and two-sided HP filters is that the one-sided filter and its variations use only past and current observations to estimate the trend at each point in time, while the two-sided filter uses the entire data sample to estimate the trend for each point in time. The two-sided filter therefore uses both past and future data observations to extract the trend, which is not a true reflection of real-time trend identification where future data observations are obviously unknown. For this reason, the one-sided filter is more appropriate for real-time decision-making.

Figure A.1 illustrates the differences between the one-sided and two-sided HP filters. The one-sided filters are applied to the real-time credit-to-GDP ratio, whereas the two-sided filter uses final revised data. We analyse both our smoothing parameter of $\lambda = 25\,600$ and the BCBS's benchmark of $\lambda = 400\,000$.

Figure A.1: Comparison: one- and two-sided HP filters



Note: One-sided credit gaps are derived from real-time data.

Source: Own calculations from SARB (2023, 2025) and SARB Core Model

With the 'benefit' of knowing the entire series and all data points at every point t , the two-sided filter extracts a significantly smoother trend (compare Figures 3, 4 and 5). This does, however, result in a structural underestimation of the credit gap (Figure A.1), compared to the gap derived from the one-sided filters, between 1980 and 1985 and again between 1993 and 2008. This may be the result of the two-sided filter smoothing over significant events, such as the early 1980s credit crunch resulting from capital flight and disinvestment from South Africa, and the post-democratisation development of domestic financial markets.

A constant, but relatively muted, credit gap is observed over the first 8–9 years of the sample under all specifications. From 1980 onwards, however, swings in the credit gap have become more pronounced and volatile, perhaps resulting from South Africa's increased integration in global financial markets and, consequently, more domestic financial vulnerabilities to international events (notably, the late-1990s emerging markets financial crises, the GFC and the COVID-19 shock).

However, as was highlighted in Figure 3, the HP2 filter seems to be late in predicting turning points in the financial cycle: both variants of the one-sided filters are better EWIs as they breach the $L = 2$ threshold well in advance of financial cycle peaks. None of these filters correctly identifies the 2016 turning point, although the HP1* filter comes closest to breaching the threshold in the middle of 2015. However, the case may be made that the 2016 downturn was not the result of unsustainable credit build-up, but rather other factors such as falling equity prices from late 2015 (SARB 2023).

Annex B: The South African financial cycle

The SARB (2023) dates the domestic business, financial and credit cycles over the period 1967–2022. Table B.1 lists the turning points and duration of the financial cycle, which suggests that the average duration of South African financial cycles is around 11 years. Over the same sample, the average duration of the business cycle is about 6 years (SARB 2023). This is in line with the empirical literature, which establishes a South African business cycle duration of 5.8 years on average (Bosch and Koch 2020; Farrell and Kemp 2020).

Table B.1: The South African financial cycle (1970–2022)

Peak	Trough	Duration (downswing)	Duration (upswing)	Duration (p-to-p)	Duration (t-to-t)
<i>Financial cycle</i>					
1972Q2	1977Q4	23	25	48	69
1984Q2	1995Q1	44	6	50	27
1996Q4	2001Q4	21	24	45	46
2008Q1	2013Q2	22	12	34	31
2016Q3	2021Q1	19	-	... ongoing	...
<i>Average:</i>		25.8	16.8	44.2	43.2

Note: Cycle durations are calculated as the number of quarters from peak to peak (p-to-p) and trough to trough (t-to-t), with a full cycle consisting of a downswing/contraction (p-to-t) and an upswing/expansion (t-to-p). Incomplete cycles at the start and end of the sample are not counted.

Source: Own calculations from SARB (2023)

The SARB's (2023) dating of the financial cycle therefore implies that the domestic financial cycle is just about twice as long as the business cycle (11 years relative to 5.8 years). This represents a notable departure from the BCBS assumption that credit and financial cycles are around three to four times longer than business cycles (Drehmann et al. 2010: 28).

Annex C: Sensitivity analyses

Annex C.1: Cycle duration and smoothing parameters

Bosch and Koch (2020) calculate a domestic financial cycle with a duration of about 17 years. Similarly, Farrell and Kemp (2020) calculate a financial cycle duration of 14–15 years. Both studies identify the financial cycle as approximately three times the length of the business cycle, which averages 5.8 years in South Africa.²⁷ Therefore, given an estimated 17-year domestic financial cycle, a λ of 125,000, corresponding to a financial cycle three times longer than the business cycle, might be more appropriate for South Africa.

However, the SARB (2023) has dated the financial cycle somewhat differently, including a short upswing in the mid-1990s, which is absent from earlier empirical work. This results in the average duration of the financial cycle calculated here being significantly shorter than originally estimated in the academic literature (11 years vs 17 years). These different results compared to the SARB's findings may be the result of different dating algorithms or procedures used by the various studies. However, detailed interrogation of these dating algorithms is beyond the scope of this paper, and we choose to use the SARB's own financial cycle dates.

Based on the average length of the domestic credit cycle over our sample of 11 years, (Table B.1), we assign $\lambda = 25\,600$ as a benchmark for our cycle estimates. Figures 3, 4 and 5 illustrate the impact of different smoothing parameters across our various filters. Notably, while the turning points in the credit gap largely overlap, the larger the choice of λ , the smoother the estimated trends, but the more volatile the cyclical components.

The two larger smoothing parameters from the preceding analyses – $\lambda = 125\,000$ and $\lambda = 400\,000$ – suggest qualitatively similar credit gaps, but with divergent magnitudes at turning points. Our preferred smoothing parameter of $\lambda = 25\,600$, consistent with a financial cycle duration of 11 years, exhibits similar characteristics. Predicted turning

²⁷ This aligns with BCBS findings that financial cycles are three to four times longer than business cycles, based on the extrapolation of the business cycle in OECD countries (Drehmann et al. 2010).

points in the financial cycle are consistent between the two higher λ specifications, and with the exception of the GFC, post-GFC recovery and post-COVID-19 recovery, the magnitudes of peaks and troughs are comparable. At the onset of the GFC in 2008, the higher smoothing parameter $\lambda = 400\,000$ yielded a higher peak, while the same filter suggests a lower post-COVID-19 trough and a more sustained and ongoing negative credit gap.

However, setting the smoothing parameter to $\lambda = 1\,600$ yields a very different credit gap under the HP1 and HP1* filters than under the other three specifications, which illustrates the importance of assigning a smoothing parameter consistent with the observed duration of the cycle.²⁸ From 1980 onwards, this gap generally reaches its turning points ahead of the other gaps, while, up until 2010, its amplitude is somewhat smaller. Moreover, this gap drastically underestimates the credit build-up to the GFC, and overestimates the subsequent credit recovery following the GFC as the credit gap does not fall by as much as it does under the one-sided specifications.

Annex C.2: Data smoothing

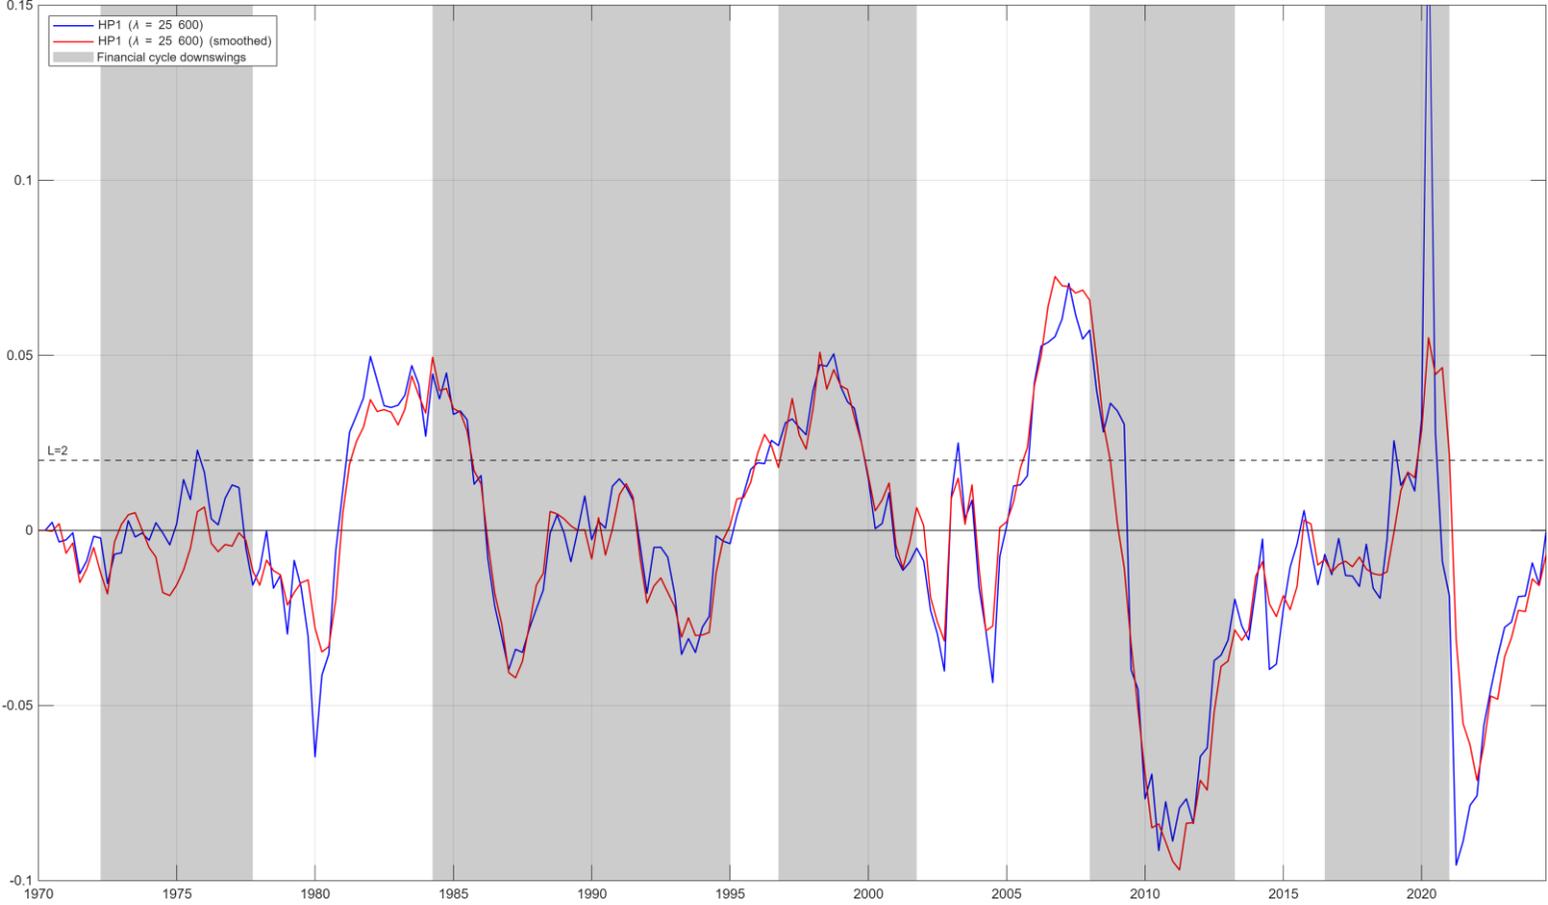
A criticism of the credit-to-GDP gap is the reliance on GDP as the denominator to normalise credit. Significant GDP fluctuations, such as large external shocks or contractions during economic downturns, can cause misleading credit gap signals. Sharp GDP declines can artificially inflate the credit gap even if there is no significant increase in credit, which may trigger inappropriate capital buffer activations. For example, during the COVID-19 pandemic, most countries experienced sharp GDP contractions. The South African economy contracted by an annualised rate of about 51% (Stats SA 2020) in the second quarter of 2020, with similar contractions observed in many other countries. This caused a massive spike in the credit-to-GDP ratio, even though credit extension did not change significantly (Figure 2). This scenario underscores the limitations of the standard credit-gap formula and the need for flexibility in interpreting the gap.

²⁸ Calibrating $\lambda = 1\,600$ implies a cycle duration of 6–8 years, which empirically does not match the South African financial cycle.

Countries adopted various methods to prevent sharp GDP declines from distorting the credit gap. Germany addressed this issue by modifying the GDP used in the credit-to-GDP gap calculation. If GDP experiences a year-on-year decline and the calculated credit gap exceeds the value from the previous quarter, the earlier value is retained. This step reduces false positives and ensures the credit gap reflects credit trends more accurately (Tente et al. 2015). Germany's approach ensured that GDP contractions did not artificially increase the gap, thereby reducing false signals. Switzerland tested smoothing GDP over a five-year moving average to minimise noise. However, this method slightly lowered predictive accuracy (Jokipii, Nyffeler and Riederer 2021). Alternative approaches to modifying the denominator, such as using real GDP per capita or other transformations for Swiss data, performed suboptimally compared to standard measures (Jokipii, Nyffeler and Riederer 2021).

The massive GDP contraction in 2020 resulted in an extremely volatile credit ratio (Figure 2) and subsequent credit gaps (see, for example, Figure 7). This may have contributed to an overshooting of the negative credit gap, with the gap exceeding -10% in early 2021 by some measures. Following the approach of Jokipii, Nyffeler and Riederer (2021), who take moving averages of GDP to smooth out extreme volatility, we construct an augmented credit-to-GDP ratio using a simple four-period moving average of real-time nominal GDP. Figure C.1 illustrates that the smoothing step reduces the volatility in the credit gap, notably during 1980 and the COVID-19 shock, while comparable turning points are predicted. The smoothed credit gap also does not breach the $L = 2$ threshold in 2002, and therefore does not give the false signal provided by the naive gap. However, it still detects the COVID-19 shock as a 'credit event'.

Figure C.1: One-sided HP filters: naive and smoothed



Note: Credit gaps are derived from real-time data. The 'smoothed' gap had underlying GDP smoothed through a four-quarter moving average to ameliorate the effect of outliers.

Source: Own calculations from SARB (2023, 2025) and SARB Core Model

Annex C.3: Forecast data

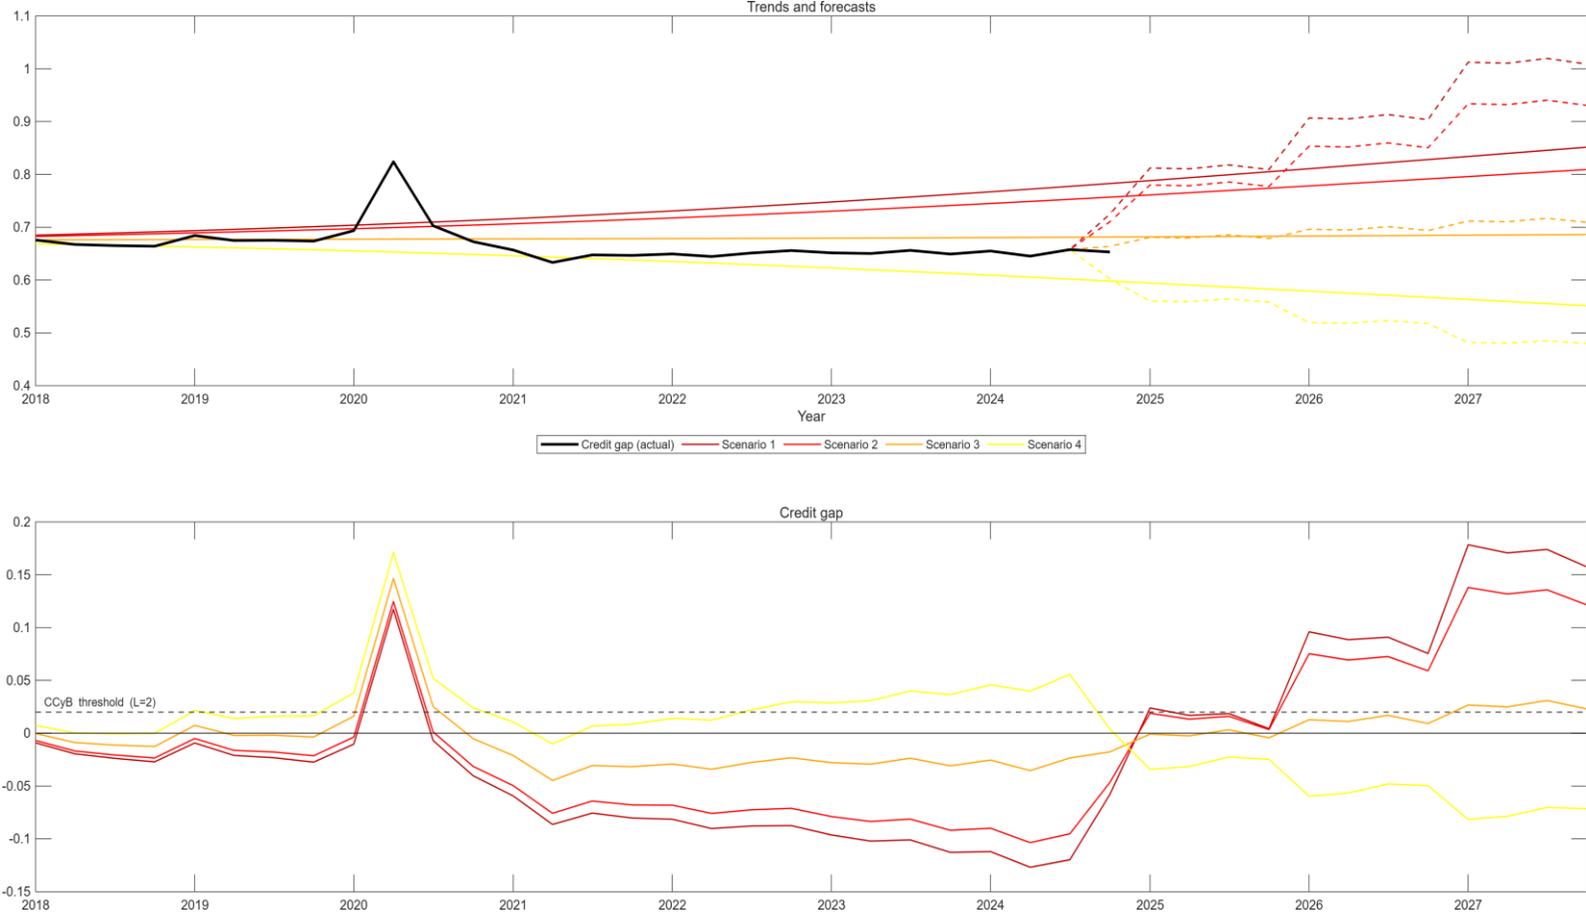
Using forecast data can potentially ameliorate the two-sided filter's end-point problem, while enabling the filter to use data from beyond the current time t to extract the trend. For example, the analysis by Alessandri et al. (2015) for Italy demonstrates that incorporating historical and forecast data improves the performance of the two-sided HP filter relative to other filtering methods. Similarly, Gerdrup, Kvinlog and Schaanning (2013) use a two-sided HP filter, incorporating historical data and a four-quarter moving average of forecasted data, to generate a credit-to-GDP gap for Norway that outperforms the BCBS recommended credit gap derived from the one-sided HP filter.

However, these presumed benefits must be tempered by uncertainty regarding the accuracy and quality of forecast data, as well as the potential for added complexity in the analyses. In South Africa, for example, Farrell (2016) finds that augmenting the data set with forecast data from an AR(4) autoregressive process enhances the stability of the filter, but this yields a credit gap that deviates significantly from actual outcomes. This highlights the importance of forecast quality in generating reliable credit gaps that can be used to predict crises. As Farrell (2016) demonstrates, the quality of forecasts has significant implications for the reliability of subsequent credit gap estimates. Therefore, in the absence of reliable forecasts, the two-sided filter may be more suited to ex post analyses.

Finally, we tested several forecast scenarios and their implications for the two-sided HP filter. Figure C.2 illustrates four simple forecast scenarios, with combinations of low, steady and high growth in nominal GDP and credit extension.²⁹ We used actual data until the third quarter of 2024, after which we used the different scenarios to forecast through the end of 2027. The darker lines represent higher credit growth relative to GDP, with the effect of increasing the credit-to-GDP ratio over time. Conversely, the yellow line represents a fall in the credit-to-GDP ratio, due to strong GDP growth relative to credit, while the orange line represents a middle ground of low GDP and moderate credit growth.

²⁹ Scenarios 1 through 4 represent forecasts of, respectively, low GDP and high credit growth, moderate GDP and high credit growth, low GDP and moderate credit growth, and high GDP and low credit growth.

Figure C.2: Forecast sensitivity: two-sided HP filters



Note: Credit gaps are derived from real-time data. The 'smoothed' gap had underlying GDP smoothed through a four-quarter moving average to ameliorate the effect of outliers.
 Source: Own calculations from SARB (2023, 2025) and SARB Core Model

It is immediately evident that the different forecasts imply a different current path for the credit gap, and even revise credit gap estimates backwards. This is evidence of the HP2 filter's end-point problem, where the addition of new data points influences the estimates of earlier trends and cycles.

The two-sided HP filter tends to try to close gaps near the end of the sample (Chen and Górnicka 2020; Anvari, Ehlers and Steinbach 2014). For this reason, perhaps, the filter tries to balance, for example, the forecasted fall in the credit gap under Scenario 4 with a mechanical upward revision of the credit gap during 2024. Concerningly though, this could trigger the activation of higher capital buffer requirements during 2023 and 2024 since the credit gap would be beyond the 2% threshold, which might perversely lead to exactly the kind of slowdown in the credit-to-GDP ratio as predicted by this forecast scenario. Similarly, there might be a self-fulfilling prophecy in the high credit growth scenarios, where the contemporaneous credit gap appears strongly negative, which might cause accompanying benign regulation to contribute to an undesirable credit boom.

References

Alessandri, P, Bologna, P, Fiori, R and Sette, E. 2015. 'A note on the implementation of the countercyclical capital buffer in Italy'. Bank of Italy Occasional Paper No. 278.

Anvari, V, Ehlers, N and Steinbach, R. 2014. 'A semi-structural approach to estimate South Africa's potential output'. *South African Reserve Bank Working Paper Series*, WP/14/08.

Basel Committee on Banking Supervision (BCBS). 2010. 'Guidance for national authorities operating the countercyclical capital buffer'. Bank for International Settlements.

BCBS. 2024. 'Implementation: range of practices in implementing a positive neutral countercyclical capital buffer'. Bank for International Settlements.

Baxter, M and King, R G. 1999. 'Measuring business cycles: approximate band-pass filters for economic time series'. *Review of Economics and Statistics* 81(4): 575–593.

Beveridge, S and Nelson, C R. 1981. 'A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the "business cycle"'. *Journal of Monetary Economics* 7(2): 151–174.

Bosch, A and Koch, S F. 2020. 'The South African financial cycle and its relation to household deleveraging'. *South African Journal of Economics* 88(2): 145–173.

Burns, A F and Mitchell, W C. 1946. *Measuring business cycles*. New York: National Bureau of Economic Research.

Chen, J and Górnicka, L. 2020. 'Measuring output gap: is it worth your time?'. IMF Working Paper, WP/20/24.

Christiano, L J and Fitzgerald, T J. 2003. 'The band pass filter'. *International Economic Review* 44(2): 435–465.

Drehmann, M, Borio, C, Gambacorta, L, Jimenez, G and Trucharte, C. 2010. 'Countercyclical capital buffers: exploring options'. *BIS Working Papers*, No. 317.

Drehmann, M, Borio, C and Tsatsaronis, K. 2011. 'Anchoring countercyclical capital buffers: the role of credit aggregates'. *BIS Working Papers*, No. 355.

Drehmann, M and Juselius, M. 2014. 'Evaluating early warning indicators of banking crises: satisfying policy requirements'. *International Journal of Forecasting* 30(3): 759–780.

Drehmann, M and Tsatsaronis, K. 2014. 'The credit-to-GDP gap and countercyclical capital buffers: questions and answers'. *BIS Quarterly Review*, March 2014, 55–73.

Drehmann, M and Yetman, J. 2018. 'Why you should use the Hodrick-Prescott filter—at least to generate credit gaps'. *BIS Working Papers*, No. 744.

Farrell, G. 2016. 'Countercyclical capital buffers and real-time credit-to-GDP gap estimates: a South African perspective'. *Journal for Studies in Economics and Econometrics* 40(1): 1–20.

Farrell, G and Kemp, E. 2020. 'Measuring the financial cycle in South Africa'. *South African Journal of Economics* 88(2): 123–144.

Gerdrup, K R, Kvinlog, A B and Schaanning, E. 2013. 'Key indicators for a countercyclical capital buffer in Norway – trends and uncertainty'. Norges Bank Staff Memo, No. 13/2013.

Giese, J, Andersen, H, Bush, O, Castro, C, Farag, M and Kapadia, S. 2014. 'The credit-to-GDP gap and complementary indicators for macroprudential policy: evidence from the UK'. *International Journal of Finance & Economics* 19(1): 25–47.

Hamilton, J D. 2018. 'Why you should never use the Hodrick-Prescott filter'. *Review of Economics and Statistics* 100(5): 831–843.

Hodrick, R J and Prescott, E C. 1997. 'Postwar US business cycles: an empirical investigation'. *Journal of Money, Credit, and Banking* 29(1): 1–16.

Jokipii, T, Nyffeler, R and Riederer, S. 2021. 'Exploring BIS credit-to-GDP gap critiques: the Swiss case'. *Swiss Journal of Economics and Statistics* 157(7): 1–19.

Jordà, Ò, Schularick, M and Taylor, A M. 2013. 'When credit bites back'. *Journal of Money, Credit and Banking* 45(s2): 3–28.

Orphanides, A and van Norden, S. 2002. 'The unreliability of output-gap estimates in real time'. *Review of Economics and Statistics* 84(4): 569–583.

Ravn, M O and Uhlig, H. 2002. 'On adjusting the Hodrick-Prescott filter for the frequency of observations'. *Review of Economics and Statistics* 84(2): 371–376.

South African Reserve Bank (SARB). 2023. *Financial stability review*. First edition. <https://www.resbank.co.za/en/home/publications/publication-detail-pages/reviews/finstab-review/2023/first-edition-2023-financial-stability-review>

SARB. 2025. 'Quarterly bulletin'. Various editions. <https://www.resbank.co.za/en/home/what-we-do/statistics/releases/online-statistical-query>

Statistics South Africa (Stats SA). 2020. 'Gross domestic product: second quarter 2020'. Statistical Release P0441.

Tente, N, Stein, I, Silbermann, L and Deckers, T. 2015. 'The countercyclical capital buffer in Germany: analytical framework for the assessment of an appropriate domestic buffer rate'. Deutsche Bundesbank.

Van Vuuren, G. 2012. 'Basel III countercyclical capital rules: implications for South Africa'. *South African Journal of Economic and Management Sciences* 15(3): 309–324.

Wolf, E, Mokinski, F and Schüler, Y S. 2020. 'On adjusting the one-sided Hodrick-Prescott filter'. *Deutsche Bundesbank Discussion Papers*, No. 11/2020.

Wolf, E, Mokinski, F and Schüler, Y S. 2024. 'On adjusting the one-sided Hodrick-Prescott filter'. *Journal of Money, Credit and Banking*.

<https://doi.org/10.1111/jmcb.13240>.