TOPICAL BRIEFING

TB/05/2025

Financial Stability Department

SOUTH AFRICAN RESERVE BANK

File ref. no.: 9/2/3

June 2025

Wongani Msiska
Department of Economics, University of Stellenbosch
Ayanda Sikhosana
Financial Stability Department, South African Reserve Bank
Cobus Vermeulen[1]
Department of Economics, University of South Africa, and Research Fellow, South African Reserve Bank

Authorised for distribution by: Dr Nicola Brink
Nicola.Brink@resbank.co.za

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

Abstract

Financial crises like the global financial crisis are rare, but often lead to bad economic outcomes. These crises usually result from excessive risk-taking during prolonged periods of growth accompanied by lax regulation. When credit grows too quickly, it can create systemic risks in the financial sector. Financial institutions tend to lend excessively during economic booms and then restrict lending during downturns, which can worsen economic recessions. To address this, the Basel Committee on Banking Supervision (BCBS) suggests using the credit-to-GDP gap as a tool to identify rising credit risks and systemic instability. When this gap exceeds certain thresholds, the countercyclical capital buffer (CCyB) can be used to strengthen the financial system against potential downturns. This paper evaluates different methods for measuring the credit-to-GDP gap, as one early warning indicator (EWI) in a suite of macroprudential policy tools, to find the most accurate real-time indicator for South Africa. The study concludes that the one-sided Hodrick-Prescott (HP) filter produces a reliable real-time

[1] Corresponding author: vermejc@unisa.ac.za.

estimate of the credit-to-GDP gap, accurately predicting 75% of financial cycle turning points, and aligning well with full-sample estimates of the credit-to-GDP gap.

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

JEL codes: C22; E32; E37; E44; G01; G28

1. Introduction

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

The BCBS further recommends using the credit-to-GDP gap, measured as the deviation of the credit-to-GDP ratio from its long run trend, as an EWI of the build-up of systemic risk. When this gap exceeds certain thresholds, the CCyB can be deployed to increase capital buffers, tempering credit growth and preventing further build-up of financial imbalances, and 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 EWI for periods of financial distress. It compares different filtering approaches to determine which provides the most accurate and reliable real-time credit gap identification for the domestic context. This is crucial for informing the South African Reserve Bank's (SARB) decisions on the CCyB, helping to pre-empt financial instability and mitigate systemic risks. The findings will contribute to refining the calculation of the domestic credit gap, ensuring it better serves South Africa's economic conditions and promoting a more resilient financial system.

The paper is structured as follows: Section 2 provides a brief background on the BCBS's guidance and recommendations around estimating the credit-to-GDP gap and how it informs the CCyB, as well as the South African experience to date. 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, while we also explore the implications for activating the CCyB based on different threshold calibrations. Section 5 concludes.

2. Background and literature review

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 or limit the banking system from amplifying economic fluctuations (Drehmann, Borio, Gambacorta, Jimenez, and Trucharte, 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, domestic and foreign sources, and includes instruments such as bonds and cross-border finance (BCBS, 2010). The BCBS guidance emphasises the superiority of the total credit measure over narrower alternatives, as supported by Drehmann and Tsatsaronis (2014), who demonstrate that credit gaps based on total credit are more effective EWIs of banking crises than narrower measures such as those based on bank credit alone. The total credit measure also ensures that systemic risk build-up in all parts of the economy is accounted for, and it mitigates the incentives for banks to divert credit supply to other parts of the financial system as would 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, $\mathrm{r}_t = \mathrm{credit}_t/\mathrm{GDP}_t$. The credit-to-GDP gap is simply the cyclical component of the credit-to-GDP ratio r_t, and is obtained by subtracting the long-term trend component of this ratio from its actual value.

The BCBS recommends using a one-sided HP filter to extract the trend. The use of a one-sided HP filter to calculate 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 is not available at the time of the CCyB decision and, therefore, is not suitable for real-time policymaking.[2]

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 EWI, the credit-to-GDP gap has some limitations[3] (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 upon which to base 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. The countercyclical capital buffer

Having a core set of indicators (inclusive of the credit gap) allows policymakers to comprehensively assess the risk environment to determine when cyclical risks are low, beginning to rise, and when they are high. 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, 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 when the gap exceeds an upper threshold ($H = 10$). Between these thresholds, the CCyB increases linearly. For example, if the maximum buffer is set at 2.5%, the capital buffer add-on will be 0% when the gap is below $L = 2$, 2.5% when the gap matches or exceeds $H = 10$, and vary proportionally between 0% and 2.5% for gaps between $L = 2$ and $H = 10$.

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

2 The question of data requirements and its real-time availability, and its implications for the choice of filter, is central to this study, and is considered in depth in Section 3.

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

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 is aimed at ensuring that threshold breaches occur at least 2-3 years before a crisis or period of distress, allowing banks sufficient time to adjust their capital levels. At the same time, the lower threshold must be high enough to avoid the 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. This balance ensures the thresholds remain effective across diverse jurisdictions.[4]

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, see Section 2.2.2) 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).

4 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

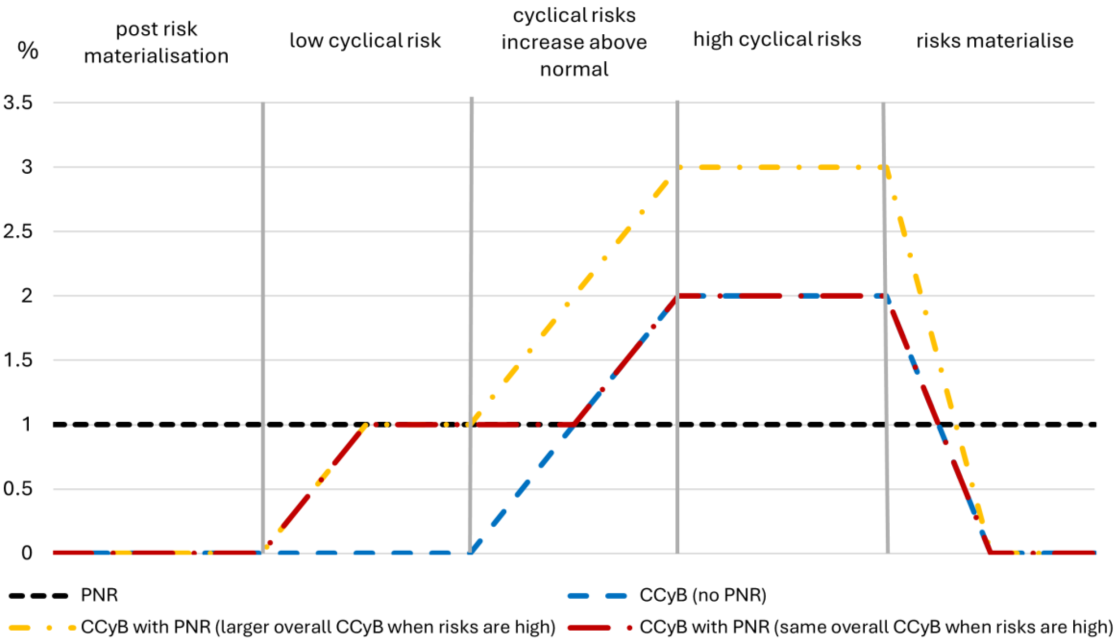

Source: Adapted from BCBS (2024)

In this case, the PCN CCyB can be seen as an early activation of the CCyB when risks increase above normal and does not increase the overall CCyB when risks are elevated. For a jurisdiction like South Africa with a PCN CCyB of 1% and a maximum level of the CCyB of 2.5% this means when cyclical risks increase above normal, banks will only be required to hold additional capital when the credit gap exceeds 6 percentage points, which translates to a CCyB of 1.25%.

For other jurisdictions, the PCN CCyB has been designed to primarily focus on covering risks unrelated to the domestic financial cycle (i.e. the orange dash-dotted line). In this case, when cyclical risks begin to rise, the CCyB increases from the level of the PCN CCyB (for example 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.

2.2.1. Implementing the CCyB in South Africa

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.[5] 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 adjustment for domestic application. This finding underscores the importance of further threshold analysis and highlights 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.

While the thresholds of L = 2 and H = 10 provide a globally validated starting point, local testing and refinement are essential to tailor these values to South Africa's specific financial stability needs. To determine whether these thresholds are optimal for South Africa, several steps could be undertaken. First, the trade-offs between different threshold combinations should be evaluated, analysing the balance between Type I and Type II errors in relation to historical banking crises or periods of financial distress. Secondly, scenario analysis should examine how these thresholds align with South Africa's economic environment, considering the duration of local financial cycles and consequently the choice of smoothing parameter (λ).

The implementation of the CCyB in South Africa highlights several practical considerations, particularly in balancing its activation frequency and potential influence on the economy. A critical discussion involves the thresholds for activation, such as whether a higher lower threshold (e.g. L = 3 or L = 4) should be considered to avoid frequent activations, which could impose operational challenges or costs on banks.

2.2.2. The Positive Cycle-Neutral CCyB

The phased introduction of the PCN CCyB from January 2025, and its planned 1% requirement starting in January 2026, demonstrates an effort to integrate the CCyB at a non-zero rate into the existing capital framework. Two options were considered:

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

substituting the PCN CCyB for other requirements (e.g. Pillar 2A) or adding it to the existing capital structure. The SARB's 2024 *Financial Stability Review* details that the latter option was chosen, prioritising resilience by increasing available funds to absorb losses during crises. However, to mitigate potential adverse effects on lending and the real economy, an extended phase-in period was adopted.

Since the PCN CCyB will be set at 1% in January 2026 while the credit gap is expected to remain below the threshold of L = 2, assessing how this will affect CCyB activation is important. In practice, the effective lower threshold for the credit gap may be closer to 6 percentage points (the midpoint), as this is the level at which the CCyB would need to increase above the neutral 1% rate. However, as mentioned above, having a set of indicators for assessing when cyclical risks increase above normal, as opposed to relying solely on the credit gap, would allow the policymaker to adjust the CCyB above the neutral rate of 1% much earlier. Therefore, with the CCyB already at 1%, when there are other indicators in addition to the credit gap the policymaker does not have to wait until the credit gap reaches 6 percentage points before adjusting the CCyB above the neutral rate of 1%.

International experience and feedback also informed the implementation strategy. For example, during the COVID-19 pandemic, jurisdictions with active CCyB's lowered their buffer rates to mitigate cyclical effects, a flexibility unavailable to South Africa due to its CCyB rate being set at 0% at the time. This underscored the importance of maintaining a capital buffer to guard against systemic risks arising from adverse shocks. While concerns persist about the CCyB's impact on lending volumes, evidence suggests negligible effects on lending rates and no clear link between higher capital requirements and negative economic growth (BCBS, 2010). Recent literature on lending patterns with regards to the activation of the CCyB demonstrates that while lending rates generally remain unaffected, lending volumes can be influenced, particularly in specific market segments or among less capitalised banks. Studies conducted in Europe during the COVID-19 crisis, for example, show that the release of CCyB's effectively supported mortgage lending without placing upward pressure on interest rates (Dursun-de Neef, Schandlbauer and Wittig, 2023). Conversely, research from Belgium indicates that CCyB activation can influence mortgage rates and house prices, but the broader credit supply effects remain limited (Damen and Schildermans, 2022). Evidence from emerging markets similarly point to modest and temporary impacts on lending volumes, with weaker banks more sensitive to buffer changes (Fang, Jutrsa, Peria, Presbitero, and Ratnovski, 2022). In summary, while early assessments like BCBS (2010) suggested minimal impact of CCyB's on lending and growth, more recent studies highlight that the effects of CCyB's can vary depending on factors such as bank capitalisation, economic conditions, and the specific design of the buffer release. That said, the overall impact of CCyB's on credit supply and macroeconomic outcomes remains inconclusive in the literature.

Finally, the decision to adopt the PCN CCyB as an addition to the existing capital framework aligns with international best practices and aims to improve South Africa's capital adequacy ratios. By setting clear limits for Pillar 2A and D-SIB requirements, the framework seeks to balance maintaining resilience and supporting economic growth. Addressing these practicalities is essential to optimising the CCyB's effectiveness as a macroprudential tool while ensuring its activation aligns with the broader goals of financial stability and economic sustainability.

The SARB has taken a cautious approach to implementing the PCN CCyB, incorporating feedback from banks and stakeholders. Concerns raised include the potential redirection of the systemic risk Pillar 2A requirement towards the CCyB, extended transition periods above the recommended 12 months, and the recent negative trend in the domestic credit gap, which could make frequent activations impractical (SARB, 2024).

2.3. The international experience

Given the varying economic conditions between countries, it is important to explore potential country-specific adjustments to the credit-to-GDP gap calculation. Countries have therefore adapted the BCBS's broad definition to align with local data availability and country-specific factors.

As suggested by the BCBS (2017), the credit gap should be calculated using a broad measure of nominal credit that encompasses all sources of debt funding to the private non-financial sector, including those from both bank and non-bank sectors, as well as from domestic and foreign sources. A notable adaptation, however, is the use of narrow credit measures, which focus primarily on credit from banking institutions. This narrower approach can offer advantages such as better timeliness and alignment with domestic financial conditions. For instance, Denmark relies on monthly data from bank and mortgage institutions for its initial credit-to-GDP gap calculations due to delays in compiling broader financial account statistics (DNB, 2016). Similarly, countries such as Australia, Belgium, and Switzerland use domestic banking system credit in their CCyB frameworks, citing its relevance to their financial environments and its superior performance in signalling systemic risks (APRA, 2017). Other countries such as France further refine their credit measures to exclude distortions created by intragroup credit within non-financial corporations to prevent organisational shifts from artificially affecting credit gap calculations (BCBS, 2017). In Germany, Tente, Stein, Silbermann, and Deckers (2015) assessed credit measures such as total debt, outstanding amount of bank loans and bank credit[6] and found that bank credit provided a superior EWI and fewer misleading signals. Therefore, Germany uses bank credit for the national buffer guide analysis.

3. Data and methodology

3.1. Data

In South Africa, the credit-to-GDP gap is calculated using credit extended to the domestic private non-financial sector by monetary institutions.[7] While closely following the BCBS's recommendations, this narrower measure prioritises efficacy and data availability but excludes credit extended by non-bank financial institutions (NBFIs), foreign sources and the informal financial sector, potentially incentivising banks to redirect credit outside the monitored banking system to avoid CCyB adjustments.

3.1.1. Data sources

Data utilised in this paper include total monthly nominal credit extended (KBP1347M) and quarterly nominal GDP (KBP6006L) over the sample period 1970Q1 to 2024Q3 from the SARB Quarterly Bulletin (QB) (SARB, 2025).[8] In each quarter the final month's observation of credit data (i.e. March, June, September and December) is taken to

[6] Measured as the sum of changes in the outstanding amount of bank loans adjusted for statistical changes.

[7] Monetary institutions include the South African Reserve Bank, the former National Finance Corporation, Corporation for Public Deposits (CPD) and the so-called 'pooled funds' of the former Public Debt Commissioners, the Land Bank, Postbank, private banking institutions (including the former banks, discount houses, equity building societies and mutual building societies).

[8] The credit aggregate is largely in line with the BCBS's recommendations; however, because bank credit comprises around 97% of total nominal credit extended, there is not much value in exploring a narrower bank-credit-only measure of credit extension.

represent credit extension for that quarter. However, these data represent final revised data, and are therefore not that useful in deriving true real-time estimates of the credit gap. This is important, as GDP data are prone to updates and revision up to several years after its first publication; therefore, final GDP data can differ significantly from what it appeared to be in 'real-time'. However, final data can 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.

3.1.2. Real-time data

We complement our full-sample final-data analysis with real-time measures of the credit-to-GDP ratio, utilising real-time data vintages sourced from the SARB's historical Core and Quarterly Projection Models going back to September 2000.[9] These vintages represent original observations of GDP and credit extension policy makers had at hand at the time; combining these vintages therefore yields the closest to a real-time series of the 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 closest to the full-sample credit gap estimates.

The influence of data revisions is significant. Figure 2 illustrates that final revised nominal GDP numbers are generally higher than its real-time counterparts,[10] while data on credit extension remain virtually untouched. For this reason, the credit-to-GDP ratio appears significantly higher in real-time than *ex post*. It follows that trends extracted in real-time will necessarily be higher than trends extracted from final data in this context; however, this does not imply that the credit gap will necessarily be bigger since the impact on the cyclical component (i.e. the difference between actual observations and its trend) is less clear. Finally, it is 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 need to be derived from data which are bound to change introduces an additional layer of uncertainty in the identification process.

---

9 We are grateful to MG Ferreira and Rowan Walter for making these data vintages available.

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

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

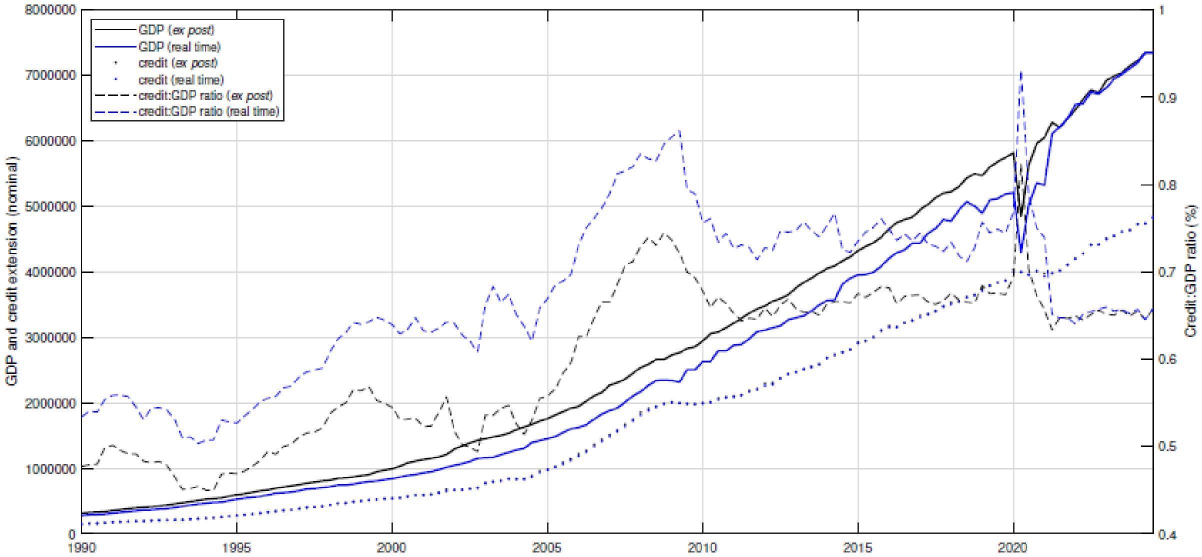

Sources: Own calculations from SARB (2025) and SARB Core Model

3.2. The South African Financial cycle

3.2.1. Dating the financial cycle

The SARB (2023b) dates the domestic business, financial, and credit cycles over the period 1967–2022. Table 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).

The SARB's (2023b) 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), and has important implications for the design of the credit gap estimator as will be discussed below.

Figure 3 presents the real-time and *ex post* credit-to-GDP ratios relative to the SARB's financial cycle. It is immediately evident that upward phases of the financial cycle are associated with an increase in the credit-to-GDP ratio, while downswings in the financial cycle are associated with a stagnant or falling credit-to-GDP ratio. Figure 3 also confirms the significant impact revisions to GDP data have on the calculated credit-to-GDP ratio alluded to above, as the real-time estimated credit ratio is permanently between 5 and 10 percentage points higher than the final ratio.[11]

Table 1: The South African financial cycle (1970-2022)

Peak	Trough	Duration (downswing)	Duration (upswing)	Duration (p-to-p)	Duration (t-to-t)
Financial cycle					
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 ...	
Average:		25.8	16.8	44.2	43.2

Source: Own calculations from SARB (2023b). 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.

3.2.2. Data volatility and revision

One of the main criticisms 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 approximately 51% (Stats SA, 2020) in the second quarter of 2020, with similar contractions observed in many other jurisdictions. 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

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

and the need for flexibility in interpreting the gap.

Countries adopted various methods to prevent sharp GDP declines from distorting the credit gap. Germany addressed this issue by modifying GDP used in the credit-to-GDP gap calculation. That is, 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, reducing false signals. Some jurisdictions, like 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 highlighted suboptimal performance compared to standard measures (Jokipii, Nyffeler and Riederer, 2021).

Figure 3: The South African financial cycle and the credit-to-GDP ratio

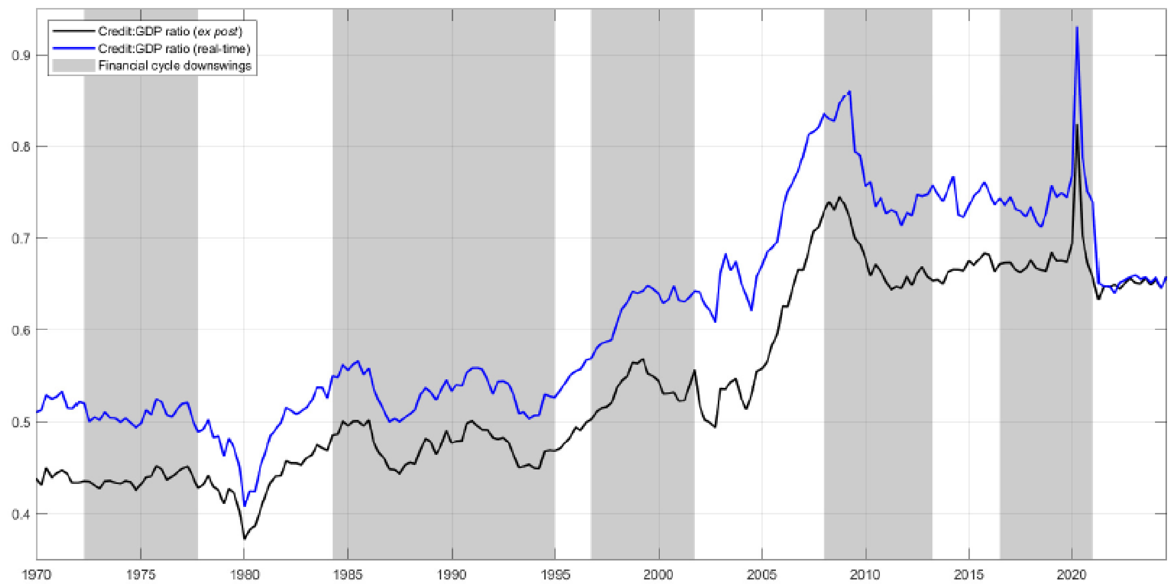

Sources: Credit cycle (SARB, 2023); credit-to-GDP ratio (own calculations from SARB, 2025 and SARB Core Model)

Furthermore, as was already alluded to, revisions to GDP data can have a significant influence on calculations of the credit-to-GDP ratio (Figure 2), which is in turn used to estimate the credit gap. However, in the South African context, data revisions seem to pose a smaller problem than the availability of new data, with the latter having a more significant impact on real-time credit gap estimates. Farrell (2016) shows that new data can substantially alter earlier estimates of the credit-to-GDP gap, raising concerns about the credit gap's reliability for policymaking.[12]

Lessons from Germany and the COVID-19 experience can be applied to enhance the reliability of the credit gap. The SARB could exercise discretion, using additional indicators such as real credit growth or financial deepening metrics to supplement gap analysis. Implementing adjustments similar to Germany's buffer guide modification could also help mitigate the impact of shocks to GDP. The BCBS also emphasises the principle of guided discretion in interpreting credit gap changes influenced by shocks to GDP. Noise tends to occur during recessions, where GDP collapses, rather than credit surges, drive the credit gap. While the mechanical use of the credit gap is discouraged, national authorities have discretion to adjust signals based on supplementary indicators or local economic contexts.

3.2.3. Recent domestic financial cycle trends

The domestic financial cycle has been in an upward phase since the second quarter of 2021 (Table 1, SARB, 2023b). However, recent analysis by the SARB indicates that the financial cycle has started trending downwards since the third quarter of 2023, with equity and house prices having started to trend downwards around the same period, while credit only started trending downwards in the third quarter of 2024. SARB (2023b) notes that credit growth, particularly for unsecured credit,[13] had remained elevated for both households and corporates due to the rising cost of living and high input costs associated with longer and more frequent stages of load shedding.

[12] Farrell (2016)'s small mean absolute revision between real-time and quasi-real gaps (1.36 percentage points) shows that data revisions account for a modest part of differences in the credit gap, while the larger revision between final and real-time gaps (4.3 percentage points) highlights the significant role of new data and end-of-sample adjustments.

[13] Unsecured credit comprises overdrafts, credit card advances and a portion of general loans and advances. In contrast, secured credit comprises mortgage advances, instalment sale credit and leasing finance, which are typically backed by an asset.

Figure 4: Credit growth

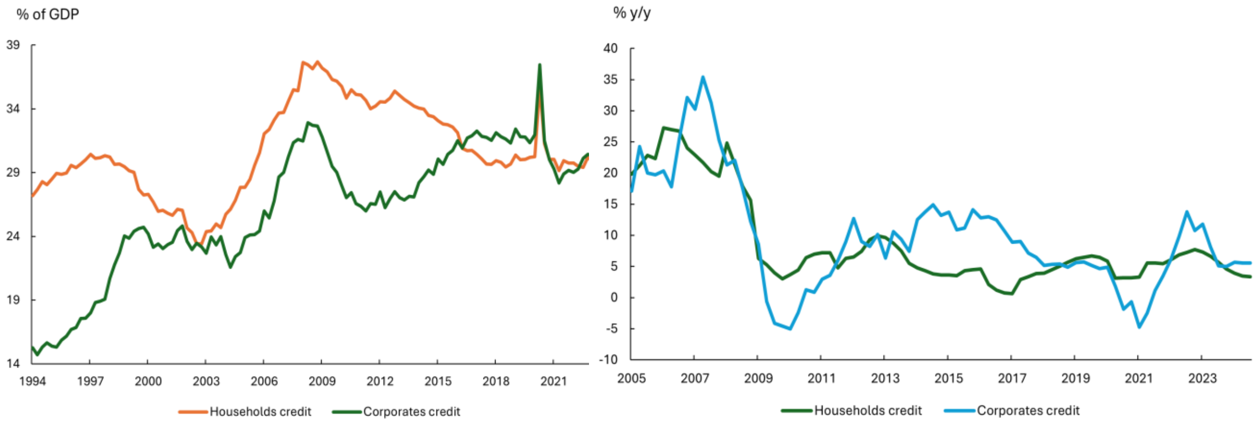

Source: SARB

As a percent of GDP, credit to corporates began rising in 2011, stabilised between 2017 and 2019 before declining sharply following the COVID-19 shock. However, since the second half of 2021, credit to corporates has resumed rising while adjusted credit to non-financial corporates[14] has declined since late 2020 (Singh, 2025). In contrast, credit to households declined from 2008 to 2018 and stabilised thereafter. Consequently, while the credit-to-GDP gap[15] has remained negative for some time, it has been narrowing since the second half of 2022 (SARB, 2023). However, on a year-on-year basis, growth in both credit to the domestic private sector and GDP has been falling since the fourth quarter of 2022, with the latter falling faster. Both have stabilised since the third quarter of 2023. Growth in credit to corporates has followed a similar trend while growth in credit to households has continued to decline since the first quarter of 2023. Growth in private consumption expenditure has also declined since its peak in the second quarter of 2021. Corporates typically borrow to fund investment, which can be seen by the co-movement between growth in investment and growth in general loans and advances, which make up more than half of total loans and advances (SARB, 2023a). Year-on-year growth in both general loans and advances as well as investment by private firms has been declining since the third quarter of 2022. Singh (2025) also notes that investment by private firms (in real rand terms) has reached levels recorded during the GFC and has not yet recovered to pre-COVID-19 levels.

---

14 Data is from the Bank for International Settlements and is adjusted for structural breaks.

15 Calculated based on credit to the domestic private sector (i.e. households and corporates), which includes financial corporates and is not adjusted for structural breaks.

Figure 5: Credit, expenditure and private investment

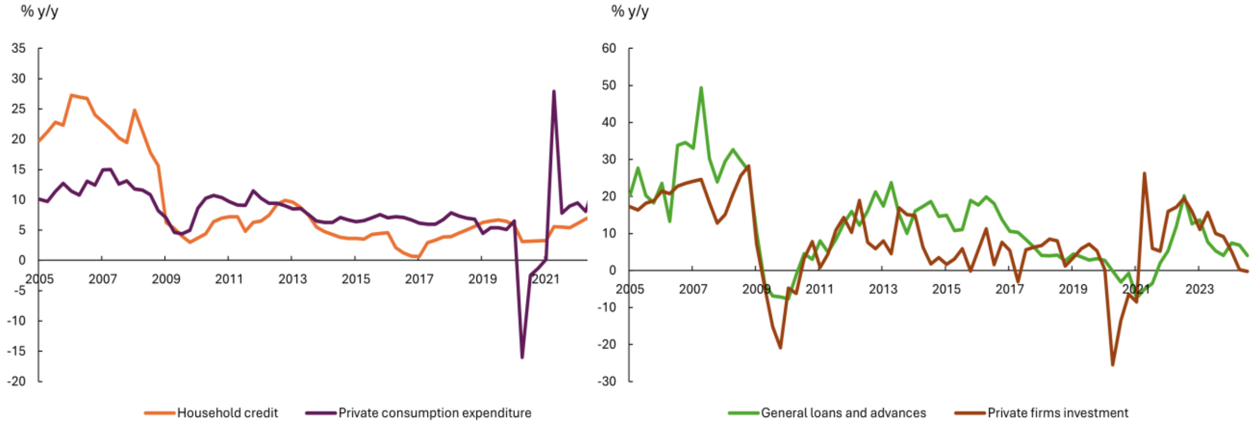

Sources: Haver and SARB

In summary, the above highlight the importance of considering a set of indicators to inform the decision on the CCyB. That is, depending on the approach adopted to estimate the credit-to-GDP gap, it may not always align with other indicators of macro-financial conditions. For example, based on the selected HP filter and smoothing parameter, which are arguably based on purely statistical considerations, this may result in a credit gap that ignores other prevailing macro-financial conditions and may potentially lead to unwarranted activation of the CCyB.

3.3. 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 3. 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 cyclical component of the credit-to-GDP series.[16]

There are two dimensions to extracting the trend and cycle from the credit-to-GDP ratio. First, a suitable filtering procedure needs to be selected, and, second, the filter should be calibrated appropriately. The choice of filter will depend mainly on data availability and the estimation window, while the calibration will depend on estimates of the

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

duration of the domestic financial cycle.

This paper focuses on 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 by no means the only way to extract trends and cycles from macroeconomic data[17] they are central to the BCBS's approach to estimating the credit gap, and is therefore a natural starting point to this analysis. The filters that are utilised here include 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 Appendix A.

3.3.1. 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:

$$L = \sum_{t=1}^{T}(y_t - \tau_t)^2 + \lambda\sum_{t=1}^{T}[(\tau_t - \tau_{t-1}) - (\tau_{t-1} - \tau_{t-2})]^2 \qquad (3.1)$$

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.[18]

The two-sided HP filter:

The conventional two-sided HP filter is widely utilised to extract trends from macroeconomic time series, as it is computationally easy and not subject to rigorous

[17] 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.

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

data requirements. However, the HP filter may result in spurious dynamics that are not found in the underlying data (Hamilton, 2018), while Orphanides and Van Norden (2002:569) laments its "pervasive unreliability of 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 datapoints 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 the use of a one-sided HP filter to extract the trend from the credit-to-GDP ratio. The key difference between the one- and two-sided HP filters is that the one-sided filter uses only past and current observations to estimate the trend at each point in time t[19] – 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 policy makers "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.[20] 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.

---

[19] 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*.

[20] While the BCBS recommends using at least 10 years of data to mitigate these issues, evidence from Switzerland (Jokipii, Nyffeler and Riederer, 2021) suggests this threshold is insufficient, as convergence in credit gap estimates often requires datasets spanning up to 30 years.

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 fulfills its primary purpose of serving as an EWI for financial or banking crises. Policymakers rely on real-time data, not *ex post* adjustments, and 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). Moreover, in many empirical applications the real-time credit gap estimates from the one-sided HP filter align closely with those from the *ex post* two-sided HP filter, validating the one-sided filter's adequacy for real-time decision-making when evaluated against the stability of 'final' estimates from the two-sided filter.

Adjusted one-sided HP filter:

Van Vuuren (2012:316) suggests that well-known problems of the two-sided filter can be "alleviated by implementing a single-sided filter". However, Wolf, Mokinski and Schüler (2024:2) contend that the one-sided HP filter "fails to eliminate low-frequency fluctuations to the same extent" than 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 dataset at once. It is more resilient against 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.[21]

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.

Alternative univariate filters and model-based approaches:

Other well-known univariate statistical filters include the Beveridge and Nelson (1981), Baxter and King (1999) and Christiano and Fitzgerald (2003) filters. However, univariate filters are all vulnerable to structural breaks, as they attempt to smooth over these events, reducing the reliability of its outputs (Jokipii, Nyffeler and Riederer, 2021). The BCBS (2010) noted that the HP filter was preferred over simple moving averages or linear trends as it gives more weight to more recent data points, enabling it to better accommodate structural breaks. However, subsequent research from Drehmann and Tsatsaronis (2014) demonstrated the long-term impact of structural breaks in skewing the credit-to-GDP ratio, resulting in unreliable patterns in the credit gap. These findings prompted the BIS to recommend the use of break-adjusted data for credit gap calculations.

Multivariate filters, by imposing economic structure and introducing additional data, may enhance the reliability of estimates in the presence of structural breaks. In Switzerland, Jokipii, Nyffeler and Riederer (2021) evaluate trend extraction alternatives such as moving averages and linear time trends, but find that these methods add complexity without significantly improving the baseline credit gap recommended by the BCBS. They advocate for model-based approaches to incorporate economic fundamentals into credit equilibrium levels, addressing a key limitation of the HP filter as a purely statistical measure. Drehmann, Borio and Tsatsaronis (2011) examine the long-run economic relationship between credit, interest rates and other macroeconomic variables using a vector error correction model (VECM) in the exploratory phase of identifying an indicator for the CCyB. However, these models introduce additional complexities around model structure and data requirements, which might not

[21] We utilise 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.

necessarily be suitable for real-time estimations. Therefore, the remainder of this analysis will focus on variants of the HP filter as a first step to pinning down a reliable real-time credit gap indicator.

3.3.2. Calibrating the smoothing parameter (λ)

Hodrick and Prescott (1997) proposed a smoothing parameter of λ = 1,600 for their original minimisation problem (eq. 3.1), based on their observations of post-war US business cycles averaging around 8 years. However, while business cycles typically average between 6 to 10 years, financial cycles usually span longer. 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

$$\lambda = \lambda_0 \left(\frac{f}{f_0}\right)^4 \tag{3.2}$$

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 financial cycles are typically 3 to 4 times longer[22] 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, Table 1), we propose setting $\lambda = 1{,}600 \times 2^4 = 25{,}600$ for our credit gap estimates. Testing the relative performance of these alternative smoothing parameters in the South African context could provide more tailored insights into the credit-to-GDP gap's reliability as an EWI of domestic systemic risk and may better reflect the country's

[22] This calculation for quarterly data yields $1{,}600 \times 4^4 = 409{,}600$, or approximately $400{,}000$.

financial dynamics. Various parameterisations of λ are duly analysed in Appendix B below.

3.4. 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 Area Under the Receiver Operating Characteristic Curve (AUROC), which evaluates the trade-offs between Type I (missed crises) and Type II (false crises) 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 unsuitable for single country analysis, especially given the relatively limited number of crisis periods in South Africa. We rely on Farrell and Kemp (2020)'s 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.[23]

Finally, we formally evaluate various real-time credit gap measures against an *ex post* credit gap extracted from final revised data through a conventional two-sided HP filter.

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 3).[24] With the benefit of hindsight, it may be argued that the two-sided HP filter

---

[23] Our sample contains only five financial cycle peaks and downswings (Table 1), which we interpret as periods of financial distress.

[24] The SARB currently identifies the credit-to-GDP gap by employing a two-sided HP filter with a smoothing parameter of $\lambda = 100,000$. Forecasts are utilised to enable the two-sided filter to use data beyond the current period t.

yields the 'best' final approximation of the credit gap. This can be motivated by the fact that the two-sided filter has the benefit of utilising *all* data points across the full sample, in contrast to the recursive nature of the one-sided filters. Moreover, the majority of data in the full sample has already been revised, so estimates of earlier trends and cycles should be more stable. Of course, as was argued above, the two-sided filter is not ideal for real-time cycle identification and policy making as it relies on unrealised future observations that are unknown to the policy maker at each point in time; however, the two-sided filter may provide a more reliable *ex post* identification of the cycle against which alternative real-time indicators can be evaluated.

Figure 6 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, at 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 λ = 1,600 implies a cycle of 7 years, while the BCBS's recommended smoothing parameter of λ = 400,000 implies a cycle of 20-25 years. We also evaluate λ = 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 own choice of λ = 100,000 for their benchmark model. Finally, we also test a smoothing parameter of λ = 25,600, which is in line with the 11-year average financial cycle duration calculated from Table 1 above.

While the different specifications of the two-sided filters yield similar turning points in the credit gap (e.g. 1980, 1984, 1998, 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 7. 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 8 or so years of the sample may point to the 'start-point' bias discussed earlier.

Figure 6: The benchmark two-sided HP filter

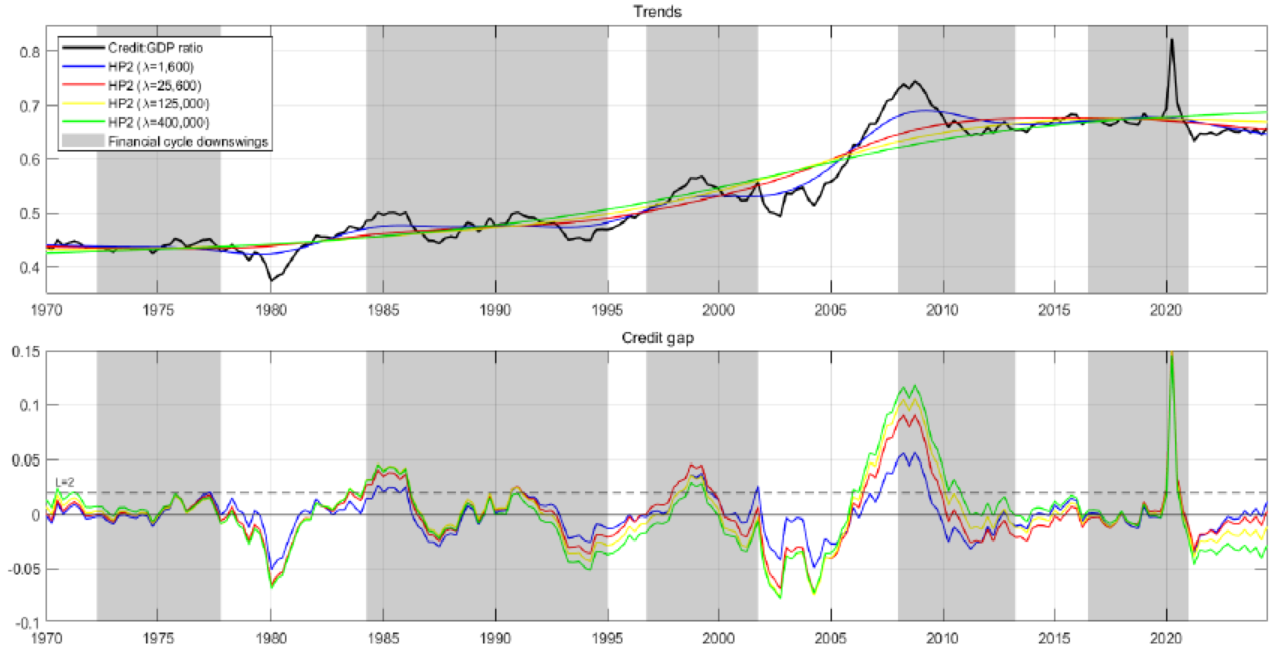

Source: Own calculations from SARB (2023b, 2025). The credit:GDP ratio is based on *final* revised data

Again, when the smoothing parameter is small (λ = 1,600), the credit gap tends to turn earlier and run ahead of other credit gaps. The turning points predicted by our choice of λ = 25,600, relative to λ = 125,000 and the BCBS recommendation of λ = 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 pronounced in the build-up to the GFC, where the BCBS parameterisation suggests the credit gap peaking at 10.4% in 2008 as compared to a peak of 7% – 8.7% under our parameterisations. This has implications for the design of the CCyB increments, as it suggests that buffers might not have been at maximum readiness before the biggest financial crisis in living history.

Figure 7: One-sided HP filter

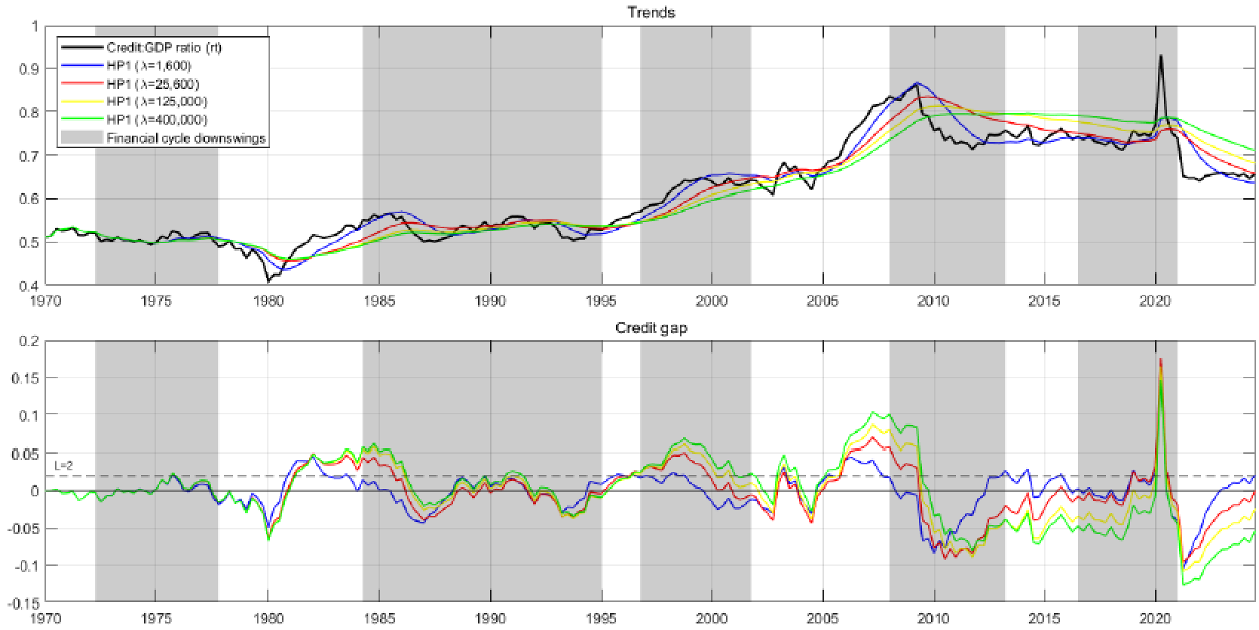

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. The credit:GDP ratio is derived from real-time data

4.1.3. Adjusted one-sided HP filter

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

The adjusted one-sided HP filter (HP1*) appears to be consistent with the standard HP1 filter (Figure 7), 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.

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

4.1.4. Rolling HP filter

Next we evaluate the rolling HP filter, again using λ = 25,600 and real-time data.[26] 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 8: Adjusted one-sided HP filter

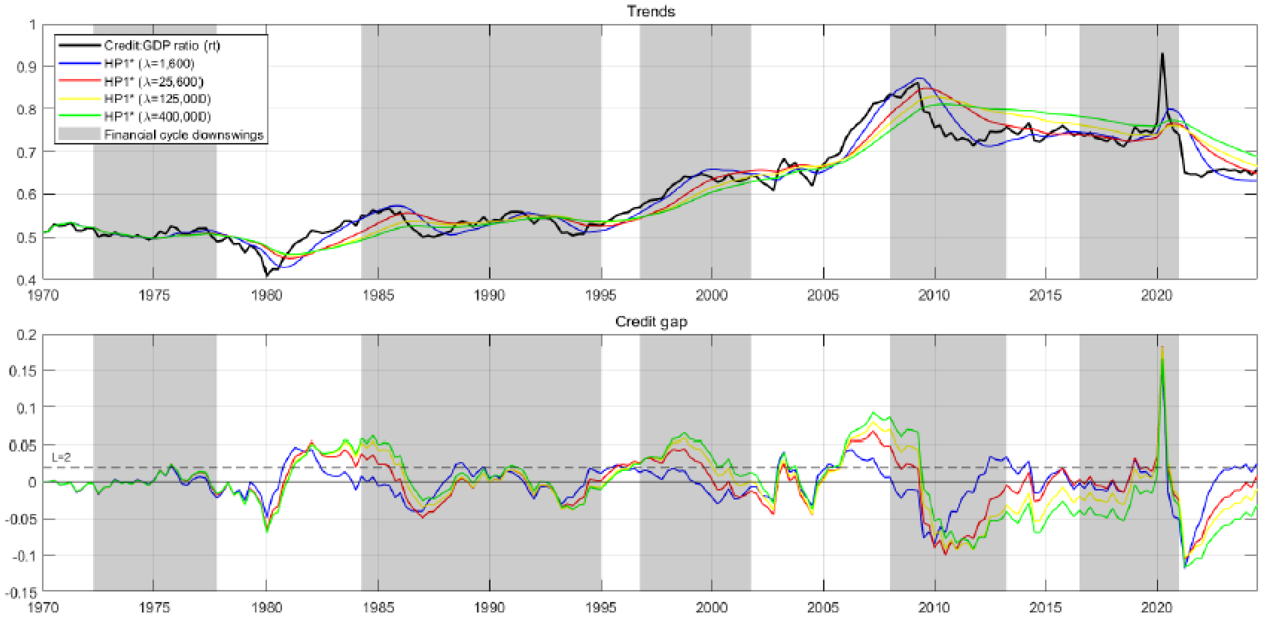

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. The credit:GDP ratio is derived from real-time data

Figure 9 illustrates that the rolling HP filter produces a significantly noisier credit gap when all the estimation windows overlap,[27] making it somewhat challenging for real-time cycle identification. However, it exhibits qualitatively similar *ex post* results to the one-sided filters with larger smoothing parameters (Figures 7 and 8), in that it identifies comparable peaks and troughs in the credit gap.

[26] 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.

[27] The longer the window W the less noisy the cycle would be; when $W \rightarrow T$ the rolling HP filter would approach the HP1 filter.

4.2. Is there a 'best' filter?

4.2.1. The credit gap and the financial cycle

From Figure 9 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 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 Appendix B.4); however, accurate identification of the credit gap then becomes vulnerable to forecast assumptions and potential errors.

Figure 9: Rolling HP filter (60-quarter rolling window)

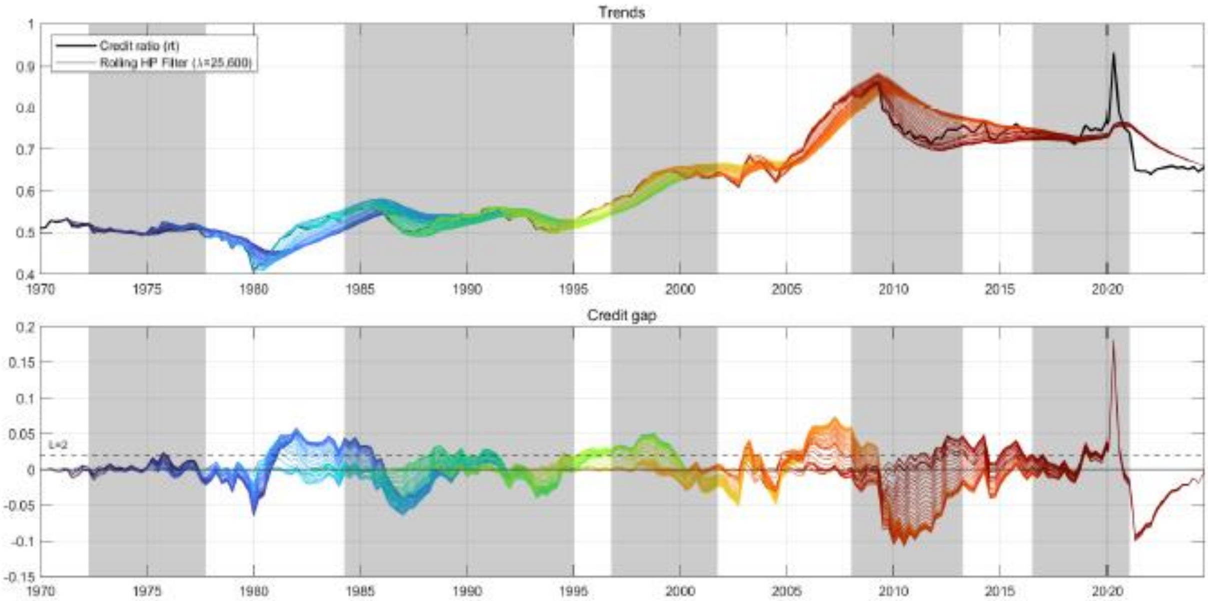

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. The credit:GDP ratio is derived from real-time data

Figures 7 and 8 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 its 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 7 and 8, from 1980 onwards both real-time HP1 filters breach the L = 2 threshold 1-2 years before the start of a financial cycle downswing. A notable exception is the financial cycle downswing that started in 2016, which was not preceded by a pronounced credit gap build-up; in fact, the credit gap remained persistently 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.[28] 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 out of the four financial cycle downturns since 1980 – 1984Q2, 1996Q4, and 2008Q1 – suggests that it is indeed suitable as 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 7 and 8), as house and equity prices started booming shortly after the rand crisis of 2001 and the small banks crisis of 2002.

4.2.2. Real-time vs *ex post* evaluation

Finally, we wish to evaluate the real-time approach which is the most consistent with the *ex post* realised credit gap. To this end, 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 current two-sided HP filter, estimated using final and

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

forecast data[29] up to 2027 and with a choice of smoothing parameter of $\lambda = 100{,}000$, and (ii) the one-sided HP filter calibrated to the BCBS default, estimated using final data[30] with a choice of smoothing parameter of $\lambda = 400{,}000$.

Figure 10 suggests that the two-sided filter benchmark is often late in predicting financial cycle peaks;[31] this is consistent with the results from the standard two-sided filter (Figure 6). The one-sided filters, on the other hand, very reliably predict financial cycle peaks. The one-sided filters based on real-time data are potentially better EWIs than its BCBS final data counterpart, as it gives an earlier warning of the 1996 downturn; however, these real-time filters also gives 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 which 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.[32]

29 These are the SARB's own forecasts from the Core Model.

30 Both *ex post* specifications utilise final revised, and not real-time, data in order to identify a *final* credit gap.

31 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.

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

Figure 10: One-sided HP filters against the SARB and BCBS benchmarks

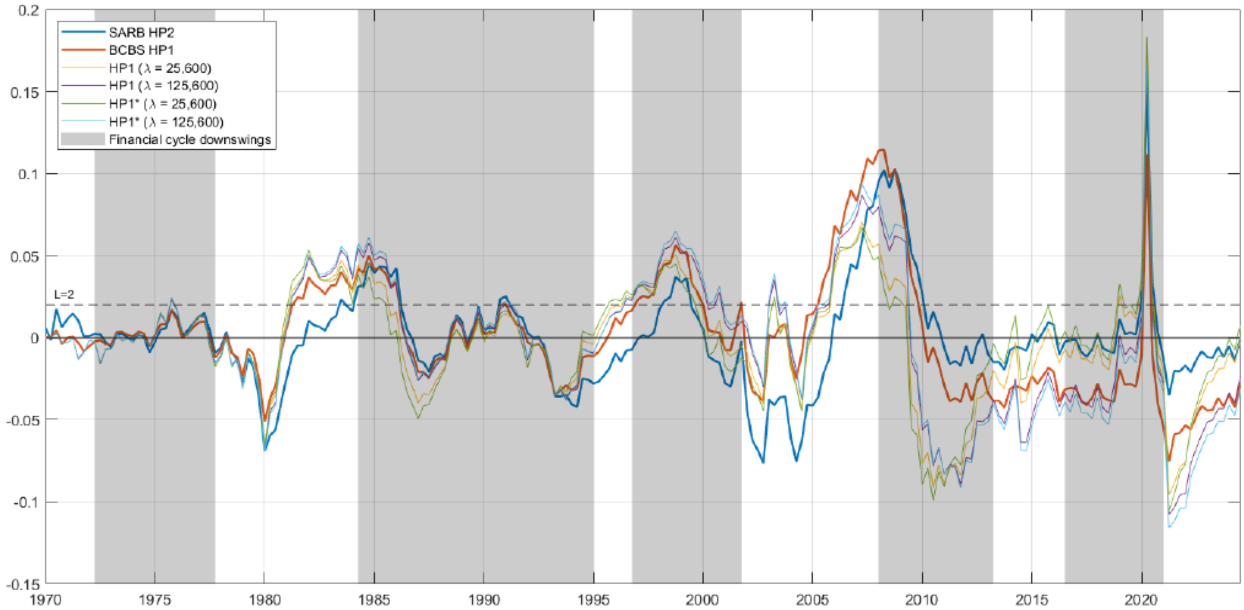

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model

Table 2 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 λ = 25,600 (filter #3), whereas the one-sided filter with λ = 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 λ's 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 λ = 1,600 is not justified by the empirical evidence on the duration of financial cycles (Table 1), and is far removed from the benchmark filters' larger smoothing parameters.

Table 2: 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, utilising 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 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

Table 3 presents the implications for how often the CCyB would be activated under various filtering approaches and threshold calibrations, while Figure 11 plots our recommended one-sided HP filter against thresholds of L = 2 and L = 3. 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.

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 (2016)'s 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.

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 4% threshold, it will be late on the 1984 and 1996 turning points (Figure 10). 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, it would probably not have been desirable to raise capital requirements in the midst of the downswings in 1976 or 1991, if the threshold was only 1%. This trade-off may, however, to some degree be ameliorated by the phasing in of the PCN CCyB.

Table 3: 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* ($\lambda = 25,600$)	53	24.2%	32	14.6%	17	7.8%
Total observations:	219					

Source: Own calculations

Naturally, the higher L, the less often 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, say L = 4 is crossed, the time window for banks to build up credit buffers under a CCyB might be too small.

4.3.2. Macroprudential and monetary policy coordination

Monetary and macroprudential policies interact in various ways, sometimes acting as substitutes (Cecchetti and Kohler, 2016) or complements. Macroprudential instruments aim to mitigate financial crises, while monetary policy focuses on inflation and growth.

Despite these distinctions, both contribute to macroeconomic stability, emphasising the need for coordination.

While monetary policy can inadvertently impact financial stability through various channels, macroprudential measures can alleviate these effects, enhancing policy effectiveness (Nier and Kang, 2016). Conversely, macroprudential tightening may slow output, but monetary policy can offset this. Aikman et al. (2019) demonstrate that the CCyB, a macroprudential tool, helps build resilience, allowing monetary policy to remain effective during financial stress.

Additionally, Jude and Levieuge (2024) illustrate that the CCyB and monetary policy complement each other, particularly in crises. The COVID-19 period demonstrated that their combined use amplified credit easing as opposed to the impact of the measures separately. Similarly, Aikman et al. (2019) highlight how using both tools improves economic stability, though trade-offs exist based on economic shocks.

While a deeper exploration of these dynamics is, however, beyond the scope of this paper, this again underscores the importance of interrogating credit gap dynamics within a broader policy context.

Figure 11: Recommended one-sided HP filter and CCyB thresholds

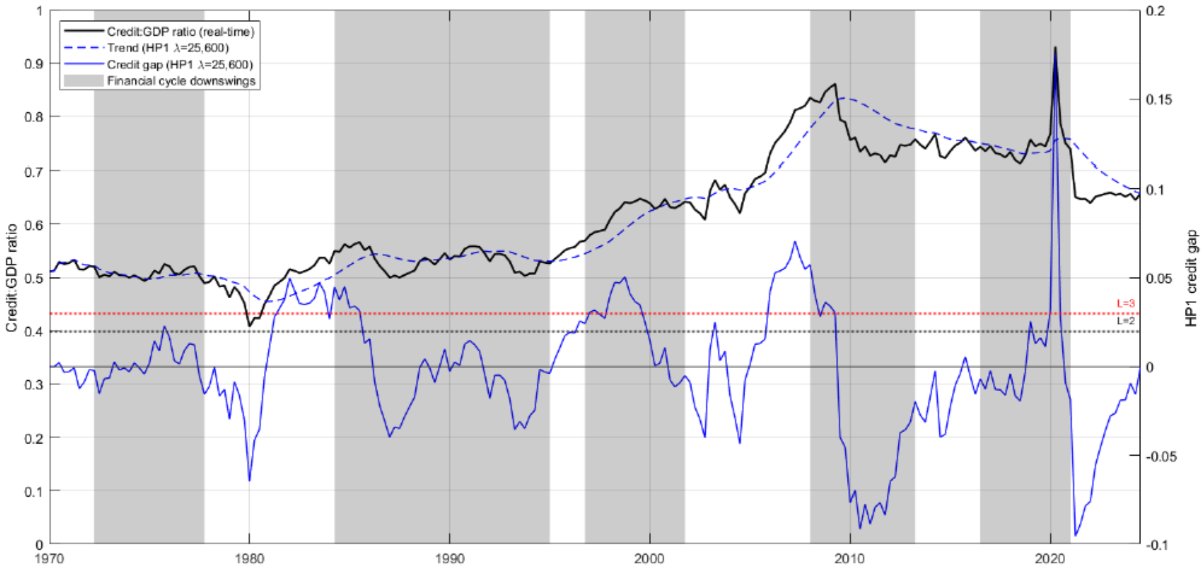

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model

South African Reserve Bank

5. Conclusion

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

The statistical analyses aimed to show how different models might produce varying 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. 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 specification credit gap correctly predicts three out of four financial cycle peaks since 1980 and matches *ex post* credit gaps from revised data. 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 with 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 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 rigidly. Judgement and discretion in capital buffer decisions remain essential.

References

Aikman, D, Giese, J, Kapadia, S and McLeay, M. 2019. *Targeting financial stability: macroprudential or monetary policy?* Number 2278. ECB Working Paper.

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 Papers* (278).

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

Australian Prudential Regulation Authority (APRA). 2017. 'A new(ish) tool in the toolkit'. Available at https://www.apra.gov.au/news-and-publications/a-newish-tool-toolkit.

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

Basel Committee on Banking Supervision (BCBS). 2017. 'Implementation: Range of practices in implementing the countercyclical capital buffer policy'. June 2017. Bank for International Settlements (BIS).

Basel Committee on Banking Supervision (BCBS). 2024. 'Implementation: Range of practices in implementing a positive neutral countercyclical capital buffer'. November 2024. Bank for International Settlements (BIS).

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.

Cecchetti, S G and Kohler, M. 2016. 'When capital adequacy and interest rate policy are substitutes (and when they are not)'. *International Journal of Central Banking* 36(September 2014): 205–231.

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.

Damen, S and Schildermans, S. 2022. 'Capital requirements, mortgage rates and house prices'. *Journal of Banking & Finance* 143(2022): 106596.

Denmark National Bank (DNB). 2016. 'Systemic risk and macroprudential policy'. Available at https://www.nationalbanken.dk/en/what-we-do/stable-financial-system/macroprudential-policy.

Drehmann, M, Borio, C, Gambacorta, L, Jimenez, G and Trucharte, C. 2010. 'Countercyclical capital buffers: exploring options'. *BIS Working Papers* July 2010(317): 1-64.

Drehmann, M, Borio, C and Tsatsaronis, K. 2011. 'Anchoring countercyclical capital buffers: the role of credit aggregates'. *BIS Working Papers* November 2011(355): 1–21.

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*.

Drehmann, M and Yetman, J. 2018. 'Why you should use the Hodrick-Prescott filter–at least to generate credit gaps'. *BIS Working Papers* September 2018(744): 1–21.

Dursun-de Neef, H Ö, Schandlbauer, A and Wittig, C. 2023. 'Countercyclical capital buffers and credit supply: Evidence from the COVID-19 crisis'. *Journal of Banking & Finance* 154(2023): 106930.

Fang, X, Jutrsa, D, Peria, S M, Presbitero, A F and Ratnovski, L. 2022. 'Bank capital requirements and lending in emerging markets: The role of bank characteristics and economic conditions'. *Journal of Banking & Finance* 135(2022): 105806.

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*. Number 13/2013. Norges Bank Staff Memo.

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*: 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.

Jude, C and Levieuge, G. 2024. 'Doubling down: The synergy of CCyB release and monetary policy easing'. Banque de France Working Paper Series September 2024(961).

Nier, E W and Kang, H. 2016. 'Monetary and macroprudential policies–exploring interactions'. A joint CBRT/BIS/IMF conference on "Macroprudential policy: effectiveness and implementation challenges". Istanbul, Turkey, 26-27 October 2015.

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.

Singh, K. 2025. 'Interest rate exposure: non-financial corporations'. Financial Stability Department Briefing Note, SARB (forthcoming).

South African Reserve Bank (SARB). 2023a. 'Drivers of corporate credit in South Africa'. SARB Occasional Bulletin of Economic Notes OBEN/23/01. Available at https://www.resbank. co.za/content/dam/sarb/publications/occasional-bulletin-of-economic-notes/ 2023/drivers-of-corporate-credit-in-south-africa-june-2023.pdf.

South African Reserve Bank (SARB). 2023b. '*Financial Stability Review*'. First Edition, 2023. Available at https://www.resbank.co.za/en/home/ publications/publication-detail-pages/reviews/finstab-review/2023/first-edition-2023-financial-stability-review.

South African Reserve Bank (SARB). 2024. '*Financial Stability Review'*. Second Edition, 2024. Available at https://www.resbank.co.za/en/home/publications/ publication-detail-pages/reviews/finstab-review/2024/second-edition.

South African Reserve Bank (SARB). 2025. 'Quarterly Bulletin'. Various editions. Available at 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, 2020. Available at https://www.statssa.gov.za.

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* 11(2020): 1–18.

Wolf, E, Mokinski, F and Schüler, Y S. 2024. 'On adjusting the one-sided Hodrick-Prescott filter'. *Journal of Money, Credit and Banking* 0(October 2024): 1–13. DOI: 10.1111/jmcb.13240.

Appendix A: Univariate filters

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:

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

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 in order to match business cycles with a typical 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 in time – 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.

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

Wolf, Mokinski and Schüler (2024, 2020) 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 utilising 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.

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

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 utilises final revised data. We analyse both our smoothing parameter of λ = 25,600 and the BCBS's benchmark of λ = 400,000.

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 6, 7 and 8). This does, however, result in a structural underestimation of the credit gap (Figure A.1), as 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.

It is interesting to observe a constant, but relatively muted, credit gap over the first 8-9 years of the sample. From 1980 onwards, however, swings in the credit gap have become more pronounced and volatile, perhaps resulting from 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 6, the HP2 filter seems to be late in predicting turning points in the financial cycle: Both 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 identify the 2016 turning point, although the HP1* filter comes closest to breaching the threshold. 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).

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

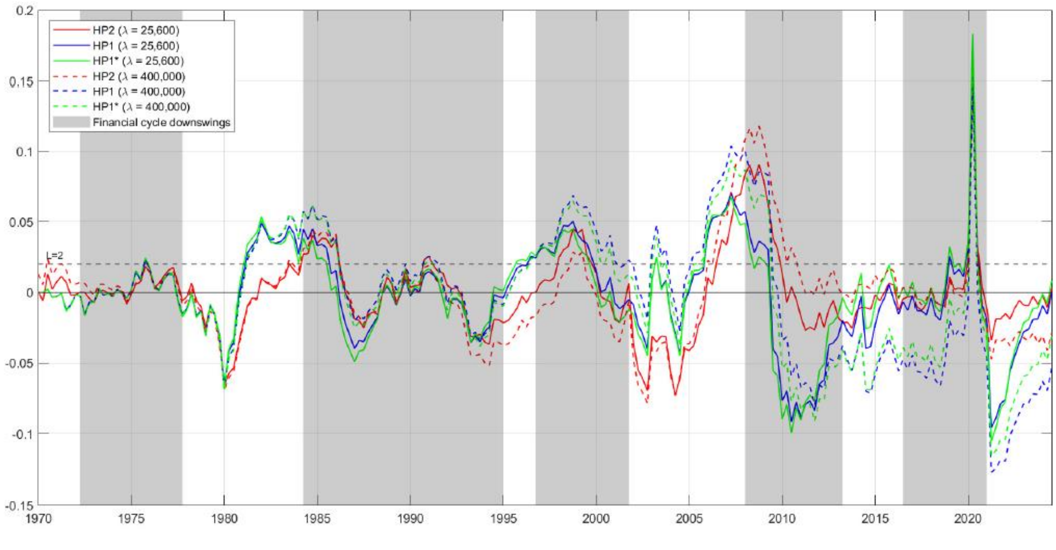

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. One-sided credit gaps are derived from real-time data

Appendix B: Sensitivity analyses

Appendix B.1: Cycle duration and smoothing parameters

Bosch and Koch (2020) calculate a 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.[33] 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 (2023b) 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 slightly 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 utilised by the various studies. However, detailed interrogation of these dating algorithms falls beyond the scope of this paper, and we choose to utilise the SARB's own financial cycle dates.

Based on the average length of the domestic credit cycle over our sample of 11 years, (Table 1), we assign λ = 25,600 as a benchmark for our cycle estimates. Figures 6, 7 and 8 have illustrated 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. Moreover, the larger λs potentially yield false signals (Type II errors) when its credit gap measures breach the L = 2 threshold in the early 1990s or early 2000s. It follows that a larger smoothing parameter tends to increase the credit gap's amplitude, and is likely to lead to a more frequent breach of the CCyB lower threshold of L = 2, requiring the SARB to more frequently deploy the CCyB.

The two larger smoothing parameters from the preceding analyses – λ = 125,000 and λ = 400,000 – suggest qualitatively similar credit gaps, but with divergent magnitudes

[33] 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).

at turning points. Our preferred smoothing parameter of $\lambda = 25{,}600$, consistent with a financial cycle duration of 11 years, exhibit similar characteristics. Predicted turning 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 magnitude of peaks and troughs are comparable. When the GFC hit 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 equal to $\lambda = 1{,}600$ yields a somewhat 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.[34] 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 in the sense that the credit gap does not fall by as much as under the one-sided specifications.

[34] Calibrating $\lambda = 1{,}600$ implies a cycle duration of 6-8 years – clearly significantly too short for the South African financial cycle.

Appendix B.2: Data smoothing

The massive GDP contraction in 2020 resulted in an extremely volatile credit ratio (Figure 3) and subsequent credit gaps (Figure 10). 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 by utilising a simple four-period moving average of real-time nominal GDP. Figure B.2 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 still 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 naïve gap.

Figure B.2: One-sided HP filters: Naïve and smoothed

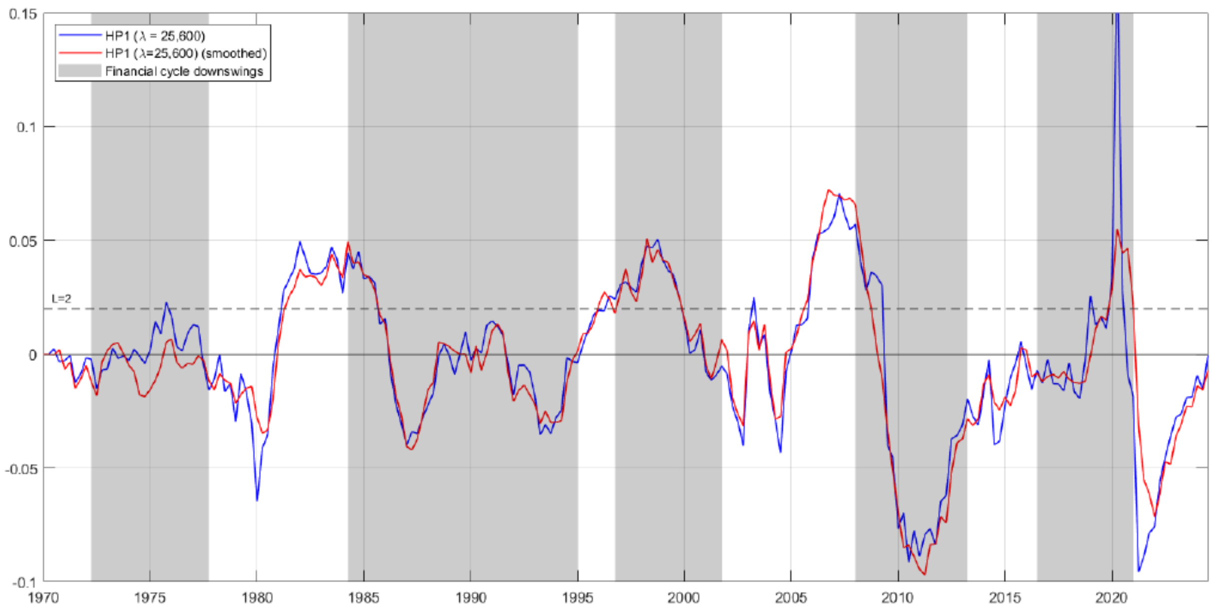

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. Credit gaps are derived from *real-time* data. The 'smoothed' gap had underlying GDP smoothed through a four-period moving average to ameliorate the effect of outliers.

Appendix B.3: Data samples

Figures B.3 and B.4 demonstrate the two- and one-sided HP filters' sensitivity to sample selection. The two-sided filter is less sensitive to the data sample, but the one-sided filter's 'start-point problem', alluded to in Section 3.3.1, is evident. Both filters perform particularly poorly during the build-up to the GFC (2005-2008) under the 2005-2019 sub-samples. But, on a positive note, each sub-sample's cycle approaches the full sample cycle over time, underscoring the notion that longer samples are associated with superior performance.

Appendix B.4: Forecast data

The use of forecast data can potentially ameliorate the two-sided filter's 'end-point problem', while at the same time enabling it to use data from beyond the current time t to extract the trend. For example, Alessandri, Bologna, Fiori, and Sette (2015)'s analysis 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.

Figure B.3: Sample sensitivity: Two-sided HP filter

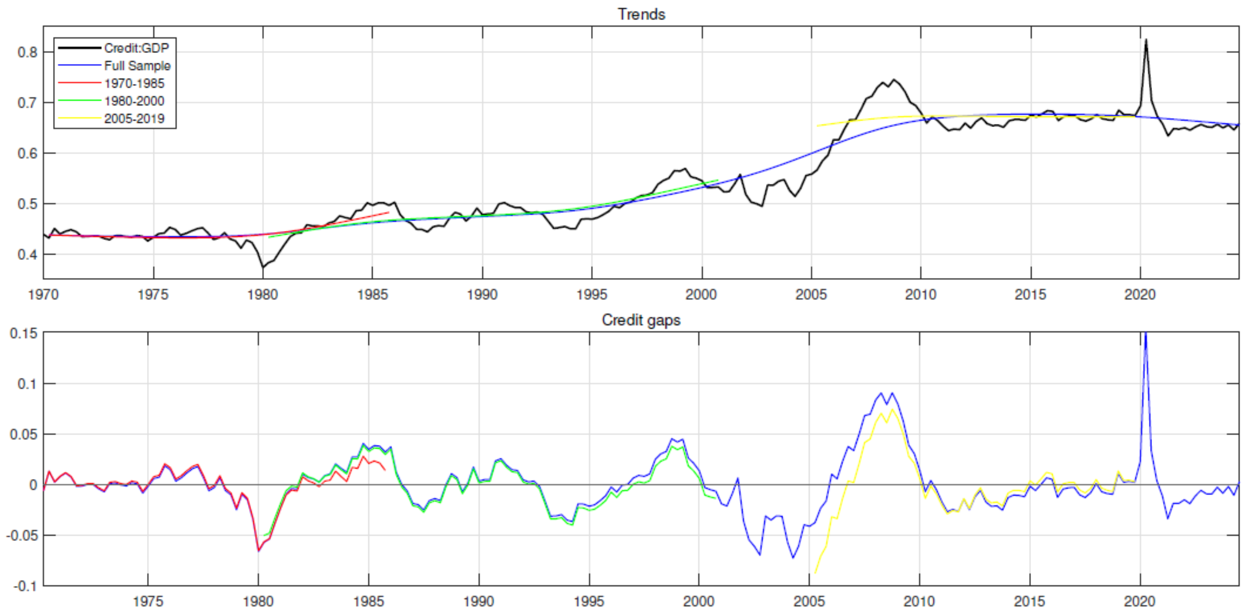

Source: Own calculations from SARB (2023b, 2025)

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 autoregressive process that relies on four previous observations 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 for crises prediction purposes. As Farrell (2016) demonstrates, the quality of the forecasts has significant implications for the reliability of the subsequent credit gap estimates. Therefore, in the absence of reliable forecasts, the two-sided filter may therefore be more suited to *ex post* analyses.

Finally, we tested a number of forecast scenarios, and its implications for the two-sided HP filter. Figure B.5 illustrates four simple forecast scenarios, with combinations of low, steady and high growth in nominal GDP and credit extension.[35] Up until the third quarter of 2024 we use actual data, and then the different scenarios are used to forecast up to the end of 2027. The darker lines represent higher credit growth relative to GDP, with the effect of increasing the credit:GDP ratio over time. Conversely, the yellow line represents a fall in the credit: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.

Figure B.4: Sample sensitivity: One-sided HP filter

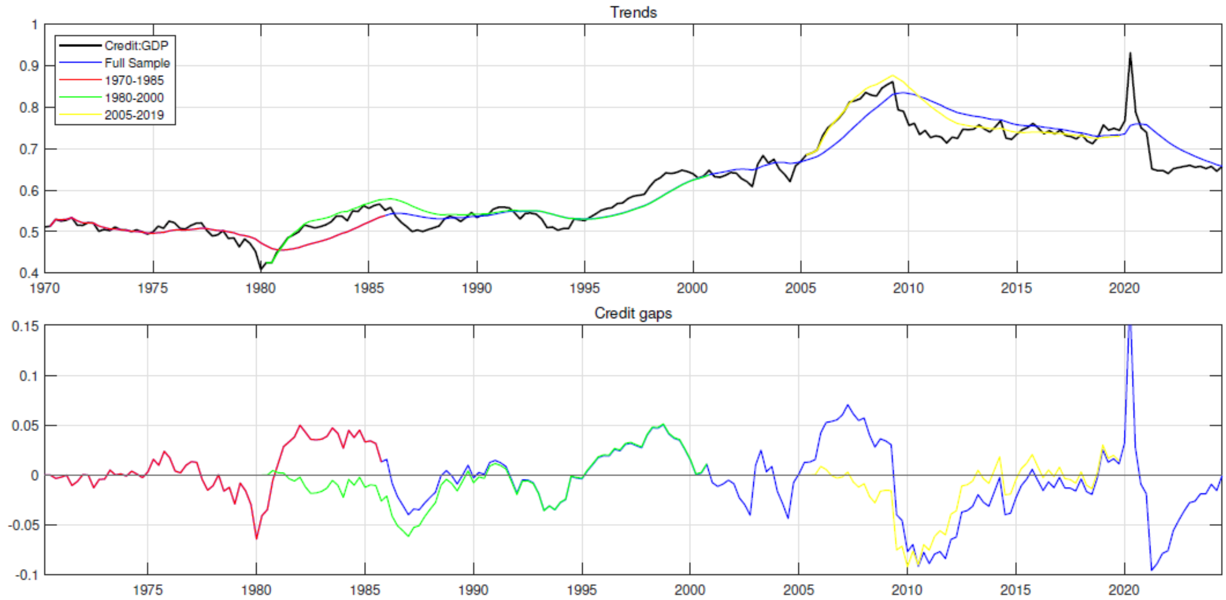

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model. One-sided credit gaps are derived from real-time data

[35] 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 B.5: Forecast sensitivity: Two-sided HP filters

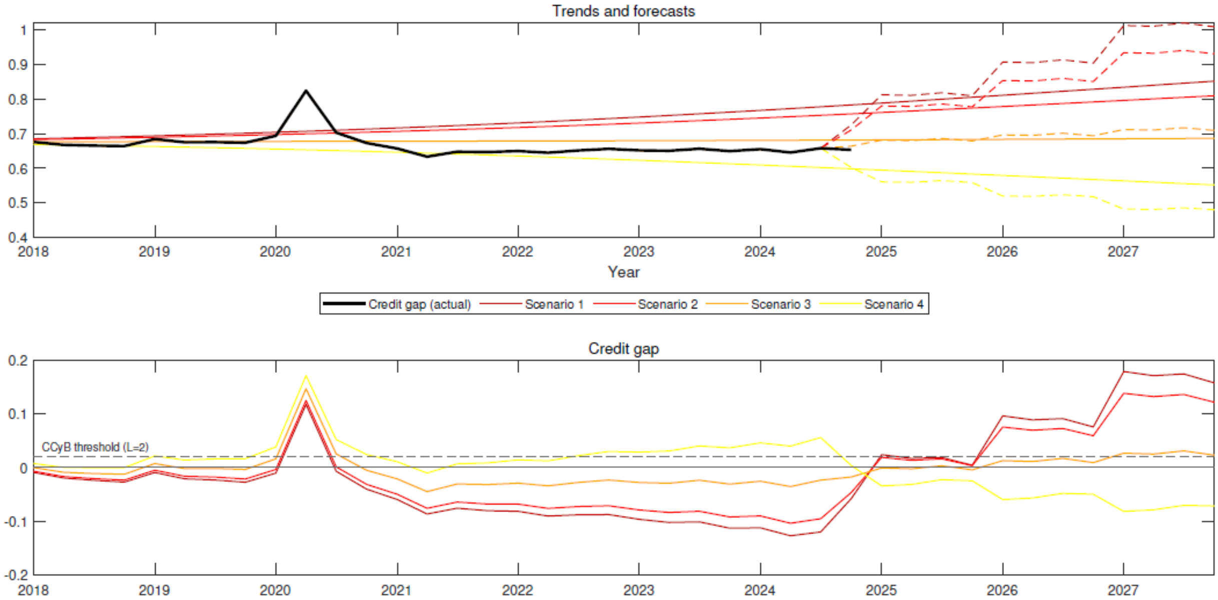

Source: Own calculations from SARB (2023b, 2025) and SARB Core Model

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

A related issue with the two-sided HP filter is its tendency to try and 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 fall in the credit gap under the 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 we are beyond the 2% threshold, which might perversely lead to exactly the kind of slowdown in the credit: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 potentially contributing to an undesirable credit boom.