

Du Rand, Gideon; Hollander, Hylton; Van Lill, Dawie

Working Paper

A deep learning approach to estimation of the Phillips curve in South Africa

WIDER Working Paper, No. 2023/79

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Du Rand, Gideon; Hollander, Hylton; Van Lill, Dawie (2023) : A deep learning approach to estimation of the Phillips curve in South Africa, WIDER Working Paper, No. 2023/79, ISBN 978-92-9267-387-1, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki,
<https://doi.org/10.35188/UNU-WIDER/2023/387-1>

This Version is available at:

<https://hdl.handle.net/10419/283775>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WIDER Working Paper 2023/79

A deep learning approach to estimation of the Phillips curve in South Africa

Gideon du Rand, Hylton Hollander, and Dawie van Lill*

June 2023

Abstract: In this study, we provide a comprehensive estimation of the contemporary Phillips curve relationship in the South African economy using a novel deep learning technique. Our approach incorporates multiple measures of economic slack/tightness and inflation expectations, contributing to the debate on the relevance of the Phillips curve in South Africa, where previous findings have been inconclusive. Our analysis reveals that long-run inflation expectations are the primary driver of inflation, with these expectations anchored around 5% historically but declining since the financial crisis. This trend suggests that the South African Reserve Bank's (SARB) monetary policy implementation has become increasingly credible in addressing high inflation and price instability. We also find that short-run expectations and real economic activity play a substantial role in explaining deviations from the long-run trend, supporting the existence of a quantifiable Phillips curve relationship in the South African economy. The deep learning methodology employed in this paper offers several advantages for policymakers. It accommodates the complex interplay of various factors influencing inflation, providing a more accurate representation of the Phillips curve relationship. The model's interpretable nature allows policymakers to discern the key drivers of inflation, facilitating more targeted and effective monetary policy interventions. Our findings can help inform strategies to combat persistently high inflation rates in South Africa, supporting the SARB's ongoing efforts to achieve price stability and sustainable economic growth.

Key words: Phillips curve, inflation, inflation expectations, output gap, monetary policy, neural network

JEL classification: E3, E58, C45

* Stellenbosch University, South Africa; corresponding author, Dawie van Lill: dvanlill@sun.ac.za

This study has been prepared within the UNU-WIDER project [Southern Africa—Towards Inclusive Economic Development \(SA-TIED\)](#).

Copyright © UNU-WIDER 2023

UNU-WIDER employs a fair use policy for reasonable reproduction of UNU-WIDER copyrighted content—such as the reproduction of a table or a figure, and/or text not exceeding 400 words—with due acknowledgement of the original source, without requiring explicit permission from the copyright holder.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-387-1

<https://doi.org/10.35188/UNU-WIDER/2023/387-1>

Typescript prepared by Siméon Rapin.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland and Sweden, as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

The trade-off between inflation and real economic growth or unemployment, often summarized by the Phillips curve, plays a central role in macroeconomic policy and has driven monetary and fiscal policy discussions for decades since A.W. Phillips first uncovered the relationship in the United Kingdom in the early 1950s. Estimating this relationship has proven difficult in many countries, however, with South Africa being no exception. One of the major challenges in estimating the Phillips curve is the presence of latent factors, such as inflation expectations and the output gap, which are not directly observable. In this paper, we offer a novel approach to estimation of the contemporary Phillips curve relationship in the South African economy by utilizing a deep learning technique that accounts for multiple measures of economic slack/tightness and inflation expectations while addressing the issue of latent factors.

The deep learning technique we employ is based on a Hemisphere Neural Network (HNN) model, as proposed by Goulet Coulombe (2022). This model constructs an approximate output gap measure by combining various economic indicators related to labour market conditions, the production side of the economy, and other national account information. In addition, it integrates information about price variables, such as price indices and their lags, as well as inflation expectations from surveys. The resulting output from this model, in its final layer, takes the form of a linearized New-Keynesian Phillips curve relationship.

Using the HNN model, we are able to overcome some of the traditional challenges in estimating the Phillips curve, such as dealing with the aforementioned latent factors like inflation expectations and the output gap. The deep learning approach allows us to capture the complex interplay of various factors that influence inflation, thereby providing a more accurate representation of the Phillips curve relationship in the South African context.

Our analysis reveals that long-run inflation expectations, historically anchored around 5%, have been the primary driver of inflation in South Africa, but have been declining since the global financial crisis of 2007–09. This trend indicates that the South African Reserve Bank’s (SARB) monetary policy stance has become increasingly credible in addressing high inflation and price instability. Furthermore, we find that short-run expectations and real economic activity play a substantial role in explaining deviations from the long-run trend, supporting the existence of a quantifiable Phillips curve relationship in the South African economy.

The deep learning methodology employed in this paper is highly interpretable, enabling policy-makers to discern the key drivers of inflation and facilitating more targeted and effective monetary policy interventions. The HNN model offers several advantages over traditional methods, as it can handle high-dimensional problems and does not rely on restrictive assumptions about linearity or limited variables, which are commonly found in state-space representations.

Our findings contribute to the ongoing debate on the relevance of the Phillips curve in South Africa, where previous research has yielded inconclusive results. By providing a comprehensive estimation of the Phillips curve relationship using a novel deep learning technique that addresses the challenges posed by latent factors, we can help inform strategies to combat persistently high inflation rates in South Africa, supporting the SARB’s ongoing efforts to achieve price stability and sustainable economic growth.

2 Evidence from South Africa

The Phillips curve is a key component of the policy discourse, both locally and internationally. The so-called ‘flattening’ of the Phillips curve after the global financial crisis has been a challenge for policy-makers in developed economies (see, for example, the recent work of Galí and Gambetti 2019). This conundrum also holds true for developing economies as evidenced in Forbes et al. (2021). In a country like South Africa with persistently high levels of inflation, it is useful to know how economic activity and inflation expectations might be placing pressure on the prevailing price level.

In this section, we provide a brief overview of some of the most recent findings in the South African literature on the Phillips curve. We place our focus on articles that investigate the Phillips curve relationship within the period of inflation targeting in South Africa. For a more in-depth discussion on the contribution of papers in this literature in the period before the early 2000s, one can reference the work of Vermeulen (2017).

One of the most comprehensive recent attempts to quantify the relationship is captured by Fedderke and Liu (2018). They use various methods to estimate and establish the Phillips curve relationship and find that the search for this relationship in South Africa has resulted in a ‘consistent failure’ to estimate it (Fedderke and Liu 2018: 197). The general consensus from the literature is that there is no conclusive evidence for effects originating from an output gap on the inflation rate (Vermeulen 2017; Fedderke and Liu 2018). Despite this lack of empirical evidence for a strong (and theoretically consistent) relationship, there are some papers that provide sound evidence, against the prevailing consensus, for the existence of a Phillips curve relationship. Articles by Dladla and Malikane (2022) and Botha et al. (2020) find a measurable impact from changes in the output gap on inflation.

The paper of Kabundi et al. (2019) is the closest to ours in terms of the methodological approach in the context of the South African literature. They use the method from Chan et al. (2016) that allows for time varying parameters in estimating the Phillips curve relationship. Kabundi et al. (2019) find that the slope of the Phillips curve flattened out in the early-to-middle 2000s, with the biggest dampening occurring after the financial crisis.

The reason for the flattening of the slope of the Phillips curve is seen as a combination of events. With the advent of inflation targeting in many countries, inflation has become more anchored which has reduced the impact of real activity on inflation. On South Africa, Kabundi et al. (2019) interpret this finding to be a result of both an increase in the credibility of the SARB to maintain their inflation target and ‘good luck’ (p. 94).¹ In contrast to Kabundi and Rapapali (2019), their findings do not point to a decreased potency of monetary policy on inflation through the real economy following the global financial crisis. In short, rather than being ‘dead’, the Phillips curve more likely exhibits non-linear relationships over time and in response to permanent versus temporary shocks, which leads to a complex lead-lag relationship between inflation and real activity (Reinbold and Wen 2020).

In our discussion section we will make direct reference to the results obtained from Kabundi et al. (2019) since it is one of the only papers that is directly comparable in the literature. Most other papers consider constant coefficients with respect to the variables in the Phillips curve equation. This means that comparison is not entirely accurate, as we are considering two different narratives.

¹ Kabundi et al. (2019) define credibility (or ‘faith in the IT regime’) to be ‘the degree to the extent of which inflation expectations are anchored to the target’ (p. 75).

3 Method

We employed a method for this paper based on the work of Goulet Coulombe (2022). The author finds that a Hemisphere Neural Network can act as a natural environment to capture the inherent nonlinearities that are present in the Phillips curve relationship. In this paper, we present the network architecture within the context of the expectations-augmented Phillips curve. The model can be expressed as follows:

$$\pi_t = \theta_t E_t(\pi_{t+1}) + \gamma_t g_t + v_t$$

where θ_t and γ_t are parameters that represent the coefficients on inflation expectations and the output gap, and may vary over time. v_t is considered a noise term.

One of the features of this network is that we can arbitrarily add new hemispheres to the existing structure of the model. These new hemispheres (which will often be called blocks) represent nonlinear combinations of the variables that comprise it. We can, as an example, add a commodity price block and reformulate the equation as follows,

$$\pi_t = \theta_t \mathcal{E}_t + \gamma_t g_t + \zeta_t c_t + v_t$$

In this equation, \mathcal{E}_t represents expectations, g_t is the output gap, and c_t represents commodity prices. As currently phrased this model is similar to a factor model with three factors. We refer to these components as hemispheres in the Hemisphere Neural Network.

We can define the hemispheres, \mathcal{H} s, broadly as $h_{t,\mathcal{E}} = \theta_t \mathcal{E}_t$, $h_{t,g} = \gamma_t g_t$, and $h_{t,c} = \zeta_t c_t$. Note that $h_{t,\mathcal{E}} \in \mathcal{H}_{\mathcal{E}}$, $h_{t,g} \in \mathcal{H}_g$, and $h_{t,c} \in \mathcal{H}_c$. We can then refer to \mathcal{H}_g , $\mathcal{H}_{\mathcal{E}}$, and \mathcal{H}_c as the expectations, real activity, and commodities prices hemispheres.

Since the architecture is constructed in such a way as to be extendable, we can easily split some of the hemispheres into sub-hemispheres. In our model this entails splitting the expectations hemisphere into short-run and long-run expectations.

There is no real limitation on the number of hemispheres that can be added. One potentially useful hemisphere that can be included for future research is a credit block that contains information on the balance sheet of banks, equity prices, and other financial market indicators. In addition, one could include an international block, which would contain exchange rate variables as well as information on terms of trade. This idea to include an international dimension is in line with the exploration of Botha et al. (2020).

3.1 Extracting the output gap and its coefficient

The field of deep learning has gained tremendous momentum in constructing models that accurately capture complex data structures. Despite its success, the utilisation of highly nonlinear approaches can result in a significant loss of interpretability, limiting our understanding of the underlying relationships within the model. To overcome this limitation and gain transparency into the ‘black box’ of fully connected neural networks, we propose imposing specific structural assumptions to ensure interpretability.

In this model, the concept of output gap is denoted by g_t . However, one should note that the neural network model does not provide direct insight into either g_t or γ_t . Rather, the architecture of the network is exclusively intended to yield information regarding their product, denoted as $h_{t,g}$, which captures the contribution of the output gap to inflation. To obtain a precise understanding of g_t or γ_t , certain

assumptions regarding their temporal evolution must be made. It is customary in the field of economics to make such assumptions with respect to the dynamics of the individual components.

A plausible approach to factorizing the variables involves setting $\gamma_t = h_\gamma(t)$ and $g_t = h_g(\mathcal{H}_g \setminus t)$. By defining g_t in this manner, the time trend is removed from the set of predictors contained within this domain. Nonetheless, the coefficient on the output gap retains its dynamic character. This particular factorization bears similarities to the assumption of a Phillips curve coefficient that follows a random walk. Consequently, the resulting output gap is composed of real activity data that are assumed to remain fixed over time, with the exception of a scalar coefficient that is allowed to vary dynamically (Goulet Coulombe 2022).

To implement this transformation, it is possible to extract a temporal dimension by introducing a time trend within an additional hemisphere. The final layer in the HNN will then be the following,

$$\hat{\pi}_t = h_{\mathcal{E}^{\text{LR}}}(t) + h_\theta(t)h_{\mathcal{E}^{\text{SR}}}(\mathcal{H}_{\mathcal{E}^{\text{SR}}} \setminus t) + h_\gamma(t)h_g(\mathcal{H}_g \setminus t) + h_\zeta(t)h_c(\mathcal{H}_c \setminus t) + v_t$$

To identify the coefficients, it is necessary to impose non-negativity constraints on the coefficient hemisphere outputs. This is achieved by subjecting the variables to an absolute value layer prior to their integration into the final layer of the network.

The HNN method is considered a supervised learning approach that does not impose a parametric law of motion on the variable of interest. Neither does it limit the output gap to be made up of a single variable, such as performed in many studies on the Phillips curve, where the output gap is provided as a filtered variable (Goulet Coulombe 2022).

In the case of the HNN method, the output gap measure, as an example, is implicitly generated as a nonlinear combination of a collection of variables that are related to real activity. The HNN method generates a sufficient statistic for the output gap that can be used to gauge the contribution to inflation.

3.2 Estimation and tuning of the model

The neural network model used in this study has fully connected networks in each hemisphere. The state and coefficient hemispheres have three layers, which is a relatively shallow network. The state hemisphere has 400 neurons, while the coefficient hemisphere has 100 neurons. The number of epochs is set at 500, and the activation functions within the network are all rectified linear unit (ReLU) functions. The learning rate for the model is 0.005, and early stopping is implemented to perform a form of ridge regularization on network weights. The data are shuffled randomly through six quarters, and the optimizer used is Adam. The dropout rate is set at 0.2, and standard scaling is applied to the predictors to improve the performance of the regression network. Each hemisphere is given equal importance by dividing the standardized block of variables by the square root of the number of variables in that hemisphere. This ensures that hemispheres with more regressors are not weighted more heavily a priori.

4 Data

The target variable is inflation, as constructed by appropriate manipulation of the consumer price index (CPI). Data for the CPI was gathered from Statistics South Africa (StatsSA), and inflation was calculated as quarter-on-quarter growth, which was then annualized. The independent variables for the model were gathered from the South African Reserve Bank's quarterly bulletin. Independent variables that are not in rate terms were transformed to induce stationarity. For the independent variables that present as rates, a procedure of annualization was followed to allow for comparison with the target variable.

Here we describe the process followed to induce stationarity for the subset of independent variables that do not enter in rate form. The testing procedure is simple. All of the relevant predictors are tested using both the Augmented Dickey Fuller (ADF) and Kwiatkowski Phillips Schmidt Shin (KPSS) test for stationarity. If the variables are considered nonstationary then a quarter-on-quarter growth rate is calculated to induce stationarity. Finally, there are variables that need to be differenced twice, those that are integrated of second order.

The model consists of four hemispheres as discussed above. The expectations hemisphere is split in two additive components, namely long-run and short-run expectations. The other two hemispheres represent real activity and commodity prices. Table 1 provides an example of the types of variables that were considered for the different hemispheres. Of particular importance is the short-run expectations component, which consists of variables that relate to different prices and also the expectations of inflation as derived from surveys conducted by the Bureau of Economic Research (BER). These surveys capture the current year and one-year-ahead inflation expectations for financial analysts, business representatives, and trade union representatives. Reid and Siklos (2021b) provide a detailed analysis strongly motivating the usage of the disaggregated expectations series, which stands in contrast to most studies that only use the aggregate (i.e. equally-weighted) expectations of the different groups.² The list of independent variables used for each hemisphere in the model is included in the appendix.

Different specifications were attempted, with different numbers of variables in each hemisphere. Specifications were robust to the addition of different variables (with results not significantly altered). Variables added beyond the baseline specification presented in the paper did not add to the explanatory power of the model. However, there is the possibility that additional variables could be included if merited.

Table 1: Defining the hemispheres

\mathcal{H}	Content
$\mathcal{E}_t^{\text{LR}}$	t (exogenous time trend)
$\mathcal{E}_t^{\text{SR}}$	Inflation expectations from BER, lags of price indexes in SARB-QB, lags of π_t, t
g_t	Labour market variables, National Accounts, t
c_t	Oil, platinum, and gold price series from SARB-QB, t

Note: BER: Bureau of Economic Research, SARB-QB: South African Reserve Bank Quarterly Bulletin.

Source: authors' elaboration.

Most of the data utilized were of quarterly frequency. However, some monthly series were utilized. These monthly series were resampled and transformed into quarterly variables. The general rule for the transformation was that rates were averaged across the quarter. When it came to flow variables, the cumulative sum was calculated.


For our model specification we added two lags of both the independent and dependent variables. In addition to the lags, moving averages of order two and four are included. This moving average inclusion is referred to as a Moving Average Rotation of X (MARX) transformation. The reason for the inclusion of the MARX transformed variables is discussed in detail in Goulet Coulombe et al. (2021). In this article they find that this moving average rotation can provide a substantial increase in the predictive ability of models that use nonlinear methods, such as the HNN model.

² Reid and Siklos (2021a) provide a comprehensive literature review on the usage of BER survey data for inflation expectations research in South Africa.

Data for the model span the period from the second quarter of 2002 until the last quarter of 2021. The reason for the selection of this sample period is data availability. As an example, a series such as inflation expectations only starts in 2002.

5 Results and discussion

Figure 1: HNN prediction of inflation vs actual inflation



Note: the blue line represents the prediction and the black line actual inflation. Recessions (business cycle turning points as retrieved from the SARB) are depicted in pink.

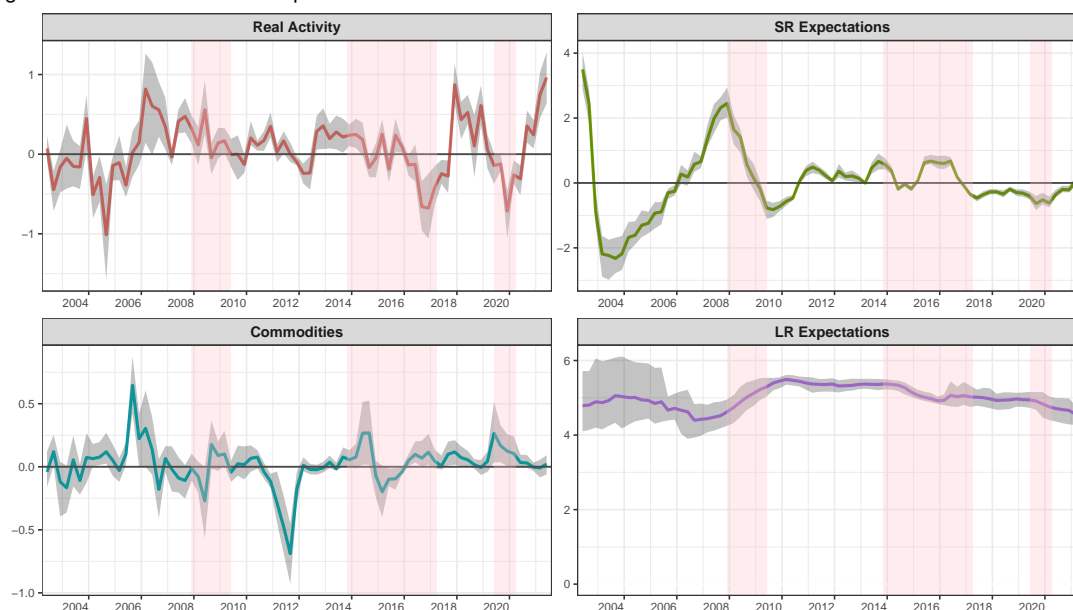
Source: authors' elaboration.

We start by considering the predictive performance of our model, as illustrated in Figure 1. The model fit is better after the financial crisis in 2008. Before the crisis, the model misses out on parts of the large reduction in inflation before 2005 and the model also does not fully reach the heights of the subsequent increase in inflation. Except for this initial sample, the model tracks inflation well.

Performance could be improved through further hyperparameter tuning, but there is a risk that we will overfit the model. One of the distinct dangers in applying neural networks is that the model tends to fit too closely to the data, since neural networks allow for an arbitrary level of nonlinearity in its construction. Because of the relatively weak fit on the period before the financial crisis, there is scope to include more explanatory variables and even different hemisphere structure. As previously mentioned, a possible extension considered is to include a credit / bank balance sheet hemisphere or even an international hemisphere that contains information about exchange rates and terms of trade. However, for this paper we focus on the baseline model with four hemispheres.

The contributions of the different hemispheres to the estimate of inflation is presented in Figure 2. The training sample for this model ends in the third quarter of 2021. The dashed line provides an indication of the out-of-sample period. Values after the dashed line are forecasts of the contribution of each component to inflation.

Figure 2: Contributions of hemispheres to inflation estimate



Note: contributions of real activity, short-run expectations, commodities, and long-run expectations to inflation over time. Dashed lines represent the beginning of the out-of-sample period. Recessions (business cycle turning points as retrieved from the SARB) are depicted in pink. The grey shading represents the 68% credible region.

Source: authors' elaboration.

5.1 Before the financial crisis

The upper left quadrant of Figure 2 shows that there is an initial negative and then corresponding positive contribution of real activity to inflation in the early 2000s just before the financial crisis occurs. At their peak the size of these movements accounts for about 1pp of inflation. Similarly commodity prices, situated in the lower left quadrant, experience an increase in 2005, which filters through to inflation.

In the upper right quadrant, we observe a sharp initial increase resulting from short-run expectations. The magnitude of the contribution in this hemisphere is much larger than the effect observed in the case of real activity (in the vicinity of 3.5pps for this initial increase). These short-run expectations contributions experience a precipitous drop (about 5pps in total) after the initial high, and then slowly start increasing, reaching a inflection point in 2007, after which the contribution becomes positive. This sharp fall in short-run expectations drives most of the downward pressure experienced on the inflation rate between 2004 and 2008. This contribution can better be observed in Figure3. These findings corroborate the literature on inflation expectations in South Africa linked to the increased credibility of the SARB and the resultant improved anchoring of inflation expectations and inflation variability itself (Kabundi et al. 2015; Reid and Siklos 2021a, 2021b).

Commodities make one significant contribution to inflation in the period before the crisis, a small increase in the run-up to the crisis. The reason for this increase might be as a result of a commodity price boom in 2005, but it is not clear that this is the reason.

During the period before the financial crisis, long-run expectations of inflation are anchored around 5%, as depicted in the bottom right quadrant of Figure 2. A slight decrease in long-run expectations is experienced before the financial crisis, but these expectations move again to above 5% in the immediate aftermath of the crisis. This is in line with findings from Kabundi and Schaling (2013), who use a basic Phillips curve estimation strategy which utilizes inflation expectations data from the BER to show that the SARB's implicit inflation target lies toward the upper bound of the inflation targeting band. The midpoint of the targeted 3% to 6% inflation band is the one mandated by the National Treasury.

However, we can clearly see that long-run expectations are at the top end of this band throughout until about 2015.

5.2 After the financial crisis

We now return to a discussion on the post-2008 sample. An interesting observation regarding the behaviour of short-run expectations is worth pointing out. The main contributor to driving down inflation in the years immediately after the financial crisis was short-run expectations. From the graph we can see that between 2008 and 2010, short-run expectations contributed to an approximate 3 – 4 percentage point decline in inflation. This was during a period where long-run expectations were on the incline, which shows the importance of short-run expectations in driving down inflation during this period.

In the period directly after the financial crisis, we observe a progression to a more muted impact of real activity on inflation until the end of 2016. This is similar to the finding of Kabundi et al. (2019), who find that demand-side shocks in the period after the financial crisis do not contribute significantly to inflation in their sample that ends in 2014. This representation is in line with the idea of a declining Phillips curve over time as found in other papers on the impact of the output gap on inflation. However, our paper spans a longer time period than most other studies, and one of the unique contributions of this method employed is the ability to identify the time-varying nature of this relationship.

Interestingly, if you were to consider averages on the contribution of short-run expectations and real activity, one would struggle to identify the time-varying impact. This is one of the major drawbacks of time series methods that impose the restriction of time-invariant coefficients, they miss out on an important dimension. Constant coefficient methods don't express the distribution of events. This can only be observed in methods that implement time-varying coefficients. If one were to construct a narrative on the nature of the Phillips curve relationship with a constant coefficient estimation, one would find different results depending on the time series sample selected. As evidenced by the results thus far, the combination of incorrect selection of output gap measure and usage of constant coefficient estimates could prove fatal to any serious argument concerning the nature of the Phillips curve relationship.

In the direct aftermath of the financial crisis, in 2012, there was a negative impact to inflation originating from the commodities sector. However, beyond that, the most prominent result in the post-2010 period is that the effect of real activity and short-run expectations on inflation has been variable and somewhat muted between 2010 and 2016. This makes sense given the nature of inflation behaviour since 2010. If one looks back to Figure 1, one observes that inflation was not particularly volatile and that it exhibits almost mean-reverting behaviour.

Thus far our results align with the findings from Reid and Siklos (2022) who find that after the financial crisis, 'inflation expectations become less volatile and tend to remain at the top of the SARB's inflation-targeting band'. Our specification allows us to disentangle the long- and short-run expectation components. We can therefore observe the long-run tendency to stay at the top of the inflation targeting band and the short-run reduction in volatility by looking at the short-run expectations component.

In Kabundi et al. (2019) they look at inflation persistence. This aligns with the long-run expectations component in our model. They also find that inflation persistence increases up until the financial crisis, and then the model showcases a plateau and slow decline of expectations. Our result extends beyond the 2014 time frame from their paper and shows a further decline in inflation persistence. In the period following 2015, long-run expectations appear to have decreased from a high of around 5% to almost 4.5% in 2020, with a further decline predicted by our model. This might be indicative of the ability of the central bank to target the middle of the band more effectively in recent years.


It is in 2019 that we see a sharp decline in inflation, which is explained in our model as a phenomenon linked to declining long-run expectations and diminished real activity. One interpretation is that the

SARB has established confidence in its ability to drive down inflation, and this has reduced the anchoring of expectations to a new target. However, the impact as observed in the long-run expectations panel is one of a slow and marginal decline.

It would be interesting to include more variables within the real activity block, or even attach another block to determine what the driving factors behind the decrease in inflation are during this period. It appears to be a combination of short- and long-run factors that have driven inflation down during this period. Kabundi et al. (2019) find that inflation in South Africa since 2008 has been less persistent. SARB's credibility has been improving, with a movement toward the midpoint of the band.

Consider the output of Figure 3, which acts to provide a better indication of the relative contribution that each of the different hemispheres makes to the fit to the inflation target. It is clear that there are still some movements in inflation that need to be explained by the model, but one needs to be careful to overfit the target with neural networks. The main contributor to inflation is definitively long-run expectations. Deviations from long-run expectations are mostly driven by changes in short-run expectations, with a supporting role played by real activity.

Figure 3: Historical decomposition of inflation estimate



Note: estimation ends in 2020, after which the contributions are projected.

Source: authors' elaboration.

Inflation has made a comeback in many developed countries and there is reason to believe that elevated inflation rates are the result of an increase in real activity. The training sample for this model does not cover the last few quarters of 2022, but one can already see from the forecast beyond the training sample that increasing inflation is predicted to be driven by real activity and potentially commodities in South Africa. In a year or two we might be able to see whether the higher inflation currently experienced since the COVID-19 pandemic (annualized inflation rates of 7 per cent and above since mid-2022) has also meant a dislodging of long-run expectations.

6 Conclusion

We have shown that it is not empirically appropriate to claim that the slope of the Phillips curve has permanently flattened. It is clear from our results that long-run expectations were historically anchored

at the top of the SARB's inflation target band, but more recently this anchoring has started to shift to the midpoint of the band. Deviations from this long-run expectation do occur quite frequently, especially in recent years, and are primarily driven by short-run expectations and changes in real activity. While the effect of the different components that make up the Phillips curve has dampened over time, especially for short-run expectations, it is not correct to claim that the relationship has disappeared. The results from our model confirm that the slope of the Phillips curve has a time-varying nature, which might have obscured results from the literature that purely employ constant coefficient techniques. The resurgence of the contribution of real activity to inflation shows that while the Phillips curve might have been dormant for a while, it is too soon to pronounce it dead.

References

- Botha, B., Kuhn, L., and Steenkamp, D. (2020). 'Is the Phillips Curve Framework Still Useful for Understanding Inflation Dynamics in South Africa?'. Working Paper 10211. Pretoria: South African Reserve Bank. Available at: <https://ideas.repec.org/p/rbz/wpaper/10211.html> (accessed in June 2023)
- Chan, J. C. C., Koop, G., and Potter, S. M. (2016). 'A Bounded Model of Time Variation in Trend Inflation, Nairu and the Phillips Curve'. *Journal of Applied Econometrics*, 31(3): 551–65. <https://doi.org/10.1002/jae.2442>
- Dladla, P., and Malikane, C. (2022). 'Inflation Dynamics in an Emerging Market: The Case of South Africa'. *Economic Analysis and Policy*, 73(March): 262–71. <https://doi.org/10.1016/j.eap.2021.11.016>
- Fedderke, J., and Liu, Y. (2018). 'Inflation in South Africa: An Assessment of Alternative Inflation Models'. *South African Journal of Economics*, 86(2): 197–230. <https://doi.org/10.1111/saje.12192>
- Forbes, K., Gagnon, J., and Collins, C. (2021). 'Low Inflation Bends the Phillips Curve around the World'. NBER Working Paper 29323. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w29323>
- Galí, J., and Gambetti, L. (2019). 'Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration'. NBER Working Paper 25476. Cambridge, MA: National Bureau of Economic Research. <https://doi.org/10.3386/w25476>
- Goulet Coulombe, P. (2022). 'A Neural Phillips Curve and a Deep Output Gap'. SSRN Electronic Journal. <https://doi.org/10.2139/ssrn.4018079>
- Goulet Coulombe, P., Leroux, M., Stevanovic, D., and Surprenant, S. (2021). 'Macroeconomic Data Transformations Matter'. *International Journal of Forecasting*, 37(4): 1338–54. <https://doi.org/10.1016/j.ijforecast.2021.05.005>
- Kabundi, A., and Rapapali, M. (2019). 'The Transmission of Monetary Policy in South Africa Before and After the Global Financial Crisis'. *South African Journal of Economics*, 87(4): 464–89. <https://doi.org/10.1111/saje.12238>
- Kabundi, A., and Schaling, E. (2013). 'Inflation and Inflation Expectations in South Africa: An Attempt at Explanation'. *South African Journal of Economics*, 81(3): 346–55. <https://doi.org/10.1111/saje.12007>
- Kabundi, A., Schaling, E., and Some, M. (2015). 'Monetary Policy and Heterogeneous Inflation Expectations in South Africa'. *Economic Modelling*, 45(February): 109–17. <https://doi.org/10.1016/j.econmod.2014.11.002>
- Kabundi, A., Schaling, E., and Some, M. (2019). 'Estimating a Phillips Curve for South Africa: A Bounded Random-Walk Approach'. *International Journal of Central Banking*, 15(2): 75–100.
- Reid, M., and Siklos, P. (2021a). 'The BER's Inflation Expectations Surveys: Findings from Academic Research'. BER Research Note 4. Stellenbosch: Bureau for Economic Research.
- Reid, M., and Siklos, P. (2021b). 'Inflation Expectations Surveys: A Review of some Survey Design Choices and their Implications'. *Studies in Economics and Econometrics*, 45(4): 283–303. <https://doi.org/10.1080/03796205.2022.2060299>
- Reid, M., and Siklos, P. (2022). 'How Firms and Experts View The Phillips Curve: Evidence from Individual and Aggregate Data from South Africa'. *Emerging Markets Finance and Trade*, 58(12): 3355–76. <https://doi.org/10.1080/1540496X.2022.2045940>
- Reinbold, B., and Wen, Y. (2020). 'Is the Phillips Curve Still Alive?' *Review*, 102(2): 121–44. <https://doi.org/10.20955/r.102.121-44>
- Vermeulen, J. C. (2017). 'Inflation and Unemployment in South Africa: Is the Phillips Curve Still Dead?' *South African Business Review*, 21(1): 20–54.

Appendix

Table A1: Benchmark model variables: real sector

Real quarterly variables	
Variable	Description
KBP7000L	Employment in the Public sector: General government
KBP7001L	Employment in the Public sector: Business enterprises
KBP7002L	Total employment in the Public sector:
KBP7003L	Employment in the private sector: Mining
KBP7004L	Employment in the private sector: Manufacturing
KBP7005L	Employment in the private sector: Construction
KBP7006L	Employment in the private sector: Trade
KBP7007L	Employment in the private sector: Financial institutions
KBP7008L	Total employment in the private sector
KBP7009L	Total employment in the non-agricultural sectors
KBP6006D	Gross domestic product at market prices
KBP6007D	Final consumption expenditure by households: Total
KBP6109D	Gross fixed capital formation: Private business enterprises (Investment)
KBP6100D	Gross fixed capital formation: General government (Investment)
KBP6008D	Final consumption expenditure by general government
KBP6009D	Gross fixed capital formation
KBP7014L	Labour productivity in the non-agricultural sectors
KBP7079L	Manufacturing: Labour productivity
KBP7013D	Total remuneration per worker in the non-agricultural sector
KBP7015L	Nominal unit labour costs in the non-agricultural sectors
Real monthly variables	
Variable	Description
KBP7082T	Manufacturing: Orders and sales: Sales
KBP7086T	Trade: Retail sales
KBP7087T	Trade: Wholesale sales
KBP7202M	Ratio of inventories to sales in manufacturing and trade
KBP7203M	BER: Constraints on current manufacturing activities: shortage of raw materials
KBP7204M	BER survey: Manufacturing, stocks of finished goods / expected demand
KBP7196M	Transnet: Total cargo handled at ports in South Africa
KBP7064T	Buildings completed
KBP7060N	Mining production: Gold
KBP7061N	Mining production excluding gold
KBP7067N	Trade: Number of new vehicles sold
KBP7068N	Electric current generated
KBP7091N	Coincident indicator of South Africa

Table A2: Benchmark model variables: price and commodities

Price variables	
Variable	Description
KBP7160N	Consumer prices of services: Housing and utilities
KBP7162N	Consumer prices of services: Health
KBP7163N	Consumer prices of services: Transport
KBP7164N	Consumer prices of services: Communication
KBP7147N	Consumer prices of goods: Clothing and footwear
KBP7145N	Consumer prices of goods: Food and non-alcoholic beverages
KBP7170N	Total consumer prices
KBP7155N	Consumer prices of goods: Total
KBP7169N	Consumer prices of services: Total
KBP7151N	Consumer prices of goods: Transport
KBP7198M	PMI: Prices
KBP7201M	Shipping rates: Baltic Dry Index
Price expectations variables	
Variable	Description
KBP7114K	Inflation expectations: Financial analysts: Current year
KBP7115K	Inflation expectations: Financial analysts: One year ahead
KBP7117K	Inflation expectations: Business representatives: Current year
KBP7118K	Inflation expectations: Business representatives: One year ahead
KBP7120K	Inflation expectations: Trade union representatives: Current year
KBP7121K	Inflation expectations: Trade union representatives: One year ahead
Commodity variables	
Variable	Description
KBP5343M	Platinum price in US dollar
KBP5344M	Brent crude oil price in US Dollar
KBP5356M	London gold price in rand
KBP5346M	Platinum price in Rand
KBP5349M	Brent crude oil price in Rand