

OBEN 1802 - September 2018

Update to the Quarterly Projection Model

Byron Botha, Franz Ruch and Rudi Steinbach

Abstract

This note discusses important changes made to the Bank's Quarterly Projection Model (QPM). Firstly, the uncovered interest parity (UIP) relationship has been modified in order to simplify the exchange rate channel in the model. Secondly, trends in relative prices have been removed from the model. As a result, it is no longer assumed that certain prices settle at inflation rates different to the inflation target. In the Phillips curves, unit labour costs have replaced real wages, as the former is more effective at explaining inflation outcomes. Finally, the model now includes a mechanism to account for the persistently high inflation expectations observed in the economy..

1. Introduction

Model development is a continuous process. In this note we highlight recent changes made to the Bank's Quarterly Projection Model (QPM). These include reducing complexity around structure, simplifying the uncovered interest parity (UIP) condition that produces forecasts of the exchange rate, addressing the problem of model versus real life inflation expectations, and improving the labour market structure.

2. Reducing QPM complexity

A model at its most fundamental level is a simplified description of reality, designed to answer questions that are of interest to its stakeholders. Adding complexity that does not fulfil this objective is undesirable as it makes the model more difficult to understand and more costly to operate - even if it increases its realism.


Two features in the existing model that have been adding unnecessary complexity to the model are relative price trends and the UIP relationship.¹ After forecasting with the model for a number of MPC meetings, and following much discussion between the forecasting team, development team, and the Policy and Research Division at large, it has become clear that these features are not value-adding.

2.1. Removing relative price trends

We start by explaining the complexity brought about by allowing for relative price trends in the model. Prices do not move together: for instance, food inflation often behaves differently relative to headline inflation. The role of relative prices in the model is twofold. First, trends in relative prices have an economic interpretation that affects the consumption patterns of consumers. When a good becomes more expensive relative to all other goods, consumers will eventually substitute away from that good. This should lead to lower demand for the relatively expensive good and, ultimately, its relative price should decrease to a point where demand is restored. However, in practise, this theoretical mechanism sometimes fails.

¹ These features of the model were published in the Bank Working Paper Series: *The Quarterly Projection Model of the SARB - WP/17/01*, hereafter referred to as the "published model".


Figure 1: The relative price of electricity



The best example of this problem is electricity prices. Following the 2008 realisation that electricity production was not sufficient to meet demand, prices of electricity rose significantly, as indicated in Figure 1. Theory would state that demand for electricity would decrease as consumers shift away from electricity to other goods, and the price would decrease. Of course demand did ultimately decrease for electricity but the price, being regulated, perversely increased further.

Second, relative price trends allow for the different components of inflation to settle at different long-term values. For instance, the steady-state inflation rate for electricity might be higher than the 4.5% target – as it has been over the past decade. Despite these relative price variations, the weighted average of the long-term inflation rates for all of the components still needs to add up to 4.5 per cent.² Figure 1 below indicates the current steady state calibration of the various subcomponents of the CPI inflation basket.

Figure 2: Current inflation targets in the QPM



This approach ought to be useful since it is a better description of the data and allows the “model agent” to have overall inflation expectations which are anchored at 4.5 per cent, without having an expectation of electricity price inflation at 4.5 per cent. This distinction ensures that policymakers understand that

² Using the weight of the components in the consumer price basket.

getting to the inflation target can mean having to steer inflation that they can control (mainly core inflation) to a rate below the midpoint of the target range to account for a higher mean in inflation they have less control over (predominantly food and energy inflation).

The problem with this setup, however, comes with the constraint already mentioned. If we set the equilibrium values of food, electricity, fuel to their long-term averages (5.5, 8, and 5.5 in the published model), to stabilise headline inflation at 4.5 per cent would imply core inflation of 4.0 per cent. For core inflation, 4.0 per cent is well below its long-term average of 5.3 per cent. In addition, underlying the core inflation steady state, are separate targets for services and core goods inflation. They are 4.3 and 3.5 per cent, respectively.

The relative price structure also complicates the forecast. If inflation is at 4.5 per cent, but the relative prices are not in equilibrium, inflation will also not be in equilibrium. The current rise in fuel price inflation – which is anticipated to exceed 20 per cent by the 3rd quarter of 2018 – serves as a useful example. If these high rates of fuel inflation were to persist, a wedge would form between the actual fuel price and its relative price trend (of 5.5 per cent growth). Eventually, the fuel price pressures should subside, but the relative fuel price could well remain above the trend. To eliminate this wedge, the model will then need to generate fuel price inflation below 5.5 per cent. This fall in fuel price inflation will put downward pressure on headline CPI, and cause it to fall below the midpoint until the wedge has dissipated. Overall, these dynamics around relative price trends lead to a very persistent inflation process.

Our solution to these problems is to do away with these relative price trends and to make the steady states of all the components of inflation 4.5 per cent. This avoids having to make arbitrary trade-offs around which sub-components of headline should be tethered to historical data. It also facilitates forecast communication, and may contribute toward influencing future outcomes to achieve the 4.5 per cent target. Finally, removing the relative price structure requires the estimation of five fewer unobserved equilibrium processes and as many “wedges” in the expectations channel, simplifying the forecasting process.

The implications of this change for the model’s dynamics are shown graphically at the end of the next section.

2.2. Simplifying the UIP relationship

In its simplest form the uncovered interest rate parity relationship says that the expected change in the exchange rate (appreciation or depreciation) is equal to the difference between home and foreign interest rates adjusted for the country risk premium. Restated in symbols we have:³

$$Es_{t+1} - s_t = (i_t - i_t^* - prem_t)/4 \quad (1)$$

where s_t is the nominal exchange rate expressed as local currency units per unit of foreign currency, i_t and i_t^* are domestic and foreign policy interest rates, and $premt$ is the domestic country risk premium. Compare that to the form in the published model:

$$\begin{aligned} \Phi_t &= e_1 Es_{t+1} + (1 - e_1) \left(s_{t-1} + \frac{2}{4} (\Delta \bar{z}_t + \pi^{target} - \pi^{target*}) \right) \\ \Phi_t - s_t &= [e_2(i_t - i_t^* - prem_t) + (1 - e_2)(\Delta \bar{z}_{t-1} + \pi_{t-1} - \pi_{t-1}^*)]/4 \end{aligned} \quad (2)$$

³ The components of the interest rate differential are divided by four since they are expressed in annualised terms, while being related to the quarterly change in the exchange rate.

Here, $\Delta\bar{z}_t$ is the real exchange rate's trend depreciation, π_t and π_t^* are domestic and foreign CPI inflation rates, while π^{target} and π^{target^*} are their respective inflation targets. If you look carefully you can see that the simple form is still embedded in Equation (2), but the increase in complexity is definitely the most striking feature. This is the result of an unintended consequence. The second equation was designed to dampen and prolong exchange rate movements and their effects on inflation, so as to more closely approximate observed outcomes.⁴ The resulting complexity of the equation was, however, under-appreciated at the time. While the simplest version of the equation remains inadequate, as it fails to capture something that is important to policy makers, some simplification is possible.

We propose the following simplification:

$$\left\{ e_1 E s_{t+1} + (1 - e_1) \left(s_{t-1} + \frac{2}{4} (\Delta\bar{z}_t + \pi^{target} - \pi^{target^*}) \right) \right\} - s_t = (i_t - i_t^* - prem_t)/4$$

The expectations mechanism here is such that expectations are formed partly by some “forward-looking” agents who have model consistent expectations, $E s_{t+1}$, and partly by other “backward-looking” agents who base their exchange rate expectations on a forecast of the currency implied by purchasing power parity, $s_{t-1} + \frac{2}{4} (\Delta\bar{z}_t + \pi^{target} - \pi^{target^*})$. It is hardly a simple equation, but some of the complexity is removed while with careful calibration a similar transmission channel is achieved.

Combined, the two modifications above lessen the persistence of inflation in the model. Although the reduction in persistence is largely due to the removal of relative price trends – as shocks to these prices no longer have a distinct channel to persist – the new UIP condition also contributes. In addition, the interest rate differential's increased role in the UIP condition leads to a stronger exchange rate response following country risk or interest rate changes, domestically or abroad. In turn, this leads to a larger inflation response than before.

3. Improving the explanatory power of the Phillips curve


The Phillips curve is perhaps the most fundamental equation in the QPM as it captures the relationship between inflation and its key drivers: inflation expectations, slack in the economy, the exchange rate, and the cost of labour. It is this relationship that monetary policy hopes to exploit to achieve its inflation objectives.

Two innovations to the equation are presented. The first is the replacement of the real wage gap with a real unit labour cost (ULC) gap to capture the impact of pricing pressures in the labour market on inflation. The second innovation is the linking of inflation expectations to observed information about the inflation expectations process. Our method is similar to that used in the UIP relationship discussed in the previous section: total inflation expectations are partly model consistent, and partly based on observed expectations – a more backward-looking process. Here we use the survey data of inflation expectations collected by the Bureau of Economic Research (BER).⁵

⁴ It is a feature of general equilibrium (GEM) models that the transmission of changes in the interest rate to inflation, through the exchange rate, are too large and are over too quickly to be consistent with the data - which is why adjustments are often required.

⁵ <https://www.ber.ac.za/BER%20Documents/Inflation-Expectations/?doctypeid=1065> [Online: 2018-06-25]

Figure 4: QPM labour market



3.1. Adding real unit labour costs

In making the decision to add unit labour costs as an operational variable in the QPM we first looked at the forecast performance of different measures of labour cost pressures onto inflation. Recent work done at the SARB finds that ULC is a better predictor of inflation than average wages, which is the current measure of labour market price pressure in the QPM. This finding motivated the replacement of average wages with ULC as the proxy for the firms' production cost of labour in the model. A second motivation for the change is that real ULC is easier to interpret, as opposed to the currently-used concept of real wages that are adjusted for medium-term productivity.

Figure 4 provides a snapshot of the labour market from the QPM's perspective. It shows how productivity declined over the 2000s, turning negative after the crisis. Adjusting real wages for this productivity trend shows how the ULC gap was inflation neutral through most of the 2000s, but became strongly inflation positive post-crisis.

Figure 5a and 5b show the contribution of current outcomes in ULC to the deviations of headline and core inflation from target.⁶ When looking at core inflation alone, one can easily see that ULC has been a potent driver of underlying inflation, keeping it persistently above the target since the global financial crisis.

⁶ In a sense it underestimates the contribution because expected future ULC pressures are not included.

Figure 5a: Headline inflation and labour cost pressures





Figure 5b: Core inflation and labour cost pressures



3.2. Explaining SAs stubbornly high inflation expectations within the QPM

In March 2016, the MPC provided guidance that the QPM should use 4.5 per cent as its inflation target. Firstly, the QPM cannot model a range of 3 to 6 per cent and requires a point target. Secondly, if average inflation were anchored closer to the mid-point, there would be more room to manoeuvre in the event of adverse supply shocks. Yet, based on surveys of longer-run inflation expectations, it does not

appear to be the inflation target in the minds of wage and price setters. This affects the realism of the model forecasts.


In the QPM, the inflation target serves as the anchor for inflation expectations, which is a fundamental component of the Phillips curve that describes the model's inflation process. Underlying the Phillips curve, are consumers and businesses whose pricing decisions are based on the belief that, in the long run, actual inflation will be whatever the inflation target is. In South Africa, however, inflation expectations have typically been closer to the upper end of the target range. As a result, there is a wedge between inflation expectations in the model and actual expectations of price setters in the economy. This allows inflation to settle at 4.5 per cent implausibly easily, underestimating the effort required from monetary policy to achieve this outcome.

In this section we motivate for an augmented Phillips curve specification that accounts for the wedge between actual inflation expectations observed in the South African economy, and those of the model. These proposed changes are supported by recent literature which argues that the inclusion of subjective expectations – such as those measured in surveys – in the Phillips curve relationship could improve the degree to which it fits actual inflation dynamics.⁷

3.2.1 A brief primer of inflation expectations

South Africa's inflation expectations have been well studied and are primarily backward looking (see Figure 6a). This property increases the short-term costs of reaching the inflation target. Reducing the short-term disinflation costs of monetary policy requires shifting these expectations so that they are forward looking and have the property of not moving up when there is a temporary shock to inflation.

Figure 6a. Actual inflation and business' inflation expectations




The exchange rate shock that occurred in 2008 provides a useful example. Figure 6a shows how inflation expectations shifted up in the period that followed. Through a process of wages adjusting higher which raised firms' costs, prices of all other goods and services in the economy rose. The second-round effects created by this process are highlighted in Figure 6b. As long as inflation expectations remain backward looking, shocks to food and energy prices (which have average 8.6%

⁷ See Coibon, O., Gorodnichenko, Y., Kamdar, R. The formation of expectations, inflation and the Phillips curve. *Journal of Economic Literature*, forthcoming.


since 2000) will become entrenched in future wage and inflation expectations. So even if we were in 2022, and had achieved inflation outcomes of 4.5% and inflation expectations shifted down, the next significant supply shock will knock inflation expectations off track and be more costly for policymakers to fight. This familiar process has accompanied most of our supply shocks, resulting in core inflation rising and then eventually moderating to about 5.5%, the expressed inflation expectations of agents.

Figure 6b. Simplified representation of second-round effects on core inflation



A second consequence of backward-looking inflation expectations is that price-setters do not pay attention to what the Bank says about future inflation. Figure 7 shows the SARB’s forecast for targeted inflation two-years-ahead against the same expectations of business. On average, business inflation expectations are 1 percentage point above SARB forecasts. Although the SARB forecasts were not accurate, they contained no systematic bias unlike business’s expectations.

Figure 7: SARB and business inflation expectations



Modeling inflation expectations with backward- and forward-looking components means that the QPM can account both for where we are now – an implied target above 4.5% and backward-looking inflation expectations – and where we want to go – forward looking price-setters and inflation at 4.5 %.

3.2.2 Breaking rational expectations


As Coibon, Gorodnichenko and Kamdar (forthcoming) argue, augmenting models with price setters' actual expectations can improve model fit. The Bureau for Economic Research's (BER) inflation expectations survey contains the requisite information. We incorporate this in the existing Phillips curve as follows:

$$\pi_t = a_1 * [b_1 * E\pi_{t+1}^{model} + (1 - b_1) * E\pi_{t+1}^{BER}] + (1 - a_1) * \pi_{t-1} + a_2 * rmc_t \quad (3)$$

where π_t is inflation, π_{t+1}^{model} is the model's expectation for future inflation, rmc_t is real marginal costs that capture inflationary pressure from slack in the economy, the real exchange rate and real unit labour costs, and a_1 , a_2 and b_1 are parameters that govern the impacts of these drivers. The modification to Equation (3) is the term π_{t+1}^{BER} , the 2-year ahead expectations from the BER survey.

Figure 8 shows the impact that this modification has on the model's projected path following an announcement to disinflate from a target of 6 to 4.5 per cent. It is assumed that the announcement to disinflate to the lower target is made in 2018Q2. The red line represents the baseline model's path, and is compared to a similar disinflation announcement where we use inflation expectations as measured by the BER (the blue and green lines). Here, the green line represents a variation of the modified model in which BER expectations are more persistent and therefore moderate more slowly.

Figure 8: Disinflation in baseline QPM vs disinflation with BER expectations




In the baseline model, price setters are rational and automatically anchor their long-run inflation expectations to the Bank's new midpoint objective. Inflation then follows. Where we use rational expectations augmented with actual expectations from BER survey data, the disinflation process is much slower. Inflation is kept higher by actual expectations, which are slow to adjust and require actual inflation to fall on a sustainable basis before moving lower.

4. Overall change to the QPM's dynamics

Having introduced all the changes discussed in this note, we are able to compare how the updated model's dynamics compare to those of the prior version. Overall the updated model's dynamics are broadly similar. This can be seen in Figure 9, where a shock to core inflation raises it by 1 per cent at the end of the first year. It is also assumed that BER expectations move to the same extent.

In response to the supply shock, monetary policy reacts by increasing the repo rate. Although the magnitude of this response is broadly similar in both versions, the repo rate returns to equilibrium more quickly as the inflation process is generally less persistent. In addition, the exchange rate and output are more responsive to the increase in the repo rate, which contribute to the relatively accelerated lowering of inflation.

Figure 9: Overall impact on model dynamics (1% shock to core inflation)



5. Conclusion

Good maps are simple and up-to-date. The same goes for economic models. This note describes some important refinements and simplifications to the QPM, aimed at making it a better guide for policymakers. Firstly, certain excessive complexities have been removed from the model's structure, such as relative price trends. While these trends allowed for the subcomponents of headline CPI to settle at steady-states that were different from the overall inflation objective of 4.5 per cent, they generally created troublesome dynamics over the forecast horizon. Secondly, the model's proxy for labour market pricing pressure has been changed from average real wages to real unit labour costs. This modification was largely motivated by research indicating that the latter measure has been a more accurate predictor of inflation. Finally, the model's inflation expectation channel has been augmented to account for actual expectations of price setters – which are not at the target of 4.5%, as previous versions of the QPM required. Overall, these changes simplify the forecasting process, and allow for more clarity in communication of the model's results.

Appendix

A 1. Labour market structure in the updated QPM

The current QPM uses a New Keynesian Phillips curve to model the dynamics of inflation. The New Keynesian Phillips curve expresses current inflation as a function of the inflation rate that price setters expect to hold in the future, past inflation outcomes, as well as price pressures that emanate from changes in the real costs of production – commonly referred to as real marginal costs. In the model, these real marginal costs are generally driven by average real wages, the real exchange rate, and the output gap.

The standard definition of average real wages is as follows:

$$\Delta r\text{wage}_t = \Delta \text{wage}_t - \pi_t,$$

and simply states that real wage growth, $\Delta r\text{wage}_t$, is nominal wage growth (Δwage_t) adjusted for inflation (π_t). More importantly, it is not the absolute growth rate of average real wages that leads to inflationary pressures, but rather the gap between real wages and their equilibrium level.

The equilibrium real wage in the model is pinned down by the theoretical condition which states that in the long run, real wage growth should equal productivity growth. In this case, productivity growth is defined as growth in output per worker, or GDP growth adjusted for employment growth.

In equilibrium, this definition of productivity growth would therefore be potential GDP growth, g_t , that is adjusted for equilibrium employment growth, $\Delta \overline{\text{emp}}_t$. The equilibrium for real wage growth is therefore specified as follows:

$$\Delta \overline{r\text{wage}}_t = \Delta \overline{\text{prod}}_t + \varepsilon_t^{\Delta \overline{r\text{wage}}},$$

where equilibrium employment growth, $\Delta \overline{\text{prod}}_t$, is defined as $(g_t - \Delta \overline{\text{emp}}_t)$, and $\varepsilon_t^{\Delta \overline{r\text{wage}}}$ is a shock that can move the equilibrium real wage growth temporarily.

Having defined both the real wage and its equilibrium, the real wage gap – which features in the real marginal cost component of the model's Phillips curves – is expressed as the difference between the average real wage and its equilibrium value:

$$\widehat{r\text{wage}}_t = r\text{wage}_t - \overline{r\text{wage}}_t.$$

In the current version of the QPM, the real wage gap is defined in terms of average real wages that are adjusted for equilibrium productivity. This measure is a slight modification to the more familiar real unit labour cost definition, which expresses real wages that are adjusted for contemporaneous productivity. However, few would think about unit labour costs from the perspective of equilibrium productivity, which is both difficult to define and often deviates significantly from the contemporaneous productivity growth that is most likely determining actual wage outcomes. By defining a productivity gap as the deviation from its equilibrium level:

$$\widehat{\text{prod}}_t = \text{prod}_t - \overline{\text{prod}}_t,$$

it is possible to transform the real wage gap into a unit labour cost gap, $\widehat{\text{rulc}}_t$, by accounting for the deviation of productivity from its equilibrium level, $\widehat{\text{prod}}_t$:

$$\widehat{\text{rulc}}_t = \widehat{r\text{wage}}_t - \widehat{\text{prod}}_t.$$