


Second-round effects on inflation, and underlying inflation

by

Franz Ulrich Ruch

*Dissertation presented for the degree of Doctor of Philosophy
(Economics) in the Faculty of Economics and Management
Sciences at Stellenbosch University*



Department of Economics,
University of Stellenbosch,
Private Bag X1, Matieland 7602, South Africa.

Promoter: Prof. Stanislaus Alexander du Plessis
Dr. Monique Brigitte Reid

December 2016

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

This dissertation includes 3 original papers published in peer-reviewed journals or books and 1 unpublished publications. The development and writing of the papers (published and unpublished) were the principal responsibility of myself and, for each of the cases where this is not the case, a declaration is included in the dissertation indicating the nature and extent of the contributions of co-authors.

Date: December 2016

Copyright © 2016 Stellenbosch University
All rights reserved.

With regard to chapter 3, “Forecasting South African Core inflation”, the nature and scope of my contribution were as follows:

Nature of Contribution	Extent of Contribution
<ol style="list-style-type: none"> 1. All forecasting models included in Tables 3.2 and 3.3 excluding those listed below explicitly implemented by co-authors. 2. Redrafted introduction. 3. Section 3.2 on the discussion of core inflation. 4. Redrafted section 3.3 on methodology. 5. Section 3.4 on data including all data collection work. 6. Redrafted section 3.5 on results. 7. Results section 3.5.1 on determining the number of factors for the Factor augmented model, and section 3.5.3 on forecasting annual core inflation. 	55%

The following co-authors have contributed to chapter 3, “Forecasting South African Core inflation”:

Name	e-mail address	Nature of Contribution	Extent of Contribution
Mehmet Balcilar	mehmet@mbalcilar.net	<ol style="list-style-type: none"> 1. Forecasting of benchmark models including random walk, structural break, autoregressive and vector autoregressive (using OLS) models. See Tables 3.2 and 3.3. 	10%
Rangan Gupta	rangan.gupta@up.ac.za	<ol style="list-style-type: none"> 1. Paper conceptualisation including sourcing matlab code to implement modelling strategies 2. Forecasting of benchmark models including random walk, autoregressive and vector autoregressive models. See Tables 3.2 and 3.3. 3. Forecasting of time-varying parameter models with time variation in intercept only. See footnote 10. 	15%

		4. Writing up part of results section including sections of page 58 discussing the importance of persistence in the literature.	
Mampho Modise	mampho.modise@treasury.gov.za	<ol style="list-style-type: none"> 1. Wrote a draft version of the introduction (section 3.1) 2. Wrote a draft version of the methodology including sections 3.3.1, 3.3.2, and 3.3.3. 3. Wrote a draft version of section 3.5 including parts of 3.5.2. 	20%

Declaration by co-authors:

The undersigned hereby confirm that

1. the declaration above accurately reflects the nature and extent of the contributions of the candidate and the co-authors to chapter 3, "Forecasting South African Core inflation"
2. no other authors contributed to chapter 3, "Forecasting South African Core inflation" besides those specified above, and
3. potential conflicts of interest have been revealed to all interested parties and that the necessary arrangements have been made to use the material in chapter 3, "Forecasting South African Core inflation" of this dissertation.

Signature	Institutional affiliation	Date
	Eastern Mediterranean University, Famagusta, Northern Cyprus	20-06-2016
	University of Pretoria	20-06-2016
	National Treasury	20-06-2016

With regard to chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics", and chapter 5, "Decomposing inflation using micro-price-level data: sticky-price inflation" the nature and scope of my contribution were as follows:

Nature of Contribution	Extent of Contribution
1. All text and figures.	90%
2. All coding in R except for what is mentioned explicitly below.	

The following co-authors have contributed to chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics", and chapter 5, "Decomposing inflation using micro-price-level data: sticky-price inflation":

Name	e-mail address	Nature of Contribution	Extent of Contribution
Neil Rankin	neilrankin@sun.ac.za	<ol style="list-style-type: none"> 1. Conceptualisation of Chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics". 2. Initial work to set-up the data in R (econometrics software). 3. Initial work to code the frequency of price change in R. 4. Comments on structure and content of chapters 4 and 5. 	10%

Declaration by co-authors:

The undersigned hereby confirm that

1. the declaration above accurately reflects the nature and extent of the contributions of the candidate and the co-authors to chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics", and chapter 5, "Decomposing inflation using micro-price-level data: sticky-price inflation"
2. no other authors contributed to chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics", and chapter 5, "Decomposing inflation using micro-price-level data: sticky-price inflation" besides those specified above, and
3. potential conflicts of interest have been revealed to all interested parties and that the necessary arrangements have been made to use the material in chapter 4, "Decomposing inflation using micro-price-level data: South Africa's pricing dynamics", and chapter 5, "Decomposing inflation using micro-price-level data: sticky-price inflation" of this dissertation.

Signature	Institutional affiliation	Date
NEIL RANKIN	STELLENBOSCH	20/06/2016



PostNet Suite 241
Private Bag X16
Constantia
7848

28 June 2016

University of Stellenbosch, Dept. of Economics
Private Bag X1
Matieland
7599

Attention: Mr Franz Ruch

Dear Franz,

CERTIFICATION: COPY-EDITING OF YOUR ECONOMICS PHD THESIS

I wish to certify that I have copy-edited the economics PhD thesis entitled 'Second-round effects to inflation and issues of core inflation'.

Further, in terms of my background and experience, I wish to point out the following.

I founded WriteArt in 2004 and have been working as a full-time professional copy-editor and proofreader since then.

My copy-editing portfolio includes many theses for the students of many universities, in South Africa and abroad. In addition, it includes numerous articles and papers for academics, also from many universities worldwide.

I have also copy-edited and proofread a variety of other types of texts and publications for a range of clients.

Please contact me if you have any queries.

Yours sincerely,

Russell

Abstract

Second-round effects on inflation, and underlying inflation

F.U. Ruch

Department of Economics,

University of Stellenbosch,

Private Bag XI, Matieland 7602, South Africa.

Dissertation: PhD (Econ)

December 2016

Supply shocks, especially food and energy price shocks, play a significant role in the evolution and dynamics of headline consumer price inflation in South Africa. Although headline inflation is the price index officially targeted by the South African Reserve Bank, monetary policy can neither control the relative price movements, nor would it be desirable for the central bank to do so. It is only when these relative price shocks affect the underlying trend rate of inflation – core inflation – through second-round effects that monetary policy has a critical role to play. The presence of second-round effects change how a central bank needs to respond to relative price shocks. Generally, a central bank can look through shocks to food and energy prices, in the absence of second-round effects, communicating clearly the reason for its inaction. However, when second-round effects are present, the central bank has to respond appropriately to ensure that inflation expectations remain anchored around the target.

Second-round effects emanate from the ability of price-setting firms and wage-setting labour to increase prices (whether through mark-ups or higher marginal costs) and wages, and therefore general prices of goods and services, in response to relative price shocks. In order for monetary policy to adequately respond to second-round effects, these effects need to be identified and quantified. Such identification depends on the definition of core inflation, the underlying trend in overall inflation, and the consequent measurement of this core inflation. This PhD dissertation contributes to the academic literature and policy discussion of second-round effects and underlying, or core, inflation in South Africa.

First, the impact of second-round effects on inflation following supply shocks will be quantified using a Structural Bayesian vector autoregressive model, with sensible zero and sign restrictions. This identification relies on a conventional exclusion-based measure of core inflation

– *headline CPI less food and energy* – that is often used in policy discussions and decision-making. The results of this model confirm the impact of wage-setters in South Africa, that changes in the prices of food, petrol and energy are accommodated and lead to strong second-round effects.

Second, monetary policy in South Africa is forward-looking and requires forecasts of inflation to set policy. To ensure the best possible estimates of core inflation are available to the central bank, we look at a host of possible models that existing literature shows to have some success in forecasting, and that cover a wide variety of new techniques. These include models that take account of large datasets of information, that address possible breaks in the inflation series as monetary policy regimes change, that address the changing relationship between variables and inflation or the structure of the economy, and that provide mechanisms to look at the importance of volatility. The myriad of forecasting techniques reveal that accounting for changing relationships improve forecasts of core inflation, while exploiting more economic information does not necessarily produce better forecasts compared with smaller models.

Last, it may be that a conventional definition of core inflation currently used by central banks – *headline CPI less food, non-alcoholic beverages and energy*, the exclusion-based measure most quoted by the South African Reserve Bank – is not the most appropriate, or theoretically consistent, to ensure best policy outcomes. To address this, a novel approach to the definition and calculation of core inflation will be followed using micro-price-level data. We study the underlying dynamics of 5,200,466 individual price quotes of goods to determine the frequency of price changes at a product level. This is used to construct a sticky-price goods inflation measure. Sticky-price goods inflation is more persistent, less volatile and correlates well with future goods inflation. The advantage of sticky-price inflation is that it grounds the concept of underlying inflation into the theoretical framework currently used by central banks to make policy decisions, and what is considered optimal policy by monetary theorists, making it an ideal core inflation candidate for the central bank.

Acknowledgements

A PhD is a long and arduous journey full of self-doubt, losing battles, taking a few steps back and suffering. I would like to express my sincere gratitude to those whose fortitude and guidance helped me win the war. First, I would like to thank my supervisors, Stan du Plessis and Monique Reid. I enjoyed the many long sessions debating the meaning of specific words, and getting my work to be a narrative rather than a jumble of words.

My co-authors on some of the chapters also deserve praise. Neil Rankin opened up the world of micro-price data to me, and guided me through the initial analysis work. Also for bringing the micro-economist eye to work that remains macro-centric. To Rangan Gupta, Mehmet Balicclair, and Mampho Modise who guided me through the arduous world of macroeconomic forecasting.

I would also like to thank my colleagues at the South African Reserve Bank, and elsewhere, who played some role in me achieving this goal: Byron Botha, Greg Farrell, Logan Rangasamy, Chris Loewald, Shaun de Jager, Rudi Steinbach, Shakill Hassan, Pamela Mjandana, Charl Jooste, Alain Kabundi and Haakon Kavli.

I would also like to thank Rashad Cassim, Chris Loewald, Shaun de Jager, Pamela Mjandana and the Governors Executive Committee for my sabbatical which made the completion of this PhD much easier. And to those colleagues at work who were forced to work harder.

A large chunk of this PhD would have also not been possible without the vision and openness of Statistics South Africa, who have entrusted academics with the micro-price-level data that is used to construct the consumer price index. This would have not been possible without the effort of Joe de Beer, Marietjie Bennett and Patrick Kelly.

Lastly, I would like to thank members of the Monetary Research group at Stellenbosch University, participants at the 2015 Economic Society of South African Conference, and the anonymous referees at Empirical Economics and Economic Research South Africa.

Dedications

This thesis is dedicated to my family who have made me who I am: Mom, Dad, Werner, Louis, and Roslynde. Also to my creator.

Contents

Declaration	i
Abstract	vi
Acknowledgements	viii
Dedications	ix
Contents	x
List of Figures	xii
List of Tables	xiv
1 Introduction	1
2 Second-round effects from food and energy prices: an SBVAR approach	11
2.1 Introduction	11
2.2 Theory of second-round effects	13
2.3 The (generalised) Phillips curve and VAR model	15
2.4 Data	17
2.5 Priors	19
2.6 Identification	20
2.7 Estimation methodology	23
2.8 Results	24
2.9 Robustness	29
2.10 Policy implications	36
2.11 Conclusion	37
3 Forecasting South African core inflation	38
3.1 Introduction	38
3.2 Core inflation	40
3.3 Methodology	50
3.4 Data	55

<i>CONTENTS</i>	xi
3.5 Results	58
3.6 Conclusion	67
4 Decomposing inflation using micro-price-level data: South Africa's pricing dynamics	71
4.1 Introduction	71
4.2 Literature review	73
4.3 The dataset	76
4.4 Properties of price dynamics	78
4.5 What factors explain the frequency of price changes?	84
4.6 Decomposing inflation	87
4.7 How important are sales in pricing conduct and flexibility?	93
4.8 Sensitivity to outliers	98
4.9 Policy implications	99
4.10 Conclusion	100
5 Decomposing inflation using micro-price-level data: Sticky-price inflation	101
5.1 Introduction	101
5.2 Literature review	102
5.3 Decomposing inflation	108
5.4 Measures of core inflation	120
5.5 Comparative performance	124
5.6 Caveats and future work	128
5.7 Conclusion	129
6 Summary	130
6.1 Second-round effects from food and energy prices: an SBVAR approach	131
6.2 Forecasting South African core inflation	132
6.3 Decomposing inflation using micro-price level data: South Africa's pricing dynamics	134
6.4 Decomposing inflation using micro-price data: Sticky-price inflation	136
Appendices	138
A Stability of the SBVAR model	139
B Forecasting performance with $p = 4$	141
C Representation of dataset	142
List of References	144

List of Figures

2.1	Contribution of food and energy prices to headline consumer inflation	17
2.2	The impact of food, petrol and electricity prices	26
2.3	The impact of ULC on core inflation	27
2.4	Model versus parameter uncertainty	30
2.5	Sensitivity to prior selection: ULC	31
2.6	Sensitivity to prior selection: core inflation	32
2.7	Sensitivity to prior hyperparameters: ULC	33
2.8	Sensitivity to prior hyperparameters: core inflation	34
2.9	Sensitivity to lag length: ULC	35
2.10	Sensitivity to lag length: core inflation	36
3.1	Estimating the number of factors	60
3.2	Estimated dynamic dimension selection (DDS) probabilities of the small, medium and large TVP-VARs	61
3.3	Result Summary: Relative performance gains	62
3.4	Cumulative sum of total predictive log-likelihood $h = 1$	65
3.5	Recursive estimates from small TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	67
4.1	Basic properties of the data	78
4.2	Distribution of price changes	80
4.3	Mean and median frequency of monthly price changes over time	82
4.4	Distribution of frequency of monthly price changes by product over time	83
4.5	Intensive and extensive margins	89
4.6	Annual inflation decomposed into price increases and price decreases	91
4.7	Proportion of sales	94
4.8	Weighted average frequency of price changes: impact of sales	96
4.9	Goods inflation including and excluding sales	98
5.1	Sticky- and flexible-price goods inflation	112
5.2	Correlation of sticky- and flexible-price inflation to overall inflation	116
5.3	Annual goods inflation by frequency groups	117
5.4	Volatility by frequency groups	118

*LIST OF FIGURES***xiii**

5.5	Sticky-price inflation at different threshold frequencies	119
5.6	Signal-to-noise ratio by threshold frequencies	120
5.7	Kernel density estimate of inflation rates at the product and category level	124
5.8	Core inflation measures	125
A.1	Maximum eigenvalue of the parameter matrix on each draw	139

List of Tables

2.1	Variables included in the VAR	18
2.2	Elliott-Rothenberg-Stock DF-GLS test (P-values)	19
3.1	Dataset used in the small, medium and large VARs	57
3.2	MSFE relative to the random walk model for core inflation	69
3.3	Sum of log predictive likelihoods relative to the random walk model for core inflation	70
4.1	Average weighted frequency of price changes by CPI category	85
4.2	Regression results of frequency: long-run coefficients	86
4.3	Decomposing inflation into intensive and extensive margins	90
4.4	Average weighted frequency of price changes by CPI category excluding sales . . .	97
5.1	Weights of sticky- and flexible-price inflation by category (threshold =27.8%) . . .	114
5.2	Properties of sticky- and flexible-price inflation	115
5.3	Mean, median and volatility	127
B.1	MSFE relative to the random walk model for core inflation	141
C.1	Micro-price dataset	143

Chapter 1

Introduction

The Billion Prices Project (BPP) was started in 2008 as a way to improve on the measurement of the consumer price index (CPI). Initially, the project was developed to obtain a true reflection of inflation in Argentina, where it was clear the statistical authorities were understating reality (Cavallo and Rigobon, 2016). Without accurate inflation, monetary policy was consequently unable to adopt an appropriate stance. Recognising the need to develop and build proxies for inflation using new sources of price information, and the importance of price information as a signal of economic outcomes, the project was expanded to 50 different countries. This thesis discusses, in a less dramatic fashion, the problem with the measurement of inflation. It challenges both the modern convention in monetary policy that focuses on headline CPI as its main target variable, and the use of exclusion-based measures – *excluding food and energy* – as an underlying measure of inflation, or core inflation.

Inflation, or the general rise in prices owing to the increase in the supply of money and credit, is unobservable. Observed prices are since used to uncover the process of inflation, but the ability to do so requires additional identifying assumptions. Cost of living indices, of which the CPI is an example, are the most well-developed, and widely used method, to measure inflation. The objective of such indices is to show “the relative change occurring in the monetary cost of those consumers’ goods which are necessary for the maintenance of a certain standard of living” (Konus, 1939:10). The identifying assumptions of the CPI are derived from welfare economics, where consumers maximise utility over a basket of consumption goods. Modern convention implies that ‘inflation’ is associated with the CPI.

The association of the process of inflation with the CPI was further entrenched in modern economic discourse when central banks elected to use headline CPI inflation as its target for monetary policy. The South African Reserve Bank (SARB) was assigned headline CPI as its target variable when inflation targeting was introduced in the early 2000s. Inflation targeting is a framework of monetary policy where an explicit target for inflation is announced, and the central bank targets forecasts of inflation to reach its goal. The rationale for using headline inflation was that a central bank should be concerned with the variable that most directly affects people’s lives. Also, headline inflation was seen as having a communication advantage, since it is easily

understood and generally accepted by the public. It is argued, however, that headline inflation should not be the ultimate goal variable of the central bank, since the bank cannot control headline inflation (Du Plessis, 2014). The normative argument, put forward by Goodfriend (2007), calls for underlying, or core, inflation to be the goal variable.

Gordon (1975) was the first to introduce the word ‘core’ inflation to distinguish between underlying inflation, caused by demand related issues (including monetary policy), and those driven by supply shocks to food and energy¹. Gordon specifically referred to underlying prices as “hard-core”. In an era of large supply-side shocks, his intention was to address how government policy, including monetary policy, should respond so as to maintain macroeconomic stability.

Subsequent theory provides two broad definitions of core inflation expounded in Roger (1998). The first, as a ‘persistence’ concept, builds on earlier work by Friedman *et al.* (1963). Friedman *et al.* (1963:25) highlights two distinct characteristics of inflation, as follows: “... a steady inflation, one that proceeds at a more or less constant rate, and an intermittent inflation, one that proceeds by fits and starts...”, the former being core inflation. The second, as a ‘generalised’ concept defined initially by Eckstein (1981:7) as “... the trend increase of the cost of the factors of production” which “... originates in the long-term expectations of inflation in the minds of households and businesses, in the contractual arrangements which sustain the wage-price momentum, and in the tax system”.

The introduction of the concept of core inflation in the 1970s recognised the need to remove relative price shocks; shocks that are supply-driven, short-lived and often large, from the measure of overall inflation so as to establish a better idea of the underlying inflation process. It took another three decades for a consensus on core inflation as the ultimate goal variable of the central bank when the normative argument was put forward by Goodfriend (2007), in the latter 2000s. He argued that monetary policy had reached a consensus that core inflation rather than headline inflation was the appropriate nominal anchor for a central bank. Core inflation is more stable and would serve as a better anchor for inflation expectations.

Part of the reaching of consensus on core inflation was the development of the theory that showed that core inflation, rather than headline inflation, led to households maximising their welfare. This ‘consensus’ model – with features that include monopolistically competitive firms who set prices in a staggered way, rational expectations, households maximising utility, and a prominent role for monetary policy – was expounded first in Goodfriend and King (1997) and Clarida *et al.* (1999). Within this model, the rationale for not targeting headline inflation is that this would require a policy response to relative price shocks that unnecessarily compounds output losses, i.e. it would force the sticky-price sector to adjust through lower demand and in doing so, decrease prices and wages. In addition, relative price shocks from goods with flexible prices, such as oil, can also be large, meaning that the output-inflation trade-off would be costly.

¹The word ‘core inflation’ does appear earlier in Schreder (1952) but its link to the modern concept is unclear (Wynne, 1999).

The optimality of core inflation is not, however, tied to the New Keynesian paradigm but is a general result in welfare economics. Inflation leads to the highest welfare loss in sectors where prices are more sticky (or more persistent), with few welfare costs when relative price shocks dissipate quickly (Walsh, 2009). Walsh (2009:30) stated that “[s]ince food and energy prices display little stickiness, responding quickly to shifts in demand and supply, there is a strong case for excluding them from the inflation rate the central bank attempts to control”. Therefore a central bank should be concerned with the money part of inflation, that which it can actually affect, and not relative price movements.

Relative price shocks, commonly from food and energy prices, play a significant role in the evolution and dynamics of headline consumer inflation in South Africa (SA). Although headline inflation is the price index officially targeted by the SARB, monetary policy can neither control the relative price movements that arise owing to these shocks, nor would it be desirable for the central bank to do so (apart from explaining their inaction), to allow the economy to adjust. It is only when these prices affect underlying, or core, inflation, through second-round effects, that monetary policy has a critical role to play. The presence of second-round effects change how a central bank needs to respond to relative price shocks. When these effects do not occur, a central bank can look through temporary shocks to food and energy prices and clearly communicate the reason for its inaction. However, when second-round effects are present, the central bank must respond appropriately to ensure that inflation expectations remain anchored around the target.

Price-setting firms and wage-setting labour have the ability to increase prices and wages in response to relative price shocks such as shocks to food or energy prices. Second-round effects capture the extent to which they increase prices² and wages, and pass these on to other goods and services, establishing the process of inflation.

The impact of second-round effects occurs through two channels, namely the cost and expectations channels. The cost channel refers to the direct impact of changes to a firm’s marginal costs owing to an increase in input costs. This assumes that the relative price shock is larger than the menu cost associated with changing output prices. As a familiar example illustrates, recent increases in electricity prices, following the 2008 energy crisis in SA, increased the cost of production by raising the price of intermediate inputs. Electricity price increases were then passed on to the consumer. Consumers would want to be compensated for higher prices and hence will adjust their expectations. The expectations channel refers to the response of wage-setting labour to a relative price shock in, for example, food and energy. If labour perceives the relative price shock to be long-lasting, or has the bargaining power to raise wages, it will raise its inflation expectations and demand higher nominal wages. An increase in the prices of other goods and services follows, either as a result of price adjustments by firms as their marginal costs increase, or due to increasing consumption.

The theoretical framework to understand the channels through which second-round effects increase underlying inflation have been formalised in Aoki (2001), Hlédik and Banka (2003),

²Firms are able to increase prices by either increasing mark-ups or marginal costs.

Bodenstein *et al.* (2008), Blanchard and Gali (2007), and Anand and Prasad (2010). Aoki (2001) provides a model that can be used as a basis from which to understand the transmission of relative price shocks to underlying inflation. He builds a two-sector dynamic stochastic general equilibrium (DSGE) model, distinguishing between a flexible- and sticky-price sector. Flexible-price goods are standardised, traded in an almost competitive market, and used as both an input into production as well as consumed by households. Sticky-price goods are differentiated and traded in a monopolistically competitive environment. The flexible-price good, which creates the relative price shock, impacts on a modified New Keynesian Phillips curve for the sticky-price good. The mechanism through which this occurs is the substitution effect between flexible-price goods and sticky-price goods.

Much of the focus of Aoki (2001), Bodenstein *et al.* (2008), and Anand and Prasad (2010) is on optimal monetary policy in the context of relative price movements. Although this is an important topic, it does not provide estimates of the size and dynamics of second-round effects in an economy. Evidence of the size and dynamics of these effects is limited to Hlédik and Banka (2003), Cecchetti and Moessner (2008), and Baumeister *et al.* (2010), who show that second-round effects are not uniform and differ by country. Rangasamy (2011) and Rangasamy and Nel (2014) provide some evidence of the existence and size of second-round effects in SA.

The first objective of this thesis is to provide a concise framework through which to estimate the impact and dynamics of second-round effects from relative price shocks on the economy. This extends the existing work on relative price shocks in SA, including Rangasamy (2011) and Rangasamy and Nel (2014), which only estimated the cost channel of second-round effects. The most important impact of relative price shocks, however, is its impact on wages, something that has not been quantified in the South African literature. We use recent advances in Bayesian estimation and structural vector autoregressive (VAR) models to estimate a Structural Bayesian VAR with plausible short- and long-run zero restrictions as well as sign restrictions, to identify the shocks. We also entrench the discussion of second-round effects into the framework used by the SARB. This includes using the common exclusion-based core inflation measure of *headline CPI less food and energy prices*. By implication relative prices are associated to food and energy.

The results confirm the impact of wage-setters in South Africa in the expectations channel, that changes in the price of food, petrol and energy are accommodated resulting in strong second-round effects. According to the Structural Bayesian VAR model, a one per cent shock to relative food and energy prices increase wages by 0.3 per cent a year after the shock. The price of other goods and services (or core inflation) increase with a maximum impact of 0.3 per cent, three quarters after the shock. This is attributed to both the cost and expectations channel.

The transmission of second-round effects to core inflation, as well as the forward-looking nature of monetary policy, means that appropriate policy responses require the best possible forecasts of core inflation. Although the SARB targets headline inflation, core inflation is by definition the part of inflation a central bank should be most concerned about. In SA, there is an

extensive literature on forecasting headline inflation including work by Moll (1999), Liu *et al.* (2009), Alpanda *et al.* (2011), Gupta and Kabundi (2011), Gupta and Steinbach (2013), and Gupta *et al.* (2015). Despite the fundamental role for core inflation, however, little work has been done to forecast this variable. International literature on this subject includes Sun (2004), Morana (2007), and Kapetanios (2004), whilst domestic literature includes only Gupta *et al.* (2015).

To expand this body of literature, the second objective of this thesis then is to provide the best possible forecast of core inflation, a topic barely covered in the South African literature despite the forward-looking nature of policy and the signal advantage of core inflation. Models for forecasting core inflation are extended in four important directions. First, recent methodological and computing gains have made it possible to increase the dimensionality of models, solving the omitted variable bias in smaller VARs, to include up to a hundred variables when analysing and forecasting macroeconomic variables. Bańbura *et al.* (2010), Giannone *et al.* (2014), and Carriero *et al.* (2015) show that increasing the number of variables used leads to better forecast accuracy, but found there are limitations. Bańbura *et al.* (2010) and Koop (2013) provide evidence that this limit is around 20 variables. We consider a number of model sizes, including those with up to 21 variables.

Second, a common assumption in simple models of analysis and forecasting is that the errors are homoscedastic. Of course, macroeconomic shocks are not. Engle (1982) first introduced heteroscedastic errors using an autoregressive conditional heteroscedastic (ARCH) process, and showed with this seminal piece, that inflation in the United Kingdom had significant, and changing, volatility especially during the 1970s. This thesis, however, will look at another version of heteroscedasticity called stochastic volatility first introduced to VARs by Uhlig (1997). Both homoscedastic and heteroscedastic error structures will be used.

Third, significant changes in the structure of the South African economy over the last four decades makes it unlikely that relationships between economic variables remained constant, or that there were not any structural breaks. Structural breaks are a significant cause of poor forecasting performance (Stock and Watson, 1996; Ang and Bekaert, 2002; Clements and Hendry, 1998; and Bauwens *et al.*, 2011). To address changing relationships, time-varying parameters are introduced. Primiceri (2005) importantly showed, using time-varying parameter models, that monetary policy has changed over time in the US. One recent example of changing relationships in the South African context is provided by Jooste and Jhaveri (2014), who showed that exchange rate pass-through to inflation is time-varying, and has declined recently. These changing relationships also matter for forecasting. Two methods are used to deal with structural breaks. First, discrete breaks are accommodated using methods introduced by Bauwens *et al.* (2011). Second, we use dynamic dimension selection (DDS), as in Koop and Korobilis (2010), allowing for switches between entirely different models to accommodate these breaks.

Fourth, the way large information sets are collated may affect the forecasting accuracy of models. So instead of estimating large VARs it may be that factor augmented VARs, where

information is combined into a smaller number of common factors which remove noise, provide better forecasts. Factor models have been shown, by Giannone *et al.* (2008) and Kabundi *et al.* (2016), to improve forecast accuracy compared to naive models at short horizons.

The results show that addressing changing dynamics, by introducing time-varying parameters, generates more accurate forecasts of core inflation. More information does not, however, necessarily mean better forecasts as smaller models outperform large models, highlighting the importance of persistence when thinking about forecasts of core inflation.

Chapters 2 and 3 of this thesis proceed with the definition of core inflation currently used by SARB, and by many other central banks, as *headline CPI less food and energy*. As a consequence, the definition of relative price shocks as *food and energy prices* becomes entrenched in how central banks think about second-round effects. In SA, although the Monetary Policy Committee (MPC) looks at a number of measures of core inflation, it generally refers to *targeted inflation less food, non-alcoholic beverages, petrol, and energy* in its deliberations on the direction that policy should take (SARB, April 2016). The SARB also often uses the exclusion-based measure of core inflation to justify its policy stance when communicating about monetary policy. In the March 2015 MPC statement of the SARB, the MPC stated that oil prices would lead to a breach of the 3-6 per cent inflation target range and that the bank would “look through these developments” (Kganyago, March 2015). Similarly, in the May 2015 statement, the MPC highlighted that “[w]hile monetary policy should generally look through supply side shocks, such as large electricity tariff increases and oil price changes, we have to be mindful of the second-round effects of such shocks” (Kganyago, May 2015).

It is not readily apparent, however, why an exclusion-based measure represents the best option for core inflation. As a consequence, a large literature focused on defining a practical measure of core inflation for monetary policy has emerged. This has seen three broad classes of measures developed including exclusion-, model-, and statistical-based methods.

The first, and most common approach, is an exclusion-based measure that is used today to define core inflation as *excluding food and energy*. This measure has its origins in the 1970s when the US economy faced volatile shocks to both food – due to significant foreign demand and drought – and energy prices – from restrictions to oil supply introduced by the Organization for Petroleum Exporting Countries (OPEC). Detmeister (2012) provides three characteristics that define an exclusion-based index: the excluded items are pre-determined, they do not change often, and the relative weights used are the same as in the overall headline price index. Exclusion-based measures are typically supported by the argument that they are thought to be more easily understood by the general public, and that they can be replicated. The disadvantage of these measures is that they exclude entire components of inflation which may include vital information regarding the underlying trend of inflation.

The second broad approach to the measurement of core inflation is the statistical approach (see, for example, Blignaut *et al.*, 2009; Rangasamy, 2009; Ruch and Bester, 2013; and Du Plessis *et al.*, 2015). One, of a range of statistical techniques is used to remove transitory noise from

(or smooth) the inflation series. Statistical methods have been the most widely used approach to generating measures of core inflation using many different techniques, mostly filters. A popular and promising measure is the trimmed means approach by Bryan and Cecchetti (1994). This measure aligns well with inflation as a monetary phenomenon and is likely to better represent underlying inflation by taking into account the excess kurtosis and positive skewness in prices. An important disadvantage of the trimmed means measure, from a theoretical perspective, is its inability to distinguish between “transient and persistent extreme price movements” (Wynne, 2008). In SA, Blignaut *et al.* (2009) calculate a number of trimmed means measures³. To address the inability of trimmed means to identify persistent price changes, Cutler (2001) introduces a persistence-weighted core inflation measure. This measure links underlying inflation to a ‘persistence’ concept as defined by Friedman *et al.* (1963), and embraces Woodford’s view that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier” (Woodford, 2003:17). Components of inflation are weighted based on their persistence, defined here by the autoregressive coefficient. Rangasamy (2009) implemented a persistence-weighted core inflation measure for SA.

The third approach to the measurement of core inflation involves the use of an economic model, based on underlying theory such as in Quah and Vahey (1995) or Cristadoro *et al.* (2005). These approaches add additional information, with economic interactions and feedback loops, to inform the path of core inflation. Core inflation is defined in Quah and Vahey (1995:1130) as “that component of measured inflation that has no medium- to long-run impact on real output”, which corresponds to Friedman’s definition of core inflation. Model-based approaches are appealing since the core inflation measure fits into a framework that ensures consistency in analysing economic interactions. However, they do not escape problems of incorrect model specification, identification and uncertainty when applied in practice.

The variety of alternative core inflation measures suggests that criteria are needed to establish which is ‘best’. Clark (2001) argues that policymakers and analysts have reached consensus on the defining properties of a good measure of core inflation. These include that it must track the components of inflation that persist for several years, help predict future headline inflation over the medium term, be less volatile, and be simple. One important omission from this list is that it must be grounded in the theory used by central banks. The appeal of this theoretical grounding is threefold. First, although many techniques can remove the higher frequency movements in headline inflation, these measures remain atheoretic and can only be judged based on the sample available. Second, aligning core inflation with theory ensures that the right identifying assumptions are used when building a practical measure of core inflation. Third, the normative objective function of the central bank is defined by core inflation in a welfare theoretic framework.

There are methods already developed that are well-grounded in theory, the most successful

³The popularity of this type of core inflation measure has meant that StatsSA now includes a trimmed means which trims five per cent off each tail at the product group level.

being model-based definitions such as those of Quah and Vahey (1995) and persistence-based measures such as those of Rangasamy (2009). A flaw of model-based measures such as those of Quah and Vahey (1995) is that core inflation is defined at a macro level, allowing only a limited set of economic relationships, such as a short-run Phillips curve and money neutrality in the long-run, in a single sector. Prices, however, are empirically strongly heterogeneous, both in the magnitude and frequency of price changes. One possible solution is to define core inflation from the perspective of pricing behaviour at a micro-price level, as is done in Bryan and Meyer (2010), Reiff and Várhegyi (2013) and Millard and O’Grady (2012). Sticky-price inflation, as defined from a micro-price product level, accounts for the heterogeneity that exists, and builds its foundation in a theory of forward-looking prices and optimal monetary policy.

Goodfriend (2007) and Walsh (2009) used the theoretical argument of core inflation to focus on the common exclusion-based inflation measure that most central banks, including SA, use when dealing with how policy will respond to relative price shocks. Walsh (2009:30) stated that “[s]ince food and energy prices display little stickiness, responding quickly to shifts in demand and supply, there is a strong case for excluding them from the inflation rate the central bank attempts to control”. Highlighting only food and energy, however, with no appreciation for all prices that may “display little stickiness” is too narrow, with little theoretical foundation to be an optimal core inflation measure. Woodford (2003:14) states that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier”. Using persistence defined as the frequency of price changes at the product level can therefore more accurately capture the theoretical argument for why core inflation would be a better nominal anchor. This measure also takes account of the heterogeneity that exists at the product level.

It is with the advent of micro-price data work starting with Bils and Klenow (2004) in the United States and Creamer and Rankin (2008) in SA that we are able to determine the frequency of price changes for the entire consumer inflation basket; i.e. determine the extent of price persistence. Inflation is the result of many unobserved adjustments. Only a fraction of prices change in a month. Some of those prices will not have changed for over a year, while others will have changed last month. Some rise and fall faster than others. Some goods are on sale, while others are not. These dynamics matter a lot in themselves, as they describe pricing behaviour. But they also matter for the economic theory forming the foundation of how these prices, and hence inflation, are predicted and forecast.

The third objective of this thesis is to introduce a decomposition of South African goods inflation into its extensive margin – *the fraction of prices changing in a specific month* – and its intensive margin – *the magnitude of price changes*. This is done using a micro-price dataset of 5,200,466 individual price quotes covering the period January 2009 to May 2015, which has never been analysed before. Decompositions of this nature provide economists with the underlying price dynamics needed to both replicate the empirical properties found in consumer prices, as well as, make choices on which models better fit this data. This decomposition also provides the information needed to define the persistence of prices, a key characteristic of core

inflation.

More persistent prices also have more forward-looking information. Firms that change prices less often generally need to take account of the likely path of future inflation when setting these prices if they want to maximise profits. For example, when an insurance company sets medical aid prices annually, it needs to take into account its expectation of future inflation. In contrast, when petrol prices change on a monthly basis, these changes are driven by contemporaneous developments in the exchange rate or the international price of oil. Prices that are sticky therefore contain more forward-looking information and can be exploited to uncover inflation expectations and underlying inflation. In South Africa (SA), consumer prices on average change every five months, with the most frequent price changes occurring every month and the least frequent occurring every 15 months (Creamer *et al.*, 2012).

The final objective of this thesis is to use micro-price data to decompose goods inflation into sticky- and flexible-price inflation. Flexible-price inflation accounts for the majority of volatility in overall goods inflation and is less persistent than sticky-price inflation. Sticky-price inflation represents underlying inflationary pressures and encapsulates forward-looking prices, is less volatile and more persistent. Combining the frequency of price changes with actual price changes at a product level allows us to censor products based on the degree of forward-looking information they contain. The usefulness of sticky-price inflation as a ‘good’ core inflation measure is then tested against a number of other core inflation measures commonly used, including trimmed means inflation, persistence-weighted inflation, and the common exclusion-based measure. The trimmed means and persistence-weighted measures are improved to take account of the heterogeneity that exists in prices at a product level. Due to the short sample period, only in-sample performance is assessed: tracking overall inflation with no clear bias and less volatility.

A number of properties make sticky-price inflation an appealing candidate for core inflation. First, unlike most practical measures of core inflation which lack theoretical underpinnings, sticky-price inflation is defined by the right identifying assumptions. It builds on the theory of forward-looking prices and fits the theoretical definitions of core inflation. Second, sticky-price inflation fits into the modern theoretical consensus used by central banks, the New Keynesian paradigm, when modelling policy choices. These include models such as Clarida *et al.* (1999), Aoki (2001) and Bodenstein *et al.* (2008). Third, and related to the second point, Goodfriend (2007) states that monetary policy reached a pre-crisis consensus that core inflation, rather than headline inflation, was the best nominal anchor for a central bank. Finally, sticky-price inflation uses micro-price data to take into account the heterogeneity that exists at the product level. This is the first core inflation measure to have this property in South Africa.

Underlying, or core, inflation has a fundamental role to play in the framework and implementation of monetary policy. This includes the desire to define core inflation as the appropriate target of monetary policy, to look through relative price shocks, to respond adequately to second-round effects, to be forward-looking, and to communicate the stance of monetary

policy. This thesis has provided evidence that: second-round effects matter in SA and need an appropriate monetary policy response to anchor inflation expectations; that forecasts of core inflation can be improved using time-varying models with limited information; and that a novel approach to the definition of core inflation which leverages micro-price data, the theory of forward-looking prices, and is entrenched in the theory of monetary policy, may be the most appropriate measure of inflation for a central bank to target.

Chapter 2

Second-round effects from food and energy prices: an SBVAR approach

2.1 Introduction

Relative price shocks, most notably from food and energy prices, complicate the implementation of monetary policy as output and prices move in opposite directions. In general, monetary policy can look through these shocks as long as there are no second-round effects through rising inflation expectations and wages, as well as core inflation in the economy. However, these price increases are entrenched in the language of wage-setters in South Africa. At the beginning of 2015 as part of negotiations for municipal workers wages, the Independent Municipal and Allied Trade Union (Imatu) general secretary Johan Koen said “[o]ur members, like the majority of South Africans, are really feeling the pinch of unprecedented increases in the costs of electricity, fuel, food and public transport”. Wage settlements include adjustments to relative price shocks leading to second-round effects.

Second-round effects emanate from the ability of price-setting firms and wage-setting labour to increase prices (through increasing mark-ups or marginal cost) and wages, and therefore, the prices of other goods and services in response to a relative price shock (to for example food or energy prices)¹. These types of shocks have been important in SA with food and energy contributing an average of 2.4 percentage points (or 39 per cent) to the average 6.1 per cent headline consumer price index (CPI) from 2000 to 2014. Second-round effects occur through two channels: cost and expectations. The cost channel refers to the effect that relative price shocks play as an intermediate input in the production of other goods and services; for example, the role of petrol in the transport of goods and services. The expectations channel refers to the impact of relative price shocks on wages. If workers, in the face of a relative price shock, believe the shock to be long-lasting, or have the bargaining power to raise wages in response to

¹In this chapter relative price shocks are defined as *those that occur from movements in food and energy prices relative to all other prices*. In a broader context, relative price shocks can occur from sources other than food and energy, but these are not studied.

relative price shocks, then underlying inflation will rise due to higher wages.

In order to measure second-round effects in SA we estimate a seven-variable Structural Bayesian VAR (SBVAR) with short- and long-run zero restrictions and sign restrictions to identify the cost and expectations channels from 1994 to 2014.

The ability to measure second-round effects adequately relies on proper indicators of relative price shocks, core inflation, wages and inflation expectations.² Of these, the most difficult to define is core inflation given its broad range of theoretical and practical definitions (see, for example, Roger (1998) for a discussion on theoretical definitions of core inflation), and properties required. In South Africa, Blignaut *et al.* (2009), Rangasamy (2009), Ruch and Bester (2013) and Du Plessis *et al.* (2015) all provide alternative measures of core inflation. Ultimately, however, for the purpose of this chapter the core measure used has to adequately capture underlying prices and the impact of second-round effects. Another complicating factor is the practicality of implementing monetary policy. Roger (1998) highlights that from an internal forecasting process as well as for the ability of the policymaker to explain policy choices to the public, it is usually necessary for a “central bank to ‘tie its colours’ to one mast”.³ Given that general consensus around core inflation, being defined as *headline CPI less food and energy*, has emerged as an appropriate measure for policymakers and analysts around the world, and the use of this type of core measure to explain policy choices in South Africa, this is the core measure of choice in this chapter. This choice also simplifies the issue of the right relative price measure to use; in this case food and energy prices. Chapter 5 attempts to resolve the theoretical and practical problems around core inflation by proposing a theoretically grounded measure using micro-price data.

The contribution of this chapter is threefold. First, we provide a concise framework through which to estimate the impact of second-round effects from relative price shocks on the economy. This extends the existing work on relative price shocks including Rangasamy (2011) and Rangasamy and Nel (2014) which only estimate the cost channel of second-round effects. However, the most important impact of relative price shocks is its impact on wages, something that hasn’t been quantified in the South African literature. Second, we use recent advances in Bayesian estimation and structural VAR models to estimate a SBVAR with plausible short- and long-run zero restrictions as well as sign restrictions to identify the shocks. Third, we entrench the discussion of second-round effects into the common framework used by the SARB to discuss these effects. This includes using the common exclusion-based core inflation measure of headline CPI less food and energy prices.

The results of this chapter confirm the impact of wage-setters in South Africa, that changes in the price of food, petrol and energy are accommodated and lead to strong second-round effects. According to the SBVAR model, shocks to relative food and energy prices increase

²In South Africa, measures of inflation expectations are sparse and from a time-series perspective, short. This makes the use of these measures in estimating the impact of relative price shocks difficult, and hence, they are not included in this analysis.

³This is equally true for a relative price measure.

wages by 0.3 per cent a year after the shock. The price of other goods and services (or core inflation) increase with a maximum impact of 0.3 per cent, three quarters after the shock. This is due to both the cost and expectations channel. The presence of second-round effects change how a central bank needs to respond to relative price shocks. Generally, when these do not occur a central bank can look through shocks to food and energy prices as they will be temporary in nature. However, when second-round effects are present, the central bank has to respond appropriately to ensure that inflation expectations remain anchored around the target.

2.2 Theory of second-round effects

Second-round effects emanate from the ability of price-setting firms and wage-setting labour to increase prices (through increasing mark-ups or marginal costs), and wages and therefore the prices of other goods and services in response to a relative price shock (Baumeister *et al.*, 2010).

The impact of second-round effects occurs through two channels, costs and expectations. The cost channel refers to the direct impact of changes to a firm's marginal costs due to an increase in input costs. This assumes that the relative price shock is larger than the menu-cost. A familiar example of this is energy prices. Recent increases in electricity prices following the 2008 energy crisis in SA increased the cost of production by raising the price of intermediate inputs. These increases were then passed on to the consumer.

The expectations channel refers to the response of wage-setting labour to a relative price shock in, for example, food and energy. If labour perceives the relative price shock to be permanent, or has the bargaining power to raise wages, it will raise its inflation expectations and demand higher nominal wages. This would increase the price of other goods and services either through the price adjustment by firms as their marginal cost increases, or through increasing consumption. If wages increase the marginal costs of firms, then this would increase the cost channel of second-round effects.

Second-round effects have been formalised in Aoki (2001), Hlédik and Banka (2003), Bodenstein *et al.* (2008), Blanchard and Gali (2007) and Anand and Prasad (2010). Aoki (2001) builds a two-sector dynamic stochastic general equilibrium (DSGE) model, with a flexible- and sticky-price sector. Flexible-price goods are standardised, traded in an almost competitive market and used as both an input in production as well as being consumed by households. This is analogous to goods such as food and energy. Sticky-price goods are differentiated and traded in a monopolistically competitive environment. This is analogous to all other goods in the economy and can be conveniently thought of as core inflation. His model can be used as a basis from which to understand the transmission of relative price shocks on inflation, and modifies the New Keynesian Phillips curve for sticky-price goods to include impacts from the relative price of flexible-price goods. The mechanism through which this occurs is the substitution effect between flexible-price goods and sticky-price goods. The relative price shock raises demand for

sticky-price goods as consumers substitute away from flexible-price goods. With the increase in demand of sticky-price goods, sellers raise their price. As a consequence the relationship between core inflation (sticky-price goods) and relative price shocks (changes in the prices of flexible-price goods) is always positive, due to a relatively elastic substitution between these two goods.

An important outcome from formalising the role of relative prices in DSGE modelling is that optimal monetary policy requires the central bank to target core inflation instead of headline inflation. Aoki (2001) found that targeting core maximises welfare and is sufficient to stabilise relative prices around their efficient level. He found that this also applies to a small open economy. This outcome is echoed in Bodenstein *et al.* (2008), who used a stylised DSGE model including a separate energy sector in order to study the impact of an adverse energy supply shock on optimal monetary policy in the context of alternative policy rules. Policy rules that focus on headline inflation in the presence of an adverse energy shock imply significantly different responses to those that focus on core inflation with headline inflation introducing significantly greater volatility. The optimality of targeting core inflation in emerging market economies is questioned, however, in Anand and Prasad (2010), who expanded the DSGE model of Aoki (2001) to include financial frictions that limit credit-constrained consumers' access to financial markets. They argue that targeting core inflation (in the sense of CPI less food and energy) in emerging market economies would not be optimal as these economies generally face higher food consumption to total consumption as well as low price and income elasticities of food. Therefore, economic agents are likely to factor in food price changes with wage negotiations affecting inflation expectations and the presence of second-round effects.

Much of the focus of Aoki (2001), Bodenstein *et al.* (2008), and Anand and Prasad (2010) is on optimal monetary policy in the context of relative price movements; an important topic but one which does not provide estimates of the size and importance of second-round effects in an economy. Hlédik and Banka (2003) take the first step in answering these questions by modelling second-round effects of supply shocks on inflation using a small dynamic rational expectations open economy model. The chapter found that second-round effects from import prices and the nominal exchange rate have a material impact on inflation, but their size is dependent on the reaction function of the central bank. Cecchetti and Moessner (2008) and Baumeister *et al.* (2010) go a step further by identifying and quantifying second-round effects. Cecchetti and Moessner (2008) directly analysed the impact of second-round effects, including from both food and energy prices, on headline inflation in a group of advanced and emerging market economies. They found that recent higher commodity prices have generally not led to strong second-round effects on inflation. Baumeister *et al.* (2010) also analysed second-round effects, but look only at the oil market (through analysing the impact of an oil supply shock) on total labour costs per employee, real consumer wages and the producer price-wage ratio. Results show that second-round effects are present in some economies (such as Switzerland and the euro area) and not in others (US) and that these effects are the key determinant of cross-country

differences in the ultimate impact of relative price shocks on inflation.

In South Africa, Rangasamy (2011) and Rangasamy and Nel (2014) provide evidence that second-round effects, or at least the cost channel, are significant. Chisadza *et al.* (2013) in contradiction suggest that second-round effects from oil shocks are not important. Rangasamy and Nel (2014) show that a shock to food prices leads to a significant positive increase in core inflation with a peak effect after five months. Rangasamy (2011) found that food prices play a significant role in inflationary episodes in South Africa, that these prices are driven mainly by domestic factors and that there exists “strong second-round impacts” that require attention from policymakers. Rangasamy (2011) argues that “core measures of inflation that exclude food price movements may not accurately reflect the underlying inflationary pressures in the economy and could compromise the attainment of the goal of price stability”. This is due to the likelihood that core inflation will no longer be an unbiased estimator of headline inflation. Walsh (2011) supported this finding showing that a core inflation measure that excludes food can misspecify inflation prompting higher inflation expectations and slow policy responses. Chisadza *et al.* (2013), on the other hand, found that an oil supply shock increases consumer prices in the short-run with no positive persistence on inflation implying no/little second-round effects. They estimated a sign restriction VAR for the impact of differentiated oil price shocks on the South African macroeconomy.

2.3 The (generalised) Phillips curve and VAR model

Theoretical models that have entrenched the role of relative price shocks, such as those of food and energy prices, into the inflation process include Aoki (2001) and Blanchard and Gali (2007), and form a useful starting point for this analysis. Assume an economy has a continuum of households that consume two types of goods: flexible-price goods and sticky-price goods. The flexible-price goods are standardised, traded in an almost competitive market and are used as both an input in production as well as being consumed by households. These are analogous to goods such as food and energy. Sticky-price goods are differentiated and traded in a monopolistically competitive environment. These are analogous to all other goods in the economy or core inflation. In the context of this chapter this is *headline CPI less food, petrol and energy*.

Assume also that there is a continuum of monopolistically competitive firms who produce differentiated domestic goods (sticky-price goods) using labour and flexible-price goods.

Prices for the sticky-price good are set according to Calvo’s (1983) scheme such that a fraction $1 - \theta$ of firms can reset prices each period with the remaining fraction of firms, θ , leaving prices unchanged. Following Smets and Wouters (2002), price inertia is introduced by assuming that domestic firms partially index prices to the previous period’s domestic inflation rate. This yields the following log-linearised sticky-price Phillips curve:

$$\pi_{S,t} = \frac{\beta}{1 + \delta\beta} E_t \pi_{S,t+1} + \frac{\delta}{1 + \delta\beta} \pi_{S,t-1} + \kappa g_t + \lambda_p \Gamma_x X_{F,t} \quad (2.1)$$

where $\pi_{S,t} = \ln \frac{P_{S,t}}{P_{S,t-1}}$ is sticky-price inflation with $P_{S,t}$ being the price index; g_t is output, $X_{F,t} = \ln \frac{P_{F,t}}{P_{S,t}}$ is the price of the flexible-price good relative to the sticky-price good (in this chapter we define the relative price as *the relative prices of food and energy relative to the prices of all other goods and services*); β is the households discount factor; δ is the degree of indexation to the previous periods inflation; and λ_p and κ are functions of other deep parameters including but not limited to β , θ , the share of labour in production, the share of the flexible price good in production of the sticky-price good, and the elasticity of substitution between domestic goods in consumption.

This can be specified in a more general form as:

$$\pi_{S,t} = C_t + \sum_{i=1}^p \rho_i \pi_{S,t-i} + x_t' \zeta + \mu_t \quad (2.2)$$

where x_t' is a vector of variables including output, the exchange rate, wages, the relative price of food and energy, imported prices, and the interest rate, and μ_t is an identically and independently distributed (i.i.d.) error.

In order to estimate the Phillips curve specified in Equation 2.2 as well as introduce a more general dynamic structure we specify a VAR(p) model as:

$$y_t = c + \sum_{j=1}^p \alpha_j y_{t-j} + \mu_t \quad (2.3)$$

where y_t is an $M \times 1$ vector of endogenous variables for $t = 1, \dots, T$; p is the lag length; μ_t is an $M \times 1$ vector of reduced-form errors assumed to be i.i.d. $N(0, \Sigma)$; and c is an $M \times 1$ vector of intercepts. Note that Equation 2.2 is the single-equation form of the core inflation Phillips curve that is estimated in Equation 2.3.

In matrix form this would be:

$$y = (I_M \otimes X) \alpha + \mu \quad (2.4)$$

where $X = [x_1, x_2, \dots, x_T]'$ is a $T \times K$ matrix with $x_t = (1, y_{t-1}', \dots, y_{t-p}')$ and $K = Mp + 1$. $\alpha = \text{vec}(A)$ is a $KM \times 1$ vector with $A = (cA_1 \dots A_2)'$ and $\mu \sim N(0, \Sigma \otimes I_M)$.

Since we are estimating a seven-variable model with a relatively short sample period it is useful to frame the estimation in the context of a likelihood function where we can introduce priors to shrink the parameter space. Following Koop and Korobilis (2010), the likelihood function can be written in two parts:

$$\alpha | \Sigma, y \sim N(\hat{\alpha}, \Sigma \otimes (X'X)^{-1}) \quad (2.5)$$

which is the distribution of α given Σ , and

$$\Sigma | y \sim W(S^{-1}, T - K - M - 1) \quad (2.6)$$

is the Wishart distribution of Σ^{-1} , with $\hat{A} = (X'X)^{-1}X'Y$ being the OLS estimate of A and $S = (Y - X\hat{A})'(Y - X\hat{A})$.


2.4 Data

In order to determine the cost and expectations effect of relative price shocks, we estimated a seven-variable VAR. Table 2.1 describes the variables used in the VAR model, their transformations and sources.

The VAR includes the price of food, petrol and energy relative to core inflation (relative food and energy); imported inflation; real GDP; the unit labour cost (ULC); CPI less food, petrol and energy (or core inflation); the repurchase rate; and the nominal effective exchange rate.

Food, petrol and energy (electricity) prices play a significant role in the evolution and dynamics of headline inflation in South Africa. These components also help consumers form inflation expectations since these items are bought regularly and are a significant part of their budget. Figure 2.1 plots the contribution of food, petrol and energy to headline inflation since 1990. Food inflation is by far the biggest contributor to inflation, averaging 1.8 percentage points of the 7.6 per cent average headline inflation figure. Petrol follows, with an average contribution of 0.5 percentage points, and electricity contributes only 0.3 percentage points to headline consumer inflation over this period. Combined, these prices explain over 30 per cent of inflation during this period.

Figure 2.1: Contribution of food and energy prices to headline consumer inflation



Conventionally, inflation-targeting central banks use headline, or overall, consumer prices as an operational target but use a core, or underlying, inflation measure to look through relative price shocks. In South Africa the monetary policy committee also looks at a number of core inflation measures; however, it generally refers to *headline CPI less food, petrol and energy* in its deliberations on the direction that policy should take (see for example Kganyago, May 2015,

Table 2.1: Variables included in the VAR

Variable	Transformation	Source
Headline CPI less food, petrol and energy (core CPI)	Percentage change over four quarters	StatsSA
Food, petrol and energy as a ratio to core CPI (relative food and energy)	Percentage change over four quarters	StatsSA
Import prices (trade weighted)	Percentage change over four quarters	SARB & own calc.
Real GDP (seasonally adjusted)	Percentage change over four quarters	SARB
Unit labour cost (ULC) defined as salaries and wages in the non-agricultural sector as a ratio to real GDP	Percentage change over four quarters	SARB & own calc.
Repurchase rate	Per cent	SARB
Nominal effective exchange rate	Percentage change over previous quarter	SARB

and SARB, April 2016). Therefore, this chapter focuses on defining the second-round effects that precipitate from food, petrol and energy (or electricity) price movements.

2.4.1 Stationarity

The seven variables of the VAR model were tested using a union-of-rejections testing strategy, as proposed in Harvey *et al.* (2009). The strategy states that the null hypothesis of a unit root must be rejected if “either $DF - QD_{\mu}$ or $DF - QD_{\tau}$ rejects”, i.e. if the quasi-differenced demeaned or detrended Dickey-Fuller test rejects the null hypothesis. The lag lengths used in the Dickey-Fuller tests are automatically selected based on the Bayesian Information Criterion (BIC), with a maximum lag length of 11.

Table 2.2: Elliott-Rothenberg-Stock DF-GLS test (P-values)

	$DF - QD_{\mu}$ intercept	$DF - QD_{\tau}$ intercept and trend
Relative food and energy	0.1877	0.0014
Imported inflation	0.0191	0.0192
Real GDP	0.0132	0.0001
Unit labour cost (ULC)	0.0110	0.0103
Core inflation	0.6256	0.0128
Repo rate	0.1189	0.0037
Nominal effective exchange rate	0.0000	0.0000

Table 2.2 shows the p – values for the quasi-differenced demeaned, which includes only an intercept term (labelled $DF - QD_{\mu}$ intercept in Table 2.2), and the quasi-differenced detrended, which in turn includes both an intercept and trend term (labelled $DF - QD_{\tau}$ intercept and trend in Table 2.2), Dickey-Fuller tests. Based on the union of rejections rule, all variables are stationary.

2.5 Priors

In the case of a seven-variable VAR with two lags, over 100 coefficients need to be estimated. Due to the likelihood of overfitting it was necessary to use priors to shrink the parameter space and ensure precise estimates. There were a number of prior types available to estimate this model. The main results presented below are based on the Minnesota prior from work by Doan *et al.* (1984) and Litterman (1986). This prior was used since it leads to simple posterior inference. We show in the robustness section that the results presented are generally robust to the choice of prior.

The Minnesota prior assumes that the coefficients of longer lag lengths are likely to have a mean of 0 with the first lag having a mean of unity. We fit a prior to the α matrix such that:

$$\alpha \sim N(\underline{\alpha}_{min}, \underline{V}_{min}) \quad (2.7)$$

with the elements of $\underline{\alpha}_{min}$ set to 0 except for the first own lag, which is set to 0.9, assuming that the data is fairly persistent but not a random walk. Empirical evidence in South Africa shows that inflation is a highly persistent series, with inflation expectations being sufficiently backward-looking.

The prior covariance matrix \underline{V}_{min} is a diagonal matrix such that the diagonals elements $\underline{V}_{i,jj}$ for equation i are:

$$\underline{V}_{i,jj} = \begin{cases} \frac{a_1}{p^2} & \text{for coefficients on own lags} \\ \frac{a_2 \sigma_{ii}}{p^2 \sigma_{jj}} & \text{for coefficients on lags of variables } i \neq j \\ a_3 \sigma_{ii} & \text{for all coefficients on exogenous variables} \end{cases} \quad (2.8)$$

where a_1 , a_2 , and a_3 are hyperparameters set to 0.5, 0.5, and 10^2 respectively; p is the lag length; and $\sigma_{ii} = S_i^2$ are the OLS estimates of the variance from an AR(p) model. A characteristic of this prior is that as the lag length increases, the variance tends to 0. In the robustness section, we show that the results of this chapter are robust to a variety of hyperparameter choices.

Alternate priors including a Diffuse, Normal-Wishart, Independent Normal-Wishart, and two versions of Stochastic Search Variable Selection (SSVS) were used for the robustness section, to assess the impact of prior choices. Monte Carlo integration was used to estimate the posterior distribution of α when using the diffuse, Minnesota and Normal-Wishart priors. The Gibbs sampler was used to estimate models with the Independent Normal-Wishart and SSVS priors. An initial burn-in phase was implemented to ensure convergence. For more details on each prior see Koop and Korobilis (2010).

2.6 Identification

Identifying the structural shocks from the reduced-form shocks estimated in Equation 2.3 requires assumptions about the relationship between these shocks. We can use a short-run impact matrix such that:

$$\mu_t = Z\varepsilon_t, \quad E\varepsilon_t\varepsilon_t' = I, \quad ZZ' = \Sigma \quad (2.9)$$

where Z is a short-run impact matrix and ε_t are the structural shocks. The information used to determine Z can come from short- and long-run zero restrictions as well as sign restrictions. Short-run zero restrictions force the impact of a shock at time t to be 0. However, from $t + 1$ the impact of the shock on the model's variables is driven by the dynamics of the model. Long-run zero restrictions ensure that there is no long-run ($t = \infty$) impact of the shock on a particular

variable. Since the infinite duration is not observable in practice, we confirmed long-run restrictions by looking at impulse-response functions (IRFs) 100-quarters ahead. Combining zero restrictions in both the short- and long-run initially required numerical optimization to solve the problem, due to its highly non-linear nature (see Gali, 1992). Rubio-Ramirez *et al.* (2010) solved this technical problem with an algorithm that finds the correct rotation matrix which satisfies these restrictions. Sign restrictions were introduced to solve under-identified VAR models, by iteratively sampling solutions to the sign restrictions that were consistent with the reduced-form VAR. This generated a distribution of impulse-response functions that can be restricted based on the sign-restrictions imposed. In Rubio-Ramirez *et al.* (2010) and Binning (2013), the imposition of sign restrictions was done using an *QR* decomposition.

An exactly-identified model requires $n(n - 1)/2$ restrictions on the impact matrix where n is the number of endogenous variables (in the case of a seven-variable VAR, this equates to 21 restrictions). However, exactly-identified models using zero restrictions generally require ‘incredible’ identifying assumptions which can be entirely ad hoc. This problem is particularly acute in larger VARs, as was the case here.

We followed an agnostic approach in this chapter, specifying an under-identified system with 14 zero restrictions on the short- and long-run impact matrix, as well as 11 contemporaneous sign restrictions derived from economic theory and following the work of Sims (1980), Blanchard and Quah (1989), Gali (1992), Baumeister *et al.* (2010), Rubio-Ramirez *et al.* (2010) and Binning (2013). We used the algorithm of Binning (2013), which further enhanced the Rubio-Ramirez *et al.* (2010) algorithm to handle short- and long-run zero restrictions, as well as sign restrictions. Since this model is under-identified, there are numerous *Z* matrices that are consistent with the reduced-form model.

The zero restrictions on the short- and long-run impact matrices are specified as:

$$f(Z,A) = \begin{bmatrix} L_0 \\ L_\infty \end{bmatrix} = \begin{array}{c} \text{Relative price} \\ \text{Imported inflation} \\ \text{Real GDP} \\ \text{ULC} \\ \text{Core inflation} \\ \text{Repo rate} \\ \text{Exchange rate} \\ \text{Relative price} \\ \text{Imported inflation} \\ \text{Real GDP} \\ \text{ULC} \\ \text{Core inflation} \\ \text{Repo rate} \\ \text{Exchange rate} \end{array} \begin{array}{c} \text{Relative price} \\ \text{Imported inflation} \\ \text{Aggregate supply} \\ \text{Wage} \\ \text{Aggregate demand} \\ \text{Monetary policy} \\ \text{Exchange rate} \end{array} \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \mathbf{0} & \times \\ \mathbf{0} & \times & \times & \times & \times & \mathbf{0} & \times \\ \mathbf{0} & \times & \times & \times & \times & \mathbf{0} & \times \\ \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times \\ \text{---} & \text{---} & \text{---} & \text{---} & \text{---} & \text{---} & \text{---} \\ \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times \\ \mathbf{0} & \mathbf{0} & \times & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \times & \times & \times & \times & \times & \times & \mathbf{0} \\ \mathbf{0} & \times & \times & \times & \times & \times & \mathbf{0} \\ \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times \end{bmatrix} \quad (2.10)$$

Where L_0 is an $m \times m$ short-run impact matrix, with $L_0 = Z$ and $L_0 L_0' = \Sigma$; and L_∞ is an $m \times m$ long-run impact matrix, $L_\infty = (I - A)^{-1} L_0$ where $A = \sum_{j=1}^p A_j$.

First, based on Blanchard and Quah (1989), we assumed that only real (aggregate supply) shocks have a long-run impact on real GDP growth. This assumption has been used in papers such as Gali (1992) and Christiano *et al.* (2006). This implies that monetary policy is neutral in the long run. Second, we assumed that relative food and energy price shocks have no long-run impact on core inflation. This assumption was based on the idea that only the expectations channel (i.e. the impact of food and energy prices on wages) can lead to a self-sustaining price spiral, with higher wages leading to higher prices. Third, since prices are sticky in the short run, we assumed that a relative food and energy price shock has no contemporaneous impact on core inflation and wages. We did not place a contemporaneous zero restriction on real GDP, as demand effects from commodity price moves are generally instantaneous. Fourth, we assumed that exchange rate shocks have no long-run impact on real GDP, unit labour cost and core inflation. Fifth, we assumed that a monetary policy shock has no contemporaneous impact on real GDP growth, ULC and core inflation. This is a common assumption in short-run restricted VAR models, starting with the work of Sims (1980).

The following contemporaneous sign restrictions were also imposed in order to identify the model:

$$S = \begin{matrix} & \begin{matrix} \textit{Relative price} \\ \textit{Imported inflation} \\ \textit{Aggregate supply} \\ \textit{Wage} \\ \textit{Aggregate demand} \\ \textit{Monetary policy} \\ \textit{Exchange rate} \end{matrix} \\ \begin{matrix} \textit{Relative price} \\ \textit{Imported inflation} \\ \textit{Real GDP} \\ \textit{ULC} \\ \pi_s \\ \textit{Repo rate} \\ \textit{Exchange rate} \end{matrix} & \begin{bmatrix} + & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times \\ - & \times & \times & + & + & \times & \times \\ \times & \times & \times & + & \times & \times & \times \\ \times & + & \times & + & + & \times & \times \\ + & \times & \times & \times & + & + & \times \\ \times & \times & \times & \times & \times & \times & \times \end{bmatrix} \end{matrix} \quad (2.11)$$

Due to contemporaneous zero restrictions, there was no need to place sign restrictions at corresponding points. We imposed the following sign restrictions. First, we assumed that a shock to relative food and energy prices would decrease real GDP due to initially higher imports as well as negative income effects on consumers. The policy rate will increase as monetary policy reacts to anticipated second-round effects. The relative price shock was assumed to be positive. Baumeister *et al.* (2010) provide empirical evidence for these sign restrictions. Second, a wage inflation shock was expected to increase real GDP, as consumers are able to increase spending, and core inflation, as inflation expectations rise. The wage shock was assumed to be positive. Third, an aggregate demand shock was assumed to increase real GDP, core inflation and the Repo rate. Fourth, a monetary policy shock was expected to increase the Repo rate.

2.7 Estimation methodology

The estimation and identification steps were both iterative. In the estimation step, IRFs were calculated from the posterior distribution, taking into account parameter uncertainty. This involved taking draws of Σ^{-1} from Equation 2.6 and finding α from Equation 2.5, conditional on this draw. In the identification step, since the VAR was under-identified, there were many possible models or Z s that satisfied the zero- and sign-restrictions. This introduced some degree of model uncertainty. To minimise this, an adequate number of zero and sign restrictions were imposed. In this chapter, we combined the estimation and identification steps into an iterative process, with each draw providing an estimate of α conditional on Σ^{-1} and a Z matrix that satisfied the zero- and sign-restrictions. The advantage of this approach is that both model and parameter uncertainty are taken into account in the IRFs. For the robustness section, results were generated using the posterior mean estimates of the parameters to highlight the role of model versus parameter uncertainty.

2.7.1 Bayesian estimation step

The Minnesota prior ensured that the posterior inference has an analytical solution involving only the normal distribution such that:

$$\alpha|y \sim N(\bar{\alpha}_{min}, \bar{V}_{min}) \quad (2.12)$$

where

$$\bar{V}_{min} = [V_{min}^{-1} + (\hat{\Sigma}^{-1} \otimes (X'X))]^{-1} \quad (2.13)$$

and

$$\bar{\alpha}_{min} = V_{min}[V_{min}^{-1}\alpha_{min} + (\hat{\Sigma}^{-1} \otimes X)'y] \quad (2.14)$$

The iteration step was augmented with a rule that discarded any draws that produced an unstable VAR, i.e. the eigenvalues of the companion form of the parameter matrix were greater than 1. The discard rate was in the region of 4 per cent, i.e. we discarded 4 draws in 100. See appendix A for more details.

2.7.2 Identification step

Following Binning (2013), the zero restrictions imposed in $f(Z, B)$ can be written as an $m \times 2m$ matrix Q_j such that:

$$Q_j f(Z, A) e_j = 0 \quad (2.15)$$

Where e_j is the j^{th} column of the $m \times m$ identity matrix. The VAR model will be exactly identified if $q_j = rank(Q_j) = m - j$ for $1 \leq j \leq m$. In order to solve for Z in an under-identified model, an orthogonal rotation matrix, P^* , was found that would rotate an initial impact matrix until the zero restrictions imposed in L_0 and L_∞ were solved. The Cholesky decomposition of the covariance matrix ($C = chol(\Sigma)$) post-multiplied by a randomly drawn orthogonal matrix, Q^* , was used for the initial short-run impact matrix ($L_0^* = CQ^*$). The j^{th} column of P^* is equal to the m^{th} column of the Q matrix from a QR decomposition of $\tilde{Q}_j = Q_j F$. The solution is $Z = CQ^*P^*$. For full details of the algorithm, see Binning (2013).

2.8 Results

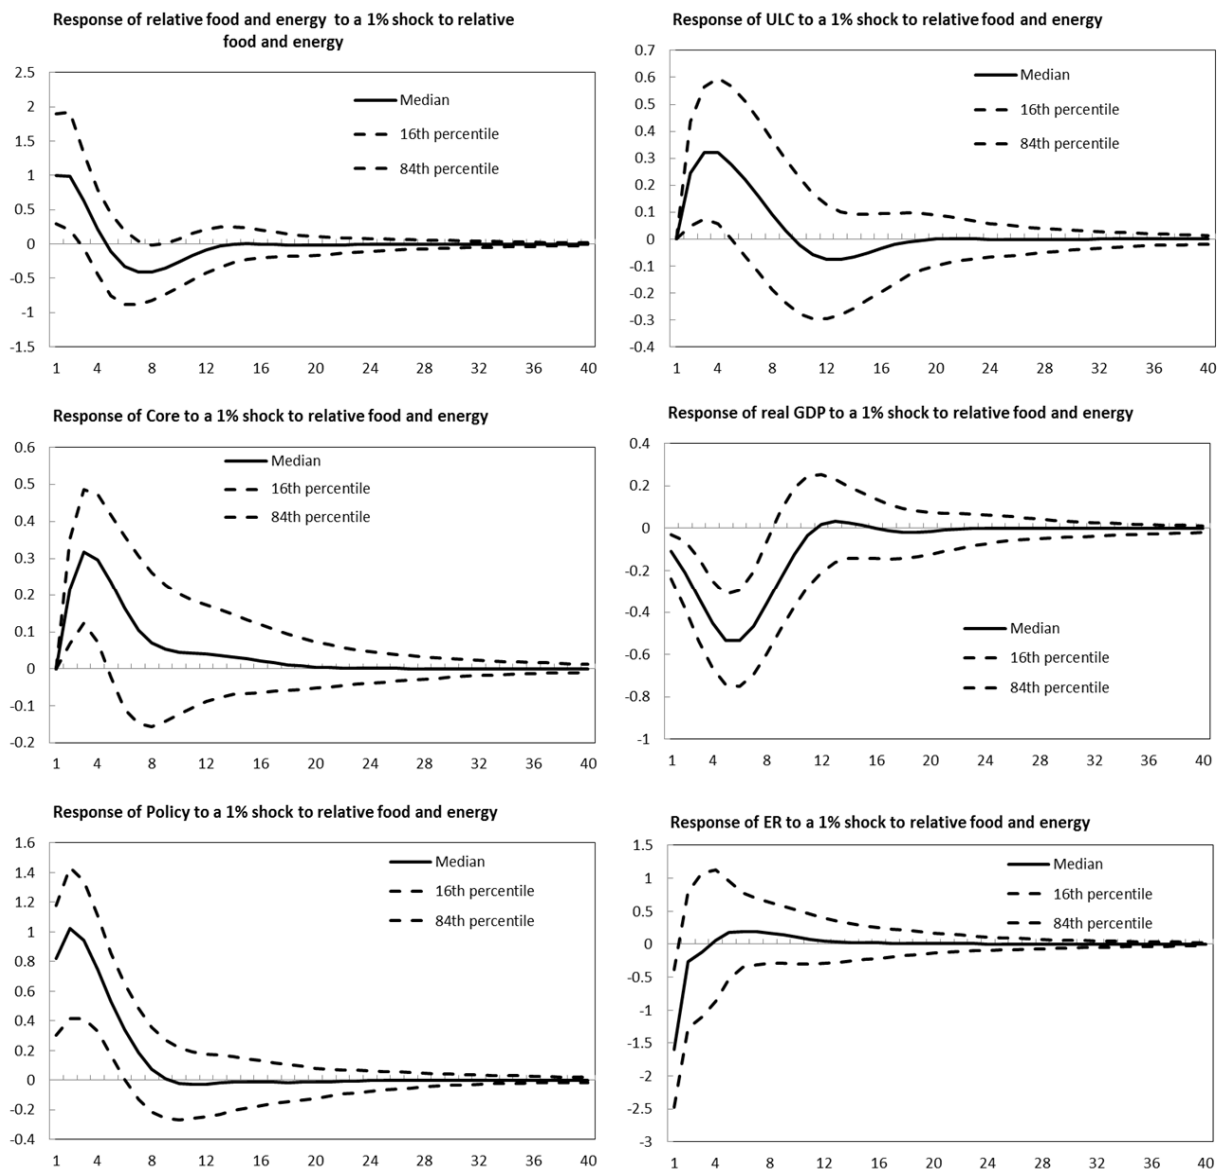
The VAR model was estimated using the Minnesota prior on data from 1994Q1 to 2014Q2. We iterated over 10,000 draws. We included a trend variable to take into account the disinflation that took place over the 1990s to the introduction of inflation targeting in 2000 as well as the concomitant decline seen in interest rates. Although Bayesian estimation mitigates the over-fitting problem, we used two lags in the estimation step due to the relatively short time period.

The Bayesian Information Criterion (BIC) suggests a lag length of one, the Likelihood Ratio test suggests a lag length of four, while the Akaike Information Criterion (AIC) suggests a lag length of 6. For the robustness section, we tested the sensitivity of results to lag length.

2.8.1 Cost and expectations channel

We highlight two specific channels through which relative food, petrol and energy prices feed through to core inflation: the cost and expectations channels. The expectations channel refers to the response of wage-setting labour to a relative price shock. If labour can successfully increase nominal wages, or price-setters their mark-up in response to a relative food and energy price shock, then there will be second-round effects. This is compounded by cost effects as the prices of other goods and services in the economy increase due to the role of food and energy in the production process. If firms pass on the higher production costs to consumers, then the cost effect can be large. Figure 2.2 plots the 40-quarter impulse-response functions of the variables of interest in the VAR to a 1 per cent shock to relative food, petrol and energy prices, including the median, 16th and 84th percentile bands. These bands include both parameter and model uncertainty.

Figure 2.2: The impact of food, petrol and electricity prices



In order to measure second-round effects, we looked at the impact of a relative food and energy price shock on ULC. Figure 2.2 indicates a strong second-round effect in post-Apartheid South Africa. A 1 per cent shock to relative food and energy increases unit labour costs by 0.32 per cent after four quarters. The shock only dissipates after ten quarters. This is only part of understanding second-round effects. It is also necessary to understand how higher labour costs feed into core inflation. Figure 2.3 shows the impact of a 1 per cent shock to ULC: Core inflation rises by 0.48 per cent two quarters after the shock.

Figure 2.3: The impact of ULC on core inflation




Figure 2.2 also shows the impact of a relative price shock on core inflation. A 1 per cent shock to relative food and energy prices has a significant impact on core inflation, with a maximum impact of 0.31 per cent three quarters after the shock. This is due to both the cost and expectations channels. The restrictions imposed ensure that only the expectations channel leads to a permanent increase in core inflation.

A relative food and energy shock also leads to a 1.6 per cent depreciation in the nominal effective exchange rate, which puts further pressure on prices. The endogenous policy response mutes the overall impact of the relative price shock on core inflation.

There is limited literature with which to compare these results. From a cross-country perspective, Cecchetti and Moessner (2008) show that in the majority of emerging market economies, there is no evidence of second-round effects. The study included South Africa; however, the authors do not provide individual country results, making it difficult to determine whether or not South Africa is an economy where second-round effects occur. Baumeister *et al.* (2010) provide some context to the impact of second-round effects, but this is only for oil. They found that second-round effects were the key determinant of cross-country differences in inflation in advanced economies. Countries like Switzerland and the euro area experienced second-round effects and, hence, higher inflation outcomes, requiring an immediate policy response. The US and Japan, on the other hand, experienced no second-round effects.

There is mixed evidence of the existence of second-round effects in the South African literature. Our results are supported by Rangasamy and Nel (2014), who show that second-round effects from food and energy prices matter, with impacts on core inflation, being similar to our estimates. However, the transmission of the peak impact of the relative price shock to core inflation is five months in Rangasamy and Nel (2014), compared with something closer to nine-

twelve months in this analysis. They also found that energy prices lead to cost effects on core inflation but are short-lived. Rangasamy (2011), using a single-equation estimate of the impact of food on non-food prices, found that for every 1 per cent increase in food inflation, non-food inflation increases 0.5 per cent twelve months after the shock. He also found that the impact is greatest 12- to 18-months after the shock. A general result in Fedderke and Liu (2016) shows that the most robust explanatory variable for headline inflation is unit labour cost and specifically the relationship between nominal wages and inflation. This shows that SA faces significant cost-push inflation and supports the finding of this chapter that relative price shocks affect unit labour cost and inflation. Chisadza *et al.* (2013) estimated a sign-restriction VAR of the impact of differentiated oil price shocks on SA and found that an oil supply shock has only a transitory impact on headline inflation, suggesting no cost or expectation effects. This is in contrast to the finding in this chapter.

2.8.2 Demand effects

For an oil-importing country such as SA, theory states that increases in food and energy prices erode disposable income; increase risk, leading to higher precautionary savings; decrease business and consumer confidence; and postpone investment and durable consumption expenditure. The presence of strong second-round effects limits the ability of firms to react to the fall in production through cost containment, forcing firms to adjust profit margins. The (un)willingness of firms to adjust to higher costs determines, in part, the extent of second-round effects. This is compounded by the need for the central bank to react to the second-round effects by raising interest rates. These downward factors are somewhat mitigated by the immediate depreciation of the exchange rate by over 1 per cent in response to the relative price shock.

Figure 2.2 shows that in response to a 1 per cent relative price shock, real GDP growth declines, peaking at 0.53 per cent after five quarters. In comparison, the domestic literature addresses only the impact of oil price shocks and is inconclusive. Chisadza *et al.* (2013) found that an oil supply shock has no significant impact on domestic output. In contrast, Nkomo (2006) shows that oil price shocks have important negative consequences for domestic output. Wakeford (2006), using a case-study approach, found that during episodes of rising oil prices, such as in 1979-80, 1990 and 2003-06, there were significant price pressures and negative output impacts. Finally, Fofana *et al.* (2009) using a macro-meso-micro analysis, found that although there is a negative and significant impact from an oil price shock, the impact is less than expected.

2.8.3 Monetary policy response

The strong second-round effects from food and energy mean that monetary policy needs to react by increasing interest rates. A 1 per cent increase in relative prices leads to an immediate increase of interest rates by 82 basis points. The real interest rate increases by about 60 basis

points. Chisadza *et al.* (2013) found that monetary policy does not react to an oil supply shock, because they found only a transitory, small impact on headline inflation. Part of the reason they did not find a significant policy response may have been the sample chosen. The VAR was estimated on data from 1975, two decades before the interest rate was used as an instrument of monetary policy.

2.9 Robustness

For this section we looked at an alternative estimation methodology, separating out the parameter and model estimation step, as well as the sensitivity of the results to changes to the priors, hyperparameters and lag length. All alternative models were estimated with 2000 iterations and a burn-in of 2000 for priors requiring the Gibbs sampler.

2.9.1 Estimation methodology

The estimation methodology followed introduced both model and parameter uncertainty to the estimate of the impact of a relative food and energy price shock. However, since each parameter draw of the Monte Carlo integration step included a model draw that satisfied the identification restrictions, it was not possible to distinguish between what is model uncertainty and what is parameter uncertainty. Therefore, an alternative strategy was followed in order to distinguish the two impacts. This involved changing the algorithm to solve first the parameter estimation step and then, using the mean parameter estimates, to draw models that satisfied the identification restrictions, i.e. shifting from one-step to two-step estimation. This is analogous to assuming that there was no parameter uncertainty, and the mean estimate was the true value of this parameter iterated over a number of models, solving the identification restrictions. We used 10,000 iterations in each step of the two-step process.

Figure 2.4: Model versus parameter uncertainty

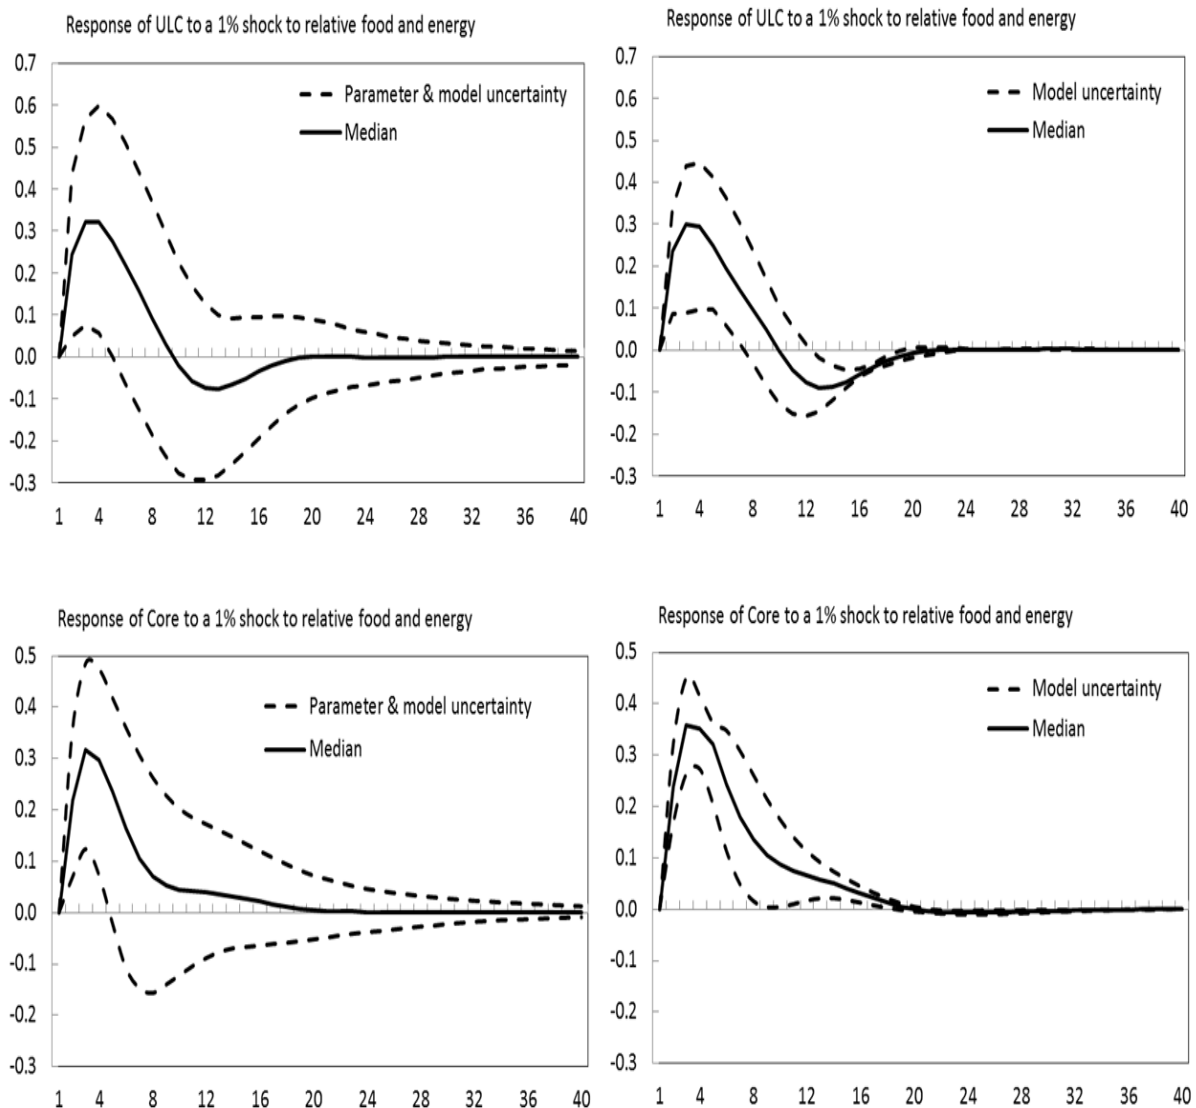


Figure 2.4 plots the median impulse response of core inflation and ULC to a 1 per cent shock to relative food and energy prices, with the left panels showing the initial results (which include both parameter and model uncertainty) and the right panels showing only model uncertainty from the two-step process. The results show that both parameter and model uncertainty contributed significantly to the variance of the coefficient estimates. Model uncertainty effectively doubled the degree of uncertainty around ULC at the peak impact while increasing uncertainty by 1.5 times at the peak impact of core. Model uncertainty had little impact on the median response. It also did not materially change the results of the IRFs, which continued to have the same overall path and magnitude, negating the need for ‘incredible’ restrictions from an exactly identified VAR model.

2.9.2 Sensitivity to priors

Prior selection requires a number of choices regarding the treatment of the covariance matrix as well as the size of the hyperparameters. For example, the drawback of the Minnesota prior used in this chapter is that although it simplified estimation, it treated the covariance matrix as known. To address this shortcoming, we looked at priors that treat the covariance matrix as unknown and introduce uncertainty in its estimation. These priors included the Independent Normal-Wishart (where the coefficients and error covariance are independent of one another) as well as the two SSVS priors. The SSVS priors implemented in this chapter are different from George *et al.* (2008) and follow the work of Koop and Korobilis (2010).

Figure 2.5: Sensitivity to prior selection: ULC

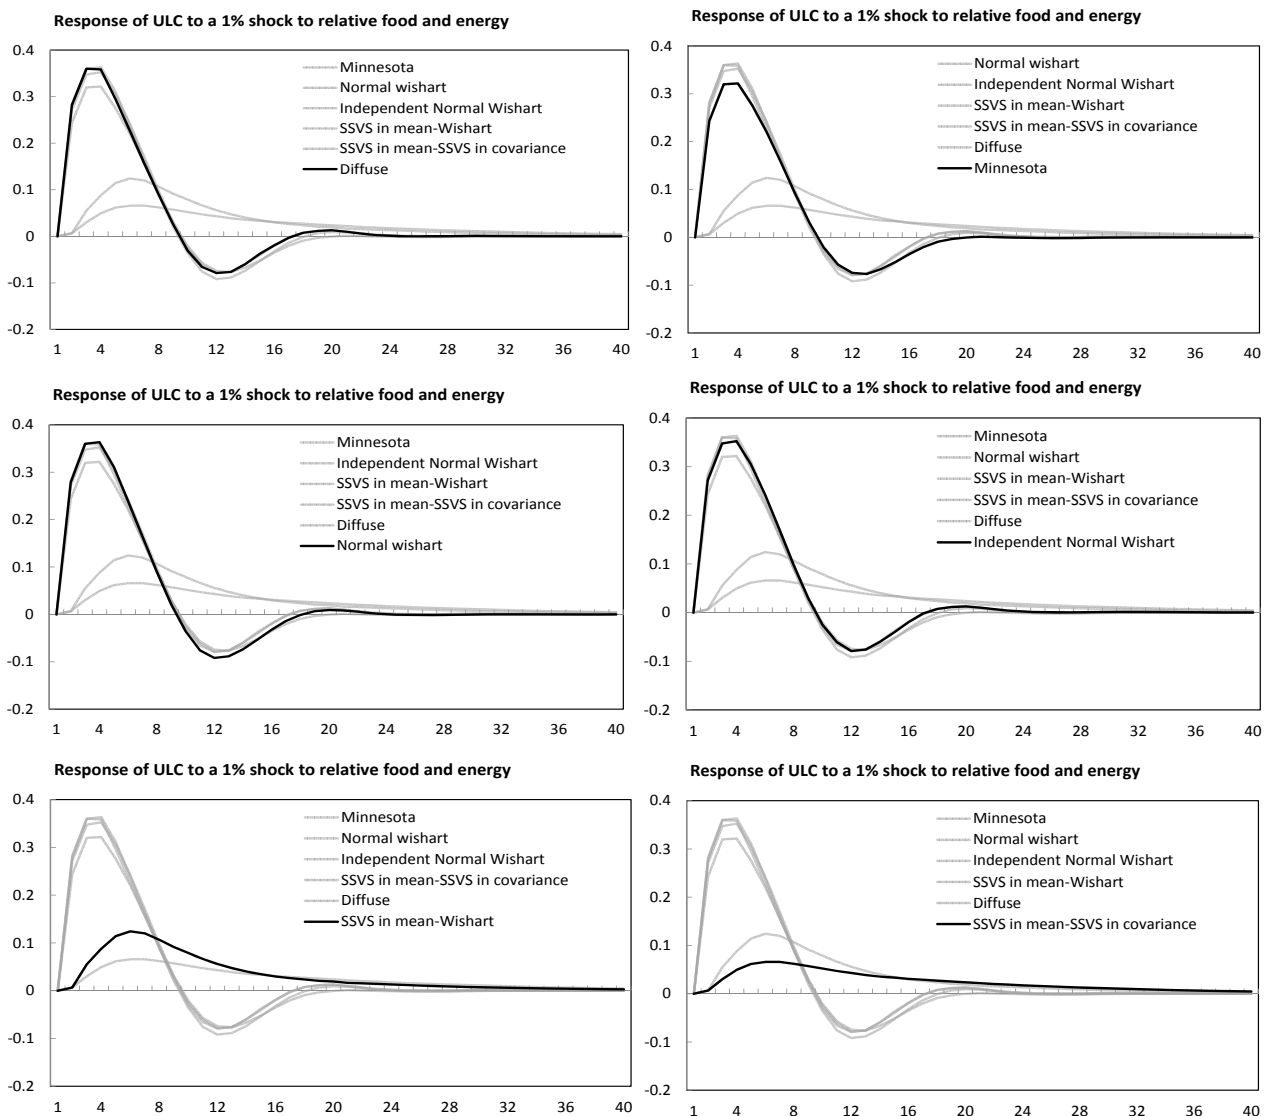


Figure 2.6: Sensitivity to prior selection: core inflation

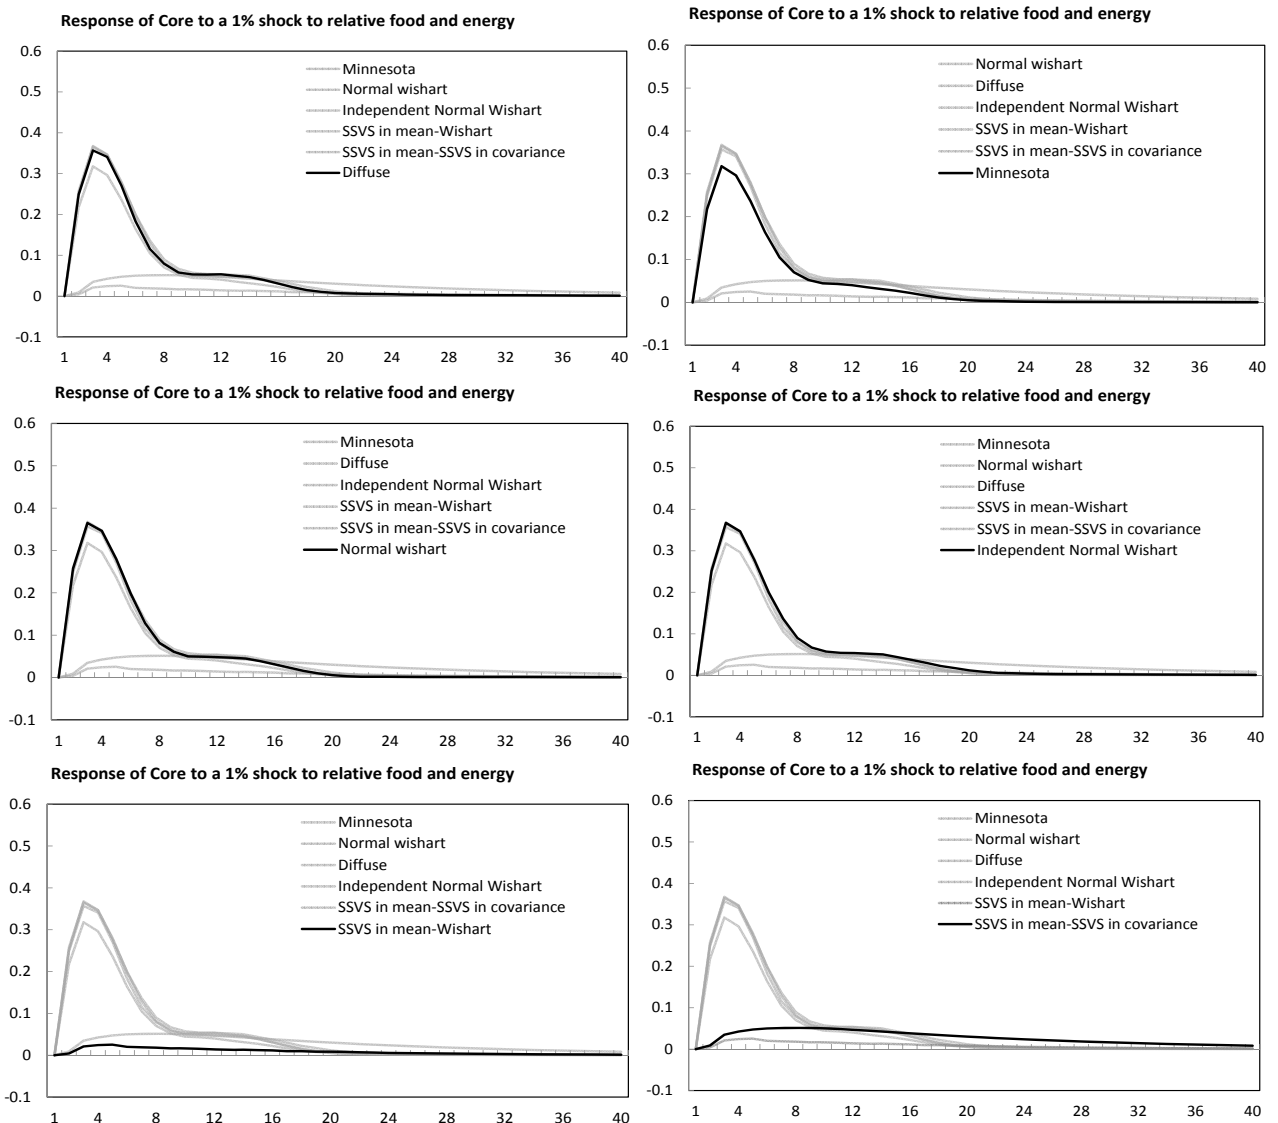


Figure 2.5 and 2.6 shows the impact of alternative prior choices on the response of core inflation and unit labour cost to a 1 per cent shock to relative food and energy prices. All priors except the two SSVS priors provide analytically similar results. In the case of the response of unit labour cost, the impacts are around the 0.3 per cent estimate of the Minnesota prior and follow a similar path. The SSVS priors flatten the impact, suggesting that a 1 per cent shock to food and energy prices leads to a 0.12 per cent impact on ULC, with the maximum impact only six to seven quarters after the shock. In the case of the response of core inflation, a similar pattern emerges, with all the priors except the SSVS priors indicating a just-above 0.3 per cent impact. The SSVS priors suggest no impact from a shock to relative food and energy.

The results of the Independent Normal-Wishart prior show that the IRFs are robust to the treatment of the covariance matrix. In both the core inflation and ULC responses, the Independent Normal-Wishart prior had effectively the same responses to those of the Minnesota prior.

2.9.3 Sensitivity to prior hyperparameters

Hyperparameters control the shrinkage on the parameter estimates by increasing or decreasing the size of the variance. In order to determine the sensitivity of the results to the hyperparameters, we re-estimated the model with combinations of:

$$(\underline{a}_1, \underline{a}_2) = [(0.1, 0.1); (0.2, 0.2); (0.3, 0.3); (0.4, 0.4); (0.5, 0.5); (0.5, 0.3)] \quad (2.16)$$

The last combination, $(0.5, 0.3)$, highlighted the impact of the assumption that own lags are more important than the lags of other predictors (Koop and Korobilis, 2010). We did not test the sensitivity of hyperparameter \underline{a}_3 as, in practice, it is diffuse, i.e. set to 10^2 .

Figure 2.7: Sensitivity to prior hyperparameters: ULC

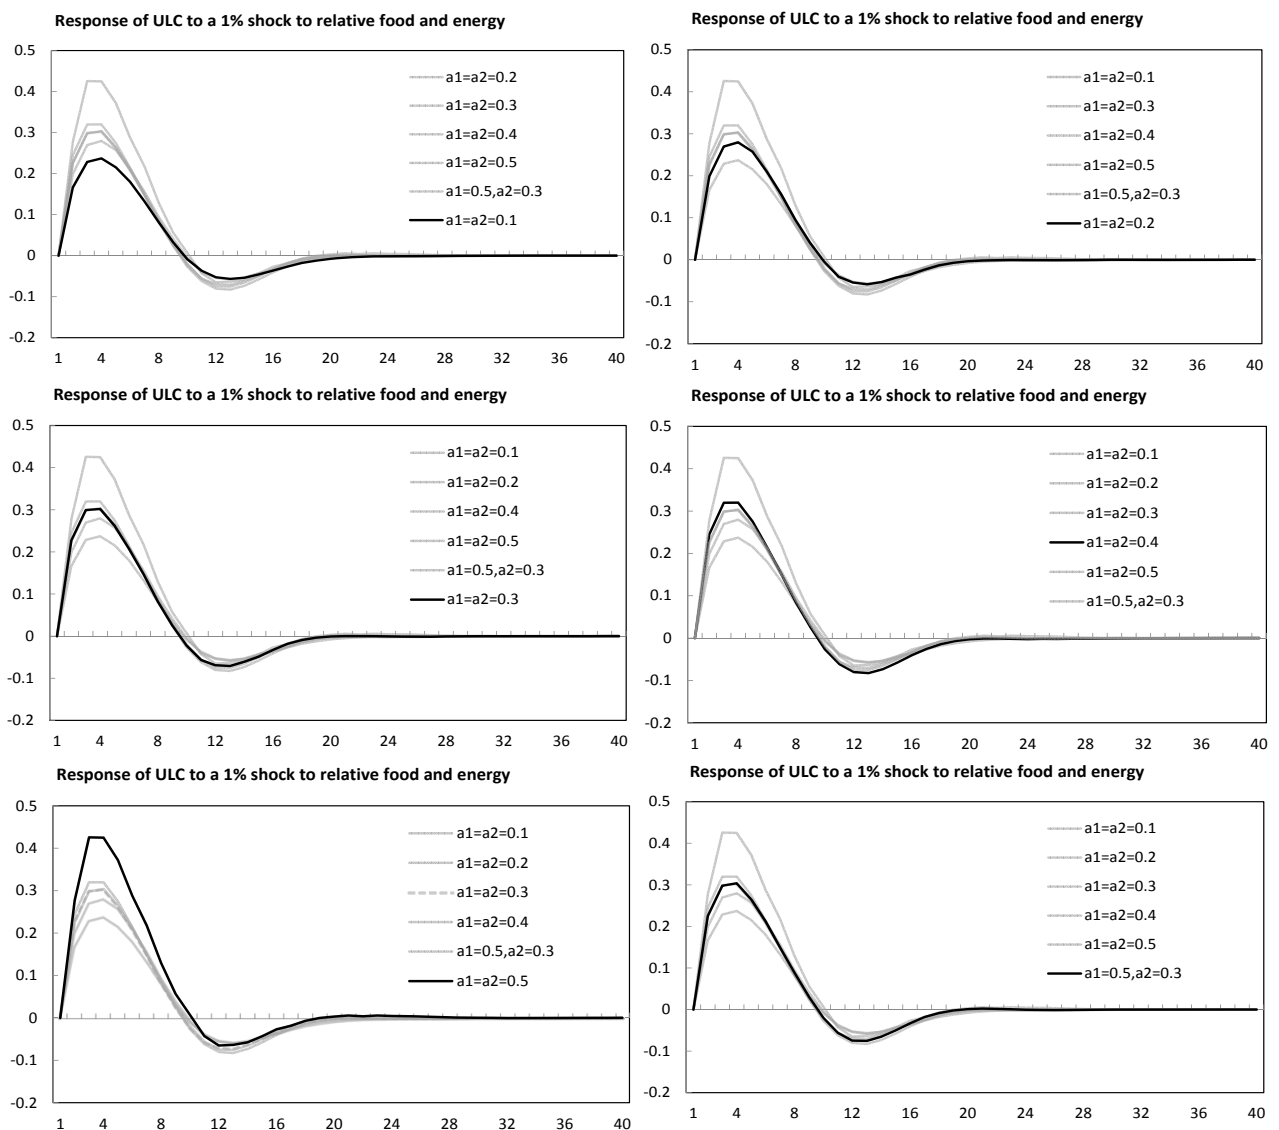


Figure 2.8: Sensitivity to prior hyperparameters: core inflation

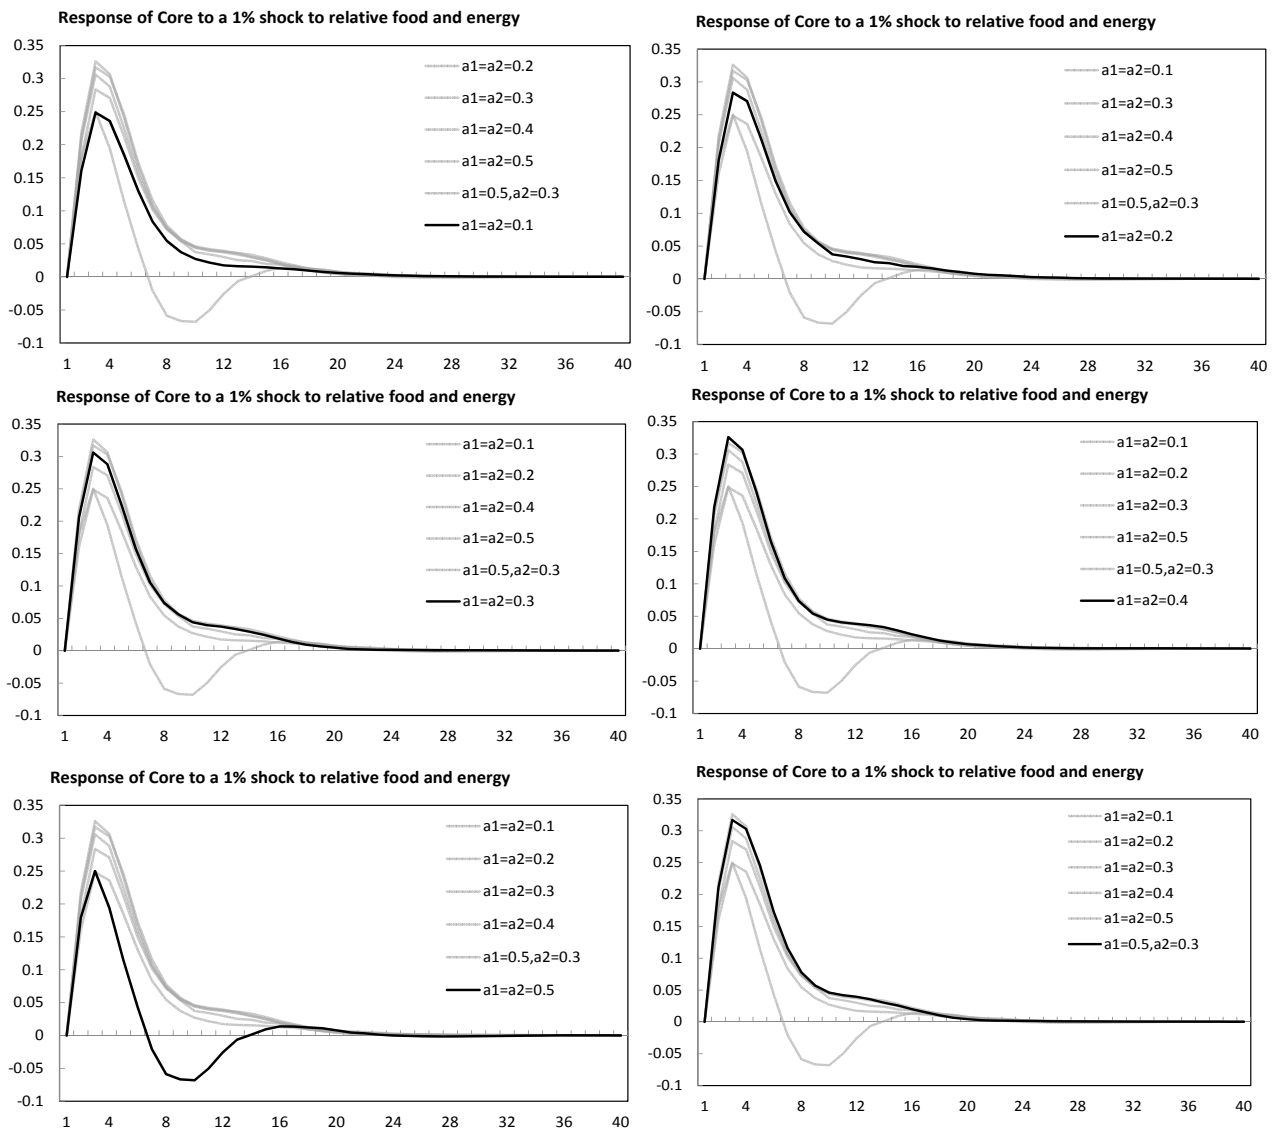


Figure 2.7 and 2.8 plots the IRFs of ULC and core inflation to different combinations of hyperparameters. The results suggest that outcomes are robust to the choice of hyperparameters. The impact on core inflation follows the same trajectory and peaks in the same quarter for all combinations. The amplitude, however, does differ, but marginally in the range of [0.25:0.32]. The outcomes for ULC are similar. The amplitude for ULC is in the range of [0.24:0.42].

2.9.4 Sensitivity to lag length

Figure 2.9 and 2.10 plots the IRFs of core inflation and ULC for lag lengths one to five. In the case of the response of core inflation, alternate lag lengths confirm that there remains an effect from a relative food and energy price shock. The maximum response, however, varies from 0.16 (for a lag length of five) to 0.31 (for a lag length of two). The maximum response also shifts from peaking one quarter after the response (for four lags) to four quarters after (in the

cases of lags three and five). In the case of the response of ULC, there is slightly more variation in the magnitude of the response, but not in its timing. The peak response varies from 0.17 per cent (for one lag) to 0.32 per cent (for two lags). The peak response, however, is four quarters after the shock for all lag lengths.

Figure 2.9: Sensitivity to lag length: ULC

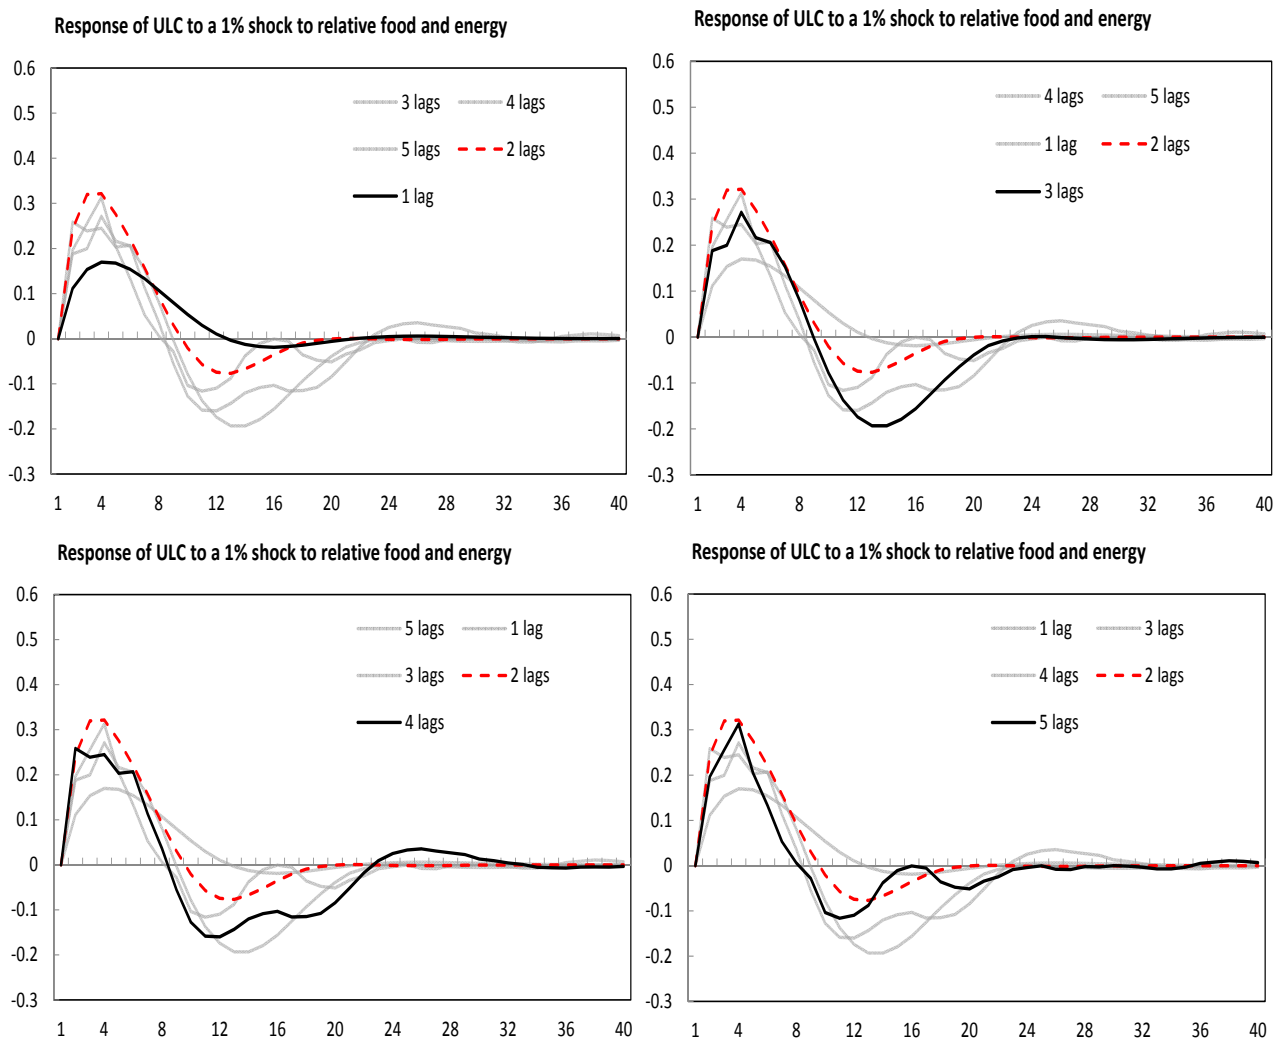
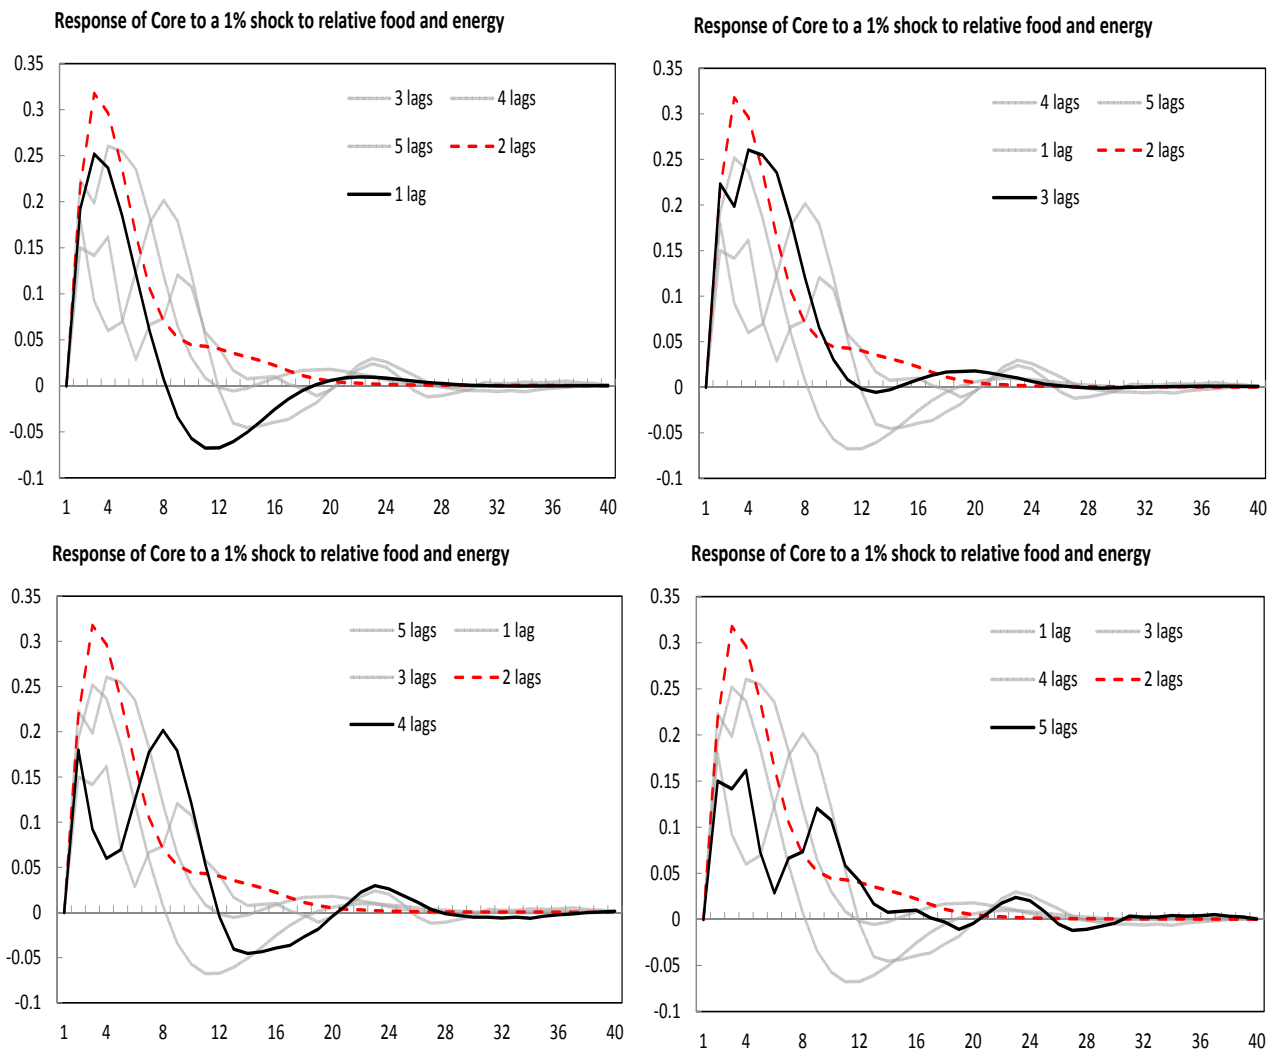


Figure 2.10: Sensitivity to lag length: core inflation



2.10 Policy implications

The results of this chapter have a number of important implications for the conducting of monetary policy which are as follows. First, relative price shocks lead to large cost and expectation effects, requiring an immediate monetary policy response to stabilise core inflation. A lack of this response could lead to persistent inflationary impacts, unanchored inflation expectations, and could result in a wage-price spiral. Second, the peak impact of second-round effects occur three-to-four quarters after the shock, suggesting that impacts take time to feed into underlying inflation, and wage-setters adjust wage demands for the next wage cycle. Third, the presence of strong second-round effects amplifies the decline in output further complicating the response of the central bank as inflation and output move further apart.

The presence of second-round effects are sub-optimal and come at a cost to the economy, or higher inflation (or both). This chapter addressed the existence of second-round effects in the SA economy, but does not answer the question as to why these second-round effects exist.

Second-round effects are the result of product and labour market structure, the process by which agents determine wage and price outcomes, and the choices of the central bank. Future research is needed to determine the causes of second-round effects, and what structural and monetary policy reforms are needed to ensure more optimal outcomes.

2.11 Conclusion

The ability of monetary policy to respond to second-round effects requires that policymakers know about the existence and magnitude of these effects. In South Africa, relative price shocks to food and energy prices are entrenched in the language and responses of wage-setters. However, much of the evidence of this remains anecdotal. In order to measure these effects, we built a SBVAR with plausible zero and sign restrictions based on economic theory, which highlights the role of both the cost and expectations channels of second-round effects.

The results of this chapter confirm the impact of wage-setters in South Africa, that changes in the price of food, petrol and energy are accommodated and lead to strong second-round effects. According to the SBVAR model, shocks to relative food and energy prices increase wages by 0.3 per cent a year after the shock. The price of other goods and services (or core inflation) increase with a maximum impact of 0.3 per cent, three quarters after the shock. This is due to both the cost and expectations channels.

The presence of second-round effects change how a central bank needs to respond to relative price shocks. Generally, when these do not occur, a central bank can look through shocks to food and energy prices, as they will be temporary in nature. However, when second-round effects are present, the central bank has to respond appropriately to ensure that inflation expectations remain anchored around the target.

Chapter 3

Forecasting South African core inflation

Contributing authors: Mehmet Balcilar, Mampho P. Modise, and Rangan Gupta

3.1 Introduction

Like many countries targeting inflation, the South African Reserve Bank (SARB) uses forecasts of headline inflation as its operational target. However, headline inflation can be volatile, making it difficult to distinguish between increases in generalised prices and relative price shocks. This volatility typically arises from a small number of goods and services, most commonly food and energy prices. Petrol prices are a good example of this type of shock, reacting quickly to changes in the international product price and the exchange rate, and having the ability to shift headline inflation by a couple of percentage points in months. However, these movements do not reflect the underlying trends in the behaviour of price-setters or demand conditions in the economy, the issues that matter for a central bank.

In the mid-2000s, prior to the financial crisis, some monetary economists including Goodfriend (2007) and Woodford (2003) argued that monetary policy had reached a consensus, namely, that core inflation rather than headline inflation was the best nominal anchor for monetary policy¹. Core inflation is more stable and would serve as a better anchor for inflation expectations. Woodford (2003:14) stated that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier”. These authors were talking about a conventional definition of core inflation as in “inflation that excludes volatile prices of such goods as food and oil” (Goodfriend, 2007:62). In South Africa, SARB generally refers to core inflation as *headline consumer prices less food, non-alcoholic beverages, petrol, and energy* (SARB, April 2016).

Only two central banks, however, target a measure of core inflation². Despite the use of

¹Monetary policy is subject to a new debate on policy frameworks, including nominal income targeting and price-level targeting, following the global financial crisis (see, for example, Woodford *et al.*, 2014).

²These are Thailand and Norway. There are a number of reasons why headline inflation has become the variable of choice for central banks, including that communication with the public is thought to be easier; wide public understanding; and that people care about a cost-of-living index, the basket of goods they actually consume,

headline inflation as the operational target, central banks depend just as heavily on core inflation in their decision-making processes. In SARB's March 2015 Monetary Policy Committee (MPC) statement, the MPC stated that oil prices would lead to a breach of the 3-6 per cent inflation target range and that the bank would "look through these developments" (Kganyago, March 2015). Similarly, in the May 2015 statement, the MPC stressed that "[w]hile monetary policy should generally look through supply side shocks, such as large electricity tariff increases and oil price changes, we have to be mindful of the second-round effects of such shocks" (Kganyago, May 2015). These statements reflect the value of core inflation in determining a path for monetary policy.

The critical importance of core inflation in the monetary policy process, especially a forward-looking process, as in South Africa, requires accurate forecasts of core inflation. To ensure that the best possible estimates of core inflation are available to the central bank, we looked at a host of possible models that the existing literature shows to have some success in forecasting, and that incorporate a wide variety of new techniques. These include models that take account of large datasets of information, that address possible breaks in the inflation series as monetary policy regimes change, that address the changing relationship between macroeconomic variables and inflation or the structure of the economy, and that provide mechanisms to look at the importance of volatility. We consider core inflation to be defined as *targeted consumer prices less food, non-alcoholic beverages, petrol, and energy*, as this is the measure used by policymakers in communicating issues surrounding underlying prices. We used targeted inflation because the target variable has changed from headline CPI less mortgage interest (CPIX) to headline CPI after the introduction of the CPI basket, based on the Classification of Individual Consumption by Purpose (COICOP) in 2009.

The first contribution of this chapter is that we employed methods for forecasting core inflation in large TVP-VARs, developed by Koop and Korobilis (2013). These models use forgetting factors for computational feasibility. Second, and in addition to the models in Koop and Korobilis (2013), we also assessed the performance power of factor-augmented VARs. As stated in Camba-Mendez and Kapetanios (2005), dynamic factor models tend to perform well in comparison to traditional measures. Third, the chapter adds value by also considering structural break models. Du Plessis *et al.* (2015) states that South African core inflation data has recently been subjected to a structural break given changes in the basket of goods and services and the methodology used in constructing this index. To deal with the structural break dilemma, we combined the Pesaran *et al.* (2006) (PPT) and the Koop and Potter (2007) (KP) methods. The basic idea of the methodology was to use the PPT prior for the break process and the KP prior in conditional mean and variance. We followed Koop and Korobilis (2012, 2013) and Stock and Watson (1999) in the selection of data, which is motivated by a basic New Keynesian model with a generalised Phillips curve. We used quarterly data starting from 1981Q1 to 2013Q4 for 21 variables that include activity variables, labour market variables, financial variables and

rather than core inflation.

other prices.

To the best of our knowledge, this is the first paper to forecast core inflation formally in South Africa. The only other relevant paper is that of Gupta *et al.* (2015), where the authors used the latent state-information recovered from a dynamic stochastic general equilibrium (DSGE) model to forecast core inflation, which, however, was not modelled within the DSGE model explicitly. This means that it was not possible to identify the variables that could help forecast core inflation. Allowing for a large number of predictors, in line with the empirical literature on forecasting inflation based on a New-Keynesian Phillips curve, we were able to determine, which variables contain predictive information for core inflation.

The main results of this chapter are that addressing changing dynamics by introducing time-varying parameters generates forecasts of core inflation that are more accurate. More information does not necessarily mean better forecasts, as small models outperform large models. In general, i) time-varying parameter models consistently outperform constant coefficient models; (ii) small time-varying parameter vector autoregressive models (TVP-VARs) outperform all other models tested; (iii) models where the errors are heteroscedastic do better than models with homoscedastic errors; (iv) models assuming that the forgetting factor remains 0.99 throughout the forecast period outperform models that allow for the forgetting factors to change with time; and (v) allowing for discrete structural breaks does not improve the predictability of core inflation.

The rest of the chapter is structured as follows: section 2 places this chapter in the context of the literature on core inflation; section 3 discusses the methodology; section 4 discusses the data; section 5 follows discusses the results; and section 6 concludes.

3.2 Core inflation

Usually, inflation is defined in two ways. “Inflation is the increase in the supply of money and credit” (Hazlitt, 1965:1), but inflation also is “a substantial *rise in prices* caused by an undue expansion in paper money or bank credit” (American College Dictionary). Jevons (1884) was the first to link the concept of inflation with changes in the prices of products, but prior to that, the term had been used to reflect changes in money supply (Hazlitt, 1965). Since prices are quoted in paper money terms, any change in the supply of money will result in a change in the prices of products as well³. According to modern convention, any discussion of ‘inflation’ is synonymous with the consumer price index⁴. This is the target variable used by SARB and many other central banks. It is used by firms and labour as a reference when setting prices and wages. The CPI has its foundation in welfare economics, based on the maximisation of consumers’ utility.

³Assuming the velocity of money remains unchanged.

⁴For example, Frisch (1977:1289) for simplicity argues that because “neither a satisfactory nor an exact definition of inflation exists as yet” for pragmatic reasons, “inflation is a continuous increase in the price as measured by the consumer price index”.

The consumer price index or cost of living measure, called the iso-utility price index, was first formalised by Konus (1939). Konus (1939:10) stated that “the true index of the cost of living...shows the relative change occurring in the monetary cost of those consumers’ goods that are necessary for the maintenance of a certain standard of living”. This is done practically using the ‘method of aggregates’, which takes the price and quantities of a consumer’s goods in amounts of normal or average consumption over time and calculates the change. Of course this method is a biased measure of cost of living changes as it is unclear which weights (quantities of goods) to use, and because of the assumption that when prices change the consumption basket does not. The size of these biases has been quantified for many economies (see, for example, Lebow and Rudd, 2003).

The CPI-based measure of inflation has not gone unchallenged. For example, Alchian and Klein (1973:173) argue that the consumer price index is “theoretically inappropriate for the purpose they are generally put” and that the actual price index should include asset prices. The debate surrounding the definition and measurement of inflation also played out in the *South African Journal of Economics* (SAJE) in the late 1970s. This started with Lewin (1977), who argued that “popular discussion” had led to the confusion of absolute and relative price changes, inflation importantly being defined as only the absolute price changes since inflation is the “*persistent increases of prices in general*” (Lewin, 1977:175). Lewin (1977:176) also highlighted the problem with measuring inflation using indices stating that “[a] Laspeyres index (such as the GDP - deflator) gives too much weight to relatively high price items whose share tends to decline and a Paasche index too little weight to those items”. These measures would lead to a convolution of absolute and relative price movements.

Botha (1977), in response, argued that Lewin did not focus on the definition of inflation but rather its symptoms, failing to view inflation in the general equilibrium framework. Botha (1977:181) quoted the monetary phenomenon of inflation, emphasising the “deterioration in the real per capita position of the individual due to monetary factors”. Lewin had missed the significance of income and expectations in the process of inflation and had misrepresented the need for inflation to be “persistent”.

The debate continued the following year when Mittermaier (1978) joined it. He argued that Lewin had treated all relative price movements as exactly offsetting and that defining inflation does not readily occur in the distinction of relative and absolute price changes⁵. Mittermaier (1978:45) recognised that part of the problem in understanding inflation is the confounding of measurement and definition. He states that the well-known problems in the measurement of inflation that make these measures inaccurate “create practical difficulties of measurement [but] do not detract *from the concept of inflation* itself”. Mittermaier (1978) used the example of substantial price increases in oil during the 1970s to try to distinguish between what could be considered “inflationary”.

⁵In his reply to Mittermaier, Lewin does not agree with this characterisation but continues to define relative price changes as exactly offsetting only now labelling them “pure”. The debate continued with Lewin, 1978c; Botha, 1978; Lewin, 1978b; Lewin, 1978a; and Mohr, 2008.

The 1970s SAJE debate on the meaning and measurement of inflation highlights more clearly another definitional and measurement problem, that of underlying, or core, inflation. The difference between relative price movements, commonly from food and energy prices, and absolute price movements, as well as the concept of ‘persistence’, goes to the heart of defining and trying to find practical measurements of the process of inflation. The reason is that monetary policy affects the process by which underlying inflation unfolds. Core inflation is actually what a central bank is interested in when setting monetary policy, rather than the common cost of living indices that are widely used today. Core inflation is a “well-defined concept of monetary inflation” (Wynne, 1999).

3.2.1 Definitions of core inflation

Core inflation is defined theoretically but is unobservable. Gordon (1975) was the first to introduce the word “core inflation” to distinguish between underlying inflation, caused by demand-related issues (including monetary policy), and those driven by the supply shocks to food and energy⁶. Gordon specifically referred to underlying prices as “hard-core”. In an era of large supply-side shocks, he was trying to address the question of how government policy, including monetary policy, should respond to maintain macroeconomic stability.

Subsequent theory provides two broad definitions of core inflation expounded in Roger (1998). The first, as a ‘persistence’ concept, builds on earlier work by Friedman *et al.* (1963). Friedman *et al.* (1963:25) highlights two distinct characteristics of inflation, as follows: “... a steady inflation, one that proceeds at a more or less constant rate, and an intermittent inflation, one that proceeds by fits and starts...”, the former being core inflation. The second, as a ‘generalised’ concept defined initially by Eckstein (1981:7) as “... the trend increase of the cost of the factors of production” which “... originates in the long-term expectations of inflation in the minds of households and businesses, in the contractual arrangements which sustain the wage-price momentum, and in the tax system”.

Core inflation differs somewhat from the cost-of-living concept. It is the concept that provides the needed identifying assumptions to reveal the process of inflation (Du Plessis, 2014). The Harmonised Index of Consumer Prices (HICP) of the euro area, for example, is not meant to be a cost-of-living index and is meant to track “the cost of actual monetary transactions” (Astin, 1999). The desire by Eurostat to define HICP as the money part of inflation aligns with the need of a central bank to focus on that part of inflation it can control.

3.2.2 Monetary policy and core inflation

Underlying, or core, inflation is a cornerstone of modern monetary policy. It represents the adequate nominal anchor, in addition to an adequate instrument and credibility, to achieve best

⁶The term “core inflation” does appear earlier in Schreder (1952) but its link to the modern concept is unclear (Wynne, 1999).

the goal of price stability. According to Goodfriend (2007), monetary policy reached the pre-crisis consensus that core inflation, rather than headline inflation, is the best nominal anchor for a central bank. Core inflation is more stable and would serve as a better anchor for inflation expectations. Goodfriend (2007:62) was referring to a conventional definition of core inflation – “inflation that excludes [the] volatile prices of such goods as food and oil”.

Part of the reaching of consensus on core inflation was the development of the theory that showed that core inflation, rather than headline inflation, led to households maximising their welfare. This ‘consensus’ model – with features that include monopolistically competitive firms who set prices in a staggered way, rational expectations, households maximising utility, and a prominent role for monetary policy – was expounded first in Goodfriend and King (1997) and Clarida *et al.* (1999). The rationale behind not targeting headline inflation is that this would require a response to relative price shocks that unnecessarily compounds output losses, i.e. it would force the sticky-price sector to adjust through lower demand and, hence, decrease prices and wages. Relative price shocks from flexible products, such as oil, can also be large, meaning that the output-inflation trade-off would be costly. This result is echoed in Aoki (2001) and Bodenstein *et al.* (2008).

There is literature that looks at whether this optimality always holds. For example, Anand and Prasad (2010) expand a stylised New Keynesian DSGE model to include financial frictions that limit credit-constrained consumers’ access to financial markets in a two-sector model, as in Aoki (2001). They argue that targeting core inflation (in the sense of CPI less food and energy) in emerging market economies would not be optimal as these economies generally face a higher food consumption to total consumption ratio as well as low price and income elasticities of food. When faced with a food price shock, a central bank targeting core inflation will cause proportionally higher food prices, leading to a windfall income gain in this sector (and hence higher wages) and rising aggregate demand. This result is premised on four empirical facts: that the proportion of credit-constrained consumers is large, that the ratio of food consumption to overall consumption is high (in the region of 40%), that the income elasticity of food is price inelastic, and that a significant proportion of wages are earned in the food sector. Although this may be true for some emerging market economies, this result is not universal and only partly true for SA.

The optimality of core inflation is not tied to the New Keynesian paradigm but is a general outcome in welfare economics. Walsh (2009) shows that inflation leads to the highest welfare loss in sectors where prices are more sticky (or more persistent), with few welfare costs when relative price shocks dissipate quickly. Walsh (2009:30) stated that “[s]ince food and energy prices display little stickiness, responding quickly to shifts in demand and supply, there is a strong case for excluding them from the inflation rate the central bank attempts to control”.

Despite the theoretical appeal of core inflation as the optimal nominal anchor, only Thailand and Norway still target it. The reason for this is that a number of practical arguments for headline inflation, and against core, have been made. These arguments include the welfare

foundation of the cost of living index, the supposed communication advantage of headline inflation, its use as a reference rate for wage determination and inflation expectations, the frequency of publication by an independent authority, and the large number of alternative core inflation measures.

The most significant argument endorsing headline inflation comes from its foundation in welfare economics, namely, that a central bank should be concerned with the variable that affects people's lives. This point is rebutted in Du Plessis (2014), who states that the claim that headline CPI is the ultimate goal variable of a central bank does not take account of the outcomes a central bank can control. An important distinction needs to be made. CPI, which is a cost of living index, is founded in welfare economics, where it is defined as some aggregate of households from a representative consumer maximising utility where market price and quantities are observed (Diewert, 1983). This statement may be construed as stating that since CPI represents a welfare-based version of the cost of living of consumers, this should be the focus of central banks. However, this is looking at the price index from the consumer's perspective and not from that of the central banks.

As stated often in many different forms (see, for example, Wynne, 1999; Bernanke, 2001; Cecchetti and Wynne, 2003; and Walsh, 2009), a central bank should be concerned with the money part of inflation, that which it can actually affect. Central banks can do nothing about relative price shocks, nor should they (apart from explaining their inaction) to allow the economy to adjust. Responding to these shocks is also likely to create more volatility (Cecchetti and Wynne, 2003). Similarly, central banks have adopted a theoretic framework (in Clarida *et al.*, 1999) in which to operate, which points to the supremacy of core inflation from a welfare maximising perspective.

A practical argument for headline inflation is its supposed communications advantage (see, for example, Svensson, 1999; Mishkin, 2007; and Roger, 2009). The argument states that headline CPI inflation, which has become a convention when thinking about inflation, is easily understood and accepted by the general public. It is also used in price and wage determination. Du Plessis (2014) argues that it is unlikely that targeting core inflation would undermine SARB's communication strategy, citing recent academic work by Rossouw and Joubert (2005) and Rossouw and Padayachee (2009) suggesting little evidence of the public's understanding of headline inflation or its acceptance of it as a proxy of inflation. Bernanke (2001:322) argued that the exclusion-based measures do not complicate communication to the public, but rather improve it by showing the "public that not every shock that raises prices will lead to a permanent increase in inflation, and that short-term changes in inflation resulting from supply shocks will be treated differently from changes driven by aggregate demand". A more recent strand of literature looking at inflation forecast disagreement shows that significant differences in inflation forecasts are explained by the gap between a conventional definition of core inflation and headline inflation, i.e. the relative prices of food and energy (Siklos, 2016). By targeting a core inflation measure, a substantial degree of disagreement in inflation forecasts can be discarded,

improving monetary policy implementation.

A criticism of targeting headline inflation is that it is subject to large and volatile relative price shocks from food and energy prices as well as from imported inflation. Interest rates are not able to deal with these relative price movements, opening monetary policy up to the ‘blunt tool’ argument. This also risks a central bank’s credibility if it is unable to communicate clearly the reasons for a breach of the inflation target from supply-side shocks. Food and energy prices explain almost 40 per cent of headline inflation in SA since the beginning of inflation targeting, and are generally responsible for breaches of the 3-6 per cent target range. When an economy faces relative price shocks of this magnitude, the desire to anchor inflation expectations to a reference measure such as headline CPI can be dangerous. This can lead to inflation expectations becoming unanchored. Of course, SARB implements its mandate in a fully flexible manner, highlighting core inflation when relative price shocks are significant (Kahn, 2009), in essence, targeting core inflation when relative price shocks are present.

Other arguments used in favour of headline inflation are also unwarranted. First, the idea that headline inflation is a better target since it is produced by an authority independent of the central bank does not recognise that the same applies to core inflation. If the central bank were to use the conventional exclusion-based measure of *excluding food and energy*, this would be based on the same underlying data. Similarly in SA, StatsSA produces a trimmed means inflation measure, and there is no reason why it cannot produce any core inflation measure that the central bank would seriously consider as its target. Second, and related to the first argument, is the frequency of, and rarely revised nature of, headline inflation. This argument holds true for most core inflation measures, and those that it does not, such as the filtered-based core inflation measures, are unlikely to be used as a nominal target.

3.2.3 Measurement of core inflation

The unobservable nature of inflation has seen exclusion-, model- and statistical-based methods all developed in order to estimate a practical measure of core inflation (see, for example, Cogley, 2002; Cristadoro *et al.*, 2005; Quah and Vahey, 1995; and Bryan and Cecchetti, 1993).

The first and most common approach is an exclusion-based measure that is used today to define core inflation as *excluding food and energy*. It has its origins in the 1970s when the US economy faced volatile shocks to both food – due to significant foreign demand and drought – and energy prices – from restrictions to the oil supply introduced by the Organization for Petroleum Exporting Countries (OPEC).

Detmeister (2012) provides three characteristics that define an exclusion-based index: the excluded items are predetermined, they do not change often, and the relative weights used are the same as in the overall headline price index. In SA, Statistics South Africa (StatsSA) publishes nine exclusion-based measures monthly together with the CPI. These generally exclude food, petrol, energy, administered prices, VAT, assessment rates and finance charges, or a combination of the above. In monetary policy deliberations, the SARB generally refers to headline

CPI excluding food, non-alcoholic beverages, petrol and energy when justifying the path of policy (SARB, April 2016).

Exclusion-based measures are typically supported by arguments that they are thought to be more easily understood by the general public and can be replicated. The disadvantage of these measures is that they exclude entire components of inflation, which may include vital information regarding the underlying trend of inflation. Rangasamy (2009) shows that when defining inflation as a persistence-weighted measure, food and energy play a significant role in the underlying trend rate of inflation. Another complicating factor is the level of aggregation at which the exclusion occurs. Lafèche (1997) and Bryan *et al.* (1997) point out that within subgroups, such as food, there may be components that are volatile and others which are persistent. Walsh (2011) states that when food consumption is high, using core inflation defined as excluding food can misspecify inflation, increasing inflation expectations and lead to poor forecasts of headline inflation. Excluding certain volatile components also increases the weight of volatile subcomponents that have not been excluded.

The second broad approach to the measurement of core inflation is the statistical approach. One, of a range of statistical techniques is used to remove transitory noise from (or smooth) the inflation series. Statistical methods have generated the most work on core inflation, using many different techniques, mostly filters. A popular and promising measure is the trimmed means approach by Bryan and Cecchetti (1994). This measure aligns well with the definition of inflation as a monetary phenomenon and is likely to represent underlying inflation better for two reasons. First, Ball and Mankiw (1995:161) show from a theoretical perspective that in a menu cost model:

[w]hen price adjustment is costly, firms adjust to large shocks but not to small shocks, and so large shocks have disproportionate effects on the price level. Therefore, aggregate inflation depends on the distribution of relative-price changes: inflation rises when the distribution is skewed to the right, and falls when the distribution is skewed to the left.

The distribution of price changes in any particular month will be affected by relative price shocks and have excess kurtosis (or fat tails), as is shown in the micro-price literature – see Bills and Klenow (2004) as an example and Creamer *et al.* (2012) for a SA specific result. This motivates the second reason why the trimmed means approach is a promising and popular measure. From a statistical perspective, if a population has excess kurtosis, then trimming the distribution will lead to a more efficient estimate of the population mean. An important disadvantage of the trimmed means measure from a theoretical perspective is its inability to distinguish between “transient and persistent extreme price movements” (Wynne, 2008). In SA, Blignaut *et al.* (2009) calculate a number of trimmed means measures. The popularity of this type of core inflation measure has meant that StatsSA now includes a trimmed means that trims five per cent off each tail at the product group level.

To address the inability of trimmed means to identify persistent price changes, Cutler (2001) introduces a persistence-weighted core inflation measure. The measure links underlying inflation to a ‘persistence’ concept as defined by Friedman *et al.* (1963) and embraces Woodford’s view that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier” (Woodford, 2003:17). Components of inflation are weighted based on their persistence, defined here by the autoregressive coefficient. Rangasamy (2009) implemented a persistence-weighted core inflation measure for SA.

Examples of filter-based approaches to defining core inflation include moving average representations of inflation, the Baxter-King filtered inflation series and Cogley’s exponential smoothing measure (Cogley, 2002). However, most filter-based measures suffer from an end-point problem, complicating their use for policy, especially in a forward-looking framework such as inflation targeting. In SA, these include a core inflation measure defined by singular-spectrum analysis in Ruch and Bester (2013), and wavelet and dynamic factor core inflation measures by Du Plessis *et al.* (2015).

The third approach to measuring core inflation involves using an economic model based on underlying theory, such as in Quah and Vahey (1995) or Cristadoro *et al.* (2005). These approaches add additional information, with economic interactions as well as feedback loops, to inform the path of core inflation. Core inflation is defined in Quah and Vahey (1995:1130) as “that component of measured inflation that has no medium- to long-run impact on real output”, which corresponds with Friedman’s definition of core inflation. Model-based approaches are appealing since the core inflation measure fits into a framework that ensures consistency in analysing economic interactions. However, they do not escape the problems of incorrect model specification, identification and uncertainty.

3.2.4 Properties of a ‘good’ core inflation measure

The extensive number of ways to define and measure core inflation does not necessarily prevent us from determining which core inflation measure is ‘best’, but they do complicate it. For example, if a comparative performance-type approach reveals one measure to be superior to all others, then that would solve the problem, as long as the set of candidate core measures is large enough to represent the likely population of core measures available, and this is true over all sample periods and circumstances. This does not happen in the literature since there is a lack of congruency as well as mutual consistency in the evaluation criteria used to define a good core measure (see Ruch and Bester, 2013 and Rich and Steindel, 2007).

Clark (2001) argues that policymakers and analysts have reached consensus on the defining properties of a ‘good’ measure of core inflation. First, this measure should track the components of inflation that persist for several years, a point also made by Blinder (1997) and Bryan and Cecchetti (1994). In order to judge core measures relative to this criterion, authors such as Detmeister (2012) look at how well a core measure compares with a trend measure, such as a 36-month centred moving average. A second relevant criterion is that a core measure should

provide as much information about this trend as possible, given each month's CPI data. Third, a core inflation measure should help predict future headline inflation over the medium run (i.e. the policy horizon). Fourth, core inflation should track headline inflation with no clear bias and be less volatile than overall headline inflation. Fifth, a core measure of inflation should be as simple as possible; if this measure is used for policy, it should be easy for the public to understand.

Not all of these criteria are mutually consistent. As was shown in Ruch and Bester (2013), there is a clear trade-off between volatility and a low in-sample root mean squared error (RMSE). As a core measure moves more like headline inflation, its RMSE lowers, but its volatility (measured by the sample variance) rises. In the extreme, headline inflation would be the preferred core inflation measure if the RMSE criterion were used.

Another evaluation criterion that creates internal inconsistency in the core inflation literature is the use of some predefined trend measure, such as a 36-month centred moving average, as a benchmark when evaluating which measure of core inflation should be considered 'best'. This approach assumes that the benchmark measures have properties that an econometrician or policymaker defines as 'best'. Using a trend measure to define core inflation guarantees only measures that are similar to the benchmark measures, with no evaluation of the benchmark measures' ability to be a good core inflation measure. Evaluating core inflation based on its ability to track trend headline inflation is present in both the SA literature (Blignaut *et al.*, 2009) and the international literature (see, for example, Clark, 2001 and Detmeister, 2012). Blignaut *et al.* (2009), for example, use various centred moving averages as well as two-sided Hodrick-Prescott filters with differing lambda values as benchmarks to proxy the trend component of inflation. Therefore, the trimmed mean measure that best reflects this trend is the appropriate trim measure for core inflation.

What should define a good core inflation measure? There are three criteria that should be used to determine what is best. First, and foremost, any good core inflation measure should be based in theory. This means that core inflation should be defined by the right identifying assumptions and should be grounded in monetary policy theory. The advantage of these criteria is that a good core inflation measure would no longer rely on sampling properties to test whether it is better than another. Second, core inflation should minimise volatility or at least have significantly lower volatility than headline inflation. The main object of a core inflation measure is to remove the relative price shocks, which by definition are volatile and determine the underlying trend.

Finally, core inflation should accurately forecast future headline over the medium run. Much of the work done up to now to determine the ability of core inflation to forecast headline inflation has focused on point forecasts (see, for example, Bryan and Cecchetti, 1994; Clark, 2001; Detmeister, 2012; and Du Plessis *et al.*, 2015). This criterion, however, should be based on its ability to forecast the centre of the distribution of headline inflation. For example, the criterion could be defined as the best forecast of the median of headline inflation over the period 18

to 30 months ahead. Since the distribution of headline inflation is likely to be skewed with excess kurtosis, due to relative price shocks, the median provides a better representation of the distribution. Similarly, the predictive likelihood, which takes account of the entire predictive density function, could provide a more accurate definition of the ability to forecast headline inflation well. Predictive likelihoods have the added advantage of providing model selection criteria among many models and weights for model averaging exercises (Warne *et al.*, 2013).

3.2.5 Forecasting core inflation

Most of the literature on core inflation has evolved around defining a practical measure of core inflation, rather than on forecasting that measure. The South African literature has similarly been focused exclusively on finding a useful measure of core inflation, rather than forecasting this measure (see Rangasamy, 2009; Blignaut *et al.*, 2009; Ruch and Bester, 2013 and Du Plessis *et al.*, 2015). There is an element to this literature that does provide some guidance on the forecastability of core inflation but only as a means to an end. One of the defining characteristics of a good core inflation measure is that it helps predict future headline inflation (see Clark, 2001 and Blinder, 1997). To this end, Nolazco *et al.* (2016), Bryan and Cecchetti (1993) and Camba-Mendez and Kapetanios (2005) internationally, and Ruch and Bester (2013) and Du Plessis *et al.* (2015) domestically are examples that look at the in- and out-of-sample performance of core inflation measures to help predict headline inflation.

However, core inflation does not only form part of the information set of headline inflation. It is by definition that part of inflation that a central bank should be concerned most about. Core inflation measures attempt to examine the component of headline inflation that is related to broad trends in economic conditions and pricing behaviour, and which are likely to be more persistent (Ranchhod, 2013). Bryan and Cecchetti (1994) further describe core inflation as a process that should be highly persistent, forward-looking and strongly linked to monetary policy dynamics. Recognising core inflation's central role in policy deliberation, Sun (2004), Morana (2007), and Kapetanios (2004) look directly at the ability to forecast core inflation. Sun (2004) proposes an approach to forecast Thailand's core inflation. He combines a short-term model that attempts to filter the forecasting power of a variety of monthly indicators based on goodness-of-fit criteria, with an equilibrium-correction model linking core inflation to its longer-run structural determinants. Morana (2007) uses a principal-components frequency-domain approach, which is suited to estimating systems of fractionally co-integrated processes, to estimate and forecast core inflation for the euro area. Kapetanios (2004) proposes using large datasets using factor models in modelling and forecasting core inflation.

There is only one paper in the South African literature that directly addresses our ability to forecast core inflation. Gupta *et al.* (2015) use the latent state-information recovered from a dynamic stochastic general equilibrium (DSGE) model to forecast core inflation, which was not modelled, however, within the DSGE model explicitly. This means that it was not possible to identify the variables that could help forecast core inflation. The domestic literature tends

to focus on forecasting headline inflation (for detailed literature reviews, see Woglom, 2005; Kanda *et al.*, 2016; and Gupta *et al.*, 2015).

3.3 Methodology

The methods used in this chapter to forecast are motivated by the desire to improve simple models in areas that have been shown to lead to bias and poor forecast performance. Models were extended in four important dimensions in an attempt to be all-encompassing. First, recent methodological and computing gains have made it possible to increase the dimensionality of models, solving the omitted variable bias in smaller VARs, to include up to a hundred variables when analysing and forecasting macroeconomic variables. Bańbura *et al.* (2010), Giannone *et al.* (2014), and Carriero *et al.* (2015) show that increasing the number of variables leads to better forecasting accuracy but that this does have its limitations. Bańbura *et al.* (2010) and Koop (2013) provide evidence that this limit is in the region of 20 variables. We considered a number of model sizes, with up to 21 variables.

Second, a common assumption in simple models of analysis and forecasting is that the errors are homoscedastic. Of course, macroeconomic shocks are not. Engle (1982) first introduced heteroscedastic errors using an autoregressive conditional heteroscedastic (ARCH) process and showed with this seminal piece that inflation in the United Kingdom had significant and changing volatility, especially during the 1970s. Fedderke and Liu (2016) highlight the “surprising” lack of work taking into account ARCH effects in SA inflation, especially given the substantial focus on the effects of the exchange rate on inflation. This chapter, however, looks at another version of heteroscedasticity called stochastic volatility first introduced to VARs by Uhlig (1997). Both homoscedastic and heteroscedastic error structures are used.

Third, significant changes in the structure of the SA economy over the last four decades have made it unlikely that relationships between economic variables remained constant or that there were not any structural breaks. Structural breaks are a significant cause of poor forecasting performance (Stock and Watson, 1996; Ang and Bekaert, 2002; Clements and Hendry, 1998; and Bauwens *et al.*, 2011). To address changing relationships, time-varying parameters were introduced. Primiceri (2005) importantly shows that monetary policy has changed over time in the US. One recent example of changing relationships in SA is provided by Jooste and Jhaveri (2014), who show that the exchange rate pass-through to inflation is time-varying and has declined recently. These changing relationships also matter for forecasting. To deal with structural breaks, two methods were used. First, discrete breaks were taken into account using methods introduced by Bauwens *et al.* (2011). Second, we used dynamic dimension selection (DDS) as in Koop and Korobilis (2010), allowing for switches between entirely different models to accommodate these breaks.

Fourth, the way large information sets are collated may affect the forecasting accuracy of models. So, instead of estimating large VARs it may be that factor augmented VARs – where

information is combined into a smaller number of common factors that remove noise – provide better forecasts. Factor models have been shown to improve forecasting accuracy compared with naive models over short horizons by Giannone *et al.* (2008) and Kabundi *et al.* (2016).

This section introduces the methodologies that were followed.

3.3.1 Large TVP-VARs

We followed the specification in Koop and Korobilis (2013) and specified the time-varying parameter vector-autoregressive model (TVP-VAR) as:

$$y_t = Z_t \beta_t + \varepsilon_t \quad (3.1)$$

and

$$\beta_{t+1} = \beta_t + \mu_t \quad (3.2)$$

where ε_t is an independently and identically distributed (i.i.d.) error with $N(0, \Sigma_t)$ and μ_t is i.i.d. $N(0, Q_t)$. ε_t and μ_t are independent of one another for all s and t . y_t for $t = 1, \dots, T$ is an $M \times 1$ vector containing observations on M time-series variables and Z_t is an $M \times k$ matrix, defined so that each TVP-VAR equation contains an intercept and p lags for each of the M variables for $k = (1 + pM)$. Following Koop and Korobilis (2013), Fagin (1964), Jazwinski (2007), and Raftery *et al.* (2005), we used forgetting factors instead of standard Bayesian statistical inference, since the latter tends to work well only with small TVP-VARs. Forgetting factors allow the Kalman filter to be run only k times, providing an accurate approximation of the likelihood function as the state vector becomes independent across models (for further details on formulating the Kalman filter see, among others, Koop and Korobilis, 2013; as well as Frühwirth-Schnatter, 2006). In estimating a TVP-VAR using forgetting factors, let $y^s = (y_1, \dots, y_s)'$ denote observations through time s . The standard Kalman filter states the following:

$$\beta_{t-1|y^{t-1}} \sim N(\beta_{t-1|t-1}, V_{t-1|t-1}) \quad (3.3)$$

The formulae for $\beta_{t-1|t-1}$ and $V_{t-1|t-1}$ are given in Frühwirth-Schnatter (2006). Further,

$$\beta_t|y^{t-1} \sim N(\beta_{t|t-1}, V_{t|t-1}) \quad (3.4)$$

where

$$V_{t|t-1} = V_{t-1|t-1} + Q_t \quad (3.5)$$

To estimate using the forgetting factor, we replaced Equation 3.5 with the following equation:

$$V_{t|t-1} = \frac{1}{\lambda} V_{t-1|t-1} \quad (3.6)$$

λ is the forgetting factor, and varies between $0 < \lambda \leq 1$. Equation 3.6 implies that observations for t periods in the past have weight λ^t in the filter estimate of β_t . This controlled the degree of time-variation of the coefficients. Equations 3.5 and 3.6 also imply that $Q_t = (\lambda^{-1} - 1)V_{t-1|t-1}$; if $\lambda = 1$, then we get constant coefficients. Raftery *et al.* (2010) set $\lambda = 0.99$, while Koop and Korobilis (2012) use $[0.8, 0.95, 0.99]$. In this chapter, we show results for $\lambda = 0.99$, and we followed the approach in Koop and Korobilis (2013) of estimating λ at each point in time⁷.

We also used a decay factor, κ , to simplify the implementation of multivariate stochastic volatility in ε_t . An exponential weighted moving average (EWMA) was used to estimate Σ_t following RiskMetrics (1996):

$$\hat{\Sigma}_t = \kappa \hat{\Sigma}_{t-1} + (1 - \kappa) \hat{\varepsilon}_t \hat{\varepsilon}_t' \quad (3.7)$$

where $\hat{\varepsilon}_t = y_t - \beta_t' Z_t$ is estimated by the Kalman filter. We set the decay factor equal to 0.96.

Although TVP-VARs work relatively well for modelling the gradual evolution of coefficients, they tend to work poorly for abrupt changes of the coefficients. One solution to this problem is allowing for switches between entirely different models to accommodate these breaks. We used methods developed in Raftery *et al.* (2010) and Koop and Korobilis (2012, 2013) for doing dynamic model averaging (DMA), which can also be used for dynamic model selection (DMS). DMA refers to the averaging of a large set of j models, weighted based on their predictive content, to forecast at a specific point in time, i.e. calculating the likelihood function for $j = 1, \dots, J$ and averaging these likelihoods to generate a forecast. This produces a probability $\pi_{t|t-1,j}$ with $j = 1, \dots, J$. $\pi_{t|t-1,j}$ varies over time, and the forecasting model can switch over time. Once the $\pi_{t|t-1,j}$ for $j = 1, \dots, J$ are obtained, they can be used either to achieve model selection or model averaging. DMS refers to when the single best model – which can change overtime, given selection over a large number of predictors – is used to forecast at each point in time, that is, selecting the model with the highest likelihood. The advantage of this approach is that optimal values for λ , κ and the VAR shrinkage parameter can be selected in a time-varying manner.

To construct a dynamic model selection, we followed the basic algorithm in Raftery *et al.* (2010) and Koop and Korobilis (2012, 2013). Given the initial condition $\pi_{0|0,j}$ for $j = 1, \dots, J$, the model prediction equation using the forgetting factor approach was derived as follows:

$$\pi_{t|t-1,j} = \frac{\pi_{t-1|t-1,j}^\alpha}{\sum_{l=1}^J \pi_{t-1|t-1,l}^\alpha} \quad (3.8)$$

with a model updating equation of:

⁷Estimating λ involves using dynamic model selection to choose a value of $\lambda \in \{0.97, 0.98, 0.99, 1\}$ at each point in time. For more details, see Koop and Korobilis (2013:9).

$$\pi_{t|t,j} = \frac{\pi_{t|t-1,j} p_j(y_t|y^{t-1})}{\sum_{l=1}^J \pi_{t|t-1,l} p_l(y_t|y^{t-1})} \quad (3.9)$$

where $p_j(y_t|y^{t-1})$ is the predictive likelihood, measuring the forecast performance. $\pi_{t|t-1,j}$ can be written as follows:

$$\pi_{t|t-1,j} \propto \prod_{i=1}^{t-1} [p_j(y_{t-i}|y^{t-i-1})] \alpha^i \quad (3.10)$$

The above equation can be interpreted as follows: if $\alpha = 0.99$, then the forecast performance five years ago receives 80 per cent as much weight as the forecast performance for the last period, but if $\alpha = 0.95$, then the weight for the forecast performance five years ago will only be 35 per cent. $\alpha = 1$ corresponds to conventional model averaging using maximum likelihood.

The forgetting and decay factors introduced help to deal with the time-varying nature of the model and negate the need for priors on the covariance matrices Q_t and Σ_t . However, equally important is how the parameters β_t are estimated. Since we are estimating large VARs and time-varying VARs, and hence could have run into overfitting problems (see Bańbura *et al.*, 2010 as well as Koop and Korobilis, 2013), we used a tight Minnesota prior for β_0 , specified in Koop and Korobilis (2013). After transforming the data to stationarity, the prior mean was set equal to $E(\beta_0) = 0$. The Minnesota prior covariance matrix for β_0 is a diagonal matrix such that $\text{var}(\beta_0) = \underline{V}$ and \underline{V}_i denotes the diagonal elements. The prior covariance matrix was then defined as:

$$\underline{V}_i = \begin{cases} \frac{\gamma}{r^2} & \text{for coefficients on } r \text{ for } r=1, \dots, p \\ \underline{a} & \text{for the intercept} \end{cases} \quad (3.11)$$

where p is the lag length. γ determines the degree of shrinkage on the VAR coefficients, as they are lagged further into the past. Generally, training samples are used to determine appropriate values of priors, as would be the case with a normal Minnesota prior. Here instead, γ is estimated in a similar way as the forgetting factors using DMS with a wide grid for $\gamma \in [10^5, 0.001, 0.005, 0.01, 0.05, 0.1]$. In practice this means that there were a number of different prior values for γ , with the optimal one being chosen by maximising the predictive likelihood. γ is small, since a large degree of shrinkage is needed to produce reasonable forecast performance in large VARs and TVP-VARs⁸. \underline{a} was set to equal 10^2 .

We also augmented the model space with models of different dimensions. In particular, we did dynamic model selection for small (including only three variables), medium (including seven variables) and large (including 21 variables) TVP-VARs. As discussed in Koop and Korobilis (2013) – and as used by Ding and Karlsson (2014) – working with TVP-VARs of different dimensions, y_t will be of different dimension, and therefore predictive densities $p_j(y_{t-1}|y^{t-i-1})$

⁸Unlike the normal Minnesota prior, which has two hyperparameters for own lags and other lags, we used one shrinkage parameter to simplify computation.

will not be comparable. This can be resolved by using the predictive densities for the small VARs (these are variables that are included in all models). In this analysis it means that the dynamic model selection is determined by the joint predictive likelihood for economic growth, core inflation and the three-month Treasury Bill rate.

3.3.2 FAVAR models

Although work by Koop and Korobilis (2013) and Bańbura *et al.* (2010) provide techniques to shrink the parameter space in order to make large VAR estimation and analysis feasible, it may be that using other methods such as data shrinkage from factor augmented VARs (FAVAR) provide better forecasts of core inflation. We therefore estimated a FAVAR model using a two-step process (as in Bernanke *et al.*, 2005). This method is simpler and easier to implement. First, the factors are estimated using principal components analysis with the number of factors determined according to the information criterion by Alessi *et al.* (2010)⁹. Next, the FAVAR model was estimated. Equations 3.1 and 3.2 then became:

$$\bar{F}_{t+1} = \bar{F}_t \beta_t + \varepsilon_t \quad (3.12)$$

and

$$\beta_{t+1} = \beta_t + \mu_t \quad (3.13)$$

where $\bar{F}_t = [F_t, x_t]$ with F_t the factors and x_t being core inflation. In the estimation step, F_t was replaced with an estimate \hat{F}_t from step one.

3.3.3 Structural break models

Structural breaks are a significant cause of poor forecasting performance (Stock and Watson, 1996; Ang and Bekaert, 2002; Clements and Hendry, 1998; and Bauwens *et al.*, 2011). SA has undergone substantial changes in its economic structure over the past four decades, including financial liberalisation with the end of apartheid, the great moderation, and disinflation from inflation rates closer to 20 per cent to rates around the SARBs upper target of 6 per cent. Similarly, the monetary policy target and the methodology used to calculate inflation have changed. Given the importance of these breaks in both levels and differences data, we also considered a structural break model to forecast core inflation. We considered a combination of the PPT and the KP priors. The model uses the PPT prior for the break process and the KP prior for conditional mean and variance. We used the same framework as in Bauwens *et al.* (2011), and a detailed discussion is presented there. We specified the linear regression model framework for the structural break models as:

$$y_t = Z_t \beta_{s_t} + \sigma_{s_t} \varepsilon_t \quad (3.14)$$

⁹We employed numerous other methods, including Bai and Ng (2002) and Onatski (2010) to ensure that we were getting the correct number of factors.

where y_t is the dependent variable; Z_t contains the lagged dependent variables or lagged exogenous variables available for forecasting y_t ; and ε_t is i.i.d. $N(0, 1)$. β_{s_t} determines the conditional mean coefficients and σ_{s_t} represents volatilities. This regression allows for β_{s_t} and σ_{s_t} to vary over time, with $s_t \in 1, \dots, K$ a random variable indicating which regime applies at time t .

We used the KP prior in conditional mean and variance, which adopts a hierarchical prior motivated by the state space literature on time-varying parameter models (discussed in detail in Bauwens *et al.*, 2011). The random walk evolution of coefficients was specified as the following:

$$\beta_j = \beta_{j-1} + \mu_j \quad (3.15)$$

where μ_j is i.i.d. $N_m(0, B_0)$, which is equivalent to $\beta_j | \beta_{j-1} \sim N_m(\beta_{j-1}, B_0)$. This meant that if a structural break occurred, the conditional mean of β_j would be drawn from a distribution with mean β_{j-1} such that the next regime would be determined by the most recent regime. The parameters β_{j-1} and B_0 were unknown and were estimated from the data.

To model the break process, we considered an approach in Chib (1998) and used in PPT. Assume that the restricted Markov process for s_t is given by:

$$Pr(s_t = i | s_{t-1} = i) = p_i \quad (3.16)$$

and

$$Pr(s_t = i + 1 | s_{t-1} = i) = 1 - p_i \quad (3.17)$$

This equation was interpreted as a hierarchical prior and implied a geometric prior distribution for $d_i = \tau_i - \tau_{i-1}$ that measures the durations of regimes. Therefore, if regime i holds at time $t - 1$, then at time t , the process can either remain in regime i with probability p_i or move to regime $i + 1$ with probability $1 - p_i$ if a break occurs. To select the number of breaks, we relied on the specification in Bauwens *et al.* (2011) and set the maximum breaks allowed to five such that $K = 1, \dots, K^{max}$.

3.4 Data

The data used was motivated by a generalised New Keynesian Phillips curve, as in Koop and Korobilis (2012) and Stock and Watson (1999). Table 3.1 provides details of the 21 variables included in the dataset, the VAR these variables were used in, as well as the transformation imposed. The data is quarterly and ranges from 1981Q1 to 2013Q4. All data was transformed to be stationary (see transformation in Table 3.1). This includes activity variables such as real GDP and capacity utilisation; labour market variables such as unit labour cost, wages and employment; financial variables such as stock returns and money stock; and other prices such as producer price inflation, oil prices and non-energy commodity prices. Note that the start and

end dates of our sample were driven purely by the available data on the various variables used, at the time of writing this chapter.

Core inflation is defined as *targeted inflation less food, non-alcoholic beverages, petrol, and energy*, obtained from data collected by StatsSA. This is the core inflation most commonly used by SARB when communicating issues of monetary policy and also the core inflation measure used in the Bank's main econometric model. Targeted inflation refers to headline CPI less mortgage interest (CPIX) prior to 2009 and as headline CPI thereafter. This takes into account the new CPI basket introduced in 2009, based on COICOP. It is seasonally adjusted. There are many practical definitions of core inflation that exist in the literature. Chapter 5 critically analyses issues of core inflation and proposes a theoretically founded core inflation measure using micro-price data.

Table 3.1: Dataset used in the small, medium and large VARs

Variable	Transformation*	Description	VAR*
RGDP	Log first diff.	Real gross domestic product at market prices (GDP)	S,M,L
CORE	Log first diff.	Headline CPI less interest on mortgages, food, petrol and electricity (SA)	S,M,L
TB3	Levels	Treasury bills: 91 days tender rate	S,M,L
NEER	Log first diff.	Nominal effective exchange rate of the rand: Average for the period – 15 trading partners	M,L
OIL	Log first diff.	Brent crude oil spot price (USD)	M,L
FORPRD	Log first diff.	Foreign wholesale price index (trade-weighted) (own calculation)	M,L
ULC	Log first diff.	Manufacturing: unit labour costs	M,L
PCE	Log first diff.	Real final consumption expenditure by households: total	L
GFCF	Log first diff.	Real gross fixed capital formation (investment)	L
JSE	Log first diff.	Johannesburg Stock Exchange (JSE) All Share Index	L
M3	Log first diff.	Money supply: M3	L
CREDIT	Log first diff.	All monetary institutions: total domestic credit extension	L
LEAD	Log first diff.	Leading indicator of all the main trading partner countries	L
RETAIL	Log first diff.	Retail sales	L
WAGES	Log first diff.	Total salaries and wages in the manufacturing sector	L
EMPL	Log first diff.	Employment in private sector (own calculation)	L
INCOME	Log first diff.	Disposable income of households	L
IP	Log first diff.	Industrial production (own calculation)	L
UTIL	Levels	Manufacturing: utilisation of production capacity – total	L
PPI	Log first diff.	Manufacturing producer price index	L
COM	Log first diff.	World Bank commodity price index: non-energy (USD)	L

*Log first diff = logged and the first difference was used, S = small VAR, M = medium VAR, L = large VAR

3.5 Results

The main results for this chapter are presented in Tables 3.2 and 3.3. These show the iterated forecasts¹⁰ for horizons 1 to 8 quarters ($h = 1, \dots, 8$) with a forecast evaluation period of 2000Q1 to 2013Q4, i.e. the starting point of the out-of-sample period corresponds to the starting quarter of the inflation-targeting era in SA monetary policy. The VAR models are estimated with $p = 1$, based on the Bayesian information criterion (BIC). In appendix B we include models with $p = 4$.

Since using iterated forecast increases the computational burden, we followed Koop and Korobilis (2013) and did the predictive simulation in two ways. First, we assumed that the VAR coefficients remain unchanged between T and $T + h$, i.e. $\beta_{T+h} = \beta_T$. Second, we assumed that these coefficients change out-of-sample and simulated from Equation 3.2, to produce draws of β_{T+h} labelled as $\beta_{T+h} \sim RW$ in the tables. Both methods provide β_{T+h} , which we used to simulate draws of y_{T+h} conditional on β_{T+h} to approximate the predictive density.

In order to determine the relative performance of the models, we compared the results with a number of benchmark models, including the common naive random walk model, a small VAR including core inflation, real GDP and interest rates, and an AR(1) model, both estimated using Ordinary Least Squares (OLS). Benchmark models are widely used and provide the simplest examples of models to forecast. We also added a structural break AR(1) model to the benchmark models to account for possible discrete breaks in core inflation over the period under review. Discrete breaks are an alternative to the more smooth adjustments provided for by the TVP models.

These benchmark models were then used to assess the relative performance of the models that became more complex in five important dimensions. First, models of different sizes were considered that took account of additional variables, which may matter for core inflation. These include an AR model with only core inflation, a small (S) VAR using three variables, a medium (M) VAR using seven variables, and a large (L) VAR using 21 variables. Second, models were considered that addressed stochastic volatility in the errors using κ to control the degree of variance. These included heteroscedastic VARs using the three dimensions, setting $\kappa = 0.96$ and homoscedastic VARs setting $\kappa = 0.6$. Third, models were considered that take into account time-varying parameters, versus the traditional constant coefficient models using the hyperparameter λ . TVP estimates can be done by setting $\lambda = 0.99$ or by allowing this to be estimated over time, $\lambda = \lambda_t$. Constant coefficient models set $\lambda = 1$. Fourth, models were considered that change the way information is collated which may affect the ability to forecast. This included FAVARs. Fifth, it might have been that different models would have provided the most accurate forecasts of core inflation at different times, and hence the use of a full approach which used all three VAR model sizes using DMS, referred to as dynamic dimension selection (DDS).

¹⁰Marcellino *et al.* (2006) show that iterated forecasts do better compared with directed forecasts.

3.5.1 Determining the number of factors in the FAVAR model

We implemented a modified Bai and Ng (2002) information criterion, as developed in Alessi *et al.* (2010), that chooses the number of factors by minimising the variance of the idiosyncratic component of the approximate factor model, subject to a penalisation in order to avoid over-parameterisation. The information criterion is:

$$\hat{r}_{c,M}^T = \underset{0 \leq w \leq r_{max}}{\operatorname{argmin}} IC_{\alpha,M}^{T*}(w) \quad (3.18)$$

where

$$IC_{\alpha,M}^{T*}(w) = \log\left[\frac{1}{MT} \sum_{i=1}^M \sum_{t=1}^T (x_{it} - \hat{\beta}_i^{(w)} \hat{F}_t^{(w)})^2\right] + cw p_a(M, T) \text{ for } a=1,2 \quad (3.19)$$


For w common factors, M is the number of variables, T the number of observations, $x_{it} - \hat{\beta}_i^{(w)} \hat{F}_t^{(w)}$ the idiosyncratic error, c an arbitrary positive real number and $p_a(M, T)$ the penalty function. The penalty function was multiplied by c since Hallin and Liška (2007) show that $p(M, T)$ leads to consistent estimation of w , the number of factors, if and only if $cp(M, T)$ does so as well.

The behaviour of $\hat{r}_{c,M}^T$ can only be determined from analysing subsamples of sizes (m_h, t_h) . For any h , we can compute $\hat{r}_{c,m_h}^{t_h}$ which is a monotonic non-increasing function in c . Therefore, there exist moderate values of c such that $\hat{r}_{c,M}^T$ converges from above to w . This has to occur independent of h for the criterion to be stable. This was measured by the variance of $\hat{r}_{c,m_h}^{t_h}$ as a function of h :

$$S_c = \frac{1}{H} \sum_{h=1}^H [\hat{r}_{c,m_h}^{t_h} - \frac{1}{H} \sum_{h=1}^H \hat{r}_{c,m_h}^{t_h}]^2 \quad (3.20)$$

We used all data included in the large VAR (excluding core inflation itself) to estimate factors for a FAVAR model. The transformed data was standardised. Figure 3.1 plots the criterion estimate for the number of factors on the y-axis and an arbitrary positive real number c on the x-axis. We ran the results over a number of subsample sizes in order to ensure that they were robust. To determine the number of factors, we had to find the first value of $\hat{r}_{c,M}^T$ where S_c was zero. The results suggested that the number of factors should be three.

Figure 3.1: Estimating the number of factors



Other methods were also used including the original Bai and Ng (2002) information criterion, and a method proposed by Onatski (2010). The Bai and Ng (2002) method did not converge, a common problem with smaller datasets¹¹. According to Onatski (2010), three factors were also chosen.

3.5.2 Forecasting performance


To evaluate the forecast performance, we use mean squared forecast errors (MSFE) and the predictive likelihood. Predictive likelihood comparisons have the advantage of comparing the forecast performance of the entire predictive density function, and not merely point forecasts. The MSFE and the predictive likelihood in Tables 3.2 and 3.3 are presented as relative to the random walk model. This means that the numbers in Table 3.2 are the ratios of a particular model specification divided by the random walk model. For Table 3.3, the results presented are the sum of log predictive likelihood of different models minus the sum of log predictive

¹¹In factor analysis, usually $M \geq 100$. See Forni *et al.* (2009) as an example of convergence problems.

likelihood obtained for the random walk model. The main results of this chapter, including the percentage gain in performance, are based on the MSFE.

DDS forecasts use the TVP-VAR of dimension with the highest probability. We therefore plotted the time-varying probabilities associated with the TVP-VAR of each dimension in Figure 3.2. Between 1981 and 1994, DMS switches between all three models, with periods where each model dominates. In general, the medium VAR tends to have the highest probability throughout this period. From 1994 onwards the small VAR dominates with the large VAR consistently having the lowest probability. This means that DMS uses the small VAR to produce forecasts of core inflation.

Figure 3.2: Estimated dynamic dimension selection (DDS) probabilities of the small, medium and large TVP-VARs

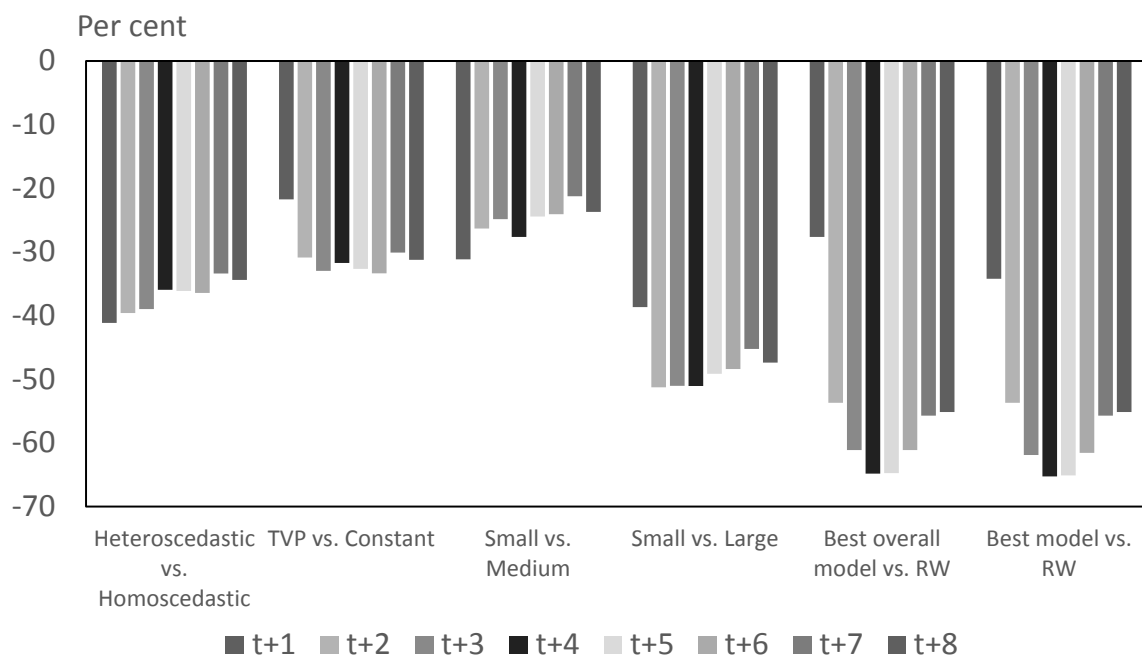


In general, most models (full model, TVP-AR, small and medium TVP-VARs and FAVAR) and different specification (excluding the VAR with homoscedastic errors) performed better than the random walk model. The models performed particularly well over longer horizons of 3-, 4-, 5-, 6-quarters ahead. The average performance of all models excluding the benchmark models improved the forecast of core inflation relative to the random walk model by between

20 and 40 per cent. The only horizon where we did not see any gains was 1-quarter ahead ($h = 1$). The full model and the TVP-AR were preferred for core inflation across all horizons and model specifications, compared to the random walk model. For the small and medium TVP-VARs as well as the FAVAR, only the VAR with homoscedastic errors was outperformed by the random walk model over some horizons. The large TVP-VAR and the benchmark models tended to compete with the random walk model since they performed better over some horizons and worse over others. The large TVP-VAR outperformed the random walk model only when $h = 4, 5, 6$.

Figure 3.3 provides a summary of the relative performance of different types of models. It compares models with heteroscedastic errors to those with homoscedastic errors, time-varying parameter models with its constant parameter counterparts, Different sizes of models, the overall best performing model – the small TVP-VAR model with $\lambda = 0.99$ and $\beta_{T+h} \sim RW$ – to the random walk model, and the best model at each forecast horizon with the random walk model. The results are discussed below. The figure reports the percentage gain in performance; i.e. Small VAR models outperform Large VAR models by about 40 per cent one-quarter ahead.

Figure 3.3: Result Summary: Relative performance gains



Small VAR models outperformed their larger variants by a significant margin. The average outperformance of the small VARs compared with the large VARs was 48 per cent while compared with the medium VAR it was 25 per cent. The best outperformance compared with the large VARs occurs at $h = 2$ at 51 per cent, while the least outperformance occurred at $h = 1$ of 39 per cent. Generally, the expectation of more information should improve forecasts of core inflation. This result was not universal across variables, as a similar exercise on real GDP

growth revealed that large VARs tend to do better than small VARs. This implies that the additional variables included in the medium and large VARs do not have any predictive power for core inflation.

The modelling and forecasting literature on South African inflation tends to suggest that what is most important in modelling inflation is persistence (Gupta and Steinbach, 2013; De Waal *et al.*, 2015), with little role from open economy features in the model. In addition, a recent study has shown that South African inflation persistence is time-varying (Balcilar *et al.*, 2016; Gupta *et al.*, forthcoming). Finally, Balcilar *et al.* (forthcoming) show that the relationship between inflation, output growth and interest rates is also theoretically nonlinearly related, based on nonlinear DSGE models. This suggests that to forecast core inflation requires a model with time-varying persistence and information from real GDP growth and interest rate. This type of model is not unrealistic if one looks at the history of South African inflation, with it being primarily driven by growth and various forms of monetary policy before 1999, and then by interest rate policies post this period, in the inflation targeting era (Gupta and Steinbach, 2013; Balcilar *et al.*, forthcoming). This result is echoed in Stock and Watson (1999), who also found that a single index of overall economic activity is the only variable that improves the forecast of the most successful univariate models. In our case, economic growth seems to function in a similar manner by representing overall economic activity.

Over all models and all horizons, time-varying parameter models outperformed constant coefficient models by an average 31 per cent¹². The minimum improvement was at $h = 1$, at 22 per cent and the best outperformance was at $h = 3, 5, 6$ of 33 per cent. The importance of time-varying parameters highlights the changing economic relationships over the past three decades and particularly since the financial crisis. Jooste and Jhaveri (2014), for example, show that exchange rate pass-through in SA has changed significantly over time with important implications for inflation. This is a particularly important result, since most models used to forecast the main macroeconomic variables by professional forecasters and the central bank are constant coefficient models, including the model used in De Jager (1998). Of course, these forecasts include judgment, but time-varying parameter models provide a better starting point. Moving to time-varying parameter models can improve forecasts by a $\frac{1}{3}$.

VARs with stochastic volatility outperform models with homoscedastic errors over all horizons, improving forecasts of core inflation by 37 per cent, with the best outperformance occurring at $h = 1$, where stochastic volatility VARs outperform by 41 per cent. However, performances are variable. Small and large homoscedastic VAR models outperform the random walk model over most horizons. Medium and FAVAR models, on the other hand, do poorly. The poor performance of the homoscedastic VAR model highlights the importance of allowing for heteroscedastic errors in getting the shape of the predictive density. In general, these results show that the models employed in this chapter provide an effective way of estimating even large

¹²We also included models with only time-varying intercept terms as an alternative TVP strategy. These models did not outperform models where all parameters were allowed to vary but did better than the constant coefficient models.

VARs with heteroscedastic errors, and choosing prior shrinkage.

The importance of time-variation in forecasting quarterly core inflation may suggest that structural breaks are present and may matter for forecasting accuracy. Two methods are used to deal with structural breaks: a discrete break AR(1) model, included as a benchmark model, and DDS. The results for both suggest that accounting for structural breaks does not improve forecasting accuracy. The AR(1) structural break model performs worse than the random walk model over 1- and 2-quarter ahead horizons and better over longer horizons. However, this model relative to all other models does not improve forecasting performance. The DDS results indicate a lot of switching between models in the 1980s and early 1990s, but the small VAR model dominates thereafter. The DDS models do not perform particularly well against the individual VAR models, suggesting structural breaks are not an important driver of forecast performance.

From Table 3.2, there are no significant gains when simulating β_{T+h} from the random walk model compared with just assuming that the VAR coefficients remain unchanged over the forecast horizon. A noticeable comparison can be made between models where $\lambda = 0.99$ and models with $\lambda = \lambda_t$. Models where the forgetting factor is pre-specified outperform models where the forgetting factors are allowed to change over time.

The BIC chose one lag for the VAR models. However, it may be that longer lag orders perform better at forecasting despite the risk of overfitting and higher parameter uncertainty. Table B.1 in Appendix B looks at the MSFE relative to the random walk model for models with lag length of four. In general, more lags improve the forecasting performance of all models by around 14 per cent compared with models with one lag. The overall relative performance of models generally mimics the main results of this chapter with a few exceptions. Models with time-varying parameters outperform models with constant coefficients by an average of 47 per cent over all horizons. Small VARs outperform large VARs. Models with stochastic volatility improve forecasts by an average of 72 per cent compared with models with homoscedastic errors.


In summary, the full model, the TVP-AR, the small and medium TVP-VAR models, as well as the FAVAR model (excluding the VAR with homoscedastic errors) all perform better on average relative to the other models and the random walk model. Specifically, the small TVP-VAR has the smallest MSFE relative to all the models employed.

With regard to the predictive likelihood results presented in Table 3.3, all VAR specifications perform significantly better than the random walk model, confirming somewhat the results presented in Table 3.2. Even in this case, models with $\lambda = 0.99$ performed better than models where $\lambda = \lambda_t$. Also, the VARs with heteroscedastic errors outperformed the VARs with homoscedastic errors. Even with the predictive likelihood, the benchmark models tended to perform poorly relative to the random walk model. Only the AR(1) structural break model performed better than the random walk model for $h = 3$ onwards. When taking the average of the models, the full model, TVP-AR, the small and medium TVP-VAR models, as well as the

FAVAR (excluding the VAR with homoscedastic errors) performed better on average relative to the other models – as is shown in Table 3.2. Similar to the results for the MSFE, the small TVP-VAR had the largest predictive likelihoods relative to all other models for all specifications.

The results presented above are average performances over the inflation-targeting period in SA. However, given the importance of volatility, it may be that forecasts do poorly during certain periods, especially in the face of large macroeconomic shocks. Figure 3.4 assesses the performance of forecasts over time: the cumulative sum of the log-likelihood function for the best relative performer, the small TVP-VAR with $\lambda = 0.99$ and $\beta_{T+h} = \beta_T$ at $h = 1$. There are two important periods where the likelihood function deteriorates significantly, namely 2001-02 and 2008-09. These correspond to a significant exchange rate shock that accelerated core inflation in 2001-02 and the global financial crisis in 2008-09.

Figure 3.4: Cumulative sum of total predictive log-likelihood $h = 1$



3.5.3 Forecast performance of annual core inflation

This chapter focuses entirely on quarter-on-quarter changes in core inflation. Policymakers, however, are more interested in persistent changes in core inflation, making the percentage change over four quarters an important policy variable. Annual changes are the sum of the previous four quarterly changes, but it is not clear that models that forecast quarterly inflation accurately will also forecast annual changes in core inflation, as these have different statistical properties. Therefore, the models were rerun to look at how well annual changes in core inflation are estimated. In this section all variables in Table 3.1, which under column head-

ing Transformation are indicated as log first different are converted to log annual changes, i.e. $\log(x_t/x_{t-4})$ ¹³.


When considering annual core inflation, most models with more complex specifications no longer did better than the benchmark models, and in most cases performed comparatively worse than the random walk model. Time-varying parameter models did slightly better at forecasting core inflation over horizons 1-, 2-, and 3-quarters ahead but on average performed equally well as their constant coefficient counterparts. Models with heteroscedastic errors still outperformed models that assume constant variance.

Size remained an important driver of forecast performance with small VARs still outperforming their large counterparts. Small VARs on average over all horizons outperformed large VARs by 37 per cent. The best outperformance occurred at $h = 2$, at 58 per cent, with the least outperformance, at $h = 8$ at 16 per cent.

The best-performing model was no longer the TVP small VAR, as was the case with quarterly inflation, but rather the FAVAR models yet only over horizons of 3-, 4-, 5-quarters ahead. These are the only models that were able marginally to beat the random walk model.

The MSFEs only give an indication of the size of the error made on a point forecast, but we were also interested in their ability to predict turning points. To put this into perspective, Figure 3.5 plots the recursive 1- to 8-quarter ahead quarterly inflation forecasts estimated from the small TVP-VAR model with $\lambda = 0.99$ and $\beta_{T+h} \sim RW$ from 2000Q1. The graph is shown in year-on-year terms. It is clear from the graph that even the best model at predicting longer horizons of core inflation cannot successfully predict turning points. There is some evidence in 2002, 2006, 2008 and 2010 that the model could have some success in getting the direction right, but these three examples are overwhelmed by the poor performance elsewhere. The dominance of persistence as an important property of core inflation suggests little additional economic information that meaningfully guides turning points in core inflation.

¹³Results are available upon request.

Figure 3.5: Recursive estimates from small TVP-VAR, $\lambda = \lambda_t$, $\beta_{T+h} \sim RW$ 

3.6 Conclusion

In this chapter, we used a suite of econometric models to forecast quarterly and annual core inflation in SA, using 21 variables for the period covering 1981Q1 to 2013Q4. The forecasts were evaluated using the MSFE and the predictive likelihood relative to the random walk model for 1- to 8-quarters ahead. We found that most VAR models (specifically the small TVP-VARs and excluding the large TVP-VARs) performed better than the random walk model and other benchmark models for both forecast evaluation methods and over all horizons. Allowing for discrete structural breaks did not improve the forecast performance for core inflation. The structural model only performed better than the random walk model for $h = 3$ onwards, but was outperformed by other models. Further, the forecasts where we allowed for heteroscedastic errors in getting the shape of the predictive density outperformed VARs with homoscedastic errors. We also found that models with $\lambda = 0.99$ performed better than models where the forgetting factors were allowed to change over time. Overall, our results imply that information on the GDP growth rate and interest rate is sufficient to forecast quarterly core inflation accurately, but the relationship between these three variables needs to be modelled in a time-varying (nonlinear) fashion.

Camba-Mendez and Kapetanios (2005) used disaggregated price indices to forecast core inflation by employing factor models. In light of this, future research could be aimed at forecasting SA core inflation using disaggregated price indices based on time-varying models, to

see if such disaggregated information on price can produce more accurate forecasts than those obtained from the GDP growth rate and interest rates.

Table 3.2: MSFE relative to the random walk model for core inflation

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8
Full Model								
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.66	0.48	0.43	0.43	0.40	0.43	0.50	0.55
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} \sim RW$	0.93	0.63	0.52	0.52	0.47	0.54	0.60	0.62
TVP-AR								
TVP-AR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.80	0.57	0.45	0.41	0.40	0.44	0.50	0.50
TVP-AR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.85	0.67	0.54	0.48	0.47	0.52	0.58	0.60
TVP-AR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.79	0.55	0.45	0.42	0.40	0.43	0.49	0.50
TVP-AR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.86	0.67	0.54	0.48	0.48	0.52	0.58	0.60
Small VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.72	0.47	0.38	0.35	0.35	0.38	0.44	0.45
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.77	0.55	0.45	0.41	0.41	0.46	0.51	0.54
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.72	0.46	0.39	0.35	0.35	0.39	0.44	0.45
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.75	0.55	0.45	0.41	0.42	0.46	0.51	0.54
VAR, Heteroscedastic	0.78	0.58	0.48	0.43	0.43	0.48	0.54	0.57
VAR, Homoscedastic	0.83	0.71	0.59	0.53	0.53	0.60	0.64	0.71
Medium VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.86	0.51	0.41	0.41	0.39	0.43	0.48	0.52
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.91	0.64	0.52	0.49	0.48	0.53	0.58	0.63
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.86	0.50	0.41	0.41	0.39	0.43	0.48	0.51
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.92	0.64	0.51	0.50	0.47	0.53	0.57	0.63
VAR, Heteroscedastic	0.93	0.68	0.55	0.53	0.51	0.57	0.61	0.67
VAR, Homoscedastic	2.33	1.64	1.31	1.15	1.10	1.24	1.26	1.38
Large VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	1.07	0.96	0.80	0.75	0.72	0.79	0.85	0.93
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	1.22	1.22	1.01	0.91	0.87	0.96	1.00	1.12
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	1.04	0.96	0.82	0.75	0.72	0.80	0.84	0.94
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	1.22	1.22	1.00	0.90	0.88	0.95	1.00	1.11
VAR, Heteroscedastic	1.30	1.31	1.07	0.97	0.92	1.03	1.07	1.19
VAR, Homoscedastic	1.55	1.00	0.80	0.69	0.68	0.75	0.77	0.80
FAVAR								
TVP-FAVAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.94	0.65	0.50	0.48	0.48	0.52	0.59	0.57
TVP-FAVAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.97	0.78	0.61	0.56	0.56	0.62	0.68	0.69
TVP-FAVAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.92	0.64	0.51	0.47	0.48	0.52	0.58	0.56
TVP-FAVAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.97	0.78	0.61	0.56	0.56	0.62	0.69	0.69
FAVAR, Heteroscedastic	0.99	0.83	0.66	0.59	0.59	0.66	0.71	0.74
FAVAR, Homoscedastic	2.07	2.27	1.83	1.57	1.53	1.71	1.74	1.93
Benchmark Models								
Random Walk	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Small VAR OLS	1.00	1.01	0.99	0.96	0.99	1.14	1.18	1.26
AR(1) OLS	1.03	1.06	1.07	1.03	1.08	1.25	1.28	1.37
AR(1) Structural Breaks	1.68	1.18	0.96	0.87	0.85	0.91	0.90	0.94
Average performance								
Excluding benchmark models	1.01	0.80	0.65	0.60	0.58	0.64	0.69	0.74
Including benchmark models	1.01	0.82	0.68	0.63	0.61	0.68	0.73	0.78

Table 3.3: Sum of log predictive likelihoods relative to the random walk model for core inflation

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8
Full Model								
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} = \beta_T$	87.7	70.4	62.8	59.4	56.8	53.6	49.7	49.7
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} \sim RW$	82.3	63.1	58.1	52.2	49.4	47.9	44.7	45.8
TVP-AR								
TVP-AR, $\lambda=0.99, \beta_{T+h} = \beta_T$	82.3	67.3	61.5	55.6	53.4	53.2	48.6	49.0
TVP-AR, $\lambda = \lambda_i, \beta_{T+h} = \beta_T$	80.6	63.5	57.7	51.9	49.8	49.1	45.2	45.2
TVP-AR, $\lambda=0.99, \beta_{T+h} \sim RW$	81.9	67.5	61.0	54.9	53.7	52.9	48.9	49.2
TVP-AR, $\lambda = \lambda_i, \beta_{T+h} \sim RW$	80.5	63.6	57.3	51.8	49.8	49.1	45.5	45.3
Small VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	86.3	71.7	65.0	59.4	56.8	55.6	51.2	51.6
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} = \beta_T$	84.6	68.5	61.7	56.1	53.6	52.2	48.4	47.8
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	85.8	71.6	64.4	59.1	55.9	55.5	51.1	50.7
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} \sim RW$	84.6	68.5	61.4	56.0	52.9	52.2	49.1	47.8
VAR, Heteroscedastic	84.1	67.3	60.7	54.8	52.3	51.1	47.4	46.6
VAR, Homoscedastic	86.1	61.8	59.0	49.8	45.0	46.4	36.6	28.3
Medium VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	85.7	68.5	62.8	55.3	53.4	52.6	49.6	48.2
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} = \beta_T$	82.4	64.3	58.5	51.5	49.4	48.3	45.5	44.1
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	85.3	68.6	62.4	55.4	53.1	52.2	48.7	47.9
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} \sim RW$	82.8	64.9	58.3	51.0	49.7	48.3	45.4	43.9
VAR, Heteroscedastic	81.5	63.1	56.9	49.6	47.7	46.7	43.7	42.7
VAR, Homoscedastic	52.0	36.1	31.5	27.1	26.3	24.5	22.0	21.5
Large VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	78.7	55.1	48.1	42.4	40.5	39.1	36.4	35.8
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} = \beta_T$	74.9	48.3	41.5	36.1	34.6	33.7	31.5	30.4
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	79.4	55.0	47.3	42.1	40.4	39.0	36.8	35.5
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} \sim RW$	74.9	48.4	41.7	36.4	34.9	34.0	32.0	31.0
VAR, Heteroscedastic	73.0	46.0	39.5	33.8	32.8	31.6	29.8	28.4
VAR, Homoscedastic	57.5	44.7	41.7	34.3	31.9	31.9	30.5	35.1
FAVAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	84.0	63.3	58.1	51.3	49.3	48.3	45.0	46.0
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} = \beta_T$	82.0	59.0	54.0	47.8	45.8	44.1	40.6	41.7
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	84.2	63.7	57.9	51.6	49.7	48.2	45.4	46.4
TVP-VAR, $\lambda = \lambda_i, \beta_{T+h} \sim RW$	82.1	58.9	53.9	47.2	45.6	44.0	40.6	41.8
VAR, Heteroscedastic	81.3	57.6	52.2	46.6	44.4	42.9	39.6	40.1
VAR, Homoscedastic	56.7	24.9	20.5	15.1	15.5	14.6	12.4	11.5
Benchmark Models								
	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8
Random Walk	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Small VAR OLS	-1.5	-3.6	-5.4	-7.4	-9.3	-10.6	-13.0	-15.8
AR(1)	-0.7	-2.2	-3.5	-5.1	-6.8	-8.0	-10.5	-13.4
AR(1) Structural Breaks	-34.1	-7.8	6.1	12.0	17.1	20.3	24.1	27.0
Average performance								
Excluding benchmark models	79.6	59.9	53.9	47.8	45.8	44.8	41.4	41.0
Including benchmark models	73.4	55.1	49.5	43.8	41.9	40.8	37.6	37.1

Chapter 4

Decomposing inflation using micro-price-level data: South Africa's pricing dynamics

Contributing author: Neil Rankin

4.1 Introduction

Inflation, or the general increase in prices, is the result of many unobserved adjustments. Only a fraction of prices change in a month. Some of those prices will not have changed for over a year, while others will have changed last month. Some rise and fall faster than others. Some goods are on sale, while others are not. These dynamics matter a lot in themselves, as they describe pricing behaviour. But they also matter for the economic theory forming the foundation of how these prices, and hence inflation, are predicted and forecast.

To address these underlying dynamics, this chapter introduces a decomposition of South African goods inflation into its extensive margin¹ – *the fraction of prices changing in a specific month* – and its intensive margin – *the magnitude of price changes*. These are further decomposed into the magnitude and fraction of prices that are increasing and decreasing. Decompositions of this nature provide economists with the underlying price dynamics needed to both replicate the empirical properties found in consumer prices and to make choices about which models better fit this data. Models of micro-founded pricing dynamics generally fall into two categories: time-dependent or state-dependent, each having significantly different implications for pricing behaviour. Time-dependent models rely on firms setting prices every n^{th} period (as in Taylor, 1999) or randomly (as in Calvo, 1983), while state-dependent models are based on firms that face a cost to change prices and hence only change prices once the change is larger than a 'menu cost' (as in Mankiw, 1985). The role and incidence of sales can also have

¹This chapter refers to inflation as *goods inflation* unless specified otherwise.

important implications for the decomposition of inflation, as sales can be a vital source of price flexibility.

The contribution of this chapter is fourfold. First, we analysed a previously not available dataset of product-level data for the goods component of the consumer price index (CPI) from 2009 to 2015. Second, we have extended the analysis of the frequency of price changes by Creamer *et al.* (2012), looking at the distribution of frequencies by product over time. This analysis reveals that the distribution of frequencies changes significantly over time and by product. Third, we decomposed South African inflation into its extensive and intensive margins to provide a more in-depth analysis of inflation dynamics. Fourth, we looked at the role sales have on price dynamics in the South African economy. The classification of prices as 'regular' or 'sale' only started to become available in the latter part of the dataset used by Creamer *et al.* (2012), and it was only in 2011 that all observations in the underlying product data was classified.

The results of the decomposition of inflation reveal the following properties in South Africa. First, the average fraction of prices changing, or the extensive margin, is 27.8 per cent, but this can vary anywhere between 37 and 18 per cent in any particular month. This implies that prices change on average every 3.6 months. The median frequency of price changes is 12.5 per cent, implying that prices change every eight months at the median. There is substantial heterogeneity between products and over time, with the distribution having moderate positive skewness and heavy tails (excess kurtosis). Second, the magnitude of price changes, or the intensive margin, averages 0.83 per cent. With 27.8 per cent of prices changing and a magnitude of price changes of 0.83 per cent, monthly goods inflation increases amounts to 0.25 per cent, or 3.0 per cent annualised, from 2009 to 2015, i.e. 0.83×0.278 . Third, the variance of monthly inflation is mainly explained by the extensive margin, or the fraction of prices changing. This suggests that inflation in South Africa is state-dependent, driven by shocks to the economy and changes in input costs, rather than time-dependent. Fourth, the variance in inflation is dominated by price increases, which explains 70 per cent. Fifth, sales prices account for only around four per cent of prices and do not materially change our assessment of the role of the intensive and extensive margins in explaining inflation.

When it comes to the role of sales in the South African consumer goods basket, the results show that on average four per cent of products are on sale. The incidence of sales has risen since 2009, from around two per cent in January 2009 to over six per cent in December 2014, and to an average five per cent for the first five months of 2015. Sales are most common in the sub-categories of Furniture and furnishings (18 per cent of products in this category were on sale), Household appliances (11.9 per cent), Audiovisual and photographic equipment (9.6 per cent), and Household textiles (7.1 per cent), while they are least common in Vehicles (0 per cent), Telephone equipment (0 per cent) and Tobacco (0.2 per cent). Despite the relatively small number of sales that occur in South Africa, they remain an important contributor to price decreases and, hence, keeping goods inflation lower. From 2009 to 2015, inflation based on product-level data weighted using expenditure weights would have been 3 per cent instead of

the actual 4.8 per cent that it was, excluding all items on sales.

The chapter proceeds as follows: In section two we provide a brief literature review and contextualise our work within the existing literature. Section three provides details on the dataset used. Section four looks at the initial properties of the intensive and extensive margins, including extending the analysis of the fraction of prices by describing its distributional properties. Section five provides an analysis of the factors that explain the extensive margin. Section six uses the information in section four to decompose goods inflation and explain its level and variation. Section seven looks at how sales affect this decomposition. Section eight considers the impact of outliers on the main results of this chapter, and section nine concludes.

4.2 Literature review

There are two common approaches to the study of price behavior: surveys and micro-data analyses². Survey approaches started with the seminal work of Blinder (1991) and Blinder *et al.* (1998), who recognised that many theories of nominal price rigidity exist but that methods to censor the correct theories were proving inconclusive. The reason is that most theories “involve unobservable variables in an essential way, or they carry no real implications other than that prices are ‘sluggish’ in some unmeasurable sense, or both” (Blinder, 1991:3). Blinder asserts that if theories state that price-setters operate in a specific way, then you can just ask them whether this is true or not. This led to similar survey-based studies in countries including Germany (for example Kohler, 1996), the United Kingdom (Hall *et al.*, 1997), Canada (Amirault *et al.*, 2006), Turkey (Şahinöz and Saraçoğlu, 2008) and the euro area as a whole (Fabiani *et al.*, 2005). Govender (2013) is the only South African study to use survey-based techniques to determine the pricing behavior of manufacturing firms.

Govender (2013) found that the median adjustment of prices by firms was once a year, while 32 per cent of firms adjusted prices twice yearly. Over half of firms used current trading conditions to determine prices, while 24 per cent of firms took a view of the near future. Close to 70 per cent of firms used time-dependent instead of state-dependent pricing rules. Also, the most common pricing strategy was a mark-up (mainly a variable percentage but also a fixed percentage) above cost. Finally, in general, manufacturing firms in South Africa follow a barometric price leadership strategy. In this situation some firms have more and better information than others and act as a type of ‘barometer’ for less informed firms on price changes. Therefore, less informed firms wait for a pricing signal from the more informed firms.

The second approach uses micro-data analyses to understand the extent of price stickiness, and started with Stigler and Kindahl (1970), who collected individual transaction prices for intermediate products used in manufacturing (other notable examples in the early literature include Carlton, 1986; Cecchetti, 1986; and Kashyap, 1995). The datasets used in these studies

²A third approach is studying aggregate price indices to determine the behaviour of prices, including in the early work by Millard and O’Grady (2012) and Means (1935)

were generally narrow. More contemporary work on micro-datasets that started looking at large datasets includes Chevalier *et al.* (2003) on supermarket scanner data, and Bils and Klenow (2004) on data collected by the Bureau of Labour Statistics (BLS) for the US consumer price index. It was with this work that the full extent of pricing dynamics could be understood and models of prices tested. Other studies using micro-data include Baharad and Eden (2004); Álvarez *et al.* (2010); Dias *et al.* (2007); and Castellet and González (2004).

Work that uses micro-price data in South Africa is limited and includes Creamer and Rankin (2008), Creamer *et al.* (2012) and Aron *et al.* (2014a). These papers' focus is limited to the behaviour of prices (as in Creamer and Rankin, 2008 and Creamer *et al.*, 2012) and exchange rate pass-through to disaggregated consumer prices (Aron *et al.*, 2014a).

Creamer *et al.* (2012) found that prices change on average every five months, that there is substantial variation in price changes among goods, that price changes tend to be big on average but still a high frequency of small changes occur, that goods prices change more frequently than services prices, that the frequency of price changes increase with inflation, and that studying pricing behaviour at a micro level provides important information for econometricians with which to calibrate models. Aron *et al.* (2014a) exploit the granularity of micro-price data to study the behaviour of individual prices to exchange rate changes.

Once micro-price data studies had provided the underlying pricing dynamics, it became possible to assess models of inflation to determine which actually fit the data and accurately characterise price persistence. There are two types of models of micro-founded pricing behaviour: time-dependent and state-dependent. Time-dependent models rely on firms setting prices every n^{th} period (as in Taylor, 1999) or randomly (as in Calvo, 1983) and treat the determinant of price rigidity as exogenous. State-dependent models are based on firms that face a cost to changing prices and hence only change prices once the change is larger than a 'menu cost' (See, for example, Sheshinski and Weiss, 1977; Mankiw, 1985; and Caplin and Spulber, 1987). The treatment of inflation based on time- or state-dependence leads to different dynamics of prices. For example, from a monetary policy perspective, time-dependent models lead to more persistent impacts of monetary shocks to the real economy. Also, inflation responds more rapidly in menu cost models. Models of inflation also make predictions about whether the extensive (*fraction of prices changing*) or intensive (*size of price changes*) margin dominates inflation outcomes when faced with a monetary shock. In Dotsey *et al.* (1999), the model predicts that the majority of the response comes from a change in the fraction of price changes, while in Golosov and Lucas Jr (2007) it is the size of price changes and the incidence of price increases and decreases.

Bils and Klenow (2004) show that underlying price dynamics are at odds with conventional time-dependent models of price stickiness (as in Calvo, 1983 and Taylor, 1999). The durations of price changes for the majority of the 123 categories of products were significantly more 'volatile and transient' than implied by these models. Also, inflation was less related to the frequency of price changes suggested in these models. Klenow and Kryvtsov (2008) show

that most models of inflation could not replicate all the stylised facts discovered using micro-price data. They found similarly that the Taylor and Calvo models do not support empirical facts, including predicting large price changes for older prices, which does not occur. In the state-dependence domain, they found that the Dotsey *et al.* (1999) model does not produce significantly large price changes and allows the extensive margin too much importance when compared with the empirical data. Similarly, the Golosov and Lucas Jr (2007) model cannot replicate enough small price changes. Álvarez (2008) takes this analysis one step further by surveying the existing literature on micro-price data and then applying these facts to 25 pricing models including sticky information, menu costs, time-dependent, costs of adjustment and customer anger models³. He similarly concludes that no model can replicate all the empirical findings from micro-price data.

The variable and often poor performance of pricing models against the empirical data led to a reformulation of these models in order to address shortcomings (see, for example, Midrigan, 2011; Nakamura and Steinsson, 2010; Gertler and Leahy, 2008; Costain and Nakov, 2011a; Costain and Nakov, 2011b; and Dotsey *et al.*, 2013). Midrigan (2011) focused on two empirical facts that models were unable to replicate – the large number of small price changes and the distribution of prices exhibiting excess kurtosis – formulating a menu-cost model that addressed these findings and therefore better replicated the large-aggregate business cycle fluctuations that were the success of time-dependent models. Costain and Nakov (2011a,b) developed a general model of state-dependent pricing that nests a continuum of models between two extremes: Calvo (1983) and fixed menu-costs. Dotsey *et al.* (2013) extended their earlier model to include “stochastic variation in productivity at the firm level” to address the outsized importance of the extensive margin.

Our contribution to this literature is as follows. First, we used a previously unavailable dataset of product-level data for the goods component of the CPI and updated the micro-price information provided in Creamer *et al.* (2012) for 2009 to 2015. Second, we extended the analysis of the frequency of price changes by Creamer *et al.* (2012), looking at the distribution of frequencies by product over time. Third, we decomposed South African goods inflation into its extensive and intensive margins to provide a more in-depth analysis of inflation dynamics. This analysis reveals that state-dependent modelling strategies of pricing better replicate pricing dynamics in SA. Fourth, we looked at the role sales have on the price dynamics in the South African economy. The classification of prices as ‘regular’ or ‘sale’ only started to become available in the latter part of the dataset used by Creamer *et al.* (2012), and it was only in 2011 that all observations in the underlying product data were classified.

³See Table 11 in Álvarez (2008) for a summary of the ability of models to replicate empirical findings.

4.3 The dataset

The micro-dataset used in this study is based on the underlying product data provided by Statistics South Africa (StatsSA) and used to produce CPI. It covers the period January 2009 to May 2015 and is an extension of the dataset used by Creamer and Rankin (2008), which included data up to December 2007. We started in 2009 to ensure a compatible dataset with no structural breaks. In January 2009, StatsSA introduced a new CPI basket based on the Classification of Individual Consumption by Purpose (COICOP). This included dropping over 600 goods and services that were collected under the old methodology – products dropped included items such as bath salts, guitar, white bread rolls, and snoek (type of fish) – while 72 were added – including CDs and DVDs, sporting tickets, and teddy bears. Our dataset includes only goods and does not provide any information on services. This is different to the dataset used in Creamer and Rankin (2008) and Creamer *et al.* (2012) because although that dataset comprised predominantly goods products it also included services. There were 5,200,466 individual price quotes in the period under review. Appendix C provides detail on the categories provided as well as a snapshot of what the dataset looks like.

In order to prepare the dataset for analysis, we include only data with an acceptable status code. This means that prices collected which were indicated as ‘Wrong item collected’, ‘Item available but not comparable’, ‘Extreme values not verified’, ‘Quality adjustment’ and ‘Available shelf price wrongly collected’ were excluded. This left 4,986,454 individual price quotes, a drop of 214,006. Despite a robust approach to processing this information at StatsSA, it is possible that due to the size of the dataset mistakes in price collection still exist after this censoring. Generally, in a dataset of this size, extreme outliers – assuming a normally distributed sample – should account for less than 1000 price quotes. To address this problem, we checked the robustness of our results to outliers. We also looked at the role of product substitutions, given the relative importance this can have on price dynamics (see Klenow and Malin, 2010); however, the dataset includes only 30 examples of this, and therefore, it was ignored. Finally, we removed all prices recorded as zero to ensure accurate frequency and magnitude calculations.

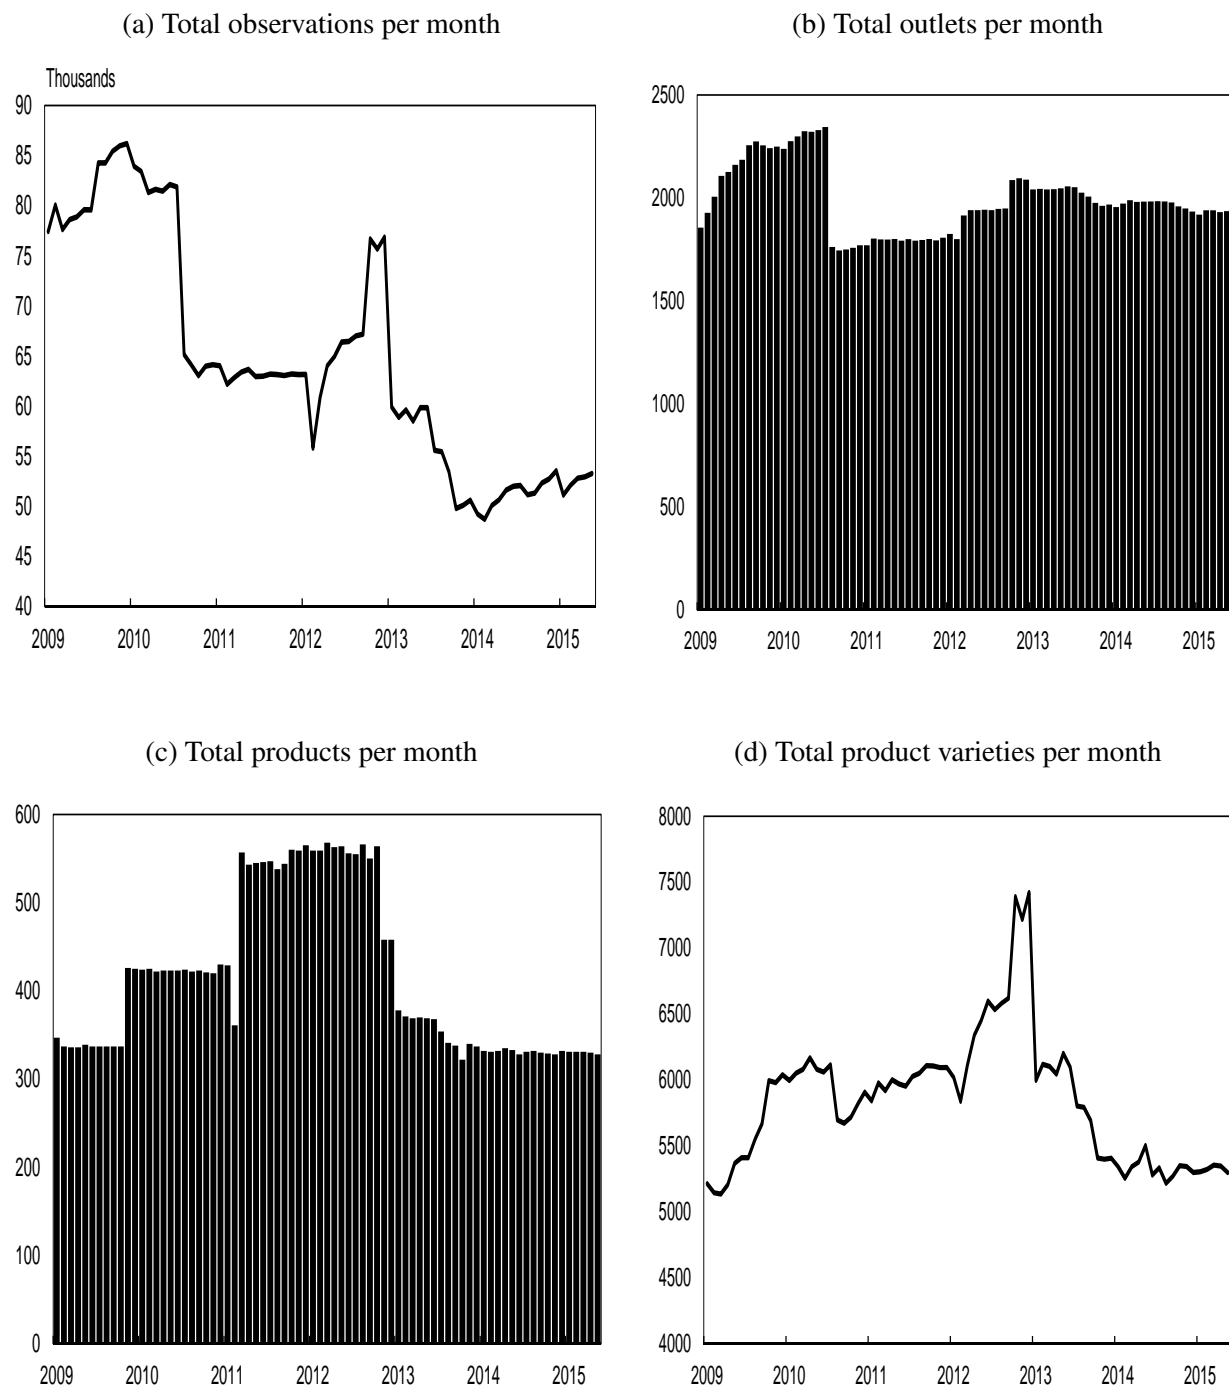
The dataset includes multiple observations of individual products from all over the country, including both rural and urban areas⁴. For example, in the entire dataset, there are over 72,000 individual price observations for a loaf of brown bread. No single product dominates the dataset, with the maximum contribution of a single product at 1.5 per cent. Products were also collected from 3,497 outlets nationwide.

Figure 4.1 shows the number of individual price records (4.1a), the number of outlets (4.1b), the number of products (4.1c), and the number of product varieties (4.1d) over time. The number of product varieties distinguishes between product brands, and their type; i.e. one-ply versus

⁴Although we compared the results to the goods CPI for all urban areas from StatsSA, we did not feel the need to drop price data from rural areas in this study, as the objective was not to replicate the CPI but to uncover the behaviour of prices in the economy.

two-ply toilet paper. Initially, the number of observations per month was over 75,000 but this drops over time to around 52,000 in 2014-15. A significant drop occurs in August 2010. The number of outlets at which prices were collected and available each month varies between 1,745 and 2,344. There is an equivalent precipitous drop in the number of outlets in mid-2010. This is most likely due to methodological changes. The data sample has on average 415 unique products per month. There is quite a lot of variation in the number of goods included in the dataset over time, initially around 347 in January 2009, before rising to 568 in March 2012, and again declining towards the end of the sample. There is also significant variation in the number of product varieties, between 5,136 and 7,422.

Figure 4.1: Basic properties of the data



4.4 Properties of price dynamics

In this section we present an update to the frequency and size of price changes in South Africa for the period January 2009 to May 2015. We also extend the analysis of frequency in Creamer *et al.* (2012) by providing a discussion of the distribution of the frequency of products over time. Creamer *et al.* (2012) provides initial estimates of the frequency of price changes from 2001 to

2007. The authors show that prices changed once every five months, that the frequency of price changes was heterogeneous, and that price changes were large on average, but that many small price changes also occur⁵.

4.4.1 Size of price changes

We define a price change in this chapter as:

$$dp_{j,k,t} = (p_{j,k,t} - p_{j,k,t-1}) \cdot 100 \quad (4.1)$$

where $p_{j,k,t}$ is the log price of a specific variety of product j at retailer k in time t . A variety of product refers to a unique brand or type of product. For example comparing one- and two-ply toilet paper from a number of different brands at a specific retailer or firm. To ensure that we compared price changes of identical products over time we created a unique identification number for each product, in a specific region, at a specific outlet, for a specific month, and of a specific type.

The price changes $dp_{j,k,t}$ are then aggregated using either the mean or median to the product level i , representing the consumption products collected using the COICOP methodology. Therefore the mean price change at product level i is:

$$dp_{i,t} = \frac{\sum_{j=1}^J \sum_{k=1}^K dp_{j,k,t}}{J + K} \quad (4.2)$$

Using the earlier example this would be the aggregation of all varieties of toilet paper as an individual product. Other examples of products in our dataset are loaf of brown bread, beef – rump steak fresh, feta cheese, sport shoes (for women), firewood, cough syrup, and printer.

We also normalised price changes by its mean, $\mu_{dp_{i,t}}$, and standard deviation, $\sigma_{dp_{i,t}}$ such that:

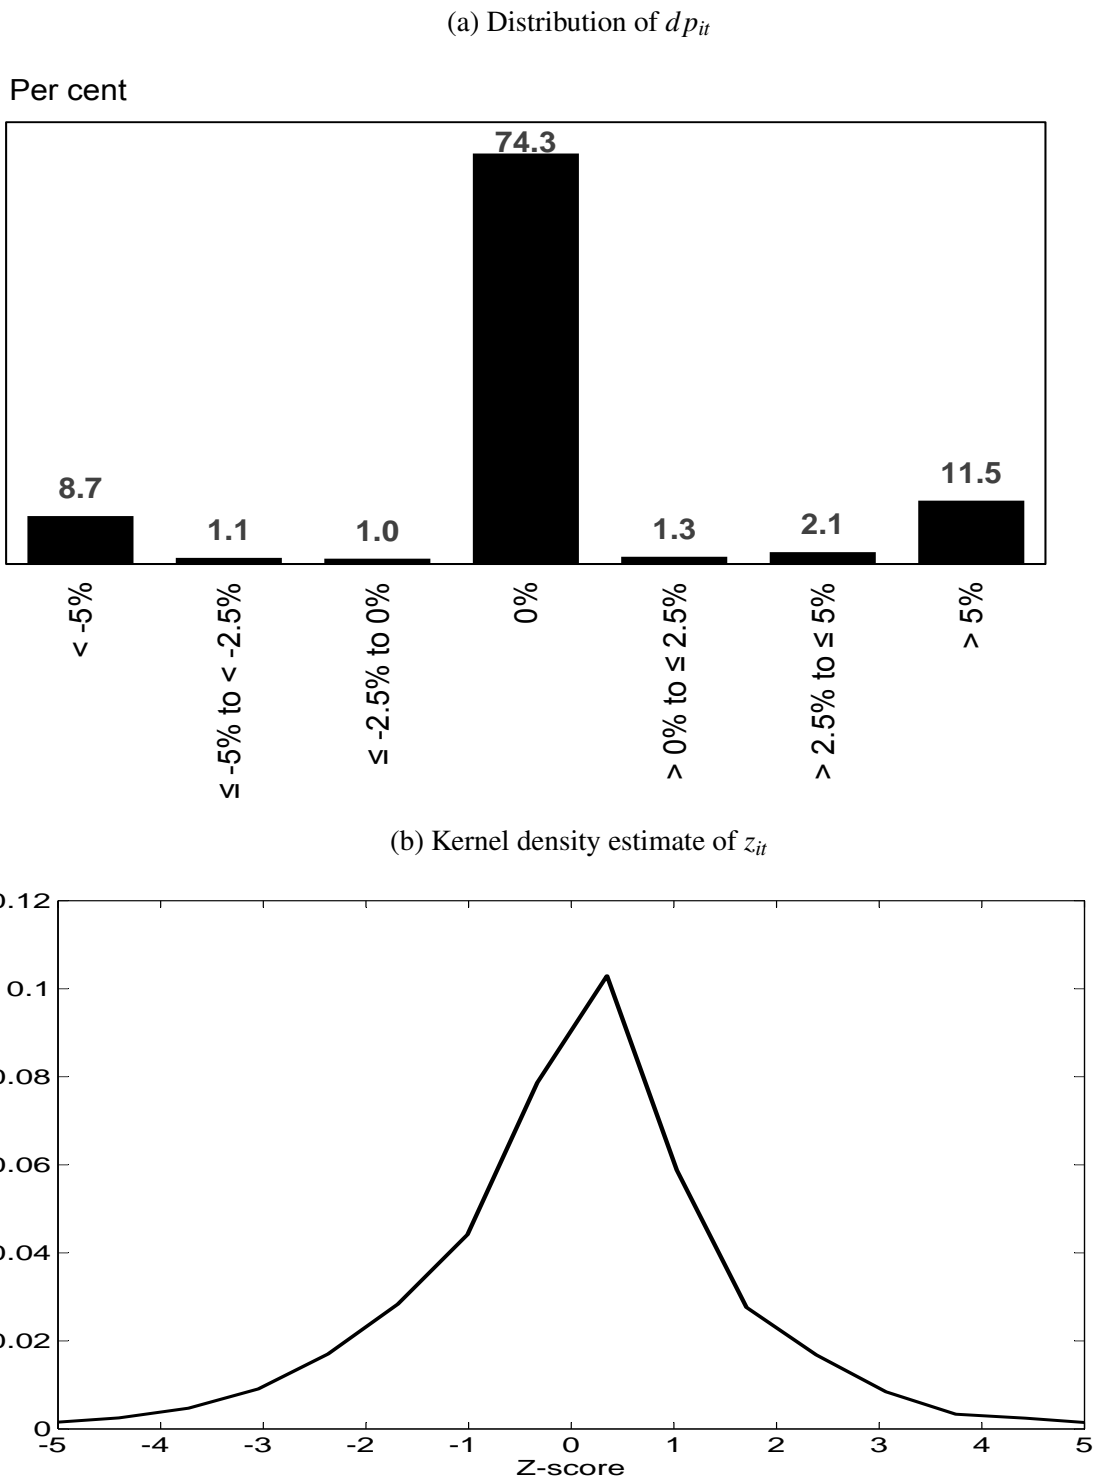
$$z_{i,t} = \frac{dp_{i,t} - \mu_{dp_{i,t}}}{\sigma_{dp_{i,t}}} \quad (4.3)$$

Figure 4.2 plots the distribution of $dp_{i,t}$ (4.2a) and $z_{i,t}$ (4.2b). It is clear from figure 4.2a that actual monthly price changes are dominated by no change (0), which accounts for 74.3 per cent of $dp_{i,t}$. Prices that do change are generally large in absolute terms. Of the number of price changes, 11.5 per cent are larger than 5 per cent while 8.7 per cent are smaller than –5 per cent. When comparing these results to Creamer *et al.* (2012), we see that the sample from 2009 to 2015 is comparable with regard to the number of 0 observations, as well as the number of large positive price changes, those above 5 per cent. However, our sample has a substantially higher proportion of negative price changes compared with the 3.89 per cent found in the CPI micro-price data from 2001 to 2007.

⁵Creamer *et al.* (2012) highlighted a number of other important properties of prices in SA, including that price changes are not synchronised to the business cycle, and that neither size nor frequency increase in the age of prices, but these are not of interest in this chapter.

Figure 4.2b plots an estimated kernel density function of $z_{i,t}$. Of the distribution of $z_{i,t}$, 87.6 per cent is within one standard deviation, 94.8 per cent is within two standard deviations, and 97.8 per cent is within three standard deviations. The standard deviation of $dp_{i,t}$ is equal to 11.4 per cent.

Figure 4.2: Distribution of price changes



4.4.2 Frequency of price changes


To calculate the frequency of price changes, we created an indicator variable $I_{i,t}$ that is equal to 1 if there was a price change and 0 otherwise. This is aggregated up to the product level from varieties j and firms k as above. For price increases, $I_{i,t}^+ = 1$ for $dp_{i,t} > 0$ and 0 otherwise, and similarly for price decreases, $I_{i,t}^- = 1$ for $dp_{i,t} < 0$ and 0 otherwise. This was then aggregated to the product level using both the mean and median. We then applied the CPI weights, dynamically, to get weighted frequency changes, i.e. we used the same weights as were used to calculate the overall CPI. Since we did not have full coverage such that the weights added up to 1, we normalised the frequency calculations.

Over the sample period, the weighted mean frequency of price change is 27.8 per cent while the median is significantly lower, at 12.5 per cent. This was calculated at the individual product level. In order to get an approximation of the duration in months between price changes we took the inverse of the frequency measure. This implies that the prices of goods changed on average every 3.6 months, while only every 8.0 months at the median. The approximate estimate of duration requires that all prices have the same expected duration, an assumption that is unlikely to hold given the heterogeneity in price changes. Another method of calculating the average duration involves taking the inverse of the frequency at the product level and aggregating that back to an overall value, as is done in Dhyne *et al.* (2006). Using this method we calculated that the weighted average duration of price changes is 6.5 months.

The frequency of goods price changes in SA increased in the sample period of 2009 to 2015 compared with the findings of Creamer *et al.* (2012). They found that on average the frequency of price changes in goods was 17 per cent from 2001 to 2007 compared with 27.8 per cent in our sample. Cross-country comparisons show that SA is now more similar to the US than Europe. In the US there is evidence that the mean frequency of prices changes is diverse but converges in the upper 20s. Bils and Klenow (2004) found an average frequency of 26.1 per cent. Klenow and Kryvtsov (2008) found that for regular prices the frequency is 30 per cent and 36 per cent for posted prices. Finally, Nakamura and Steinsson (2008) found that the frequency of price changes including sales was between 26 and 28 per cent. In Europe, Dhyne *et al.* (2006) found that over the period 1996 to 2001, prices changes occur 15.1 per cent of the time for all euro area countries, excluding Ireland and Greece. An important cross-country finding in Klenow and Malin (2010) is that prices in Europe tend to change the least frequently, followed by the US, and then high inflation developing countries such as Brazil, Mexico and Sierra Leone.

Figure 4.3 plots the weighted mean and median frequency of price changes per month from 2009 to 2015. The frequency of price changes has a slight downward slope, falling from an average 37.4 per cent in March 2009 to 29.4 per cent in May 2015. The decline in the frequency is more clear when looking at the median, with a significant break occurring in the second half of 2012 and early 2013.

Figure 4.3: Mean and median frequency of monthly price changes over time



An important distinction worth highlighting is the difference between mean and median frequencies. In SA, we found that the median frequency change was significantly lower than the mean. This implies that the distribution of frequency of price changes are skewed to the left or has a positive skew. When looking at frequency by product over time, we note that skewness averages 1.36. In the US, Bils and Klenow (2004) found that the mean frequency is also higher than the median, but the spread is not as large: 20.9 per cent compared to 26.1 per cent. This empirical fact highlights important heterogeneity in the frequency of price changes by product.

4.4.3 Heterogeneity in the frequency of price changes


Another way to look at the heterogeneity in the frequency of price changes is to look at the unweighted distribution of this frequency by product over time. Figure 4.4 provides a 3-D kernel density estimate of the frequency of CPI products from 2009 to 2015⁶. The x-axis is the frequency estimate by products, the y-axis is the date, and the z-axis is the density function.

Figure 4.4 shows that the underlying distribution of frequencies changes over time and is multi-modal. Generally there is clustering around 0 and 1 with products that do not change and those that always change price in a specific month. There is also clustering around the median, which averages 19 per cent over the sample (this is unweighted). Between March 2011 and May 2012 there was a significant rise in the number of products that do not change price and, similarly, those that do. This means that although the proportion of products at the mode was

⁶This graph was generated in Matlab using the `ksdensity` function. The bandwidth for kernel smoothing was set to 0.05. The estimate was also generated with supports at -0.2 and 1.2 such that the function used the transformation $\text{Log}\left(\frac{x-(-0.2)}{1.2-x}\right)$.

higher during this period (around 26 per cent compared with 11 per cent for the entire sample), there was also more dispersion in the frequencies of products. There is nothing necessarily significant during that period from a macroeconomic perspective that may shed light on this phenomenon.

Figure 4.4: Distribution of frequency of monthly price changes by product over time



Three measures can help highlight important heterogeneity over time in the underlying distribution of price changes. First, the distribution is always positively skewed with an average skewness of 1.36. This, however, moves in the range of a moderate positive skew, 0.59, in March 2014, to a maximum value of 2.08 in December 2013. The distributions are also always heavy-tailed, with an average kurtosis of 5.8. Again this varies significantly, from a distribution that is almost normal in December 2013, with a kurtosis of 3.16 (remember this is the same date as the lowest positive skew) and a maximum of 10.2. Finally, the coefficient of variation averages 0.84 over time, with a minimum value of 0.59 and a maximum of 1.15. The coefficient of variation is higher during March 2011 and May 2012, when we observe significantly more 0 and 1 observations.

Further heterogeneity occurs at the category level. Table 4.1 provides information on the frequency of price changes in aggregate as well as increases and decreases, and the weight coverage, for 23 sub-groups of the CPI. Only goods are included in this analysis, and therefore, categories such as Operation of vehicles do not include the price of a major or minor service, for example.

The subcategory with the most frequent price changes is Vehicles (39.7 per cent) followed by Miscellaneous goods (39.4 per cent) and, unsurprisingly, Food (34.1 per cent). This implies that Vehicle prices changed once every 2.5 months. The categories with the least frequent price changes are Footwear (7.0 per cent), Clothing (7.4 per cent), and Hotel and restaurant goods (9.0 per cent). This implies that the category with the least frequent price changes occur once every 14.4 months. These results are generally consistent with the findings in Creamer *et al.* (2012), who also found that Other goods and services, Food and Transport experienced the most frequent price changes, while Communication, Footwear, and Clothing experienced the least frequent price changes. The categories are generally the same, but the ordering does differ slightly between periods.

All sub-categories, except three, experienced a higher frequency of price increases versus decreases. The three exceptions are Furniture and furnishings, Vehicles and Telephone equipment. The incidence of price decreases was highest in Vehicles where 25.5 percentage points of the 39.7 per cent price changes were decreases. One explanation for the higher incidence of price decreases in these categories could be sales. This is true for Furniture and furnishings in which goods were on sale 18 per cent of the time in this dataset. However, the Vehicles and Telephone equipment sub-categories did not have any recorded sales items. The high incidence of price changes in Vehicles is due to Used vehicle prices declining over time. The fraction of Used vehicle prices changing was 88 per cent, dominated by price decreases, while new car prices changed 21 per cent of the time, dominated by price increases. Telephone equipment was likely to have a higher incidence of price decreases, since this is an area where technological gain has rapidly decreased prices over time. In Creamer *et al.* (2012) only one category experienced a higher frequency of price decreases, which was Footwear.

4.5 What factors explain the frequency of price changes?

What explains frequency over time? We used the determinants of a simple micro-founded mark-up model of prices, as in Campa and Goldberg (2005), and extended it to CPI as in Aron *et al.* (2014a) to build an equation of frequency. The regression model we estimated is:

$$fr_t = c + \beta_\pi \sum_{i=1}^6 \Delta \log(cpi_{t-i}) + \beta_{er} \sum_{i=1}^6 \Delta \log(er_{t-i}) + \beta_{ulc} \sum_{i=1}^6 \Delta \log(ulc_{t-i}) + \beta_y \sum_{i=1}^6 \Delta \log(y_{t-i}) + \beta_{\pi^*} \sum_{i=1}^6 \Delta \log(cpi_{t-i}^*) + \gamma X_t \quad (4.4)$$

Table 4.1: Average weighted frequency of price changes by CPI category

	Frequency of price changes (%)	Frequency of price increases (%)	Frequency of price decreases (%)	2013 CPI weight
Food	34.1	19.9	14.1	12.3
Non-alcoholic beverages	26.1	15.7	10.4	1.2
Alcoholic beverages	18.1	13.1	5.0	4.0
Tobacco	20.3	17.1	3.2	1.5
Clothing	7.4	4.0	3.4	2.8
Footwear	7.0	3.8	3.2	1.3
Maintenance and repair of dwelling	16.7	10.6	6.1	1.0
Other fuels	25.1	15.1	10.0	0.0
Furniture and furnishings, carpets and other	25.8	12.9	13.0	0.5
Household textiles	16.7	8.7	8.0	0.6
Household appliances	25.8	13.8	12.0	0.6
Glassware, tableware and household utensils	13.9	7.5	6.5	0.1
Tools and equipment for house and garden	17.0	9.8	7.2	0.1
Goods for routine household maintenance	28.2	16.3	11.9	0.5
Medical products, appliances and equipment	20.3	13.4	6.9	0.4
Vehicles	39.7	14.1	25.5	6.0
Operation of vehicles	27.5	17.7	9.9	0.6
Telephone equipment	27.9	8.0	19.8	0.1
Audiovisual and photographic equipment	23.1	9.9	13.2	0.7
Other recreation equipment	17.8	10.8	7.0	0.8
Newspapers, books and stationery	14.9	9.5	5.4	0.0
Hotel and restaurant	9.0	7.1	1.9	2.5
Miscellaneous goods	39.4	21.6	17.8	1.9

where fr_t is the frequency of price change, cpi_t is an index of seasonally adjusted monthly domestic consumer goods prices multiplied by 100, cpi_t^* is foreign prices proxied by foreign wholesale prices weighted according to export and import weights, er_t is the spot rate of the rand against the US dollar, y_t is domestic demand proxied by the volume of manufacturing production, ulc_t is unit labour cost proxied by an interpolated manufacturing ULC, and X_t are other variables of interest, including the number of observations each month and seasonality. Since the factors that explain frequency are unlikely to have strong contemporaneous relationships, we implemented lags of order six. To overcome any problems with degrees of freedom, we followed Aron *et al.* (2014b) and use 'parsimonious longer lags' (PLL)⁷. In the equation above, we used $\Delta_3 \log(cpi_{t-3})$ and $\Delta_6 \log(cpi_{t-6})$ to replace lags from three to six. $\Delta_3 \log(cpi_{t-3})$ is the three-monthly change in cpi , lagged three periods.

Table 4.2 presents the long-run coefficient results of an ordinary least squares (OLS) regression of equation 4.4 with heteroscedastic and autocorrelation consistent (HAC) standard errors. All coefficients that are significant at a 5 per cent level of significance are in bold. The long-run coefficient values (6-months) and their significance are based on a Wald test for all the variables of interest. The overall adjusted R^2 is 0.7, with an F-statistic of 5.76. The errors are serially uncorrelated, with no evidence of heteroscedasticity according to the Breusch-Pagan-Godfrey test. The errors are also normally distributed according to the Jarque-Bera test.

Table 4.2: Regression results of frequency: long-run coefficients

	Coefficient	Std. error	t-statistic
$\Delta \log(cpi)$	1.44	1.80	0.80
$\Delta \log(er)$	0.65	0.20	3.26
$\Delta \log(ulc)$	3.65	1.31	2.79
$\Delta \log(cpi^*)$	4.00	1.24	3.22
$\Delta \log(y)$	-1.40	0.71	-1.98
Adjusted R^2	0.70		
F-statistic	5.76		
DW-stat	2.37		
JB-stat	0.61		

The regression results show that the exchange rate, unit labour cost, foreign inflation and demand all have a significant relationship with frequency, at a 5 per cent level of significance. The significance of these variables suggests that both domestic and imported pricing pressures increase the frequency of price changes. On the domestic side, a 1 per cent increase in the exchange rate leads to a 0.65 per cent increase in frequency. A 1 per cent increase in domestic unit labour costs increases the extensive margin by 3.65 per cent. On the foreign side, a 1

⁷See the online Appendix 3 from Aron *et al.* (2014b) for more information.

per cent increase in foreign inflation increases the extensive margin by 4.0 per cent. Domestic demand is also significant but has the opposite sign to what is expected, suggesting an increase in production lowers the frequency of price changes.

There is some evidence of seasonality in price changes, with January, February, March, April, May, August, September, October and November all being significant at a 5 per cent level of significance. January and March are by far the most significant seasonal dummies, with the highest t-statistic values. January is likely to be a month during which most firms following a time-dependent pricing strategy would consider changing prices. March is likely significant from a frequency perspective, since taxes for certain items, especially alcoholic beverages and tobacco, rise during this month. The results are not affected by the number of observations since it is not a significant predictor of frequency.

4.6 Decomposing inflation

4.6.1 The intensive and extensive margins

Having the underlying price data used to construct CPI inflation allowed us to decompose inflation into an extensive margin (EM), or *the fraction of items that change price in a particular month*, and an intensive margin (IM), or *the average magnitude of price changes*, as in Klenow and Kryvtsov (2008). The advantage of this decomposition is threefold. First, this allows us to determine whether inflation is driven by the actual average size of changes or how often prices change. Second, and related to the first, is the ability to determine what drives the variance in inflation. Third, inflation is the outcome of many different price increases and decreases as well as changing and unchanging prices all occurring in a particular month. This decomposition can aid in understanding these dynamics over time and enriching the picture we have of aggregate consumer goods inflation.

Inflation can be decomposed into IM ($dp_{i,t}$) and EM ($fr_{i,t}$) such that:

$$\begin{aligned}\pi_t &= \sum_i \omega_{i,t}(p_{i,t} - p_{i,t-1}) \\ &= fr_t \cdot dp_t \\ &= \sum_i \sum_t \omega_{i,t} I_{i,t} \cdot \frac{\sum_i \sum_t \omega_{i,t} (p_{i,t} - p_{i,t-1})}{\sum_i \sum_t \omega_{i,t} I_{i,t}},\end{aligned}\tag{4.5}$$

where $p_{i,t}$ is the log price of product i at time t , π_t is monthly inflation, $I_{i,t}$ is an indicator that is equal to 1 if there was a price change and 0 otherwise, and $\omega_{i,t}$ is the CPI weights. Equation 4.5 decomposes inflation into the fraction of items changing price in month t , $fr_t = \sum_i \sum_t \omega_{i,t} I_{i,t}$, and the weighted average magnitude of price changes, $dp_t = \frac{\sum_i \sum_t \omega_{i,t} (p_{i,t} - p_{i,t-1})}{\sum_i \sum_t \omega_{i,t} I_{i,t}}$.

Table 4.3 reports the results for the components of inflation, showing the mean and standard deviation for each component, as well as their correlation to inflation, and the results of a regression analysis. We regressed each component successively on monthly inflation and report

the coefficient value, its significance and the R^2 . The standard errors are HAC. Regression results represent simple correlations and should be interpreted with caution.


Figure 4.5 replicates Figure 9 in Klenow and Kryvtsov (2008) and shows the 12-month moving average for the extensive margin (fr_t), scaled for ease of reference, the intensive margin (dp_t) as well as the annual inflation rate weighted using CPI weights provided by StatsSA. Monthly goods inflation over the sample period from February 2009 to May 2015 averages 0.25 per cent (or 3.0 per cent on an annual basis) with a standard deviation of 0.32 per cent (see Table 4.3). This is due to an average 27.8 per cent fraction of prices changing every month (EM), and a 0.83 per cent average monthly increase in prices (IM).

It is clear from the graph that the intensive margin, or average monthly increases in prices (dp_t), is a significantly more important driver of inflation. Regressing dp_t on inflation indicates that a 1 per cent increase in the intensive margin increases inflation by 0.29 per cent and explains 97 per cent of the variation in inflation according to the R^2 .

The fraction of prices changing (fr_t), on the other hand, plays a less important role in the level of inflation. The correlation between fr_t and inflation is 0.53. Regressing fr_t on inflation shows that a 1 per cent increase in the frequency of price changes increases the monthly inflation rate by 0.04 per cent, with an R^2 of 0.28. It varies between 18 and 37 per cent with a coefficient of variation of 0.14.

Over the sample period, it was twice as likely for the fr_t to change by 1 per cent (85.5 per cent of the time) as it was for dp_t to change by 1 per cent. Also, dp_t was much more strongly correlated with inflation at 0.98. Over the period under study monthly inflation increased by as much as 3.2 per cent and declined by a maximum of -1.3 per cent. The coefficient of variation is significantly larger for dp_t , at 1.3.

Figure 4.5: Intensive and extensive margins



4.6.2 Price increases and decreases

A fraction of prices are changing, and aggregate prices are rising each month on average, but this is due to an interplay of goods price increases and decreases. An important finding in the literature on pricing behaviour is that a distinction between price increases and decreases matters for the dynamics of inflation. Nakamura and Steinsson (2008) show that level of inflation in the US, for two decades to 2008, is driven by price increases (pos_t) and not decreases (neg_t). This is contradicted in Klenow and Kryvtsov (2008), who found that not only do price decreases matter but they tend to explain 40 – 50 per cent of inflation variance. Gagnon (2009) takes this analysis further, by showing that price decreases, rather than the magnitude of price increases was a key driver of dynamics between high- and low-inflation environments, based on Mexican micro-price data.

In order to determine the contribution of price decreases and increases, we decomposed inflation further into the frequency of price increases and decreases, combined with the size of those increases and decreases. Hence, inflation in Equation 4.5 can also be represented as:

$$\begin{aligned}\pi_t &= fr_t^+ \cdot dp_t^+ + fr_t^- \cdot dp_t^- \\ &= pos_t + neg_t,\end{aligned}\tag{4.6}$$

This is based on the decomposition of fr_t into those prices that are increasing and decreasing:

$$fr_t = fr_t^+ + fr_t^- = \sum_i \sum_t \omega_{i,t} I_{i,t}^+ + \sum_i \sum_t \omega_{i,t} I_{i,t}^- \tag{4.7}$$

where $I_{i,t}^+ = 1$ if $p_{i,t} > p_{i,t-1}$ and $I_{i,t}^- = 1$ if $p_{i,t} < p_{i,t-1}$, else it is equal to 0. A similar decomposition for dp_t can be calculated where $dp_t^+ = \frac{\sum_i \sum_t \omega_{i,t} I_{i,t}^+ (p_{i,t} - p_{i,t-1})}{\sum_i \sum_t \omega_{i,t} I_{i,t}^+}$. Finally, we can define the positive part of inflation as $pos_t = fr_t^+ \cdot dp_t^+$ and the negative component as $neg_t = fr_t^- \cdot dp_t^-$.

Table 4.3 shows the contribution of the decomposition of the intensive and extensive margins into its positive and negative components. Both price increases and decreases occur each month, but price decreases are less frequent than price increases. Over the sample, 14.8 percentage points of price changes were due to price increases (fr_t^+) compared with 12.9 percentage points due to price decreases. The standard deviation of the frequency of price increases is also larger at 3.8 per cent compared with 2.9 for price decreases. As expected, the frequency of price increases are also more strongly correlated with inflation at 0.8, while price decreases are weakly and negatively correlated at -0.3 . The opposite signs of correlation between price increases and decreases lower the overall correlation of fr_t to inflation. The regression results support the relative importance of the frequency of price increases to decreases, with R^2 for frequency of price increases at 0.64 compared with 0.1 for the frequency of price decreases. A 1 per cent increase in fr_t^+ (fr_t^-), leading a 0.069 per cent increase (-0.037 per cent decrease) in monthly inflation, but these differences are not large.

Table 4.3: Decomposing inflation into intensive and extensive margins

Variable	Mean,	Standard deviation,	Correlation	Regression on π_t		
	%	%		with π_t	Coefficient	S.E
π_t	0.25	0.32				
fr_t	27.8	4.01	0.53	0.043	0.006	0.28
dp_t	0.83	1.12	0.98	0.285	0.009	0.97
fr_t^+	14.81	3.77	0.81	0.069	0.004	0.64
fr_t^-	12.89	2.85	-0.32	-0.037	0.014	0.1
dp_t^+	4.18	0.98	0.75	0.248	0.022	0.56
dp_t^-	-2.98	0.75	0.65	0.284	0.043	0.43
pos_t	0.63	0.25	0.95	1.223	0.061	0.9
neg_t	-0.38	0.12	0.73	2.031	0.167	0.54


The frequency of price changes has increased in SA compared with the results found in Creamer *et al.* (2012) for the period 2001 to 2007, but price increases remain more prevalent than decreases. They found that for goods, the mean frequency of price increases was 10.8 per cent and for price decreases 6.1 per cent. Our results are similar to what Klenow and Kryvtsov (2008) found for the US. Over the period 1988 to 2004, the mean frequency of price increases was 15 per cent compared with 11.5 per cent for decreases.

Moving to the intensive margin, the average magnitude of price increases was 4.2 per cent, and price decreases -3.0 per cent. dp_t^+ is more correlated with inflation than dp_t^- , and similarly more volatile. dp_t^+ and dp_t^- are also less correlated to monthly inflation than dp_t . Regression results suggest that a 1 per cent increase in dp_t^+ (dp_t^-) increases inflation by about 0.25 (0.28) per cent, with dp_t^+ having a higher R^2 at 0.56 compared with 0.43.

An analysis of pos_t and neg_t supports the finding that price increases mattered more than price decreases to the dynamics of inflation during this period. Figure 4.6 plots the marginal impact of pos_t and neg_t to inflation on an annual basis from 2010 to 2015. The graph indicates that there were both large increases and decreases that combined to form aggregate inflation movements. Over this period, as inflation rose, the magnitude of price increases rose and the magnitude of price decreases fell. According to Table 4.3, pos_t is significantly more strongly correlated to inflation at 0.95 compared with neg_t at 0.73. The impact, however, of a 1 per cent rise in pos_t tends to be smaller, leading to a 1.2 per cent increase in monthly inflation compared with a 2.0 per cent increase due to neg_t . Pos_t nevertheless explains more of the variation in inflation according to the R^2 .

The decomposition above of inflation into its intensive and extensive margins provides information on the level of inflation but does not explain its variation. The drivers of the variation in inflation are vital to determining which models of micro-founded pricing behaviour best replicate the properties of price dynamics in SA – time-dependent or state-dependent models.

Figure 4.6: Annual inflation decomposed into price increases and price decreases



4.6.3 Decomposing the variation in inflation

In this section we look at the importance of IM and EM as drivers of the variance of inflation. In order to do this we followed Klenow and Kryvtsov (2008) and used a first-order Taylor expansion of $\pi_t = fr_t \cdot dp_t$ around the sample means \overline{fr} and \overline{dp} such that:

$$var(\pi_t) = var(dp_t) \cdot \overline{fr}^2 + var(fr_t) \cdot \overline{dp}^2 + 2 \cdot \overline{fr} \cdot \overline{dp} \cdot cov(fr_t, dp_t) + O_t \quad (4.8)$$

The higher-order terms are expressed as O_t and are functions of fr_t . The advantage of this variance decomposition is that it allows us to determine whether inflation is driven by the intensive margin (IM term = $var(dp_t) \cdot \overline{fr}^2$), and hence inflation is time-dependent, or the extensive margin (EM term = $var(fr_t) \cdot \overline{dp}^2 + 2 \cdot \overline{fr} \cdot \overline{dp} \cdot cov(fr_t, dp_t) + O_t$), and state-dependent. Over the full sample available, the intensive margin accounts for only 26.5 per cent of the variance in monthly inflation, while 73.5 per cent is due to the extensive margin. When we exclude sales from the sample, the intensive margin accounts for only 21.9 per cent of the variation in inflation. The importance of the extensive margin in explaining inflation variance in SA suggests that prices are state-dependent. This contrasts with the survey results of Govender (2013) who found that 70 per cent of manufacturing firms in SA use time-dependent pricing strategies (of course this was only for manufacturing firms, whereas our sample covers many firms including wholesalers and retailers).

These results also contrast with the results of Klenow and Kryvtsov (2008) who found that the intensive margin explains the majority of inflation variance, between 86 and 113 per cent for different samples, for the United States between 1988 and 2004. Gagnon (2009) found that in Mexico from 1994 to 2004, the extensive margin (or frequency of price changes) explains the majority, close to 60 per cent, of the variance of inflation. This is particularly acute during periods of high inflation, in the specific case of rates above 10-15 per cent, the extensive margin explains over 60 per cent of the variation in inflation. This suggests that countries that have higher inflation rates would tend to be better described by state-dependent models of price rigidities.

How much do price decreases and increases contribute to the variance of inflation? In order to calculate this, we decomposed the variance of inflation such that:

$$var(\pi_t) = var(pos_t) - cov(pos_t, neg_t) + var(neg_t) - cov(pos_t, neg_t) \quad (4.9)$$

where the contribution of positive price changes is $var(pos_t) - cov(pos_t, neg_t)$ and negative price changes is $var(neg_t) - cov(pos_t, neg_t)$. According to this decomposition, the majority of the variance of inflation is explained by positive price movements. Of the variance in overall inflation, pos_t explains 69.7 per cent of it, while neg_t explains 30.3 per cent. This result is confirmed by the regression analysis in Table 4.3, where pos_t has an R^2 of 0.9 compared with neg_t of 0.54. When we exclude sales items from the analysis, the results tend to increase the importance of price increases, with pos_t explaining 75.4 per cent of the inflation variance.


4.7 How important are sales in pricing conduct and flexibility?

A question that remains unanswered in South Africa is the role of sales in pricing conduct and flexibility. This question was impossible to answer from a consumer prices perspective prior to 2006, as the statistical authorities did not provide information on whether an item was on sale or not. From March 2006 onwards, this information is available, but the dataset does suffer from missing data values, making it difficult to determine whether products were indeed on sale. For example, in 2008, of the 1,248,255 individual observations, 410,741 observations or 32.9 per cent did not indicate whether a product was on sale or not. This gradually improved to 18.6 per cent for 2009 and 10 per cent for 2010. By 2011 all observations indicated whether or not a product was on sale. Creamer *et al.* (2012) do provide some evidence for the role of sales in price frequency changes, highlighting that since sales information started being provided, the frequency of price changes has risen. The causal link to sales remains elusive given the short sample period available. This section looks at the role of sales in pricing conduct.

It is important to note that sales usually last for a much shorter period than a month, the period of time we have between observations. Similarly, there are alternative sales strategies differentiating between the size and length of a sale. This discrepancy means that we are likely to understate the true frequency of price changes. It also means that we are likely to overstate the impact of sales on the level of inflation. This is why sales are excluded from the official inflation numbers released by StatsSA.

Figure 4.7 plots the proportion of items classified as sales in the dataset. There are two lines: the proportion including all data where a large subset of items are not classified, and the proportion that include only data points that are classified either 'sale' or 'regular'. The total number of items that are classified under 'sales' according to the 'Price type' code averages between 4.2 and 4.4 per cent from January 2009 to May 2015, depending on the treatment of unclassified items. However, the actual proportion of items on sale has increased gradually over time. In January 2009, the proportion of items on sale was 2 (1.5) per cent based on only those items classified (the entire dataset including items not classified), rising gradually to 4.4 (4) per cent in 2010, before peaking at 6.1 per cent in December 2014.

Figure 4.7: Proportion of sales



The proportion of sales is low compared with Klenow and Kryvtsov (2008) who showed that about 11 per cent of prices in the US BLS data for the CPI was for items on sale. In Norway, Wulfsberg and Ballangrud (2009) showed that sales account for only three per cent of observations in the Norwegian CPI. Dhyne *et al.* (2006) do not report the incidence of sales, but showed that frequency results were not that sensitive to these changes. These findings are of course not directly comparable, as they include both goods and services, whereas we have only goods data. However, analysis shows that goods, especially food, dominate sales items. Including services is likely to lead to a downward shift in the proportion of sales, suggesting that these numbers would be higher if they excluded services. It is also important to highlight the challenge in doing cross-country comparisons. Statistical agencies' modus operandi differ by country in the reporting of sales⁸.

Table 4.4 shows the impact of sales on the frequency of aggregate price changes, price increases and price decreases under the column heading Freq. Next to these in columns labelled Change is the difference between the frequency including and excluding sales, i.e. a positive number indicates a drop in the frequency. The last two columns show the proportion of sales to total sales and the proportion of sales within each category.

The distribution of sales among commodities is dominated by food goods (second last column in Table 4.4), which account for 47.4 per cent of all classified sales in the dataset. Despite this dominance, excluding sales does not decrease the frequency of price changes, nor increase


⁸One way research has resolved this is to introduce filters. Nakamura and Steinsson (2008) use a V-shaped sales indicator to identify products on sale.

the magnitude of price changes, mostly in the food category. The next categories with the most number of sales in order are Furniture and furnishings (8.3 per cent), Miscellaneous goods (7.5 per cent), Household appliances (6.3 per cent), Audiovisual and photographic equipment (5.5 per cent), and Non-alcoholic beverages (4.3 per cent).

Next we look at the proportion of sales within each CPI category. Furniture and furnishings has the highest rate of sales, with goods on sale 18 per cent of the time. This category is also most affected by sales, with the frequency of price changes dropping by 11.5 percentage points, to 14.4 per cent, of which 7 percentage points are due to price decreases. The next few categories also experiencing higher proportions of within-category sales are Household appliance (11.9 per cent), Audiovisual and photographic equipment (9.6 per cent), Household textiles (7.1 per cent), Goods for routine household maintenance (5.7 per cent), and Miscellaneous goods (5.1 per cent). Categories that have little to no sales items include Vehicles (0 per cent), Telephone equipment (0 per cent), Tobacco (0.2 per cent), and Hotel and restaurant goods (0.2 per cent).

Sales impact on the frequency of price changes. In the full sample, the frequency of price changes is 27.8 per cent but declines to 24.5 per cent when sales are excluded. The implied duration of price changes therefore rises from 3.6 months to 4.1 months when excluding sales. The median frequency drops from 12.5 to 10.6 per cent. Using an alternative method to calculate the average duration by taking the inverse of the frequency at the product level and aggregating that back to an overall value as in Dhyne *et al.* (2006), yields an average duration of 7.9 months, compared with 6.5 months when including sales products. The impact of sales on frequency is low compared with US data, where Nakamura and Steinsson (2008) found that excluding sales decreases the frequency of price changes by roughly half. Figure 4.8 shows the evolution of frequencies over time with sales included and excluded.

Figure 4.8: Weighted average frequency of price changes: impact of sales



The conventional way of thinking about the impact of monetary policy is through using models where the frequency of price changes plays a central role, such as in Calvo (1983). If prices are relatively inflexible, then a monetary policy shock has large real effects, and if these prices are flexible, then the real effects of a monetary policy shock are small. Creamer *et al.* (2012) have already showed that the micro-data contradicts the Calvo assumption of price changes and that the frequency of price changes occur more often than what was estimated in models of monetary policy, such as in Steinbach *et al.* (2009). As a consequence, they found that monetary policy should be more aggressive but less persistent.

The incidence and impact of sales on the frequency of price changes may also have important monetary policy implications. Kehoe and Midrigan (2008) show that the treatment of temporary versus permanent price changes can make a big difference to whether prices are considered flexible or inflexible and modelling choices about this. Our contribution is to show that in SA the role of sales, which are temporary price changes, does not substantially change the frequency of price changes in the aggregate. This means that more simple menu-cost models, which do not include an extension, as in Kehoe and Midrigan (2008), can more adequately replicate the properties of pricing conduct in SA compared with countries where sales have a larger impact. We also found that the frequency of price changes among goods has risen compared with the findings in Creamer and Rankin (2008), suggesting that prices may be even more flexible from a modelling perspective. Two issues complicate this analysis. First, services have become more important in the overall CPI basket over time, with a weighting of 50.14 per cent currently compared with 45.8 per cent in 2011. Second, the proportion of sales has been


Table 4.4: Average weighted frequency of price changes by CPI category excluding sales

	Frequency of price changes (%)		Frequency of price increases (%)		Frequency of price decreases (%)		Proportion of sales to total (%)	Proportion of sales in each category (%)
	Freq	Change	Freq	Change	Freq	Change		
	Food	29.3	4.8	18.1	1.9	11.2	2.9	47.4
Non-alcoholic beverages	20.2	5.8	13.2	2.4	7.0	3.4	4.3	4.6
Alcoholic beverages	16.2	1.8	12.4	0.7	3.9	1.1	2.2	2.0
Tobacco	20.1	0.2	17.0	0.0	3.0	0.1	0.1	0.2
Clothing	5.6	1.8	4.0	0.0	1.6	1.8	3.0	1.3
Footwear	5.4	1.6	3.8	0.0	1.6	1.6	1.4	1.4
Maintenance and repair of dwelling	15.4	1.3	10.1	0.4	5.3	0.8	1.0	1.5
Other fuels	24.7	0.4	15.1	0.1	9.7	0.3	0.1	0.6
Furniture and furnishings, carpets and other	14.4	11.5	8.4	4.5	6.0	7.0	8.3	18.0
Household textiles	10.4	6.4	6.2	2.5	4.1	3.9	3.0	7.3
Household appliances	15.9	9.9	9.8	4.0	6.1	5.8	6.3	11.9
Glassware, tableware and household utensils	10.2	3.7	6.1	1.4	4.1	2.3	1.2	3.0
Tools and equipment for house and garden	13.9	3.2	8.7	1.1	5.2	2.0	0.6	2.6
Goods for routine household maintenance	21.5	6.6	13.5	2.8	8.0	3.9	2.7	5.7
Medical products, appliances and equipment	16.4	3.9	11.8	1.7	4.6	2.3	1.3	2.6
Vehicles	39.6	0.0	14.1	0.0	25.5	0.0	0.0	0.0
Operation of vehicles	26.2	1.4	17.2	0.5	9.0	0.9	0.7	1.9
Telephone equipment	25.3	2.5	7.4	0.6	17.9	1.9	0.8	0.0
Audiovisual and photographic equipment	16.0	7.1	7.4	2.6	8.7	4.5	5.5	9.6
Other recreation equipment	14.4	3.4	9.5	1.3	4.9	2.1	2.3	2.9
Newspapers, books and stationery	12.9	2.0	8.6	0.9	4.3	1.1	0.3	1.9
Hotel and restaurant	8.8	0.2	7.0	0.1	1.8	0.1	0.1	0.2
Miscellaneous goods	33.3	6.0	19.1	2.4	14.2	3.6	7.5	5.1

increasing over time (only gradually) but if this continues, the role of temporary price changes may need to be reassessed.

Despite the small proportion of sales in the overall dataset, its impact on the level of inflation seems outsized. Figure 4.9 plots annual inflation weighted using dynamic expenditure weights from StatsSA. It is clear that there is a significant divergence between inflation that does and does not account for the role of sales. Monthly inflation including sales data averages 0.25 per cent compared with 0.4 per cent excluding sales. On an annualised basis, the average difference is 1.8 percentage points from 3.0 to 4.8 per cent.

Figure 4.9: Goods inflation including and excluding sales



Sales also impact on the decomposition of inflation variance. Excluding sales tends to predictably magnify the role of price increases, raising their importance to 75.4 per cent compared with 69.7 per cent in the entire dataset. Similarly, excluding sales lowers the variance of the intensive margin more than that of the extensive margin, increasing the variation explained by the extensive margin to 78.1 per cent compared with 73.5 per cent.

4.8 Sensitivity to outliers

Outliers can occur due to very large price changes or due to mistakes in the dataset. This section looks at the possible impact of outliers on the main results of this chapter. Caution is needed in this section, as the price data is complicated by two regularities that are defining characteristics.

First, there are sales that can substantially change the price of an item. Second, the distribution of price changes has excess kurtosis, which we wish to maintain as a property.

We define outliers as any price change that is greater than four standard deviations away from the mean⁹. This equates to 43,676 observations or 0.88 per cent of the sample. Removing outliers had a small impact on the extensive margin, decreasing the average frequency of price changes by only 0.5 percentage points, to 27.2 per cent. It similarly did not significantly change the incidence of price increases or decreases, taking away an equal amount from the frequency of both. The exclusion of outliers did have a variable impact on the categories of CPI, with the largest impact being on Miscellaneous goods, where the frequency of price changes declined by 4.1 percentage points to 35.3 per cent. The next most affected category was Furniture and furnishings with a decline of 1.3 percentage points; thereafter all categories change by less than 1 percentage point. The categories least affected were Vehicles (0 per cent change), Alcoholic beverages (−0.1 per cent change), and Tobacco (−0.1 per cent change).

Removing outliers also had a small effect on the intensive margin. The average price change (dp_t) was 0.1 percentage points higher, at 0.92 per cent, with both increases and decreases being moderated to leave the aggregate change small. The average price increase was 3.8 per cent, compared with 4.2 per cent, while the average price decrease was −2.4 per cent compared with −3.

Moving to the variance of inflation, excluding outliers increased the variation explained by the extensive margin slightly more to 74.2 per cent compared with 73.5 per cent while the variation explained by price increases rose to 70.8 per cent, less than 1 percentage point higher.

4.9 Policy implications

This chapter has a number of policy implications. First, and most prominently, goods inflation in SA is best represented by state-dependent pricing models, rather than time-dependent models. Generally this means that models used in a policy setting, such as Steinbach *et al.* (2009), that are based on time-dependent pricing rules, may not accurately capture the type of pricing dynamics in South African consumer prices. With the micro-price-level data, we are now able to choose and calibrate models of pricing behaviour that align to these micro-price outcomes. This is left to future research. Second, the prominence of time-dependent models suggest that even if these models may not capture all the price dynamics occurring at the micro-level, they can still be calibrated based on this micro-price data. This is done in Creamer *et al.* (2012) for example. However, an important outcome in this chapter is that the frequency of price changes is not constant over time. Therefore the parameter calibration in Creamer *et al.* (2012) is no longer appropriate as the frequency of price changes has risen. In practice, models can be calibrated more frequently based on the results of the micro-price data of the time. Third, sales can have important impacts on the flexibility of prices in an economy. Hence, the modelling

⁹Results for this section are available on request.

choices surrounding the inclusion of sales can lead to significantly different outcomes for pricing behaviour. We have shown, however, that sales do not have a large impact on the frequency of price changes in the South African context. Fourth, uncovering the granularity of pricing behaviour at a micro-price-level widens the SARBs understanding of pricing dynamics, and hence, its ability to implement policy in future. Finally, the results of this chapter show that prices are sticky or downwardly rigid.

4.10 Conclusion

The underlying dynamics of prices matter for how we choose to model inflation, think about price stickiness, contextualise inflation outcomes, and conduct monetary policy. To these ends this chapter provides a decomposition of inflation into its extensive and intensive margins. We show that the extensive margin in South Africa from 2009 to 2015 averaged 27.8 (median is 12.5) per cent, but this can vary anywhere between 37 (25.6) and 18 (3.1) per cent in any particular month. The magnitude of price changes, or the intensive margin, in a particular month averages 0.83 per cent. Multiplying the extensive and intensive margins meant that monthly inflation averaged 0.25 per cent, or 3.0 per cent annualised. The variance of monthly inflation is explained mainly by the extensive margin, or the fraction of prices changing, which accounts for 73.5 per cent of the variance. This suggests that inflation in South Africa is state-dependent rather than time-dependent. Inflation is also dominated by price increases, which explains 70 per cent of the variation in inflation.

We also found that sales do not substantially change the frequency of price changes but do have an important impact on their level. On average, four per cent of products were on sale. The incidence of sales has risen since 2009, from around two per cent in January 2009 to over six per cent in December 2014, and an average five per cent for the first five months of 2015. Sales were most common in the sub-categories of Furniture and furnishings (18 per cent of products in this category were on sale), Household appliances (11.9 per cent), Audiovisual and photographic equipment (9.6 per cent), Household textiles (7.1 per cent) and were least common in Vehicles (0 per cent), Telephone equipment (0 per cent), Tobacco (0.2 per cent). Despite the relatively small number of sales that occur in South Africa, they remain an important contributor to price decreases, and hence keeping inflation lower. Over the period from 2009 to 2015, goods inflation based on the product level would have been 3.0 per cent, instead of the actual 4.8 per cent when it excluded all sales items.

Chapter 5

Decomposing inflation using micro-price-level data: Sticky-price inflation

Contributing author: Neil Rankin

5.1 Introduction

Some prices are stickier than others. In South Africa (SA), consumer prices on average change every five months, with the most frequent price changes occurring every month and the least frequent occurring every 15 months (Creamer *et al.*, 2012). Firms that change prices less often generally need to take account of the likely path of future inflation when setting these prices, if they want to maximise profits. For example, when an insurance company sets medical aid prices on an annual basis, it needs to decide what it expects inflation to be over that period, to ensure that its price is optimal. In contrast, when petrol prices change on a monthly basis, these changes are driven by contemporaneous developments in the exchange rate or the international price of oil. Therefore, prices that are sticky contain more forward-looking information and can be exploited to uncover inflation expectations and underlying, or core, inflation.

It was not until the advent of the micro-price data work starting with Bils and Klenow (2004) in the United States (US) and Creamer and Rankin (2008) in SA that we were able to determine the frequency of price changes for the entire consumer inflation basket, i.e. to determine the extent of price persistence. With the micro-price data, this chapter decomposes goods inflation into flexible- and sticky-price inflation measures for SA from 2008 to 2015.

The advantage of sticky-price inflation is that it grounds the concept of underlying inflation in the theoretical framework currently used by central banks to make policy decisions, and what is considered optimal policy. According to Goodfriend (2007), monetary policy reached the pre-crisis consensus that core inflation, rather than headline inflation was the best nominal anchor for a central bank. Core inflation is more stable and would serve as a better anchor for

inflation expectations. New Keynesian models such as Clarida *et al.* (2002), Aoki (2001), and Bodenstein *et al.* (2008) show that targeting core (or domestic) inflation rather than headline CPI leads to households maximising their welfare. Walsh (2009) shows more generally that inflation leads to the highest welfare loss in sectors where prices are more sticky (or more persistent), with few welfare costs when relative price shocks dissipate quickly. This means that targeting a measure of underlying inflation that is defined by the persistence of prices, such as a sticky-price inflation measure, is optimal.

The contribution of this chapter is fourfold. First, we use micro-price data in SA to decompose goods inflation into flexible- and sticky-price inflation, and define the properties of each. Second, we show that it is possible to build a measure of underlying inflation based on sticky-prices, which is both theoretically appealing and fits into the type of modelling and policy analysis done at central banks. By doing this we add to the argument put forward in Du Plessis (2014) that core inflation should actively be considered as the best nominal anchor for an emerging market central bank. Third, we use the dimensionality available at the product level to improve two important core inflation measures in the existing literature that are considered to be good alternatives to the more common exclusion-based measure: persistence-weighted core inflation (first developed in SA by Rangasamy, 2009) and trimmed means inflation. Fourth, we compare sticky-price inflation to our candidate core inflation measures to provide an initial analysis of relative historical performance. Future work should focus on extending the sample period and should include services data to achieve a full exposition of the merits of sticky-price inflation.

The chapter proceeds as follows: section 2 contextualises this chapter relative to the current literature on core inflation and to the theory of forward-looking prices; section 3 decomposes goods inflation into sticky- and flexible-price indices; section 4 improves the existing measures of core inflation; section 5 provides an initial analysis of the relative performance of core inflation measures and section 6 concludes.

5.2 Literature review

5.2.1 Core inflation measures in the literature

There are two broad definitions of core inflation expounded in Roger (1998). The first of these is as a ‘persistence’ concept building on earlier work by Friedman *et al.* (1963). Friedman *et al.* (1963:25) highlights two distinct characteristics of inflation “...a steady inflation, one that proceeds at a more or less constant rate, and an intermittent inflation, one that proceeds by fits and starts...”, the former being core inflation. The second of these as a ‘generalised’ concept, defined initially by Eckstein (1981:7) as “...the trend increase of the cost of the factors of production” which “...originates in the long-term expectations of inflation in the minds of households and businesses, in the contractual arrangements which sustain the wage-price momentum, and in

the tax system”. Although the two definitions approach core inflation differently, there are important commonalities. First, both concepts attempt to get to the steady or underlying trend in inflation. Second, core inflation is related to the underlying trends in the behaviour of price-setters, or demand conditions in the economy, rather than supply-side shocks.

The unobservable nature of core inflation has seen exclusion-, model- and statistical-based methods all developed in order to estimate a practical measure of core inflation (see, for example, Cogley, 2002; Cristadoro *et al.*, 2005; Quah and Vahey, 1995; and Bryan and Cecchetti, 1993).

The first and most common approach is an exclusion-based measure that is used today to define core inflation as *excluding food and energy*. It has its origins in the 1970s when the US economy faced volatile shocks to both food – due to significant foreign demand and drought – and energy prices – from restrictions to oil supply introduced by the Organization for Petroleum Exporting Countries (OPEC). Detmeister (2012) provides three characteristics that define an exclusion-based index: the excluded items are predetermined, they do not change often, and the relative weights used are the same as in the overall headline price index. Exclusion-based measures are typically supported by arguments that they are thought to be more easily understood by the general public and can be replicated. The disadvantage of these measures is that they exclude entire components of inflation, that may include vital information regarding the underlying trend of inflation.

The second broad approach to the measurement of core inflation is the statistical approach. One of a range of statistical techniques is used to remove transitory noise from (or smooth) the inflation series (see Blignaut *et al.*, 2009; Rangasamy, 2009; Ruch and Bester, 2013; and Du Plessis *et al.*, 2015). Statistical methods have generated the most work on core inflation using many different techniques, mostly filters. A popular and promising measure is the trimmed means approach by Bryan and Cecchetti (1994). This measure aligns well with inflation as a monetary phenomenon and is likely to represent underlying inflation better, by taking into account the positive skewness in prices. An important disadvantage of the trimmed mean measure from a theoretical perspective is its inability to distinguish between “transient and persistent extreme price movements” (Wynne, 2008). In SA, Blignaut *et al.* (2009) calculate a number of trimmed means measures. The popularity of this type of core inflation measure has meant that StatsSA now includes a trimmed mean, which trims five per cent off each tail at the product group level.

To address the inability of trimmed means to identify persistent price changes, Cutler (2001) introduces a persistence-weighted core inflation measure. The measure links underlying inflation to a ‘persistence’ concept as defined by Friedman *et al.* (1963) and embraces Woodford’s view that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier” (Woodford, 2003:17). Components of inflation are weighted based on their persistence, defined here by an autoregressive coefficient. Rangasamy (2009) implemented a persistence-weighted core inflation measure for SA.

The third approach to the measurement of core inflation involves the use of an economic model based on underlying theory such as in Quah and Vahey (1995) or Cristadoro *et al.* (2005). These approaches add additional information, with economic interactions as well as feedback loops, to inform the path of core inflation. Core inflation is defined in Quah and Vahey (1995:1130) as “that component of measured inflation that has no medium- to long-run impact on real output”, which corresponds to Friedman’s definition of core inflation. Model-based approaches are appealing since the core inflation measure fits into a framework that ensures consistency in analysing economic interactions. However, they do not escape problems of incorrect model specification, identification and uncertainty.

The various alternative core inflation measures suggest that criteria are needed to establish which is ‘best’. Clark (2001) argues that policymakers and analysts have reached consensus on the defining properties of a good measure of core inflation. These include that it must track the components of inflation that persist for several years, help predict future headline inflation over the medium term, be less volatile, and be simple. One important omission from this list is that it must be grounded in the theory used by central banks. The appeal of this theoretical grounding is threefold. First, although many techniques can remove the higher frequency movements in headline inflation, these measures remain atheoretic and can only be judged based on the sample available. Second, aligning core inflation with theory ensures that the right identifying assumptions are used when building a practical measure of core inflation. Third, the normative objective function of the central bank is defined by core inflation in a welfare theoretic framework.

There are methods already developed that fit the theory well, the most successful being model-based definitions such as those of Quah and Vahey (1995) and persistence-based measures such as those of Rangasamy (2009). A flaw of model-based measures such as those of Quah and Vahey (1995) is that defining core inflation is done at a macro level, allowing only a limited set of economic relationships, such as a short-run Phillips curve and money neutrality in the long-run, in a single sector. Prices, however, are empirically strongly heterogeneous, both in the magnitude and frequency of price changes. One possible solution to this is to define core inflation from the perspective of pricing behaviour at a micro-price level, as is done in Bryan and Meyer (2010), Reiff and Várhegyi (2013) and Millard and O’Grady (2012). Sticky-price inflation as defined from a micro-price product level accounts for the heterogeneity that exists, and builds its foundation in a theory of forward-looking prices and optimal monetary policy.

5.2.2 Monetary policy and core inflation as the nominal anchor

Underlying, or core, inflation is a cornerstone of modern monetary policy. It represents the adequate nominal anchor, in addition to an adequate instrument and credibility, to achieve the goal of price stability. According to Goodfriend (2007), monetary policy reached the pre-crisis consensus that core inflation rather than headline inflation was the best nominal anchor for a central bank. Core inflation is more stable and would serve as a better anchor for inflation

expectations. Goodfriend (2007:62) was referring to a conventional definition of core inflation – “inflation that excludes volatile prices of such goods as food and oil”.

Part of reaching consensus on core inflation was the development of the theory that showed that core inflation, rather than headline inflation, led to households maximising their welfare. This ‘consensus’ model – with features that include monopolistically competitive firms who set prices in a staggered way, rational expectations, households maximising utility, and a prominent role for monetary policy – was first expounded in Goodfriend and King (1997) and Clarida *et al.* (1999). The rationale behind not targeting headline inflation is that this would require a response to relative price shocks that unnecessarily compounds output losses, i.e. forces the sticky-price sector to adjust through lower demand and hence decrease prices and wages. Relative price shocks from flexible-price products such as oil can also be large, meaning that the output-inflation trade-off would be costly. This finding is echoed in Aoki (2001) and Bodenstein *et al.* (2008).

The optimality of core inflation is not tied to the New Keynesian paradigm but is a general outcome in welfare economics. Walsh (2009) shows that inflation leads to the highest welfare loss in sectors where prices are more sticky (or more persistent), with few welfare costs when relative price shocks dissipate quickly. Walsh (2009:30) also stated that “[s]ince food and energy prices display little stickiness, responding quickly to shifts in demand and supply, there is a strong case for excluding them from the inflation rate the central bank attempts to control”.

Despite the theoretical appeal of core inflation as the optimal nominal anchor, only Thailand and Norway still target it. The reason for this is that a number of practical arguments for headline inflation, and against core, have been made. These arguments include the welfare foundation of the cost of living index, the supposed communication advantage of headline inflation, its use as a reference rate for wage determination and inflation expectations, the frequency of publication by an independent authority, and the large number of alternative core inflation measures.

The most significant argument endorsing headline inflation comes from its foundation in welfare economics, that a central bank should be concerned with the variable that affects people’s lives. This point is rebutted in Du Plessis (2014), who states that the claim that headline CPI is the ultimate goal variable of a central bank does not take account of the outcomes a central bank can control. As stated often in many different forms (see, for example, Wynne, 1999; Bernanke, 2001; Cecchetti and Wynne, 2003; and Walsh, 2009), a central bank should be concerned with monetary inflation that excludes relative price shocks. Central banks can do nothing about relative price shocks and responding to these shocks is likely to create more volatility (Cecchetti and Wynne, 2003). Similarly, central banks have adopted a theoretic framework (in Clarida *et al.*, 1999) in which to operate that points to the supremacy of core inflation from a welfare maximising perspective.

A practical argument for headline inflation is its supposed communications advantage (see, for example, Svensson, 1999; Mishkin, 2007; and Roger, 2009). This argument states that

headline CPI inflation, which has become a convention when thinking about inflation, is easily understood and accepted by the general public. It is also used in price and wage determination. Du Plessis (2014) argues that it is unlikely that targeting core inflation would undermine the South African Reserve Bank's (SARB's) communication strategy, citing recent academic work by Rossouw and Joubert (2005) and Rossouw and Padayachee (2009) suggesting little evidence of the public's understanding of headline inflation or its acceptance of it as a proxy of inflation. Bernanke (2001:322) argues that the exclusion-based measures do not complicate communication to the public but rather improves it by showing the "public that not every shock that raises prices will lead to a permanent increase in inflation, and that short-term changes in inflation resulting from supply shocks will be treated differently from changes driven by aggregate demand". A more recent strand of literature looking at inflation forecast disagreement shows that significant differences in inflation forecasts are explained by the gap between a conventional definition of core inflation and headline inflation, the relative prices of food and energy (Siklos, 2016). By targeting a core inflation measure, a substantial degree of disagreement in inflation forecasts can be discarded, improving monetary policy implementation.

A criticism of targeting headline inflation is that it is subject to large and volatile relative price shocks from food and energy prices as well as from imported inflation. Interest rates are not able to deal with these relative price movements, opening monetary policy up to the 'blunt tool' argument. This also risks a central bank's credibility if it is unable to communicate clearly the reasons for a breach of the inflation target from supply-side shocks. Food and energy prices explain almost 40 per cent of headline inflation in SA since the beginning of inflation targeting and are generally responsible for breaches of the 3-6 per cent target range. When an economy faces relative price shocks of this magnitude, the desire to anchor inflation expectations to a reference measure such as headline CPI can be dangerous. This can lead to inflation expectations becoming unanchored. Of course, SARB implements its mandate in a fully flexible manner, highlighting core inflation when relative price shocks are significant (Kahn, 2009), in essence targeting core inflation when relative price shocks are present.

An argument against targeting core inflation comes from Walsh (2011) and Rangasamy (2011), who state that excluding food from core inflation misspecifies underlying inflation, prompting higher inflation expectations and slowing policy responses. This outcome is linked to two results. First, that not all food is created equal and, hence, excluding all food is an undesirable property of core inflation. Second, that food inflation plays a significant role in inflation expectations and therefore provides important signals to policymakers. The second point is particularly relevant when second-round effects are present. The dilemma raised by Walsh (2011) and Rangasamy (2011) can be solved by moving away from the conventional exclusion-based definition of core inflation to one that embraces an appropriately theoretical definition. Similarly, if second-round effects are present and well understood, responding to forecasts of core inflation will ensure that the central bank adequately responds to changing inflation expectations.

Goodfriend (2007) and Walsh (2009) used the theoretical argument of core inflation to focus on the common exclusion-based inflation measure that most central banks, including SA, use when dealing with how policy will respond to relative price shocks. Walsh (2009:30) stated that “[s]ince food and energy prices display little stickiness, responding quickly to shifts in demand and supply, there is a strong case for excluding them from the inflation rate the central bank attempts to control”. But highlighting only food and energy with no appreciation for all prices that may “display little stickiness” is too narrow, with little theoretical foundation to be an optimal core inflation measure. Woodford (2003:14) states that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier”. Therefore using persistence defined as the frequency of price changes at the product level can more accurately capture the theoretical argument for why core inflation is a better nominal anchor. This measure also takes account of the heterogeneity that exists at the product level.

5.2.3 Theory of forward-looking prices

To show that forward-looking prices contain information about a firm's inflation expectations, we start with a simple model of consumers and producers maximising utility and profits. Consumers maximise utility subject to a budget constraint, yielding a familiar Dixit-Stiglitz demand function such that:

$$C_t(i) = \frac{P_t(i)^{-\theta}}{P_t} C_t \quad (5.1)$$

where $C_t(i)$ is the demand for product i produced by firm i , $C_t = \int_0^1 (C_t(i)^{1-\theta})^{\frac{1}{\theta-1}}$ is aggregate consumption, $P_t(i)$ is the nominal price of product i , and P_t is the aggregate price level.

Heterogeneous firms operate in a monopolistically competitive market, producing differentiated goods using labour such that:

$$Y_t(i) = A_t(i)L_t(i) \quad (5.2)$$

where $Y_t(i)$ is output of firm i , $A_t(i)$ is exogenous technology and $L_t(i)$ is labour available at a wage rate, w . We assume that log technology evolves as an autoregressive process with mean zero and constant variance. We also assume that the market structure in individual goods markets are unchanging.

Following Karadi and Reiff (2012), the model includes some form of price stickiness, which can take on the form of time dependence as in Calvo (1983) or state dependence (menu-cost-based) as in Golosov and Lucas Jr (2007). The form of price stickiness does not change the intuition of the model's result. Each firm maximises the discounted sum of all future profits, subject to its exogenous technology and wages, with the per-period profit function specified as:

$$\Pi(P_t(i), A_t(i)) = (P_t(i) - \frac{w_t}{A_t(i)})Y_t(i) \quad (5.3)$$

Price stickiness means that the firm's pricing decision becomes dynamic. If the firm changes its prices in period t , whether because it is randomly chosen to change its price or pays the menu cost, then the value of the firm would be

$$V^C(A_t(i)) = \max_{P_t(i)} \{ \Pi(P^*(i), A_t(i)) + \beta E_t V(P^*(i), \sum_{l=1}^{\infty} A_{t+l}(i)) \} \quad (5.4)$$

which is the dynamic optimum¹. Firms that face price stickiness will therefore set the current price to maximise future profits, rather than just focusing on current profits.

There are two components to equation 5.4: the static optimum, the price that maximises current profits

$$P^{*S}(A(i)) = \Pi(P^*(i), A_t(i)) = \frac{\theta}{\theta-1} \frac{w}{A(i)}$$

and the forward-looking optimum, the price that maximises the future value of the firm:

$$P^{*F}(A(i)) = \beta E_t V(P^*(i), \sum_{l=1}^{\infty} A_{t+l}(i))$$

Since the profit function is monotonic, the dynamic optimum will always fall between the static optimum and the forward-looking optimum. Therefore, we can define the degree of forward-lookingness to be:

$$D = \frac{P^*(A(i)) - P^{*S}(A(i))}{P^{*F}(A(i)) - P^{*S}(A(i))} \quad (5.5)$$

where $0 \leq D \leq 1$, since the dynamic optimum will always fall between the static and forward-looking optimum. Note, if prices are fully flexible, then firms can change prices in every period and will set the optimal price to maximise current profits, i.e. $D=0$. If however, prices are sticky, then firms would have to base current pricing decisions on maximising future profits as well. In the extreme case where firms can change prices only once, then $D=1$. This provides the link between the frequency of price changes and the extent of forward-looking information. For more details see Karadi and Reiff (2012) and Reiff and Várhegyi (2013).

5.3 Decomposing inflation

To construct a core inflation measure that adequately addresses the theory of forward-looking prices depends on our ability to distinguish between prices that are sticky and those that are flexible. It was with the advent of the micro-price data work of Bils and Klenow (2004) on CPI in the US and Creamer and Rankin (2008) on SA CPI that categorising consumer products became possible. Combining the frequency of price changes with actual price changes at a product level allows us to censor products based on their degree of forward-looking information, to construct

¹In a menu-cost model, the firm would maximise current profits less the cost imposed when changing prices, usually a function of labour input.

a flexible- and sticky-price index. The sticky-price inflation measure has the potential to define core inflation, which would be optimal for monetary policy, measurable, timely, and defined by the theoretical definition of core inflation. This section discusses the underlying micro-price data, constructs flexible- and sticky-price inflation measures, and looks at the basic properties of these measures.

5.3.1 Micro-price data and calculating frequency of price changes

The micro-dataset used in this study is based on the underlying product data provided by Statistics South Africa (StatsSA) and was used to produce CPI. It covers the period of January 2009 to May 2015, and is an extension of the dataset used by Creamer and Rankin (2008), which included data up to December 2007. Our dataset includes only goods and does not provide any information on services. There are 5,200,466 individual price quotes in the period under review. In order to prepare the dataset for analysis, we included only data with an acceptable status code. This means that prices collected which were indicated as ‘Wrong item collected’, ‘Item available but not comparable’, ‘Extreme values not verified’, ‘Quality adjustment’, and ‘Available shelf price wrongly collected’ were excluded. This left 4,986,454 individual price quotes, a drop of 214,006.

To calculate the frequency of price changes, we created an indicator variable $I_{j,k,t}$ that was equal to 1 if there was a price change and 0 otherwise. Consider a retailer or firm k that sells a variety of a product j at time t . A variety of product refers to a unique brand or type of product, for example, comparing one- and two-ply toilet paper from a number of different brands at a specific retailer or firm. To ensure that we compared price changes of identical products over time, we created a unique identification number for each product, in a specific region, at a specific outlet, for a specific month, and of a specific type. The data was then sorted based on this identification number and a price change was then calculated as:

$$dp_{j,k,t} = (p_{j,k,t} - p_{j,k,t-1}) \cdot 100 \quad (5.6)$$

where $p_{j,k,t}$ is the log price of a specific variety of product j at retailer k in time t . We then applied the indicator variable to the micro-data such that:

$$I_{j,k,t} = \begin{cases} 1 & \text{for } dp_{j,k,t} \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.7)$$

The indicator variable $I_{j,k,t}$ was then aggregated using both the mean and median to the product level i , representing the consumption products collected using the Classification of Individual Consumption by Purpose (COICOP) methodology. In this dataset there were 410 individual products. Therefore, the mean frequency at product level i is:

$$I_{i,t} = \frac{\sum_{j=1}^J \sum_{k=1}^K I_{j,k,t}}{J + K} \quad (5.8)$$

Using the earlier example, this would be the aggregation of all varieties of toilet paper as an individual product.

To get the mean and median of the entire dataset, we applied expenditure weights as calculated by StatsSA at product level i . Since we did not have full coverage such that weights added up to 1, we had to normalise the frequency calculations.

Over the sample period, the weighted mean frequency of monthly price changes is 27.8 per cent. Taking the inverse of this (i.e. $1/0.278$) to get an approximation of duration implies that goods prices changed on average every 3.6 months at the mean. This approximation of duration, however, requires that all prices have the same expected duration, an assumption that is unlikely to hold given the heterogeneity of price changes. Therefore, another method for calculating the average duration is to take the inverse of the frequency at the product level and aggregate that back to an overall value as in Dhyne *et al.* (2006). Using this method, we calculated that the weighted average duration of price changes is 6.5 months. According to the weighted median frequency of monthly price changes, goods prices changed 12.5 per cent of the time or approximately every 8.0 months.

5.3.2 Sticky- and flexible-prices

The by-product frequency calculations provide two important sources of information. First, this is the foundation of decomposing inflation into its sticky and flexible components. Second, the mean and median frequency changes give us a natural starting point at which to censor the product-level data. The goods micro-price data we had included 39.72 percentage points of the 49.86% weight of goods, based on 2013 weights, in the overall CPI. An important exclusion from the goods products was petrol prices, which account for 5.67% of the basket weight. To include this product, we assumed based on the behaviour of the petrol price that it has a frequency of 0.9, i.e. it effectively changes almost every month. Calculating frequency at an aggregated level (once all the individual and unique price quotes all over the country are aggregated cross-sectionally) is generally not possible for products as there is a substantial aggregation bias, for example, a food product that from the micro-price data has a true frequency of price change of 0.18 but has a frequency of price change of 1 when looking at that product aggregated cross-sectionally. This aggregation bias does not hold for petrol prices given the homogeneity of the product and how the price is determined nationwide. The inclusion of petrol increases the coverage of our micro-price data to 45.39% of the overall CPI basket and over 90% of the goods component.

We decomposed inflation into two components, sticky- and flexible-price inflation such that:

$$\begin{aligned}\pi_t &= \pi_t^S + \pi_t^F \\ &= \sum_{i=1}^n \omega_{i,t} \cdot I_i^S \cdot \pi_{i,t} + \sum_{i=1}^n \omega_{i,t} \cdot I_i^F \cdot \pi_{i,t}\end{aligned}\tag{5.9}$$

where π_t is the month-on-month inflation rate; π_t^S is sticky-price inflation; π_t^F is flexible-price inflation; I_t^S is an indicator equal to 1 if $I_i \leq T$ and 0 otherwise, I_t^F is an indicator equal to 1 if $I_i > T$ and 0 otherwise; T is an arbitrary threshold value for the frequency of price changes determining which prices are considered sticky or flexible; $\omega_{i,t}$ are the expenditure weights determined by StatsSA for each sub-index; and $i = 1, \dots, n$ is products, in this case $n = 410$.

We simplified the calculation of inflation above (as in Blignaut *et al.*, 2009) such that it is the weighted sum of inflation rates and not a modified Laspeyres-type index employed by StatsSA. This means that the construction of goods CPI here will differ slightly from what StatsSA reports. See Aron and Muellbauer (2004) for a detailed discussion of the methodology used by StatsSA, and Du Plessis *et al.* (2015) for an update. We also normalised the sticky- and flexible-price component weights ($\omega_{i,t}$) to 1 individually to aid representation.

Determining which products should be considered flexible or sticky required a decision on the threshold value T . We followed Reiff and Várhegyi (2013), Bryan and Meyer (2010) and Millard and O'Grady (2012) and calculated three threshold frequency values including the weighted mean (27.8%) and median (12.5%) of frequency changes, as well as 15%, which Reiff and Várhegyi (2013:7) show ensures that “the extent of forward-lookingness is always more than 60 percent”. We mainly present the results for the weighted mean frequency threshold since this ensures sufficient coverage of products on either side of the decomposition.

Figure 5.1 plots the month-on-month seasonally adjusted and annualised (5.1a), and the year-on-year (5.1b), changes in goods CPI against sticky- and flexible-price inflation based on the weighted mean frequency change of 27.8%. The seasonal adjustment was done on the aggregated sticky- and flexible-price indices using the X-13ARIMA-SEATS Seasonal Adjustment. It is clear from both 5.1a and 5.1b that the division of inflation into these two sub-components generates a more volatile flexible-price inflation series with peaks and troughs significantly higher than both sticky-price and overall goods inflation, and a more persistent sticky-price inflation index that captures underlying inflationary trends. The average duration of price changes in sticky-price inflation is 5.9 months compared with an average duration of 1.8 months for flexible-price inflation. In comparison, the average duration of sticky-price inflation increases to 12 months when the threshold is at the median of 12.5%, while the duration of flexible-price changes rises to 2.5 months.

Figure 5.1: Sticky- and flexible-price goods inflation

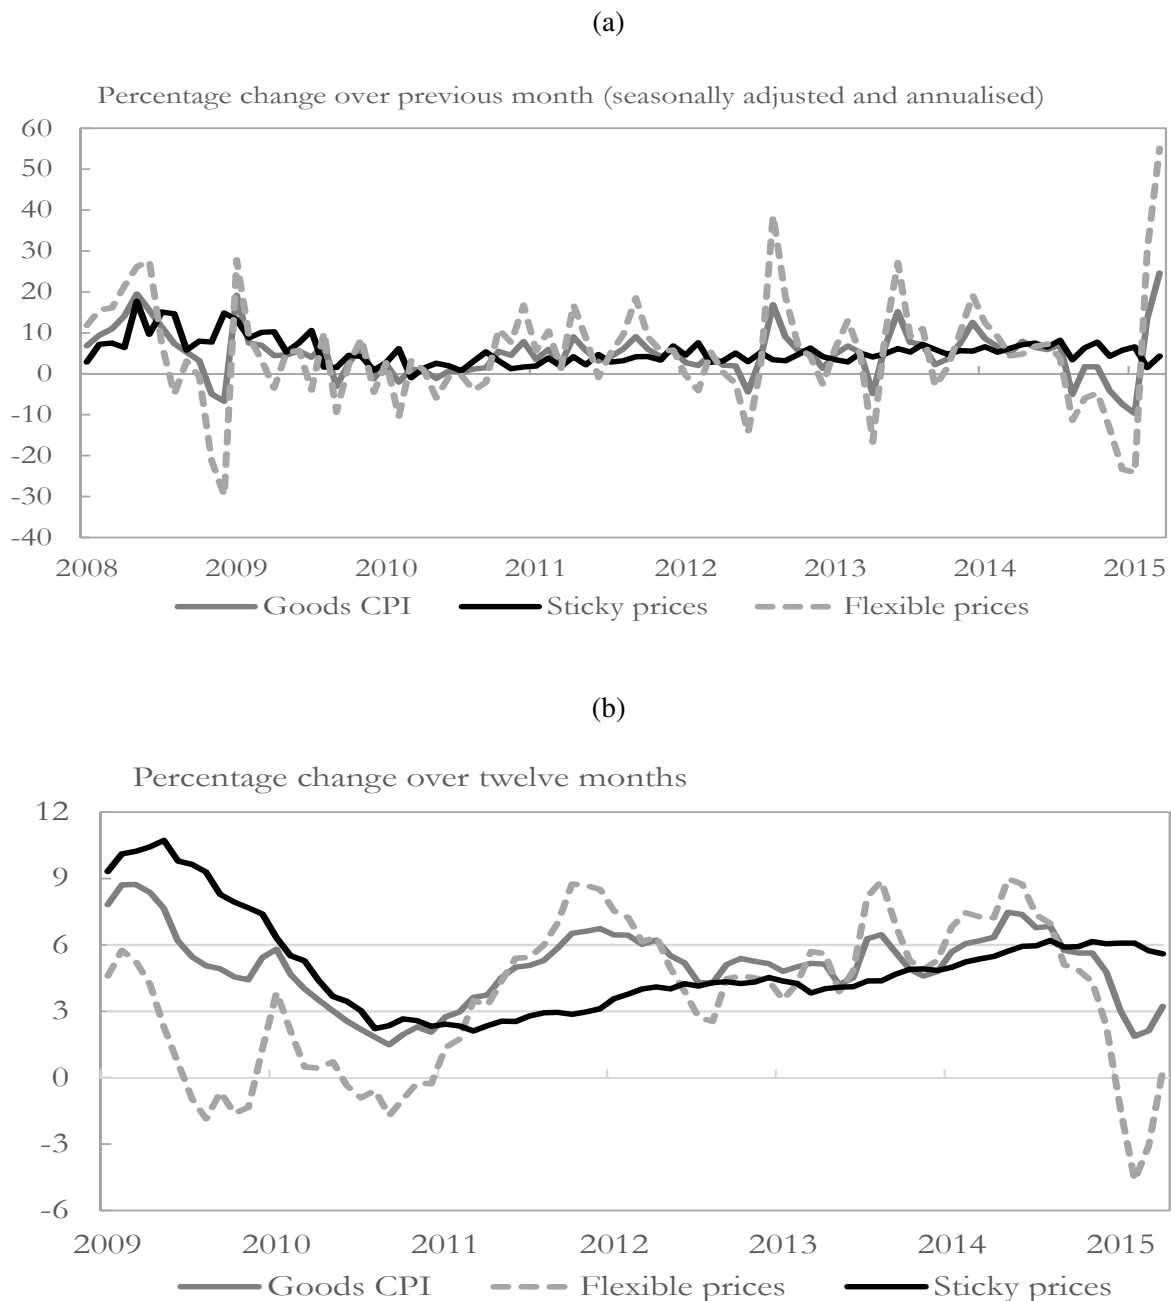


Table 5.1 shows the decomposition of goods inflation, sticky- and flexible-price inflation into its category components based on the 2013 expenditure weights. Sticky-price goods inflation based on the mean weighted frequency of 27.8% has a weight of 24.5 percentage points (in total CPI), made up mainly of New vehicles (5.29%), Alcoholic beverages (3.97%), Clothing (2.76%), Hotel and restaurant goods (2.54%) and Food (1.55%)². Flexible-price inflation, on

²The sticky-price inflation measure based on the median frequency value has a total weight of 7.4%, while the threshold frequency value has a weight of 7.9%.

the other hand, has a total weight of 20.86 percentage points, of which close to 80% comprises food (10.73%) and petrol (5.67%). The next biggest category is miscellaneous goods, which account for 1.47% of flexible-price inflation.

One important finding from constructing the mean weighted sticky-price inflation is the relative importance of certain components of food as well as non-alcoholic beverages, categories usually excluded from a conventional exclusion-based measure. In the mean weighted sticky-price inflation measure, about 10 per cent of food items are classified as core goods, with price changes occurring among these goods every 4.3 months on average. The stickiest food product changes prices only once every 9.3 months. Assuming homogeneity among food products – and excluding them from the common exclusion-based measure used by SARB when defining core inflation – may result in important information on underlying prices being excluded. This can be detrimental for two reasons. First, Anand and Prasad (2010) find that emerging market economies are generally faced with higher food consumption to total consumption as well as low price and income elasticities of food. Therefore, economic agents are likely to factor in food price changes with wage negotiations, affecting inflation expectations. Second, Rangasamy (2011) argues that “[c]ore measures of inflation that exclude food price movements may not accurately reflect the underlying inflationary pressures in the economy and could compromise the attainment of the goal of price stability.” Further validation comes from Walsh (2011), who finds that a core inflation measure that excludes food can misspecify inflation, leading to higher inflation expectations and slow policy responses.

Table 5.2 provides some basic properties of the sticky- and flexible-price inflation measures, based on month-on-month annualised changes, including their standard deviation and persistence (measured by a basic AR(1) equation). We expect flexible-price inflation to contribute most to the volatility in goods inflation, and this is clearly the case at all threshold values. For the weighted mean frequency value of 27.8%, flexible-price inflation has a standard deviation almost four times as large as sticky-price inflation, at 13.1 per cent compared with 3.4 per cent. The standard deviation of flexible-price inflation does decrease with the lower threshold values as the duration of price changes included in the flexible-price inflation measure decline. Flexible-price inflation is also much less persistent than sticky-price inflation, with an autoregressive term of 0.42 compared with 0.88. As the threshold value dividing sticky and flexible prices decreases, flexible-price inflation becomes more persistent, but sticky-price inflation remains relatively unchanged.

CHAPTER 5. DECOMPOSING INFLATION USING MICRO-PRICE-LEVEL DATA:
STICKY-PRICE INFLATION

Table 5.1: Weights of sticky- and flexible-price inflation by category (threshold =27.8%)


	Total goods excluding petrol	Total goods including petrol	Sticky-prices	Flexible-prices
Food	12.28	12.28	1.55	10.73
Non-alcoholic beverages	1.22	1.22	0.64	0.58
Alcoholic beverages	3.97	3.97	3.97	0.00
Tobacco	1.48	1.48	1.48	0.00
Clothing	2.76	2.76	2.76	0.00
Footwear	1.31	1.31	1.31	0.00
Maintenance and repair of dwelling	1.03	1.03	1.03	0.00
Other fuels	0.05	0.05	0.00	0.05
Furniture and furnishings, carpets and other	0.53	0.53	0.20	0.33
Household textiles	0.59	0.59	0.59	0.00
Household appliances	0.60	0.60	0.52	0.08
Glassware, tableware and household utensils	0.10	0.10	0.10	0.00
Tools and equipment for house and garden	0.08	0.08	0.06	0.01
Goods for routine household maintenance	0.54	0.54	0.10	0.44
Medical products, appliances and equipment	0.39	0.39	0.39	0.00
Vehicles	5.99	5.99	5.29	0.70
Operation of vehicles	0.65	6.32	0.21	6.11
Telephone equipment	0.13	0.13	0.00	0.13
Audiovisual and photographic equipment	0.69	0.69	0.45	0.24
Other recreation equipment	0.84	0.84	0.84	0.00
Newspapers, books and stationery	0.03	0.03	0.03	0.00
Hotel and restaurant	2.54	2.54	2.54	0.00
Miscellaneous goods	1.92	1.92	0.45	1.47
Total weight	39.72	45.39	24.53	20.86

Table 5.2: Properties of sticky- and flexible-price inflation

	Goods CPI	Sticky prices	Flexible prices
<i>27.8% Threshold</i>			
Standard deviation	5.89	3.37	13.15
AR-coefficient	0.63	0.88	0.42
<i>15% Threshold</i>			
Standard deviation		2.81	7.29
AR-coefficient		0.87	0.57
<i>12.5% Threshold</i>			
Standard deviation		2.50	6.88
AR-coefficient		0.86	0.57

An important differentiation between sticky- and flexible-price inflation is that sticky prices are set based on firms' inflation expectations in a forward-looking manner. This means that sticky-price inflation should be better correlated with future goods inflation compared with flexible prices. To test this hypothesis, Figure 5.2 looks at the cross-correlogram between goods inflation at $t+1$ to $t+12$ months ahead against contemporaneous sticky-price and flexible-price inflation, at the mean threshold, on a year-on-year basis. It is clear from the graph that sticky-price inflation has a strong positive correlation with goods inflation, with the correlation peaking at $t+3$ at 0.65 and remaining positive. Flexible-price inflation, on the other hand, has a strong contemporaneous correlation with goods inflation at 0.72, but this diminishes quickly and turns negative at $t+5$.

Figure 5.2: Correlation of sticky- and flexible-price inflation to overall inflation




5.3.3 Distributional properties

Censoring the data at specific threshold values gives an indication of broad-based measures of sticky- and flexible-price inflation but lacks granularity over the entire distribution. To look at the properties of annual goods inflation at different threshold values, Figure 5.3 plots 20 per cent rolling windows of threshold values, starting with prices that do not change often, between 0 and 20 per cent frequency of price changes, and rolling forward by 5 percentage points. Hence, the observation at 0.2 is annual inflation extracted from prices, with a frequency between 0 and 20 per cent, and at 0.3 it is between 10 and 30 per cent. Figure 5.3 shows that as product prices become more flexible, their prices also become more volatile. Products that fall between a frequency of 0 and 20 per cent have an annual inflation range of 2.6 to 8.9 per cent, with an interquartile range (IQR) of 2.3 per cent. As product prices become more flexible, these values increase. The inflation rates of products that have frequencies of between 30 and 50 per cent range from -0.7 to 16.3 per cent, with an IQR of 3.6 per cent. The most price-flexible products in our dataset falling between frequencies of 70 and 92 per cent, have inflation rates as low as -27.3 per cent and as high as 13 per cent with an IQR of 7.8 per cent³.

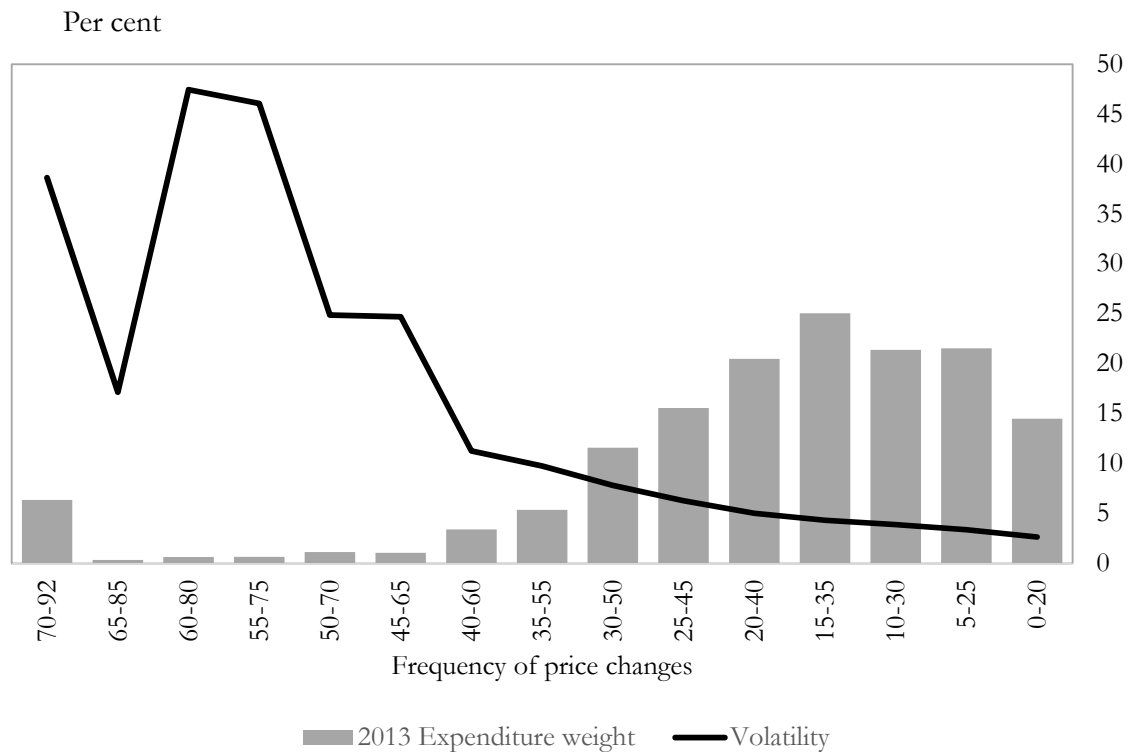
³The largest frequency change among the products studied is 91.7 per cent, so for convenience these are included in the last frequency grouping.

Figure 5.3: Annual goods inflation by frequency groups




An important outcome of more flexible prices is rising volatility. To gauge how this changes over frequencies, Figure 5.4 plots the standard deviation from the groupings in Figure 5.3 as well as the expenditure weights of each group (based on 2013 weights). The standard deviation starts at 2.7 per cent for products with frequencies of between 0 and 20 per cent. This falls between the volatility range of sticky-price inflation at the mean and median thresholds. Volatility slowly rises as products become more flexible, increasing to 5 per cent for products with frequencies of 20 and 40 per cent, 7.8 per cent for 30 and 50 per cent, and ending at 38.6 per cent for products between 70 and 92 per cent. The graph also shows how the expenditure weights of goods products change over the frequency groups. Products with a frequency of 0 and 20 per cent account for 14.5 percentage points of the goods basket; between 20 and 40 per cent, 20.5 percentage points; and 70 to 92 per cent, 6.3 percentage points.

Figure 5.4: Volatility by frequency groups



To get an idea of what sticky-price inflation would look like over the different threshold values, Figure 5.5 plots sticky-price inflation as the threshold value increases from 10 per cent (prices changing every 10 months) to 30 per cent, just above the mean threshold. To determine the bounds for this exercise, we looked at the distribution of frequency of price changes by products, as well as its expenditure weight in the CPI. Based on the 410 products that have frequency information, the minimum frequency of price changes is 3.4 per cent and the maximum is 91.8 per cent. We needed to ensure that more than one product represents sticky-price inflation. At 5 per cent, there are 17 products that make up 4.1 per cent of the products, but only 0.46 percentage points based on expenditure weights. Therefore, we looked for the first frequency change that covers at least ten per cent of total expenditure weights; this occurs at a threshold value of 10 per cent.

Figure 5.5: Sticky-price inflation at different threshold frequencies



It is clear from Figure 5.5 that there is a strong correlation between various sticky-price inflation measures and that different threshold frequencies do not create entirely different outcomes. Inflation tends to rise as the threshold value increases. Sticky-price inflation at a 10 per cent frequency threshold generally creates the lower boundary, while those at a frequency of 25 per cent or lower generally create the upper boundary. The maximum difference between sticky-price measures in Figure 5.5 occurs at the first peak in 2009, at 4.5 percentage points, while the smallest difference occurs in 2014, with a range of 0.7 percentage points.

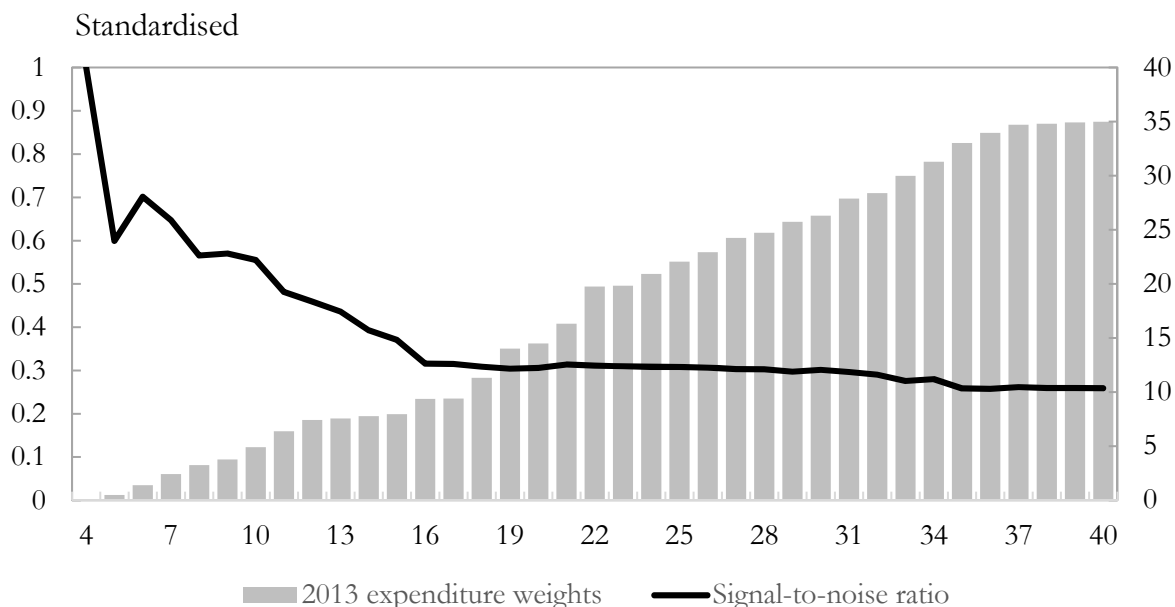
5.3.4 Optimising sticky-price inflation

The choice as to where to censor the data between sticky- and flexible-prices is somewhat arbitrary when thinking about sticky-prices from a core inflation perspective. Sticky-price inflation should represent underlying prices in an economy with maximum forward-looking information but should also have enough products included to ensure an adequate signal of overall goods inflation. This section looks at which threshold frequency provides the best signal of overall goods inflation, what the minimum number of products is at which this is defined and which products represent this signal.

To determine an ‘optimal’ threshold value, we implemented a signal-to-noise ratio test to determine at which point sticky-price inflation represents the optimal signal of overall goods inflation. Examples using the signal-to-noise ratio (SNR) in the literature on core inflation include Mankikar and Paisley (2004), Walsh (2011) and Bullard (2011). Figure 5.6 shows the signal-to-noise ratio normalised to 1 at its minimum, and decreasing as the signal improves, as well as the sum of expenditure weights for frequency values from 4 to 40 per cent. The

graph starts at 4 per cent, since the lowest frequency by product is 3.2 per cent. To determine the optimal threshold, we looked for an inflection point where the least amount of information required maximises the signal-to-noise ratio. This occurs at a frequency of 16 per cent, where the sum of expenditure weights is equal 9.37 percentage points. In the sample of frequencies shown above, the signal-to-noise ratio actually reaches a maximum at a threshold value of 36 per cent; however, the marginal gain in SNR does not justify the higher information requirement.

Figure 5.6: Signal-to-noise ratio by threshold frequencies



The optimal sticky-price inflation measure is close to the 15 per cent threshold suggested by Reiff and Várhegyi (2013:7), which ensures that “the extent of forward-lookingness is always more than 60 percent”. Optimal sticky-price inflation is made up of 169 products and accounts for 9.37 percentage points of the CPI basket. The product categories that account for over 80 per cent of this weight are Clothing (29.5 per cent), Hotel and restaurant goods (27.2 per cent), Footwear (14.1 per cent), and Alcoholic beverages (13.2 per cent).

5.4 Measures of core inflation

In this section we improve two alternative core inflation measures in the existing literature by recognising the heterogeneity in prices and doing the analysis at a product level. The first of these is the persistence-weighted measure introduced for SA by Rangasamy (2009), which has a sufficiently appealing theoretical foundation, as highlighted in Du Plessis (2014). The second of these is the trimmed means inflation measure, first introduced by Bryan and Cecchetti (1994).

5.4.1 Persistence-weighted core inflation

One possible alternative core inflation measure that defines inflation as a ‘persistence’ concept is the persistence-weighted core inflation measure implemented by Cutler (2001) and Rangasamy (2009). Persistence-weighted core inflation aligns with core inflation as defined by Friedman *et al.* (1963) and Woodford’s view that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier” (Woodford, 2003:17). We improved this measure by calculating persistence at the product level, in this case for 410 individual products, rather than at the category level (33 categories) as in Rangasamy (2009). This further disaggregation should provide a richer and more accurate measure of core inflation.

To determine persistence we followed Rangasamy (2009) and defined an autoregressive model:

$$\pi_{it}^m = \sum_{j=1}^p \beta_{ij} \pi_{i,t-j}^m + \varepsilon_{it} \quad (5.10)$$

Where π_{it}^m is demeaned inflation at time t of product $i = 1, \dots, n$, ε_{it} is an error term, and p is the optimal lag length based on the Akaike Information Criterion (AIC) with a maximum lag length of 4. Persistence is then defined as the absolute value of the sum of the autoregressive coefficients such that:

$$P_i = \left| \sum_{j=1}^p \beta_{ij} \right| \quad (5.11)$$

Two weighting schemes were used in order to determine core inflation including weighting based on just persistence (labelled CoreP) and the average of persistence and the expenditure weights provided by StatsSA (labelled CorePC):

$$\begin{aligned} \pi_t^{CoreP} &= \sum_{i=1}^n P_i \cdot \pi_{it} \\ \pi_t^{CorePC} &= \sum_{i=1}^n (\omega_{it} + P_i) / 2 \cdot \pi_{it} \end{aligned} \quad (5.12)$$

CoreP is based on the weighting scheme proposed by Cutler (2001). The problem with this scheme is that it may unduly exaggerate the importance of outliers. One possible way to deal with this problem is to use the persistence measure scaled by a product’s relative importance in the expenditure basket, as proposed by Babetskii *et al.* (2007) (CorePC). Figure 5.8 plots the two persistence-weighted core measures defined in Equation 5.12.

One possible problem with the methodology as applied in Rangasamy (2009) is that the core inflation measures calculated using the persistence weighting scheme did not appear to have lower volatility. This problem was rectified in the application above, as there is more differentiation among the persistence of products that make up the overall CPI.

5.4.2 Trimmed means core inflation

Bryan and Cecchetti (1994:195) first suggested a trimmed means core inflation measure as a solution to “the measurement of aggregate inflation as a monetary phenomenon”. There are two reasons why trimmed means may provide a better representation of underlying inflation. First, Ball and Mankiw (1995:161) show from a theoretical perspective that in a menu cost model “[w]hen price adjustment is costly, firms adjust to large shocks but not to small shocks, and so large shocks have disproportionate effects on the price level. Therefore, aggregate inflation depends on the distribution of relative-price changes: inflation rises when the distribution is skewed to the right, and falls when the distribution is skewed to the left”. This phenomenon means that the distribution of price changes in any particular month will – as is shown in the micro-price literature (see Bils and Klenow (2004) as an example and Creamer *et al.* (2012) for a South Africa-specific result) – be affected by relative price shocks and have excess kurtosis (or fat tails). This motivates the second reason. From a statistical perspective if a population has excess kurtosis, then trimming the distribution will lead to a more efficient estimate of the population mean. An important disadvantage of the trimmed means measure from a theoretical perspective is its inability to distinguish between “transient and persistent extreme price movements” (Wynne, 2008). The popularity of this type of core inflation measure has meant that StatsSA now includes a trimmed means that trims 5 per cent off each tail at the product-group level.

To improve on the trimmed means measures calculated by Blignaut *et al.* (2009), we introduced a product-level trimmed means core inflation measure. The product-level calculation has important advantages. First, we did not introduce any aggregation bias that would occur at the category level. Lafèche (1997) and Bryan *et al.* (1997) highlight that within sub-groups such as food, there may be components that are volatile and others that are persistent. The sticky-price inflation measure above also showed that blindly excluding entire categories does not take account of the type of heterogeneity that exists within categories and across products. The product-level trimmed mean measure therefore embraces the actual heterogeneity that is found in micro-price data and trims from the ‘true’ distribution of price changes, rather than from an aggregated version. Second, unlike as occurs with an exclusion-based measure of core inflation, we did not restrict the subset of products that may experience relative price shocks (e.g. to food and energy).


The trimmed means measures were calculated as follows. The month-on-month changes in each individual product were calculated. These price changes are ranked in ascending order with their associated expenditure weights (ω_{it}). The cross-sectional distribution of price changes was then calculated using the cumulative density function. Price changes that fell into the $t_1\%$ of the lower tail and $(1 - t_2)\%$ of the upper tail were trimmed. The trimmed means inflation rate for each month is then the weighted sum of the left over products using the associated expenditure weights normalised to 1.

This methodology differs slightly from Blignaut *et al.* (2009), as we used month-on-month

inflation rates, rather than year-on-year rates. We chose two trimmed means measures for our analysis, including a symmetric trim of 15 per cent off each tail, as in Bryan and Cecchetti (1994), and an asymmetric trim of 24 and 17 per cent respectively off the top and bottom of the tail, as in Blignaut *et al.* (2009), to account for the positive skewness inherent in the micro-price data. The period under review is too short to determine an optimal trim based on its fit to a trend variable such as the 36-month moving average as in Blignaut *et al.* (2009). Even if this were possible, the desirability to do this remains questionable as these benchmark measures provide only one way to remove relative price shocks and in no way have been proven to be optimal. Our intention in this chapter was to look at the value of defining a sticky-price inflation measure, and therefore, we used common trimmed means measures in the literature as our benchmarks.

To highlight the aggregation bias, Figure 5.7 plots the kernel density estimate of month-on-month price changes at the category level, the 44 categories provided by StatsSA, and at the product level, over 300 products for September 2013. The by-category distribution has a positive skew of 0.76, while the by-product distribution is more positively skewed at 1.56. The by-product distribution also has fatter tails, with kurtosis at 14.5 compared with 3.8 for the by-category distribution.

Figure 5.7: Kernel density estimate of inflation rates at the product and category level

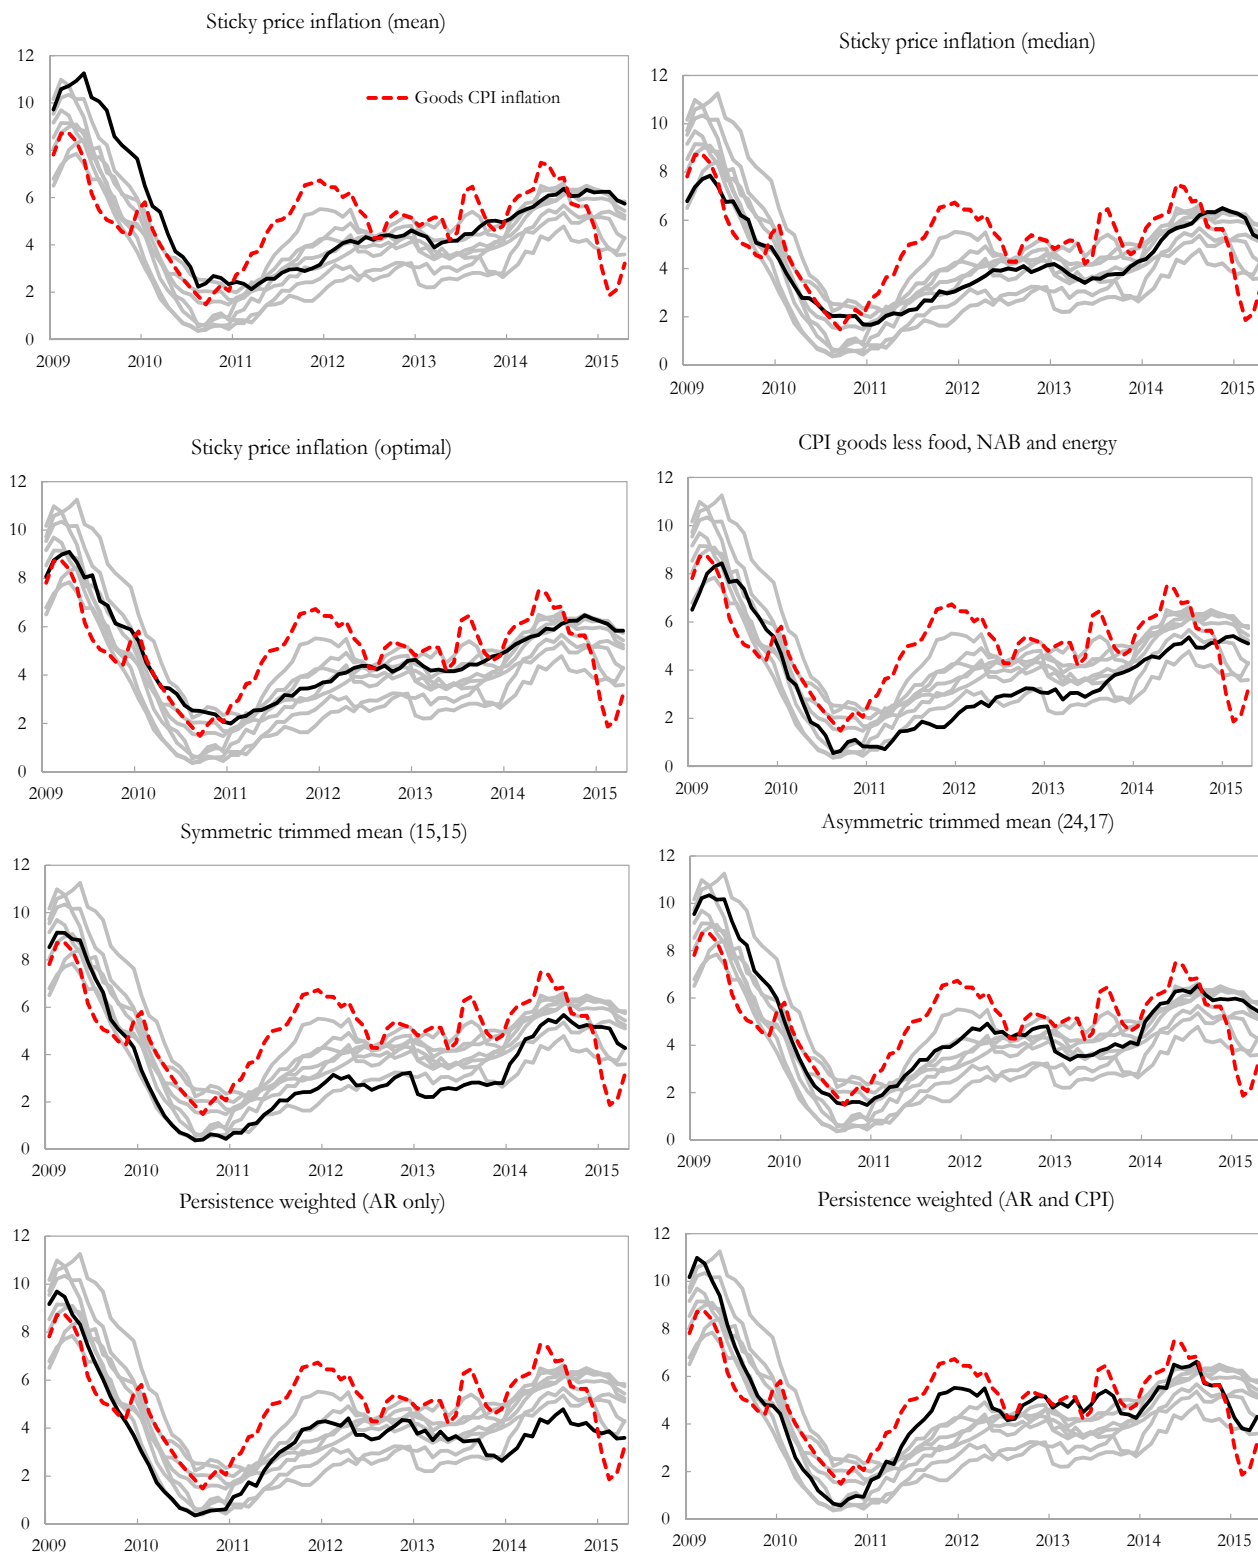


5.5 Comparative performance

In this section we compare our candidate core inflation measures with the properties of a good core measure and against one another. Due to the short sample period, we analysed only in-sample performance, tracking overall inflation with no clear bias and less volatility. Figure 5.8 plots the eight core inflation measures that we analysed including the sticky-price inflation measures censored at the mean (27.8%), median (12.5%), and optimal (16%) frequencies of price change. These three measures were compared with the common exclusion-based measure used by central banks: goods CPI excluding food and non-alcoholic beverages, petrol and energy, which account for 24.6 percentage points of the CPI basket. They were also compared with two versions each of the improved trimmed means and persistence-weighted measures, including a symmetric trimmed means measure with 15% of monthly price changes trimmed off each tail; an asymmetric trimmed means measure where 24% and 17% of the lower and upper tails are trimmed, as in Blignaut *et al.* (2009); a persistence-weighted core inflation where persistence is the sum of the autoregressive coefficients at the product level (AR only); and a

persistence-weighted core inflation measure based on the average of the persistence weights and expenditure weights in the CPI (AR and CPI).

Figure 5.8: Core inflation measures



5.5.1 In-sample performance

For core inflation to be an unbiased predictor of overall inflation, it needs to have the same mean in the long run. In Table 5.3, we compare the mean of month-on-month annualised core inflation measures to overall goods inflation. A t-test allowing for unequal variances was used to compare whether the means were statistically different from overall goods inflation. Goods inflation from 2008M02 to 2015M05 has a mean of 5.39 per cent. All core inflation measures have a lower mean value with a minimum value of 3.83 per cent for the symmetric trimmed means measure. Based on the t-test, the exclusion-based goods measure, CPI goods excluding food, non-alcoholic beverages (NAB) and energy prices, the symmetric trimmed means inflation measure and the persistence-weighted inflation measure based on persistence only (AR only) have statistically different means from overall goods inflation.

Another measure of unbiasedness is the mean error. A value close to zero indicates no clear upward (–) or downward (+) bias. The last column in Table 5.3 shows the mean error for each core inflation measure. The core inflation measure with the lowest mean error is sticky-price inflation (mean) at 0.02 per cent, indicating no clear bias during the sample period. The next lowest mean error is attributed to sticky-price inflation at the optimal threshold, with a value of 0.31 per cent. The worst mean error is the symmetric trimmed mean measure, which has a positive bias of 1.69 per cent, meaning that on average the core inflation measure is 1.69 percentage points below overall goods inflation.

Testing unbiasedness at the mean is appropriate if the underlying distribution of price changes is normally distributed. This assumption is, however, not necessarily true for inflation, which often is positively skewed and has excess kurtosis given the importance of relative price shocks. This property exists at the micro-price level and is likely to exist over time as well. Therefore, for core inflation to be unbiased, it should also have equal medians. To test this we used the Kruskal-Wallis test for equal medians. Table 5.3 shows that the median for overall goods inflation is 5.33 per cent. As was the case with the mean, all core inflation measures also have a lower median value. Based on the test, the three core inflation measures that were found to have unequal means from overall goods inflation also have unequal medians. However, now sticky-price inflation at the median threshold also has a statistically different median from goods inflation.

A significant barrier to determining which core inflation measure is best is the lack of out-of-sample forecastability of headline inflation. The current sample size does not allow an adequate assessment of how well core inflation forecasts future headline inflation over the policy horizon. Two facets of these criteria that have been inappropriately implemented in the literature, however, need to be assessed. First, much of the work done up to now to determine the ability of core inflation to forecast headline inflation has focused on point forecasts using root means squared errors (see, for example, Bryan and Cecchetti, 1994; Clark, 2001; Detmeister, 2012; and Du Plessis *et al.*, 2015). Point forecasts provide little information on core inflation's actual forecast ability since they are too narrow. Instead, these criteria should be based on the ability

Table 5.3: Mean, median and volatility

	Mean	Test for equality of means	Median	Test for equality of medians	Standard deviation	Test for equality of variances	Mean error
Goods inflation	5.39		5.33		5.36		
Sticky price inflation (mean)	5.39	-0.01	4.90	0.17	3.37	2.5***	0.02
Sticky price inflation (median)	4.38	1.6	4.24	3.93**	2.50	4.6***	0.82
Sticky price inflation (optimal)	5.00	0.6	4.65	0.95	2.61	4.2***	0.31
CPI goods less food, NAB and energy	3.84	2.3**	3.54	7.15***	2.97	3.3***	1.36
Symmetric trimmed mean (15,15)	3.83	2.3**	3.71	8.05***	3.34	2.6***	1.69
Asymmetric trimmed mean (24,17)	5.11	0.4	4.86	0.42	3.30	2.6***	0.37
Persistence weighted (AR only)	4.01	2.0**	3.62	6.63***	3.20	2.8***	1.61
Persistence weighted (AR and CPI)	5.04	0.5	4.43	0.75	4.14	1.7***	0.61

*** represents a 1% level of significance, ** a 5% level of significance, and * a 10% level of significance

to forecast the centre of the distribution of headline inflation. One possible way to implement this would be to look at the predictive likelihood, which takes account of the entire predictive density function, over the policy horizon. Predictive likelihoods have the added advantage of providing model selection criteria among many different models and weights for model averaging exercises (Warne *et al.*, 2013). Another option is to find the best forecast of the median of headline inflation over the period 18 to 30 months ahead. Since the distribution of headline inflation is likely to be skewed with excess kurtosis, due to relative price shocks, the median provides a better representation of the distribution. Acknowledging this possibility, Table 5.3 looks at both the mean and median as a test of unbiasedness.

5.6 Caveats and future work

There are three important caveats to this chapter, all requiring future work. First, the micro-price dataset provided covers only about eight years of prices, partly due to an important structural break in the CPI methodology. As a result, determining whether or not sticky-price inflation provides any substantial benefit, beyond the theoretical argued here, above the current subset of core inflation measures remains unanswered. Future work entails combining the micro-price dataset studied in Creamer and Rankin (2008) and Creamer *et al.* (2012) with this dataset. This would involve linking the previous International Trade Classification (ITC) methodology to the new COICOP methodology to classify household expenditure.

Second, the dataset provided by StatsSA included only goods products but no services. This limited the coverage of our study to only 50% of the overall consumer basket. One possible problem that may arise in the services sector is the process by which certain items are surveyed by StatsSA, introducing technical frequency changes that may or may not be appropriate. For example, StatsSA provides a table (Table F - Survey schedule for non-monthly surveys) in its monthly CPI release that indicates which goods and services are surveyed on a quarterly (gymnasium fees, funeral expenses, domestic workers' wages, private-sector hospitals, and taxi fares, for example), biannual (including medical aid, television licences, electricians, plumbers, and municipal charges for utilities), annual (private-sector doctors and dentists, rugby tickets, school and university tuition fees, university boarding fees, and stamps, for example), and ad hoc basis (local bus fares). These are surveyed in this way partly because StatsSA knows the prices are unlikely to change between surveys. StatsSA does state in the CPI publication that "[a]dditional surveys are conducted for these items when Stats SA is aware of significant price changes outside regular survey months".

Only once a sticky-price inflation measure can be calculated over a long enough sample and covers the entire consumer basket can we answer the question of whether this theoretically founded measure provides superior performance to the existing core inflation measures.

The third caveat relates to the market structure under which goods prices are determined. The methodology applied here assumes that the market structure is unchanging. Markets, for

example, that face anti-competitive behaviour, in the form of cartels, will change the forward-looking nature of prices. The impact of changing market structure depends on the speed of this change, and its importance in the overall consumer basket. Future work should look at how a constant market assumption affects the interpretation of forward-looking prices in individual goods markets.

5.7 Conclusion

Using the micro-price data for consumer goods in South Africa from 2008 to 2015, we decomposed goods inflation into its sticky- and flexible price components. Flexible-price inflation is more volatile than overall goods inflation and sticky-price goods inflation, and accounts for the majority of the volatility in overall goods inflation. Sticky-price inflation is more persistent and less volatile than overall goods inflation and the flexible-price inflation measure.

This chapter shows that it is possible to construct a theoretically coherent definition of core inflation based on the concept of price flexibility using micro-price data. We have also improved two existing measures in the literature including trimmed means and persistence-weighted inflation, by calculating these at the product level. In-sample tests show that all core inflation measures have lower volatility than overall goods inflation, but not all have equal means. CPI goods less food, non-alcoholic beverages and energy; the symmetric trimmed means measure and persistence-weighted (AR only) inflation measures have statistically different means from overall goods inflation. All measures except sticky-price inflation (mean) also have a downward bias compared with goods inflation.

Future work should focus on extending the sample period and including services to explore fully the value of sticky-price inflation as a 'good' core inflation measure.

Chapter 6

Summary

Monetary policy is tasked in section 224 (1) of the Constitution of the Republic of South Africa “to protect the value of the currency in the interest of balanced and sustainable economic growth”. Functionally government has set the SARB the requirement to target headline consumer inflation between 3 and 6 per cent, using the repurchase rate as the instrument. Despite headline consumer inflation being the target variable, the SARB relies heavily on signals of underlying inflation – *those that are persistent and exclude relative price movements* – to inform and guide policy. Underlying, or core, inflation is the optimal nominal anchor based on modern economic theory and welfare economics (Goodfriend, 2007). This thesis looked at the role second-round effects have on underlying inflation and wages, the best ways to forecast underlying inflation, and the dynamics and behaviour of inflation from a micro-price perspective.

This thesis contributed to the available literature in the following ways. In chapter 2, the dynamics and size of second-round effects were determined in a coherent theoretical framework using a Structural Bayesian vector autoregressive model with plausible short- and long-run zero and sign restrictions. This chapter showed that second-round effects are present in SA, and require an appropriate monetary policy response to ensure that inflation expectations do not become unanchored.

For a forward-looking central bank that targets consumer price inflation, it is vital to have the most accurate forecasts of underlying inflation. Chapter 3 looked at alternative ways of forecasting core inflation including: models that account for large datasets, that address possible breaks in the inflation series due to changes in monetary policy regimes, that address the changing relationship between variables and inflation or the structure of the economy, and that provide mechanisms to look at the importance of volatility. Most of the literature on core inflation generally looks at finding practical measures rather than forecasting it. This chapter showed that accounting for changing relationships meant better forecasts of core inflation while exploiting more economic information did not necessarily produce better forecasts compared with smaller models.

In chapter 4, goods inflation was decomposed into its extensive margin – *the fraction of prices changing* – and its intensive margin – *the magnitude of price changes* – using micro-price

data. This chapter contributes to the literature in the following ways. First, it studied a dataset of product level data for the goods component of the CPI that was previously unavailable. Second, it extended the analysis of the behaviour of prices in Creamer and Rankin (2008). We looked at the properties of the distribution of frequencies by product over time. Third, by decomposing inflation into its extensive and intensive margins, this chapter provided an entirely new view of goods inflation dynamics. This decomposition showed that goods inflation in SA is likely to be state-dependent, driven by macroeconomic shocks, rather than time-dependent. Fourth, the role of sales in the flexibility of prices in SA is studied. Sales do not affect the overall frequency of price changes but do play an important role in price decreases and keeping inflation at a lower level. The micro-price data analysis done in this chapter provided the starting point for the construction of a theoretically-founded core inflation measure in chapter 5.

Finally, chapter 5 provided a decomposition of goods inflation into its sticky- and flexible-price components. Flexible-price inflation is more volatile than overall goods inflation and sticky-price goods inflation, and accounts for the majority of the volatility in overall goods inflation. Sticky-price inflation is more persistent and less volatile than overall goods inflation and the flexible-price inflation measure. This chapter also went beyond the conventional definition of core inflation, the exclusion-based measure used by most central banks, and argued that exploiting the forward-looking nature of prices, the consensus around core inflation as the best nominal anchor, and the micro-price data, makes sticky-price inflation an advantageous core inflation measure for a central bank.

6.1 Second-round effects from food and energy prices: an SBVAR approach

Second-round effects emanate from the ability of price-setting firms and wage-setting labour to increase prices (by increasing mark-ups or marginal costs) and wages, and therefore prices of other goods and services in response to a relative price shock (to for example food or energy prices). These types of shocks have been important in SA with food and energy contributing an average of 2.4 percentage points (or 39 per cent) to the average 6.1 per cent headline CPI from 2000 to 2014. Second-round effects occur through both the cost and expectations channels. The cost channel refers to the effect that relative price shocks play as an intermediate input in the production of other goods and services. For example, the role of petrol in the transport of goods and services. The expectations channel refers to the impact of relative price shocks on wages. If workers, in the face of a relative price shock, believe the shock to be long-lasting, or if they have the bargaining power to raise wages, then this will cause underlying inflation to rise.

In SA, relative price shocks to food and energy prices are entrenched in the language and responses of wage-setters. The ability of monetary policy to respond to second-round effects requires that policymakers know about the existence and magnitude of these effects. In order to measure second-round effects, we estimated a seven-variable Structural Bayesian VAR (SB-

VAR) with short- and long-run zero restrictions, and sign restrictions, to identify the cost and expectations channel from 1994 to 2014. To simplify the analysis we stuck to a conventional definition of core inflation, used by the SARB in policy communication, as *headline CPI less food and energy*.

The contribution of this chapter is threefold. First, we provided a concise framework through which to estimate the impact of second-round effects from relative price shocks on the economy. This extended the existing work on relative price shocks including Rangasamy (2011) and Rangasamy and Nel (2014), that only estimated the cost channel of second-round effects. However, the most important impact of relative price shocks are their impact on wages, something that has not been quantified in the South African literature. Second, we used recent advances in Bayesian estimation and structural VAR models to estimate an SBVAR with plausible short- and long-run zero restrictions, as well as sign restrictions, to identify the shocks. Third, we entrenched the discussion of second-round effects into the common framework used by the SARB to discuss these effects. This includes using the common exclusion-based core inflation measure of *headline CPI less food and energy prices*.

The results of chapter 2 confirm the impact of wage-setters in South Africa. Changes in the price of food, petrol and energy are accommodated resulting in strong second-round effects. According to the SBVAR model, a one per cent shock to relative food and energy prices increase wages by 0.3 per cent a year after the shock. The price of other goods and services (or core inflation) increase with a maximum impact of 0.3 per cent, three quarters after the shock. This can be attributed to both the cost and expectations channels. The presence of second-round effects change how a central bank needs to respond to relative price shocks. Generally, in the absence of second-round effects, a central bank can look through shocks to food and energy prices as they will be temporary in nature. However, when second-round effects are present, the central bank has to respond appropriately to ensure that inflation expectations remain anchored around the target.

6.2 Forecasting South African core inflation

Like many countries targeting inflation, the SARB relies on forecasts of headline inflation as its operational target. Headline inflation, however, can be volatile, making it difficult to distinguish increases in generalised prices from relative price shocks. This volatility typically arises from a small number of goods and services, most commonly food and energy prices. Petrol prices are a good example of this type of shock, reacting quickly to changes in the international product price and the exchange rate, having the ability to shift headline inflation by a couple of percentage points in months. However, these movements do not reflect the underlying trends in the behaviour of price-setters, or demand conditions in the economy, the issues that matter for a central bank.

Underlying, or core, inflation is likely to be the most important variable for monetary pol-

icy. It is considered to be the optimal nominal anchor as it is stable, excludes relative price shocks, and reflects underlying trends in the behaviour of price-setters and demand conditions in the economy (see, for example, Goodfriend, 2007). In the mid-2000s, monetary economists including Goodfriend (2007) and Woodford (2003) argued that monetary policy had reached a pre-crisis consensus that core inflation rather than headline inflation was the best nominal anchor for monetary policy. Woodford (2003:14) stated that “central banks should target a measure of ‘core’ inflation that places greater weight on those prices that are stickier.” These authors were talking about a conventional definition of core inflation as in “inflation that excludes volatile prices of such goods as food and oil” (Goodfriend, 2007:62). In SA, the SARB generally refers to core inflation as *headline consumer prices less food, non-alcoholic beverages, petrol, and energy*.

Despite its importance there is sparse literature on estimating and forecasting core inflation in SA, with the focus still on finding a practical measure of core inflation. This chapter looked at a host of possible models that existing literature shows to have some success in forecasting, extending the models in four important directions.

First, recent methodological and computing gains have made it possible to increase the dimensionality of models, solving the omitted variable bias in smaller VARs, to include up to a hundred variables when analysing and forecasting macroeconomic variables. Bańbura *et al.* (2010), Giannone *et al.* (2014), and Carriero *et al.* (2015) show that increasing the number of variables used improved forecast accuracy, but that there is a limit to the benefits of size. Bańbura *et al.* (2010) and Koop (2013) provided evidence that this limit is around 20 variables. This chapter considered a number of model sizes up to 21 variables.

Second, a common assumption in simple models of analysis and forecasting is that the errors are homoscedastic. Of course macroeconomic shocks are not. Forecasts from both homoscedastic and heteroscedastic error structures were considered using stochastic volatility first introduced to VARs by Uhlig (1997).

Third, significant changes in the structure of the South African economy over the last four decades make it unlikely that relationships between economic variables remained constant or that there were not any structural breaks. Structural breaks are a common cause of poor forecasting performance (Stock and Watson, 1996; Ang and Bekaert, 2002; Clements and Hendry, 1998; and Bauwens *et al.*, 2011). To address changing relationships, time-varying parameters are introduced. Jooste and Jhaveri (2014) showed that exchange rate pass-through to inflation in SA is time-varying and that it has declined recently. These changing relationships also matter for forecasting. To deal with structural breaks two methods are used. First, discrete breaks are taken into account using methods introduced by Bauwens *et al.* (2011). Second, we used dynamic dimension selection (DDS), as in Koop and Korobilis (2010), allowing for switches between entirely different models to accommodate these breaks.

Fourth, the way large information sets are collated may affect the forecasting accuracy of models. So instead of estimating large VARs it may be the case that factor augmented VARs,

where information is combined into a smaller number of common factors which remove noise, provide better forecasts. Factor models have been shown by Giannone *et al.* (2008) and Kabundi *et al.* (2016) to improve forecast accuracy compared to naive models at short horizons.

The main results of this chapter are that addressing changing dynamics by introducing time-varying parameters generated more accurate forecasts of core inflation. More information did not necessarily improve forecasts as smaller models outperformed larger models. In general, i) time-varying parameter models consistently outperform constant coefficient models (ii) small time-varying parameter vector autoregressive models (TVP-VARs) outperform all other models tested; (iii) models where the errors are heteroscedastic do better than models with homoscedastic errors; (iv) models assuming that the forgetting factor remains 0.99 throughout the forecast period outperform models that allow for the forgetting factors to change with time; and (v) allowing for discrete structural breaks does not improve the predictability of core inflation.

6.3 Decomposing inflation using micro-price level data: South Africa's pricing dynamics

Inflation, or the general increase in prices, is the result of many unobserved adjustments. Only a fraction of prices change in a month. Some of those prices will not have changed for over a year, while others will have changed last month. Some rise and fall faster than others. Some goods are on sale, while others are not. These dynamics matter a lot in themselves, as they describe pricing behaviour. But they also matter for the economic theory forming the foundation of how these prices, and hence inflation, are predicted and forecast.

To address these underlying dynamics, this chapter introduces a decomposition of South African goods inflation into its extensive margin – *the fraction of prices changing in a specific month* – and its intensive margin – *the magnitude of price changes*. These are further decomposed into the magnitude and fraction of prices that are increasing and decreasing. Decompositions of this nature provide economists with the underlying price dynamics needed to both replicate the empirical properties found in consumer prices and to make choices about which models better fit this data. Models of micro-founded pricing dynamics generally fall into two categories: time-dependent or state-dependent, each having significantly different implications for pricing behaviour. Time-dependent models rely on firms setting prices every n^{th} period (as in Taylor, 1999) or randomly (as in Calvo, 1983), while state-dependent models are based on firms that face a cost to change prices and hence only change prices once the change is larger than a 'menu cost' (as in Mankiw, 1985). The role and incidence of sales can also have important implications for the decomposition of inflation, as sales can be a vital source of price flexibility.

The contribution of this chapter is fourfold. First, we analysed a previously not available dataset of product-level data for the goods component of the consumer price index (CPI) from 2009 to 2015. Second, we have extended the analysis of the frequency of price changes by

Creamer *et al.* (2012), looking at the distribution of frequencies by product over time. This analysis reveals that the distribution of frequencies changes significantly over time and by product. Third, we decomposed South African inflation into its extensive and intensive margins to provide a more in-depth analysis of inflation dynamics. Fourth, we looked at the role sales have on price dynamics in the South African economy. The classification of prices as ‘regular’ or ‘sale’ only started to become available in the latter part of the dataset used by Creamer *et al.* (2012), and it was only in 2011 that all observations in the underlying product data was classified.

The results of the decomposition of inflation reveal the following properties in South Africa. First, the average fraction of prices changing, or the extensive margin, is 27.8 per cent, but this can vary anywhere between 37 and 18 per cent in any particular month. This implies that prices change on average every 3.6 months. The median frequency of price changes is 12.5 per cent, implying that prices change every eight months at the median. There is substantial heterogeneity between products and over time, with the distribution having moderate positive skewness and heavy tails (excess kurtosis). Second, the magnitude of price changes, or the intensive margin, averages 0.83 per cent. With 27.8 per cent of prices changing and a magnitude of price changes of 0.83 per cent, monthly goods inflation increases by 0.25 per cent, or 3.0 per cent annualised, from 2009 to 2015, i.e. 0.83×0.278 . Third, the variance of monthly inflation is mainly explained by the extensive margin, or the fraction of prices changing. This suggests that inflation in South Africa is state-dependent, driven by shocks to the economy and changes in input costs, rather than time-dependent. Fourth, the variance in inflation is dominated by price increases, which explains 70 per cent. Fifth, sales prices account for only around four per cent of prices and do not materially change our assessment of the role of the intensive and extensive margins in explaining inflation.

When it comes to the role of sales in the South African consumer goods basket, the results show that on average four per cent of products are on sale. The incidence of sales has risen since 2009, from around two per cent in January 2009 to over six per cent in December 2014, and to an average five per cent for the first five months of 2015. Sales are most common in the sub-categories of Furniture and furnishings (18 per cent of products in this category were on sale), Household appliances (11.9 per cent), Audiovisual and photographic equipment (9.6 per cent), and household textiles (7.1 per cent), while they are least common in Vehicles (0 per cent), Telephone equipment (0 per cent) and Tobacco (0.2 per cent). Despite the relatively small number of sales that occur in South Africa, they remain an important contributor to price decreases and, hence, keeping goods inflation lower. From 2009 to 2015, inflation based on product-level data weighted using expenditure weights would have been 3 per cent instead of the actual 4.8 per cent that it was, excluding all items on sales.

6.4 Decomposing inflation using micro-price data: Sticky-price inflation

Some prices are stickier than others. In South Africa (SA), consumer prices on average change every five months, with the most frequent price changes occurring every month and the least frequent occurring every 15 months (Creamer *et al.*, 2012). Firms that change prices less often generally need to take account of the likely path of future inflation when setting these prices, if they want to maximise profits. For example, when an insurance company sets medical aid prices on an annual basis, it needs to decide what it expects inflation to be over that period, to ensure that its price is optimal. In contrast, when petrol prices change on a monthly basis, these changes are driven by contemporaneous developments in the exchange rate or the international price of oil. Therefore, prices that are sticky contain more forward-looking information and can be exploited to uncover inflation expectations and underlying, or core, inflation.

Using the micro-price data for consumer goods in SA from 2008 to 2015, chapter 5 decomposes goods inflation into its sticky- and flexible-price components. Flexible-price inflation is more volatile than overall goods inflation and sticky-price goods inflation, and accounts for the majority of the volatility in overall goods inflation. Sticky-price inflation is more persistent and less volatile than overall goods inflation and the flexible-price inflation measure.

The contribution of this chapter is fourfold. First, we use micro-price data in SA to decompose goods inflation into flexible- and sticky-price inflation, and define the properties of each. Second, we show that it is possible to build a measure of underlying inflation based on sticky-prices, which is both theoretically appealing and fits into the type of modelling and policy analysis done at central banks. By doing this we add to the argument put forward in Du Plessis (2014) that core inflation should actively be considered as the best nominal anchor for an emerging market central bank. Third, we use the dimensionality available at the product level to improve two important core inflation measures in the existing literature that are considered to be good alternatives to the more common exclusion-based measure: persistence-weighted core inflation (first developed in SA by Rangasamy, 2009) and trimmed means inflation. Fourth, we compare sticky-price inflation to our candidate core inflation measures to provide an initial analysis of relative historical performance. Future work should focus on extending the sample period and should include services data to achieve a full exposition of the merits of sticky-price inflation.

The advantage of sticky-price inflation is that it grounds the concept of underlying inflation in the theoretical framework currently used by central banks to make policy decisions, and what is considered optimal policy. According to Goodfriend (2007), monetary policy reached the pre-crisis consensus that core inflation, rather than headline inflation was the best nominal anchor for a central bank. Core inflation is more stable and would serve as a better anchor for inflation expectations. New Keynesian models such as Clarida *et al.* (2002), Aoki (2001), and Bodenstein *et al.* (2008) show that targeting core (or domestic) inflation rather than headline

CPI leads to households maximising their welfare. Walsh (2009) shows more generally that inflation leads to the highest welfare loss in sectors where prices are more sticky (or more persistent), with few welfare costs when relative price shocks dissipate quickly. This means that targeting a measure of underlying inflation that is defined by the persistence of prices, such as a sticky-price inflation measure, is optimal.

When comparing sticky-price inflation to other core inflation measures, in-sample tests showed that all core inflation measures have lower volatility than overall goods inflation but not all had equal means. The CPI goods less food, non-alcoholic beverage and energy, symmetric trimmed means, and persistence-weighted (AR only) measures, had statistically different means from overall goods inflation. All measures, except sticky-price inflation (mean), also had a downward bias compared to goods inflation. Given the short sample period, and the lack of services data, a full exposition of the merits of sticky-price inflation is left for future research.


Appendices

Appendix A

Stability of the SBVAR model

The estimation methodology involves iteratively solving the model to determine the posterior distributions of the parameters, as well as satisfy the under-identifying restrictions (Zs). To ensure that we did not draw results from an unstable VAR, on each iteration the algorithm checked the maximum eigenvalue of the companion form of the parameter matrix. If this eigenvalue was greater than or equal to 1, the draw was discarded. Figure A.1 shows a histogram of the maximum eigenvalue of the 10,000 draws, with a vertical line indicating the cut-off. About 4 per cent of draws were discarded in the main results of this chapter. The discard rate remained relatively constant to the number of iterations run. The mean posterior estimate of the parameter matrix has a maximum eigenvalue of 0.8423. Hence our VAR is stable.

Figure A.1: Maximum eigenvalue of the parameter matrix on each draw



Appendix B

Forecasting performance with $p = 4$

Table B.1: MSFE relative to the random walk model for core inflation

	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8
Full Model								
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.78	0.61	0.64	0.75	0.64	0.61	0.68	0.78
TVP-VAR-DDS, $\lambda=0.99, \beta_{T+h} \sim RW$	0.77	0.65	0.68	0.78	0.63	0.60	0.72	0.78
TVP-AR								
TVP-AR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.66	0.43	0.36	0.31	0.26	0.31	0.38	0.42
TVP-AR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.67	0.44	0.37	0.32	0.28	0.33	0.40	0.44
TVP-AR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.67	0.43	0.36	0.31	0.26	0.31	0.38	0.42
TVP-AR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.67	0.45	0.38	0.33	0.29	0.33	0.40	0.44
Small VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.67	0.44	0.36	0.31	0.27	0.31	0.37	0.41
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.66	0.43	0.35	0.31	0.27	0.31	0.36	0.42
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.67	0.44	0.36	0.32	0.27	0.31	0.37	0.41
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.66	0.43	0.35	0.30	0.27	0.31	0.37	0.41
VAR, Heteroscedastic	0.65	0.43	0.35	0.30	0.27	0.31	0.36	0.41
VAR, Homoscedastic	0.73	0.56	0.56	0.55	0.56	0.65	0.72	0.83
Medium VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.73	0.41	0.36	0.34	0.27	0.31	0.37	0.43
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.69	0.39	0.35	0.34	0.26	0.30	0.37	0.43
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.76	0.41	0.35	0.35	0.27	0.30	0.37	0.43
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.70	0.39	0.36	0.33	0.26	0.30	0.37	0.43
VAR, Heteroscedastic	0.69	0.39	0.35	0.34	0.26	0.30	0.37	0.44
VAR, Homoscedastic	2.00	1.83	1.76	1.70	1.81	2.02	2.09	2.32
Large VAR								
TVP-VAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.89	0.53	0.49	0.45	0.41	0.46	0.56	0.66
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.78	0.52	0.46	0.43	0.38	0.42	0.49	0.56
TVP-VAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.87	0.55	0.49	0.46	0.41	0.46	0.57	0.67
TVP-VAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.77	0.53	0.46	0.43	0.38	0.42	0.50	0.57
VAR, Heteroscedastic	0.79	0.54	0.48	0.45	0.40	0.45	0.51	0.59
VAR, Homoscedastic	3.04	2.27	1.89	1.72	1.75	1.97	2.02	2.27
FAVAR								
TVP-FAVAR, $\lambda=0.99, \beta_{T+h} = \beta_T$	0.81	0.50	0.41	0.36	0.31	0.37	0.43	0.45
TVP-FAVAR, $\lambda = \lambda_t, \beta_{T+h} = \beta_T$	0.76	0.50	0.41	0.35	0.30	0.36	0.42	0.44
TVP-FAVAR, $\lambda=0.99, \beta_{T+h} \sim RW$	0.82	0.51	0.41	0.36	0.31	0.38	0.43	0.46
TVP-FAVAR, $\lambda = \lambda_t, \beta_{T+h} \sim RW$	0.77	0.50	0.40	0.36	0.30	0.36	0.42	0.45
FAVAR, Heteroscedastic	0.75	0.50	0.41	0.35	0.30	0.37	0.42	0.45
FAVAR, Homoscedastic	1.61	1.47	1.44	1.42	1.34	1.52	1.59	1.78
Benchmark Models								
Random Walk	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Small VAR OLS	1.01	1.05	1.07	1.09	1.11	1.13	1.16	1.18
AR(1) OLS	1.02	1.05	1.08	1.11	1.15	1.18	1.21	1.23
AR(1) Structural Breaks	1.73	1.26	1.02	0.91	0.87	0.89	0.87	0.89
Average performance								
Excluding benchmark models	0.88	0.62	0.55	0.51	0.47	0.53	0.59	0.67
Including benchmark models	0.92	0.67	0.61	0.57	0.53	0.59	0.65	0.71

Appendix C

Representation of dataset

Table C.1 provides an example of what the micro-price dataset looks like. These are fictitious to ensure the confidentiality of the data is maintained. It includes the following information: the outlet a product was collected at, the province and region, a unique commodity code for each product, the price of the product, an item status code describing the status of the product (for example if it is out of season or substituted), an item unit code describing its attribute such as weight, size, or unit number, a commodity sub-code describing brands, and the price type code indicating whether the product is on ‘sale’ or ‘regular’ price.

For example, in the third row of table C.1 we have a box of 38 painkillers collected in the city of Polokwane, in Limpopo province, that is on sale at R21.99 in January 2015. In order to ensure that we compare this specific product over time we create a unique identification (ID) code, which is a function of the outlet, commodity, item unit and commodity sub-codes. So for this product the unique ID would be 139061110111651. The data is then sorted and price changes are calculated.

Table C.1: Micro-price dataset

Outlet	Survey	Province	Region	Commodity		Commodity Code	Item Status Code	Item Status Description	Current Price	Item Code	Item Unit	Item Unit Description	Price Type	Price	Description
				Code	Description										
130	201207	7	10	SHIRT - BUSINESS - MEN	3121004	1	25	Item available	350	651	Each	Each	R		Regular
130	201405	7	10	SHIRT - BUSINESS - MEN	3121004	1	25	Item available	450	651	Each	Each	R		Regular
1390	201501	9	1	PAIN KILLERS	6111001	1	23	Item available	21.99	167	Box Of 38	Box Of 38	S		Sale
1390	201501	9	1	PAIN KILLERS	6111001	2	23	Item available	17.5	165	Box Of 24	Box Of 24	R		Regular
1	201308	6	10	TOOTH BRUSH	12131007	1	20	Item available	13.95	651	Each	Each	R		Regular
59	201103	1	2	SLEEPWEAR - GIRLS	3123023	1	20	Item available	99.99	651	Each	Each	R		Regular

List of References

- Alchian, A.A. and Klein, B. (1973). On a correct measure of inflation. *Journal of Money, Credit and Banking*, vol. 5, no. 1, pp. 173–191.
- Alessi, L., Barigozzi, M. and Capasso, M. (2010). Improved penalization for determining the number of factors in approximate factor models. *Statistics & Probability Letters*, vol. 80, no. 23, pp. 1806–1813.
- Alpanda, S., Kotzé, K. and Woglom, G. (2011). Forecasting performance of an estimated DSGE model for the South African economy. *South African Journal of Economics*, vol. 79, no. 1, pp. 50–67.
- Álvarez, L.J. (2008). What do micro price data tell us on the validity of the New Keynesian Phillips curve? *Economics: The Open-Access, Open-Assessment E-Journal*, vol. 2.
- Álvarez, L.J., Burriel, P. and Hernando, I. (2010). Price-setting behaviour in Spain: evidence from micro PPI data. *Managerial and Decision Economics*, vol. 31, no. 2-3, pp. 105–121.
- Amirault, D., Kwan, C.C. and Wilkinson, G. (2006). *Survey of price-setting behaviour of Canadian companies*. Bank of Canada.
- Anand, R. and Prasad, E.S. (2010). Optimal Price Indices for Targeting Inflation Under Incomplete Markets. *NBER Working Paper Series*.
- Ang, A. and Bekaert, G. (2002). Regime switches in interest rates. *Journal of Business & Economic Statistics*, vol. 20, no. 2, pp. 163–182.
- Aoki, K. (2001). Optimal monetary policy responses to relative-price changes. *Journal of Monetary Economics*, vol. 48, no. 1, pp. 55–80.
- Aron, J., Creamer, K., Muellbauer, J. and Rankin, N.A. (2014a). Exchange rate pass-through to consumer prices in South Africa: Evidence from micro-data. *Journal of Development Studies*, vol. 50, no. 1, pp. 165–185.
- Aron, J., Farrell, G., Muellbauer, J. and Sinclair, P. (2014b). Exchange rate pass-through to import prices, and monetary policy in South Africa. *Journal of Development Studies*, vol. 50, no. 1, pp. 144–164.
- Aron, J. and Muellbauer, J. (2004). Construction of CPIX data for forecasting and modelling in South Africa. *South African Journal of Economics*, vol. 72, no. 5, pp. 884–912.
- Astin, J. (1999). The European Union Harmonised Indices of Consumer Prices (HICP). *Statistical Journal of the United Nations Economic Commission for Europe*, vol. 16, no. 2, 3, pp. 123–135.

- Babetskii, I., Coricelli, F. and Horváth, R. (2007). *Measuring and explaining inflation persistence: disaggregate evidence on the Czech Republic*. Czech National Bank.
- Baharad, E. and Eden, B. (2004). Price rigidity and price dispersion: Evidence from micro data. *Review of Economic Dynamics*, vol. 7, no. 3, pp. 613–641.
- Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, vol. 70, no. 1, pp. 191–221.
- Balcilar, M., Gupta, R. and Jooste, C. (2016). Analysing South Africa's Inflation Persistence Using an ARFIMA Model with Markov-Switching Fractional Differencing Parameter. *The Journal of Developing Areas*, vol. 50, no. 1, pp. 47–57.
- Balcilar, M., Gupta, R. and Kotze, K. (forthcoming). Forecasting South African Macroeconomic Variables with a Markov-Switching Small Open-Economy Dynamic Stochastic General Equilibrium Model. *Empirical Economics*.
- Ball, L. and Mankiw, N.G. (1995). Relative-price changes as aggregate supply shocks. *The Quarterly Journal of Economics*, pp. 161–193.
- Bañbura, M., Giannone, D. and Reichlin, L. (2010). Large Bayesian vector autoregressions. *Journal of Applied Econometrics*, vol. 25, no. 1, pp. 71–92.
- Baumeister, C., Peersman, G. and Van Robays, I. (2010). The economic consequences of oil shocks: differences across countries and time. *Inflation in an era of relative price shocks*, Reserve Bank of Australia, pp. 91–128.
- Bauwens, L., Koop, G., Korobilis, D. and Rombouts, J.V. (2011). A comparison of forecasting procedures for macroeconomic series: the contribution of structural break models. *CIRANO-Scientific Publications*, , no. 2011s-13.
- Bernanke, B. (2001). *Inflation targeting: lessons from the international experience*. Princeton University Press.
- Bernanke, B.S., Boivin, J. and Elias, P. (2005). Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *The Quarterly Journal of Economics*, pp. 387–422.
- Bils, M. and Klenow, P.J. (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy*, vol. 112, no. 5, pp. 947–985.
- Binning, A. (2013). Underidentified SVAR models: A framework for combining short and long-run restrictions with sign-restrictions. *Available at SSRN 2350094*.
- Blanchard, O.J. and Gali, J. (2007). The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s? *NBER Working Paper Series*.
- Blanchard, O.J. and Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *The American Economic Review*, vol. 79, no. 4, pp. 655–673.

- Blignaut, Z., Farrell, G., Munyama, V. and Rangasamy, L. (2009). A note on the trimmed mean measure of core inflation in South Africa. *South African Journal of Economics*, vol. 77, no. 4, pp. 538–552.
- Blinder, A., Canetti, E.R., Lebow, D.E. and Rudd, J.B. (1998). *Asking about prices: a new approach to understanding price stickiness*. Russell Sage Foundation.
- Blinder, A.S. (1991). Why are prices sticky? preliminary results from an interview study. *The American Economic Review*, pp. 89–96.
- Blinder, A.S. (1997). Measuring short-run inflation for central bankers – commentary. *Federal Reserve Bank of St. Louis Review*, , no. May, pp. 157–160.
- Bodenstein, M., Erceg, C.J. and Guerrieri, L. (2008). Optimal monetary policy with distinct core and headline inflation rates. *Journal of Monetary Economics*, vol. 55, pp. S18–S33.
- Botha, D.J.J. (1977). On the definition of inflation: A comment. *South African Journal of Economics*, vol. 45, no. 3, pp. 179–182. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1977.tb00529.x>
- Botha, D.J.J. (1978). On the definition of inflation: A rejoinder. *South African Journal of Economics*, vol. 46, no. 1, pp. 57–59. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1978.tb00550.x>
- Bryan, M.F. and Cecchetti, S.G. (1993). The consumer price index as a measure of inflation. *Economic Review*, , no. Q IV, pp. 15–24.
- Bryan, M.F. and Cecchetti, S.G. (1994). Measuring core inflation. In: *Monetary Policy*, pp. 195–219. The University of Chicago Press.
- Bryan, M.F., Cecchetti, S.G., Wiggins, I. and Rodney, L. (1997). Efficient inflation estimation. *NBER Working Paper*, , no. w6183.
- Bryan, M.F. and Meyer, B. (2010). Are some prices in the CPI more forward looking than others? we think so. *Economic Commentary*, vol. 5, no. 2010, pp. 2010–2.
- Bullard, J. (2011). Measuring inflation: the core is rotten. *Federal Reserve Bank of St. Louis Review*, vol. 93, no. 4, pp. 223–233.
- Calvo, G.A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, vol. 12, no. 3, pp. 383–398.
- Camba-Mendez, G. and Kapetanios, G. (2005). Forecasting euro area inflation using dynamic factor measures of underlying inflation. *Journal of Forecasting*, vol. 24, no. 7, pp. 491–503.
- Campa, J.M. and Goldberg, L.S. (2005). Exchange rate pass-through into import prices. *Review of Economics and Statistics*, vol. 87, no. 4, pp. 679–690.
- Caplin, A.S. and Spulber, D.F. (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics*, vol. 102, no. 4, pp. 703–725.

- Carlton, D.W. (1986). The rigidity of prices. *The American Economic Review*, vol. 76, no. 4, pp. 637–658.
- Carriero, A., Clark, T.E. and Marcellino, M. (2015). Bayesian VARs: specification choices and forecast accuracy. *Journal of Applied Econometrics*, vol. 30, no. 1, pp. 46–73.
- Castellet, I.H. and González, L.J.A. (2004). Price setting behaviour in Spain: stylised facts using consumer price micro data. *Documentos de trabajo del Banco de España*, , no. 22, pp. 1–61.
- Cavallo, A. and Rigobon, R. (2016). The billion prices project: Using online prices for measurement and research. *The Journal of Economic Perspectives*, vol. 30, no. 2, pp. 151–178.
- Cecchetti, S.G. (1986). The frequency of price adjustment: a study of the newsstand prices of magazines. *Journal of Econometrics*, vol. 31, no. 3, pp. 255–274.
- Cecchetti, S.G. and Moessner, R. (2008). Commodity prices and inflation dynamics. *BIS Quarterly Review*, pp. 55–66.
- Cecchetti, S.G. and Wynne, M.A. (2003). Inflation measurement and the ECB’s pursuit of price stability: a first assessment. *Economic Policy*, vol. 18, no. 37, pp. 395–434.
- Chevalier, J.A., Kashyap, A.K. and Rossi, P.E. (2003). Why don’t prices rise during periods of peak demand? evidence from scanner data. *The American Economic Review*, vol. 93, no. 1, pp. 15–37.
- Chib, S. (1998). Estimation and comparison of multiple change-point models. *Journal of Econometrics*, vol. 86, no. 2, pp. 221–241.
- Chisadza, C., Dlamini, J., Gupta, R. and Modise, M. (2013). The Impact of Oil Shocks on the South African Economy. *University of Pretoria, Department of Economics*.
- Christiano, L.J., Eichenbaum, M. and Vigfusson, R. (2006). Alternative procedures for estimating vector autoregressions identified with long-run restrictions. *Journal of the European Economic Association*, vol. 4, no. 2–3, pp. 475–483.
- Clarida, R., Gali, J. and Gertler, M. (1999). The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature*, vol. 37, pp. 1661–1707.
- Clarida, R., Gali, J. and Gertler, M. (2002). A simple framework for international monetary policy analysis. *Journal of Monetary Economics*, vol. 49, no. 5, pp. 879–904.
- Clark, T.E. (2001). Comparing measures of core inflation. *Economic Review—Federal Reserve Bank of Kansas City*, vol. 86, no. 2, p. 5.
- Clements, M. and Hendry, D. (1998). *Forecasting economic time series*. Cambridge University Press.
- Cogley, T. (2002). A simple adaptive measure of core inflation. *Journal of Money, Credit and Banking*, pp. 94–113.

- Costain, J. and Nakov, A. (2011a). Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics*, vol. 58, no. 6, pp. 646–665.
- Costain, J. and Nakov, A. (2011b). Price adjustments in a general model of state-dependent pricing. *Journal of Money, Credit and Banking*, vol. 43, no. 2-3, pp. 385–406.
- Creamer, K., Farrell, G. and Rankin, N. (2012). What price-level data can tell us about pricing conduct in South Africa. *South African Journal of Economics*, vol. 80, no. 4, pp. 490–509.
- Creamer, K. and Rankin, N.A. (2008). Price Setting in South Africa 2001-2007-stylised facts using consumer price micro data. *Journal of Development Perspectives*, vol. 1, no. 4, pp. 93–118.
- Cristadoro, R., Forni, M., Reichlin, L. and Veronese, G. (2005). A core inflation indicator for the euro area. *Journal of Money, Credit and Banking*, pp. 539–560.
- Cutler, J. (2001). Core inflation in the UK. *External MPC Unit Discussion Paper*.
- De Jager, S. (1998). *The framework of the monetary block's equations in the South African Reserve Bank econometric model*. South African Reserve Bank.
- De Waal, A., Van Eyden, R. and Gupta, R. (2015). Do we need a global VAR model to forecast inflation and output in South Africa? *Applied Economics*, vol. 47, no. 25, pp. 2649–2670.
- Detmeister, A.K. (2012). What should core inflation exclude?
- Dhyne, E., Alvarez, L.J., Le Bihan, H., Veronese, G., Dias, D., Hoffmann, J., Jonker, N., Lünnemann, P., Rumler, F. and Vilmunen, J. (2006). Price changes in the euro area and the United States: Some facts from individual consumer price data. *The Journal of Economic Perspectives*, vol. 20, no. 2, pp. 171–192.
- Dias, D.A., Marques, C.R. and Silva, J.S. (2007). Time-or state-dependent price setting rules? evidence from micro data. *European Economic Review*, vol. 51, no. 7, pp. 1589–1613.
- Diewert, W.E. (1983). *The theory of the cost of living index and the measurement of welfare change*. Department of Economics, University of British Columbia.
- Ding, S. and Karlsson, S. (2014). Model averaging and variable selection in VAR models.
- Doan, T., Litterman, R. and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, vol. 3, no. 1, pp. 1–100.
- Dotsey, M., King, R.G. and Wolman, A.L. (1999). State-dependent pricing and the general equilibrium dynamics of money and output. *Quarterly Journal of Economics*, pp. 655–690.
- Dotsey, M., King, R.G. and Wolman, A.L. (2013). Inflation and real activity with firm-level productivity shocks.
- Du Plessis, S. (2014). Targeting core inflation in emerging market economies. In: *Fourteen years of inflation targeting in South Africa and the challenge of a changing mandate*, pp. 53–68. South African Reserve Bank Conference Series.

- Du Plessis, S., du Rand, G. and Kotzé, K. (2015). Measuring core inflation in South Africa. *South African Journal of Economics*, vol. 83, no. 4, pp. 527–548.
- Eckstein, O. (1981). *Core inflation*. Prentice Hall.
- Engle, R.F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, pp. 987–1007.
- Fabiani, S., Druant, M., Hernando, I., Kwapil, C., Landau, B., Loupias, C., Martins, F., Mathä, T., Sabbatini, R. and Stahl, H. (2005). The pricing behaviour of firms in the euro area: New survey evidence. *Banque De France*.
- Fagin, S.L. (1964). Recursive linear regression theory, optimal filter theory, and error analysis of optimal systems. In: *IEEE International Convention Record*, vol. 12, pp. 216–245.
- Fedderke, J.W. and Liu, Y. (2016). Inflation in South Africa: An Assessment of Alternative Inflation Models. *Economic Research Southern Africa*.
- Fofana, I., Chitiga, M. and Mabugu, R. (2009). Oil prices and the South African economy: A macro-meso-micro analysis. *Energy Policy*, vol. 37, no. 12, pp. 5509–5518.
- Forni, M., Giannone, D., Lippi, M. and Reichlin, L. (2009). Opening the black box: Structural factor models with large cross sections. *Econometric Theory*, vol. 25, no. 05, pp. 1319–1347.
- Friedman, M., Friedman, M. and Friedman, M. (1963). *Inflation: causes and consequences*. Asia Publishing House.
- Frisch, H. (1977). Inflation theory 1963-1975: A “second generation” survey. *Journal of Economic Literature*, pp. 1289–1317.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov-switching models*. Springer Science & Business Media.
- Gagnon, E. (2009). Price setting during low and high inflation: Evidence from Mexico. *The Quarterly Journal of Economics*, vol. 124, no. 3, pp. 1221–1263.
- Gali, J. (1992). How well does the IS–LM model fit postwar US data? *The Quarterly Journal of Economics*, pp. 709–738.
- George, E.I., Sun, D. and Ni, S. (2008). Bayesian stochastic search for VAR model restrictions. *Journal of Econometrics*, vol. 142, no. 1, pp. 553–580.
- Gertler, M. and Leahy, J. (2008). A Phillips Curve with an Ss Foundation. *Journal of Political Economy*, vol. 116, no. 3.
- Giannone, D., Lenza, M., Momferatou, D. and Onorante, L. (2014). Short-term inflation projections: A Bayesian vector autoregressive approach. *International Journal of Forecasting*, vol. 30, no. 3, pp. 635–644.

- Giannone, D., Reichlin, L. and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, vol. 55, no. 4, pp. 665–676.
- Golosov, M. and Lucas Jr, R.E. (2007). Menu costs and phillips curves. *Journal of Political Economy*, vol. 115, no. 2.
- Goodfriend, M. (2007). How the world achieved consensus on monetary policy. *The Journal of Economic Perspectives*, vol. 21, no. 4, p. 47.
- Goodfriend, M. and King, R. (1997). The new neoclassical synthesis and the role of monetary policy. In: *NBER Macroeconomics Annual 1997, Volume 12*, pp. 231–296. MIT Press.
- Gordon, R.J. (1975). Alternative responses of policy to external supply shocks. *Brookings Papers on Economic Activity*, vol. 1975, no. 1, pp. 183–206.
- Govender, N. (2013). Price setting behaviour of manufacturing firms in South Africa. *MBA Research Report*.
- Gupta, R., Jooste, C. and Ranjbar, O. (forthcoming). The Changing Dynamics of South Africa's Inflation Persistence: Evidence from a Quantile Regression Framework. Tech. Rep..
- Gupta, R. and Kabundi, A. (2011). A large factor model for forecasting macroeconomic variables in South Africa. *International Journal of Forecasting*, vol. 27, no. 4, pp. 1076–1088.
- Gupta, R., Kanda, P.T., Modise, M.P. and Paccagnini, A. (2015). DSGE model-based forecasting of modelled and nonmodelled inflation variables in South Africa. *Applied Economics*, vol. 47, no. 3, pp. 207–221.
- Gupta, R. and Steinbach, R. (2013). Forecasting Key Macroeconomic Variables of the South African Economy: A Small Open Economy New Keynesian DSGE-VAR Model. *Economic Modelling*, vol. 33, no. 1, pp. 19–33.
- Hall, S., Walsh, M. and Yates, A. (1997). How do UK companies set prices?
- Hallin, M. and Liška, R. (2007). Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association*, vol. 102, no. 478, pp. 603–617.
- Harvey, D.I., Leybourne, S.J. and Taylor, A. (2009). Unit root testing in practice: dealing with uncertainty over the trend and initial condition. *Econometric Theory*, vol. 25, no. 03, pp. 587–636.
- Hazlitt, H. (1965). *What you should know about inflation*. Ludwig von Mises Institute.
- Hlédik, T. and Banka, Č.N. (2003). *Modelling the Second-Round Effects of Supply-Side Shocks on Inflation*. Czech National Bank.
- Jazwinski, A.H. (2007). *Stochastic processes and filtering theory*. Courier Corporation.
- Jevons, W. (1884). *Investigations into currency and finance*.

- Jooste, C. and Jhaveri, Y. (2014). The Determinants of Time-Varying Exchange Rate Pass-Through in South Africa. *South African Journal of Economics*, vol. 82, no. 4, pp. 603–615.
- Kabundi, A., Nel, E. and Ruch, F. (2016). Nowcasting Real GDP growth in South Africa. *SARB Working Paper Series*, , no. WP/16/01.
- Kahn, B. (2009). Challenges of inflation targeting for emerging-market economies: The South African case. *Challenges for Monetary Policy-makers in Emerging Markets*, p. 123.
- Kanda, P.T., Balcilar, M., Bahramian, P. and Gupta, R. (2016). Forecasting South African inflation using non-linear models: a weighted loss-based evaluation. *Applied Economics*, pp. 1–16.
- Kapetanios, G. (2004). A note on modelling core inflation for the UK using a new dynamic factor estimation method and a large disaggregated price index dataset. *Economics Letters*, vol. 85, no. 1, pp. 63–69.
- Karadi, P. and Reiff, A. (2012). Large shocks in menu cost models. *European Central Bank*.
- Kashyap, A.K. (1995). Sticky prices: New evidence from retail catalogs. *The Quarterly Journal of Economics*, vol. 110, no. 1, pp. 245–274.
- Kehoe, P.J. and Midrigan, V. (2008). Temporary price changes and the real effects of monetary policy. *NBER Working Paper Series*, p. 14392.
- Kganyago, L. (March 2015). Statement of the monetary policy committee. *South African Reserve Bank*. Available at: <https://www.resbank.co.za/Lists/News%20and%20Publications/Attachments/6662/MPC%20statement%20March%202015.pdf>
- Kganyago, L. (May 2015). Statement of the monetary policy committee. *South African Reserve Bank*. Available at: <https://www.resbank.co.za/Lists/News%20and%20Publications/Attachments/6729/MPC%20statement%20May%202015.pdf>
- Klenow, P.J. and Kryvtsov, O. (2008). State-dependent or time-dependent pricing: Does it matter for recent US inflation? *The Quarterly Journal of Economics*, vol. 123, no. 3, pp. 863–904.
- Klenow, P.J. and Malin, B.A. (2010). Microeconomic evidence on price-setting. *NBER Working Paper Series*, p. 15826.
- Kohler, A. (1996). Nominale preisrigiditat auf gutermärkten: Eine empirische überprüfung neukeynesianischer erklärungsansätze.
- Konus, A.A. (1939). The problem of the true index of the cost of living. *Econometrica*, vol. 7, no. 1, pp. 10–29. ISSN 00129682, 14680262.
Available at: <http://www.jstor.org/stable/1906997>
- Koop, G. and Korobilis, D. (2010). *Bayesian multivariate time series methods for empirical macroeconomics*. Now Publishers Inc.

- Koop, G. and Korobilis, D. (2012). Forecasting inflation using dynamic model averaging. *International Economic Review*, vol. 53, no. 3, pp. 867–886.
- Koop, G. and Korobilis, D. (2013). Large time-varying parameter VARs. *Journal of Econometrics*, vol. 177, no. 2, pp. 185–198.
- Koop, G. and Potter, S.M. (2007). Estimation and forecasting in models with multiple breaks. *The Review of Economic Studies*, vol. 74, no. 3, pp. 763–789.
- Koop, G.M. (2013). Forecasting with medium and large Bayesian VARs. *Journal of Applied Econometrics*, vol. 28, no. 2, pp. 177–203.
- Lafèche, T. (1997). Mesures du taux d'inflation tendanciel. Tech. Rep., Bank of Canada.
- Lebow, D.E. and Rudd, J.B. (2003). Measurement error in the consumer price index: where do we stand? *Journal of Economic Literature*, vol. 41, no. 1, pp. 159–201.
- Lewin, P. (1977). On the definition and measurement of inflation. *South African Journal of Economics*, vol. 45, no. 3, pp. 175–178. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1977.tb00528.x>
- Lewin, P. (1978a). Indexation and inflation: A note. *South African Journal of Economics*, vol. 46, no. 3, pp. 194–196. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1978.tb00574.x>
- Lewin, P. (1978b). On the Definition of Inflation: Further Reply to Professor Botha. *South African Journal of Economics*, vol. 46, no. 1, pp. 60–61. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1978.tb00551.x>
- Lewin, P. (1978c). On the Definition of Inflation: Reply to Professor Botha. *South African Journal of Economics*, vol. 46, no. 1, pp. 51–56. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1978.tb00549.x>
- Litterman, R.B. (1986). Forecasting with Bayesian vector autoregressions – five years of experience. *Journal of Business & Economic Statistics*, vol. 4, no. 1, pp. 25–38.
- Liu, G., Gupta, R. and Schaling, E. (2009). A New-Keynesian DSGE model for forecasting the South African economy. *Journal of Forecasting*, vol. 28, no. 5, pp. 387–404.
- Mankikar, A. and Paisley, J. (2004). Core inflation: a critical guide. *Bank of England. Quarterly Bulletin*, vol. 44, no. 4, p. 466.
- Mankiw, N.G. (1985). Small menu costs and large business cycles: A macroeconomic model of monopoly. *The Quarterly Journal of Economics*, vol. 100, no. 2, pp. 529–537.
- Marcellino, M., Stock, J. and Watson, M. (2006). A Comparison of Direct and Iterated AR Methods for Forecasting Macroeconomic Series h-Steps Ahead. *Journal of Econometrics*, vol. 135, p. 499.
- Means, G.C. (1935). Industrial prices and their relative inflexibility.

- Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, vol. 79, no. 4, pp. 1139–1180.
- Millard, S. and O’Grady, T. (2012). What do sticky and flexible prices tell us?
- Mishkin, F.S. (2007). *Monetary policy strategy*.
- Mittermaier, K. (1978). On the meaning of “inflationary”: Comment. *South African Journal of Economics*, vol. 46, no. 1, pp. 44–47. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.1978.tb00547.x>
- Mohr, P. (2008). On inflation. *South African Journal of Economics*, vol. 76, no. 1, pp. 1–15. ISSN 1813-6982.
Available at: <http://dx.doi.org/10.1111/j.1813-6982.2008.00158.x>
- Moll, P. (1999). Money, interest rates, income and inflation in South Africa. *South African Journal of Economics*, vol. 67, no. 1, pp. 15–30.
- Morana, C. (2007). A structural common factor approach to core inflation estimation and forecasting. *Applied Economics Letters*, vol. 14, no. 3, pp. 163–169.
- Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, pp. 1415–1464.
- Nakamura, E. and Steinsson, J. (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics*, vol. 125, no. 3, pp. 961–1013.
- Nkomo, J. (2006). Crude oil price movements and their impact on South Africa. *Journal of Energy in Southern Africa*, vol. 17, no. 4, p. 25.
- Nolazco, J.L., Pincheira, P. and Selaive, J. (2016). The evasive predictive ability of core inflation. Tech. Rep..
- Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics*, vol. 92, no. 4, pp. 1004–1016.
- Pesaran, M.H., Pettenuzzo, D. and Timmermann, A. (2006). Forecasting time series subject to multiple structural breaks. *The Review of Economic Studies*, vol. 73, no. 4, pp. 1057–1084.
- Primiceri, G.E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, vol. 72, no. 3, pp. 821–852.
- Quah, D. and Vahey, S.P. (1995). Measuring core inflation. *The Economic Journal*, pp. 1130–1144.
- Raftery, A.E., Gneiting, T., Balabdaoui, F. and Polakowski, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, vol. 133, no. 5, pp. 1155–1174.
- Raftery, A.E., Kárný, M. and Ettler, P. (2010). Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill. *Technometrics*, vol. 52, no. 1, pp. 52–66.

- Ranchhod, S. (2013). Measures of New Zealand core inflation. *Reserve Bank of New Zealand Bulletin*, vol. 76, no. 1.
- Rangasamy, L. (2009). Inflation persistence and core inflation: The case of South Africa. *South African Journal of Economics*, vol. 77, no. 3, pp. 430–444.
- Rangasamy, L. (2011). Food inflation in South Africa: Some implications for economic policy. *South African Journal of Economics*, vol. 79, no. 2, pp. 184–201.
- Rangasamy, L. and Nel, E. (2014). Reconsidering the role of food prices in South African headline inflation. *Agrekon*, vol. 53, no. 4, pp. 16–37.
- Reiff, Á. and Várhegyi, J. (2013). Sticky price inflation index: An alternative core inflation measure.
- Rich, R.W. and Steindel, C. (2007). A comparison of measures of core inflation. *Economic Policy Review*, vol. 13, no. 3.
- RiskMetrics (1996). Technical document. *New York–JP Morgan*.
- Roger, S. (1998). Core inflation: concepts, uses and measurement. *Reserve Bank of New Zealand Discussion Paper*, , no. G98/9.
- Roger, S. (2009). Inflation targeting at 20: Achievements and challenges. *IMF Working Papers*, pp. 1–31.
- Rossouw, J. and Joubert, F. (2005). Supporting an inflation targeting policy with the measurement of inflation credibility. *South African Journal of Economics*, vol. 73, no. 2, pp. 321–336.
- Rossouw, J. and Padayachee, V. (2009). Measuring inflation credibility: Results of a first representative South African sample. *South African Journal of Economics*, vol. 77, no. 2, pp. 314–331.
- Rubio-Ramirez, J.F., Waggoner, D.F. and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, vol. 77, no. 2, pp. 665–696.
- Ruch, F. and Bester, D. (2013). Towards a measure of core inflation using singular spectrum analysis. *South African Journal of Economics*, vol. 81, no. 3, pp. 307–329.
- Şahinöz, S. and Saraçoğlu, B. (2008). Price-setting behaviour in Turkish industries: Evidence from survey data. *The Developing Economies*, vol. 46, no. 4, pp. 363–385.
- SARB (April 2016). Monetary policy review. *South African Reserve Bank*.
Available at: <https://www.resbank.co.za/Lists/News%20and%20Publications/Attachments/7241/MPRApr2016.pdf>
- Schreder, H.X. (1952). Impact of business conditions on investment policies. *The Journal of Finance*, vol. 7, no. 2, pp. 138–173.

- Sheshinski, E. and Weiss, Y. (1977). Inflation and costs of price adjustment. *The Review of Economic Studies*, vol. 44, no. 2, pp. 287–303.
- Siklos, P.L. (2016). Forecast disagreement and the inflation outlook: New international evidence. *Institute for Monetary and Economic Studies, Bank of Japan*.
- Sims, C.A. (1980). Macroeconomics and reality. *Econometrica: Journal of the Econometric Society*, pp. 1–48.
- Smets, F. and Wouters, R. (2002). Openness, imperfect exchange rate pass-through and monetary policy. *Journal of Monetary Economics*, vol. 49, no. 5, pp. 947–981.
- Steinbach, M., Mathuloe, P. and Smit, B. (2009). An open economy New Keynesian DSGE model of the South African economy. *South African Journal of Economics*, vol. 77, no. 2, pp. 207–227.
- Stigler, G.J. and Kindahl, J.K. (1970). *The behavior of industrial prices*, vol. 90. National Bureau of Economic Research New York.
- Stock, J.H. and Watson, M.W. (1996). Evidence on structural instability in macroeconomic time series relations. *Journal of Business & Economic Statistics*, vol. 14, no. 1, pp. 11–30.
- Stock, J.H. and Watson, M.W. (1999). Forecasting inflation. *Journal of Monetary Economics*, vol. 44, no. 2, pp. 293–335.
- Sun, T. (2004). *Forecasting Thailand's core inflation*. 4-90. International Monetary Fund.
- Svensson, L.E. (1999). Price stability as a target for monetary policy: Defining and maintaining price stability. *NBER Working Paper*, , no. w7276.
- Taylor, J.B. (1999). Staggered price and wage setting in macroeconomics. *Handbook of Macroeconomics*, vol. 1, pp. 1009–1050.
- Uhlig, H. (1997). Bayesian vector autoregressions with stochastic volatility. *Econometrica: Journal of the Econometric Society*, pp. 59–73.
- Wakeford, J. (2006). The impact of oil price shocks on the South African macroeconomy: History and prospects. In: *Accelerated and Shared Growth in South Africa: Determinants, Constraints and Opportunities*, (18-20 October 2006). *The Birchwood Hotel and Conference Centre, Johannesburg, South Africa*.
- Walsh, C.E. (2009). Inflation targeting: What have we learned? *International Finance*, vol. 12, no. 2, pp. 195–233.
- Walsh, J.P. (2011). *Reconsidering the role of food prices in inflation*. International Monetary Fund.
- Warne, A., Coenen, G. and Christoffel, K. (2013). Predictive likelihood comparisons with DSGE and DSGE-VAR models. *European Central Bank*.

- Woglom, G. (2005). Forecasting South African Inflation. *South African Journal of Economics*, vol. 73, no. 2, pp. 302–320.
- Woodford, M. (2003). Interest and price.
- Woodford, M. *et al.* (2014). Monetary policy targets after the crisis. *IN, Akerlof, GA, Blanchard, O., Romer, D. and Stiglitz, J.(eds.), What have we Learned*, pp. 55–62.
- Wulfsberg, F. and Ballangrud, A.E. (2009). Price adjustments and inflation: Evidence from consumer price data in Norway 1975–2004. *Norges Bank WP*, vol. 11, p. 2009.
- Wynne, M. (1999). Core inflation: a review of some conceptual issues.
- Wynne, M. (2008). How should central banks define price stability? *Available at SSRN 1734932*.