South African Reserve Bank Working Paper Series WP/23/07

Rationality and biases: insights from disaggregated firm-level inflation expectations data

Monique Reid and Pierre Siklos

Authorised for distribution by Konstantin Makrelov

22 August 2023

SOUTH AFRICAN RESERVE BANK

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

South African Reserve Bank Working Papers are externally refereed.

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

Enquiries relating to the Working Paper Series can be addressed to:
Head: Economic Research Department
South African Reserve Bank
P O Box 427
Pretoria 0001

Tel. +27 12 313 3911

Rationality and biases: insights from disaggregated firm-level inflation expectations data

Monique Reid* and Pierre Siklos†

Abstract

In this paper, we reflect on the controversial concept of 'rational expectations' and point out that the meaning of 'rational' has changed over time. After briefly reviewing the literature, we describe the disaggregated firm-level data from the Bureau for Economic Research in South Africa. Our empirical investigation focuses on inflation expectations. The firm-level data are unique in breadth, scope and time span. We focus on these data, which are considered to be more representative of price setters than financial analysts or households. We compare these results with analysis of financial analysts, who have traditionally been the subject of these types of analyses and who tend to be relatively more rational. We find that while neither the inflation expectations of the business sector nor those of the financial analysts are rational in the strict sense of the term, both respond quickly to changes in underlying macroeconomic and financial conditions. There is evidence that inflation forecast errors stem partly from errors made in forecasting variables such as wages, the exchange rate and interest rates. We reach this conclusion because the survey data used also require respondents to forecast related economic variables. In addition, we find that important socio-economic factors such as firm size, the position of the respondent, and the industry a firm belongs to have significant effects on inflation expectations.

JEL classification

E31, E37, E58, E66

Key words

Inflation expectations, rational expectations, survey, disaggregated data, firms

---

* South African Reserve Bank (SARB) and Department of Economics, University of Stellenbosch, South Africa. Corresponding author, email: monique.reid@resbank.co.za.

† Wilfrid Laurier University, Balsillie School of International Affairs, Waterloo, Canada, and Research Fellow, University of Stellenbosch, South Africa and the SARB.

1. Introduction

Monetary economists have long recognised that they make decisions in a dynamic setting and that the success of monetary policy depends not only on their own behaviour, but also on that of the public. The public are thinking agents[1] rather than passive recipients of policy. In this sense, expectations matter and inflation is a strategic outcome. This widely accepted position was initiated by influential developments in theory (Friedman (1968), Phelps (1967), Kydland and Prescott (1977), Lucas (1976) and others) and reinforced by historical experiences such as stagflation during the Great Inflation of the 1970s.

Policymakers since the 1970s/1980s were motivated to increase their focus on expectations[2] for two broad reasons. Firstly, policymakers aimed to manage the expectations of the public to strengthen support for monetary policy institutions. Central banks should serve society and, within a democratic system, they are ultimately accountable to the public. This institutional credibility supports the second motivation, which is to improve the effectiveness of monetary policy. Even before the global financial crisis (GFC), there was growing agreement that the expectations channel of the monetary transmission mechanism was central to monetary policy (Woodford 2005). However, the crisis increased this focus on the role of expectations because of the need for unconventional monetary policy tools in the face of the ineffectiveness of conventional monetary policy at the zero lower bound.

---

1 We avoid the term 'rational' here deliberately, to delay to later in this paper the discussion of what rationality and subsequent deviations from full rationality mean. The point here is just that we do expect the public to respond to some extent (monetary policy is not a game against nature).

2 It is important to note that the idea that inflation expectations drive inflation has been challenged more recently by Rudd (2021), on both theoretical and empirical grounds. A large part of this paper focuses on the fact that some theory central to capturing the impact of inflation expectations on inflation and the economy (such as the New Keynesian Phillips curve) are not well supported by the data. While there is a wealth of other recent research that also finds that our current theories and models are far from adequate (Coibion, Gorodnichenko and Kumar (2018), De Grauwe and Ji (2019) and others), most of these authors do not conclude that expectations are not important to determining the dynamics of inflation. Furthermore, Rudd (2021) emphasises that these theories tend to rely on short-term expectations, rather than long-term expectations, which he argues receive more attention from policymakers and for which he concedes (without enthusiasm) that there is more empirical evidence. While Rudd adds to the complaint by Tarullo (2017) that policymakers are forced to make policy decisions without a "working theory of inflation", the author has found it difficult to convince policymakers (such as Bernanke (2015)) that the expectations channel of the transmission mechanism is not important.

While the motivations to incorporate expectations in economic analysis have been widely accepted for decades, in practice this required explicit theory about expectations and decisions about how to model them. These details have proven to be challenging. The rational expectations revolution, which has had an enormous effect on macroeconomic modelling since the 1970s, was precisely focused on incorporating expectations in a way that was consistent with the theory or model of the decision-maker (model-consistent rather than ad hoc). However, among the criticisms directed at mainstream macroeconomic models following the GFC, one of the fiercest was of the conventional use of the rational expectations assumption by economic modelers.

In addition, the very active empirical literature has found overwhelming evidence from micro-level expectations survey data across the world that expectations deviate from rational expectations “in systematic and quantitatively important ways including forecast-error predictability and bias” (Coibion, Gorodnichenko and Kumar 2018: 1).

Despite these serious reasons to question the way that modelers were modelling expectations, the least sensible response would be to ignore the role of expectations. The fact remains that monetary policy is made in a dynamic setting and operates with long and variable lags (Friedman 1968), which means that expectations matter in some form. Even if you drop the strict assumption of rational expectations, as long as we recognise that the public responds as thinking agents (monetary policy is not a game against nature), then monetary policy is strategic in nature. Policymakers must take a position on the likely future path of the economy to implement monetary policy and rely on models to offer some consistency (as one of the inputs into their decision-making process). Svensson (1997) argued that inflation targeting should rather be called inflation forecast targeting to reflect its inherently forward-looking nature.

However, the need to continue to implement policy in a forward-looking manner is not equivalent to blind support for the traditional ways of modelling expectations or slavish reliance on models in the implementation of policy. Tarullo (2017) captures the significance of this vulnerability for policymaking by lamenting the lack of a “working theory of inflation”. Policymakers need to have some sense of the implications of the occasionally pragmatic decisions about how to define and model expectations made by theorists and modelers, for two main reasons. These decisions impact macroeconomic models that inform policy (i.e. they affect the accuracy of the models), and they are often used to structure reasoning about how responsive the public will be to policy actions and communication.

In summary, expectations matter and monetary policy is implemented in a strategic setting. Therefore, researchers need to be explicit about how they incorporate expectations into theory and economic models. This task has been difficult in practice and was a major source of criticism of macroeconomics after the GFC. While substantial resources are being devoted to the task in academia, the current state of our understanding of how expectations are formed and how to best model these at an aggregate level often leaves policymakers uncertain. The policy implications are not always trivial under the policy frameworks currently adopted by central banks across much of the developed and less developed world.

In this paper, we reflect on the controversial concept of 'rational expectations' and point out that the meaning of 'rational' has changed over time. After briefly reviewing the literature, we introduce and describe the disaggregated firm-level data from the Bureau for Economic Research (BER) in South Africa. We continually compare the firm-level results with those of the financial analysts (FAs), because traditionally research has focused on the views of analysts or professional forecasters, who tend to be relatively more informed. The analysis of the FAs therefore provides a benchmark against which to understand the firm-level results.

Our empirical investigation focuses on inflation expectations, rather than expectations of other macroeconomic variables. The BER firm-level data are unique in breadth, scope and time span. Because the survey is at the firm level, the data are viewed as being theoretically more representative of price setters than FAs or households. We conclude with a summary.

Although in this paper we do not explore in depth the contribution of the financial economics literature to our understanding of expectations formation, we note that it does offer insights about potential reasons for bias in FA forecasts.[3] A detailed analysis

---

[3] FAs may have strategic reasons for the biased forecasts (Ashiya 2009), rather than being

of this literature would lengthen this paper and, given that the primary focus is on the firm-level data, we refer interested readers to, for example, Bordalo, Gennaioli and Shleifer (2022) for a review of the literature. What we do note is that these incentives do not all affect bias in the same direction, making it difficult to draw conclusions about reasons for the final level of bias observed. Since non-financial firms are not in the business of forecasting inflation or other macroeconomic variables, firms are less likely to be subject to the kind of incentives facing FAs as captured in the financial economics literature. In this paper, we analyse whether the forecast bias of firms differs across firm size (perhaps signalling greater access to professional forecasts) or according to the industry that firms belong to. There are likely to be other reasons for the observed bias, but due to data limitations we restrict ourselves to analyses of these two dimensions.

Briefly, we conclude that while inflation expectations of the business sector and from FAs are not rational in the strict sense of the term, both respond fairly quickly to changes in underlying macroeconomic and financial conditions. Indeed, there is evidence that inflation forecast errors stem partly from errors made in forecasting variables such as wages, the exchange rate and interest rates. We are able to reach this conclusion because the survey data used are unique in requiring respondents to forecast not just inflation but also related economic variables. We also conclude that important socio-economic factors such as firm size, the position of the respondent, and the industry that a firm belongs to have significant effects on inflation expectations.

2. The evolution of the inflation expectations literature

Curtin (2019) argues that while there is substantial disagreement about how expectations are formed and how they should be treated in theory, all social sciences incorporate expectations in their theories of decision-making. The concept of inflation expectations is central to both microeconomic and macroeconomic theory, particularly through the assumption that people aim to maximise utility. However, in macroeconomics, expectations attracted particular attention when leading macroeconomists, such as Keynes (1936), Friedman (1968) and Phelps (1967), began

uninformed or incapable of making more accurate forecasts. For example, the signalling hypothesis suggests that some FAs may offer forecasts that differ from the consensus view to signal confidence (Ashiya and Doi 2001).

to consider their implications for theory and modelling. We will not attempt to provide a comprehensive review of this very extensive literature. Instead, we will rely on a selection of research to capture the general motivation for the path the seminal literature took and focus briefly on the current position in the literature. It will be clear that the most recent literature is both active and contested.

Cagan (1956) (a student of Friedman) first introduced the concept of adaptive expectations to explicitly incorporate an expectations formation process in his money demand model. Under adaptive expectations, agents in the model update their expectations after considering their past forecast errors. In this sense, the information they use is historical and they assume some partial adjustment to their forecast errors. This backward-looking approach was criticised for not incorporating all the information available to the agent at the time that the decision was made (i.e. for allowing systematic errors (Sargent and Wallace 1976).

The rational expectations revolution of the 1970s reignited the concept of rational expectations introduced by Muth (1961) and demonstrated its broad significance. Muth (1961) argued that decision-makers would make decisions in line with their own informed view of the world (they would make *model-consistent* decisions). These decision-makers would use all the information available to them to maximise their utility (Sargent and Wallace 1976), which the proponents emphasised made the models forward-looking. Proponents did not deny that individuals sometimes made mistakes, but argued that on average these mistakes would not persist and so, at an aggregate level, modelers could model expectations as being efficient, based on the information available to individuals at the time. This still meant that the assumption implied that decision-makers know the distribution of all the stochastic shocks they face. Coibion, Gorodnichenko and Kumar (2018) label this form of expectations “full information rational expectations” (FIRE).

It is worth mentioning that opposition to the concept of rational expectations was present from the start within economics (Simon (1959) and Pesaran (1987)) as well as in the other social sciences. While economics in general did not respond quickly to these criticisms, microeconomists did start to incorporate the findings of psychology more readily, reflected in the rapidly growing popularity of behavioural economics and the award of the Nobel Memorial Prize in Economic Sciences to Daniel Kahneman in 2002 and Richard Thaler in 2015. Macroeconomics was slower to incorporate the insights from behavioural economics, but the GFC drew a lot of attention to the rational expectations assumption widely used in macroeconomic models.

Research responses within macroeconomics have varied, with some explicitly dissociating from rational expectations (De Grauwe and Ji 2019) and some presenting their suggested departures from rationality as still being broadly within the rational expectations framework (Curtin 2019). However, as these alternatives are explored, the distinction between what is and is not a version of rational expectations is becoming less precise. Notably, many of the differences between models have concentrated on what information the decision-makers use, in line with the crucial information assumptions in the Lucas (1972) islands model. Adaptive expectations assumed backward-looking individuals use their own previous expectations errors to update their expectations. In contrast, rational expectations assumed that individuals were forward-looking and used all the information available to them. However, De Grauwe and Ji (2019: 14) label this distinction "an illusion", reasoning that "the future is unknown, by [rational expectations]-agents also".[4]

Some of the alternatives to the rational expectations assumption suggest that we can still assume decision-makers on average use information in a rational manner, but they face various informational frictions. Sticky information implies that information moves slowly through the full population, rational inattention suggests that decision-makers only update their expectations periodically and, under the assumption of noisy information, people receive the true signal with an error. Apart from these models that capture informational frictions, Coibion, Gorodnichenko and Kumar (2018) group other alternatives to FIRE into 'bounded rationality' (capturing model misspecification, where decision-makers lack the capacity to efficiently process the information for various reasons) and learning models (where the decision-maker has full information and full ability to process the information efficiently but has imperfect information about the economy).

---

[4] The statistical distribution of these shocks is still based on historical information (De Grauwe and Ji 2019).

Expectations as presented by the psychologists was seen as the antithesis of rationality. Kahneman and Tversky (1979) present the deviations from full and efficient use of information by decision-makers as failures of rational decision-making. Similarly, according to behavioural macroeconomists De Grauwe and Ji (2019: 14), “we will depart from the [rational expectations]-assumption and we are not ashamed about it”.

The failures of rationality presented by the psychologists (attributed to heuristics and biases) were focused on the expectations formation process. In contrast, Curtin (2019) argues that macroeconomists have traditionally used the precision of the outcome at an aggregate level, rather than the rationality of the decision-making process, as the criterion to justify the use of rational expectations in economic models.

It is more natural for microeconomics to incorporate the findings of psychology into its theories and models given its focus on individuals and smaller groups of decision-makers. The focus on the process of expectation formation is more natural at a disaggregated level. In contrast, macroeconomic models typically aim to represent the whole economy. After all, a single policy rate is the principal instrument used by monetary policymakers to influence the entire economy. Recognition that the micro foundations of many macroeconomic models and theories are inadequate is increasingly reflected in the intense research interest in incorporating heterogeneity in modelling and use of disaggregated data in empirical research. The problem for policymakers is that if we have limited understanding of the expectations formation process (if we can say little about the behaviour that causes the final outcome), then policymakers have less insight into how to use policy to influence these behaviours. Policy is usually aimed at groups of individuals, not a theoretical representative agent.

3. The BER’s surveys

Since 2000, the BER has surveyed trade unions, businesses and FAs on a quarterly basis. The surveys aim to ensure good representation of a cross-section of the South African economy.[5] Table 1 provides an overview of the number of observations in the

---

[5] Since each respondent is identified only by an ID number, we are also able to establish that there are only a very small number of duplicate respondents surveyed over time. More precisely, 7.5% of trade union respondents, 6.5% of businesses and 5.1% of FAs are duplicates over the complete sample. There are many respondents that appear many times in the sample, but few that are present for the entire sample of over two decades. It is worth noting that that we are referring to continuity at the level of the firm (in the case of businesses) and union (in the case of the labour

surveys of the business sector and FAs for the full sample considered in this study, namely 2000Q2–2022Q4.[6] In what follows, the terms 'firm' and 'business' are used interchangeably, as are 'trade union' and 'labour'. We retain the FA – financial analyst – designation throughout.

As stated earlier, in this paper we focus primarily on the expectations of the firms and compare these with those of the FAs who are treated as a kind of benchmark. We choose not to analyse the trade union group. While we could argue that trade unions capture the wage-setting behaviour within an economy, they over-represent sectors with more organised labour, the number of respondents in the sample is much smaller than for the firms,[7] and there is less confidence that the individual responding to the survey is a senior decision-maker in the organisation (reducing the likelihood that the respondent's view represents the trade union).

Table 1: Mean forecast errors in the BER survey, 2000Q2–2022Q2: annualised

	Business sector				Financial analysts			
Year	T0	T1	T2	5A	T0	T1	T2	5A
2000	-0.80 *1.73*	-1.60 *2.10*	1.66 *2.69*		0.69 *0.80*	-0.78 *0.97*	3.66 *1.25*	
2001	-0.53 *1.02*	2.75 *1.27*	-0.50 *1.63*		-0.33 *0.52*	4.03 *0.59*	0.85 *0.87*	
2002	1.22 *1.95*	-1.63 *1.83*	-5.87 *2.06*		0.60 *1.43*	-1.23 *1.68*	-4.52 *1.66*	
2003	-2.59 *1.76*	-6.56 *2.08*	-4.39 *2.43*		-1.52 *1.27*	-2.74 *1.65*	-2.11 *1.02*	
2004	-4.24 *1.03*	-2.54 *1.15*	1.54 *1.42*		-0.94 *1.18*	-1.08 *1.15*	-0.41 *1.01*	
2005	0.11 *1.45*	0.87 *1.60*	3.00 *1.80*		-0.10 *0.51*	-0.19 *0.90*	2.45 *1.19*	
2006	0.14 *0.87*	2.24 *1.02*	6.04 *1.17*		0.09 *0.47*	1.75 *0.92*	6.01 *0.66*	
2007	1.53 *0.86*	5.36 *1.00*	1.58 *1.18*		0.82 *0.63*	5.57 *0.88*	2.49 *0.75*	
2008	1.44 *2.01*	-1.84 *1.87*	-1.17 *2.10*		0.70 *1.75*	0.53 *1.32*	-1.15 *0.87*	
2009	-1.98 *2.02*	-1.46 *2.40*	-3.52 *2.62*		0.34 *0.60*	-1.39 *0.49*	-0.66 *0.86*	

organisations), not at the level of the individual respondent. The data do include a question about the position held by the person that answers on behalf of the institution, but this individual is not tracked. In other words, if the individual holding the particular position of CEO in a particular company changes, the dataset would not track this.

6 Information about the distribution of observations by year and quarter is reported in Reid and Siklos (2022).

7 The number of trade union responses is typically less than 10 per quarter.

2010	-1.57 *1.68*	-1.83 *1.53*	-1.43 *1.65*		-0.62 *0.61*	-0.34 *0.73*	-0.16 *0.70*	
2011	-0.54 *1.05*	-0.37 *1.23*	-0.49 *1.41*	-0.87 *1.85*	0.01 *0.30*	-0.03 *0.50*	0.15 *0.60*	-0.16 *0.74*
2012	-0.46 *0.84*	-0.67 *1.03*	-0.40 *1.31*	0.99 *1.40*	-0.17 *0.36*	0.30 *0.49*	0.69 *0.50*	0.02 *0.54*
2013	-0.49 *0.75*	-0.37 *0.92*	-1.94 *1.03*	-1.19 *1.21*	-0.14 *0.16*	0.49 *0.41*	-0.93 *0.44*	-0.26 *0.55*
2014	-0.10 *0.65*	-1.75 *0.89*	0.02 *0.95*	-1.94 *1.18*	-0.07 *0.16*	-1.05 *0.29*	0.91 *0.39*	-0.48 *0.52*
2015	-1.57 *0.78*	-0.09 *0.83*	-1.21 *1.04*	-1.69 *1.20*	-0.16 *0.39*	0.48 *0.51*	-0.17 *0.56*	-0.66 *0.61*
2016	0.22 *0.92*	-0.91 *0.93*	-1.54 *1.06*	-1.83 *1.20*	-0.13 *0.25*	-0.68 *0.47*	-0.87 *0.42*	-1.21 *0.43*
2017	-1.08 *0.72*	1.76 *0.76*	-2.32 *0.90*	-1.84 *1.11*	-0.18 *0.27*	-0.55 *0.31*	-1.28 *0.36*	-1.00 *0.49*
2018	-0.96 *0.64*	-1.59 *0.70*	-0.47 *0.82*		-0.23 *0.23*	-1.10 *0.30*	-1.97 *0.36*	
2019	-0.89 *0.78*	-1.97 *0.84*	-0.85 *0.97*		-0.32 *0.30*	-1.01 *0.40*	-0.35 *0.41*	
2020	-1.11 *0.98*	-1.27 *1.09*	1.06 *1.17*		-0.17 *0.51*	0.44 *0.43*	1.78 *0.45*	
2021	0.27 *0.77*	1.44 *0.95*			0.26 *0.25*	1.71 *0.32*		
2022	0.53 *0.99*				0.58 *0.63*			

Note: The above calculations are based on micro survey results and assume forecast errors are evaluated in terms of annual inflation. The first line gives the mean; the second line the standard deviation of forecast errors. Forecast errors are observed inflation (annualised) less individual current-year forecasts (T0), one year ahead (T1), two years ahead (T2) and average expected inflation over a five-year period (beginning with the current year). Results for 2022 are for the first two quarters only. The estimates for column 5a in 2017 are based on data until the end of 2021. Observed inflation is defined so as to match the forecast horizon in question. For example, a current-year forecast in 2020 is matched with observed inflation over the same year, a one-year-ahead forecast made in 2020 requires knowledge of observed inflation in 2021, and so on for the remaining forecast horizons with the exception of the five-year forecast which is an average of inflation expected over a five-year period; $FE_{t,h} = \pi_t - F_{t,t+h}$.

Beyond the usual questions asking for (headline) inflation and economic growth (percent change in real GDP) forecasts, the survey is notable in at least three respects. First, it also asks for forecasts for a wide range of key macrofinancial variables. For example, the survey sent to trade unions and firms requests forecasts for the prime interest rate,[8] wage and salary growth, and the rand/US dollar exchange rate. Surveys of FAs add questions that elicit expectations about growth rates in the M3 money stock,[9] the yield on long-term government bonds and capacity utilisation in the

8 That is, the interest rate charged by commercial banks for loans to their best customers.

9 M3 is a broad money supply measure that includes notes, coins, commercial bank deposits, time

manufacturing sector (i.e. percentage utilisation of production capacity). Second, in addition to current-year forecasts, one-year-ahead forecasts are recorded for all the variables, except inflation for which two- and five-years-ahead (since 2011Q2) horizons are also included. Finally, unlike virtually all other firm-level forecasts we are aware of (e.g. see Coibion et al. 2020 and Reid and Siklos 2022) the BER dataset generates a time series that is more than 20 years long.

The precise wording for the inflation question is: "What do you expect average headline inflation rate to be during the year?"[10] For the longer-term inflation expectations question, respondents are asked, "What do you expect the average CPI [consumer price index] inflation rate to be over the next five years?" Respondents are then asked to fill in boxes for the current calendar year and the next two. The phrasing of the question for the other series surveyed is comparable. The way the question is phrased may suggest a focus on the view of the individual, rather than the official position of the institution that the respondent represents. However, the dataset records the forecasts against the identity of the institution rather than the individual responding, which implies that the view of this senior decision-maker represents the view of the institution. The only way in which this assumption is tested, to some extent, is by recording the position of the individual respondent in the institution. This allows us to test if the survey responses differ systematically based on the position of the respondent.

There is also some 'priming' in the survey question because respondents are provided with average inflation rates (actual inflation outturns) for the previous calendar year as well as the mean inflation rate for the last five years. The survey is a fixed-event survey, which means that the forecast horizon is determined by an event which is fixed in time. In a particular year, respondents are asked once a quarter to forecast inflation for the full year, which means that the actual horizon being forecast is shorter in the later quarters. Between 2000 and 2003, the quarterly surveys were conducted in February, May, August and October. Since that time, the February and October surveys were

deposits, money market funds and other liquid financial assets.

10 As measured by the annualised percentage change in the CPI. Between 2000 and 2008 both the CPIX (CPI, excluding mortgage costs) and the CPI were surveyed. Thereafter, only the CPI data were collected. CPIX includes the cost of shelter but not the investment portion of housing investment. Instead, a measure of the imputed rent is included. Throughout the empirical portion of this study, only CPI data are considered.

shifted to March and November. The timing of the remaining two surveys is unchanged. Respondents are anonymous to researchers.

4. Test specifications and empirical estimates

4.1 Definitions and test specifications

We begin with some notation.[11] Let π_t represent headline inflation (i.e. CPI) in quarter *t*. Since the survey data are for a calendar year, the inflation rate is the percentage change in the CPI over the calendar year in question. We also separately analyse and compare forecasts made by the two major groups covered in the BER survey, namely FAs and the business sector (B). Next, define expected inflation, $\pi_{t,t+h}^{E(j)}$, as the expected inflation rate collected at quarter *t* for the forecast horizon *t+h* where *h* is the current year, next year and two calendar years ahead. Alternatively, we write *t+h* = 0, 1 and 2, where the numerical values refer to calendar years in the future. The superscript *j* identifies the source of the forecast (i.e., *j*= FA, B).[12]

Given how the longer-run inflation expectations question is framed, we define $\bar{\pi}_t^{5ya(j)}$ as the mean five-years-ahead observed CPI inflation rate in quarter *t*, while $\bar{\pi}_t^{E,5ya(j)}$ is the expected inflation rate over the next five years (i.e. over the next 20 quarters). Consequently, forecast errors (expressed as FE below) can be expressed as:

$$FE_{t,t+h}^{(j)} = \pi_t - \pi_{t,t+h}^{E(j)} \tag{1}$$

$$FE_t^{5ya(j)} = \bar{\pi}_t^{5ya} - \bar{\pi}_t^{E,5ya(j)} \tag{2}$$

depending on the horizon one is interested in. While the study's focus concerns the behaviour of inflation expectations, we also utilise other forecasts generated by the

---

11 Recall that much of the survey data are of the fixed-event variety. Transforming from the fixed-event inflation definition used in the survey to a fixed-horizon forecast more commonly encountered in expectations data generated from models (e.g. as in the SARB's forecasts or an AR(1) model) do not greatly affect the results except possibly when inflation is volatile. Illustrations of the two types of inflation definitions are shown in Figure 1. More generally, many professional forecasts (e.g. Consensus Economics) are of the fixed-event variety.

12 Although the focus of our tests is the performance of expectations against outturns in inflation, we briefly mention two alternative benchmarks, namely forecasts for an AR(1) model and forecasts of the SARB. Econometric results for these benchmarks are available on request.

survey to test the sources of bias in inflation forecasts, as explained below. Hence, the vectors

$$\Omega^B = [\dot{y}^{E(B)}_{t,t+h}, i^{E(B)}_{t,t+h}, R^{E(B)}_{t,t+h}, \dot{W}^{E(B)}_{t,t+h}]'$$

and

$$\Omega^{FA} = [\dot{y}^{E(FA)}_{t,t+h}, i^{E(FA)}_{t,t+h}, R^{E(FA)}_{t,t+h}, \dot{W}^{E(FA)}_{t,t+h}, \dot{M}^{(FA)}_{t,t+h}, CU^{(FA)}_{t,t+h}, i^{LT(FA)}_{t,t+h}]'$$

represent the other variables for which business and FAs are asked to provide forecasts. The FA group is required to forecast three additional economic and financial variables relative to respondents in the business sector. Leaving out subscripts and superscripts for simplicity, $\dot{y}$ is real GDP growth, *i* is the prime interest rate, *R* is the rand exchange rate vis-à-vis the US dollar expressed in domestic currency units, $\dot{W}$ is (nominal) wage growth, $\dot{M}$ is M3 money growth, *CU* is capacity utilisation, and i^{LT} is the interest rate on South African government bonds with a maturity of 10 years and longer.

Rationality tests have a long lineage, going back at least to Mincer and Zarnowitz (1969), but tests have been extended and updated by, among others, West and McCracken (1998). More recently, Rossi and Sekhposyan (2016) have considered whether deviations from rationality might be impacted by structural breaks in the underlying test equation (see also Rossi and Soupré 2017). More formally, we write:

$$FE^{(j)}_{t,t+h} = \delta_0 + \delta_1 \pi^{E(j)}_{t,t+h} + \eta^{(j)}_{t,t+h} \qquad (3)$$

where all terms were previously defined. To avoid unnecessary notation, we only consider the cases where *h=0,1,*2 *and 5*. Equation (3) asks whether expectations of inflation influence FE. For example, a change in inflation forecasts may inadequately follow changes in observed inflation and only subsequently show up in changes in forecast errors.

Under the null hypothesis of rationality, forecast errors ought to be zero mean with a constant variance. In other words, if we write the null as H_o: $\delta_0 = 0$ and $\delta_1 = 0$, then

$\eta_{t,t+h}^{(j)}$ is a forecast error.[13] Hence, forecast errors are not systematically related to how expectations are formed. Rossi and Sekhposyan (2016) present a test (called the fluctuation rationality test) which is used to test whether equation (3) is robust to structural breaks. If forecast errors are sensitive to inflation expectations, and the relationship is subject to change over time (i.e. coefficient estimates are not constant), then clearly expectations are not rational in the sense defined above since this information should have been incorporated into the forecast. Accordingly, it is called the fluctuation rationality test. The test equation is estimated over rolling samples. We set each sample to roll at 10 quarters in length.

Equation (3) identifies the presence of bias, but not its sources. Therefore, we also ask whether forecast errors might have been influenced by data available to respondents when forecasts were made. Accordingly, we estimate the following version of equation (3), which is written:

$$FE_{t,t+h}^{(j)} = \gamma_0 + \gamma_1 \pi_{t,t+h}^{E(j)} + \boldsymbol{\gamma_2 \Omega_{t,t+h}^{(j)}} + \boldsymbol{\gamma_3 X_{t-1}} + \boldsymbol{\gamma_4 \Lambda_t^{(m)}} + \zeta_{t,t+h}^{(j)} \qquad (4)$$

where $\boldsymbol{\Omega_{t,t+h}^{(j)}}$ was defined above and is a vector of other variables that are forecasted by survey respondents. The vector $\boldsymbol{X_{t-1}}$ parallels the vector of other forecasts collected from survey respondents, but where the lagged observed values of the series being forecast are considered. Finally, we also consider the possibility that there are socio-economic characteristics that might bias the results. They are summarised by the vector $\boldsymbol{\Lambda}$. These were described in the preceding section but essentially consist of personal descriptors of the respondents (e.g. their roles at the firm) or characteristics that identify the industry and size of the firm where the forecasts originate. Traditionally, estimates of specifications such as equation (4) set $\gamma_1 = 0$ and $\boldsymbol{\gamma_2} = 0$ and we adopt this approach below in part because a sufficient test for biases in forecasts is to test whether known macrofinancial and socio-economic characteristics, that is, whether $\boldsymbol{\gamma_3} \neq \mathbf{0}$ or $\boldsymbol{\gamma_4} \neq \mathbf{0}$, are sufficient to reject rationality. The consequence of relaxing this assumption is, however, briefly discussed below.

13 The original Mincer-Zarnowitz regression is, in our notation, traditionally written $\pi_t = \beta_0 + \beta_1 \pi_{t,t+h}^{(j)} + \epsilon_t$.

Other than the dataset itself, we provide two novel perspectives on the behaviour of forecast errors and the consequences of potential biases in inflation expectations. First, as described below, we show how forecast errors can be highly sensitive according to the distance from the mean in the distribution of inflation forecast errors. In the first instance we explore inflation forecast errors at the quartile level and estimate a simple Vector Autoregression (VAR) consisting of forecast errors and observed inflation as endogenous variables and one lag in the observed values of the vector of variables for which respondents are also asked to provide forecasts. If $\pi(\tau)_t$ is inflation for quartile τ, then the VAR is written:

$$\boldsymbol{Y(\tau)}^{j}_{t,t+h} = \boldsymbol{B(L)Y(\tau)}^{(j)}_{t-1,t+h} + \boldsymbol{D\Omega}^{(j)}_{\boldsymbol{t-1}} + v(\tau)^{(j)}_{t,t+h} \tag{5}$$

where $\boldsymbol{Y} = [FE, \pi]'$, dropping superscripts and subscripts to economise on notation. All other terms have been defined and $v(\tau)$ is residual. The lag length of the VAR (i.e. determined by B(L)) is chosen according to Schwarz Information Criterion. The VAR allows us to trace the impact of inflation and forecast error shocks not only on their own history, but also on each other in a companion test to the one described in equation (3).

Finally, in an extension of equation (5), we estimate quantile regressions based on individual respondents' expectations of inflation and the other series they are asked to forecast. The estimated specification is written:

$$FE(\tau)^{j}_{t,t+h} = \theta_0 + \boldsymbol{\theta_1 X_{t-1}} + \boldsymbol{\theta_2 \Lambda_t} + \psi^{j}_{t,t+h} \tag{6}$$

where all terms that were previously described represent individual respondents in group *j*. We can employ estimates of equation (6) to ask how observed lagged values of the variables in the survey impact forecast errors at different quantiles. We provide estimates for each decile.

4.2 Stylised facts

Tables 1 and 2 and Figure 1 provide some stylised facts about the forecast errors made by the B and FA survey respondents. The data shown are based on equations (1) and

(2). With one exception, detailed below, we rely on observed inflation – the traditional benchmark used to estimate forecast errors.

Table 1 shows the mean and standard deviation of forecast errors by year, forecast horizon and group surveyed. On a year-by-year basis, mean forecast errors are smaller for the FA group than its B counterparts. While the data in Table 1 do not provide sources for the differences between the two groups, a plausible explanation is that respondents in the B sector employ a range of information about input and output prices to form their inflation expectations. FAs typically consider a wide array of macrofinancial data in forming their inflation outlook. On an absolute value basis, FA outperforms B forecasts in terms of the average yearly size of forecast errors. It is only at the two-year horizon that forecast errors made by B sector respondents begin to be more competitive with those of the FA group, though the latter continues to outperform the former group more regularly over the sample period. While the econometric evidence presented later cannot provide conclusive answers, it supports this observed stylised fact.

Another stylised fact from Table 1 is that both groups have tended to overestimate future inflation roughly two-thirds of the time since 2000. When the five-years-ahead average inflation forecast is considered, the degree of overestimation in both groups is substantially higher, exceeding observed inflation in almost 90% of the years where data are available (i.e. 2011–2017).

Figure 1: Time series of forecast errors from the BER survey

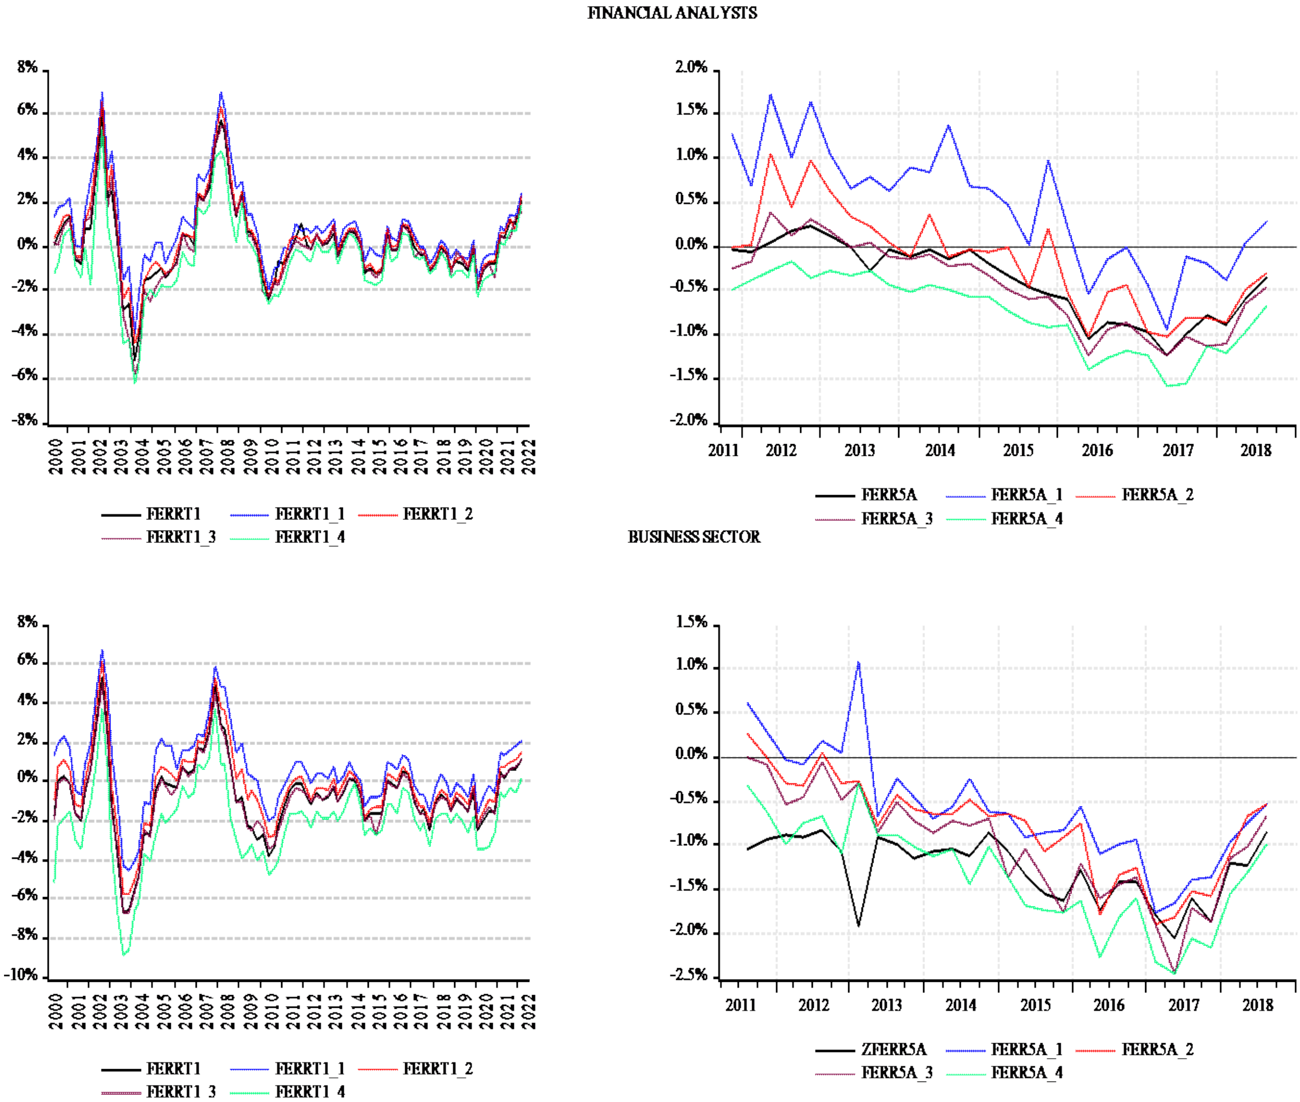

Note: See text for details.

Finally, the standard deviation of B sector forecasts is always higher than for the FA group. This has to do partly with differences in the number of observations (see Table 2) as well as possibly the wider set of backgrounds and education of the B sector respondents. Unfortunately, we do not observe the educational backgrounds or other personal details of any of the respondents in both groups. However, we can observe the industry, firm size and position of the individual respondents. We use this information below.

Next, Table 2 provides estimates of forecast errors again for the B and FA groups separately, except that the focus now shifts to the full sample and a breakdown by quartile in the distribution of errors. The skewness and kurtosis of forecast errors are also shown by horizon, quartile and group of respondents. The table provides some summary statistics in the form of the forecasts' mean percent error (MPE) and root mean squared error (RMSE). Lastly, by way of comparison, parts (b) and (c) of the table provide the same set of summary statistics when the benchmarks are the SARB's own forecasts (available only for the one- and two-year horizons) and an AR(1) model.[14]

If we define pessimistic forecasts as ones that are to the left of zero mean forecast errors (i.e. the lower two quartiles),[15] then forecast errors by both the B and FA groups tend to rise as the forecast horizon increases. In contrast, the optimistic forecasts – that is, forecasts to the right of the mean of forecast errors – tend to be skewed relatively further away from the mean than their pessimistic colleagues. Perhaps most interesting is that when the data for the full sample are considered by quartile, average forecast errors made by FA are no longer lower than those of their B sector counterparts. For example, for the five-years-ahead horizon, B sector forecast errors (in absolute value) are lower than ones obtained from the FA group. It is almost the same for the current-year (T0) forecast errors. These results are also clearly seen from the MPE and RMSE calculations.

Two other striking results are apparent from parts (b) and (c) of Table 2. When the benchmark shifts from observed inflation to the SARB's inflation forecasts or estimates from an AR(1) model, forecast errors are smaller for the full and fourth quartile for the FA group. A plausible interpretation is that SARB, AR(1) and FA group forecasts over the 2000–2022 period are more alike than forecasts from the B group. It is worth considering this result again when discussing some of the policy implications of our results.

---

14 Estimates for the AR(1) model can be found in Table 2 and it should be noted that the SARB's forecast are fixed-horizon forecasts.

15 Since forecast errors to the left of the mean are negative, this means that observed inflation is less than expected inflation, and the reverse when forecast errors are positive. Those who underestimate inflation are optimists about future inflation while overestimates of inflation are suggestive of pessimism about the future course of inflation.

Table 2: Summary statistics: forecast errors in the BER survey over the full sample, 2000Q2–2022Q2

a) Benchmark: observed inflation

Full sample	**Business**				**Financial analysts**			
Horizon	T0	T1	T2	5a	T0	T1	T2	5a
Mean	-0.07	-0.07	0.19	-0.38	-0.64	-0.86	-0.95	-1.29
Standard deviation (SD)	1.14	2.35	2.61	0.44	1.66	2.78	2.93	0.36
Skew	-0.72	0.60	0.74	-0.40	-0.76	0.12	0.59	-0.65
Kurtosis	7.87	5.54	4.05	1.80	5.44	4.22	3.39	2.23
First quartile								
Mean	0.44	0.80	1.09	0.48	0.57	0.51	-0.26	-0.58
SD	1.21	2.34	2.70	0.69	1.69	2.54	3.18	0.65
Skew	0.02	0.93	0.89	-0.11	-0.07	0.49	0.56	0.48
Kurtosis	7.69	5.40	3.76	2.25	5.40	4.07	3.22	3.17
Second quartile								
Mean	0.06	0.29	0.63	-0.15	-0.15	-0.29	-0.54	-0.79
SD	1.12	2.34	2.58	0.58	1.68	2.63	3.10	0.57
Skew	-0.67	0.72	0.82	0.24	-0.59	0.36	0.45	-0.35
Kurtosis	7.57	5.57	3.98	2.40	5.38	4.38	3.24	2.41
Third quartile								
Mean	-0.11	-0.06	0.04	-0.45	-0.73	-0.91	-0.89	-1.01
SD	1.11	2.44	2.50	0.50	1.66	2.75	3.10	0.62
Skew	-0.84	0.67	0.76	-0.12	-0.85	0.13	0.37	-0.28
Kurtosis	7.90	5.53	4.08	1.79	5.81	4.42	3.20	2.45
Fourth quartile								
Mean	-0.43	-0.62	-0.54	-0.76	-1.73	-2.19	-1.31	-1.35
SD	1.21	2.43	2.65	0.43	1.74	2.97	3.21	0.59
Skew	-1.14	0.16	0.39	-0.45	-0.97	-0.20	0.23	-0.11
Kurtosis	7.41	5.44	4.42	1.91	5.63	4.42	3.21	2.13
MPE	**Business**				**Financial analysts**			
Full sample	-20.19	-8.83	-5.18	-8.31	-48.49	-41.43	-38.12	-25.92

First quartile	-6.16	14.20	18.12	8.25	-16.93	-3.99	-19.54	-12.49
Second quartile	-15.78	-0.90	1.52	-4.04	-36.67	-36.14	-27.78	-16.63
Third quartile	-20.33	-14.52	-9.12	-9.95	-51.46	-42.68	-37.29	-21.11
Fourth quartile	-31.78	-29.99	-24.64	-15.98	-75.94	-76.93	-52.25	-27.79
RMSE	**Business**				**Financial analysts**			
Full sample	5.85	6.26	6.48	5.49	6.36	7.26	7.64	6.40
First quartile	5.34	5.61	5.78	4.63	5.12	5.94	7.21	5.70
Second quartile	5.74	6.05	6.25	5.25	5.85	6.67	7.41	5.91
Third quartile	5.91	6.37	6.55	5.55	6.44	7.30	7.69	6.13
Fourth quartile	6.23	6.93	7.14	5.86	7.48	8.54	8.07	6.47
Observations	89	85	81	28	89	85	81	28

b) Benchmark: SARB Monetary Policy Committee forecasts

Full sample	**Business**				**Financial analysts**			
Horizon	T0	T1	T2	5a	T0	T1	T2	5a
Mean		-0.61	-0.99			0.21	0.001	
SD		1.11	0.96			0.64	0.43	
Skew		-0.52	-0.44			1.04	-0.07	
Kurtosis		3.68	3.51			7.07	3.11	
First quartile								
Mean		0.64	-0.29			0.85	0.81	
SD		0.96	1.26			0.80	0.69	
Skew		1.41	0.31			1.95	0.54	
Kurtosis		5.57	4.84			9.06	2.48	
Second quartile								
Mean		-0.10	-0.54			0.39	0.20	
SD		0.93	1.13			0.68	0.49	
Skew		0.52	-0.42			1.86	0.36	

Kurtosis		4.03	4.93			9.88	2.98	
Third quartile								
Mean		-0.67	-0.88			0.06	-0.16	
SD		1.06	1.14			0.56	0.42	
Skew		-0.54	-1.01			0.03	0.09	
Kurtosis		4.06	4.51			5.21	3.04	
Fourth quartile								
Mean		-1.81	-1.23			-0.38	-0.59	
SD		1.28	1.25			0.61	0.46	
Skew		-1.47	-1.50			-0.67	-0.30	
Kurtosis		5.66	5.39			6.34	2.64	
MPE								
Full sample		-11.19	-18.98			3.41	-0.22	
First quartile		11.69	-5.22			15.09	15.29	
Second quartile		-1.78	-10.88			6.70	3.64	
Third quartile		-12.02	-16.61			0.71	-3.30	
Fourth quartile		-33.03	-23.03			-7.14	-4.56	
RMSE								
Full sample		6.65	6.30			5.73	5.23	
First quartile		5.01	5.68			4.74	4.45	
Second quartile		5.73	5.89			5.19	5.03	
Third quartile		6.32	6.23			5.51	5.39	
Fourth quartile		7.48	6.59			5.96	6.82	
Observations		70	70			70	70	

c) Benchmark: 'naïve' AR(1) forecast

Full sample	**Business**				**Financial analysts**			
Horizon	T0	T1	T2	5a	T0	T1	T2	5a
Mean	-0.64	-0.79	-0.83	-0.84	-0.08	0.13	0.28	0.06
SD	1.60	1.33	1.15	0.38	1.80	0.88	0.50	0.17
Skew	-0.71	-0.44	-0.27	-0.67	-0.92	-0.68	0.24	-0.09
Kurtosis	3.41	2.99	2.86	6.50	4.83	4.61	2.57	3.03
First quartile								
Mean	0.55	0.58	-0.10	-0.13	0.42	0.85	1.15	0.94
SD	1.33	1.06	1.40	0.36	1.70	0.89	0.60	0.34
Skew	0.42	1.08	0.32	0.95	-0.60	0.61	0.73	0.44
Kurtosis	3.73	4.84	3.96	4.28	5.08	4.79	2.98	2.58
Second quartile								
Mean	-0.15	0.22	-0.40	-0.35	0.04	0.35	0.51	0.31
SD	1.47	1.13	1.32	0.29	1.83	0.87	0.47	0.27
Skew	-0.21	0.15	-0.15	0.61	-0.88	-0.16	0.52	0.54
Kurtosis	3.44	3.45	3.43	3.34	4.90	4.89	2.24	2.72
Third quartile								
Mean	-0.73	-0.83	-0.74	-0.57	-0.12	-0.002	0.04	0.01
SD	1.59	1.33	1.37	0.32	1.86	0.88	2.50	0.16
Skew	-0.81	-0.51	-0.56	0.30	-0.93	-0.98	0.76	-0.39
Kurtosis	3.75	3.10	3.08	3.08	4.83	4.96	4.08	3.37
Fourth quartile								
Mean	-1.70	-2.40	-1.14	-0.91	-0.44	-0.77	-0.41	-0.30
SD	1.87	1.63	1.57	0.28	1.89	1.00	0.77	0.16
Skew	-1.34	-0.94	-0.91	0.77	-1.06	-1.50	-0.96	0.91
Kurtosis	4.47	3.28	3.40	2.71	4.62	5.37	5.50	2.86
MPE								
Full sample	-11.74	-14.95	-15.13	-15.19	-1.54	2.26	5.09	1.16
First quartile	9.88	10.38	-1.93	-2.42	7.55	15.43	20.85	16.88
Second quartile	-2.88	-4.15	-7.34	-6.27	0.65	6.25	9.31	5.54

Third quartile	-13.34	-15.32	-13.59	-10.29	-2.38	-0.15	2.53	0.11
Fourth quartile	-31.17	-13.91	-20.88	-16.41	-8.15	-14.12	-7.53	-5.47
RMSE								
Full sample	6.34	6.42	6.43	6.40	5.86	5.44	5.25	5.49
First quartile	5.13	5.04	5.78	5.70	5.35	4.73	4.41	4.63
Second quartile	5.84	5.83	6.05	5.91	5.75	5.22	5.02	5.25
Third quartile	6.42	6.47	6.39	6.12	5.92	5.57	5.39	5.55
Fourth quartile	7.44	8.05	6.82	6.47	6.23	6.43	5.96	5.86
Observations	88	88	87	28	88	88	87	28

Note: See note to Table 1. The SARB's Monetary Policy Committee forecasts are only available for the one-year and two-years-ahead horizons. Observed inflation (adjusted for the horizon of the forecast) less BER forecasts by individual respondents. Parts (b) and (c) change the benchmark to SARB one-year- and two-years-ahead horizons (fixed horizon forecasts) and an AR(1) model applied to the whole sample (2000Q2–2022Q2) for observed inflation.

Finally, Figure 1 plots the time series of forecast errors based on the next calendar year forecasts and the average inflation forecast over the next five years. There are a few additional stylised facts that are less apparent from the ones surrounding Tables 1 and 2. First, forecast errors spike around the time of the rand and banking crises of 2001 and 2002 and again during the GFC. Next, the forecast errors for one calendar year ahead clearly show a sharp and continuous rise as the coronavirus crisis and its aftermath take hold, beginning in 2020. While the overall profile of forecast errors is comparable at the one-year-ahead horizon for both FA and B groups surveyed, the gap across the quartiles is noticeably higher among the B sector than in its FA counterparts. There is also a long period, from 2011 to 2019, when forecast errors are small and fluctuate around zero. For the most part, this interpretation applies to almost all quartiles and for both the FA and B groups. At the macro level these were relatively calm periods, which may go some way to explaining the superior forecast performance of both groups.

Turning to the five-years-ahead forecast horizon, both FA and B groups underestimate long-term inflation rates at the beginning of the available sample (i.e. 2011–2013), although the B group begins to overestimate inflation earlier and to a greater extent than FAs over the last several years. Notice also that differences in forecast errors across the quantiles are quite visible for both FA and B groups. Therefore, optimists and pessimists can and do differ in both groups of forecasters.

4.3 Forecast biases and their sources: some econometric estimates

4.3.1 Aggregated survey data

The stylised facts suggest that, over time, forecast errors do not behave as the strict rational expectations hypothesis would predict. Figures 2 and 3 rely on equation (3) as the starting point to test and pinpoint breakdowns in the usual test of rationality in inflation expectations. Figure 2 provides a graphical representation of the Rossi-Sekhposyan (2016) test statistic for the FA group at all forecast horizons. The null hypothesis is whether there is no structural break in equation (3). Figure 3 repeats the same exercise for the B group of forecasters. The dating of rejections of forecast rationality is summarised in a table below each figure.

We identify four conclusions based on the results in the two figures. At all forecast horizons, and for both groups, the year 2020 is the source of a break. To the extent that the coronavirus crisis was unexpected, and its economic impact unpredictable, this result is hardly a surprise. Although the number of times the null hypothesis of the Rossi-Sekhposyan test (see previous discussion) is rejected is largely the same at all horizons, the period over which forecast rationality is rejected can vary greatly. In essence, the longer the horizon the longer the test statistic exceeds the critical value shown by the horizontal dashed lines in Figures 2 and 3.

Figure 2: Fluctuation rationality test: financial analysts

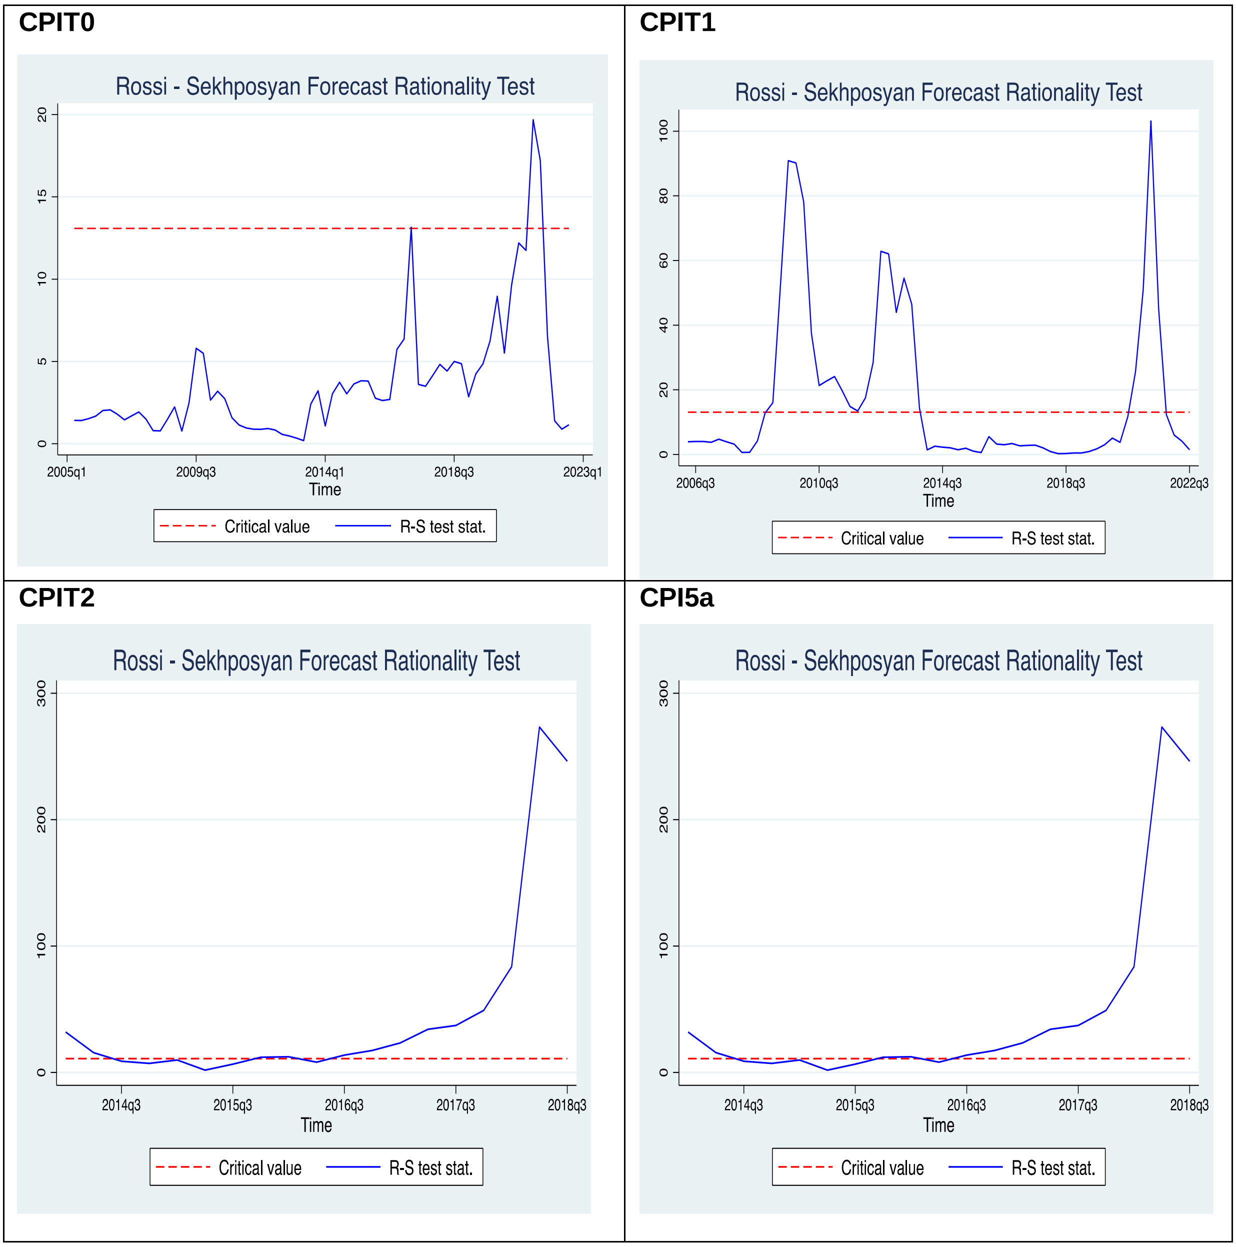

Dating of breaks

T0	2016Q4, 2021Q1–2021Q2
T1	2008Q4–2013Q4, 2020Q4–2021Q2
T2	2013Q4–2018Q3, 2019Q2–2022Q2
5a	2013Q3–2013Q4, 2016Q2–2018Q2

Note: The text provides details of the test equation and statistics.

Figure 3: Fluctuation rationality test: business sector

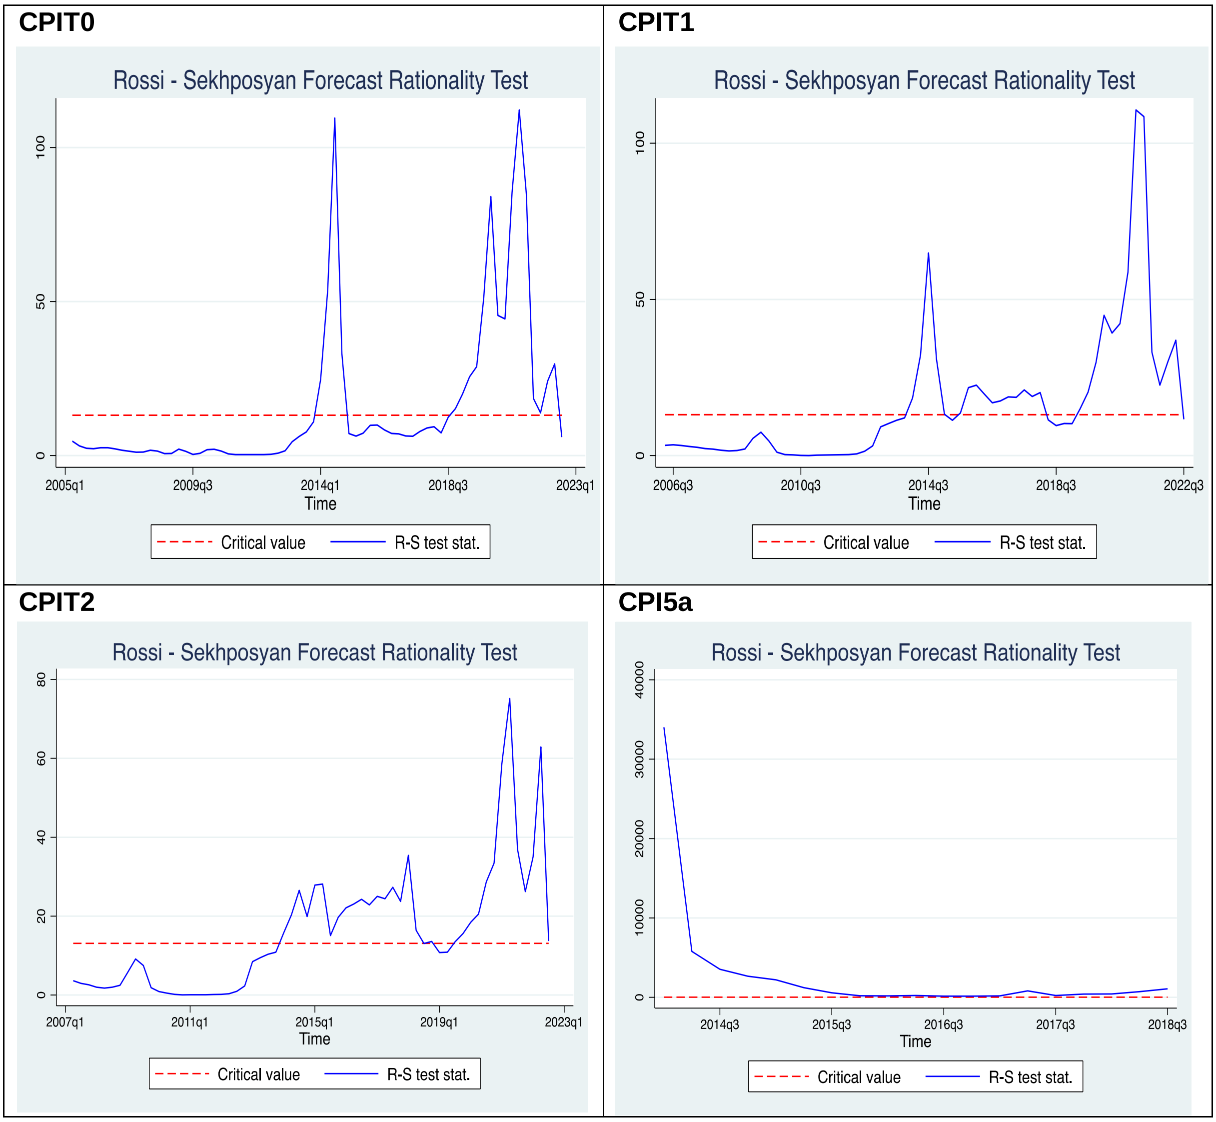

Dating of breaks

T0	2013Q4–2014Q3, 2018Q3–2022Q1
T1	2013Q4–2017Q4, 2019Q1–2022Q2
T2	2013Q4–2018Q3, 2019Q1–2022Q2
5a	Entire sample rejects H_0

Note: See note to Figure 2.

It is also clear for both groups of forecasters that each horizon highlights different events that could explain the rejection of rationality. For example, the GFC and the euro area sovereign debt crises (2008–2009 and 2011–2013, respectively) impact FA forecasts for the next year but not the two-years-ahead forecast horizon. So, it appears that the FAs are influenced by external crisis events, but that they do not expect these impacts to last longer term. In contrast, rationality is not rejected as a consequence of either external crisis among the B group of forecasters. There is a one-time rejection of the null among the B group for the one-year-ahead horizon in 2014Q3 and then again in 2018. It is likely that the 2018 event was at least partly influenced by large exchange rate movements in response to political events. The election of President Cyril Ramaphosa in December 2017 was greeted with improved business sentiment and a stronger rand, which reversed dramatically after a debate surrounding expropriation of land without compensation gained momentum in 2018Q2. The rand depreciated from over R12/US$ in December 2017 to almost R15/US$ in September 2018. The rejection of rationality in 2014Q3 is less obvious. Rejections are similar for both groups at the two-years-ahead horizon. Finally, differences between the two groups emerge at the five-years-ahead horizon with a near complete rejection of rationality for the full sample of the B group while rejections become persistent for the FA group beginning in 2016.

Clearly, both groups of forecasters fail the rationality test but, perhaps more importantly, breaks in the relationship based on equation (3) can be highly sensitive according to whether one focuses on the FA or the B sector. This suggests not only that each sector may miss some clues that might have improved their forecasts, but that different and important economic events reveal their inattention, leading them to forecast differently, and incorrectly.

We turn now to discussing selected estimates based on the VARs specified in equation (5). Figures 4 and 5 plot the impulse responses for the FA group while Figures 6 and 7 display the impulse responses for the B sector group. The impulse responses shown are for the T1 (next calendar year) (top portion of the figures) and 5a (average over 5 years) cases (bottom portion of the figures) only. As shown in Table 2, the behaviour of expectations can be highly sensitive to the location of respondents along the distribution of forecast errors.

While the VARs defined in equation (5) were estimated for all quartiles, Figures 4 and 5 highlight the differences in behaviour between expectations in the left and right tails of the distribution of forecast errors, that is, for the first and fourth quartiles. The impulse responses are based on estimates via local projections.[16] For comparison, we also plot the point estimates of the impulse responses that would be obtained if the VARs were estimated in the usual fashion using a Cholesky decomposition with inflation and forecast errors as endogenous variables in that order.

In the case of FAs, all shocks displayed have temporary effects. A positive shock to forecast errors among respondents in the first quartile[17] is associated with higher future inflation. While the local projections are silent about the mechanism that generates such a result, it is plausible to think that a rise in forecast errors presages a rise in future inflation. There is some visual evidence of this phenomenon in Figure 1. On the other hand, a positive shock to inflation does lead to lower forecast errors in the future. Hence, a rise in inflation may be seen as a wake-up call prompting better inflation forecasting performance (possibly through more attentiveness).

What Figure 4 cannot say, however, is the precise source of this improvement. The behaviour of respondents in the fourth quartile, namely forecasters who tend to underestimate inflation the most, differs from those who overestimate inflation when the impulse responses of inflation to forecast errors are considered. At first, the former rises when there is a positive shock to inflation, but in later quarters the sign of the relationship changes. When cumulated (results not shown), the total impact is not different from zero. In contrast, as with the left tail of the distribution of forecast errors, a positive shock to forecast errors sends a signal of higher future inflation. In other words, individuals who tend to overestimate inflation tend to incorporate the shock in a more permanent manner than those who tend to underestimate inflation. It is likely

---

16 Local projections is a technique proposed by Jordà (2005) which simplifies and improves how impulse responses (i.e. numerical estimates of the response to economic shocks) are estimated. To conserve space, own impulse responses are not shown (i.e. the response of a shock to inflation or forecast error on themselves). In common with most macroeconomic series, and the stylised facts in Tables 1 and 2, there is considerable persistence in the responses of forecast errors to their past and inflation to its own history. This result holds for both FA and B groups. The omitted impulse responses are in Figures 4, 5 and 6.

17 A reminder that these respondents generate the largest negative forecast errors.

that there will always be some proportion of the respondents that is more pessimistic. From a policy point of view, it may be useful to monitor the extent to which the distribution changes as this may be a signal of changes in inflationary pressure. It may also be useful to monitor whether certain communications efforts appear to influence this distribution.

Figure 4: Impulse responses: financial analysts

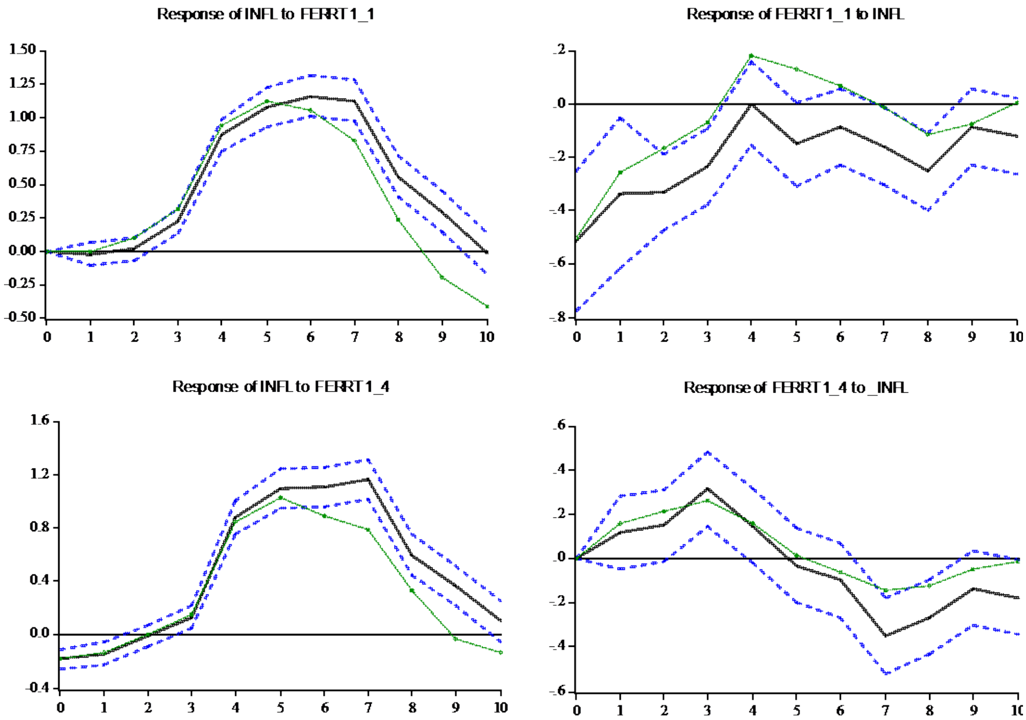

Figure 5 repeats the same exercise for the five-year-ahead (5a) average inflation forecasts. The results in this case stand in sharp contrast with the ones shown in Figure 4. There is no statistically significant link between forecast errors and inflation. There is only a very small positive response of forecast errors from a rise in the average five-year-ahead inflation rate (at the first and second quarters). These results suggest that central banks that rely on longer-term forecasts from FA alone can be easily misled into thinking that expectations are anchored, whereas the expectations of B (price setters) behave differently.

Figure 5: Impulse responses: financial analysts

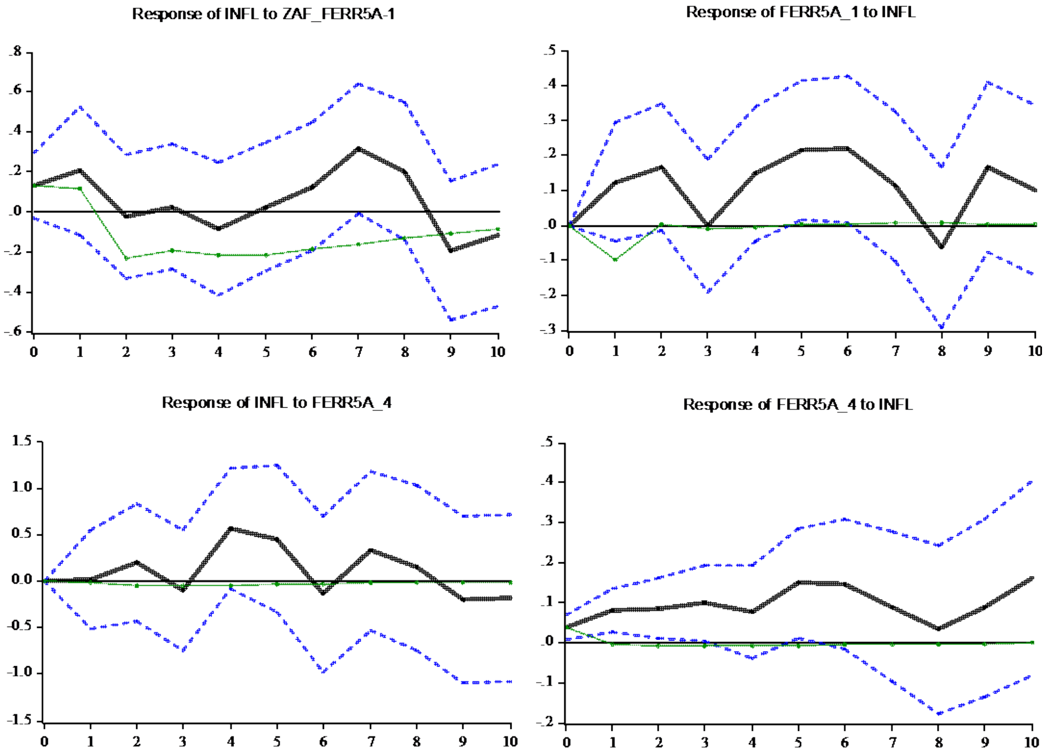

We now turn our attention to the impulse responses for the B group of forecasts and forecast errors in Figures 6 and 7. To avoid repetition, we focus on the differences between the B and FA forecasters' responses to shocks. At the first quartile for the T1 horizon, the responses are similar for both groups. There is only a small difference in the size of the impulse responses at the peak (i.e. around six or seven quarters). Similarly, for the respondents in the fourth quartile (bottom portion of Figure 6) the responses by the B forecasters are comparable to the ones estimated for the FA group, although, again, the size of the largest responses is slightly smaller for the B respondents. The major difference in responses between the two groups takes place when forecast error reactions to an inflation shock (bottom right-hand-side impulse responses) are considered. Unlike with the FA group, there is a sharp, but temporary, decline in forecast errors for the B group when inflation unexpectedly increases. Hence, at least among the respondents that tend to underestimate future inflation, there appears to be some learning that takes place.

Figure 6: Impulse responses: business sector

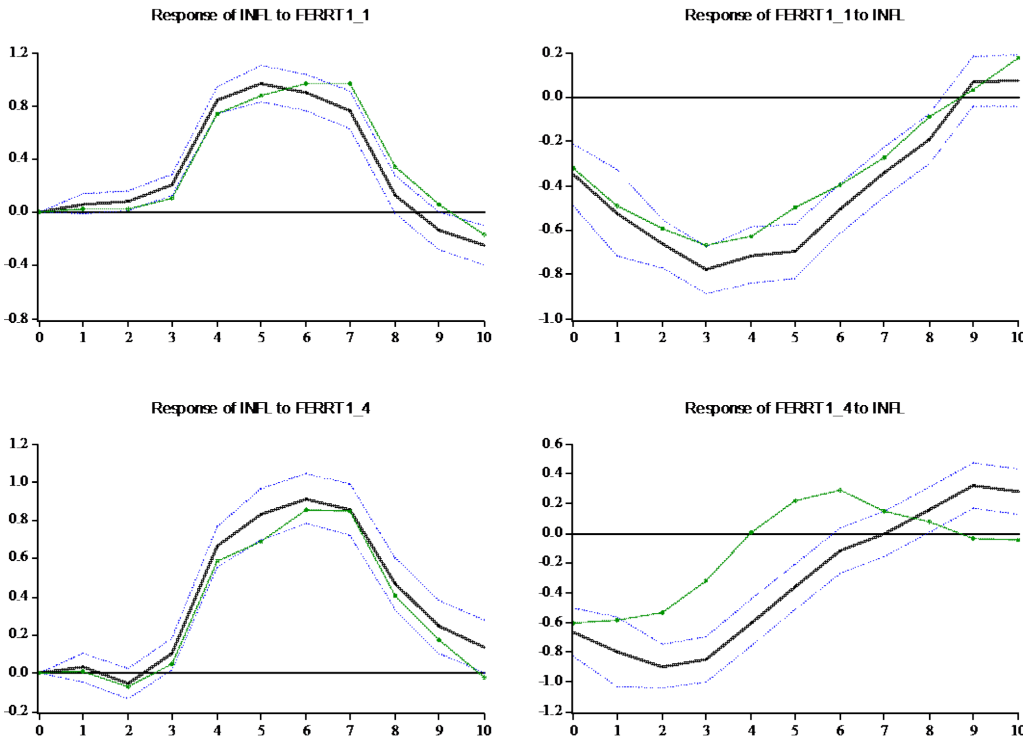

Figure 7: Impulse responses: business sector

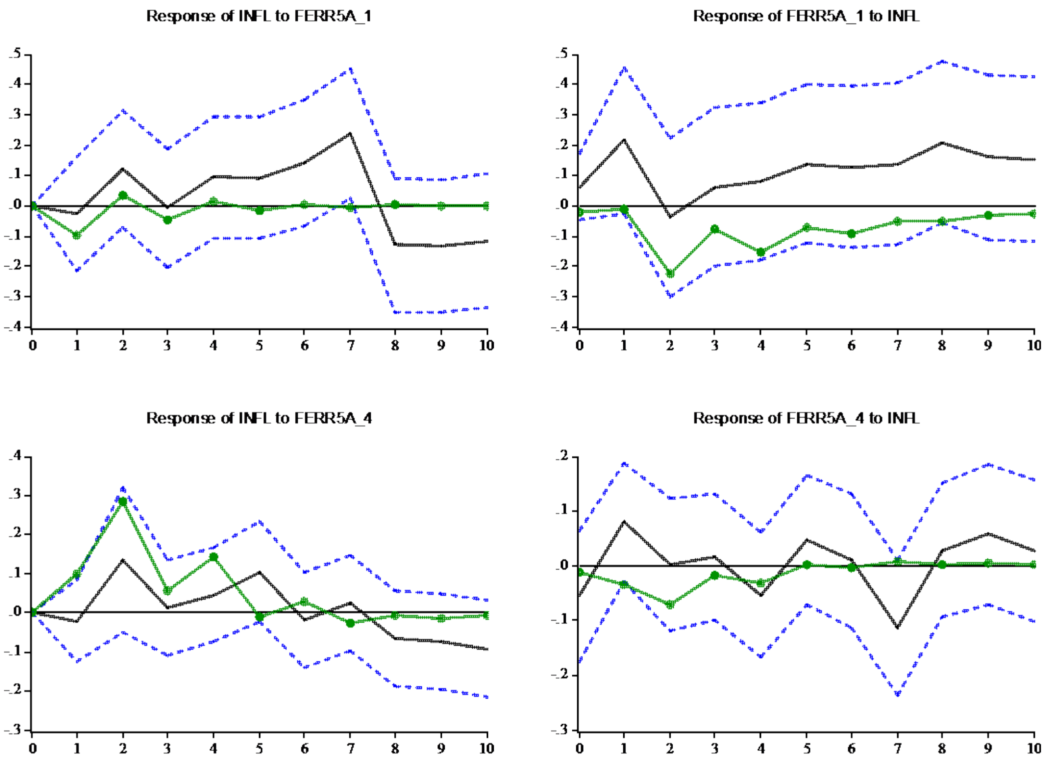

Figure 7 repeats for the B group the exercise conducted for the FA group shown in Figure 5. For the five-years-ahead forecasts there is essentially no difference between the FA and B forecasters, with one exception. While FA forecast errors are seen as rising very slightly in the event of an inflation shock, there is no response whatsoever to the same shock among the B group of respondents to the BER survey. While inflation and forecast errors are strongly related for short horizon forecasts,[18] the local projections for the five-years-ahead horizon are far less informative. Clearly, something else is driving forecasts that the VARs are unable to identify. It is worth highlighting that the results presented so far represent estimates based on data aggregated from individual responses. As such, we cannot learn whether there are socio-economic or other factors that can also contribute to our understanding of forecast errors, nor can we address the impact of aggregating individual-level data.

4.3.2 Individual-level survey data

Table 3 presents estimates of quantile regressions relying on individual-level data for both the FA and B groups. The estimates shown are the median in the distribution of forecast errors based on equation (6). This is followed by a discussion of more detailed estimates at different quantiles paralleling the approach taken with the aggregated survey data. To conserve space, we restrict the discussion to the one-year-ahead and average five-years-ahead forecast errors.

The discussion is divided into two parts. If we examine the lagged outcome values for some of the series that also form part of the survey (the variables contained in Ω^B and Ω^{FA}, as described in section 4.1), we observe in both groups that attention to their history could have reduced forecast errors. In particular, this applies to the wage growth variable for both groups, while incorporating the history of the prime rate would also have reduced forecast errors in the B group of respondents. In the case of the remaining variables in the survey, including past outturns in the series would have led to a deterioration in forecast performance. This applies to real GDP growth, inflation, lagged forecast errors, the rand exchange rate, money growth and the long-term interest rate.[19] Finally, and for the most part, the same variables bias forecasts in the B group at both the one-year-ahead and five-years-ahead horizons. In the case of the

---

18 Broadly speaking, conclusions are similar for the T0 and T2 horizons (not shown).

19 For convenience, the vector of fundamental macrofinancial time series includes variables for which the B group is not asked to provide forecasts. They are: capacity utilisation, M3 money growth and the long-term interest rate. It is plausible that any, or all, of these variables might be incorporated into the information set of B forecasters.

FA respondents, lagged wage growth and real GDP growth have no impact on forecasts of average inflation over the next five years, while inclusion of the rand would produce a deterioration of the same forecast.

Table 3: The determinants of BER survey forecast errors: selected quantile regression estimates at the median

Dependent variable: forecast errors	**Financial analysts**		**Business sector**	
Determinants	Horizon T1	Horizon 5a	Horizon T1	Horizon 5a
Constant	-3.61 (.41)*	4.613 (1.16)*	-1.834 (.20)*	-1.834 (.20)*
Prime rate	0.007 (.03)	0.034 (.03)	-0.275 (.09)*	-0.275 (.01)*
Real GDP growth	0.019 (.007)*	0.022 (.01)	0.090 (.01)*	0.090 (.01)*
Wage growth	-0.034 (.01)*	-0.029 (.02)	-0.071 (0.01)*	-0.071 (.01)*
Rand exchange rate	-0.036 (.02)	0.099 (.03)*	0.071 (.01)*	0.071 (.01)*
CPI inflation	0.446 (.04)*	-0.237 (.09)*	0.731 (.01)*	0.731 (.01)*
Forecast error lagged	0.435 (.04)*	0.408 (.05)*	0.169 (.01)*	0.169 (.01)*
Capacity utilisation	-0.0004 (.00004)	-0.049 (.01)*	0.000 (0.000)	-0.000 (.000)
M3 growth	0.022 (.007)*	-0.016 (.02)	0.076 (.003)*	0.076 (.003)*
Long-term interest rate	0.101 (.03)*	0.255 (.066)*	-0.005 (.01)	-0.005 (.01)
Socio-economic determinants				
Small firms	0.056 (.10)	0.954 (.15)*	0.038 (.03)	0.038 (.03)
Medium-sized firms	0.094 (.07)	-0.163 (.07)**	0.085 (.03)*	0.085 (.03)*
Large firm	0.100 (.06)†	-0.222 (.065)*	0.169 (.03)*	0.169 (.03)*
Banks	0.135 (.05)*	0.001 (.07)	NA	NA
Advisors	0.156 (.04)*	-0.059 (.05)	NA	NA
Time	0.037 (.01)*	-0.216 (.04)*	-0.048 (.01)*	-0.048 (.01)*
Q1	0.023 (.06)	0.305 (.07)*	0.089 (.03)*	0.089 (.03)*
Q2	0.004 (.06)	0.391 (.09)*	-0.006 (.03)	-0.006 (.03)
Q3	0.037 (.06)	0.416 (.10)*	0.017 (.02)	0.017 (.02)
Financial manager			-0.038 (.04)	-0.038 (.04)
CEO			-0.011 (.04)	-0.011 (.04)

Production manager			-0.124 (.06)**	-0.124 (.06)**
Other			0.169 (.08)**	0.169 (.08)**
Agriculture			-0.432 (.13)*	-0.432 (.13)*
Mining			-0.477 (.137)*	-0.477 (.14)*
Manufacturing			-0.398 (.130)*	-0.398 (.13)*
Construction			-0.486 (.134)*	-0.486 (.13)*
Wholesale and retail			-0.545 (.13)*	-0.545 (.13)*
Transportation			-0.414 (.14)*	-0.414 (.14)*
Finance and real estate			-0.332 (.13)*	-0.332 (.13)*
Community and social services			-0.327 (.14)**	-0.327 (.14)**
No. of observations	1381	540	22570	22570
Pseudo-R^2	0.60	0.39	0.38	0.38

Turning to the socio-economic and other determinants of forecast errors, we observe that – at the median for the FA group of forecasters – large institutions, banks and financial advisors[20] have poorer forecast records than the rest. At the five-years-ahead forecast horizon, smaller firms also display poorer forecast records than the rest, while medium-sized and larger firms outperform others at the median. For the B group, only large and medium-sized firms generate relatively poorer forecasts, while neither firm size nor the position of the person who completes the survey has any statistical impact, at least for the median respondents, on forecast performance. This finding is surprising and requires further investigation. For the B group of respondents, we can also identify industries that perform better (or worse) than the others, conditional on the others in the regression. Forecasts from the agricultural, mining, manufacturing, construction, wholesale and retail, transportation, finance, and community service sectors deliver reductions in forecast errors than the remaining sectors, other things being equal, and the extent of the reduction is comparable across these industries.[21] Unfortunately, we have too little information about, for example, whether or not information sets might differ across industries. In line with the literature on subjective expectations of households (Weber et al. 2022), it is very likely that firms consider their own circumstances in forming their expectations rather than relying on the official CPI experience, and therefore firms in different industries have varied information sets. This too requires more attention in future research.[22] Note, however, that other biases in forecast performance noted earlier would offset the benefits of position or industry considered.

Finally, it is worth noting two other conclusions from Table 3. First, as the duration of the inflation-targeting regime has increased (i.e. the time variable), forecast errors have generally improved across groups and forecast horizons. Hence, the longer that the inflation-targeting policy strategy has been in place, the easier it has become to forecast inflation. Once lower and stable inflation became more persistent, there was greater scope to reduce forecast errors. Some credit is due to the SARB because it has succeeded in delivering inflation rates within the inflation target band. This

---

[20] This group also includes financial brokers and investment managers.

[21] This is confirmed by a Wald test (not shown).

[22] An extension not considered here, for example, would be to interact industry and firm size. However, we have no prior information about the likely forecast performance of, say, small versus large firms in a particular industry.

suggests a credibility dividend. Of course, as we now know, it helps if the central bank is able to navigate the shocks hitting the economy. Large shocks that generate more volatile inflation rates (e.g. the coronavirus pandemic and its aftermath) may well reverse this result. It also needs to be pointed out that we are unable to observe from the dataset the extent to which the parameters of the inflation-targeting regime are understood by survey respondents or how SARB communication may have led to lower forecast errors over time.

Second, there is some evidence of seasonality in the forecasts and these worsen forecast performance. For example, at the median, forecasts generated in the first quarter tend to be relatively larger for both groups of forecasters. This might reveal something about how forecasts are revised or not across the four quarters of the survey. Similarly for the FA group, second- and third-quarter forecasts are poorer than year-end (i.e. fourth-quarter) forecasts at the five-years-ahead horizon.[23]

Next, paralleling the approach used to describe and analyse aggregated survey data, we proceed with estimates from the quantile regressions (equation (6)). The quantile regression results are shown in Figures 8 and 9. Figures 8a and 8b display the coefficients by quantiles for the FA group; Figures 9a and 9b repeat the exercise for the B group of respondents. To conserve space, the focus of our discussion is on comparing the response of FA and B group forecast errors to a selection of macrofinancial fundamentals at different quantiles. Note that the results are based on individual data rather than the aggregated data used to date.

[23] Readers will have noticed that equations (4) and (6) are similar except that (4) also considers forward-looking survey questions as potential determinants of forecast errors. Alternatively, if we impose the restriction that $\gamma_1 = 0, \boldsymbol{\gamma}_2 = 0$, then equations (4) and (6) are essentially the same, other than the former is specified for aggregated data while the latter is estimated with individual data. We also estimated a version of equation (6), which relaxes these restrictions. In effect we ask how forecasts are related to each other, a topic beyond the scope of this paper. Suffice to say, however, that forecasts at all horizons are improved by asking respondents to think ahead about different macrofinancial variables at different horizons. Results are available on request.

Figure 8a: Quantile regressions (one-year horizon): financial analysts

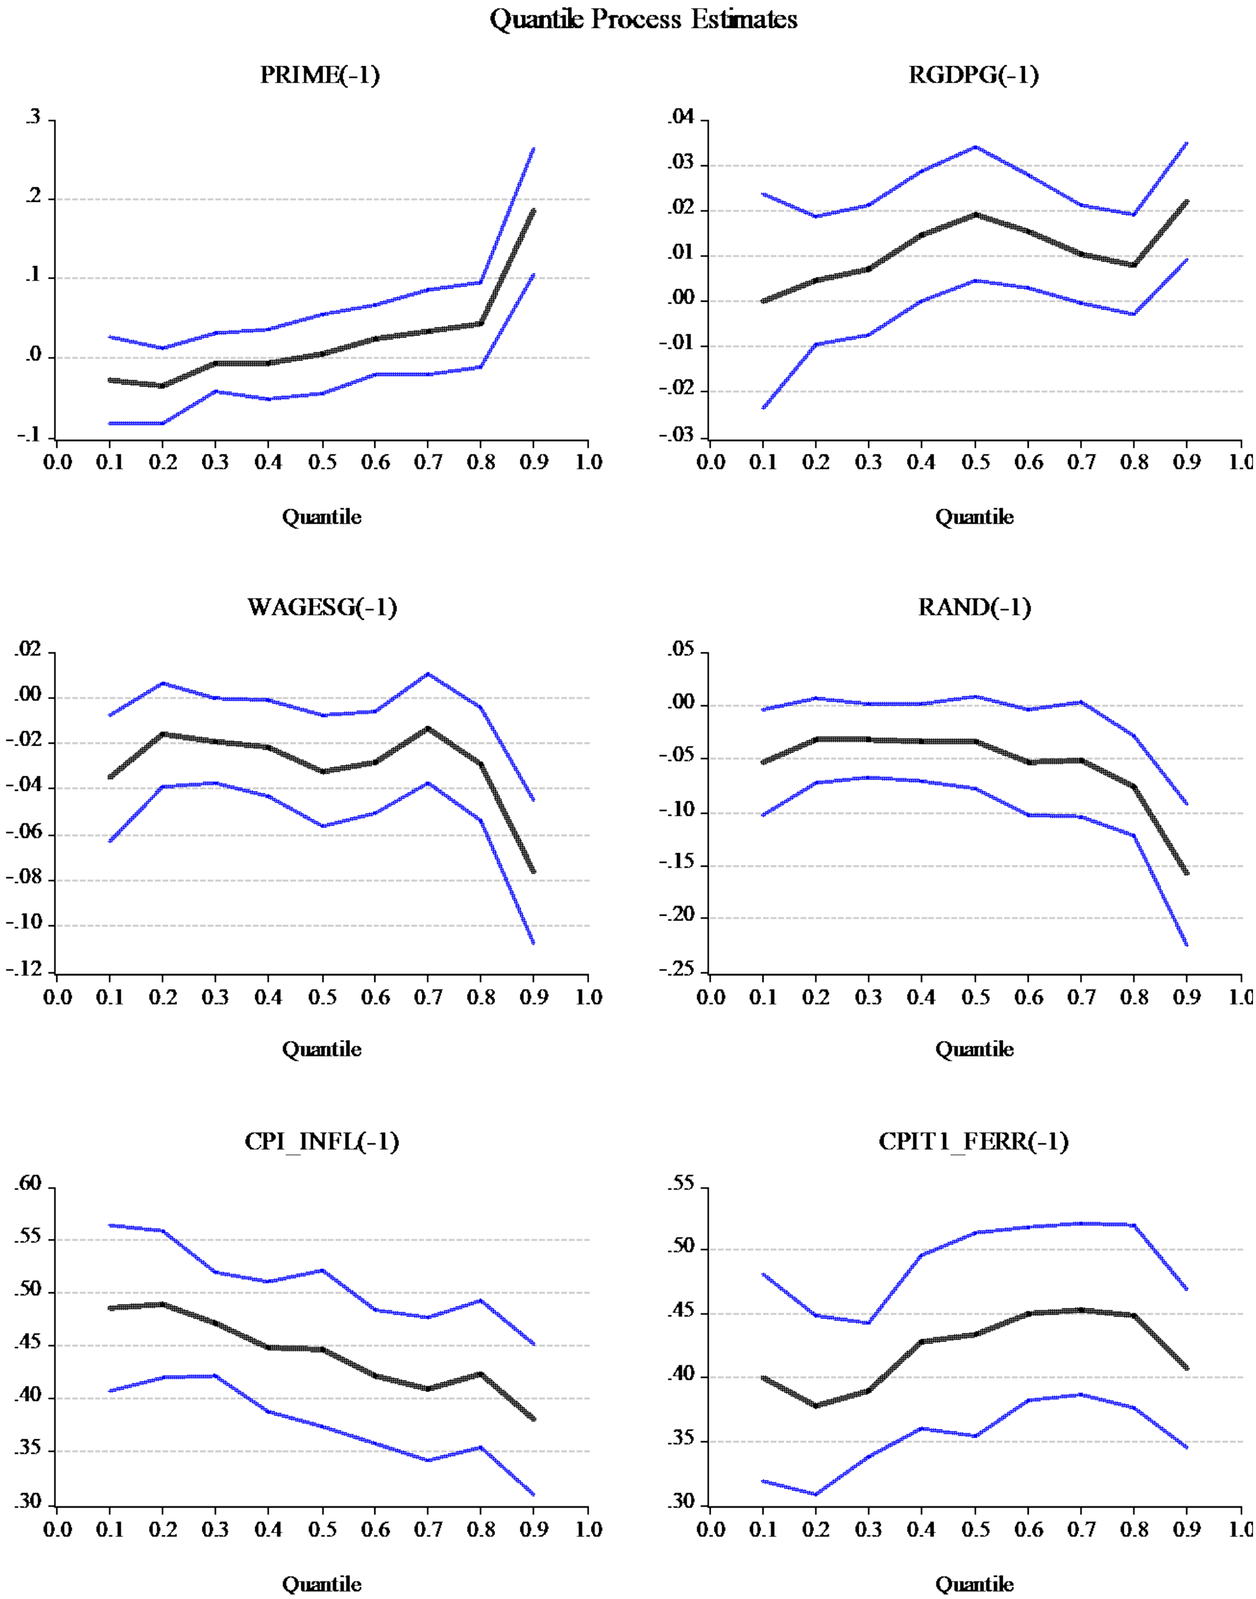

Figure 8b: Quantile regressions (five-year horizon): financial analysts

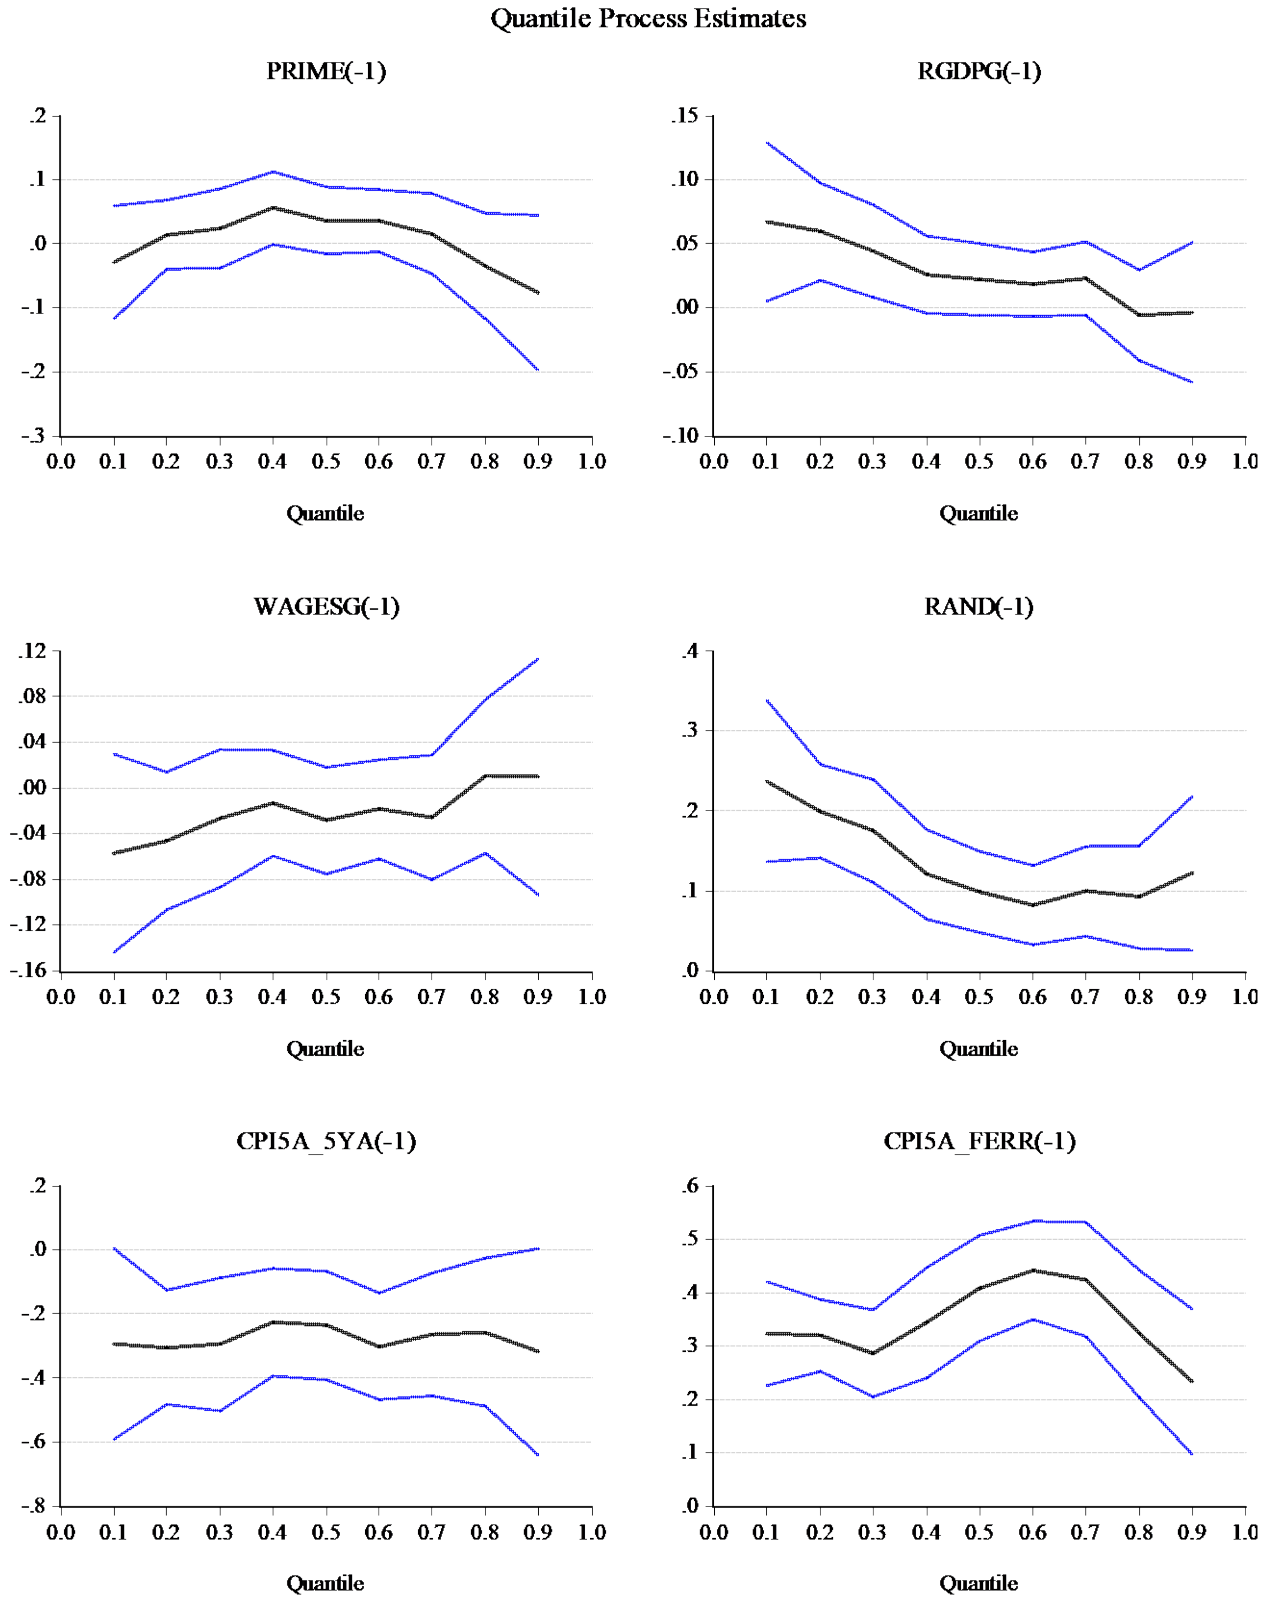

Figure 9a: Quantile regressions (one-year horizon): business sector

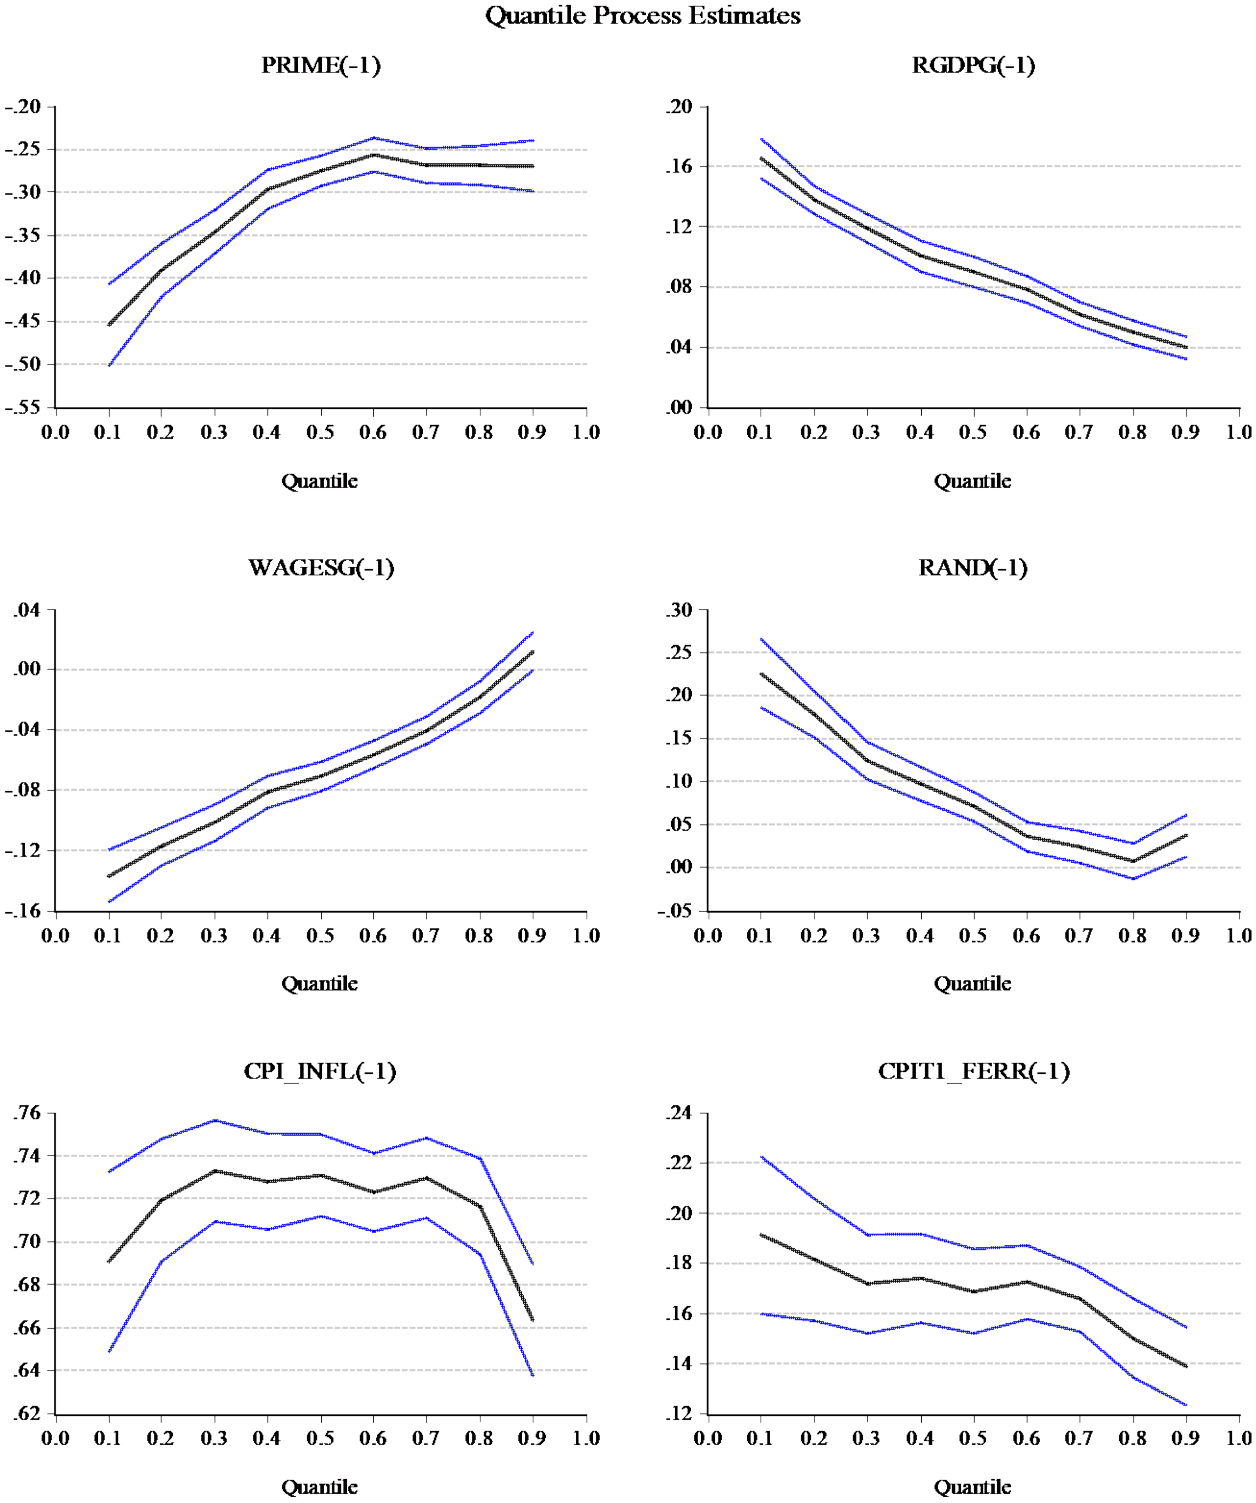

Figure 9b: Quantile regression (five-year horizon): business sector

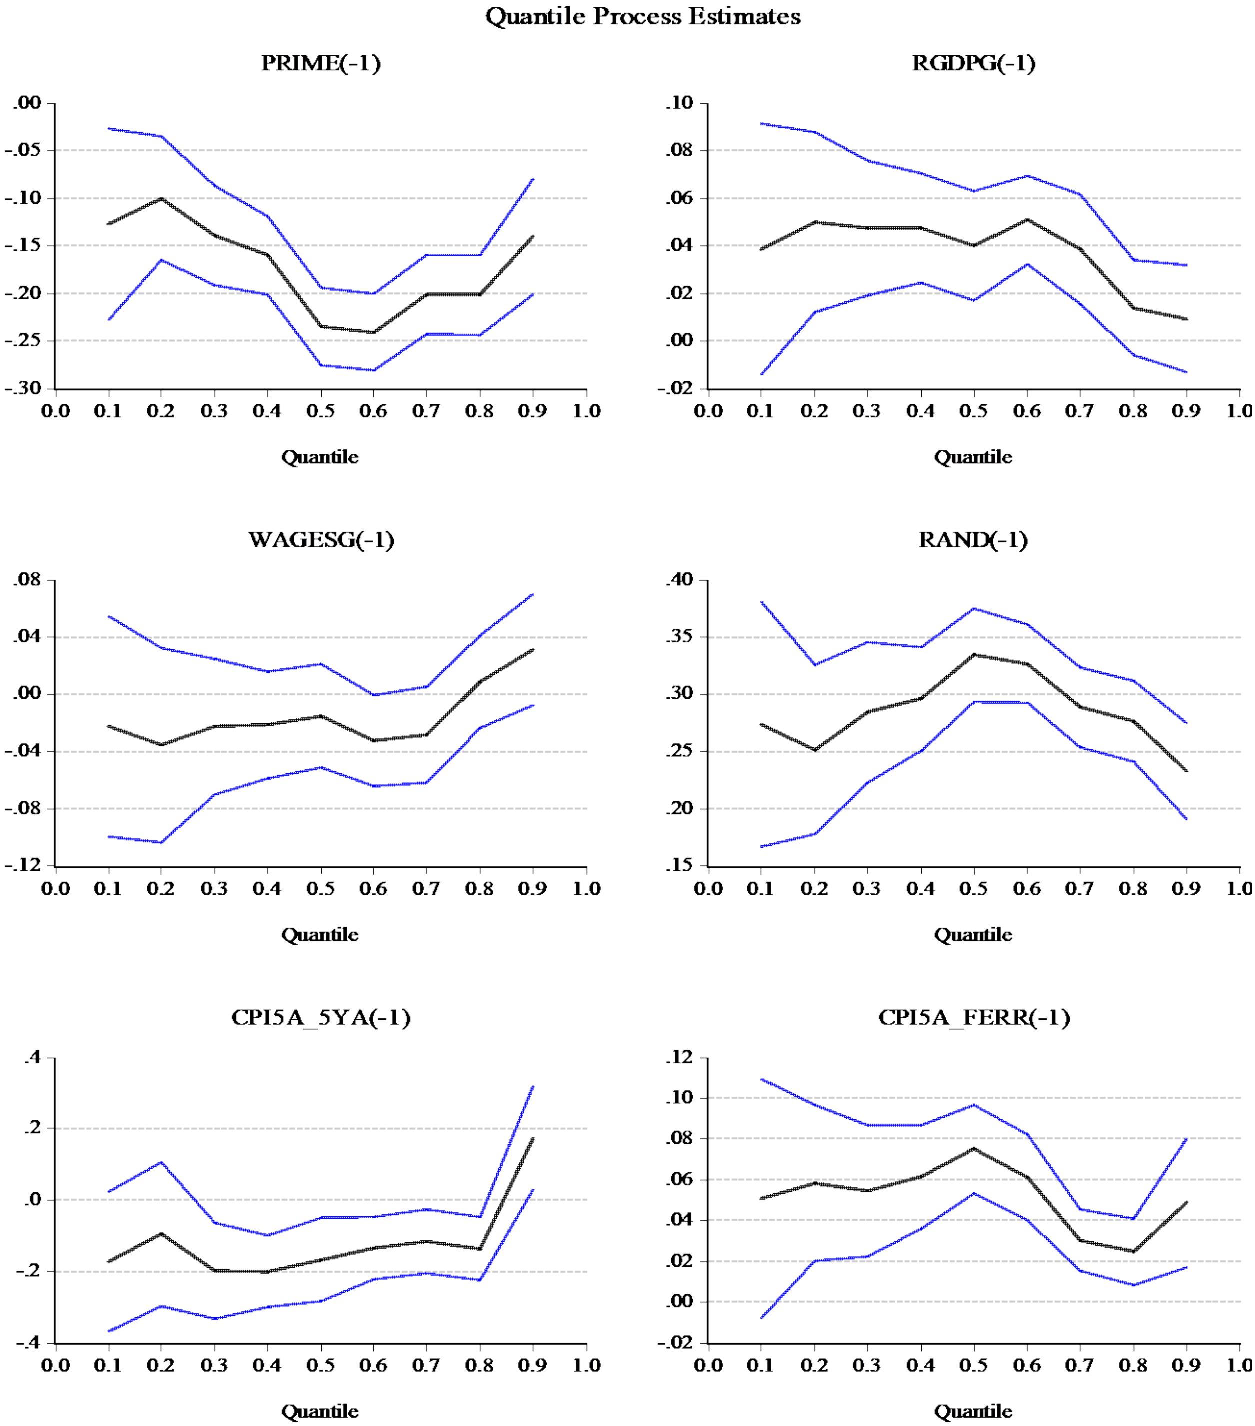

Differences between the two groups are striking. Consider Figures 8a and 9a, which investigate how the six selected determinants impact forecast errors in the FA and B groups, respectively. First, the FA group's forecast errors rise when the prime interest rate increases at the higher quantiles, that is, among the group of respondents that make the largest forecast errors. Otherwise, there is essentially no impact on forecast errors. Turning to the B group, higher interest rates reduce forecast errors at all quantiles. However, the size of the response is much larger at the lower quantiles than at higher quantiles. These results suggest that interest rate developments impact forecast performance so that monetary tightening or loosening does influence expectations. Interestingly, at the one-year horizon, GDP growth changes have no effect on forecast error performance in the case of FAs. In contrast, a rise in economic growth generates higher forecast errors for the one-year horizon relative to the five-year horizon. A plausible explanation is that better economic conditions perhaps reduce the incentive to improve inflation forecasts.

Differences between the FA and B groups are also stark when the remaining variables are considered. Hence, lagged wage growth reduces forecast errors in the FA group. The same variable also generates lower forecast errors at all quantiles for the B respondents, but the size and persistence of the effect is considerably larger. Clearly, the B group is highly responsive to wage growth – as one would expect. The level of the rand has little impact on forecast performance in the FA group except at the highest quantiles, where the poorest forecasts reside (i.e. at the 0.8 and 0.9 quantiles). A depreciating rand is a wake-up call that leads to lower forecast errors. In contrast, a rise in the rand raises forecast errors at almost all quantiles of the B group, although the impact is smallest among the forecasters that produce the highest forecast errors. If a rise in the rand is akin to a positive inflation shock, this result fits well with the estimates based on local projections discussed above.

The bottom two figures show the responses of forecast errors to past inflation and past forecast errors. Here the results for the FA and B groups are more similar. In both groups higher inflation raises forecast errors, contrary to the results based on local projections. This is suggestive of the ‘wisdom of the crowds’ argument since it appears that, once aggregated, forecast errors improve when inflation is rising (ie. the group collectively forecasts better than individuals do). The impact of higher inflation is considerably larger among the B group of forecasters and persists at all quantiles, whereas the same response drops significantly at the higher quantiles for the FA group of forecasters. Furthermore, and this result mirrors the ones shown in Figures 4 and 6 especially, there is considerable persistence in forecast errors although the degree of persistence is higher at all quantiles for the FA group.

Finally, we consider the quantile regression estimates for the longer-term inflation forecasts in Figures 8b and 9b. The first result of note is that there are far fewer differences between the FA and B groups in terms of the response to various sources of change in economic fundamentals on forecast errors. Nevertheless, a few differences do remain. First, at the longer forecast horizon, interest rates as measured by the prime rate have no impact on forecast errors in the FA group, while – as was the case at the one-year-ahead forecast horizon – a rise in the prime rate reduces forecast errors of the B group.[24] Moving on to the other determinants shown in Figures 8b and 9b, the responses to real GDP growth, wage growth, the rand, lagged inflation and lagged forecast errors are comparable in the two groups. That said, the responses to the rand and lagged forecast errors are larger for the B respondents than their professional counterparts.

Other than differences across forecast horizons and the types of individuals asked to provide forecasts, arguably the most notable result is that forecasts by both groups could have been substantially improved if the immediate history of all the variables shown had been incorporated into their inflation forecasts. This suggests that they are not fully incorporating the most recent historical data. We cannot, however, determine the reasons for this. The respondents may generally be inattentive to the new information, the information may be sticky (i.e. takes time to reach decision-makers) or the respondents may not have a good understanding of relationships between these macroeconomic factors. It is notable that both groups are asked to forecast inflation as well as other sets of macroeconomic variables, but they are only giving priming in the inflation forecast question, not when they forecast the other macro variables. This priming may override the respondents' reasoning about the macroeconomic relationships between the variables being forecast.

We conclude by summarising our results and suggesting some extensions to the existing research.

[24] As discussed above, long-term interest rate changes do influence the FA group's forecast errors.

4. Conclusions, survey design and suggestions for future research

While motivations for incorporating expectations in economic analysis have been widely accepted for decades, in practice this has required explicit theory about expectations and decisions about how to model them. These details have proved challenging. If expectations are not rational, how do they behave across groups and across time? In this paper, we have added to this literature, relying on individual-level data of inflation expectations of firms and financial analysts in South Africa.

In line with traditional tests of rationality, we began by estimating the forecast errors for both the FA and B groups surveyed by the BER, at each forecast horizon. From our preliminary data analysis, we identified five stylised facts. These suggest that, over time, the forecast errors of both groups do not behave as the strict rational expectations hypothesis would predict. Within the context of our cautionary warning about how the term 'rationality' has evolved, we therefore stated that the rationality of both groups is limited in the sense that they do not incorporate the available information efficiently. All we can say is that the different groups surveyed do display evidence of being forward-looking. However, until a metric can be developed that permits us to quantify the degree to which forecasters are forward-looking, and how this changes over time, the best strategy for a central bank is to be transparent, clear and, where possible, reduce the information frictions facing the public that contribute to increasing forecast errors.

We then asked if there are biases in these forecast errors and explored some of the potential sources of any bias. To do this, we extended the standard test of rationality, using the Rossi and Sekhposyan (2016) test on the aggregate data, which allows structural breaks in the rationality – that is, it allows us to see if the rationality changes over time. Our results suggest that the failures of rationality across the two groups are not just a matter of degree. They are different in character and suggest that the two groups are attentive to different information and perhaps sensitive to different incentives.

Dividing respondents into quartiles based on the degree of optimism or pessimism reflected by their forecast errors, we next estimated a series of VARs. Generally, the results reflect some difference in the reactions of respondents in different groups (across and within the FA and B groups), particularly within the forecast errors in response to an inflation shock. The impulse response function using the five-years-ahead average inflation shows no notable relationship between inflation and forecast errors for the FA or B groups. We conclude that something is driving the forecasts that the VARs are unable to identify.

Next, we analysed the disaggregated (individual-level) data, using quantile regressions, to determine whether a series of macroeconomic variables could have improved the forecasts. Again, we also considered different responses between optimists and pessimists. We find that for both the FA and B groups, greater attention to past forecast errors as well as some macroeconomic factors (CPI especially) would have improved their forecasts. In the case of the FA group, forecasts could also have been improved by incorporating more information about the long-term interest rate, whereas the B group could have improved its forecasts by paying more attention to the prime rate.

Finally, socio-economic factors are included to investigate whether any of these contribute to our understanding of forecast errors. Counterintuitively, we find that respondents from large FAs and large firmsB both make poorer forecasts than those from medium and small institutions.

For future research, it would be valuable to explore the different characterisation of the patterns in rationality as displayed using the Rossi and Sekhposyan (2016) test. For example, persistent rising forecast errors at the average five-years-ahead horizon that appear for the FAs but not the B group from 2017 are notable. It would also be worth exploring the reason for the counterintuitive result that large firms make poorer forecasts than medium and small institutions. Another implication of our findings relates to the BER survey itself. While continuity in surveys is desirable, additional analysis should lead to suggestions to improve the survey to assist policymakers in uncovering sources of bias or misinformation in the preparation of forecasts, without losing the valuable data already available. Other policy implications might well generate new strategies for how the SARB communicates its outlook to firms and the wider public.

References

Ashiya, M. 2009. 'Strategic bias and professional affiliations of macroeconomic forecasters'. *Journal of Forecasting* 28: 120–130.

Ashiya, M and Doi, T. 2001. 'Herd behavior of Japanese economists'. *Journal of Economic Behavior and Organization* 46(3): 343–346.

Bernanke, B. 2015. *The courage to act*. New York: W.W. Norton & Co.

Bordalo, P, Gennaioli, N and Shleifer, A. 2022. 'Overreaction and diagnostic expectations in macroeconomics'. *Journal of Economic Perspectives* 36 (summer): 223–244.

Cagan, P. 1956. 'The monetary dynamics of hyperinflation.' In *Studies in the quantity theory of money*, edited by M. Friedman. Chicago: University of Chicago Press. 25–117.

Coibion, O, Gorodnichenko, Y and Kumar, S. 2018. 'How do firms form their expectations? New survey evidence'. *American Economic Review* 108(9): 2 671–2 713.

Coibion, O, Gorodnichenko, Y, Kumar, S and Pedemonte M. 2020. 'Inflation expectations as a policy tool?' *Journal of International Economics* 124 (May): 103297.

Curtin, R T. 2019. *Consumer expectations: micro foundations and macro impact*. Cambridge: Cambridge University Press.

De Grauwe, P and Ji, Y. 2019. *Behavioural macroeconomics: theory and policy*. Oxford: Oxford University Press.

Friedman, M. 1968. 'The role of monetary policy'. *American Economic Review* 58(1): 1–17.

Jordà, O. 2005. ‘Estimation and inference of impulse responses by local projections’. *American Economic Review* 95(1): 161–182.

Kahneman, D and Tversky, A. 1979. ‘Prospect theory: an analysis of decision under risk’. *Econometrica* 47(2): 263–292.

Keynes, J M. 1936. *The general theory of employment, interest and money*. 1st ed. London: Macmillan.

Kydland, F E and Prescott, E C. 1977. ‘Rules rather than discretion: the inconsistency of optimal plans’. *Journal of Political Economy* 85(3): 473–492.

Lucas, R E. 1972. ‘Expectations and the neutrality of money’. *Journal of Economic Theory* 4(2): 103–124.

Lucas, R E. 1976. ‘Econometric policy evaluation: a critique’. *The Phillips curve and labor markets*, edited by K Brunner and A H Meltzer. North-Holland. 19–46.

Mincer, J and Zarnowitz, V. 1969. ‘The evaluation of economic forecasts’. In *Economic forecasts and expectations: analysis of forecasting behavior and performance* (out of print), edited by J. Mincer. *National Bureau of Economic Research.* 3–46.

Muth, J F. 1961. ‘Rational expectations and the theory of price movements’. *Econometrica* 29(3): 315–335.

Pesaran, H M. 1987. ‘Global and partial non-nested hypotheses and asymptotic local power.’ *Economic Theory* 3(1): 69–97.

Phelps, E S. 1967. ‘Phillips curves, expectations of inflation and optimal employment over time’. *Economica* 34(3): 254–281.

Reid, M and Siklos, P. 2022. ‘How firms and experts view the Phillips curve: evidence from individual and aggregate data from South Africa’. *Emerging Markets Finance and Trade* 58(12): 3355–3376.

Rossi, B and Sekhposyan, T. 2016. ‘Forecast rationality tests in the presence of instabilities, with applications to Federal Reserve and survey forecasts’. CEPR discussion papers 11391.

Rossi, B and Soupré, M. 2017. ‘Implementing tests for forecast evaluation in the presence of instabilities’. *Stata Journal* 17(4): 850–865.

Rudd, J B. 2021. ‘Why do we think that inflation expectations matter for inflation? (and should we?)’ *Finance and Economics Discussion Series 2021-062*, 1–27.

Sargent, T J and Wallace, N. 1976. ‘Rational expectations and the theory of economic policy’. *Journal of Monetary Economics* 2(2): 169–183.

Simon, H A. 1959. ‘Theories of decision-making in economics and behavioural science’. *American Economic Review* 49(3): 253–283.

Svensson, L. 1997. ‘Inflation forecast targeting: implementing and monitoring inflation targets’. *European Economic Review* 41(6): 1111–1146.

Tarullo, D K. 2017. ‘Monetary policy without a working theory of inflation’. Hutchins Center Working Paper #33 at Brookings.

Weber, M, D'Acunto, F, Gorodnichenko, Y and Coibion, O. 2022. ‘The subjective inflation expectations of households and firms: measurement, determinants, and implications’. *Journal of Economic Perspectives* 36(3): 157–84.

West, K D and McCracken, M W. 1998. ‘Regression-based tests of predictive ability’. *International Economic Review* 39(4): 817–840.

Woodford, M. 2005. ‘Central bank communication and policy effectiveness’. Proceedings – Economic Policy Symposium – Jackson Hole, Federal Reserve Bank of Kansas City, August 25–27, 399–474.