From deliberate sample to representative sample: pilot study for the BER inflation expectations survey

Monique Reid, Dieter von Fintel and Anis Foresto

ERSA working paper 918

September 2025

The views expressed are those of the author(s) and do not necessarily represent those of the funder, ERSA or the author's affiliated institution(s). ERSA shall not be liable to any person for inaccurate information or opinions contained herein.

From deliberate sample to representative sample: pilot study for the BER inflation expectations survey

Abstract

With the adoption of inflation targeting in South Africa in 2000, the Bureau for Economic Research (BER) began to collect inflation expectations survey data on behalf of the South African Reserve Bank. This respected survey is rich by international standards and has contributed valuable insights to policy, academic and private sector analyses. International trends towards greater reliance on microdata within macroeconomics are, however, placing slightly different demands on the survey, and access to complementary datasets has offered new opportunities to enhance it. In this pilot study, we link the inflation expectations survey data of firms to a spatial tax panel dataset. We investigate whether the survey sample adequately represents the structure of the South African economy and offer a series of survey weights to be added to the micro dataset. The results show that the BER has maintained an adequate level of representativity over the life of the survey, but we recommend that sample weights be estimated periodically to ensure that representativity is ensured institutionally. The sample weights can also support targeted recruitment in future. Finally, through careful documentation we hope to enable other researchers to pursue questions that benefit from linking the inflation expectations data with other datasets.

JEL classification: E31, D84, C81, C83

Keywords: Inflation expectations, business surveys, administrative data

1. Introduction

The importance of inflation expectations for explaining the dynamics of inflation is not uncontested (Rudd 2022), but it remains the dominant view of the mainstream that expectations are central to the process of inflation (Adrian 2023). Inflation targeting is designed around the premise that clear and credible communication should be used to anchor the inflation expectations of the public, improving the efficiency of monetary policy and reducing the cost of maintaining low and stable inflation.[1]

With the adoption of inflation targeting in South Africa in 2000, the Bureau for Economic Research (BER) was commissioned by the South African Reserve Bank (SARB) to collect inflation expectations survey data and provide information viewed as central to operationalising the monetary policy framework. Acknowledging from the outset that the differences in how distinct groups within society collect information and set prices may be economically significant, it was decided that inflation expectations would be collected for four separate social groups – financial analysts, the business sector, trade unions and households.

While inflation expectations surveys of financial specialists (financial analysts, professional forecasters) and households are common internationally, comparable surveys of firms or trade unions are rare. In this paper, we focus on the inflation expectations of firms. This choice is in line with notable international interest in this group, supported by the argument that this group may capture important information about the price-setting process in an economy.[2] South Africa is privileged to have quarterly data for firm-level inflation expectations for over two decades, unlike most other countries. This South African firm-level (business

---

[1] This cost is typically represented by the sacrifice ratio.

[2] A number of papers argue that firms offer a lot of promise for capturing the price-setting behaviour that drives inflation (Candia, Coibion and Gorodnichenko 2021). Coibion and Gorodnichenko (2015) argue that household data (which they show is more like firm-level expectations than that of professional forecasters) can be used to explain the missing disinflation/inflation experience in the United States (US). Using disaggregated household data for the US, Binder (2015) also finds that individuals with certain characteristics (high-income, highly educated males) play a larger role in inflation dynamics than do the expectations of other groups of consumers or of professional forecasters.

Questions about which group of people has the most impact on inflation or on the inflation expectations of others are valuable and should be answered with the support of both empirical analysis and theory. However, answering these questions is not necessary to justify the addition of weights to the BER firm-level survey. If the case can be made that the inflation expectations of firms are worth measuring to learn more about the process of inflation, then any measures to ensure the continued quality of the data are justified.

The South African literature that provides insight into the heterogeneity of inflation expectations, the relative rationality of the expectations of different groups and the extent to which information diffuses from one group to the next is less extensive than in many advanced economies. Readers who wish to learn more about this literature are encouraged to begin with Crowther-Ehlers (2019) and references therein.

sector) data is also rich in that the survey respondents are asked to forecast inflation for different horizons and to forecast other macroeconomic variables in the same survey.[3]

There is potential to enhance the survey more than two decades after its launch. In this paper, we conduct a pilot study to test a proposal aimed at sharpening the extent to which the firm (business sector) category's survey sample accurately represents the population of firms in the South African economy. This should be viewed as part of a natural progression of research in line with international developments (Candia, Coibion and Gorodnichenko 2022).[4] Macroeconomic researchers and policymakers are currently far more focused than before on asking questions using disaggregated microeconomic data (Israël and Tisso 2021), necessitating greater accuracy of subcomponents of the data. The availability of administrative tax data for use in research, well over a decade after the BER survey was first designed, also enables this next step.

Since 2000, the BER has used deliberate sampling for the firm-level inflation expectations survey, with the aim of creating a panel that is roughly representative of the population of firms in South Africa. The BER has not had access to formal information about the composition of the economy (the Business Registrar is not published by Statistics South Africa), and the initial aim of the survey was to track the aggregate over time rather than to allow analysis of disaggregated data. Recently, administrative data has become available that allows us to formally increase our confidence about the extent to which the sample is representative of the structure of the South African economy, as well as to contribute to maintaining this representativity in the face of declining response rates. While it is obviously impossible to recollect or reconstruct the historical sample to be representative of the sectoral and geographic composition of the economy, administrative data can be used to add weights to the historical sample. Specifically, we create post-stratification weights for the deliberate sample to better reflect each firm's contribution to the broader economy, and thus their contribution to the inflation expectations index. Auxiliary administrative data from the South African Revenue Service (SARS) are used to represent formal sector firms at the level of firm size, industrial sector and municipality. The BER survey data are reweighted to reflect this

---

3 Forecasts at different horizons allow researchers to capture the views of survey respondents about the dynamics of inflation into the future, while the collection of forecasts for other macroeconomic variables allows researchers to investigate how the forecasters likely reason about the interaction between these macroeconomic variables and inflation into the future.

4 For further details on the design and methodology of inflation expectations surveys globally, see Candia, Coibion and Gorodnichenko (2022: 35–39).

composition. Reweighting cannot correct all biases that arise in the original survey design, however, such as unit non-response.

This pilot study uses a spatial tax panel dataset (administrative data) to determine the characteristics of the South African economy along a particular set of dimensions. Three firm-level characteristics were selected to link the spatial tax panel data to the BER inflation expectations data. Each characteristic posed practical challenges, and our pragmatic choices regarding each are presented as transparently as possible. Once satisfied that the three firm descriptors were accurate, we constructed inverse propensity weights to reweight the BER data. Practically, no adjustment to the data collection itself is necessary, as a column is simply added to the dataset containing a weighting for each observation. To increase the confidence with which results using the data can be generalised, these weights can then be applied to the data collected by the BER before interpretation.

The contributions of this paper are threefold. Firstly, we prepared the BER dataset so that it is comparable to the tax panel data along three dimensions – sector, firm size and region. This involved creating a ‘region’ descriptor with the assistance of the BER and testing its ability to locate firms accurately within administrative geographic regions. We also needed to convert the industrial classification codes from the SIC 5 codes used in the BER survey to the SIC 7 codes used in the tax panel data. The choice to work with the newer SIC 7 classification was guided by future versions of administrative data that will be used to recalibrate the survey data and by the fact that over time this will become the dominant version. Beyond the construction of sample weights in this paper, the preparation of these variables can also contribute to many applied questions that benefit from linking the BER data to the spatial tax panel data or other datasets. The second contribution of this paper is to use these three firm-descriptor variables to construct sample weights to increase our confidence in the degree to which the BER data are representative of the structure of the South African economy. The third contribution is a comparison of the results of the deliberate and post-stratified (reweighted) data, providing some empirical evidence for the accuracy of the sampling in the BER survey and evaluating the extent to which the survey data can be used for analyses that rely on the disaggregated data. The weights could also offer insights into sample design, with the potential to draw firms from a sampling frame maintained by StatsSA and/or SARS and provide information about whether other relevant information could be collected in the survey to enable micro-level analysis. The results of our analysis suggest that the BER’s deliberate sample has performed well over the period 2014 to 2021. While the weighted means that do not adjust for *item* non-response lie consistently below the unweighted means, the difference between the two is not statistically significant for any of the horizons. Taking item non-response into account, however, brings the weighted and unweighted means very close to each other. While this is testament to the BER's knowledge of the South African economy, it is always better to ensure that the methodology guards against slippage caused by continual changes in the composition of the economy, sustained institutional knowledge required for maintaining the survey, or the impact of declining response rates.

2. The BER inflation expectations survey data

2.1 Emerging demands and new opportunities

The BER inflation expectations data have been collected and published by the BER on behalf of the SARB since the third quarter of 2000, surveying four different stakeholder groups in the economy. The survey is widely respected and closely monitored by both the SARB and the private sector, featuring regularly in the SARB's formal communications (e.g. monetary policy statements and monetary policy reviews) and in news articles and reports by private sector financial commentators. The data have also enabled numerous deeper analyses of the South African economy, contributing to studies of forecast rationality, of the degree to which inflation expectations are anchored, and of Phillips curves, evaluations of central bank communication and identification of the factors that drive inflation expectations.[5]

The inflation expectations survey of firms, specifically, is a rich quarterly survey conducted continually since 2000. In addition to forecasts of inflation at fixed event horizons,[6] respondents are asked to forecast a range of other macroeconomic variables, as captured in Table 1. In addition, the dataset contains three firm-level characteristics about each respondent, as per Figure 1.

[5] See Reid and Siklos (2021) for a survey of the academic literature that uses this data.
[6] The five-year-ahead horizon was only added to the survey from 2012.

Table 1: Firm survey questions

Firms: survey questions
Inflation: current year, year ahead, two years ahead, five years ahead*
GDP: current year, year ahead
Interest rate (prime): current year, year ahead
Wages: current year, year ahead
ZAR/US$ exchange rate: current year, year ahead

Note: * The five-year-ahead inflation forecast question was only added to the survey in 2012.
Source: Own construction from BER survey data.

Figure 1: Firm characteristics collected as part of the BER inflation expectations survey

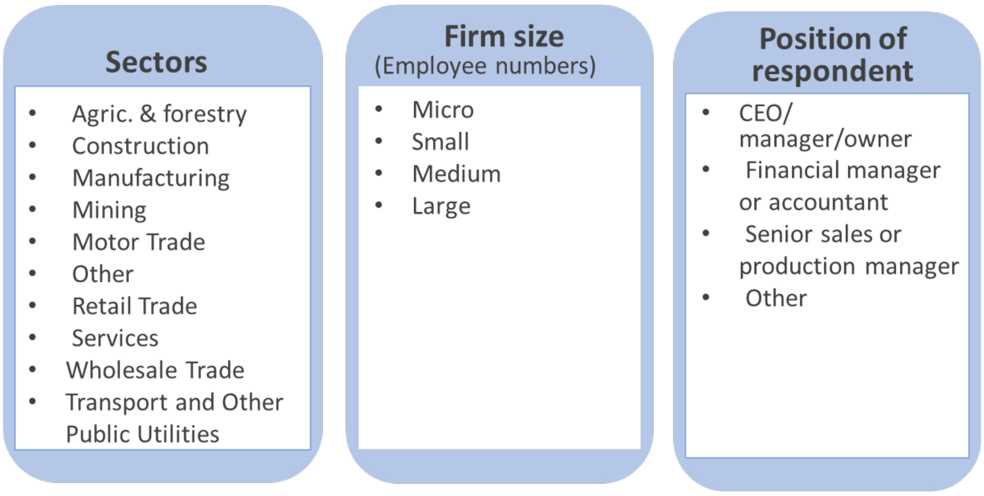

Source: Own construction from BER survey data.

Until the mid-2010s the vast majority of the users of the BER survey data used the data in aggregate form, with the exception of Ehlers and Steinbach (2007).[7] Thereafter, the increased availability of microeconomic data in South Africa and the international trend towards analysis using microeconomic data to answer macroeconomic questions has led to the desire and opportunity to do so in South Africa too.

2.2 The deliberate sample

The BER uses deliberate or purposive sampling[8] to ensure sufficient respondents are surveyed from each industrial sector. To ensure the data are representative of the structure of the South

[7] The increased use of disaggregated data is partly because researchers outside the SARB have only recently been granted access to the data (researchers not employed by the SARB can now apply for access), and partly because of the recent increased interest internationally in answering macroeconomic questions with disaggregated data.

[8] The BER used records of the Bureau for Market Research at Unisa to assemble a new panel of businesspeople in 2000 (Kershoff and Smit 2005). Over the years, Introye, Interactive Direct and own telephone directory internet searches were used to maintain the panel. Throughout, the BER strived to

African economy, it does not use a systematic sampling design (such as drawing from a sampling frame, stratification or weighting). The BER did not have access to formal information about the composition of the South African economy, because Statistics South Africa does not publish the Business Register. The primary shortcoming of deliberate sampling is that a sample design is not used, it is difficult to sample in this way without a list of available firms and when the response rate is not high. Furthermore, given that this is a panel, repeated interviewing of the remaining respondents also means that the sample becomes non-random and non-representative of the population over time. Firms that exit the sample do so non-randomly and are replaced by new deliberate samples.

It is important to balance the pursuit of representative sampling with the need to follow the sample longitudinally in order to track individual behaviour over time (enabling the estimation of Calvo parameters, for example).[9] While weighting helps to make samples representative when a non-random sample is used, weighting should be distinguished from sample design. This paper poses questions about sampling but does not fully answer them. We speculate about possibilities, but more thinking is needed to develop a future strategy, especially given that the sample sizes are declining and may be doing so selectively. While both sampling design and reweighting will likely become important in the future, this paper's focus is on reweighting.

The BER's objective is to estimate inflation expectations for the three institutional groups. In the case of the business sector, it was not originally the BER's intention to report the results reliably at the disaggregated level (such as in different sectors or regions) but rather to deliver a valid measurement of aggregate inflation expectations over time. Random selection is crucial for a once-off quantitative estimation of a statistic's level (e.g. a value of Rx, or y number of individuals), whereas a panel (even assembled and maintained in a deliberate sampling manner) is more appropriate for observing changes in a statistic over time, as long as sufficient cross-sectional variation is maintained by refreshing the sample and there are no significant changes to the biases that could result from potential compositional changes of the survey in

---

separate participants in its various surveys – so an inflation expectation survey participant would not also be recruited for the business tendency survey to avoid survey fatigue.

9 It is an open question how the longitudinal element should be maintained in the inflation expectations survey. Retrospective data suggest that only 15.6% of firms have been sampled repeatedly for up to 10 years, so achieving a truly long-run panel element in this context is challenging and may become even more so in light of declining sample sizes. This paper does not fully address questions about sampling. The focus is rather on whether existing samples can be reweighted to produce representative statistics, which can inform the design of future rounds. The reweighting exercise implicitly corrects for the selective entry and exit of firms that could compromise the reliability of aggregate measures of inflation expectations. Reweighting rebalances estimates to be representative despite selective attrition.

different periods. Compositional biases that remain constant over time do not necessarily compromise estimates of time trends unless the panel survey follows a completely unrepresentative part of the economy.

The recent availability of administrative tax data in South Africa, however, provides a way to gauge the composition of the economy along a set of firm-level characteristics. This can then be used to construct weights for the sample that confirm or improve cross-section level estimates. The question of whether any discrepancy between the weighted and unweighted data has an economically significant (sizeable) impact on the data and interpretations using this data can only be settled empirically.

If the user of the data is most interested in the changes in the aggregate level of expectations over time, then the weightings may have limited impact if the biases introduced by selecting the sample of firms also remain constant over time. The weightings are, however, likely to improve confidence in analyses using the disaggregated data and interpretations about the dynamics of this data that look beyond the first moment. For example, by monitoring how expectations vary across industries, we may find some indications of whether elevated expectations are becoming more broad-based. If a shock affects the exchange rate, it would be unsurprising to find a rise in short-horizon expectations in sectors exposed to the immediate effects.[10] If, however, expectations in other sectors that are isolated from the initial effects also begin to rise, then this would provide some evidence that expectations are becoming broad-based or entrenched (there are second-round effects) (Amaral, Kruger and Reid 2023).

Figure 2 presents the sample sizes of the BER dataset for 2014–2021 broken down by sector (using the SIC 5 codes, based on the international industrial classification, which we elaborate on in section 3.3 below). It is important to note that the BER has relied on its own industrial classification to create and interpret this dataset. The BER classification involves fewer categories, so it would not have been aiming to ensure that each of the categories in Figure 2 was represented. It is also worth noting that the international classification on which the SIC 5 codes are based is updated periodically, as the industrial composition of economies is changing over time – new industries emerge through innovation and changing demands as others disappear.

While we do not draw any strong conclusions from this simple graphical representation, Figure 2 does reveal that if we use the SIC codes to enable comparability with other South

[10] We do not take a firm position on the speed of the survey participants response. We are only commenting on a hypothesis that there is a heterogenous reaction. This kind of work enables us to test that sort of question in future research.

African surveys, some industries are under-represented. While under-sampling of homogeneous groups is of limited consequence, this is not the case for heterogeneous groups. Together with careful sample design (which is beyond the scope of this paper), sample weights will correct where needed to ensure that aggregates are representative. We can also see that the decreasing survey response rate (an international trend) exaggerates this challenge. Graphically, this figure suggests additional challenges when analysing components of the survey rather than the aggregate data alone.

Figure 2: Sectoral frequency in the BER data by year[11]

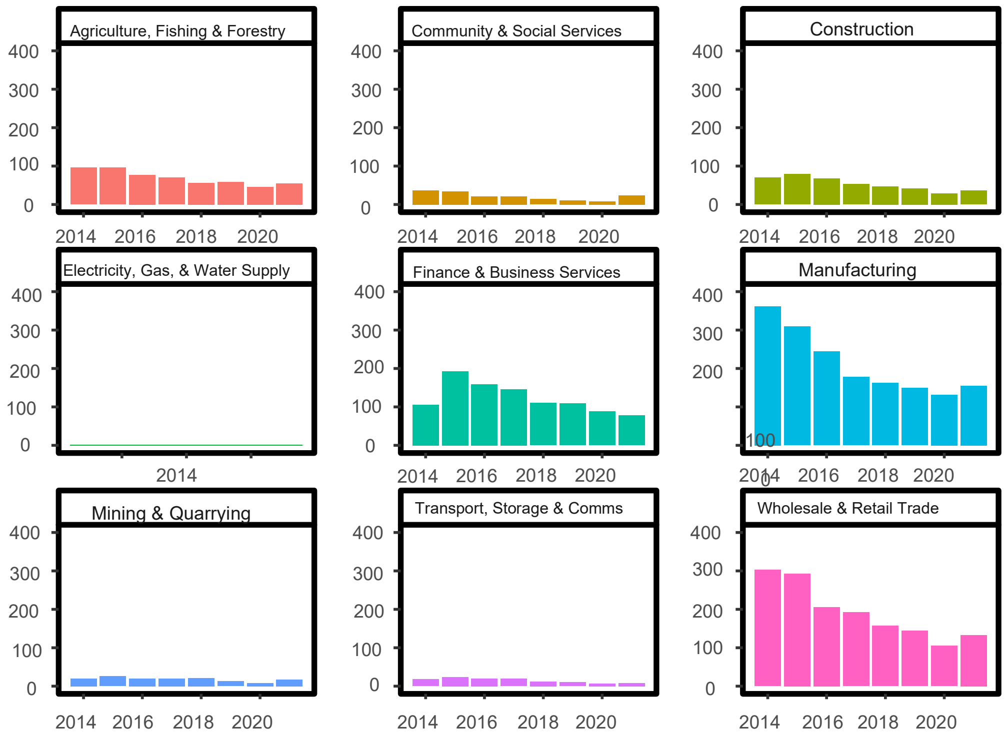

Source: Authors' calculations, based on BER data (2014–2021) classified according to SIC 5

3. Choice of datasets and points of contact between them

The starting point of this paper was to select a dataset that could enable us to characterise the structure of the population of firms in the South African economy. South African tax microdata, available from the South African National Treasury (NT) and SARS since 2014, is naturally the first option (Pieterse, Gavin and Kreuser 2018). However, researchers need to physically visit a secure venue in Pretoria (South Africa) to work with the data, which does create barriers to its use, particularly if researchers need to revisit the raw data more than once during the

[11] We include a figure on the sectoral composition of the tax data in the annexure.

study. It is also worth noting that the tax data are from firms that have submitted employee IRP5s and do not include the informal sector, self-employed people and sole proprietors. However, implicitly the BER survey does not target the informal sector or self-employed people, and its inflation expectations are intended to be representative of firms who employ workers, so the composition of both datasets is reasonably comparable.

In 2023, the Human Sciences Research Council (HSRC) launched a spatial tax panel dataset, which offers much freer access to tax data (Spatial Economic Activity Data 2023). The dataset contains administrative data from 2014 and is an aggregated version of the micro-firm and employee tax records in the formal sector, housed in the NT's secure lab. Key statistics were aggregated at municipal level and at sub-municipal level in metropolitan areas, providing anonymity for firms and employees but creating a time series of economic evidence with a geographic lens. This dataset has proven to be extremely useful for producing regular statistics about local labour market trends, inequality and spatial variations in economic activity in South Africa, where spatial inequalities define the structure of the economy (Harrison et al. 2023). The spatial tax panel data were therefore used in this study.

Having selected the spatial tax panel dataset ("the tax data") to link to the BER inflation expectation dataset ("the BER data"), the next step was to identify firm-level characteristics through which to match the two. The BER collects three characteristics for each firm respondent (Figure 1): the industrial sector in which it operates (*sector*); the size of the firm by number of employees (*firm size*); and the position of the respondent in the firm (*position*).[12] The first two are relevant for our objective of characterising the structure of the South African economy and can be linked with matching characteristics in the tax data. Weighting by firm size gives greater emphasis to responses from firms that employ larger shares of the population and are more labour-intensive, while weighting by sector ensures that all industrial sectors of the economy are adequately represented in the statistics.

We decided to try to include a third firm-level characteristic, '*region*', which is, by definition, available in the spatial panel. However, classifying survey respondents according to municipal boundaries required assistance from the BER because survey respondents' anonymity is guaranteed. Each of the firm-level characteristics used to link the datasets posed some level of challenge and at times required judgement calls. We aim to be as transparent as possible about these judgement calls in the discussions that follow.

---

[12] See Kershoff (2002) and Kershoff and Smit (2005) for further details about the BER survey.

i. Region

The *region* variable is part of the tax data but not formally part of the BER data. Much of the survey data records specific text addresses that can be located within municipal boundaries, but these addresses are not released to researchers for the sake of confidentiality. The BER assisted by releasing an additional column that indicates the municipality within which a firm is located. The choice to include *region* as a firm-level characteristic was motivated by the fact that it is a natural link given the structure of the spatial tax data and that regional differences across South Africa may inform inflation dynamics.

After considering various methods of adding the *region* characteristic, it was decided to use the postal code of survey respondent addresses to identify their municipality using Open Streetmap's Neonatim encoder. The BER guarantees the anonymity of the firms that take part in the survey (including anonymity from the SARB), so we required the BER's assistance to create this variable, which is time-consuming work.[13]

A second challenge was that the tax data recorded region by an alphabetic municipal code, which sometimes does not have a one-to-one match with the postal codes recorded in the BER records. While postal codes are informative for spatial analysis in South Africa, many span multiple standard administrative units, such as municipalities and provinces, complicating the choice of where to classify a particular address (Lombaard 2004).[14]

Finally, the postal codes of some respondents were missing. The BER collected the firms' physical addresses until 2021, because some firms still chose to have the surveys mailed to them through the postal system (Figure 3). The number of addresses collected then decreased suddenly over the pandemic period as most respondents elected to rather respond to the survey via email. Compared to the period between 2014 and 2021, the sudden decline in postal addresses in 2022 illustrates a change in how respondents chose to receive the survey questionnaire. We performed diagnostic tests on the samples with and without postal codes, including a t-test of the means for current-year inflation expectations and a Chi-squared test on the industrial distribution in each sample.

---

13 We acknowledge and are grateful for this contribution by the BER.
14 We are grateful to Andrew Nell for advice about the postal codes.

Figure 3: Proportion of respondents with postal codes

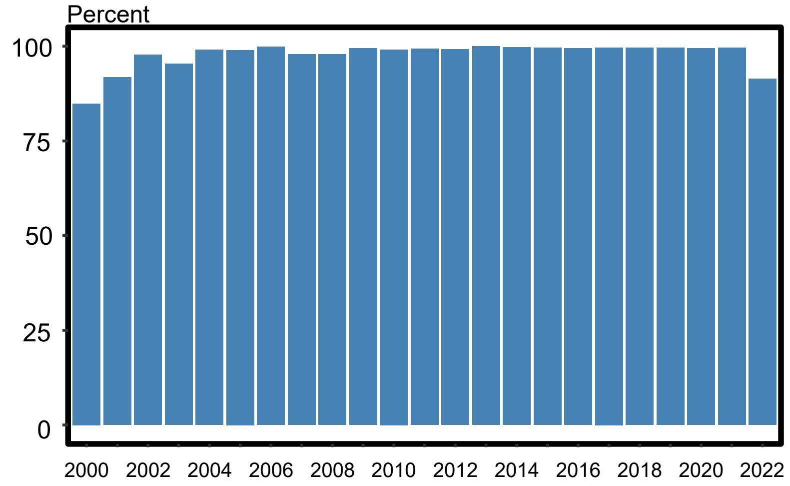

Source: Authors' calculations, based on BER data

We compared the inflation expectations of the part of the sample that contains the postal codes (Figures 4a, 4c, 4e and 4g) with the part of the sample that does not (Figures 4b, 4d, 4f and 4h) to assess the degree of bias that could be caused by omitting those firms without postal codes.[15] The figures illustrate that the distributions are marginally different at each horizon. The variation for the 2014 to 2021 sample without postal codes reduces significantly, suggesting the degree of missingness is less severe. Figures 4a and b: T0 inflation expectations with postal codes (left) and without (right)

Figures 4a and b: T0 inflation expectations with postal codes (left) and without (right)

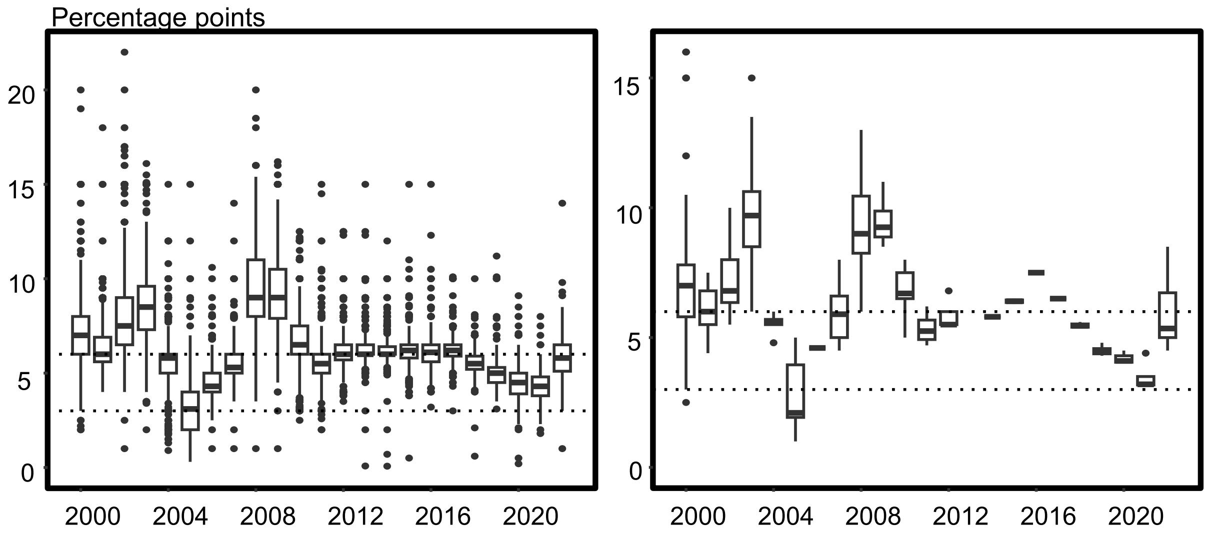

[15] The dotted horizontal lines in Figures 3a–h represent the upper and lower bounds of the inflation-targeting band. There is significant variation in both samples at all horizons from 2000 to 2009. From 2010 onwards, expectations at all horizons anchor near the upper bound, gradually falling towards the midpoint of the band from 2017.

Figures 4c and d: T+1 inflation expectations with postal codes (left) and without (right)

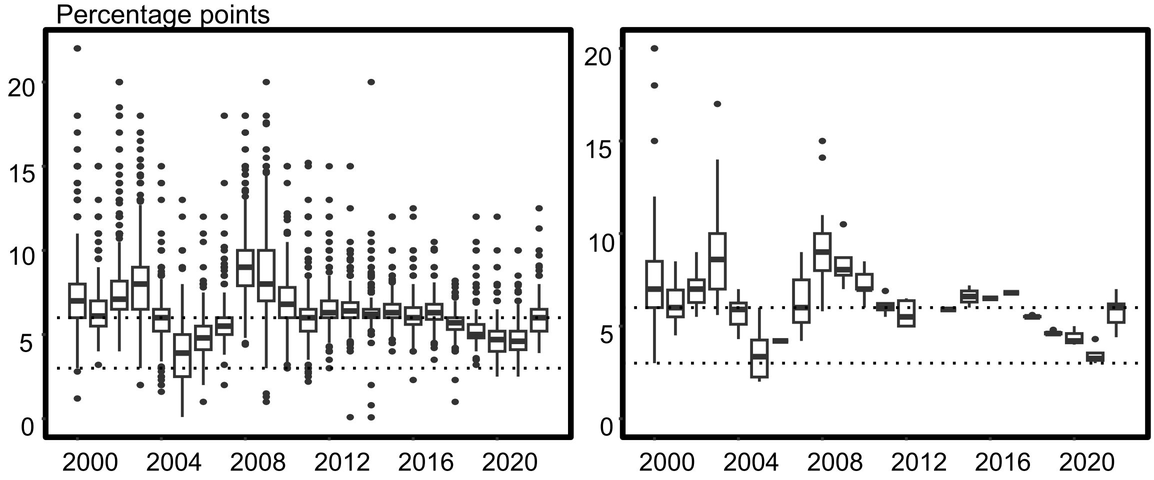

Figures 4e and f: T+2 inflation expectations with postal codes (left) and without (right)

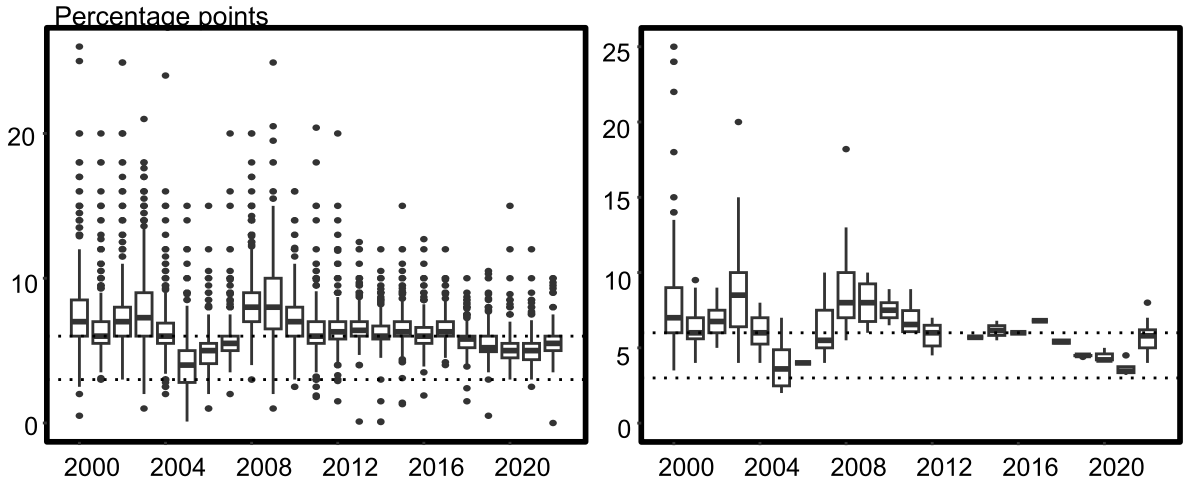

Figures 4g and h: T+5 inflation expectations with postal codes (left) and without (right)

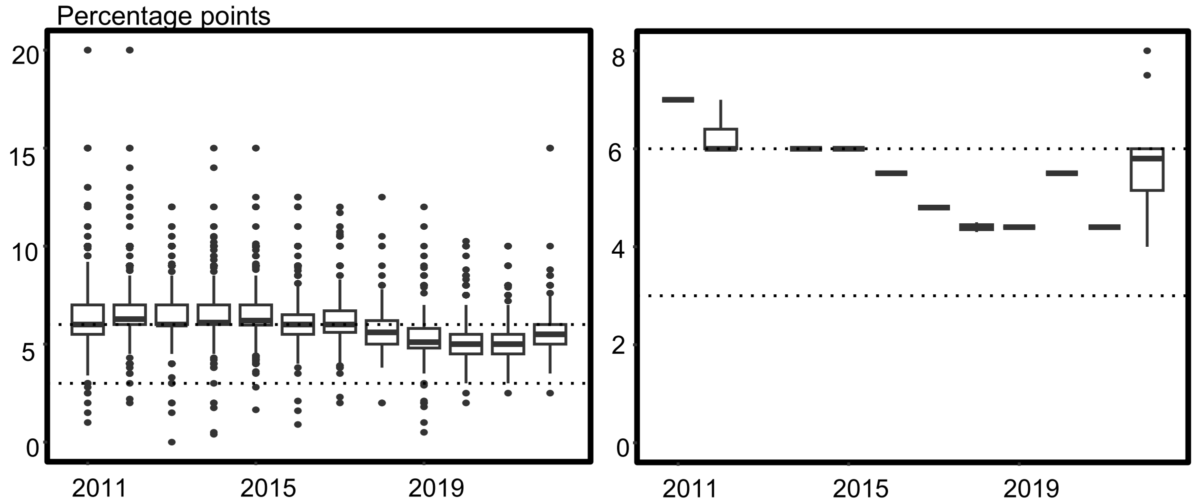

Source: BER inflation expectations data (2000–2022).
Note: Dotted lines are the upper and lower bounds of the target band.

The results of the formal tests are reported in Tables 2 and 3. In line with Figures 4a–h, the estimates of both t-tests and Chi-squared tests suggest that the distributions for the sample with postal codes are different from the sample without postal codes. This is evident in the overwhelming rejection of the null hypothesis in both tests. We are unable to use the responses without recorded postal codes, and that would therefore introduce systematic sample selection bias. However, the degree of missingness is generally very low, so excluding these observations is unlikely to introduce substantial bias. The bias is likely to be more pronounced in the earlier years of the sampling period and post-pandemic, when more surveys were electronically enumerated (Figure 3).

Table 2: T-test for mean difference in sample with postal codes vs without

Year	T-value	Year	T-value
2000	0.8685	2012	1.5267
2001	2.1401**	2015	-2.3220
2002	2.8147**	2017	-10.3545***
2003	-4.4623***	2018	1.7519
2004	0.5842	2019	4.5219
2005	1.2886	2020	1.2235
2007	-1.6072	2021	2.6534*
2008	0.0904	2022	0.2402
2009	-0.5893		
2010	-0.2801		
2011	0.5551		

Note: Results for 2006, 2013-2014 and 2016 were omitted from the table as there were too few observations in the Missing sample to conduct the test. Values with ***, **, * reject the null hypothesis at the 10, 5, and 1 % level, implying the means between the samples are different.
Source: Authors' calculations.

Table 3: Chi-squared test for industrial distribution in sample with postal codes vs without

Entire sample (2000–2022)	Windowed sample (2014–2021)
4419.1***	1242.4***

Note: Values with *** reject the null hypothesis that the distributions are the same at the 1% level.
Source: Authors' calculations.

To exclude the period when postal-code reporting tapered off and to match the availability of the tax records, we therefore focus on the sample period between 2014 and 2021. Over our chosen sample period, 22 observations are without postal codes. We then consider the sectoral composition of the observations with missing postal codes (Figure 5) over the 2014 to 2021 sample period. Most of the missing postal codes are from sector 30 (manufacture of food, beverage and tobacco products), with code 33[16] appearing only once. The number of missing postal codes for each year is low, so we conclude that the problem of missingness is not prominent in this part of the sample.

Figure 5: Sectoral frequency by year (without postal code)

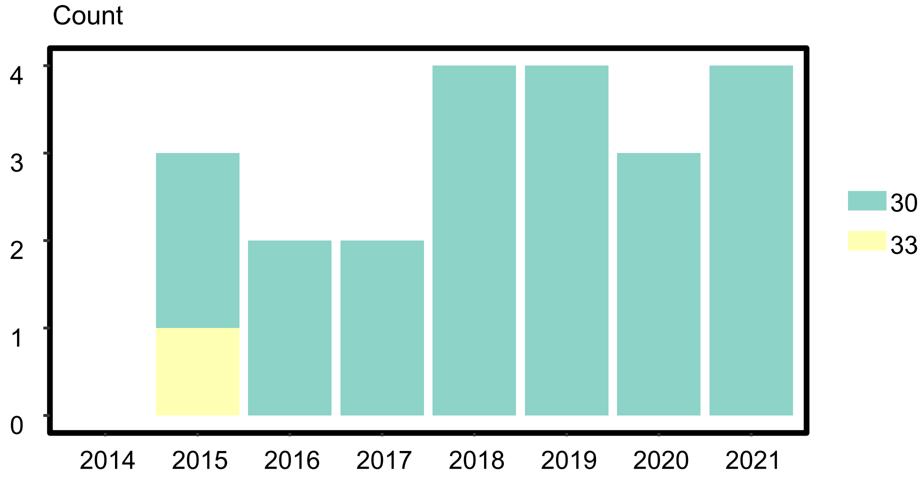

Source: Authors' calculations, based on BER data

Analysis beyond 2022 will require one of two adjustments. The ideal would be for the BER to collect information on respondents' municipalities regardless of whether physical address information is required or not. The second option would be to adjust sampling weights by the inverse of the propensity score of reporting a postal code or not – it is likely that characteristics such as firm size and industry can be used to generate propensity scores for reporting this information.

A third consideration in using region as a post-stratification variable is the so-called "headquarter bias" inherent in firms' address information. Both survey respondents and firms filing their taxes commonly use the registered address of the firm's headquarters. However, firms may also operate in stores, franchises and plants across the country. The distribution of a firm's operations is not captured in the single address reported in survey and administrative records, so economic activity in the large metropolitan areas where firms are more likely to register their operations may be easily overstated. While there is some evidence for headquarter bias in the tax records, we assume it has a similar geographic pattern in the survey data (Von

16 Manufacture of coke, refined petroleum products and nuclear fuel; manufacture of chemicals and chemical products; manufacture of rubber and plastic products.

Fintel 2023). Supposing this assumption is valid, a similar bias in both sets of records circumvents bias in estimating reweighting factors, as they affect the numerator and denominator of the weights in similar ways.

ii. Firm size by number of employees

The second firm characteristic used to link the two datasets is firm size, as measured by the number of people employed. For our purposes, firm turnover would have been a preferable indicator of firm size, as this would better capture the economic influence of the firm. In contrast, the number of people employed by a firm captures both firm size and the labour intensity of the industry. However, the BER does not record firm size by turnover in the survey.[17]

iii. Industrial category

The most consequential firm characteristic for the construction of these weights is probably the industrial category each firm survey respondent is in. The BER uses its own industry classification, which is obviously useful for its own applications, but it is difficult to compare or link this dataset with one that uses the Standard Industrial Classification (SIC). For a previous research application (Reid and Siklos, 2022), the BER generously added SIC 5 codes to their dataset and included information in their formal documentation to show how these two classifications compare (see Table 4).

Table 4: BER sectoral conversion table

Industry description	BER	SIC 5 (2-digit)
Agriculture	100	11
Mining	110	20
Manufacturing	130	30–39
Electricity & water	140	42
Construction	120	50
Wholesale, retail, motor, hotels, restaurants	150–180	61–64
Transport & communication	140	71–75
Finance, real estate & business services	180	82–88
Community, social & personal services	180–190	91–99

Note: *Standard Industrial Classification of all Economic Activities*, 5th edition, 1993, adapted for SA
Source: Bureau for Economic Research (2023)

[17] Some firms are likely to be less comfortable reporting turnover than number of employees, and the pressure to maintain survey response rates might cause hesitation about adding this question. Using a Likert scale for the answer would offer a middle ground.

To link the BER survey to the tax data we must overcome an additional hurdle, as the tax dataset uses the updated version (SIC 7) of this international code. This challenge applies to most other macro datasets in South Africa, as well as Statistics South Africa's survey data, as most still use the SIC 5 codes. Budlender and Ebrahim's (2020) concordance table guides researchers on how to compare the SIC 5 and SIC 7 codes, but they caution that the need for discretion remains.

Our use of the SIC 7 codes is limited to the one-digit or section level of disaggregation, while the SIC 5 codes supplied by the BER were at the two-digit division level. Conversion to a similar level of disaggregation in SIC 7 would require more discretion and judgement, which could result in uncertainty. We thus use the less disaggregated section level, as we have more confidence in it to correctly classify survey respondents.

To systematise the conversions necessary for this paper, we constructed a conversion table tailored to our purposes (Table 5). This table assisted us to make the classifications and makes our choices transparent. The table differs in two ways from Budlender and Ebrahim's (2020) concordance table. It is simplified by using only the section-level (one-digit) codes, and manual adjustments were made in sections where the match between SIC 5 and SIC 7 was less obvious.

Table 5: Industry conversion table

SIC 5 division (2-digit)	SIC 7 section (1-digit)
11	A
20	B
30–39	C
42	E
50	F
61–63	G
64	I
71–75	H
82–83	K
84	L
85	N
86	J
87–88	M
91	O
92	P
93	Q
94	E
95	S
96	R
99	S

Note: Adapted from Budlender and Ebrahim (2020)

SIC 7 codes are more disaggregated than SIC 5, so a SIC 5 code may be split into more than one SIC 7 code. Given that the BER recorded SIC 5 codes at a two-digit level, this information sometimes allowed us to place a firm into the correct SIC 7 division. Some categories remained challenging and are labelled "judgement calls".

Judgement calls

Only four SIC 5 divisions in the BER data required judgement calls after the concordance table (Budlender and Ebrahim 2020) had been used. Table 6 records these calls and indicates what information these calls are based on. The first column records the four SIC 5 divisions that required the judgement calls. The second column captures the percentage of the full BER sample that contains these observations, so we can see how big the potential 'problem' is. Figure 5 shows that about 10% of the BER sample is affected, so the impact is not negligible.

Table 6: Judgement calls for industrial classification

SIC 5: Level 2	BER sample %	Percentage of firms in each SIC 7 section (based on Budlender and Ebrahim (2020))						
84	0.53	F (3)	L (96)	N (1)				
86	0.23	C (4)	F (1)	G (3)	J (60)	N (2)	S (30)	
88	9.44	H (0.5)	J (0.6)	K (0.5)	M (61)	N (37)	Q (0.3)	R (0.1)
94	0.55	E (96)	N (4)					

Source: Own construction, based on BER data, Budlender and Ebrahim's (2020) concordance table and BER (2023).

The right-hand side of Table 6 (column three and its sub-components) shows the SIC 7 codes these judgement calls could fall into (based on Budlender and Ebrahim (2020)). The percentage of observations in each row that fall into a particular SIC 7 code is in parentheses next to the SIC 7 codes. So assuming there was no difficulty mapping the firms from SIC 5 to SIC 7 (using Budlender and Ebrahim (2020)) and the BER sampling process is adequately representative within sectors, column three shows, for instance, what percentage of the SIC 5 codes for 84 end up in each of the SIC 7 categories.

One or two cells are highlighted in each row of column three. If we were to guess which division a firm should fit into, selecting these highlighted categories would deliver the correct choice at least 90% of the time if the survey respondents are randomly selected from the population of firms. For example, the first row shows that 96% of SIC 5 division 84 ends up in SIC 7 section L. Consequently, if we were to allocate all of SIC 5 division 84 to SIC 7 section

L, we would be correct. Following this reasoning, SIC 5 divisions 84 and 94 are low risk. However, they only make up a very small proportion of the BER sample that requires judgement calls.

Two different SIC 7 sections likely correspond to the judgement call observations SIC 5 divisions 86 and 88, so it is difficult to categorise these observations into a single SIC 7 one-digit code. Of the two SIC 5 divisions, 88 is much more important, as it makes up by far the largest proportion of the judgement call observations in the BER dataset. We categorise all SIC 5 division 86 (computer and related activities) as SIC 7 section J (information and communication), but there is a higher chance of misclassification. As this SIC 5 division constitutes only 0.23% of the BER sample, this may require judgement calls on 13 of the 5 534 total observations, meaning we are likely to misclassify 5. From Table 6, we can see that SIC 7 sections M or N are both very likely to match these 88 SIC 5 division observations.

In the second column of Table 7, we focus on SIC 5 division 88 alone. We list the set of SIC 7 sections that correspond with SIC 5 division 88 (first column), but we also include the official labels of each to understand the categories better. In summary, SIC 5 division 88 can be linked with seven sections from SIC 7 that are not easily distinguishable without further information.

Table 7: SIC 5 division 88

SIC 7	SIC 5
H Transportation and storage	88 – other business activities (where the SIC 5 level 1 classification is "Financial intermediation, insurance, real estate and business services")
J Information and communication	
K Financial and insurance activities	
M Professional, scientific and technical activities (e.g. attorney, auditing)	
N Administrative and support service activities (e.g. renting items, hiring equipment, renting sporting equipment, renting flowers, food processing renting)	
Q Human health and social work activities	
R Arts, entertainment and recreation	

Fortunately, when the BER added the SIC 5 divisions to their dataset, they updated their official documentation to include a table (replicated above as Table 4) that assists researchers converting between BER sectoral classifications and SIC 5 divisions. This table shows that

BER code 180 captures services that the BER split into three types of services when compared to the SIC 5 codes. Specifically, Table 5 shows that an observation designated both BER code 180 and SIC 5 division 88 refers to “Finance, real estate and business services” rather than services in “wholesale, retail, motor, hotels and restaurants” or “community, social and personal services”. This observation would therefore be placed in SIC 7 section M in Table 5.

4. Construction of the inverse propensity weights

4.1 Basic approach

We created post-stratified weights in the survey data for the purposes of this pilot study. The weights ensure that firm-level and sectoral total employment in each municipality are the same as total employment in each municipal-sector combination in the tax data after individual firms are reweighted in calculating statistics of interest. In essence, the weights correct for differential survey shares of sector-specific employment in each municipality to match the administrative records as follows:

$$w_{ms} = \frac{total\ employment_{ms}^{SARS}}{total\ employment_{ms}^{BER}} \tag{1}$$

where m indexes municipalities and s SIC 7 sector.

4.2 Adjustments for non-response

A number of adjustments to the basic equation (1) are required to complete the process of reweighting. Firstly, we have to adjust for *unit* non-response, where firms in particular sectors, locations and of particular sizes may systematically decline to respond to the BER survey in its entirety, creating a biased sample. Standard practice would be to adjust the weights to reflect this survey non-response. However, given that we are not privy to information on the number or characteristics of the firms that declined to respond, we reweight the enumerated data to reflect the population data. Our assumption is that the BER topped up samples for non-response until they judged responses to be adequate in number and representation. A future consideration would be to document all sampling procedures and keep records of non-respondents to make full adjustments. Secondly, we should adjust for *item* non-response – firms that responded to the survey but omitted information on some items or questions. This factor can be addressed in the current setting, with two types of item non-response of concern. The first occurs when information is missing on the variables used to construct weights (location, sector and firm size), while the second occurs when information is missing on the outcome variables of interest (inflation expectations at various horizons).

In the first instance, our largest potential contention is with the location data. But given that we are not concerned about the extent of missing location data in the existing surveys we analyse, we do not make any additional adjustments for location non-response. However, if the location data do not improve in future, or should any specific item response deteriorate, the weights will be adjusted according to

$$w_{ms}^{adjusted} = w_{ms} \times \frac{1}{P(non - miss|x)} \tag{2}$$

where $P(non - miss|x) = \log i\, t(x)$ is the propensity score for the particular question being reported as non-missing in the data, generated from a logistic regression with the vector x as covariates from the survey data. Generating propensity scores is dependent on a vector of attributes that determines whether firms report important information or not in the survey data or whether it is possible to match firms to a municipality and/or sector.

In the second instance, respondents may choose to report their expectations at some horizons but not at others. For instance, as documented below, respondents appear more certain of their assessments of inflation at short horizons, while they are less likely to make predictions at longer horizons. One approach to solving this is similar to the approach used for equation (2), while another is to produce question-specific weights. For the latter, reweighting is done based on a non-missing sample at the *T+0* horizon, and a separate reweighting is done based on a non-missing sample at the *T+j* horizons. We use this approach below.

We also calculate adjustment factors to account for municipalities and SIC 7 sectors not sampled in the BER survey but present in the SARS population data. We therefore adjust the post-stratified weights for municipalities and SIC 7 sectors that are not sampled at all in the BER survey to match the population *provincial* margins. We estimate the adjustments for municipalities and sectors as follows:

$$adjustment = \frac{empl_{SARS_{all\, provinces}}}{empl_{SARS_{excl\, missing}}} \tag{3}$$

where the numerator is the total population for that province or sector in the tax data and the denominator is the population excluding the municipalities and sectors missing in the BER sample.

4.3. Raking methodology

Given the small number of covariates we use to reweight, each of the $m \times s$ municipality-sector employment totals in the survey data can be calibrated to the $m \times s$ cells in the tax data using formula (1). However, should additional variables be used to reweight, the statistical power to do so at 'cell' level would begin to lack. In more complex cases, raking algorithms are appropriate. Instead of recalibrating so that each covariate *combination* in the survey is reweighted to the SARS cell-level totals, weights are chosen so that only marginal totals of the *individual* covariates in the survey data of each post-stratification variable match the tax totals individually. An optimal allocation within each cell is then estimated algorithmically while maintaining the marginal totals of each post-stratification variable in the SARS data. For instance, the maximum entropy estimator (Wittenberg 2010) provides an optimal allocation between the dimensions to ensure that marginal totals are correct along each chosen dimension. In this application, we use the famous CALMAR algorithm implemented in the R software by Rebecq (2016). We therefore ensure that municipal employment totals are reflected in the SARS data, as are sectoral employment totals. The allocation of sector shares uniquely within municipalities is determined by the CALMAR algorithm.

To calibrate weights for respondents using a raking algorithm, each observation requires an initial weight. We assign each observation an initial weight of 1, treating each respondent in our sample homogenously.

5. The impact of calibrated weights on estimating aggregate inflation expectations

Our initial assumption was that the differential between the unweighted and weighted means would increase with the time horizon, suggesting a significant increase in item non-response by firms in the BER sample at the *T*+2 and *T*+5 horizons. We therefore estimate the non-response rates in inflation expectations at each horizon, as shown in Figures 6a–d. The non-response rate for *T*+0 is between 1 and 2 %, with significant increases in 2018 and 2021. This low overall non-response rate is consistent with respondents feeling comfortable to comment on contemporary price pressures. The jump in non-response at this horizon for 2021, however, is likely to have been caused by practical challenges brought about by the COVID-19 pandemic. At *T*+1 the non-response rate is quite stable at 4% to 6 % throughout the sample period. The non-response rates at *T*+2 and *T*+5 are higher on average, at 6% and 7% respectively. These results suggest that as the expectation horizon increases, there is less certainty about future inflation outcomes, and respondents become less likely to offer their opinions. Our initial assumption that the increase in the deviation at longer time horizons is caused by a sharp change in the non-response rate appears to be correct, as there is a significant increase in *T*+5 non-response.

Figure 6: Non-response by expectation horizon

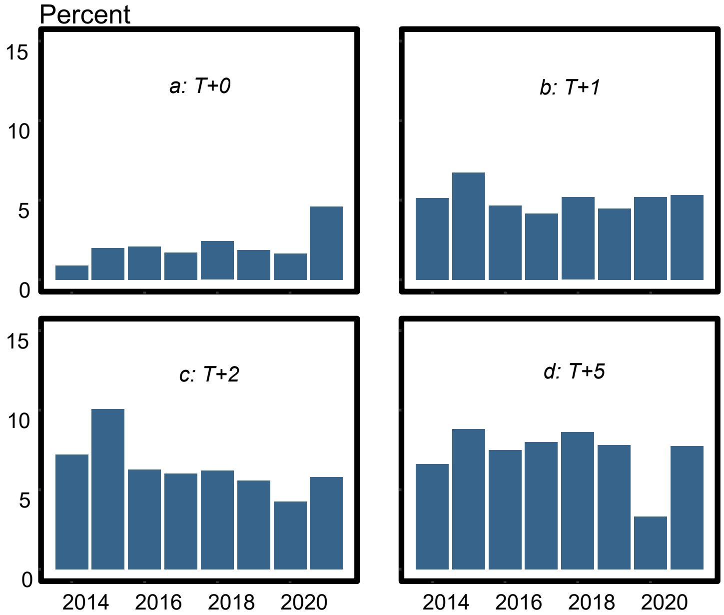

Source: Authors' calculations, based on BER data

Focusing on the T+5 horizon, where the deviation is greatest, we investigate whether the non-response rate varies across regions. Figure 7 plots the non-response for T+5 expectations by province, showing it to be stable across many of the provinces. Item non-response rates are significantly higher in rural provinces, such as Free State and Northern Cape. Limpopo appears to be an outlier, as non-response rates are high across the entire 2014–2018 period.

Figure 7: T+5 non-response rates by province

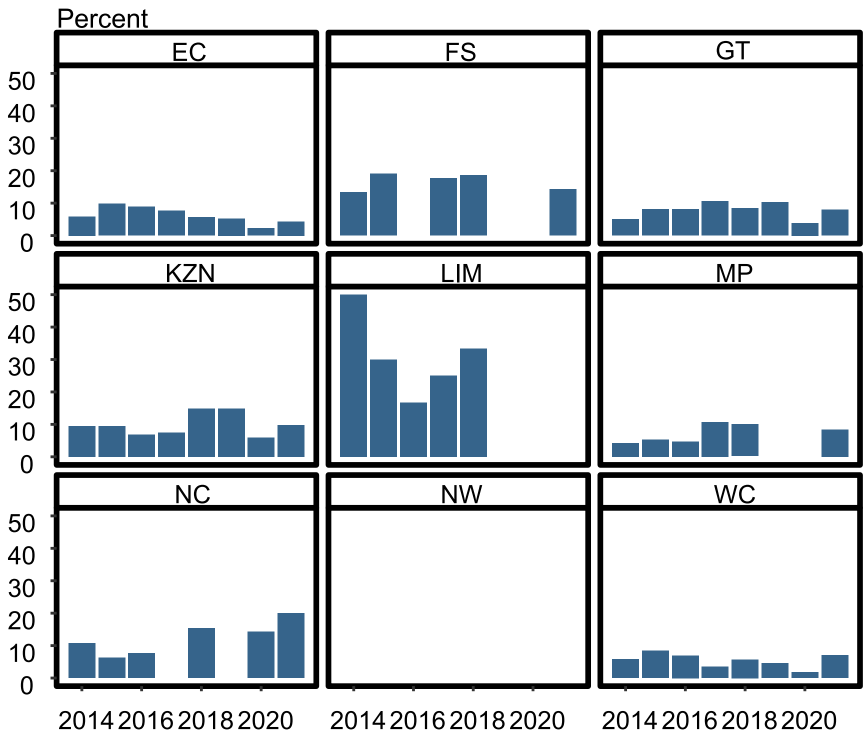

Source: Authors' calculations, based on BER data

Figures 8a–d illustrate the impact of our calibrated weights on average inflation expectations in the sample period for each of the four expectation horizons. Each figure shows three time series: (i) an unweighted mean; (ii) a mean estimated using weights that does not account for item non-response (using weights that can be applied to all variables but that disregard missingness patterns in specific items); and (iii) a mean estimated using weights that are specific to each time horizon and take the missingness into account in the specific outcome. Differences between (i), (ii) and (iii) show whether the BER survey represents the same structure as reflected in the SARS records or not. Differences between (ii) and (iii) highlight the importance of item non-response.

We are able to make some initial observations (i) the weighted mean inflation expectations that ignore item non-response are below the unweighted means, although the differences are rarely statistically significant even if they could be economically meaningful; (ii) the gap between unweighted and these weighted mean inflation expectations increases with the time horizon; (iii) however, accounting for item non-response removes the horizon-specific anomalies, and the level is much closer to unweighted trends. Overall, then, the sample used by the BER has been representative of the economy at large, and item-specific weights are necessary to remove sample selection biases.

Figure 8a: T0 mean inflation expectations (unweighted, weighted and adjusted weights)

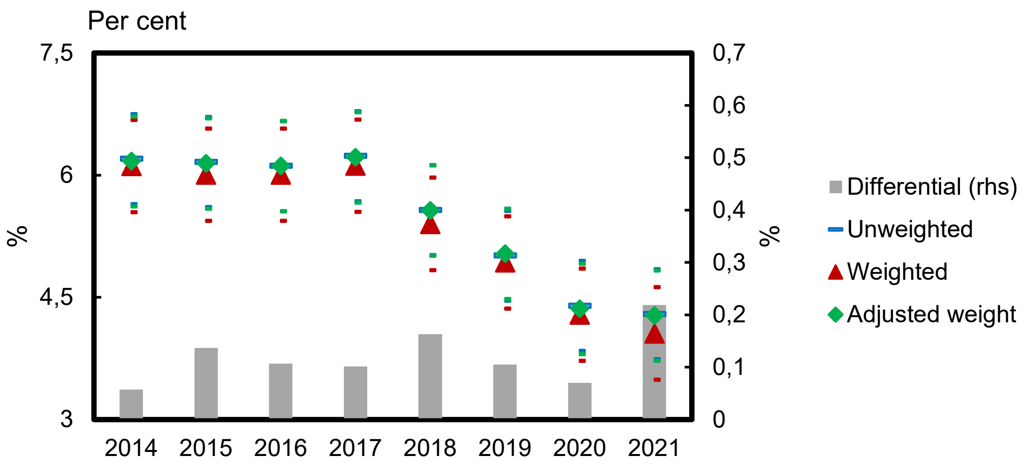

Figure 8b: T+1 mean inflation expectations (unweighted, weighted and adjusted weights)

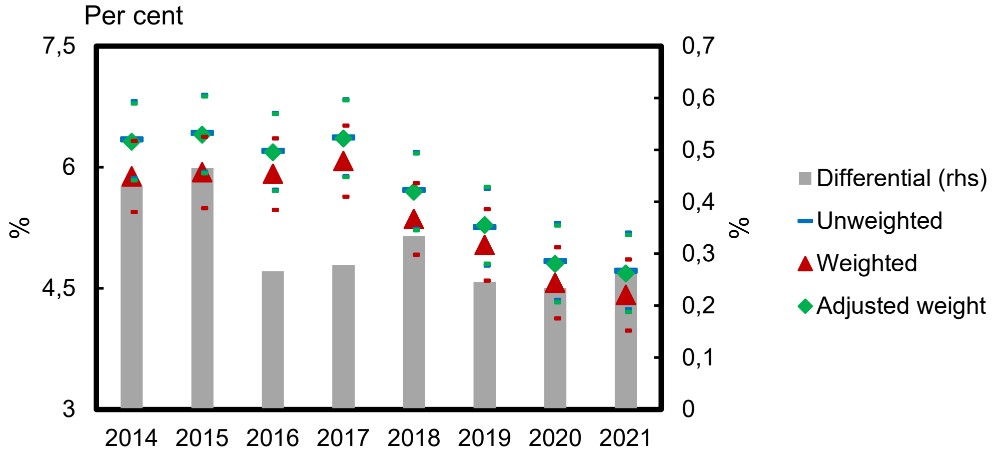

Figure 8c: T+2 mean inflation expectations (unweighted, weighted and adjusted weights)

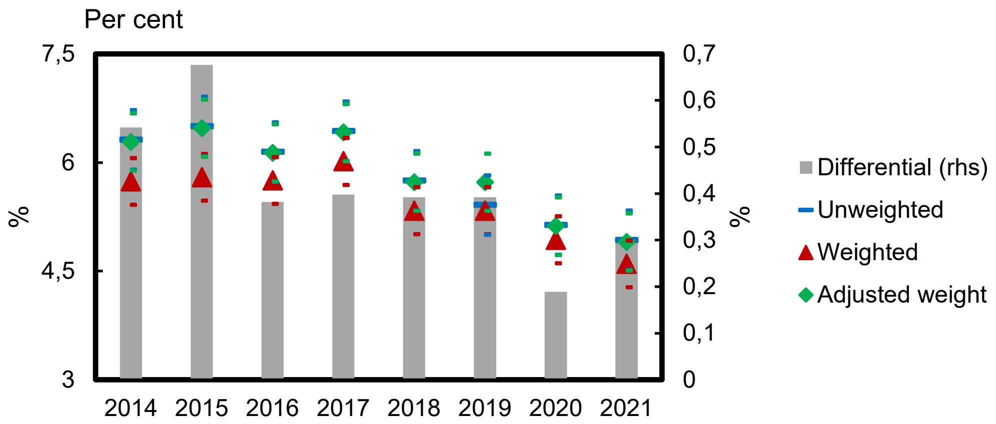

Figure 8d: T+5 mean inflation expectations (unweighted, weighted and adjusted weights)

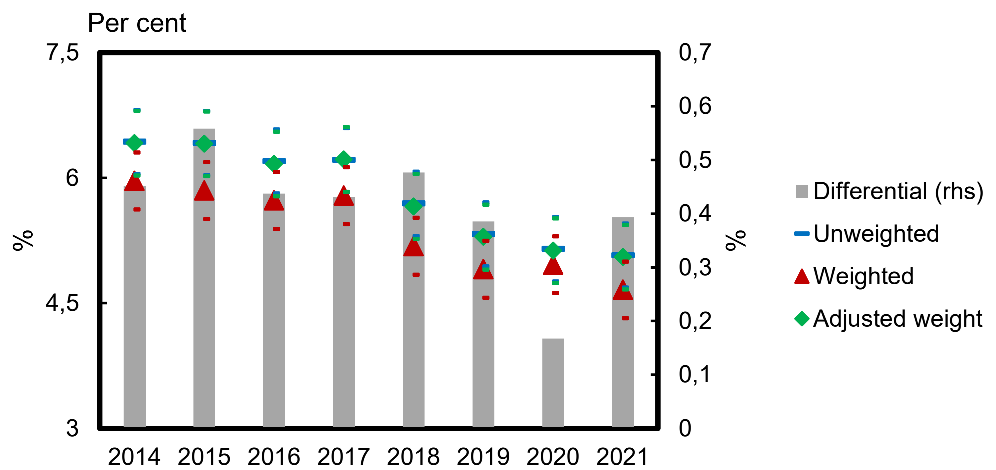

Source: Authors' calculations, based on BER data

Note: Dotted blue, red and green lines are 90% confidence intervals. Grey bars are the difference between adjusted weights and the full sample weights.

The overlap between the confidence bands of the weighted and unweighted expectations at all four horizons suggests that the differences between the various approaches are not statistically significant. This overlap is more pronounced at the T+0 and T+1 horizons, whereas at the longer horizons, especially T+5, the overlapping of the bands is minor by comparison. Figure 8 shows that adjustments for item non-response address these specific anomalies.

6. Recommendations

The focus of this paper is the calculation of sample weights, not the design of a comprehensive sampling strategy. Weighting is applied to a sample after it has been collected to ensure that it is representative when a non-random or deliberate sample is used. Our first recommendation is to add these sample weights to the survey. In contrast, a sampling strategy is used to collect information from an appropriate sample in such a way that the correct information is available to enable generalisation of conclusions (often with weights applied) to the full population. However, sample design and sample weights interact with one another and jointly influence the quality of the survey data. In this sense, our analysis and reweighting exercise has provided some insights relevant for sample design, but we do not claim to have offered a comprehensive sample design strategy.

6.1 Sampling

To ensure a good survey sample, we need to consider both the size and composition of the sample. This was not a significant concern for the BER survey's original purpose of evaluating changes in aggregate inflation expectations over time. The impact of a sample that is slightly biased in some way is more pronounced when researchers use the microdata to explore questions about sub-components of the population.

With regard to sample size, declining survey response rates are an international challenge, caused at least in part by survey fatigue (Dutwin and Buskirk 2021; Stedman et al. 2019). In the face of smaller sample sizes, we need to be even more intentional about the composition of the sample. For instance, we do not adjust for unit non-response – when respondents refuse to participate in the survey altogether – which may be selective in its own right. To adjust for this one would have to have details about all the firms that declined to participate and view this non-response in light of a sampling frame.

The application of sample weights could assist both by adjusting the final results ex post and by enabling more targeted sample recruitment ex ante. It would be useful to identify which risks are amplified when response rates are low and to sample in a methodologically rigorous manner to limit these risks. It is also valuable to be transparent about the limitations of the dataset. For example, the sample's size and composition may not allow conclusions to be drawn with confidence about regional differences. Figure 9 shows the geographical composition of the BER's sample. The survey respondents are predominantly located in Gauteng, KwaZulu Natal and both the Eastern and Western Cape. The other five provinces are underrepresented by comparison, particularly North West and Limpopo.

Figure 9: Geographical composition of survey respondents

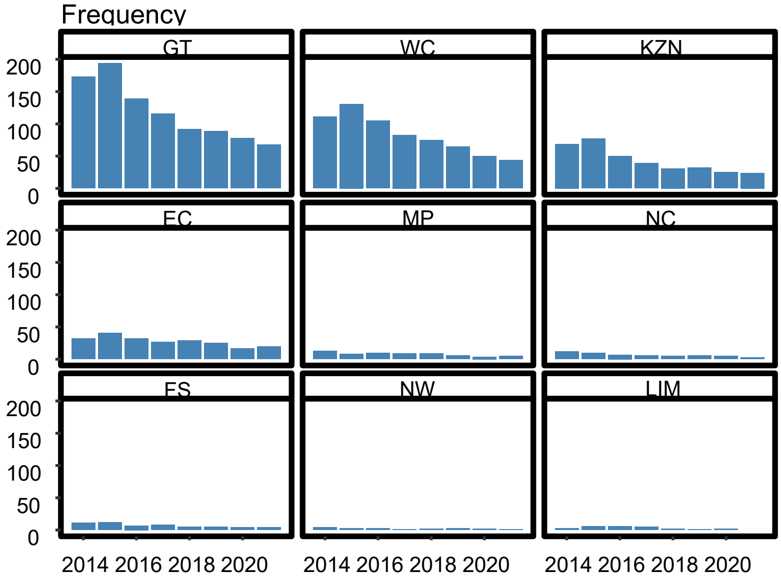

Source: Author's calculations, based on BER data

It is conceivable that, with the exception of the Free State, the regional composition of the sample approximately captures the contribution of each province to aggregate economic activity in South Africa, but this is open to error. If the weights were going to be applied, it would be necessary to follow the common practice of oversampling portions of the population that are smaller before applying survey weights. This would reduce the chances of overstating the characteristics of a single observation. From a practical point of view, it is not possible to upweight a segment that is missing. Given that the BER was not initially focused on looking into regional differences in inflation expectations, less attention has been devoted to geographic representation.[18] If, however, regional price pressures play a role in the overall inflation outlook, it may become necessary to oversample the underrepresented areas of the country and collect information about all the provinces in which a firm does business (as per recommendation 3 in section 6.2).

[18] The BER particularly aimed to ensure that provinces that made a notable contribute to gross domestic product were well represented.

6.2 Additional features to support research with the microdata

Another notable trend is the desire to link surveys such as this one to other datasets, and to administrative data in particular. In this paper, we have contributed to linking the BER inflation expectations surveys to the spatial tax dataset in South Africa. The BER could support this by collecting additional details about its survey respondents. This need not lengthen the quarterly survey, as these details would only be collected when the respondent is registered and possibly checked every couple of years, in the case of respondents that remain in the survey for a long time.

We believe that these characteristics should be guided by the characteristics of the administrative data, and the spatial tax data in particular. The spatial tax data are freely accessible online, and the vision of the team creating them is both ambitious and credible. The ability to link BER data to this dataset is therefore likely to present a large number of opportunities.

We recommend expanding the BER firm-level characteristics to include the following:

1. Both the SIC 5 and SIC 7 industrial classification codes. Following the practice used by SARS, survey respondents could be asked to describe their business activity. The BER could use this to classify each firm according to both the SIC 5 and SIC 7 classifications. This information would need to be recorded when a firm is first registered to take part in the survey, with the details of firms that remain active in the survey being confirmed every five years to track any changes.
2. Firm size by turnover. This might be more of a challenge, as firms may be reluctant to provide this information. One option would be to use turnover ranges (in line with the formal classifications by SARS) rather than precise numbers.
3. Municipality and postal code of the respondent's head office. This could be supported by a second question asking for all the provinces in which the firm does business.

7. Policy implications

Monitoring inflation expectations data is central to the inflation-targeting process, so central banks continuously aim to gather the best information they can about the evolution, factors or processes that drive inflation expectations. Efforts to continue testing the data and models that are central to policy contribute to the credibility of the policy process that relies on them. The

BER has offered quality data to the SARB for decades. This paper's reweighting exercise confirms that the representativity of the surveys between 2014 and 2021 has been good, and it offers additional ways to improve this representativity.

Recent experience has also highlighted the value of analysing disaggregated data. The interactions between supply and demand shocks during and after the pandemic were not taken seriously enough, and the resulting inflation caught many countries, particularly in the advanced world, off guard. In the years ahead, climate-related shocks and geopolitical tensions mean that supply shocks are likely to continue to present real threats. These shocks often start in particular regions and sectors but may eventually transmit to other parts of the economy. The ability to analyse disaggregated inflation expectations data may offer earlier signals and deeper insights into the dynamics of inflation.

It is appropriate for a central bank to look through supply shocks if these do not become entrenched. The initial shock is likely to affect the inflation expectations of firms in industries immediately affected by the shock, but if these price pressures start being reflected in the expectations of industries further removed from the initial impact, this would be early evidence that the shock will be less transitory. Many macroeconomists today recognise that much can be learnt from looking beyond the mean, reflected in the proliferation of research using microdata to answer macroeconomic questions. This paper is a contribution to enabling this kind of work in South Africa.

8. Conclusion

The BER survey is well respected in South Africa. It is rich by international standards and provides admirable time series, covering over two decades. The survey is regularly used for policy analysis as well as by the private sector and media to provide some insights into inflation expectations and the dynamics of inflation itself. The notable increase in attention to inflation expectations surveys in recent years has inevitably resulted in suggestions to improve their design. In this paper, we investigate the extent to which the sample is unrepresentative and calculate survey weights to ensure the continued precision of the survey composition and to support more targeted sample recruitment in the face of high non-response rates.

The BER has relied on deliberate sampling and well-informed intuition about which sectors need to be better represented when new survey respondents are recruited. In aggregate (as the data have been almost exclusively used since its inception), low levels of sample misrepresentativeness are unlikely to substantially change the result. If the researcher's main interest is in changes to this aggregate figure, then the impact would be muted. However, increased interest in the disaggregated data – in order to ask questions about how different industrial sectors are responding or to more rigorously understand how the expectations of other sub-groups are formed – magnifies the impact of any non-representativeness.

Enabled by access to administrative data via the spatial tax data project, we have estimated survey weights, which can be added to the survey dataset ex post. These weights offer confidence in the representativity of the dataset along three characteristics (region, industrial sector and firm size by employment numbers). Knowledge of the degree to which the current sample differs from the actual structure of the economy enables more targeted recruitment of new survey respondents. Ideally, we recommend that the BER move towards a sample design pulled from a sampling frame of SARS data, which is now available. In addition to offering a formal way of maintaining the representativity of the survey, this will also allow researchers to link this dataset to other datasets in the country, enabling rigorous research into topics that to date have been inaccessible.

We further recommend that the BER collect information about the business activity undertaken by each firm as it is recruited. The BER should then use this information to categorise each firm according to both the SIC 5 and SIC 7 classifications so that future researchers can link datasets with either set of codes. It would also be useful to know a firm's size by turnover rather than employment numbers. Finally, we also recommend that some metric of regional activity be collected at registration. Besides enabling us to link the BER survey data to the spatial panel dataset, this would enable us to provide insights into regional inflation patterns.

Annexure

Figure 10: Sectoral frequency in the tax data by year

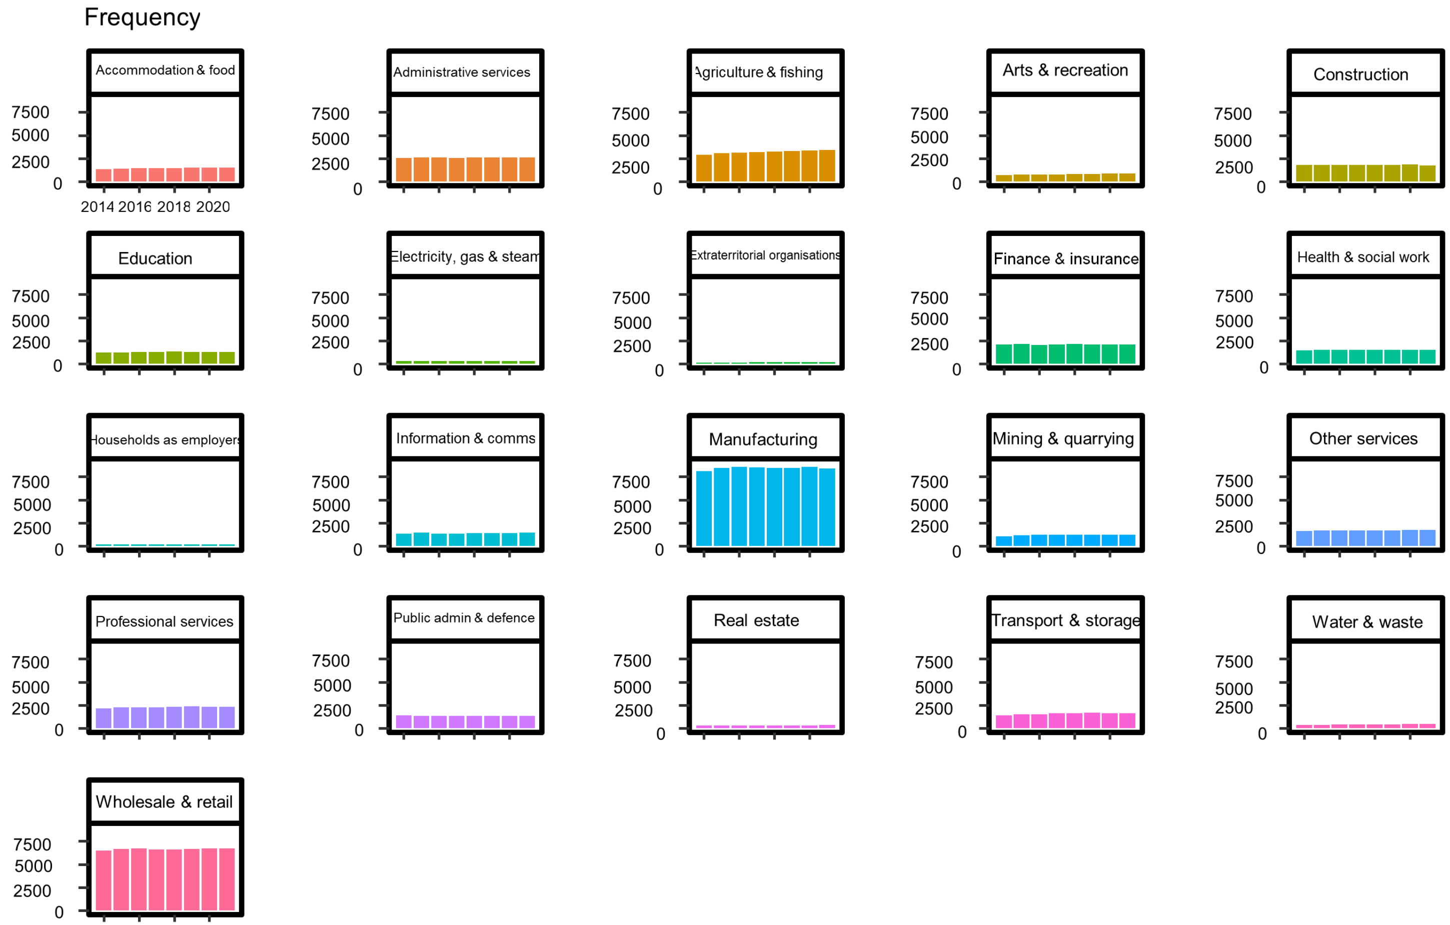

Source: Author's calculations based on SARS tax data, classified according to SIC 7

References

Adrian, T. 2023. 'The role of inflation expectations in monetary policy'. Remarks at the Deutsche Bundesbank Symposium.

Amaral, A, Kruger, M and Reid, M. 2023. 'Looking to the price setters: what can we learn from firm-level inflation expectations data?' *South African Reserve Bank Economic Note Series*, No. EN/2023/16.

Binder, C. 2015. 'Whose expectations augment the Phillips curve?' *Economics Letters* 136(C): 35–38.

Budlender, J and Ebrahim, A. 2020. 'Industry classification in the South African tax microdata'. SA-TIED Working Paper No. 134, August.

Bureau for Economic Research. 2023. 'BER/SARB Inflation Expectations Survey micro data (individual de-personalised responses): description'. Unpublished internal documentation about the data.

Candia, B, Coibion, O and Gorodnichenko, Y. 2021. 'The inflation expectations of U.S. firms: evidence from a new survey'. *National Bureau of Economic Research Working Paper Series,* No. 28836.

Candia, B, Coibion, O and Gorodnichenko, Y. 2022. 'The macroeconomic expectations of firms'. *National Bureau of Economic Research Working Paper Series*, No. 30042.

Coibion, O and Gorodnichenko, Y. 2015. 'Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation'. *American Economic Journal*: *Macroeconomics* 7(1): 197– 232

Crowther-Ehlers, N. 2019. 'Inflation expectations in South Africa: non-rational, intertemporal and idiosyncratic heterogeneity represented by a term structure approach'. PhD thesis, Stellenbosch University.

Dutwin, D and Buskirk, T D. 2021. 'Telephone sample surveys: dearly beloved or nearly departed? Trends in survey errors in the era of declining response rates'. *Journal of Survey Statistics and Methodology* 9(3): 353–380.

Ehlers, N and Steinbach, R. 2007. 'The formation of inflation expectations'. *South African Reserve Bank Working Paper Series No. WP/07/06*, November.

Harrison, K, Nell, A, Turok, I and Visagie, J. 2023. 'A breakthrough in spatial economic data'. In *Cities economic outlook 2023: insights into South Africa's spatial economy from tax data* (Pretoria: Spatial Economic Activity Data, South Africa), 3–8.

Israël, J and Tisso, B. 2021. 'Incorporating micro data into macro policy decision-making'. Bank for International Settlements, IFC Bulletin 53.

Kershoff, G. 2002. 'Conducting inflation expectations surveys in South Africa'. *South African Journal of Economics* 70(3): 1–16.

Kershoff, G J and Smit, B W. 2005. 'Conducting inflation expectation surveys in South Africa'. *South African Journal of Economics* 70(3): 205–212.

Lombaard, M. 2004. South African postcode geography. Available at: https://api.semanticscholar.org/CorpusID:130886761.

Pieterse, D, Gavin, E and Kreuser F. 2018. 'Introduction to the South African Revenue Service and National Treasury Firm-Level Panel'. *South African Journal of Economics* 86(1): 6–39.

Rebecq, A. 2016. 'Icarus: an R package for calibration in survey sampling'. *R Package version 0.3.0*.

Reid, M B and Siklos, P L. 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.

Reid, M and Siklos, P. 2021. 'The BER's inflation expectations surveys: findings from academic research'. *Bureau for Economic Research*, Research Note No. 4.

Rudd, J B. 2022. 'Why do we think that inflation expectations matter for inflation? (And should we?)'. *Review of Keynesian Economics* 10(1): 25–45.

Spatial Economic Activity Data, South Africa. Available at: https://spatialtaxdata.org.za/download-data-landing (accessed: 12 September 2023).

Stedman, R C, Connelly, N A, Heberlein, T A, Decker, D J and Allred, S B. 2019. 'The end of the (research) world as we know it? Understanding and coping with declining response rates to mail surveys'. *Society & Natural Resources 32*(10): 1139–1154.

Von Fintel, D. 2023. 'Monitoring South Africa's metropolitan economies: a survey of the data landscape'. Economic Research South Africa, Discussion Document No. 13.

Wittenberg, M. 2010. 'An introduction to maximum entropy and minimum cross-entropy estimation using Stata'. *The Stata Journal* 10(3): 315–330.