Can households predict where the macroeconomy is headed?

Kamil Kladívko and Pär Österholm*

Kamil Kladívko is Associate Senior Lecturer in Economics at the School of Business, Örebro University. Pär Österholm is Professor of Economics at the School of Business, Örebro University, and is an affiliated researcher at the National Institute of Economic Research.

Survey data of households' expectations of macroeconomic variables might provide useful information to those who analyse or forecast the economy. In this article, we evaluate whether households can predict in which direction inflation and the unemployment rate will move over the coming year. The analysis is conducted using monthly Swedish data from the National Institute of Economic Research's *Economic Tendency Survey* over the period from January 1996 until August 2019. Our results indicate that households can forecast in what direction the unemployment rate is headed, but they fail to predict the direction of future inflation.

1 Introduction

For an economic policymaker, such as a central bank, the expectations of the economy's agents might be of interest from several perspectives. For example, long-run inflation expectations could be informative regarding the credibility of the inflation target. Other types of expectations, such as short- or medium-term expectations of GDP growth, inflation, wage growth or the unemployment rate, can provide useful input for policymakers since the actions of the agents – and thereby macroeconomic outcomes – tend to depend on the agents' expectations. Yet another aspect is that expectations might be good forecasts that the policymaker could take into account in order to improve its own forecasts. Various properties of the expectations might also reveal how expectations are formed and evidence of deviations from rationality, for example, could affect how a policymaker both conducts policy and communicates. Accordingly, it is not surprising that quite some effort is put into measuring agents' expectations. For example, in Sweden, two surveys are conducted on a monthly basis which (among other things) address the issue of inflation expectations; households are interviewed in the National Institute of Economic Research's *Economic Tendency Survey* ('*Konjunkturbarometern*') and money-market participants are interviewed in a survey commissioned by Sveriges Riksbank, commonly referred to as the *Prospera Survey*.[1]

In this article, we analyse households' survey expectations. The reason for this focus is the simple fact that households constitute an important part of the economy; for example, household consumption's share of GDP is approximately 45 per cent in Sweden. More specifically, we evaluate households' directional forecasts of inflation and the unemployment rate in Sweden. This is done using monthly data from the *Economic Tendency Survey*.

In conducting this analysis, we follow a line of research that can be seen as being concerned with the forecasting properties of the expectations themselves; see, for example,

* We are grateful to Marianne Nessén, André Reslow and Ulf Söderström for valuable comments.

1 Businesses are also interviewed regarding their inflation expectations in the *Economic Tendency Survey*. However, this is done at a quarterly frequency. In a similar manner, employee organisations, employer organisations, manufacturing companies and trade companies are interviewed in the *Prospera Survey* at a quarterly frequency. The *Economic Tendency Survey* is conducted by Origo Group. The *Prospera Survey* is conducted by Kantar Sifo.

Batchelor and Dua (1989), Thomas (1999), Trehan (2015), and Berge (2018). That is, we are interested in whether the households can predict where the economy is headed. Our focus is accordingly different from that of the fairly voluminous literature which utilises household survey data in order to generate model-based macroeconomic forecasts.[2] Another novel aspect of our article is that we study directional forecasts. While it is not uncommon for such forecasts to be analysed in the macroeconomic literature – see, for example, Ash et al. (1998), Greer (2003), Thomas and Grant (2008), Baghestani et al. (2015), and Driver and Meade (2019) – it is nevertheless the case that numerical forecasts tend to be the focus in the overwhelming majority of empirical studies.[3] No rigorous analysis has previously been conducted on the directional forecast data that we study in this article. Our study should hence bring new information to policymakers and others who analyse and forecast the Swedish economy.

2 Data

We use monthly data from the *Economic Tendency Survey* ranging from January 1996 to August 2019.[4] In this survey, 1,500 randomly sampled Swedish households are interviewed each month.[5] The respondents are asked a number of questions related to their own economic situation and the Swedish economy at an aggregate level. This is Sweden's most important household survey and it is part of the European Commission's *Joint Harmonised EU Programme of Business and Consumer Surveys*.

As a key survey in Sweden, data from the *Economic Tendency Survey* have of course been analysed previously. For example, based on micro-level data, Jonung (1981), Jonung and Laidler (1988), Batchelor and Jonung (1989), Palmqvist and Strömberg (2004), and Hjalmarsson and Österholm (2019, 2020, 2021) have investigated various aspects of perceived inflation, inflation expectations, mortgage-rate expectations and housing-price expectations. Aggregate time series from the survey – such as confidence indicators or the mean expectation of a variable – are also commonly used for macroeconomic forecasting and analysis; see, for example, Hansson et al. (2005), Assarsson and Österholm (2015), Hjalmarsson and Österholm (2017), and Jönsson (2020).

In this article, we analyse data on the two questions in the survey that concern directional forecasts – that is, questions 6 and 7 in the survey. These questions pertain to inflation and the unemployment rate. Their formulations, and the answers available to the respondents, are given below:[6]

> Question 6: Compared to the situation today, do you think that in the next 12 months prices in general will...?
>
> i. Increase faster
> ii. Increase at the same rate
> iii. Increase at a slower rate
> iv. Stay about the same
> v. Fall slightly
> vi. Don't know

2 See, for example, Carroll et al. (1994), Easaw and Herevi (2004), Dreger and Kholodilin (2013), Assarsson and Österholm (2015), and Campelo et al. (2020) for just a few contributions.

3 Additional studies addressing directional forecasts of inflation or the unemployment rate include Sinclair et al. (2010), Ahn and Tsuchiya (2016), Chen et al. (2016), Ahn (2018), Pierdzioch et al. (2018), and Sosvilla-Rivero and Ramos-Herrera (2018).

4 The survey started out as a quarterly survey in 1973. Since 1993 it has been conducted on a monthly basis. The starting date for the time series studied here is January 1996.

5 The number of respondents in the survey has varied over time. During the sample that we are employing, it has ranged between 1,500 and 2,100. The present number of respondents is 1,500 per month.

6 It should be noted that question 6 has the phrasing stated here if the respondent's 'perceived inflation now' (which is question 5 in the survey) is positive. If the respondent's 'perceived inflation now' is non-positive, the phrasing of the question is adjusted somewhat in order to make it consistent with non-positive inflation today.

Question 7: How do you think the level of unemployment in the country will change over the next 12 months? Will it...?

i. Increase sharply
ii. Increase slightly
iii. Remain the same
iv. Fall slightly
v. Fall sharply
vi. Don't know

We evaluate the survey expectations against the outcomes for CPI inflation and the unemployment rate (seasonally adjusted, age group 16 to 64 years); the last available observation for the outcomes is from August 2020.

In order to econometrically analyse the forecasting performance of the survey data, we generate a directional forecast. This is achieved by first taking the balance, b_t, of the share of respondents (as a percentage) that at time t predicted an increase, minus the share that predicted a decrease. This balance is similar to diffusion indices that are commonly generated from survey data; see, for example, OECD (2000) and Pinto et al. (2020). We then turn the balance into a directional forecast, x_t, according to the rule $x_t = 1$ (indicating an increase) if $b_t > 0$ and $x_t = 0$ (indicating a decrease) if $b_t \leq 0$.[7]

To construct the balance for the inflation question, we take the share of respondents choosing the first alternative among the possible answers, minus the total share choosing the third, fourth and fifth alternatives. Two things deserve to be pointed out concerning this issue. The first is that the question is phrased in terms of prices rather than inflation. This might add a layer of complication if the respondent is used to thinking in terms of inflation. This is not unlikely to be the case in Sweden, since formal inflation targeting was introduced in 1993 and communication typically concerns inflation (rather than the price level).[8] Second, only one of the possible answers implies that inflation will increase, whereas three alternatives imply that inflation will decrease. This feature has the possibility of skewing the respondents' answers due to the so-called *end aversion bias*, which means that respondents tend to avoid the endpoints of a response scale and prefer alternatives closer to the midpoint.[9] For the unemployment rate, the balance is generated as the share of respondents choosing the first two alternatives minus the share choosing the fourth and fifth.

As an illustration of how the balance and forecasts are constructed, consider the unemployment-rate question in January 1996. 6 per cent of respondents answered that the unemployment rate would 'increase sharply', 33 per cent answered 'increase slightly', 42 per cent answered 'remain the same', 16 per cent answered 'fall slightly' and 0 per cent answered 'fall sharply'; finally, 3 per cent answered 'don't know'. The balance is given as $b_{January\ 1996} = 6 + 33 - 16 - 0 = 23$ and the directional forecast accordingly becomes $x_{January\ 1996} = 1$, indicating that respondents predicted an increase in the unemployment rate.

The shares of responses (six for each variable) over time are shown in the top panels of Figures 1 and 2; the bottom panels of Figures 1 and 2 show the resulting balances along with the actual rates of inflation and unemployment.

7 We have removed the possibility of having 'unchanged' as a category by merging $b_t = 0$ and $b_t < 0$. This is reasonable though as $b_t = 0$ in only three cases for inflation and two cases for the unemployment rate.

8 Formally, the Riksbank announced in January 1993 that the target for monetary policy would be 2 per cent inflation, starting in 1995.

9 This bias is related to the more general behavioural phenomenon *extremeness aversion*; see, for example, Neumann et al. (2016) for a discussion.

Figure 1. Shares of different answers concerning inflation, and related balance and inflation rate
Shares (top panel) in per cent. Balance (bottom panel, left axis) in percentage points. Inflation (bottom panel, right axis) in per cent

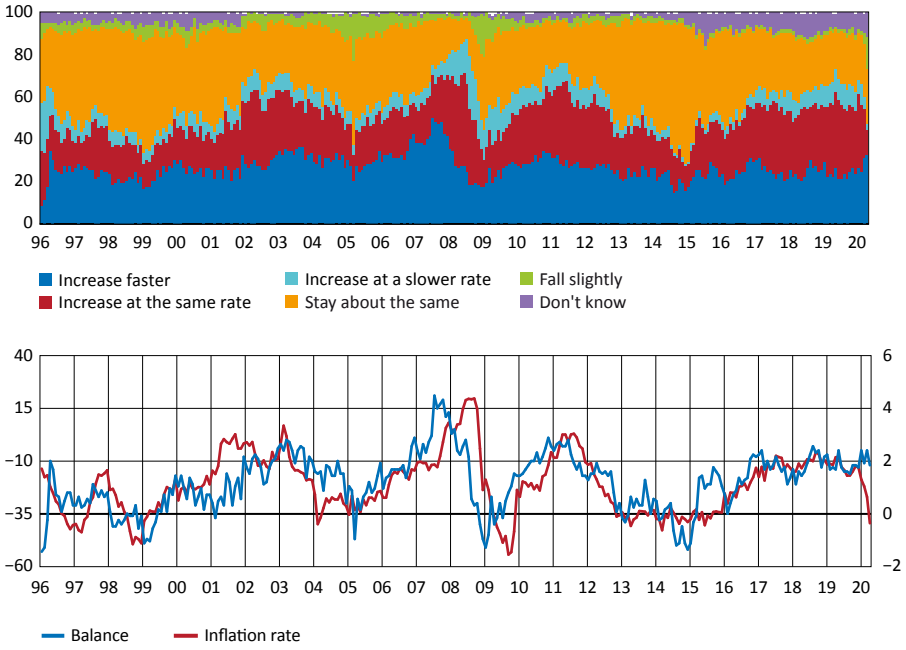

Sources: National Institute of Economic Research, Macrobond and authors' calculations

Figure 2. Shares of different answers concerning the unemployment rate, and related balance and unemployment rate
Shares (top panel) in per cent. Balance (bottom panel, left axis) in percentage points. Unemployment rate (bottom panel, right axis) in per cent

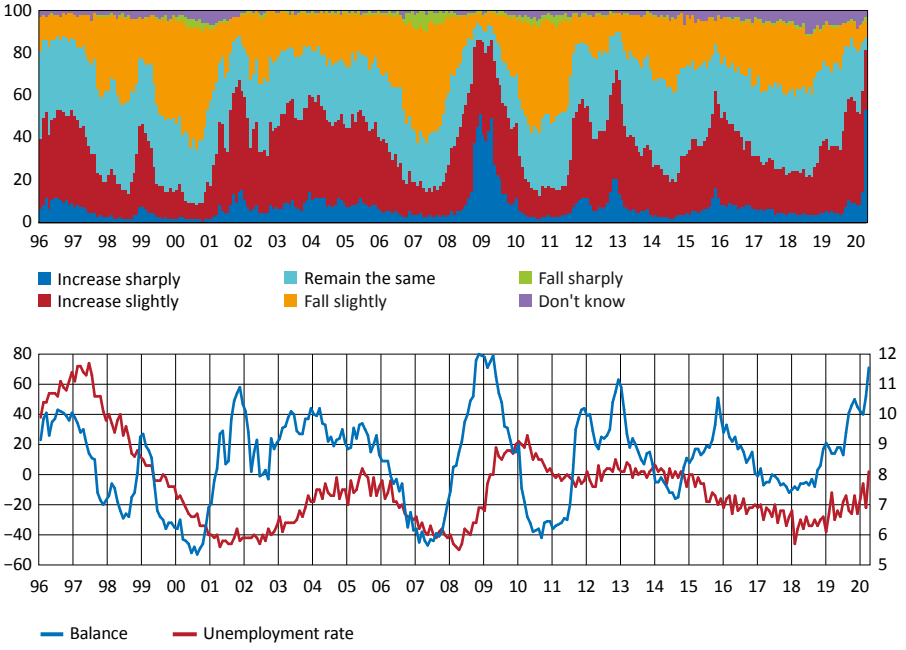

Sources: National Institute of Economic Research, Macrobond and authors' calculations

Looking at Figure 1, it can be seen that the share of respondents that say that prices will 'increase faster' (which indicates the opinion that inflation will increase) has been 26 per cent on average. A substantially higher share can be found in 2007, when it was in the interval 38 to 50 per cent. This was a time period when inflation in Sweden was rising and there was a discussion about increased inflationary pressure; see, for example,

Sveriges Riksbank (2007). However, the share predicting an increase in inflation came down substantially in 2008; interestingly, this process started well before the financial crisis hit its peak in the autumn. In addition, the share of respondents answering that prices will 'increase at the same rate' – that is, that inflation will stay the same – has historically had a similar average, namely 23 per cent. It is noteworthy though that, on average, only seven per cent of respondents have said that prices will 'increase at a slower rate', whereas 35 per cent have chosen the alternative that prices will 'stay about the same'. If inflation is above zero – which it typically has been – the latter statement is also a statement about falling inflation but a more specific one. Finally, the share that says that prices will 'fall slightly' has, on average, been small – approximately four per cent. Concerning the balance regarding the inflation question, a striking feature is the fact that the series almost never takes on positive values. In fact, the balance is positive only in eleven cases, implying that it is very rare that a majority of the households forecast increasing inflation.

Turning to Figure 2, this shows a fair bit of variation over time in the shares for the unemployment rate. For example, for a few months in the year 2000, less than 10 per cent of the respondents said that the unemployment rate would increase; in December 2008, this figure peaked at 86 per cent. It can be noted that during the financial crisis and its more immediate aftermath many respondents also said that the unemployment rate would 'increase sharply', making this period stand out. In general, most of the variation is due to changes in the three central alternatives ('increase slightly', 'remain the same' and 'fall slightly'). Very few respondents – approximately one per cent on average over time – suggest that the unemployment rate will 'fall sharply'. It is reasonable that households seem reluctant to predict a sharply falling unemployment rate. Stylised facts regarding business cycles do not suggest that unemployment rates tend to decrease rapidly. The balance for the unemployment rate question is more centred around zero and appears to have a clearer cyclical pattern than the balance for the inflation question.

3 Empirical analysis

We now turn to an empirical analysis of our data and first give a graphical illustration. Figures 3 and 4 show the actual rates of inflation and unemployment, along with the directional forecasts and the actual directional changes. In each figure, Panel A displays the directional forecast, x_t, that was calculated from the balance, b_t. Forecasts of an increase ($x_t = 1$) are indicated with dark grey and forecasts of a decrease ($x_t = 0$) are indicated with light grey. Panel B records the actual directional change of the forecasted variable. It should be noted that this has been aligned with the forecast origin date – that is, at a given date, it indicates the directional change between that date and twelve months later. The actual directional change, y_t, is coded analogously to x_t, namely $y_t = 1$ if the variable increases over the twelve-month horizon and $y_t = 0$ otherwise. Note that the value of y_t becomes known at time $t + 12$. An increase ($y_t = 1$) is indicated with dark grey and a decrease ($y_t = 0$) is indicated with light grey. Finally, panel C captures the match of the directional forecast with the actual directional change (and also displays the actual rates of inflation and unemployment). Correctly forecasted directions are reported using dark and light grey shaded areas, while incorrectly forecasted directions are reported using white areas. The dark grey areas correspond to the case when an increase in the variable was correctly forecasted and the light grey areas indicate when a decrease in the variable was correctly forecasted. Observe that the correctly forecasted directions in panel C follow from the intersection of directions in panels A and B.

The share of correct forecasts is 46 per cent for inflation and 62 per cent for the unemployment rate (see Table 1). It is noteworthy how an increase in inflation is almost never correctly forecasted; in almost all cases where the outcome was an increase in inflation, a decrease had been predicted. This is, of course, related to the fact pointed out above, namely that households almost always predict a decrease in inflation.

Figure 3. Directional forecast of inflation
Inflation (panel C) in per cent

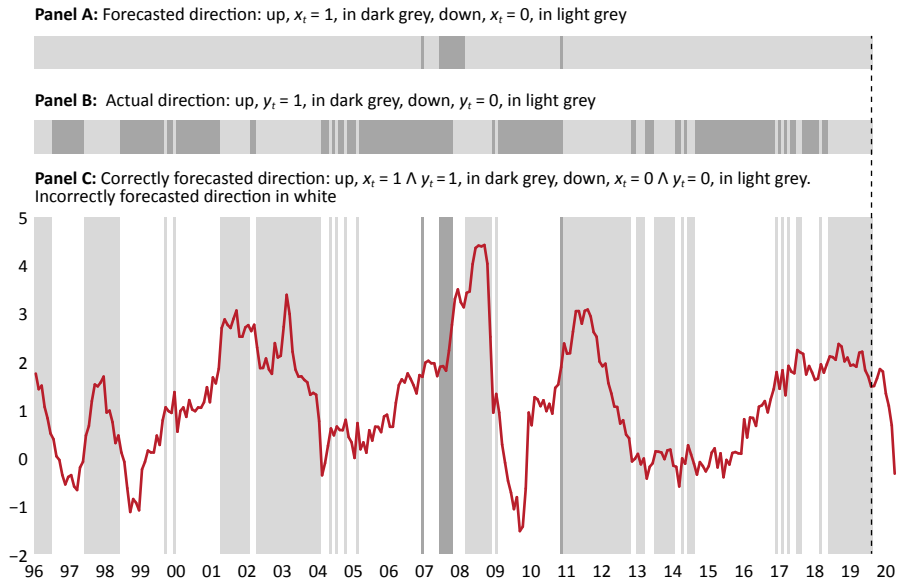

Note. Panel A: Forecasted directional change aligned with the forecast origin date. The vertical dashed line indicates the last forecast that can be evaluated. Panel B: Actual directional change of inflation aligned with the forecast origin date. Panel C: The red line gives CPI inflation (year-on-year). Correctly forecasted direction in dark and light grey; incorrectly forecasted direction in white.
Sources: Macrobond and authors' calculations

Figure 4. Directional forecast of the unemployment rate
Unemployment rate (panel C) in per cent

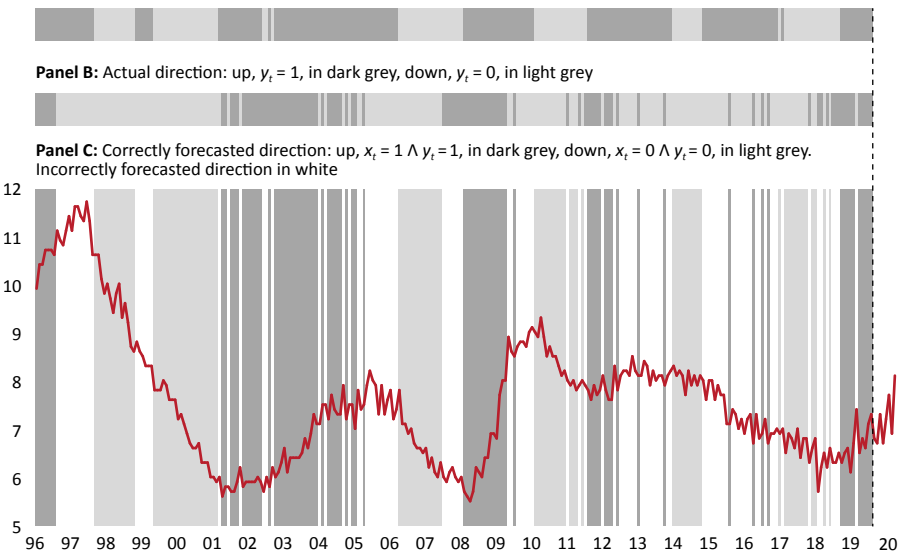

Note. Panel A: Forecasted directional change aligned with the forecast origin date. The vertical dashed line indicates the last forecast that can be evaluated. Panel B: Actual directional change of the unemployment rate aligned with the forecast origin date. Panel C: The red line gives the unemployment rate (in per cent). Correctly forecasted direction in dark and light grey; incorrectly forecasted direction in white.
Sources: Macrobond and authors' calculations

We assess the accuracy of the directional forecasts by employing the Pesaran and Timmermann (1992) test with a Newey-West correction for the presence of serial correlation. This test is effectively a test of independence between the directional forecast x_t and the actual directional change y_t; see the Appendix for a detailed description of the test.

The 2×2 tables to test the independence between the forecast and the realized directional change for our two variables, as well as the test results, are reported in Table 1.

The test statistic, t_{PT}^{NW}, clearly confirms what is suggested by the figures. The null hypothesis of independence of forecasted and actual direction cannot be rejected for inflation but is forcefully rejected for the unemployment rate. We accordingly conclude that Swedish households are unable to forecast where inflation is headed, whereas they have highly significant ability in forecasting the direction of the unemployment rate.

Table 1. Results from the directional accuracy test.

	Inflation		Unemployment rate	
	Actual up $y_t = 1$	Actual down $y_t = 0$	Actual up $y_t = 1$	Actual down $y_t = 0$
Forecast up $x_t = 1$	7	4	86	89
Forecast down $x_t = 0$	149	124	20	89
Proportion of correct forecasts, $\hat{P}$	0.46		0.62	
Estimated expected proportion of correct forecasts, $\hat{P}_{H0}$	0.45		0.47	
Test statistic, t_{PT}^{NW}	0.43		3.47	
p-value	0.664		<0.001	

Note. The top part of the table gives the 2×2 contingency tables of 284 forecasts to test the independence between the households' forecast and actual direction. $\hat{P}$ is the proportion of correct forecasts (see equation (4) in the Appendix), whereas $\hat{P}_{H0}$ is an estimate of the expected proportion of correct forecasts under the null hypothesis of independence (see equation (6)). t_{PT}^{NW} is the Pesaran-Timmermann (1992) test statistic with Newey-West correction (see equations (10)–(12)).

The fact that households have some success in predicting the direction of the unemployment rate is not completely surprising. While macroeconomic forecasting by no means is a trivial exercise, the unemployment rate appears to have fairly distinctive cyclical swings (as can be seen from Figures 2 and 4). It likely also helps that the unemployment rate is a concept to which it should be reasonably easy for households to relate.

The failure when it comes to predicting the direction of inflation is perhaps no surprise either. Given the somewhat mixed evidence in the previous literature, no unambiguous conclusions can be drawn regarding different agents' ability to forecast the direction of inflation. Our results are nevertheless in line with recent studies that point to households not being successful at this task; see, for example, Ahn and Tsuchia (2016) and Ahn (2018).[10] It should be kept in mind, however, that this international evidence is based on households that face economic environments that are quite different to that in Sweden. We believe that contributing factors to the failure are the phrasing of the question and the available answers, which were discussed above. One should also consider that a substantial part of the investigated sample comes from a period where inflation may have been quite difficult to predict. This is related to the fairly widespread claim that the connection between the real economy and inflation in many countries is weaker today than previously or, put differently, that the Phillips curve has become flatter; see, for example, Bean (2006), Gaiotti (2010), Kuttner and Robinson (2010), IMF (2013), and Occhino (2019).[11] It accordingly does not seem unreasonable to conclude that the conditions under which households have been forecasting inflation have, at least in parts, been non-trivial.

10 In contrast, some studies indicate that professional forecasters are somewhat more successful at forecasting the direction of inflation; see, for example, Chen et al. (2016) and Sosvilla-Rivero and Ramos-Herrera (2018).

11 This is not an undisputed claim though; see, for example, Fitzgerald et al. (2013) and Berger et al. (2016). For some additional recent contributions concerning the Phillips curve, see, for example, Coibion and Gorodnichenko (2015), Blanchard (2016), Leduc and Wilson (2017), and Karlsson and Österholm (2020).

4 Concluding remarks

In this article, we have shown that Swedish households have statistically significant ability in forecasting the direction of the unemployment rate but that they fail in forecasting where inflation is headed. Despite the failure regarding the directional forecasts of inflation, it can still be worth monitoring these expectations since flawed expectations can still contain useful information to a policymaker, for example. Of course, it is also of interest to know that the expectations have shortcomings.

The finding that the households fail in forecasting the direction of inflation can, to some extent, probably be explained by the fact that inflation objectively has been difficult to predict during a substantial part of the analysed sample. However, we believe that another relevant aspect is that the phrasing of the question in the survey and the answers available are somewhat problematic. The question is phrased in terms of prices, which might complicate things for a respondent who is used to thinking in terms of inflation. Concerning the answers, the respondents' choices could be affected by the fact that only one of the alternatives implies that inflation will increase. To conclude, it does not seem unlikely that the inflation question might be perceived as complicated by the respondents and it could be the case that some respondents do not have sufficient 'economic literacy' to pass this hurdle. This is something that designers of household surveys ought to keep in mind. Considering that this question is part of the *Joint Harmonised EU Programme of Business and Consumer Surveys* (European Commission, 2016), this is likely to be an issue of relevance beyond the Swedish context.

References

Ahn, Young Bin (2018), 'Directional accuracy of urban consumers' inflation forecasts in China', *Emerging Markets Finance and Trade,* vol. 54, pp. 1414–1424.

Ahn, Young Bin and Yoichi Tsuchiya (2016), 'Directional analysis of consumers' forecasts of inflation in a small open economy: Evidence from South Korea', *Applied Economics*, vol. 48, pp. 854–864.

Ash, J. Colin K., David J. Smyth and Saeed M. Heravi (1998), 'Are OECD forecasts rational and useful? A directional analysis', *International Journal of Forecasting*, vol. 14, pp. 381–391.

Assarsson, Bengt and Pär Österholm (2015), 'Do Swedish consumer confidence indicators do what they are intended to do?', *Applied Economics Quarterly*, vol. 61, pp. 391–404.

Baghestani, Hamid, Mohammad Arzaghi and Ilker Kaya (2015), 'On the accuracy of Blue Chip forecasts of interest rates and country risk premiums', *Applied Economics*, vol. 46, pp. 113–122.

Batchelor, Roy A. and Pami Dua (1989), 'Household versus economist forecasts of inflation: A reassessment', *Journal of Money, Credit and Banking*, vol. 21, pp. 252–257.

Batchelor, Roy A. and Lars Jonung (1989), 'Cross-sectional evidence on the rationality of the mean and variance of inflation expectations', In: Grunert, Klaus G. and Folke Ölander (eds.), *Understanding Economic Behavior*, Kluwer Press, Boston.

Bean, Charles (2006), 'Globalisation and inflation', *Bank of England Quarterly Bulletin*, vol. 46, pp. 468–475.

Berge, Travis J. (2018), 'Understanding survey-based inflation expectations', *International Journal of Forecasting*, vol. 34, pp. 788–801.

Berger, Tino, Gerdie Everaert and Hauke Vierke (2016), 'Testing for time variation in an unobserved components model for the U.S. economy', *Journal of Economic Dynamics and Control*, vol. 69, pp. 179–208.

Blanchard, Olivier (2016), 'The Phillips curve: Back to the '60s?', *American Economic Review*, vol. 106, pp. 31–34.

Blaskowitz, Oliver and Helmut Herwartz (2014), 'Testing the value of directional forecasts in the presence of serial correlation', *International Journal of Forecasting*, vol. 30, pp. 30–42.

Campelo, Aloisio, Viviane Seda Bittencourt and Marco Malgarini (2020), 'Consumers confidence and households consumption in Brazil: Evidence from the FGV survey', *Journal of Business Cycle Research*, vol. 16, pp. 19–34.

Carroll, Christopher D., Jeffrey C. Fuhrer and David W. Wilcox (1994), 'Does consumer sentiment forecast household spending? If so, why?', *American Economic Review*, vol. 84, pp. 1397–1408.

Chen, Qiwei, Mauro Costantini and Bruno Deschamps (2016), 'How accurate are professional forecasts in Asia? Evidence from ten countries', *International Journal of Forecasting*, vol. 32, pp. 154–167.

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

Dreger, Christian and Konstantin A. Kholodilin (2013), 'Forecasting private consumption by consumer surveys', *Journal of Forecasting*, vol. 32, pp. 10–18.

Driver, Ciaran and Nigel Meade (2019), 'Enhancing survey-based investment forecasts', *Journal of Forecasting*, vol. 38, pp. 236–255.

Easaw, Joshy Z. and Saeed M. Heravi (2004), 'Evaluating consumer sentiments as predictors of UK household consumption behaviour. Are they accurate and useful?', *International Journal of Forecasting*, vol. 20, pp. 671–681.

European Commission (2016), *Joint Harmonised EU Programme of Business and Consumer Surveys. User Guide*, European Commission, Brussels.

Fitzgerald, Terry, Brian Holtemeyer and Juan Pablo Nicolini (2013), 'Is there a stable Phillips curve after all?', *Federal Reserve Bank of Minneapolis Economic Policy Paper* 13-6.

Gaiotti, Eugenio (2010), 'Has globalization changed the Phillips curve? Firm-level evidence on the effect of activity on prices', *International Journal of Central Banking*, vol. 6, pp. 51–84.

Greer, Mark (2003), 'Directional accuracy tests of long-term interest rate forecasts', *International Journal of Forecasting*, vol. 19, pp. 291–298.

Hansson, Jesper, Per Jansson and Mårten Löf (2005), 'Business survey data: Do they help in forecasting GDP growth?', *International Journal of Forecasting*, vol. 21, pp. 377–389.

Hjalmarsson, Erik and Pär Österholm (2017), 'Households' mortgage-rate expectations – More realistic than at first glance?', *Sveriges Riksbank Economic Review*, no. 2, pp. 57–64.

Hjalmarsson, Erik and Pär Österholm (2019), 'A micro-data analysis of households' expectations of mortgage rates', *Economics Letters*, vol. 185, Article 108693.

Hjalmarsson, Erik and Pär Österholm (2020), 'Heterogeneity in households' expectations of housing prices – Evidence from micro data', *Journal of Housing Economics*, vol. 50, Article 101731.

Hjalmarsson, Erik and Pär Österholm (2021), 'Anchoring in surveys of household expectations', *Economics Letters*, vol. 198, Article 109687.

IMF (2013), *World Economic Outlook*, April 2013.

Jönsson, Kristian (2020), 'Machine learning and nowcasts of Swedish GDP', *Journal of Business Cycle Research*, vol. 16, pp. 123–134.

Jonung, Lars (1981), 'Perceived and expected rates of inflation in Sweden', *American Economic Review*, vol. 71, pp. 961–968.

Jonung, Lars and David Laidler (1988), 'Are perceptions of inflation rational? Some evidence for Sweden', *American Economic Review*, vol. 78, pp. 1080–1087.

Karlsson, Sune and Pär Österholm (2020), 'A note on the stability of the Swedish Phillips curve', *Empirical Economics*, vol. 59, pp. 2573–2612.

Kuttner, Ken and Tim Robinson (2010), 'Understanding the flattening Phillips curve', *North American Journal of Economics and Finance*, vol. 21, pp. 110–125.

Leduc, Sylvain and Daniel J. Wilson (2017), 'Has the wage Phillips curve gone dormant?', *Federal Reserve Bank of San Francisco Economic Letter* 2017-30.

Neumann, Nico, Ulf Böckenholt and Ashish Sinha (2016), 'A meta-analysis of extremeness aversion', *Journal of Consumer Psychology*, vol. 26, pp. 193–212.

Newey, Whitney K. and Kenneth D. West (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica*, vol. 55, pp. 703–708.

Newey, Whitney K. and Kenneth D. West (1994), 'Automatic lag selection in covariance matrix estimation', *Review of Economic Studies*, vol. 61, pp. 631–653.

Occhino, Filippo (2019), 'The flattening of the Phillips curve: Policy implications depend on the cause', *Federal Reserve Bank of Cleveland Economic Commentary* 2019-11.

OECD (2000), *Main Economic Indicators. Sources and Definitions*, OECD, Paris.

Palmqvist, Stefan and Lena Strömberg (2004), 'Households' inflation opinions – A tale of two surveys', *Sveriges Riksbank Economic Review*, no. 4, pp. 23–42.

Pesaran, M. Hashem and Alan G. Timmermann (1992), 'A simple nonparametric test of predictive performance', *Journal of Business and Economic Statistics*, vol. 10, pp. 461–465.

Pierdzioch, Christian, Monique B. Reid and Rangan Gupta (2018), 'On the directional accuracy of inflation forecasts: Evidence from South African survey data', *Journal of Applied Statistics*, vol. 45, pp. 884–900.

Pinto, Santiago, Pierre-Daniel G. Sarte and Robert Sharp (2020), 'Learning about consumer uncertainty from qualitative surveys: As uncertain as ever', *International Journal of Central Banking*, vol. 16, pp. 47–99.

Sinclair, Tara M., Herman O. Stekler and Lindsay Kitzinger (2010), 'Directional forecasts of GDP and inflation: A joint evaluation with an application to Federal Reserve predictions', *Applied Economics*, vol. 42, pp. 2289–2297.

Sosvilla-Rivero, Simón and María del Carmen Ramos-Herrera (2018), 'Inflation, real economic growth and unemployment expectations: An empirical analysis based on the ECB survey of professional forecasters', *Applied Economics*, vol. 50, pp. 4540–4555.

Sveriges Riksbank (2007), Separate minutes of the Executive Board, Meeting on 29 October 2007.

Thomas, Lloyd B. (1999), 'Survey measures of expected U.S. inflation', *Journal of Economic Perspectives*, vol. 13, pp. 125–144.

Thomas, Lloyd B. and Alan P. Grant (2008), 'The accuracy and rationality of US and Australian household inflation forecasts: A comparative study of the Michigan and Melbourne Institute surveys', *Economic Record*, vol. 84, pp. 237–252.

Trehan, Bharat (2015), 'Survey measures of expected inflation and the inflation process', *Journal of Money, Credit and Banking*, vol. 47, pp. 207–222.

Appendix

In this appendix, we provide the details of the econometric test employed in Section 3 to assess the accuracy of the directional forecasts.

The directional forecast is denoted x_t and the actual directional change y_t. Note that any directional variable is Bernoulli distributed. We define

(1) $$P_x = \mathbb{P}(x_t = 1) \text{ and } P_y = \mathbb{P}(y_t = 1),$$

where $\mathbb{P}$ is the probability function. We further introduce the variable z_t which takes on the value 1 if the forecast is correct, and the value 0 if the forecast is wrong. By using the conjunction operator, $\wedge$, we write $z_t = 1$ if $x_t = 1 \wedge y_t = 1$ or $x_t = 0 \wedge y_t = 0$, and $z_t = 0$ if $x_t = 1 \wedge y_t = 0$ or $x_t = 0 \wedge y_t = 1$ (see panel C in Figures 3 and 4). The probability of $z_t = 1$ is thus given by

(2) $$P = \mathbb{P}(x_t = 1, y_t = 1) + \mathbb{P}(x_t = 0, y_t = 0).$$

Under the null hypothesis that x_t and y_t are independent – that is, if x_t has no power to predict y_t – then it follows from the definition of independence that the probability of $z_t = 1$ is given by

(3) $$P_{H0} = P_x P_y + (1 - P_x)(1 - P_y).$$

The probability P is efficiently estimated as the proportion of correct directional forecasts in a data set with T observations, and thus the estimate is given by

(4) $$\hat{P} = T^{-1} \sum_{t=1}^{T} z_t.$$

Under the null hypothesis of no predictive power, $T\hat{P}$ has a binomial distribution with expected value TP_{H0} and variance $TP_{H0}(1 - P_{H0})$. In the case in which P_x and P_y are known, one can use the approximate test for the Bernoulli parameter P. The test statistic is asymptotically standard normal and given by

(5) $$t = \frac{\hat{P} - P_{H0}}{\sqrt{T^{-1} P_{H0}(1 - P_{H0})}}.$$

For example, assuming a symmetric random walk behaviour implies $P_y = 0.5$ since an up-move and a down-move of the forecasted variable are equally likely. It is then natural for any forecast to have $P_x = 0.5$. In this case the test statistic simplifies to $t = \sqrt{T}(2\hat{P} - 1)$. However, in practice, P_x and P_y are not known and need to be estimated from sample data. Their efficient estimates are given by $\hat{P}_x = T^{-1} \sum_{t=1}^{T} x_t$ and $\hat{P}_y = T^{-1} \sum_{t=1}^{T} y_t$, and consequently P_{H0} is replaced by

(6) $$\hat{P}_{H0} = \hat{P}_x \hat{P}_y + (1 - \hat{P}_x)(1 - \hat{P}_y).$$

Pesaran and Timmermann (1992) derive that in this case the test of predictive performance of x_t can be based on

(7) $$t_{PT} = \frac{\hat{P} - \hat{P}_{H0}}{\sqrt{\widehat{\text{var}}(\hat{P}) - \widehat{\text{var}}(\hat{P}_{H0})}},$$

where $\widehat{\text{var}}(\hat{P}_{H0}) = T^{-1}(2\hat{P}_y - 1)^2 \hat{P}_x(1 - \hat{P}_x) + T^{-1}(2\hat{P}_x - 1)^2 \hat{P}_y(1 - \hat{P}_y) + 4T^{-2} \hat{P}_y \hat{P}_x(1 - \hat{P}_y)(1 - \hat{P}_x)$ and $\widehat{\text{var}}(\hat{P}) = T^{-1} \hat{P}_{H0}(1 - \hat{P}_{H0})$. The t_{PT} test statistic is asymptotically standard normal.

The suggested approach implicitly assumes that the forecast and actual process are serially independent. However, serial correlation if often present in economic applications. Blaskowitz and Herwartz (2014) suggest a Newey-West correction for the directional accuracy test we consider. First note that independence of Bernoulli variables x_t and y_t is equivalent to zero covariance between x_t and y_t. It then follows that

$$P - P_{H0} = 2\text{cov}(x_t, y_t), \tag{8}$$

where $\text{cov}(\cdot,\cdot)$ denotes the covariance operator.[12]

Consequently $\widehat{\text{cov}}(x_t, y_t) = 0$ if and only if $t_{PT} = 0$. We can thus alternatively test for zero covariance between the directional forecast x_t and the actual directional change y_t. We follow the exposition in Blaskowitz and Herwartz (2014) and decompose

$$x_t = P_x + w_t \text{ and } y_t = P_y + v_t, \tag{9}$$

where w_t and v_t are binary zero mean random errors which may be serially correlated. It follows that the null hypothesis of $\text{cov}(x_t, y_t) = 0$ is equivalent to $\mathbb{E}[w_t\, v_t] = 0$, where $\mathbb{E}[\cdot]$ denotes the expectation operator. To bring the model to data we estimate $\hat{w}_t = x_t - \hat{P}_x$, $\hat{v}_t = y_t - \hat{P}_y$, and $\overline{wv} = \widehat{\text{cov}}(x_t, y_t) = T^{-1}\sum_{t=1}^{T}\hat{w}_t\hat{v}_t$. The test of predictive performance is then based on the test statistic

$$t_{PT}^{NW} = \frac{\widehat{\text{cov}}(x_t, y_t)}{\sqrt{T^{-1}\hat{s}_T^{NW}}}, \tag{10}$$

which is asymptotically standard normal, and where $\hat{s}_T^{NW}$ is the heteroscedasticity and autocorrelation consistent variance estimator (Newey and West, 1987) for $\widehat{\text{cov}}(x_t, y_t)$. In particular,

$$\hat{s}_T^{NW} = \widehat{\text{cov}}(w_t v_t, w_t v_t) + 2\sum_{g=1}^{G}\left(1 - \tfrac{g}{G+1}\right)\widehat{\text{cov}}(w_t v_t, w_{t+g} v_{t+g}), \tag{11}$$

$$\widehat{\text{cov}}(w_t v_t, w_{t+g} v_{t+g}) = \sum_{t=1}^{T-g}(\hat{w}_t\hat{v}_t - \overline{wv})(\hat{w}_{t+g}\hat{v}_{t+g} - \overline{wv}), \tag{12}$$

and the truncation lag G is – as is commonly done in the literature – set equal to the integer part of $4(T/100)^{2/9}$, that is, we set $G = 5$; see Newey and West (1994) for a technical discussion.

12 Note that $\text{cov}(x_t, y_t) = \mathbb{E}[x_t y_t] - \mathbb{E}[x_t]\mathbb{E}[y_t] = \mathbb{P}(x_t = 1, y_t = 1) - \mathbb{P}(x_t = 1)\mathbb{P}(y_t = 1)$. Introduce $a_t = 1 - x_t$, $b_t = 1 - y_t$ and note that $\text{cov}(a_t, b_t) = \text{cov}(x_t, y_t)$. Since $P - P_{H0} = \mathbb{P}(x_t = 1, y_t = 1) + \mathbb{P}(x_t = 0, y_t = 0) - \mathbb{P}(x_t = 1)\mathbb{P}(y_t = 1) - \mathbb{P}(x_t = 0)\mathbb{P}(y_t = 0)$, the result follows.