Weighted Averages of Individual Measures of Uncertainty for the European Central Bank's Survey of Professional Forecasters

Víctor López-Pérez (Universidad Politécnica de Cartagena, Spain)[1]

Abstract

This paper explores to what extent the estimates of aggregate macroeconomic uncertainty calculated with survey data change when relatively higher weights are given to the data submitted by the most accurate forecasters. It is found that weighted estimates of aggregate uncertainty differ significantly from the simple averages used in the SPF literature. The differences are statistically significant and economically relevant. In particular, weighted estimates indicate a much larger increase in uncertainty than the simple averages since the start of the financial crisis. Moreover, while the unweighted estimates of aggregate uncertainty have stayed rather flat since 2010, the weighted estimates display significant variation. The latter is much more consistent with the shocks that have hit the euro area in the last five years, like the sovereign debt crisis and the recession in 2012. It is also more in line with the volatility displayed by the uncertainty indicators from financial markets, like the VSTOXX index.

Keywords: score, weights, uncertainty, Survey of Professional Forecasters, European Central Bank.

JEL classification: D80, D84.

[1] I thank Benjamin Beckers, Maximilian Podstawski and participants at the 10th Warsaw International Economic Meeting (Warsaw, 2015) and the Lupcon Finance and Economics Conference (Frankfurt, 2015) for helpful comments. E-mail: victor.lopez@upct.es. Telephone number: (+34) 868071259. Fax number: (+34) 968325781.

1. Introduction

Macroeconomic outcomes are the result of millions of decisions taken by economic agents worldwide, and the economic literature tries hard to understand the determinants of these decisions. One of these determinants is the degree of uncertainty in the economy, which is especially important for savings and investment decisions. If uncertainty is large, consumers are expected to save more for precautionary reasons (Caballero 1990) and risk-averse investors may delay irreversible investment plans (Leahy and Whited 1996).

The evolution of uncertainty over time is thereby of interest for academics and researchers, who need estimates of uncertainty to investigate the links between uncertainty and economic outcomes. It is also of interest for policy-makers, who need to closely monitor the available estimates of uncertainty, anticipate the effects on the economy of changes in uncertainty and take the appropriate policy actions to achieve their policy objectives (Bloom 2009).

Data from surveys may be useful to know the degree of macroeconomic uncertainty perceived by survey respondents. The density forecasts of macroeconomic variables by professional forecasters are particularly valuable as they combine the expertise of highly-skilled professionals with the heterogeneity of views that naturally comes from survey methods. Consequently, measures of uncertainty from surveys like the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters, the European Central Bank's Survey of Professional Forecasters (ECB's SPF) and the Bank of England's Survey of External Forecasters have gained prominence in economic and policy discussions during the recent past.

The existing literature has explored different measures of uncertainty constructed from survey data. Available measures of uncertainty are the standard deviation of point forecasts, typically known as "disagreement" (Neamtiu, Shroff, White and Williams 2014), the variance of the average individual density forecast (ECB 2014), the average standard deviation of the individual density forecasts (Giordani and Soederlind 2003), the root mean subjective variance or RMSV (Batchelor and Dua 1996), the implied RMSV (Boero, Smith and Wallis 2008), the average or median inter-quartile range of the individual density forecasts (Engelberg, Manski and Williams 2011, Rich, Song and Tracy 2012), the average entropy of the individual density forecasts (Rich and Tracy 2010, Wallis 2006), and the average Gini index of the individual density forecasts (López-Pérez 2015).

All these measures have one thing in common: they assign the same weight to the information obtained from each forecaster, without taking forecasting performance into consideration. Put differently, these estimates of aggregate uncertainty do not give more weight to the information submitted by a forecaster who always performs better than the average. Similarly, they do not give less weight to the responses by a forecaster that always underperforms.

Therefore, the unweighted estimates of aggregate uncertainty used in the literature may be biased because they may implicitly give too much weight to underperforming forecasters. These forecasters frequently use only one, two or three bins for their density forecasts, even when the forecast horizon is very long.[2] This behaviour is not consistent with the dynamics of the forecasted variables in the recent past. They probably send these "overconfident" density forecasts because they do not have time to compute model-based forecasts, they do not use the density forecasts for purposes other than the SPF, and they do not perceive any cost from submitting inaccurate density forecasts.[3]

---

[2] For instance, one third of the density forecasts of GDP growth two years ahead submitted between 1999 Q1 and 2014 Q3 have non-zero probabilities in three bins or less.

[3] Overconfidence is defined here as "excessive precision in one's beliefs" (Moore and Healy 2008).

The behaviour of these forecasters may contaminate the estimates of aggregate uncertainty computed with the SPF data. In particular, it may lead to too small changes in the simple average of the individual estimates of uncertainty because a subset of forecasters keeps submitting density forecasts with too low probability in the tails, irrespectively of the changes to the economic environment. This may be the reason why the unweighted estimates of aggregate uncertainty tend to be rather flat over relatively long periods.

Take for instance the period since 2010. The unweighted estimates of aggregate uncertainty stay basically unchanged until 2014 (see the black solid lines in Figure 1). Is it reasonable to believe that uncertainty has not changed much when the euro area economy went from robust growth around 2% in 2010 to a recession in 2012? Is it reasonable to believe that uncertainty has not changed when the sovereign debt crisis erupted in 2010, when the Greek sovereign debt was restructured in 2012 or when the president of the ECB announced in June 2012 that the central bank would protect the euro area "whatever it takes"? Based on measures of uncertainty derived from financial market data, the answer is a resounding no. See for instance the behaviour of an indicator of financial market volatility in the euro area, the VSTOXX index (Figure 2). It is anything but constant after 2010 and suggests that the unweighted estimates of aggregate uncertainty obtained from the ECB's SPF density forecasts may be sending wrong signals.

One possible solution is to ignore the SPF density forecasts and use measures of uncertainty that rely on point forecasts only (e.g. disagreement). This solution, however, throws away very valuable information embedded in the density forecasts submitted by survey participants with very good forecasting records. The alternative is to discount more the information submitted by the underperforming forecasters while giving more weight to the forecasters that perform better than the average.

In this paper, I assign performance-based weights to the information obtained from each forecaster. Forecasting performance is assessed with a strictly proper scoring rule: the logarithmic score of the density forecasts (Gneiting and Raftery 2007). The forecasters with better scores are assigned higher weights. I use these weights to compute estimates of aggregate uncertainty as weighted averages of individual measures of uncertainty.

I am not aware of any other study that has applied this procedure to construct estimates of uncertainty from survey data. The literature on forecast combination has explored the use of performance-based weights to obtain weighted averages of point forecasts (Stock and Watson 2004, Capistrán and Timmermann 2009, Smith and Wallis 2009, Genre, Kenny, Meyler and Timmermann 2013). Others have used weights to obtain aggregate density forecasts as combinations of individual density forecasts (Hall and Mitchell 2005 and 2007, Geweke and Amisano 2010, Jore, Mitchell and Vahey 2010, Kascha and Ravazzolo 2010, Billio, Casarin, Ravazzolo and van Dijk 2013). What this paper proposes is that, for the measurement and estimation of uncertainty, performance-based weights may be used to obtain weighted averages of individual measures of uncertainty. Therefore, individual density forecasts are not combined into an aggregate density forecast in this paper. Instead, an individual estimate of uncertainty is obtained from each individual density forecast and then a weighted average of them is computed. This exercise is conducted with ECB's SPF data from 1999 Q1 to 2014 Q3.

Is there any reason to believe that the simple average of the individual measures of uncertainty might be different than a weighted average based on forecasting performance? Yes, if there is correlation between forecast performance and individual uncertainty. Kenny, Kostka and Masera (2014), found a positive relationship at the individual level between forecasting performance and the variance of the density forecasts: the best-performing respondents submitted density forecasts with more probability in the tails. Moreover, Boero, Smith and Wallis (2015) found that the degree of uncertainty embedded in the density forecasts at the individual level is very persistent.

These two results suggest that the most uncertain or "prudent" forecasters tend to outperform the rest. If the highest weights are given to these "prudent" forecasters when aggregating their individual estimates of uncertainty, the resulting estimates of aggregate uncertainty may be higher than the unweighted estimates.

As a preview of the main result of the paper, weighted estimates of aggregate uncertainty differ significantly from the simple averages used in the SPF literature. The differences are statistically significant and economically relevant. In particular, weighted estimates indicate a much larger increase in uncertainty than the simple averages since the start of the financial crisis. Moreover, while the unweighted estimates of aggregate uncertainty have stayed rather flat since 2010, the weighted estimates display significant variation. The latter is much more consistent with the shocks that have hit the euro area in the last five years, like the sovereign debt crisis and the recession in 2012. It is also more in line with the volatility displayed by the uncertainty indicators from financial markets, like the VSTOXX index.

The rest of the paper is structured as follows. Section 2 describes the data used in the analysis. Section 3 presents the estimates of aggregate uncertainty from SPF data when the individual measures of uncertainty are weighted according to each participant's forecasting performance. Section 4 concludes, discussing the implications of the findings and proposing directions for future research.

2. The data

To obtain performance-based weights, I need to compute scores by forecaster. Therefore, I need to compare the forecasts with the realisations of the forecasted variables.

2.1. Density forecasts

This paper is about measuring uncertainty and thus I will use data from density forecasts. The density forecasts used in this paper are obtained from the ECB's SPF, which has been conducted quarterly since 1999 Q1. 102 forecasters have participated at least once in the survey, although average participation is around 60 forecasters per round. The panel is unbalanced, as many forecasters do not reply sometimes, while others have left the panel and have been replaced with new panellists.

The SPF surveys density forecasts of the year-on-year inflation rate, the year-on-year GDP growth rate, and the unemployment rate, all for the euro area. Forecasters are asked to distribute probabilities among a set of predefined bins for each variable.[4] The forecast horizons used in this paper are rolling horizons one and two years ahead. Therefore, the SPF provides data on the subjective probabilities that individual forecasters assigned to different macroeconomic events.[5] For instance, the data from the ECB's SPF webpage indicates that, in October 2013, forecaster number 1 assigned 70% probability to the September 2014 inflation rate in the euro area being between 1.5% and 1.9%.

---

[4] See http://www.ecb.europa.eu/stats/prices/indic/forecast/html/index.en.html for a full description of the survey.

[5] Details on the bins available to the SPF forecasters and on the forecast horizons surveyed in each SPF round can be obtained from the document "ECB Survey of Professional Forecasters (SPF): description of SPF dataset", available here: http://www.ecb.europa.eu/stats/prices/indic/forecast/shared/files/dataset_documentation_csv.pdf??8b0b9ba730b2241d43fec92dacd2944d.

For the analysis of SPF density forecasts, an assumption needs to be made on the probability placed in open-ended bins (e.g. less than 0.0%). These bins are much less informative than closed bins (e.g. from 0.0% to 0.4%) for the measurement of uncertainty because of their infinite width. Previous studies have typically assumed that open-ended bins have the same or double width than closed bins.[6] This assumption may lead to underestimate uncertainty, especially if open-ended bins contain relatively large probabilities. For instance, forecaster 52 in round 2009 Q1 assigned 100% probability to GDP growth one year ahead being lower than -1.0%. Given that the width of the closed bins in the ECB's SPF is 0.5%, is it reasonable to assume that he/she assigned 100% probability either to the [-1.5%,-1.1%] bin or to the [-2.0%,-1.1%] bin when his/her point forecast was -2.9%?

To avoid drawing wrong inference from uninformative data, any density with at least one open-ended bin that cannot *reasonably* have the same width than the closed bins is removed from the sample. Specifically, any density with 50% more probability in an open-ended bin than in any other bin with non-zero probability is excluded. This resulted in the exclusion of 224 inflation densities, 492 GDP-growth densities and 165 unemployment densities for the sample used in this paper: 1999 Q1 – 2014 Q3.[7,8] For the remaining density forecasts, 5265 for inflation, 4861 for GDP growth and 4847 for unemployment, open-ended bins are treated as having the same width than closed bins. Thus, the next sections do not make any distinction between closed and open-ended bins: all bins are assumed to be closed and assumed to have the same width.

From each of these density forecasts, I compute an estimate of individual uncertainty. In previous work I have proposed the Gini index of a density forecast as a measure of individual uncertainty. This measure is better than others, like the variance of the density forecast or its interquartile range, because of two reasons. First, the Gini index take its maximum value when the forecaster allocates the same probability to every bin (i.e. when the degree of uncertainty is the highest). Second, the Gini index does not require to make any assumption on how the probability allocated to each bin is distributed within the bin (see López-Pérez 2015 for details).

Borrowed from the literature on income and wealth inequality, the Gini index (Gini 1955) is based on the Lorenz curve (Lorenz 1905). This curve is typically used to represent how much wealth is in the hands of the poorest x% of the population. The Lorenz curve may also be applied to the measurement of uncertainty by representing the cumulative probability allocated to the x% least likely bins of a density forecast. Under minimum uncertainty, with 100% probability placed in a single bin, the Lorenz curve would be zero from the first bin to the one before the last. Then, it would jump to 100% in the last bin. Under maximum uncertainty, with the same probability allocated to every bin, the Lorenz curve would increase in regular steps from the first bin to the last.

From the Lorenz curve of the individual density forecast, the calculation of the individual Gini index is straightforward. The Gini index of a density forecast is the distance between the Lorenz curve under maximum uncertainty and the Lorenz curve of the density forecast:

$$G = -\frac{\sum_{k=1}^{K}(x_k - lc_k)}{\sum_{k=1}^{K} x_k} \qquad [1]$$

[6] See Batchelor and Dua (1996) for an example.
[7] For completeness, densities with 100% probability in open-ended bins are also excluded.
[8] Densities that verified this criterion but have less than 1% probability in open-ended bins are not removed from the sample. Otherwise, all computer-generated densities with support from -∞ to ∞ (e.g. a normal density function) would be excluded.

where K is the number of bins, x is the nx1 vector (1/n, 2/n,…, 1)', and lc is the nx1 vector of ordinates from the Lorenz curve of the density forecast. As the original Gini index would decline with uncertainty, its sign has been changed to make it an increasing function of uncertainty.

2.2. Realisations of the forecasted variables

The realisations of the forecasted variables are retrieved from the ECB's Euro Area Real-Time Database.[9] This database collects vintages of many macroeconomic variables as they appeared in each issue of the ECB Monthly Bulletin.[10]

Real-time data is used in this paper because ECB's SPF participants tried to forecast inflation, GDP growth and unemployment as defined by the statistical methodology existing at the time of the production of the forecast. If the latest vintage of data were used instead, differences between forecasts and realisations would not only arise because of forecast errors. They may also be caused by subsequent methodological changes to the calculation of the forecasted variable that led to backward revisions in the historical time series of the variable.

At the time of retrieving the data from the Real-Time Database (July 2015), inflation data was available until September 2014, real GDP growth data until 2014 Q3, and unemployment data until August 2014. Due to the length of the forecast horizons in the ECB's SPF, scores may be computed up to 2013 Q4 and 2012 Q4 for density forecasts of inflation one and two years ahead respectively, 2014 Q1 and 2013 Q1 for density forecasts of GDP growth one and two years ahead, and 2013 Q4 and 2012 Q4 for density forecasts of unemployment one and two years ahead.

3. Weighted averages of individual measures of uncertainty for the ECB's SPF

This section presents estimates of aggregate uncertainty computed with data from the ECB's SPF. These estimates are weighted averages of individual measures of uncertainty, instead of the simple unweighted averages used in the literature. The weight assigned to each forecaster is based on his/her forecasting performance, which is assessed by the logarithmic scoring rule. This rule is one of the four strictly proper scoring methods highlighted by Gneiting and Raftery (2007).

The logarithmic score is one if not the most popular scoring rule (see Hall and Mitchell 2005 and Kascha and Ravazzolo 2010, for examples of uses of the logarithmic score with survey data). The logarithmic score by forecaster i in period t is:

$$S_{it} = \log p_{rit} \qquad [2]$$

where p_{rit} is the probability assigned by the forecaster to the bin that includes the realisation of the forecasted variable. The logarithmic score takes a value of zero if the forecaster assigned all the probability to the bin where the realisation fell and takes a value of minus infinity if the forecaster placed zero probability in that bin. Each forecaster has a different score for each forecasted variable and forecast horizon, i.e. the score for inflation forecasts one year ahead by forecaster 1 is likely to differ from his/her score for inflation forecasts two years ahead.

[9] http://sdw.ecb.europa.eu/browse.do?node=4843526
[10] See Giannone, Henry, Lalik and Modugno (2010) for details.

Individual weights will be assigned on the basis of each forecaster's average performance over a period of time. Average performance by forecaster *i* at time *t* is computed as the simple average of his/her individual scores from period *t-W* to period *t*.

$$\overline{S}_{it} = \frac{1}{1+W}\sum_{w=0}^{W} S_{it-w} \qquad \text{with } W \geq 0 \qquad [3]$$

For instance, if W=1, the individual average score is computed over the current and the previous survey rounds. Obviously, the individual average score cannot be computed for the first W survey rounds. If a forecaster has not participated in all the 1+W survey rounds, the average is computed over the rounds when he/she participated. The weight assigned to forecaster *i* in survey round *t* is assumed to be: [11]

$$W_{it} = \frac{e^{\overline{S}_{it-h}}}{\sum_{j=1}^{J} e^{\overline{S}_{jt-h}}} \qquad [8]$$

J is the number of forecasters that participated in survey round *t*. This guarantees that the sum of the weights of the participating forecasters equals one. *h* is the length of the forecast horizon (in quarters) of the variable of interest: *h*=4 for one-year-ahead forecasts and *h*=8 for two-years-ahead forecasts. This guarantees that the weights can be computed in real time because the scores are available after four and eight quarters for one- and two-years ahead forecasts respectively.

Intuitively, the best forecasters, i.e. those with the highest average logarithmic scores, will receive the highest weights. These weights may then be used to obtain an estimate of aggregate uncertainty as the weighted average of the individual measures of uncertainty. This is the point at which this paper departs from the existing SPF literature, which uses *simple* averages. Figure 1 shows the resulting estimates of aggregate uncertainty for the three main macroeconomic variables (the expected inflation rate, the expected rate of growth of real GDP and the expected unemployment rate) surveyed in the ECB's SPF. As two different forecast horizons are considered in this paper (one and two years ahead), six different charts appear in Figure 1.

Each chart shows a solid black line, which is the time series of the simple average of the individual Gini indices of uncertainty. The two dashed lines are the bounds of the 95% confidence interval around the simple average. The red line is the weighted average of the individual Gini indices of uncertainty.

For the estimates of uncertainty based on density forecasts of inflation, the weighted and unweighted estimates are similar until the start of the financial crisis. Then, the weighted estimates indicate a much larger increase in uncertainty than the unweighted estimates. After the initial increase, the unweighted estimates stay rather flat, while the weighted estimates indicate a sharp decrease in uncertainty, followed by another, more moderate increase around 2011-2012 when fears of a potential break-up of the euro area mounted. After the president of the ECB announced in the summer of 2012 that the institution would do "whatever it takes" to preserve the integrity of the euro area, the weighted estimates of uncertainty declined significantly. None of these movements can be observed in the unweighted estimates.

[11] The results are also robust to the use of the optimal weights that minimise the Kullback-Leibler information criterion as in Conflitti, De Mol and Giannone 2012. These results are available from the author upon request.

The estimates of uncertainty based on density forecasts of unemployment show similar dynamics. The unweighted estimates stay rather flat after the initial increase in 2008-2009. The weighted estimates show more interesting dynamics, with uncertainty declining sharply in 2010 when GDP in the euro area started to grow around 2% a year. Uncertainty increased again in 2012 when the euro area economy went into a recession. Finally, the weighted estimates of uncertainty fell again in 2014 when GDP growth climbed above zero again. Weighted estimates of uncertainty based on density forecasts of GDP growth show a similar picture, although those based on one year ahead forecasts are more volatile.

Interestingly, the weighted estimates of uncertainty track much better than the unweighted ones the dynamics of uncertainty extracted from financial market indicators. Figure 2 shows the standardised 12 and 24-month VSTOXX indices of stock market volatility in the euro area. The M-shaped dynamics of these indices between 2008 and 2013 are remarkably similar to those of the weighted estimates of uncertainty described above. The unweighted estimates are not able to replicate these fluctuations.

The main results described above are robust to the value of W. I have tried with values ranging from zero to eleven quarters. Very high values of W increase the chances of giving zero weight to most forecasters, while low values of W induce more volatility in the individual weights over time. The values used in Figure 1 try to balance these two effects: W is set equal to 1 for inflation forecasts one year ahead and unemployment forecasts two years ahead, 4 for unemployment forecasts two years ahead, 5 for inflation and GDP-growth forecasts two years ahead, and 6 for GDP-growth forecasts one year ahead. Results for different values of W are available from the author upon request.

Even with these relatively low values of W, there are a few occasions when the average score is minus infinity for all the participants in a survey round. These episodes are concentrated around the start of the financial crisis, which most participants did not foresee. The weighted estimates of uncertainty based on the logarithmic score cannot be computed in such occasions and are replaced by linear interpolations between the previous and the next weighted estimates available. Weighted estimates based on different proper scoring rules, like the Brier score, the rank-probability score or the spherical score do not suffer from this drawback and the main results of the paper are robust to the scoring rule used.

An interesting feature of the weighted estimates of aggregate uncertainty is that their volatility is higher than the volatility of the unweighted estimates. This is because of two reasons. First, relative performance is changing over time, especially during the most turbulent periods. It is during these periods when the best forecasters outperform the rest. Consequently, the weights assigned to the best forecasters tend to increase in more turbulent times. In these periods, the weighted estimate of aggregate uncertainty deviates from the simple average and moves closer to the individual measures of uncertainty by the best forecasters.

The second reason for the higher volatility of the weighted estimates compared to the unweighted estimates is that changes in participation from one round to the next cause changes in the weights. This is because the denominator in equation [8] varies, even if relative performances did not change much. In previous work, I found that the unbalanced nature of the SPF panel of participants adds volatility to the estimates of aggregate uncertainty (López-Pérez, 2015). The changes in the weights used here amplify this effect.

Another interesting result from the comparison between weighted and unweighted estimates of aggregate uncertainty is that, when the former deviates from the latter, it is mostly to indicate higher uncertainty. Figure 1 shows how frequently the weighted estimate of aggregate uncertainty crosses the upper bound of the 95% confidence interval around the simple average. However, the weighted estimate very rarely crosses the lower bound of the confidence interval. In other words, simple averages of individual measures of uncertainty computed with data from

the ECB's SPF may produce estimates of aggregate uncertainty that are frequently too low. The reason is that the same-weight rule assigns too much weight to "overconfident" forecasters. These forecasters' predictions are typically worse than the average forecast, reducing the signal-to-noise ratio of the estimations of aggregate uncertainty. When the weight assigned to these forecasters is lowered according to their forecasting performance, the estimates of aggregate uncertainty are frequently significantly higher, both statistically and economically.

4. Conclusion

The SPF literature has developed many different measures of aggregate uncertainty but all of them assign the same weight to each forecaster's individual estimates of uncertainty, irrespectively of his/her forecasting performance. This paper deviates from the existing literature and assigns a higher weights to the individual estimates of uncertainty by the most accurate forecasters, and lower weights to the individual estimates by the worst forecasters.

The main result of the paper is that weighted estimates of aggregate uncertainty differ significantly from the simple averages used in the SPF literature. The differences are statistically significant and economically relevant. In particular, weighted estimates indicate a much larger increase in uncertainty than the simple averages since the start of the financial crisis. Moreover, while the unweighted estimates of aggregate uncertainty have stayed rather flat since 2010, the weighted estimates display significant variation. The latter is much more consistent with the shocks that have hit the euro area in the last five years, like the sovereign debt crisis and the recession in 2012. It is also more in line with the volatility displayed by the uncertainty indicators from financial markets, like the VSTOXX index.

These results have relevant implications for the analysis of uncertainty with the ECB's SPF dataset. The unweighted estimates of aggregate uncertainty used in the literature may be biased because they implicitly give too much weight to underperforming forecasters. These forecasters frequently use only one or two bins for their density forecasts, even when the forecast horizon is very long. This behaviour is not consistent with the dynamics of the forecasted variables in the recent past. They probably send these "overconfident" density forecasts because they do not have time to compute model-based forecasts, they do not use the density forecasts for purposes other than the SPF, and they do not perceive any cost from submitting inaccurate density forecasts.

The behaviour of these forecasters contaminates the estimates of aggregate uncertainty computed with SPF data. In particular, it leads to too small changes in the simple average of the individual estimates of uncertainty because a subset of forecasters keep submitting density forecasts with too low probability in the tails, irrespectively of the changes to the economic environment. That is the reason why the unweighted estimates of aggregate uncertainty tend to be rather flat over relatively long periods.

One possible solution is to ignore the SPF density forecasts and use measures of uncertainty that rely on point forecasts only (disagreement). This solution, however, throws away the very valuable information embedded in the density forecasts submitted by survey participants with very good forecasting records. The alternative is to use the kind of weighted estimates of aggregate uncertainty described in this paper, with the aim to discount more the information submitted by the underperforming forecasters while giving more weight to the forecasters that perform better than the average.

Future research would revisit the link between uncertainty and macroeconomic outcomes using weighted estimates of aggregate uncertainty. In particular, the robustness of the existing results in the literature may be explored when weighted estimates of aggregate uncertainty are used instead of simple averages.

5. Bibliography

Batchelor, R.A., Dua, P. (1996). Empirical measures of inflation uncertainty: a cautionary note. Applied Economics, 28, 333-341.

Billio, M., Casarin, R., Ravazzolo, F., van Dijk, H.K. (2013). Time-varying combinations of predictive densities using nonlinear filtering. Journal of Econometrics, 177, 213-232.

Bloom, N. (2009). The impact of uncertainty shocks. Econometrica, 77, 623-685.

Boero, G., Smith, J., Wallis, K.F. (2008). Uncertainty and disagreement in economic prediction: the Bank of England Survey of External Forecasters. The Economic Journal 118, 1107-1127.

Boero, G., Smith, J., Wallis, K.F. (2015). The measurement and characteristics of professional forecasters' uncertainty. Forthcoming in the Journal of Applied Econometrics.

Caballero, R.J. (1990). Consumption puzzles and precautionary savings. Journal of Monetary Economics 25, 113-136.

Capistrán, C., Timmermann, A. (2009). Forecast combination with entry and exit of experts. Journal of Business and Economic Statistics, 27, 428-440.

Conflitti, C., De Mol, C., Giannone, D. (2012). Optimal combination of survey forecasts. ECARES Working Paper, 2012-023.

ECB (2014). Fifteen years of the ECB Survey of Professional Forecasters. ECB Monthly Bulletin, January, 55-67.

Engelberg, J., Manski, C.F., Williams, J. (2011). Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition. Journal of Applied Econometrics, 26, 1059-1078.

Geweke, J., Amisano, G. (2010). Optimal prediction pools. Journal of Econometrics, 164, 130–141.

Genre, V., Kenny, G., Meyler, A., Timmermann, A. (2013). Combining expert forecasts: can anything beat the simple average? International Journal of Forecasting, 29, 108-121.

Giannone, D., Henry, J., Lalik, M., Modugno, M. (2010). An area-wide real-time database for the euro area. ECB Working Paper, 1145.

Gini, C. (1955). Memorie di metodologica statistica. Rome. EV Veschi.

Giordani, P., Soederlind, P. (2003). Inflation forecast uncertainty. European Economic Review, 47, 1037–59.

Gneiting, T., Raftery, A.E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 477, 359-378.

Hall, S.G., Mitchell, J. (2005). Evaluating, comparing and combining density forecasts using the KLIC with an application to the Bank of England and NIESR 'fan' charts of inflation. Oxford Bulletin of Economics and Statistics, 67, 995-1033.

Hall, S.G., Mitchell, J. (2007). Combining density forecasts. International Journal of Forecasting, 23, 1–13.

Jore, A.S., Mitchell, J., Vahey, S.P. (2010). Combining forecast densities from VARs with uncertain instabilities, Journal of Applied Econometrics, 25, 621–634.

Kascha, C., Ravazzolo, F. (2010). Combining inflation density forecasts. Journal of Forecasting, 29, 231-250.

Kenny, G., Kostka, T., Masera, F. (2014). Density characteristics and density forecast performance: a panel analysis. Empirical Economics, doi: 10.1007/s00181-014-0815-9.

Leahy, J.V., Whited, T.M. (1996). The effect of uncertainty on investment: some stylized facts. Journal of Money, Credit and Banking, 28, 64-83.

López-Pérez, V. (2015). Measures of macroeconomic uncertainty for the ECB's survey of professional forecasters. In Donduran, M., Uzunöz, M., Bulut, E., Cadirci, T.O., Aksoy, T. (eds.), Proceedings of the 1st Annual International Conference on Social Sciences (600-614). Also available at http://hdl.handle.net/10317/4248.

Lorenz, M.O. (1905). Methods of measuring the concentration of wealth. Publications of the American Statistical Association, 9, 209–219.

Moore, D.A., Healy, P.J. (2008) The trouble with overconfidence. Psychological Review, 115, 502–517.

Neamtiu, M., Shroff, N., White, H.D., Williams, C.D. (2014). The impact of ambiguity on managerial investment and cash holdings. Journal of Business Finance & Accounting, 41, 1071-1099.
Rich, R., Song, J., Tracy, J. (2012). The measurement and behavior of uncertainty: evidence from the ECB Survey of Professional Forecasters. Federal Reserve Bank of New York Staff Report, 588.
Rich, R., Tracy, J. (2010). The relationships among expected inflation, disagreement, and uncertainty: evidence from matched point and density forecasts. Review of Economics and Statistics, 92, 200-207.
Smith, J., Wallis, K.F. (2009). A simple explanation of the forecast combination puzzle. Oxford Bulletin of Economics and Statistics, 71, 331-355.
Stock, J.H., Watson, M.W. (2004). Combination forecasts of output growth in a seven-country data set. Journal of Forecasting, 26, 405-430.
Wallis, K.F. (2006). A note on the calculation of entropy from histograms. MPRA Paper, 52856.

6. Appendix: Figures

Fig. 1 Simple and weighted averages of individual measures of uncertainty

a) From density forecasts of inflation one year ahead

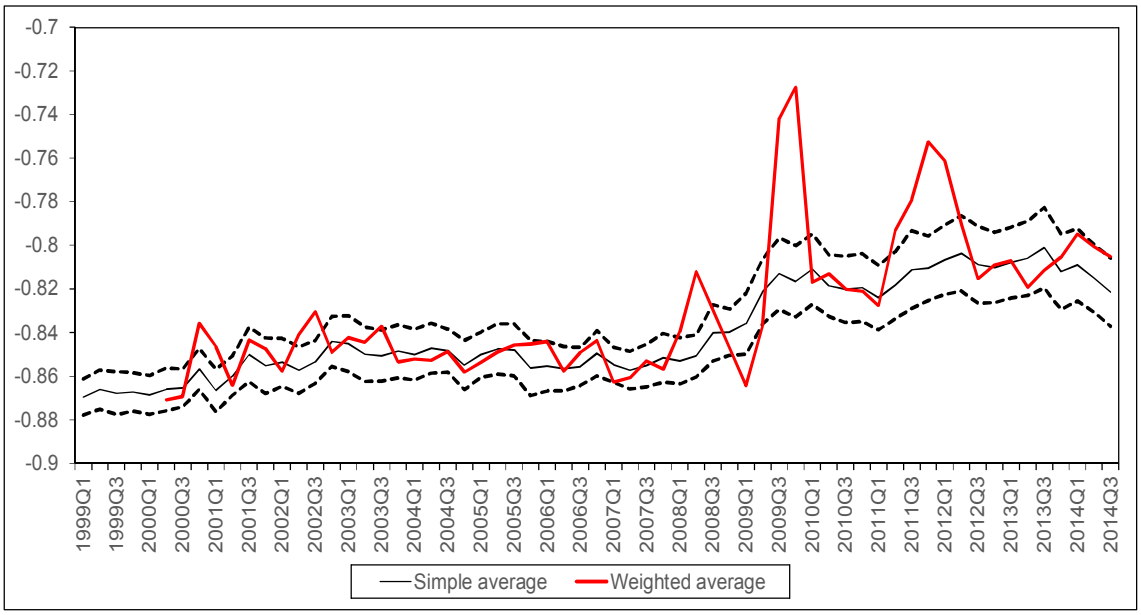

b) From density forecasts of inflation two years ahead

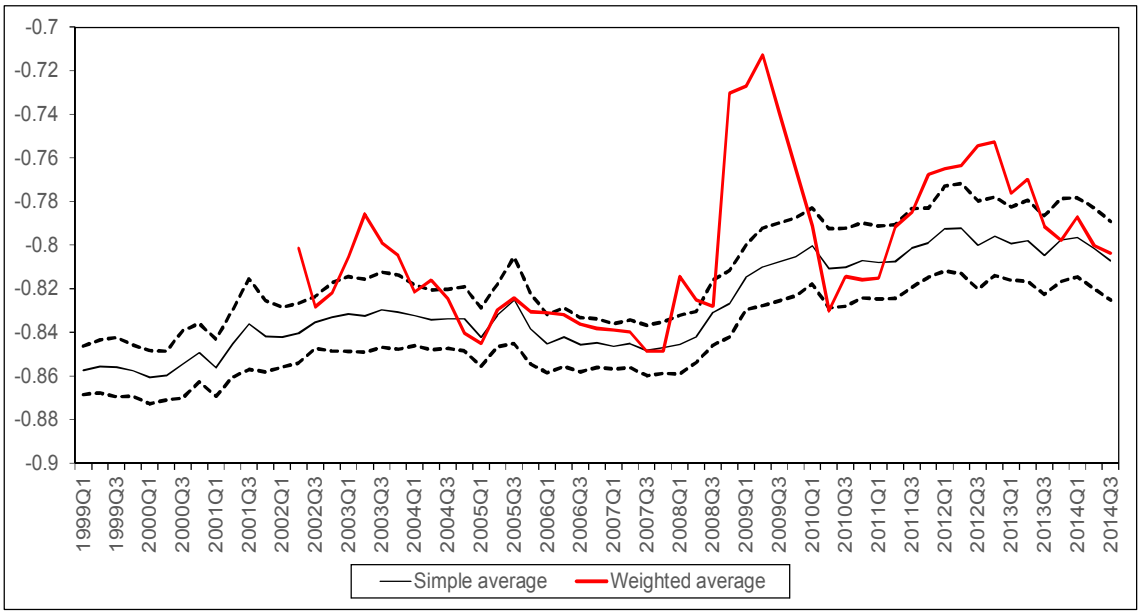

c) From density forecasts of GDP growth one year ahead

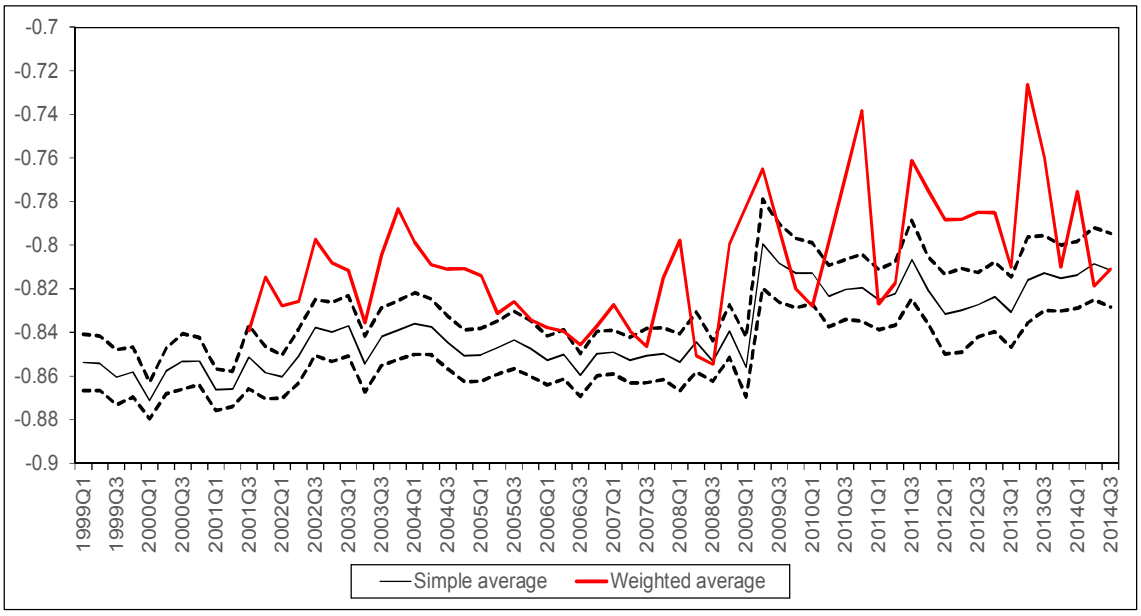

d) From density forecasts of GDP growth two years ahead

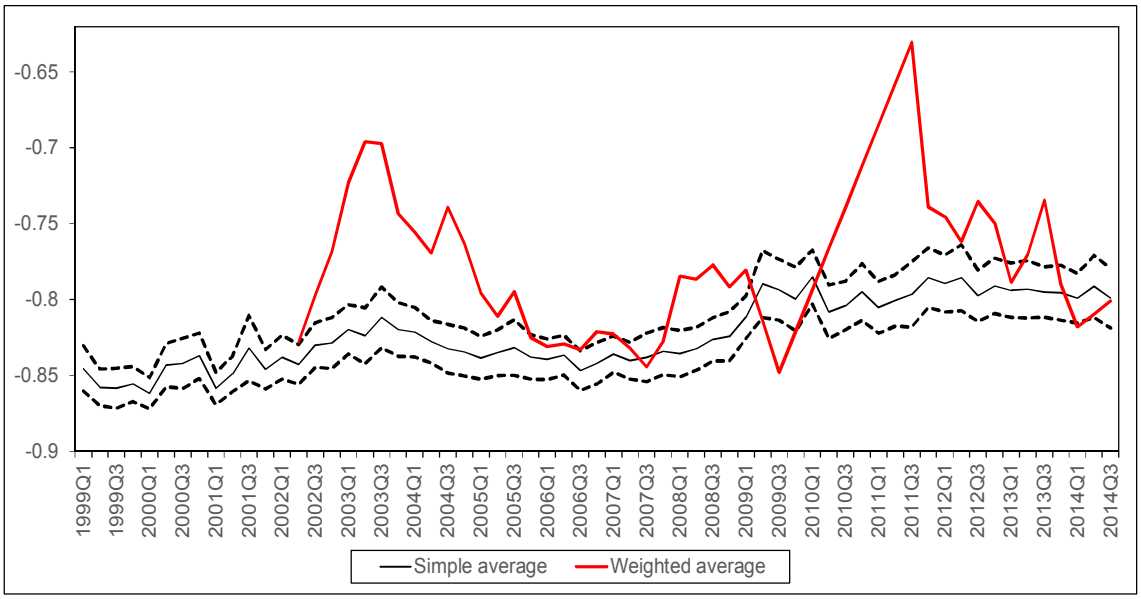

e) From density forecasts of unemployment one year ahead

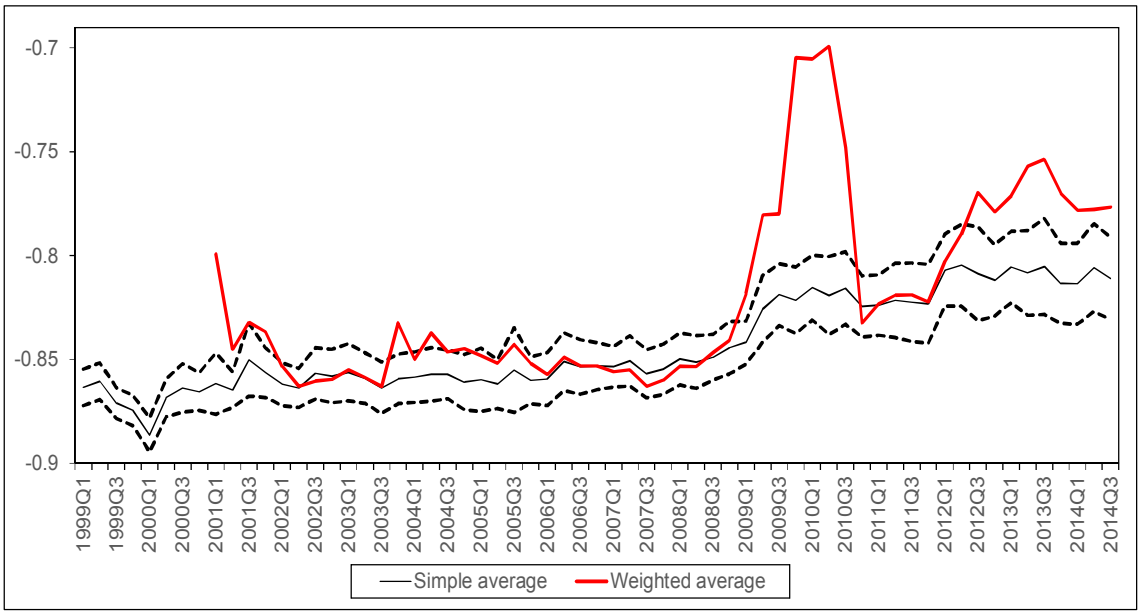

f) From density forecasts of unemployment two years ahead

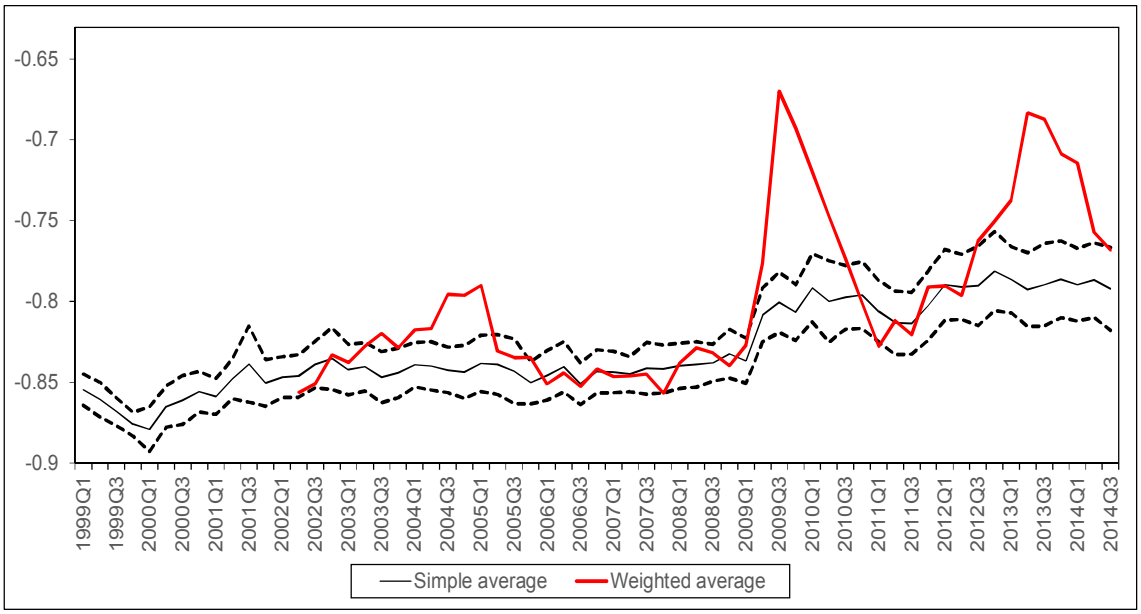

Fig. 2 Standardised estimates of uncertainty from European financial markets

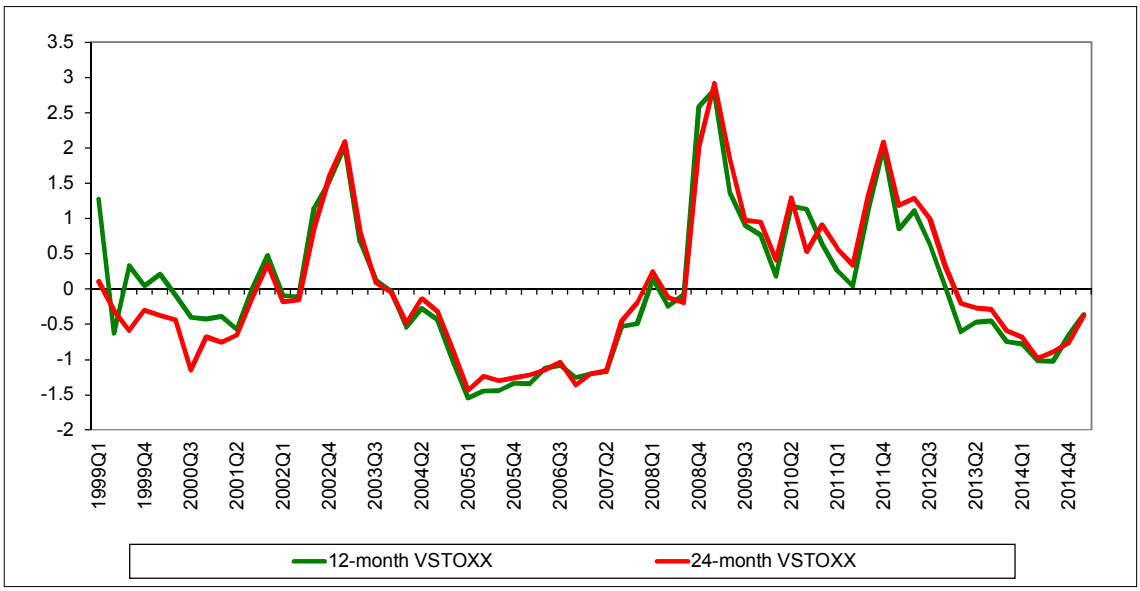

Source: stoxx.com and own calculations.