Erasmus University Rotterdam
Erasmus School of Economics
Econometrics and Operations Research
Bachelor's Thesis

Combining Macroeconomic Forecasts

A Quantile Regression Combination Approach

Author: Dorian James Hein
Student Number: 471379
Supervisor: Prof. Dr. DJC van Dijk
Second Assessor: JA Oorschot
Date of Submission: 4th of July, 2021

Abstract

Quantile regressions are performed on key macroeconomic variables to assess the risk of substantial economic change. Additionally, probabilistic densities are fitted to account for a larger range of probabilistic outcomes. The inclusion of M3 money growth and the University of Michigan Inflation Survey as conditioning variables is facilitated to improve forecast accuracy. The inclusion of the additional variables leads to a stronger accommodation of economic risk in the models and more accurate density forecasts. Optimally weighed forecast combination based on tick-loss minimization is preferred in some cases and can improve constructed densities. They slightly outperform the benchmark average combination, measured by predictive scoring. The out-of-sample analysis finds that constructing optimal weights works well for GDP growth and unemployment. Improvements for inflation are less pronounced, possibly due to limited initialization of the sample. Comparatively small differences in loss-function scores warrant further investigation of equally weighed quantile forecasts.

Contents

1 Introduction

In this bachelor's thesis, I answer the question:

What role do money supply and the University of Michigan inflation survey play in improving quantile forecasts of macroeconomic variables?

In the light of COVID-19, policymakers around the globe have taken unprecedented steps to support economies that were affected by stay-at-home orders, forced closures, and lockdowns. In particular, the most developed nations have poured trillions of dollars of stimulus into their economies with the goal of preventing economic recession or crisis. With the introduction of this stimulus, largely based on newly created money, voices of concern have grown with regards to inflation and unemployment. Macroeconomic theory suggests that the new creation of money is the most significant factor causing inflation, as Leeper writes that "if [an] economy is at or near potential, a stimulus to demand raises inflation through the usual Phillips curve mechanisms" [1]. On the forefront of this stand the United States with a 39.1 percent money supply growth year-on-year. When considering unemployment, the US jobs report, which is released on the first Friday every month, has become a pivotal point for investors to guide their intuition to the overall health of the US economy. Some voices of concern have appeared in the United States with regards to inflation, for example the New York Times pointing out the "steepest year-over-year jump [of the Consumer Price Index] in 13 years" [2]. In this work, I investigate the role of money growth on other macroeconomic variables.

Macroeconomic variables like inflation, unemployment, and the risks surrounding their future development are of crucial interest to policy makers and the private sector. For a private business, for example, the state of the economy is an important factor with regards to the decision of hiring additional workers. Financial institutions pay close attention to these variables to assess the health of the economy and make decisions about future policies like quantitative easing or security purchases. In connection to that, accurate forecasts of these variables play an integral role.

To assess up- and downside risk as well as uncertainty in these variables, I employ quantile regression, as outlined in the 2019 paper of Adrian et al. [3]. The regression is based around the Survey of Professional Forecasters' (SPF) median consensus forecasts, conditional on the information at the time when the forecast is made. Historically, economic models to forecast these variables provided point forecasts that did not consider the variables' entire distribution and were largely based on vector autoregressive models. Recent literature suggests that the inclusion of financial conditions like bond spreads can improve the accuracy of forecasts of these variables. In their 2021 paper, Adrian et al. find that the inclusion of financial conditions in the conditioning variables leads to significant improvements in the forecast accuracy of GDP growth and unemployment and modest improvements for inflation [4]. I replicate their findings and extend their approach in two ways: First, connecting to the role that growth of money supply and different inflation surveys might have on inflation, I include them as additional conditioning variables in the models.

The second differentiating factor is my use of a forecast combination technique, where one forecast is crystallize from different individual ones. Because many individual researches or investors might arrive at different forecasts for different reasons, a combination of their forecasts might account for uncertainty or asymmetric information. This leads to the second part of the research question, namely

What are the benefits of combining models with different conditioning variables in a quantile regression setting?

Quantile forecasts from models conditioned on money growth are combined with forecasts from models based on the University of Michigan inflation survey. To investigate this, point and density forecasts are constructed using a combination technique based on tick-loss minimization. The performance is investigated by comparing the performance of optimally combined forecasts with a benchmark average combination. I begin with a presentation of the data and the move to a discussion of the methodology. Thereafter, the results are presented before arriving at the conclusion.

2 Data

2.1 SPF & NFCI

Data from the survey of professional forecasters (SPF) is used for real output, unemployment, and inflation. The SPF is published every quarter and is the oldest survey of macroeconomic forecasts of that form in the United States [5]. It surveys professional forecasters on a large number of key variables connected to the state of the United States' economy. The data collection began in 1968 through the American Statistical Association and the National Bureau of Economic Research and was taken over by the Federal Reserve Bank of Philadelphia in 1990. On their website, the individual, mean, and median forecasts are released to the public at no charge. For each variable of interest, the forecasters provide quarterly point forecasts ranging from one up to four quarters ahead. Like Adrian et al. (2021), I use the median forecasts for quarter-over-quarter real GDP growth, the quarterly average unemployment rate and the quarter-over-quarter GDP price index inflation.

In addition, financial conditions are included into the conditioning set in the form of the Federal Reserve Bank of Chicago's National Financial Conditions Index (NFCI). The index provides a standardized value of financial tight- or looseness, where a low value denotes looser financial conditions and a higher value a tighter financial environment. The collection began January of 1971 and a complete description of the index's methodology employed can be found in Brave and Butters [6].

2.2 UMICH Inflation Survey

To get additional insights on the inflationary movements, I use an additional data-set regarding this phenomenon from the University of Michigan. The University of Michigan Inflation Expectation survey is a measurement of price rise expectations. It is published monthly and the collection started in January of 1978 and the latest publication that considered is May 2021. For illustrative purposes, its quarterly values are shown together with the SPF in Figure 1. I chose the quarterly points as the monthly values for the corresponding quarter, that is February for the first quarter, May for the second, and so on. We observe that in the period from 1980s, the SPF estimates a higher inflation that the University of Michigan and in the last 30 years, the SPF has become more modest and consistently estimates lower inflation than UMICH does. Inclusion of this data-set is sensible as it provides a different view on inflation, and especially in recent times might be a refreshing look to consider.

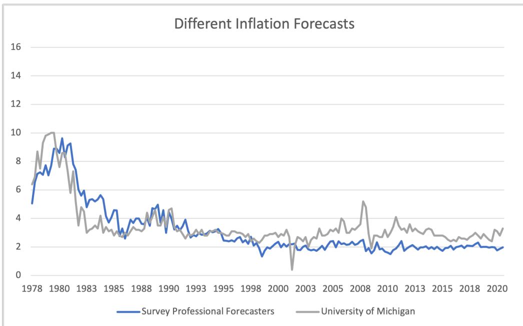

Figure 1: University of Michigan inflation survey (grey), compared to the SPF (blue)

2.3 M3 Money Growth

As a final variable, I consider a measure that is related to inflation through macroeconomic theory: Money growth. As outlined earlier, an increase in circulating money has been linked to inflation and there has been a clear trend of new money creation in the United States in the last year. We consider the growth of the money supply M3, a measure which includes cash and checking deposits, saving deposits, money market securities, mutual funds and short term repurchase agreements as well as larger liquid assets. It occurs that this measure revolves around a fixed mean and has some higher values occurring incidentally. On average, M3 increased by 0.58 percent per month and the most outstanding values all occurred in 2020, with April logging a 6.42 percent increase and May a 4.98 increase. Other notable increases occurred in 1975 and 1983. Figure 2 shows the plots of the time series of the macroeconomic variables, with inflation and unemployment in relation to M3 growth.

Figure 2: Historical time series for all variables, in relation to the median SPF forecasts, money supply (for unemployment and inflation, one quarter ahead), and the University of Michigan survey (for inflation four quarters ahead)

3 Methodology

The methodology is divided into three parts. First, I outline some principles of forecast combinations and then we move to the inner workings of quantile regression. Next, the construction of probabilistic forecasts and predictive distributions is discussed. This is complemented by the theory of the out-of-sample evaluation.

3.1 Combination of Forecasts

To a decision maker like government agencies or business executives, multiple forecasts regarding the same variable might be available. The SPF alone collects the responses of around 40 individuals. That means that in practice, a large amount of forecasts for macroeconomic variables exist. Different forecasters arrive at different forecasts for a variety of reasons, so it is a relevant topic to consider a possible combination of them. We might, for one, just consider one particular forecast, like it appears to be done in the vindication of lockdown measures. On the other hand, we can also consider a combination of the different available forecasts. These might be either obtained through different techniques with regards to their regression analysis or from similar techniques that consider different explanatory variables. Forecasting expert Allan Timmerman writes that the combination of different forecasts has been empirically useful in many fields, including inflation [7]. Therefore, the choice to do so seems intuitive for this research. Combining forecasts is sensible since different forecasters might have asymmetric access to information and different assumptions that they base their models on. The fact that differences might arise for reasons that are not entirely transparent warrants a combination of these forecasts even more.

Let $\hat{y}^1$ be the forecast arising from the first conditioning set of variables, $\hat{y}^2$ from the second set, and so forth. Then, there exists multiple ways to combine these forecasts. For one, consider the linear combination w of n forecasts:

$$w = \xi_1 \hat{y}^1 + \xi_2 \hat{y}^2 + ... + \xi_n \hat{y}^n,$$

with ξ_i the linear weight allocated to forecast i, such that $\sum_{i=1}^{n} \xi_i = 1$. A simple example would be assigning equal weights to all forecasts such that $\xi_1 = \xi_2 = ... = \xi_n = \frac{1}{N}$, which results in the arithmetic mean of all forecasts. This mean will further be referred to as the *average combination* and is used throughout this work as a benchmark to test the optimal combination against. In my analysis, I restricted the weights ξ_i to be non-negative. While this is not strictly necessary, I decided to do so on the basis of simplicity and recommend a further investigation of potentially negative weighed forecasts for further research. There are different ways to establish a preferred combination of the forecasts and the majority of literature bases it on some loss function that depends on the forecast error in which the combined forecast results. The optimal combination w^* then minimises the expectation of the loss function given the information available at that time. Formally,

$$w^*_{t+h,t}= \text{argmin}_{w_{t+h,t}}\colon E[L(e^c_{t+h,t}(w_{t+h,t}))|F_t],$$

where the expectation E is taken over the conditional distribution of combined forecast error e^c. Frequently, the mean squared error (MSE) is taken as a first loss function, defined as $L(y_{t+h}, \hat{y}_{t+h,t}) = (y_{t+h} - \hat{y}_{t+h,t})^2$. For combining quantiles, the mean squared error is not a very useful measure as it reduces to ordinary least squares. Instead, Giacomini and Komunjer suggest the use of a tick-loss function that takes the quantiles into account [8]. For the αth quantile, the loss is defined as $L(y_{t+h}, \hat{y}_{t+h,t}) = (\alpha - I(e_t < 0))e_t$, with $I()$ the indicator function. We accordingly find the optimal weights by minimizing the sum of all observations of the tick-loss function for a quantile α.

3.2 Quantile Regression

Quantile regression is a method that is frequently used to assess up- and downside risk, especially for macroeconomic variables. The ability to create forecasts at the tails of probability distributions is of crucial importance to monitor the potential occurrence of rare events, like the transition from financial up- to downturn. The 5th and 95th percentiles are of special interest as they capture the edges of forecasts that are easily overlooked with other types of analysis.

Let y_{t+h} be the observed value of a variable of interest. For inflation and GDP growth, y_{t+h} represents the annualized average growth between quarter t and quarter $t+h$. For unemployment, it is the average unemployment rate in quarter $t+h$. Similarly to Adrian et al., I perform quantile regression like in Koenker and Bassett [9] on the SPF forecast errors of the three different variables. That is, quantile regression is performed on $e^{SPF}_{t+h|t} = y_{t+h} - \hat{y}^{SPF}_{t+h|t}$, where $\hat{y}^{SPF}_{t+h|t}$ is the h-quarter ahead median SPF forecast for y_{t+h}. Contrary to their approach, I run two quantile regressions for each variable, conditioned on two different sets of variables. For $i = 1, 2$, define the τ-th quantile $Q^i_{e_{t+h|t}}(\tau) = \inf\{q \in R | F_e(q|x^i_t) \geq \tau\}$, where F_e is the h-quarter ahead SPF forecast error distribution conditional on the variables of model i and x^i the set of variables of that model. The quantile regression coefficients for model i, β^i are chosen to minimize the sum of quantile weighted absolute residuals, that is

$$\hat{\beta^i}_\tau = \text{argmin}_{\beta^i_\tau \in R} \sum_{t=1}^{T-h} (\tau * I(e^{SPF}_{t+h|t} > x^{i'}_t \beta^i_\tau)|e^{SPF}_{t+h|t} - x^i_t \beta^i_\tau|$$
$$+(1\text{-}\tau) I(e^{SPF}_{t+h|t} < x^{i'}_t \beta^i_\tau)|e^{SPF}_{t+h|t} - x^i_t \beta^i_\tau|)$$

As a result, the quantile prediction $\hat{Q}^i(\tau)$ gives a linear estimate of the τ-th quantile of e^{SPF} conditional on the variables of model i. Denote by $\mathbf{\hat{Q}^i}$ the matrix containing all quantile predictions resulting from model i. The matrix containing all optimal linear quantile combinations is then $\mathbf{Q}^*$, given by $\Xi^*_1 \mathbf{\hat{Q}^1} + \Xi^*_2 \mathbf{\hat{Q}^2}$, where Ξ^*_i is the diagonal matrix containing the optimal weights ξ assigned to each individual quantile of model i.

3.3 Predictive distributions

Each quantile regression now provides estimates for specific quantiles of the variables of interest, conditional on a set of variables. For applications that only require certain quantiles (e.g. growth at risk) these estimates suffice while for other applications that require probabilistic ranges where variables could fall require a full conditional probability distribution. The approach from Adrian et al. is followed to fit skewed t-distribution to the estimated conditional quantiles. I will first do this for the general models and for the combined model and the combined simple average model. The general probability density function is given by:

$$f(y, \mu, \sigma, \alpha, v) = \frac{2}{\sigma} t(\frac{y-\mu}{\sigma}; v) T(\alpha(\frac{y-\mu}{\sigma}) \sqrt{\frac{v+1}{v+(\frac{y-\mu}{\sigma})}}; v+1),$$

with μ a location parameter, σ a scale parameter, α a shape parameter, and v the degrees of freedom. t and T denote the *pdf* and *cdf* of the student's t distribution.

Like outlined in Adrian et al., this general type of distribution is widely used as it allows for many distributive properties like skewness and custom tails. For each quarter, distributions are fitted such that they minimize the squared differences between the implied quantiles and the quantile regression estimates. Formally:

$$\text{argmin}_{\mu,\sigma,\alpha,v} \sum (\hat{Q}_{yt+h}(\tau) - F^{-1}(\tau, \sigma, \alpha, v)),$$

with F^{-1} the quantile of the distribution.

To compare the probability distributions, a baseline version that does not include additional variables is constructed, that is, only based on the historical distribution of errors. The plots of the predictive distributions is then presented for two specific quarters in the results section.

3.4 Out-of-sample evaluation

To evaluate how well the combined models work, an out-of-sample analysis is conducted. This is done from the period of 1993Q1 to the current 2021Q2 with the training sample from 1973Q1 to 1992Q4. Here, we compare the performance of the optimally combined model with the average combined model. For each quarter, predictive predictive densities were re-estimated like outlined in 3.3 so that they are now fitted to the optimally or average combined quantiles, respectively. The accuracy of the out-of-sample density forecasts is then compared by the predictive scores $PS_{\hat{f}_{y_{t+h}|I_t}} = \hat{f}_{y_{t+h}|I_t}(y^0_{t+h})$ assigned to each forecast. Additionally, the probability integral transforms $PIT_{\hat{F}_{y_{t+h}|I_t}}(y^0_{t+h}) = \hat{F}_{y_{t+h}|I_t}(y^0_{t+h})$ are computed, representing the fitted *cdf* evaluated at the realized value. The PITs are uniformly distributed if the calibration of the predictive distributions is appropriate.

4 Results

4.1 Quantile Regression

This section shows the plots of the quantile regressions based on different variables and discusses the results. Some of the findings replicate Adrian et al. while others are completely new.

4.1.1 Inflation

<u>One Quarter Ahead</u>

Figure 3 shows the plot of the quantile regression of the one quarter ahead inflation forecast errors, based on M3 money growth. We see that the explanatory impact of M3 on inflation is not necessarily strong and the quantiles seem to only barely be related to the explanatory variable. This might be connected to a lagged effect of money growth in inflation. Different lags where tried but there seems to be a very low effect of the **short term** money growth on inflation. As I deem the usefulness of M3 inclusion low here, I abstain from conducting further analysis on this set and will instead focus on the four-quarter ahead inflation forecast.

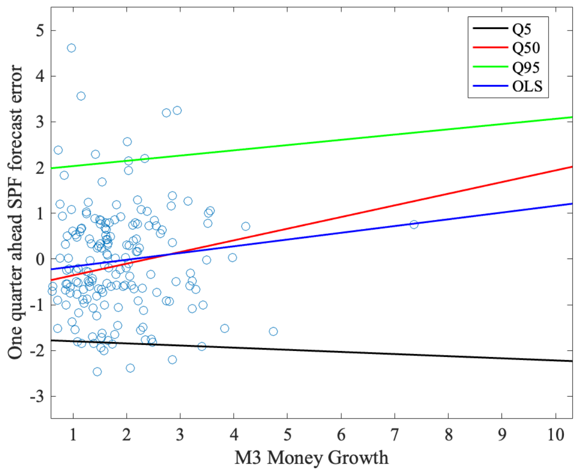

Figure 3: Quantile regression: Inflation one quarter ahead

Four Quarters Ahead

Figure 4 shows the quantile regression results of the four-quarter-ahead forecast errors on inflation based on the University of Michigan survey and the M3 Money growth. These findings are particularly interesting. The UMICH survey strongly influences the 5th quartile while the

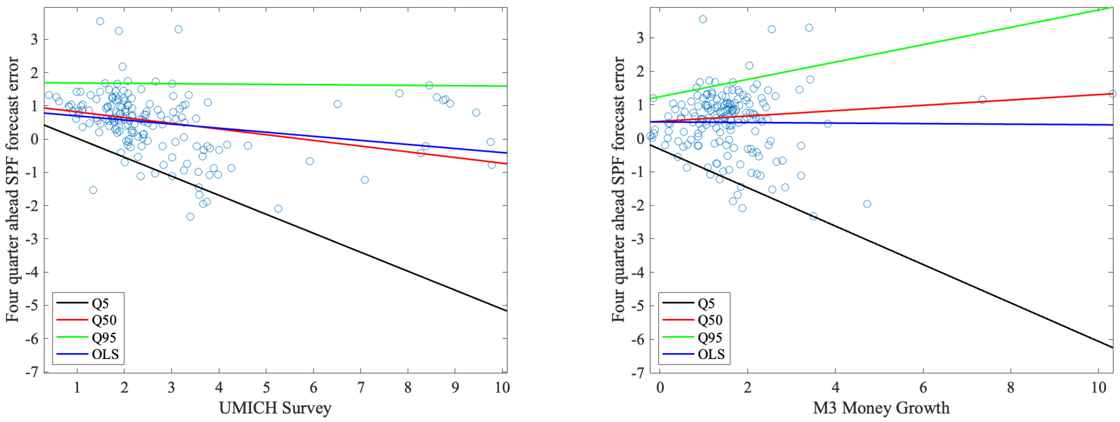

Figure 4: Quantile regression: Inflation four quarters ahead, conditioned on M3 (left) and NFCI (right)

95th quantile is completely flat. This means that some of the strongest underestimations of the inflation by the SPF are connected to higher values of the University of Michigan survey. This gives rise to the possibility that the University of Michigan survey incorporates additional perspectives into the inflation expectations. There might be some implications of possible endogeneity as one survey might be used to predict the other. While I abstain from a further discussion on this, I recognize the potential shortcomings that this could cause. With regards to M3, we see that the M3 yields similar relations for the 5th quantile as UMICH, while the 95th quantile shows a small upwards slope. Since these two models show some clear similarities for the 5th quantile and quite some differences for the others, I conduct a combination of forecast analysis for them.

4.1.2 Unemployment

Figure 5 shows the results of the quantile regressions of the SPF forecast errors on one quarter ahead unemployment, based on the M3 money growth and the NFCI. While the graph on the right is a new result, the left part corresponds to the findings from Adrian et al. Note the slight differences in the slopes of the quantiles to Adrian et al.: this is due to the inclusion of newer observations that were not included in their work. We see that between NFCI and M3, the results are quite different. In particular, the slope of the 95th quantile is much steeper when based on the NFCI and the slope of the 5th quantile much more negative when based on the M3

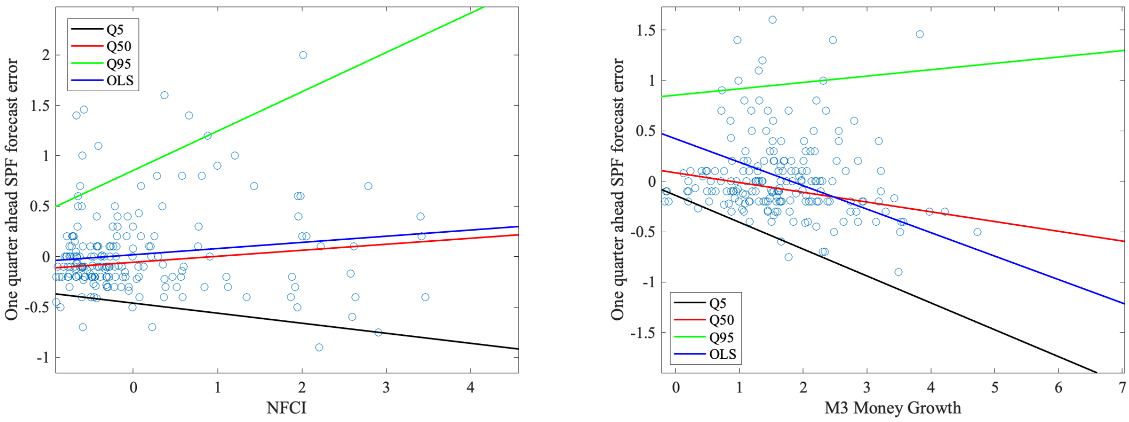

Figure 5: Quantile regression: Unemployment one quarter ahead, conditioned on UMICH (left) and M3 (right)

money growth. We also see that unemployment does seem to be much more immediately affected by a growth in money supply, compared to inflation. Lastly, we see quite some differences in the median and OLS estimations. Conditioned on the NFCI, these are almost flat, while for M3 they are slightly downward sloping. This implies that the median forecast errors could be negatively correlated to M3 money growth. Since the differences between these regressions are quite substantial, I also conduct a combined quantile aproach for this.

4.1.3 GDP growth

As the last result, we consider GDP growth at the one quarter ahead horizon. While the analysis one the left of Figure 6, conditioned on the NFCI, corresponds to the results from Adrian et al., the incorporation of M3 is again new. We observe a slightly different slope of the 95th quantile line due to the inclusion of more recent data, in particular some outliers in 2021. Compared with the NFCI, conditioning on M3 money growth lead to much higher slopes, both positively for the 95th quantile and negatively for the 5th quantile. This implies that very large errors of the SPF can be stronger associated with absolute higher values in M3 growth. This makes intuitive sense, as a rapid increase usually comes in a time of economic turbulence. This confirms Adrian et al.'s findings that during economic uncertainty, forecast accuracy declines. In addition, the changes in the errors cannot only be explained by financial conditions like they outlined but also by money growth, in fact even on a more rapid scale. Since the differences between the regressions conditioned on M3 and NFCI are of interest, we also conduct a forecast combination approach for this data. The coefficients for all other quantiles can be found in Appendix A.1.

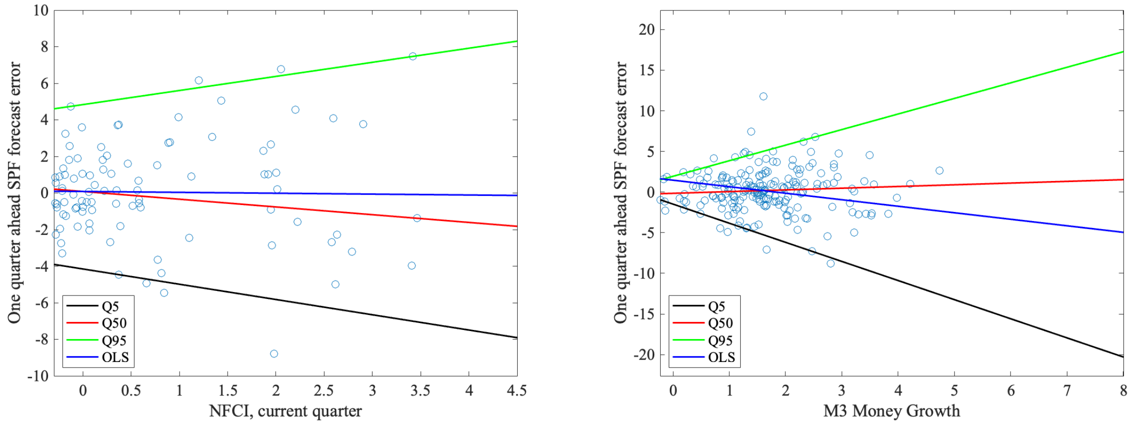

Figure 6: Quantile regression: GDP growth one quarter ahead, conditioned on NFCI (left) and M3 (right)

4.2 Forecast combination

Based on the previous findings, it is deemed a sensible assumption to use forecast combination for three of the data-sets. I combine two models with different explanatory variables for GDP growth (1QA), unemployment (1QA), and inflation (4QA). Like outlined in the methodology, there is assumed to be a linear combination with weights ξ_1 and ξ_2 corresponding to the weight of the particular model's forecast. The best linear combination was found by minimizing the asymmetric linear loss function (*tick loss*): $T_\tau(e_{t+1}) = (\tau - I(e_{t+1} < 0)e_{t+1}$ over all observations. Table 1 shows the optimal weights and loss function values resulting from these weights. For comparison, the loss function values of the average combination with $\xi_1 = \xi_2 = 0.5$ are also shown. The optimization was done over the entire range of realized values and the coefficients therefore present a certain upper-bound solution that cannot quite be achieved in out-of-sample estimation. Running out-of-sample estimation where optimal weights are either determined for a certain training period and then kept fixed or continously updated are others method that can be investigated to assess the performance of combinations in real-time. In 4.3, the out-of-sample technique is employed.

The bold entries denote the cases where the tick-loss criterion resulted in a forecast combination. Note that the relative differences in the tick-loss function values occur due to the different relative sizes of the dependent variables. We see that in almost all cases, a combination of forecasts is preferred. In the majority of those, a smaller weight is attributed to M3 than to the other variable. This can be explained by the fact that the conditioning on M3 causes very steep quantiles that already incorporate large changes at smaller weights. It is interesting to note that for quite some of the data, a certain model is preferred 100 percent over another, for a specific quantile. Adopting the notation from [8], if a conditional quantile q_1 from one model is completely preferred over the q_2 from another model, q_1 encompasses q_2. This has interesting implications for for the distribution fitting which is done next.

y	i	j	τ	ξ_i^*	ξ_j^*	$\sum_t(T_{\xi_i^*,\xi_j^*})$	$\sum_t(T_{0.5,0.5})$
$GDP1A$	$NFCI$	$M3$	$5th$	0	1	78.249	82.501
$\boldsymbol{GDP1A}$	$NFCI$	$M3$	$25th$	**0.45**	**0.55**	193.658	193.678
$\boldsymbol{GDP1A}$	$NFCI$	$M3$	$75th$	**0.92**	**0.08**	193.758	193.937
$GDP1A$	$NFCI$	$M3$	$95th$	0	1	74.535	77.815
$UNEMP1A$	$NFCI$	$M3$	$5th$	0	1	12.353	13.939
$\boldsymbol{UNEMP1A}$	$NFCI$	$M3$	$25th$	**0.19**	**0.81**	25.546	25.595
$\boldsymbol{UNEMP1A}$	$NFCI$	$M3$	$75th$	**0.64**	**0.36**	26.840	26.842
$\boldsymbol{UNEMP1A}$	$NFCI$	$M3$	$95th$	**0.71**	**0.29**	16.782	16.835
$\boldsymbol{INFL4A}$	$UMICH$	$M3$	$5th$	**0.49**	**0.51**	11.028	11.035
$\boldsymbol{INFL4A}$	$UMICH$	$M3$	$25th$	**0.9**	**0.1**	30.215	30.482
$\boldsymbol{INFL4A}$	$UMICH$	$M3$	$75th$	**0.86**	**0.14**	32.399	32.481
$\boldsymbol{INFL4A}$	$UMICH$	$M3$	$95th$	**0.82**	**0.18**	13.012	13.039

Table 1: Optimal Combination Weights, Entire Sample

4.3 Predictive Distributions

Next, the predictive fitted skewed t-distributions of the combined models together with the benchmark historical models are shown. This is done for two periods: In the 2nd quarter of 2020, a very uncertain quarter, and in the fourth quarter of 2017, a more stable quarter. Comparing the distributions to the average distribution was omitted as the distributions show only slight differences to the optimally distributed forecast.

We see that the combined forecasts appears to be more accommodative of a large range of outcomes and during times of uncertainty, its mean is closer to the real value. This could be due to the fact that it contains information from two forecasts, where one might paint a more pessimistic picture for the macroeconomic variables. We see that for inflation and unemployment, the distributions have a very strong peak at their mean and the differences to the historical distribution is less pronounced. For these two variables, it appears that during times of higher certainty, the combined forecast does not show great differences to the historical one.

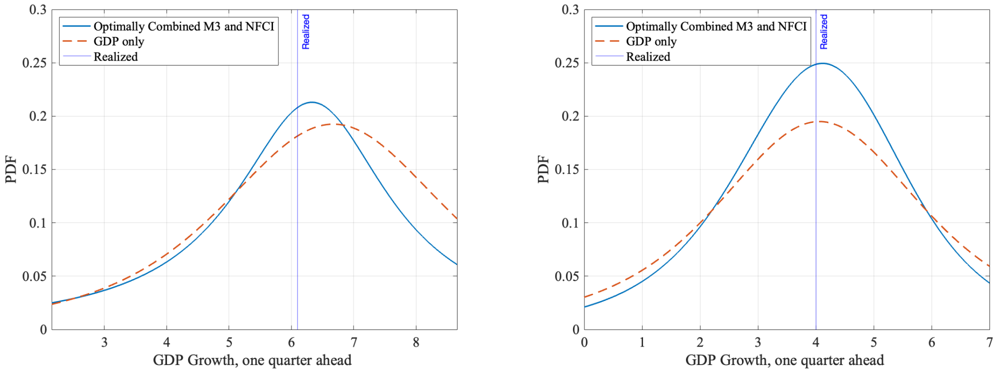

Figure 7: Optimally constructed predictive distributions (blue) and GDP-only distributions (red) for GDP growth, 2020Q2 (left), 2017Q4(right),

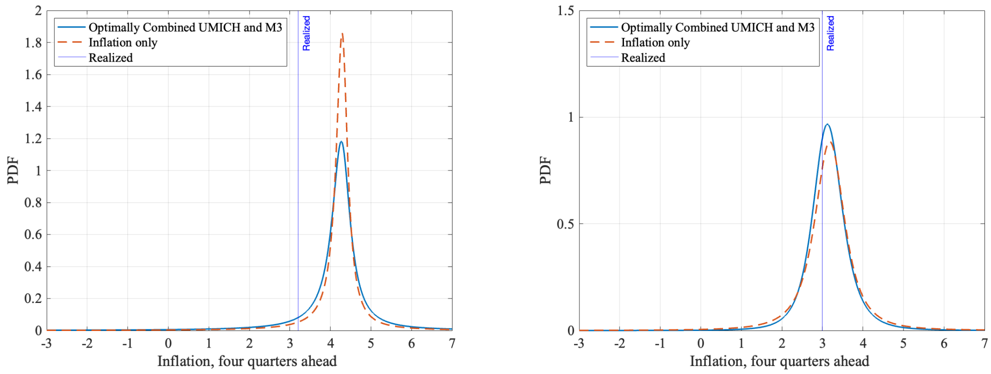

Figure 8: Optimally constructed predictive distributions (blue) and Inflation-only distributions (red) for Inflation, 2020Q2 (left), 2017Q4(right),

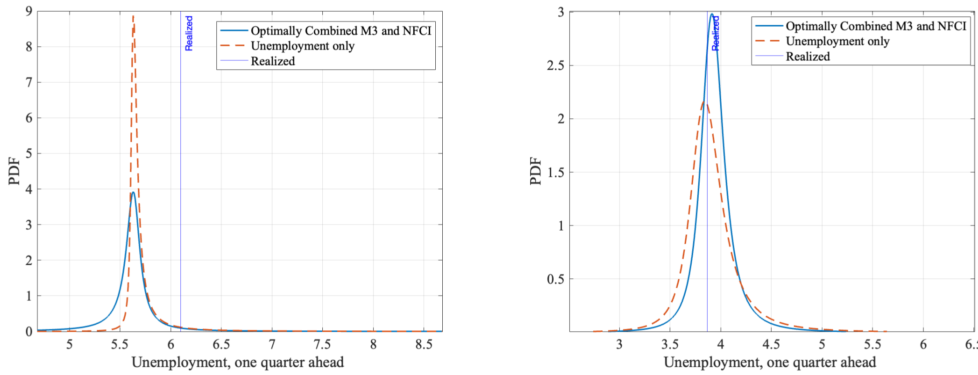

Figure 9: Optimally constructed predictive distributions (blue) and Unemployment-only distributions (red) for Unemployment, 2020Q2 (left), 2017Q4(right),

4.4 Out-of-Sample Estimation

Out-of-sample estimation is performed to evaluate the real-time performance of the models. In particular, I compare the performance of the optimally combined forecasts with the average combined forecasts. This is done in the form of out-of-sample density forecasting, beginning in Q3 of 1992. The optimal combination weights ξ_i are found by minimizing the tick-loos function over the all observations preceding the out-of-sample period. Table 2 shows these optimal weights, with sum of the tick-losses here the sum over all quantiles - the results for the individual quantiles can be found in Appendix A. The out-of-sample point forecasts are shown in blue for illustration in Figure 10. It can be seen that for unemployment and inflation they are almost indistinguishable to the in-sample. For GDP growth, they perform a little worse and have problems capturing the 95th quantile. This indicates that upside GDP risk is more difficult to accomodate with this model. The individual weights are interesting in so far that M3 is frequently encompassed by other variables, especially for the 75th and 95th quantile. Like mentioned before, this could be due to the fact that its quantiles are of larger magnitute or already incorporated in another fashion. The absolute differences of the tick-loss are very small, raising some questions about whether the optimization of weights is justified.

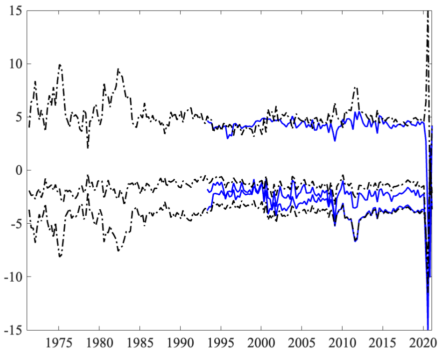

a) GDP growth

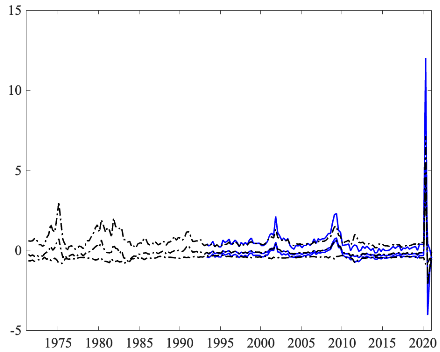

b) Unemployment

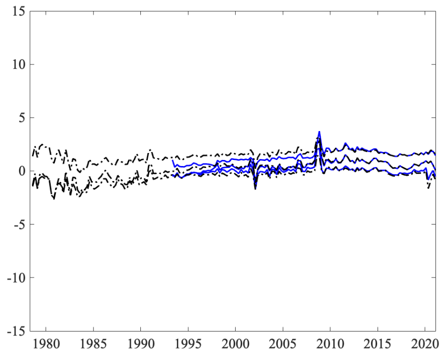

c) Inflation

Figure 10: Out-of-Sample Point Forecasts.

y	i	j	$\xi^*_{i,5}$	$\xi^*_{i,25}$	$\xi^*_{i,75}$	$\xi^*_{i,95}$	$\sum_t(T_{\xi^*_i,\xi^*_j})$	$\sum_t(T_{0.5,0.5)}$
$GDP1A$	$NFCI$	$M3$	0.97	0.11	1	1	242.294	247.210
$UNEMP1A$	$NFCI$	$M3$	0.79	0.71	0.65	1	23.903	24.078
$INFL4A$	$UMICH$	$M3$	0.46	1	1	1	45.924	47.047

Table 2: Optimal weights, Used for Out-of-Sample Analysis

The shown weights were then used to fit the distributions and compute the PITs for the out-of-sample period. Figure 11 shows the predictive scores of both models. We see that the optimally combined model consistently scores higher than the average model. This implies that the distribution derived from the optimally combined quantiles is frequently more accurate than the average combined one. However, the differences are not very large. Figure 12 shows the probability integral transforms of the models. For GDP growth, the optimal is slightly better as it never leaves the confidence bands. For unemployment, neither model leaves the confidence bands, indicating only very slight differences in the models. Inflation is frequently outside the confidence band for both models. This indicates that the out-of-sample accuracy is high for the optimal and average combination of GDP growth and unemployment, respectively, but not for inflation. This most likely stems from the fact that the training period is too short for inflation. Because the collection of the UMICH survey only starts in 1978, quite some initial observations could not be used. It is recommended to try different sample sizes for out-of-sample analysis but this result is presented nonetheless so illustrate the potential shortcomings. Figure 11 plots the differences in log-scores for the forecasts from the equally weighed and optimal forecasts. We see that for GDP growth the optimal combination consistently scores higher. For unemployment, the scores of both models are similar with some scattered improvements for the optimal combination. For inflation, there is no clear improvement of the optimal combination. This is possibly due to a misalignment of the out-of-sample periods.

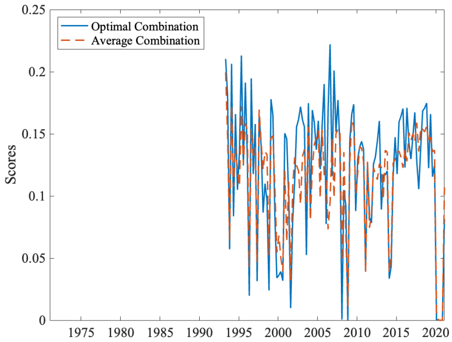

a) GDP growth

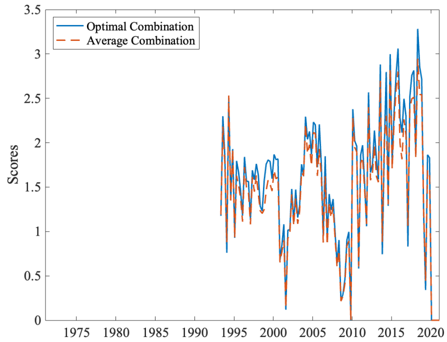

b) Unemployment

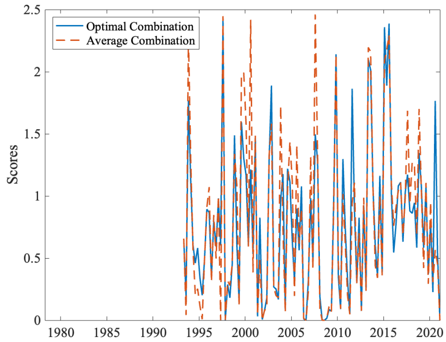

c) Inflation

Figure 11: Forecasting log-scores for the the Equally Weighed and Optimal Combination.

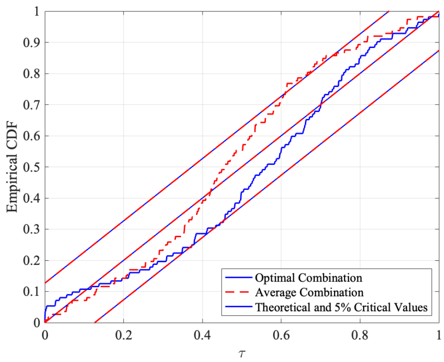

a) GDP growth

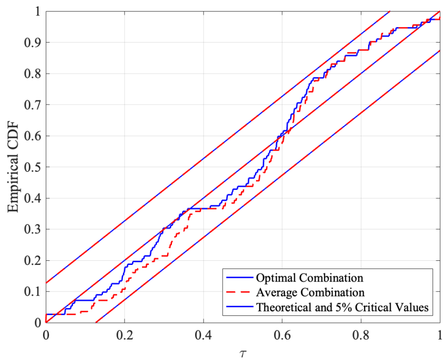

b) Unemployment

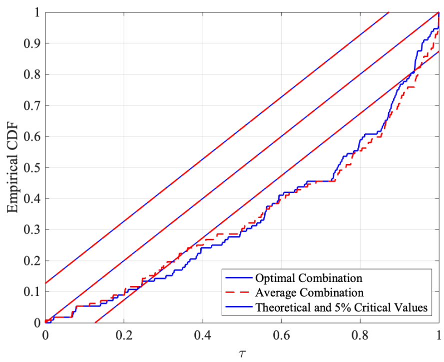

c) Inflation

Figure 12: Probability Integral Transforms. Confidence bands accounting for sample uncertainty are adopted from Rossi & Sekhposyan [10].

5 Conclusion

In this thesis, the impact of money growth and the UMICH inflation survey was assessed in terms of the improvement of forecasts that were initially based on financial conditions. Both variables show interesting interpretations for quantile regressions of macroeconomic variables. The combination of these variables leads to visible improvements for constructing predictive densities for GDP growth. For inflation and unemployment, these effects were less significant. Additionally, forecast combination for tick-loss minimization and equally weighed forecasts was conducted to investigate the implication for density forecasts. It was found that both techniques slightly improve density forecasts and out-of-sample results for unemployment and GDP growth. In particular, during time of economic turbulence, the inclusion and combination of the additional variables lead to a greater accommodation of up- and downside risk in the densities of the models. The differences during times of economic stability are less pronounced, especially for inflation and unemployment. Out-of-sample estimation for performed to compare the performance between the equally weighed and optimally weighed forecast combinations. The differences between the optimally weighed and equally weighed forecast were found to not be strongly pronounced and the equally weighed combination produces strong forecasts. The tick-loss differences between optimally and equally weighed models are small and require further research. For brevity, researches can choose equally weighed quantiles as they produce forecasts that are similarly accurate. The computation of optimal weights of a low number of quantiles is however not computationally complex and should be conducted to see if these differences increase with dimensionality. Probability integral transformations and forecast scores were computed which indicate that the optimally weighed model is slightly better specified and produces better forecasts for GDP growth and unemployment. Further research should be directed at finding strong short-term inflation indicators. Additionally, investigating the role of equally weighed forecast combinations and the trade-off between accuracy and model complexity in this regard is encouraged.

References

[1] Eric M Leeper. Fiscal policy and inflation: Pondering the imponderables. Working Paper 9506, National Bureau of Economic Research, February 2003.

[2] Neil Irwin. Inflation is here. what now? *The New York Times*, May 2021.

[3] Tobias Adrian, Nina Boyarchenko, and Domenico Giannone. Vulnerable growth. *American Economic Review*, 109(4):1263–89, April 2019.

[4] Patrick A. Adams, Tobias Adrian, Nina Boyarchenko, and Domenico Giannone. Forecasting macroeconomic risks. *International Journal of Forecasting*, 37(3):1173–1191, 2021.

[5] Dean Croushore. Introducing: The survey of professional forecasters. *Business Review*, November/December 1993:3–15.

[6] Scott Brave and R. Butters. Diagnosing the financial system: Financial conditions and financial stress. *International Journal of Central Banking*, 8:191–239, 06 2012.

[7] Allan Timmermann. Forecast Combinations. In G. Elliott, C. Granger, and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 1 of *Handbook of Economic Forecasting*, chapter 4, pages 135–196. Elsevier, 2006.

[8] Raffaella Giacomini and Ivana Komunjer. Evaluation and combination of conditional quantile forecasts. *Journal of Business & Economic Statistics*, 23(4):416–431, 2005.

[9] Roger Koenker and Gilbert Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.

[10] Barbara Rossi and Tatevik Sekhposyan. Alternative tests for correct specification of conditional predictive densities. *Journal of Econometrics*, 208(2):638–657, 2019.

Appendix

I: Quantile Regression Coefficients for all Quantiles

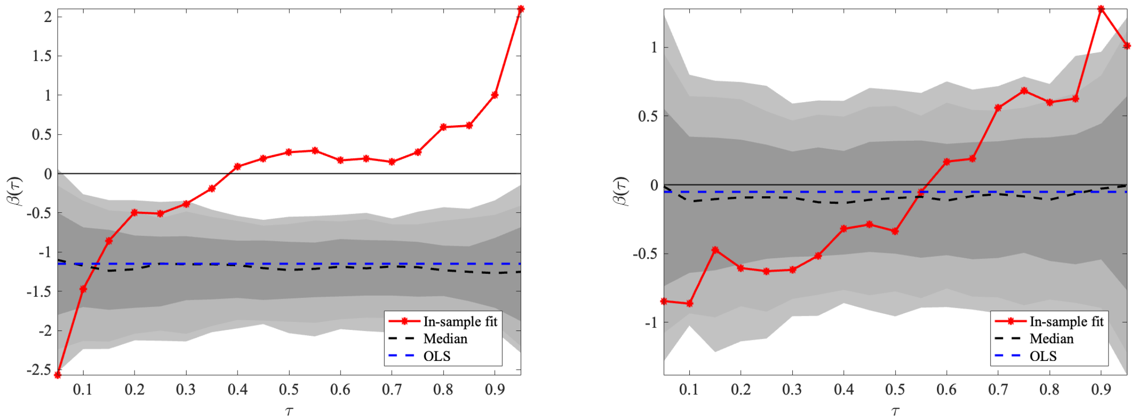

a) Quantile coefficients for GDP growth one quarter ahead, conditioned on NFCI (left) and M3 (right)

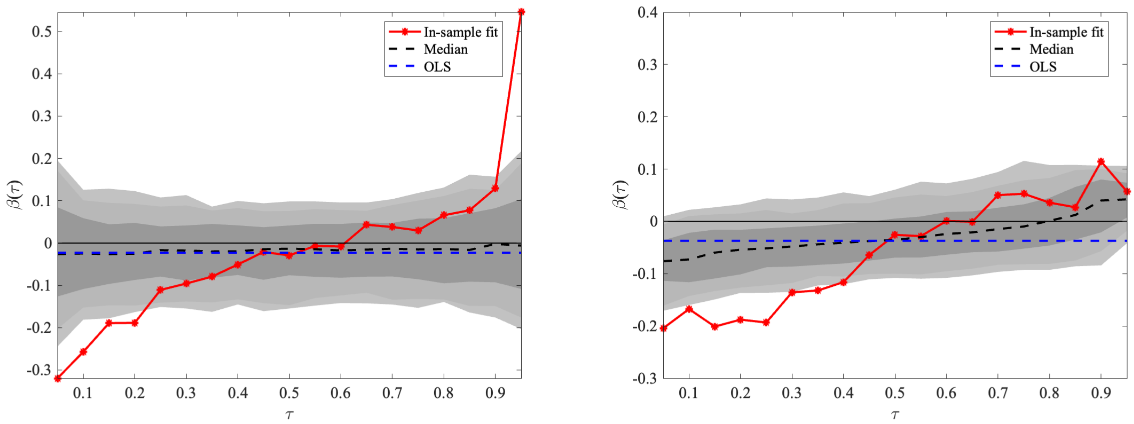

b) Quantile coefficients for inflation, four quarters ahead, conditioned on UMICH (left) and M3 (right)

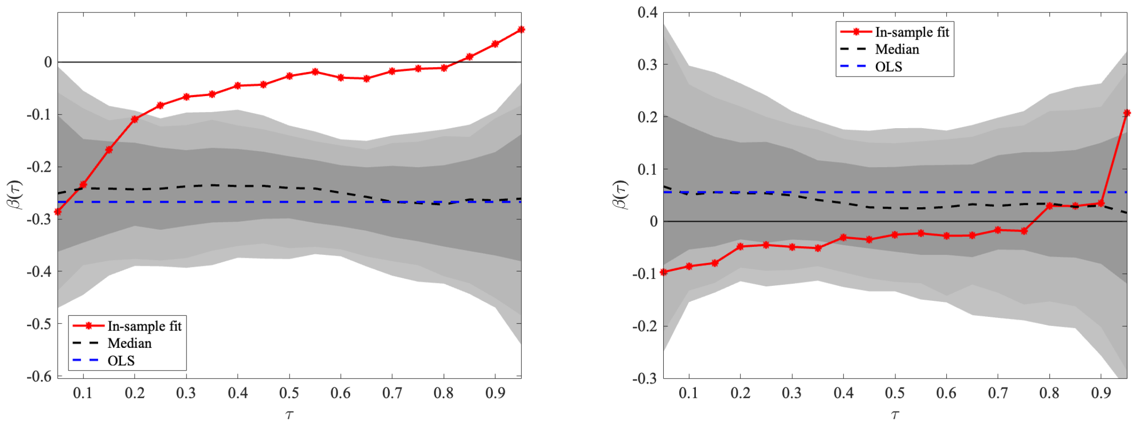

c) Quantile coefficients for unemployment one quarter ahead, conditioned on NFCI (left) and M3 (right)

II: Individual Tick-Loss Values, Out of Sample

y	i	j	τ	ξ_i^*	ξ_j^*	$\sum_t(T_{\xi_i^*,\xi_j^*})$	$\sum_t(T_{0.5,0.5})$
$GDP1A$	$NFCI$	$M3$	5	0.97	0.03	26.830	27.943
$GDP1A$	$NFCI$	$M3$	25	0.11	0.89	86.496	86.520
$GDP1A$	$NFCI$	$M3$	75	1	0	92.236	93.608
$GDP1A$	$NFCI$	$M3$	95	1	0	36.732	39.139
$UNEMP1A$	$NFCI$	$M3$	5	0.79	0.21	3.158	3.219
$UNEMP1A$	$NFCI$	$M3$	25	0.71	0.29	7.617	7.626
$UNEMP1A$	$NFCI$	$M3$	75	0.65	0.35	9.068	9.071
$UNEMP1A$	$NFCI$	$M3$	95	1	0	4.060	4.162
$INFL4A$	$UMICH$	$M3$	5	0.46	0.54	4.912	4.945
$INFL4A$	$UMICH$	$M3$	25	1	0	17.177	17.532
$INFL4A$	$UMICH$	$M3$	75	1	0	17.532	17.845
$INFL4A$	$UMICH$	$M3$	95	1	0	6.303	6.724